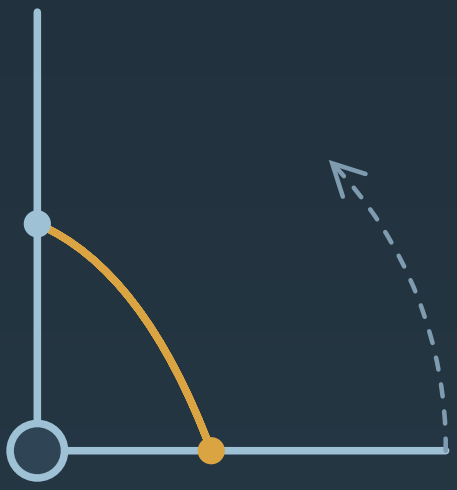


FORCE IS GENERATED. THE TENDON TRANSMITS IT.

Human Biology for the Medical Maker

Anatomy, physiology, and collaborative design
for software developers and practical makers

A SIXTEEN-WEEK SELF-STUDY PROGRAMME

Human Biology for the Medical Maker

Chris

Version 1.0, 2026-10-02

Table of Contents

Colophon . . . 1
Preface: what this book is for . . . 2
What you will be able to do . . . 2
What this book deliberately does not do . . . 3
The running project . . . 3
The device and software sections . . . 4
A note on tone . . . 4
How to use this book . . . 5
The shape of a chapter . . . 5
A repeatable weekly routine . . . 5
Spaced repetition, concretely . . . 6
Building the book . . . 6
Conventions used . . . 7
Part I: Reading the body . . . 8
1. Reading the body and understanding living systems . . . 9
1.1. Why this chapter exists . . . 9
1.2. Anatomy answers *where*; physiology answers *how* . . . 9
1.2.1. Levels of organisation . . . 9
1.3. A shared coordinate system . . . 10
1.3.1. Directional terms . . . 10
1.3.2. Planes and sections . . . 12
1.3.3. Body cavities and regions . . . 13
1.4. Cells maintain an internal environment . . . 14
1.4.1. Tonicity: why this is not academic . . . 15
1.4.2. Energy, and why oxygen matters . . . 15
1.4.3. Enough chemistry to not be lost . . . 15
1.5. Homeostasis is ongoing regulation . . . 16
1.5.1. Negative feedback . . . 16
1.5.2. Positive feedback . . . 17
1.5.3. Variation, and why "normal" is a trap . . . 17
1.6. Engineering bridge: scales, units, and language . . . 18
1.6.1. Orders of magnitude . . . 18
1.6.2. Units you will meet . . . 18
1.6.3. Reading medical words . . . 19
1.7. Devices, software, and training technology . . . 20
1.7.1. Clinical devices . . . 20
1.7.2. Software . . . 21
1.7.3. Training and simulation . . . 22

1.8. Worked example: an ambiguous instruction . . . 23
1.9. Workshop exercise . . . 24
1.10. Questions . . . 24
1.11. Answers and reasoning . . . 25
1.12. Interactive tools . . . 25
1.13. Video companion and reading . . . 26
1.14. Exit criteria . . . 26
1.15. What this chapter does not establish . . . 26
2. Foundations of chemistry: atoms to reactions, and how chemists work . . . 27
2.1. Why this chapter exists . . . 27
2.2. Atoms: what everything is made of . . . 27
2.2.1. Inside an atom . . . 28
2.2.2. Electron shells and why atoms react . . . 28
2.2.3. Ions and electronegativity . . . 28
2.2.4. Isotopes . . . 29
2.3. The periodic table . . . 29
2.3.1. The elements of the body . . . 30
2.3.2. The elements of devices . . . 31
2.4. States of matter and gases . . . 31
2.4.1. Solids, liquids, gases, and changes between them . . . 31
2.4.2. The gas laws . . . 32
2.5. Amounts and concentrations . . . 32
2.5.1. The mole . . . 32
2.5.2. Concentration units, and why medicine uses so many . . . 33
2.5.3. Dilution . . . 34
2.5.4. Solutions and solubility . . . 34
2.6. Chemical reactions . . . 34
2.6.1. Equations and stoichiometry . . . 34
2.6.2. Energy in reactions . . . 35
2.6.3. Rates and kinetics . . . 36
2.6.4. Equilibrium . . . 37
2.7. Acids, bases, and salts . . . 37
2.8. Oxidation, reduction, and electrochemistry . . . 38
2.8.1. Redox . . . 38
2.8.2. Electrochemical cells . . . 38
2.8.3. Batteries . . . 39
2.9. Organic chemistry in one sitting . . . 40
2.9.1. Functional groups . . . 40
2.9.2. Isomers and chirality . . . 41
2.9.3. Polymers, briefly . . . 41
2.10. Chemical safety . . . 41

2.10.1. Hazard communication . . . 41
2.10.2. Controls, in order of preference . . . 42
2.10.3. The workshop's own chemistry . . . 42
2.11. Engineering bridge: chemical arithmetic . . . 43
2.11.1. Unit discipline with moles . . . 43
2.11.2. Ideal gas arithmetic for cylinders . . . 44
2.11.3. Accelerated ageing and the Arrhenius rule . . . 44
2.11.4. Faraday's law: how much metal moves . . . 44
2.11.5. Battery energy . . . 44
2.12. The field in practice: what chemists do, and the tools they use . . . 45
2.12.1. The branches that touch health . . . 45
2.12.2. The analytical workflow . . . 46
2.12.3. Measuring and preparing . . . 46
2.12.4. Separating . . . 47
2.12.5. Identifying and quantifying . . . 47
2.12.6. Handling, containment, automation, and data . . . 48
2.12.7. Software . . . 49
2.12.8. From molecule to medicine . . . 49
2.13. Devices, software, and training technology . . . 49
2.13.1. Clinical devices built on general chemistry . . . 49
2.13.2. Software . . . 50
2.13.3. Training and simulation . . . 51
2.13.4. Where a maker contributes . . . 51
2.14. Worked example: "can you make us a dispenser for the chlorine tablets?" . . . 52
2.15. Workshop exercise . . . 53
2.16. Questions . . . 54
2.17. Answers and reasoning . . . 54
2.18. Interactive tools . . . 55
2.19. Video companion and reading . . . 56
2.20. Exit criteria . . . 56
2.21. What this chapter does not establish . . . 57
3. Chemistry and molecular biology for the maker . . . 58
3.1. Why this chapter exists . . . 58
3.2. Atoms, bonds, and why water behaves as it does . . . 58
3.2.1. The bonds you will meet . . . 58
3.2.2. Water . . . 59
3.3. pH, buffers, and why the range is so narrow . . . 60
3.3.1. Why pH matters to a device . . . 60
3.4. The four classes of biological molecule . . . 61
3.4.1. Proteins, in more detail . . . 61
3.4.2. Enzymes . . . 62

3.4.3. Lipids and membranes . . . 63
3.4.4. Carbohydrates . . . 63
3.4.5. Nucleic acids, briefly . . . 63
3.5. Bioenergetics: where the energy comes from . . . 63
3.6. Cell signalling and pharmacology . . . 64
3.7. Engineering bridge: the chemistry you meet as a maker . . . 65
3.7.1. Polymer chemistry, practically . . . 65
3.7.2. Leachables, extractables, and biocompatibility . . . 66
3.7.3. Metals and corrosion . . . 67
3.7.4. Textiles and glass . . . 68
3.7.5. Adhesion and surface chemistry . . . 69
3.8. Devices, software, and training technology . . . 69
3.8.1. Clinical devices built on this chemistry . . . 69
3.8.2. Software . . . 70
3.8.3. Training and simulation . . . 71
3.9. Worked example: why the handle got sticky . . . 72
3.10. Workshop exercise . . . 73
3.11. Questions . . . 73
3.12. Answers and reasoning . . . 74
3.13. Interactive tools . . . 75
3.14. Video companion and reading . . . 75
3.15. Exit criteria . . . 75
3.16. What this chapter does not establish . . . 76
4. Foundations of biology: life, cells, and how biologists work . . . 77
4.1. Why this chapter exists . . . 77
4.2. What counts as life . . . 78
4.3. The organising ideas . . . 78
4.3.1. Cell theory . . . 78
4.3.2. Evolution by natural selection . . . 78
4.4. The tree of life . . . 79
4.4.1. Prokaryotic versus eukaryotic cells . . . 79
4.5. The eukaryotic cell, part by part . . . 80
4.5.1. Scale . . . 82
4.6. Membranes and transport, briefly . . . 82
4.7. From DNA to protein: the central dogma in working detail . . . 82
4.7.1. DNA and its replication . . . 83
4.7.2. Transcription and processing . . . 83
4.7.3. Translation and the genetic code . . . 83
4.7.4. Gene regulation and epigenetics . . . 84
4.8. The cell cycle, division, and death . . . 84
4.9. Specialisation and stem cells . . . 86

4.10. Microbes at the cell level . . . 86
4.10.1. Bacteria . . . 86
4.10.2. Viruses . . . 87
4.11. The immune system at the cell level . . . 87
4.12. Metabolism, development, and ageing, in one paragraph each . . . 88
4.13. Engineering bridge: cells by the numbers . . . 88
4.13.1. Why cells are small: surface and volume . . . 88
4.13.2. Why tissues need blood vessels: diffusion time . . . 89
4.13.3. Exponential growth: bacteria, cell culture, and PCR . . . 89
4.13.4. The genome as information . . . 89
4.13.5. Seeing small things: resolution and signal . . . 90
4.13.6. Spinning things down: relative centrifugal force . . . 90
4.14. The field in practice: what biologists do, and the tools they use . . . 90
4.14.1. A laboratory workflow . . . 92
4.14.2. Tools for seeing . . . 92
4.14.3. Tools for growing . . . 93
4.14.4. Tools for separating and preparing . . . 93
4.14.5. Tools for measuring and counting . . . 94
4.14.6. Tools for reading and editing genes . . . 95
4.14.7. Tools for storing, automating, and keeping records . . . 95
4.14.8. Software . . . 96
4.14.9. Biosafety: containment levels . . . 96
4.15. Devices, software, and training technology . . . 97
4.15.1. Clinical and laboratory devices, with maker entry points . . . 97
4.15.2. Software . . . 99
4.15.3. Training and simulation . . . 99
4.16. Worked example: twelve gel imagers for a teaching lab . . . 100
4.17. Workshop exercise . . . 101
4.18. Questions . . . 102
4.19. Answers and reasoning . . . 103
4.20. Interactive tools . . . 104
4.21. Video companion and reading . . . 104
4.22. Exit criteria . . . 104
4.23. What this chapter does not establish . . . 105
5. Tissues, skin, and contact with an object . . . 106
5.1. Why this chapter exists . . . 106
5.2. The four tissue families . . . 106
5.2.1. Epithelium . . . 107
5.2.2. Connective tissue: the engineer’s family . . . 108
5.2.3. Muscle and nervous tissue, briefly . . . 109
5.3. Skin is an organ . . . 110

5.3.1. Epidermis . . . 111
5.3.2. Dermis . . . 112
5.3.3. Hypodermis (subcutaneous tissue) . . . 112
5.3.4. What skin does . . . 112
5.4. Pressure is only the beginning . . . 112
5.4.1. Pressure, friction, shear, and microclimate . . . 113
5.4.2. Soft tissue is viscoelastic, non-linear, and anisotropic . . . 114
5.5. Inflammation and tissue repair . . . 115
5.6. Engineering bridge: the interface as a design problem . . . 116
5.6.1. Intact skin, compromised skin, and breached skin . . . 116
5.6.2. A contact design checklist . . . 116
5.6.3. Design moves, and what each one costs . . . 117
5.7. Devices, software, and training technology . . . 118
5.7.1. Clinical devices . . . 118
5.7.2. Software . . . 119
5.7.3. Training and simulation . . . 120
5.8. Worked example: the awkward handle . . . 121
5.9. Workshop exercise . . . 122
5.10. Questions . . . 122
5.11. Answers and reasoning . . . 123
5.12. Interactive tools . . . 124
5.13. Video companion and reading . . . 124
5.14. Exit criteria . . . 124
5.15. What this chapter does not establish . . . 125
6. Bones, joints, and movement . . . 126
6.1. Why this chapter exists . . . 126
6.2. Bone as a material and as a tissue . . . 126
6.2.1. Structure . . . 126
6.2.2. Mechanical properties, for intuition . . . 128
6.3. The skeleton you actually need . . . 129
6.3.1. The upper limb in detail . . . 129
6.4. Joints . . . 130
6.4.1. Classification . . . 130
6.4.2. Anatomy of a synovial joint . . . 131
6.4.3. The six synovial joint types . . . 133
6.5. Describing motion precisely . . . 135
6.5.1. Active versus passive movement . . . 136
6.5.2. Typical ranges, for orientation only . . . 137
6.6. Muscle: how force is actually produced . . . 137
6.6.1. Structure . . . 138
6.6.2. Types of contraction . . . 138

6.6.3. Length–tension and force–velocity . . . 139
6.6.4. Agonist, antagonist, synergist, stabiliser . . . 139
6.7. Engineering bridge: levers, moment arms, and statics . . . 140
6.7.1. Torque . . . 140
6.7.2. Worked statics: the idealised elbow lever . . . 140
6.7.3. Including the limb's own weight . . . 142
6.7.4. Lever classes . . . 142
6.7.5. Grip, pinch, and the problem with strength figures . . . 142
6.7.6. Anthropometry and percentiles . . . 143
6.8. Devices, software, and training technology . . . 144
6.8.1. Clinical devices . . . 144
6.8.2. Software . . . 145
6.8.3. Training and simulation . . . 146
6.9. Worked example: moving the cord attachment . . . 147
6.10. Workshop exercise . . . 147
6.11. Questions . . . 148
6.12. Answers and reasoning . . . 148
6.13. Interactive tools . . . 149
6.14. Video companion and reading . . . 150
6.15. Exit criteria . . . 150
6.16. What this chapter does not establish . . . 151
7. Nerves, sensation, and accessible design . . . 152
7.1. Why this chapter exists . . . 152
7.2. The nervous system, organised . . . 152
7.3. How a neuron works . . . 153
7.3.1. Structure . . . 153
7.3.2. The resting membrane potential . . . 154
7.3.3. The action potential . . . 154
7.3.4. Synapses and the neuromuscular junction . . . 155
7.4. Sensation . . . 155
7.4.1. The somatosensory receptors . . . 155
7.4.2. Proprioception . . . 157
7.4.3. Pain, and why it is not a sensor reading . . . 157
7.4.4. Reflexes . . . 157
7.5. Engineering bridge: biosignals . . . 158
7.5.1. Psychophysics: the numbers behind "noticeable" . . . 159
7.6. Decomposing a task: perception, cognition, action, feedback . . . 159
7.7. Designing for variation, not for a diagnosis . . . 161
7.7.1. Accessible feedback: redundancy, not substitution . . . 161
7.8. Devices, software, and training technology . . . 162
7.8.1. Clinical devices . . . 162

7.8.2. Software . . . 163
7.8.3. Training and simulation . . . 164
7.9. Worked example: two identical cord ends . . . 166
7.10. Workshop exercise . . . 166
7.11. Questions . . . 167
7.12. Answers and reasoning . . . 167
7.13. Interactive tools . . . 169
7.14. Video companion and reading . . . 169
7.15. Exit criteria . . . 169
7.16. What this chapter does not establish . . . 170
8. The head and neck: eye, ear, nose, throat, and mouth . . . 171
8.1. Why this chapter exists . . . 171
8.2. The special senses, and the nerves that carry them . . . 171
8.3. The eye . . . 173
8.3.1. Structure . . . 173
8.3.2. The light path . . . 174
8.3.3. Refractive error . . . 174
8.3.4. The retina . . . 174
8.3.5. Where vision fails, and why the pattern matters . . . 175
8.3.6. How the eye protects itself . . . 175
8.4. The ear . . . 176
8.4.1. Three parts, three media . . . 176
8.4.2. The sound path . . . 176
8.4.3. The cochlea . . . 177
8.4.4. Two kinds of hearing loss . . . 177
8.4.5. Balance . . . 177
8.5. The nose and sinuses . . . 178
8.6. The throat: pharynx and larynx . . . 178
8.6.1. The crossing . . . 178
8.6.2. The larynx and the voice . . . 179
8.6.3. Swallowing . . . 180
8.6.4. When the upper airway is bypassed . . . 181
8.7. The mouth . . . 181
8.8. The face and neck as a place devices sit . . . 182
8.9. Engineering bridge: optics, acoustics, and airflow . . . 182
8.9.1. Dioptres . . . 182
8.9.2. Visual angle and legibility . . . 183
8.9.3. Decibels . . . 183
8.9.4. Two pieces of ear engineering . . . 184
8.9.5. Airflow in the upper airway . . . 185
8.10. Devices, software, and training technology . . . 185

8.10.1. Clinical devices . . . 185
8.10.2. Software . . . 187
8.10.3. Training and simulation . . . 187
8.11. Worked example: the label at the back of the room . . . 189
8.12. Workshop exercise . . . 190
8.13. Questions . . . 190
8.14. Answers and reasoning . . . 191
8.15. Interactive tools . . . 192
8.16. Video companion and reading . . . 192
8.17. Exit criteria . . . 193
8.18. What this chapter does not establish . . . 193
9. Circulation, breathing, and the wider body . . . 194
9.1. Why this chapter exists . . . 194
9.2. The heart and the two circuits . . . 194
9.2.1. Chambers and valves . . . 194
9.2.2. Vessel types . . . 195
9.2.3. The cardiac cycle . . . 196
9.2.4. The conduction system and the ECG . . . 197
9.3. Engineering bridge: haemodynamics . . . 198
9.3.1. The Ohm’s law analogue . . . 198
9.3.2. Poiseuille’s law, and why arterioles matter so much . . . 198
9.3.3. Pressure, and what a blood pressure reading is . . . 199
9.3.4. Blood itself . . . 199
9.4. Breathing . . . 200
9.4.1. Anatomy . . . 200
9.4.2. Mechanics of ventilation . . . 201
9.4.3. Lung volumes . . . 203
9.4.4. Gas exchange . . . 203
9.5. What the vital signs actually measure . . . 204
9.5.1. How a pulse oximeter works, and where it fails . . . 205
9.5.2. The four layers: phenomenon, sensor, display, interpretation . . . 206
9.6. The other systems: a map, not a textbook . . . 207
9.6.1. Renal and urinary . . . 207
9.6.2. Digestive . . . 208
9.6.3. Endocrine . . . 208
9.6.4. Immune and lymphatic . . . 208
9.6.5. Reproductive . . . 209
9.7. Devices, software, and training technology . . . 209
9.7.1. Clinical devices . . . 209
9.7.2. Software . . . 211
9.7.3. Training and simulation . . . 212

9.8. Worked example: the classroom mock monitor . . . 213
9.9. Workshop exercise . . . 214
9.10. Questions . . . 214
9.11. Answers and reasoning . . . 215
9.12. Interactive tools . . . 216
9.13. Video companion and reading . . . 216
9.14. Exit criteria . . . 217
9.15. What this chapter does not establish . . . 217
10. Kidneys, gut, and hormones: the internal regulators . . . 218
10.1. Why this chapter exists . . . 218
10.2. The kidneys and urinary tract . . . 218
10.2.1. Anatomy . . . 218
10.2.2. The nephron: a filter followed by a refinery . . . 219
10.2.3. What the kidneys regulate . . . 220
10.2.4. Measuring kidney function . . . 221
10.2.5. The bladder and continence . . . 222
10.2.6. Catheters . . . 222
10.3. The digestive system . . . 223
10.3.1. The tube and its glands . . . 223
10.3.2. Stomach . . . 224
10.3.3. Small intestine . . . 224
10.3.4. Large intestine . . . 224
10.3.5. Liver, gallbladder, and pancreas . . . 225
10.3.6. When the tube cannot be used normally . . . 225
10.4. The endocrine system . . . 226
10.4.1. Hormones as a broadcast channel . . . 226
10.4.2. The hypothalamus and pituitary: the control hierarchy . . . 228
10.4.3. The main glands . . . 228
10.4.4. Calcium . . . 229
10.4.5. Glucose: the most engineered loop in the body . . . 229
10.5. Engineering bridge: filters, transit, loops, and units . . . 230
10.5.1. Clearance: the kidney as a mass-transfer device . . . 230
10.5.2. Dialysis . . . 230
10.5.3. Transit, surface, and residence time . . . 231
10.5.4. Hormonal loops as control systems . . . 231
10.5.5. Units: where this chapter’s numbers bite . . . 232
10.6. Devices, software, and training technology . . . 233
10.6.1. Clinical devices . . . 233
10.6.2. Software . . . 234
10.6.3. Training and simulation . . . 235
10.7. Worked example: the drainage-bag hanger . . . 235

10.8. Workshop exercise . . . 236
10.9. Questions . . . 237
10.10. Answers and reasoning . . . 237
10.11. Interactive tools . . . 238
10.12. Video companion and reading . . . 238
10.13. Exit criteria . . . 239
10.14. What this chapter does not establish . . . 239
11. How disease works: from mechanism to symptom . . . 241
11.1. Why this chapter exists . . . 241
11.2. The language of disease . . . 241
11.3. Cell injury and cell death . . . 242
11.3.1. Why oxygen supply is the commonest failure . . . 242
11.3.2. Two ways to die . . . 243
11.4. Inflammation . . . 243
11.5. Infection, sepsis, and shock . . . 244
11.5.1. Infection and the host response . . . 244
11.5.2. Sepsis . . . 244
11.5.3. Shock . . . 244
11.5.4. When negative feedback turns positive . . . 245
11.6. Thrombosis and embolism . . . 245
11.7. Blocked pipes: atherosclerosis, heart attack, and stroke . . . 246
11.7.1. Atherosclerosis . . . 246
11.7.2. Heart attack . . . 247
11.7.3. Stroke . . . 247
11.8. Failing pumps and airways . . . 248
11.8.1. Heart failure . . . 248
11.8.2. Asthma and COPD . . . 248
11.9. Diabetes and its complications . . . 249
11.10. Cancer . . . 250
11.10.1. What makes a cell cancerous . . . 250
11.10.2. Staging and treatment . . . 250
11.11. Immunity turned inward, and wear . . . 251
11.11.1. Autoimmunity . . . 251
11.11.2. Degenerative disease . . . 251
11.12. Genes, inheritance, and genetic disease . . . 252
11.12.1. From code to protein, briefly . . . 252
11.12.2. Patterns of inheritance . . . 252
11.12.3. Genetic testing and pharmacogenomics . . . 252
11.13. Engineering bridge: failure, growth, and risk . . . 253
11.13.1. Disease as failure mode — and where the analogy breaks . . . 253
11.13.2. Exponential growth . . . 253

11.13.3. Relative and absolute risk . . . 254
11.14. Devices, software, and training technology . . . 254
11.14.1. Clinical devices . . . 254
11.14.2. Software . . . 255
11.14.3. Training and simulation . . . 256
11.15. Worked example: a foot-inspection aid . . . 257
11.16. Workshop exercise . . . 258
11.17. Questions . . . 259
11.18. Answers and reasoning . . . 260
11.19. Interactive tools . . . 261
11.20. Video companion and reading . . . 261
11.21. Exit criteria . . . 261
11.22. What this chapter does not establish . . . 262
12. Across the lifespan, and the mind . . . 263
12.1. Why this chapter exists . . . 263
12.2. Children are not small adults . . . 263
12.2.1. Size and proportion . . . 263
12.2.2. The airway . . . 264
12.2.3. Physiology by age . . . 264
12.2.4. Development and behaviour . . . 265
12.2.5. Consent, assent, and the family . . . 265
12.2.6. Why paediatric devices are a field of their own . . . 266
12.3. Pregnancy, briefly . . . 266
12.4. Older adults . . . 266
12.4.1. Ageing is gradual, uneven, and individual . . . 266
12.4.2. Frailty . . . 267
12.4.3. Falls . . . 268
12.4.4. Multimorbidity and polypharmacy . . . 268
12.4.5. Designing for older users . . . 269
12.5. The mind . . . 269
12.5.1. Cognition: the resources every interface spends . . . 269
12.5.2. Dementia and delirium . . . 270
12.5.3. Mental health . . . 270
12.5.4. Learning disability, autism, and neurodiversity . . . 271
12.5.5. Stress and performance . . . 271
12.5.6. Ligature risk: a design requirement from mental-health settings . . . 272
12.6. Engineering bridge: scaling, distributions, and limits . . . 272
12.6.1. Scaling with body size: weight versus surface area . . . 272
12.6.2. Growth charts are distributions . . . 272
12.6.3. Designing across the whole range . . . 273
12.6.4. Choice time and working memory . . . 273

12.6.5. The physics of a fall . . . 274
12.7. Devices, software, and training technology . . . 274
12.7.1. Clinical devices . . . 274
12.7.2. Software . . . 275
12.7.3. Training and simulation . . . 276
12.8. Worked example: the elbow model goes to two new places . . . 277
12.9. Workshop exercise . . . 278
12.10. Questions . . . 279
12.11. Answers and reasoning . . . 279
12.12. Interactive tools . . . 280
12.13. Video companion and reading . . . 280
12.14. Exit criteria . . . 281
12.15. What this chapter does not establish . . . 281
Part II: Putting an object into a clinical world . . . 283
13. Infection control and choosing materials . . . 284
13.1. Why this chapter exists . . . 284
13.2. Microbiology, in the amount you need . . . 284
13.2.1. The chain of infection . . . 285
13.2.2. Routes of transmission . . . 285
13.3. Cleaning, disinfection, and sterilisation are three different things . . . 286
13.3.1. Spaulding classification . . . 286
13.3.2. Steam sterilisation, since it is the common one . . . 287
13.4. The finished object is the design . . . 287
13.4.1. What actually varies: additive manufacturing . . . 288
13.4.2. What actually varies: other methods . . . 289
13.4.3. Design for cleanability . . . 290
13.5. Engineering bridge: comparing candidate materials . . . 291
13.6. Devices, software, and training technology . . . 294
13.6.1. Clinical devices and equipment . . . 294
13.6.2. Software . . . 295
13.6.3. Training and simulation . . . 296
13.7. Worked example: the pivot cover . . . 297
13.8. Workshop exercise . . . 298
13.9. Questions . . . 299
13.10. Answers and reasoning . . . 299
13.11. Interactive tools . . . 301
13.12. Video companion and reading . . . 301
13.13. Exit criteria . . . 301
13.14. What this chapter does not establish . . . 302
14. Making it: choosing the technology and the method . . . 303
14.1. Why this chapter exists . . . 303

14.2. Start from the problem, not the tool . . . 303
14.2.1. The questions that choose the method . . . 303
14.2.2. The options before fabrication . . . 304
14.3. The technology families . . . 305
14.3.1. Software . . . 306
14.3.2. Electronics and embedded systems . . . 306
14.3.3. Additive manufacturing . . . 307
14.3.4. CNC machining and other subtractive methods . . . 308
14.3.5. Moulding, casting, and forming . . . 309
14.3.6. Metalwork . . . 309
14.3.7. Glasswork . . . 310
14.3.8. Textiles and soft goods . . . 311
14.3.9. Wood and other workshop materials . . . 312
14.3.10. Combining methods . . . 312
14.4. From prototype to product . . . 313
14.4.1. Process validation . . . 313
14.4.2. Supplier controls and material certificates . . . 313
14.4.3. Why the hospital may prefer to buy . . . 314
14.4.4. What a maker hands over . . . 314
14.5. Engineering bridge: costs, tolerances, surfaces, and energy . . . 314
14.5.1. The cost model . . . 314
14.5.2. Tolerance stack-up . . . 315
14.5.3. Surface roughness and cleanability . . . 315
14.5.4. Battery energy . . . 316
14.5.5. Process capability, in brief . . . 316
14.6. Devices, software, and training technology . . . 316
14.6.1. Where making happens in healthcare . . . 316
14.6.2. Software and design tools . . . 317
14.6.3. Training and simulation made by each method . . . 318
14.7. Worked example: thirty elbow models for a teaching programme . . . 318
14.8. Workshop exercise . . . 320
14.9. Questions . . . 320
14.10. Answers and reasoning . . . 321
14.11. Interactive tools . . . 322
14.12. Video companion and reading . . . 322
14.13. Exit criteria . . . 323
14.14. What this chapter does not establish . . . 323
15. Measurement, imaging, and the diagnostic laboratory . . . 324
15.1. Why this chapter exists . . . 324
15.2. Medical imaging . . . 324
15.2.1. The modalities . . . 324

15.2.2. Image quality, and the trade you cannot escape . . . 325
15.2.3. DICOM, and the path from scan to physical model . . . 326
15.3. Laboratory medicine . . . 326
15.3.1. The measurement principles worth knowing . . . 327
15.3.2. Point-of-care testing . . . 328
15.4. Engineering bridge: what makes a measurement trustworthy . . . 328
15.4.1. Diagnostic performance: the other sensitivity and specificity . . . 329
15.4.2. ROC curves and the threshold decision . . . 330
15.5. Devices, software, and training technology . . . 330
15.5.1. Clinical devices . . . 330
15.5.2. Software . . . 331
15.5.3. Training and simulation . . . 332
15.6. Worked example: the request for a model from a scan . . . 333
15.7. Workshop exercise . . . 334
15.8. Questions . . . 335
15.9. Answers and reasoning . . . 335
15.10. Interactive tools . . . 336
15.11. Video companion and reading . . . 337
15.12. Exit criteria . . . 337
15.13. What this chapter does not establish . . . 337
16. Diagnostic devices: how they work and what they are for . . . 339
16.1. Why this chapter exists . . . 339
16.2. The map: which question, which device . . . 339
16.3. Radiation, ionising and not . . . 340
16.4. Projection radiography: the plain X-ray . . . 341
16.4.1. How X-rays are made . . . 341
16.4.2. How the body makes the image . . . 341
16.4.3. The radiographic family . . . 343
16.5. Computed tomography (CT) . . . 343
16.5.1. Many shadows, solved backwards . . . 343
16.5.2. Hounsfield units . . . 344
16.5.3. What CT is for, and what it costs . . . 345
16.6. Magnetic resonance imaging (MRI) . . . 346
16.6.1. The physics in five steps . . . 346
16.6.2. Choosing the contrast: weighting and sequences . . . 347
16.6.3. What MRI is for, and what it costs . . . 348
16.6.4. MR safety: the field is always on . . . 349
16.7. Ultrasound . . . 349
16.7.1. Pulse and echo . . . 350
16.7.2. Frequency is a trade between detail and depth . . . 350
16.7.3. Doppler: measuring movement . . . 351

16.7.4. What ultrasound is for, and what it costs . . . 352
16.8. Nuclear medicine: SPECT and PET . . . 352
16.8.1. Gamma camera and SPECT . . . 352
16.8.2. PET . . . 353
16.8.3. The standardised uptake value . . . 354
16.8.4. What nuclear medicine is for, and what it costs . . . 355
16.9. Endoscopy and direct visualisation . . . 355
16.10. Physiological measurement . . . 356
16.10.1. The heart . . . 356
16.10.2. Blood pressure . . . 357
16.10.3. The lungs . . . 357
16.10.4. Sleep . . . 358
16.10.5. Nerves, muscle, and brain . . . 358
16.10.6. Metabolism and temperature . . . 358
16.11. Tissue diagnosis: the biopsy . . . 358
16.12. Choosing a test . . . 359
16.12.1. The modalities side by side . . . 360
16.13. Engineering bridge: the mathematics underneath . . . 360
16.14. Devices, software, and training technology . . . 361
16.14.1. Clinical devices and the maker’s entry points . . . 361
16.14.2. Software . . . 362
16.14.3. Training and simulation . . . 363
16.15. Worked example: the positioning model that went to the MRI room . . . 364
16.16. Workshop exercise . . . 365
16.17. Questions . . . 366
16.18. Answers and reasoning . . . 366
16.19. Interactive tools . . . 367
16.20. Video companion and reading . . . 368
16.21. Exit criteria . . . 368
16.22. What this chapter does not establish . . . 369
17. Treatment and life-support devices: how they work and what they are for . . . 370
17.1. Why this chapter exists . . . 370
17.2. Life support, and why it is different . . . 370
17.3. Mechanical ventilation . . . 371
17.3.1. Breathing in reverse . . . 371
17.3.2. The circuit . . . 371
17.3.3. Modes, in one distinction . . . 372
17.3.4. Without a tube: non-invasive support and oxygen . . . 373
17.3.5. Who operates it . . . 373
17.4. Infusion pumps and syringe drivers . . . 373
17.4.1. Two mechanisms . . . 373

17.4.2. How they fail . . . 374
17.4.3. Smart pumps and their limits . . . 374
17.5. Defibrillators, pacemakers, and implantable defibrillators . . . 375
17.5.1. Defibrillation . . . 375
17.5.2. The AED: usability as a life-saving feature . . . 375
17.5.3. Pacemakers and ICDs . . . 376
17.6. Dialysis and extracorporeal circuits . . . 376
17.6.1. Haemodialysis . . . 376
17.6.2. ECMO and cardiopulmonary bypass . . . 377
17.7. Energy in the operating theatre . . . 377
17.7.1. Electrosurgery . . . 377
17.7.2. Lasers . . . 378
17.7.3. Surgical fire . . . 378
17.7.4. Minimally invasive and robotic surgery . . . 379
17.8. Radiotherapy . . . 379
17.8.1. Killing cells with dose . . . 379
17.8.2. The linear accelerator . . . 380
17.8.3. Immobilisation: the maker-shaped object in the room . . . 380
17.9. Implants . . . 381
17.10. Drug delivery . . . 383
17.10.1. Pharmacokinetics in four words . . . 383
17.10.2. Routes and the devices that serve them . . . 384
17.11. Rehabilitation and assistive therapy devices . . . 385
17.12. Engineering bridge: the numbers inside the machines . . . 386
17.12.1. The lung as an RC circuit . . . 386
17.12.2. Defibrillator energy . . . 386
17.12.3. Infusion arithmetic and the unit hazard . . . 386
17.12.4. Half-life arithmetic . . . 387
17.12.5. Dialysis clearance . . . 387
17.12.6. Dose fall-off . . . 387
17.13. Devices, software, and training technology . . . 387
17.13.1. Clinical devices and the maker’s entry points . . . 387
17.13.2. Software . . . 389
17.13.3. Training and simulation . . . 389
17.14. Worked example: the request for an emergency ventilator . . . 391
17.15. Workshop exercise . . . 392
17.16. Questions . . . 393
17.17. Answers and reasoning . . . 393
17.18. Interactive tools . . . 395
17.19. Video companion and reading . . . 395
17.20. Exit criteria . . . 395

17.21. What this chapter does not establish . . . 396
18. Discovering a clinical need . . . 397
18.1. Why this chapter exists . . . 397
18.2. Who is who . . . 397
18.3. Start with the task as it is performed . . . 399
18.3.1. Map the whole task, not the difficult moment . . . 399
18.3.2. Interviewing . . . 400
18.4. Write a need without choosing the answer . . . 401
18.4.1. Then gather the constraints . . . 401
18.5. Separate evidence from interpretation . . . 402
18.6. A sample conversation . . . 403
18.7. Engineering bridge: filtering needs . . . 403
18.7.1. Intellectual property, briefly . . . 404
18.8. Worked example: from request to brief . . . 404
18.9. Workshop exercise . . . 406
18.10. Questions . . . 406
18.11. Answers and reasoning . . . 407
18.12. Interactive tools . . . 408
18.13. Video companion and reading . . . 408
18.14. Exit criteria . . . 408
18.15. What this chapter does not establish . . . 408
19. Requirements, hazards, and useful tests . . . 410
19.1. Why this chapter exists . . . 410
19.2. Intended use, users, and environment . . . 410
19.3. Turning adjectives into requirements . . . 411
19.4. Hazards, hazardous situations, and harm . . . 412
19.4.1. The risk control hierarchy . . . 412
19.4.2. Every control creates new questions . . . 413
19.4.3. Non-physical harms count . . . 413
19.4.4. A lightweight hazard table . . . 413
19.5. Designing a test around one question . . . 415
19.5.1. Worked example, with an invented result . . . 415
19.6. Verification and validation . . . 416
19.7. Traceability: linking object to evidence . . . 416
19.7.1. Versioning a physical object . . . 417
19.8. Engineering bridge: how the professional process is shaped . . . 417
19.8.1. Software-specific practices . . . 418
19.8.2. Hardware-, material-, and process-specific practices . . . 419
19.9. Worked example: three requirements with their limits . . . 420
19.10. Workshop exercise . . . 420
19.11. Questions . . . 421

19.12. Answers and reasoning . . . 421
19.13. Interactive tools . . . 422
19.14. Video companion and reading . . . 422
19.15. Exit criteria . . . 423
19.16. What this chapter does not establish . . . 423
20. Health data, medical software, and security . . . 424
20.1. Why this chapter exists . . . 424
20.2. When is software a medical device? . . . 425
20.2.1. Intended use decides . . . 425
20.2.2. SaMD and SiMD . . . 425
20.2.3. The IMDRF risk framework . . . 425
20.2.4. The shape of the main regimes . . . 426
20.3. The software lifecycle: IEC 62304 . . . 427
20.3.1. Safety classes . . . 427
20.3.2. SOUP: your dependencies, formally . . . 428
20.3.3. Configuration management and problem resolution . . . 428
20.3.4. Agile is allowed . . . 428
20.4. What goes wrong: lessons written in harm . . . 428
20.4.1. Therac-25 . . . 428
20.4.2. Infusion pumps and the arithmetic of a keypad . . . 429
20.4.3. Units and silent conversions . . . 430
20.4.4. Alarms . . . 430
20.5. Health data . . . 430
20.5.1. What counts . . . 430
20.5.2. The identifiability spectrum . . . 430
20.5.3. The legal frameworks in outline . . . 431
20.5.4. Governance in practice . . . 432
20.6. Interoperability: making systems understand each other . . . 432
20.6.1. Codes and units carry the meaning . . . 433
20.7. Security . . . 434
20.7.1. Why hospitals are targets . . . 434
20.7.2. What regulators now expect of device makers . . . 434
20.7.3. Secure development, applied . . . 435
20.8. AI and machine learning in medicine . . . 435
20.9. Engineering bridge: safety-critical habits in code . . . 436
20.9.1. Units in the type system . . . 436
20.9.2. Input validation that respects the user . . . 436
20.9.3. Audit trails . . . 437
20.9.4. k-anonymity arithmetic . . . 437
20.9.5. Threat modelling with STRIDE . . . 438
20.9.6. Fault trees for software . . . 438

20.10. Devices, software, and training technology . . . 439
20.10.1. Clinical systems and software-intensive devices . . . 439
20.10.2. Tools you can use today . . . 440
20.10.3. Training and simulation . . . 440
20.11. Worked example: the companion app that grew . . . 441
20.12. Workshop exercise . . . 443
20.13. Questions . . . 443
20.14. Answers and reasoning . . . 444
20.15. Interactive tools . . . 445
20.16. Video companion and reading . . . 445
20.17. Exit criteria . . . 446
20.18. What this chapter does not establish . . . 446
21. Making a model that teaches . . . 448
21.1. Why this chapter exists . . . 448
21.2. Begin with what the learner must do . . . 448
21.2.1. Bloom’s taxonomy, briefly . . . 448
21.2.2. Miller’s pyramid: knowing is not doing . . . 449
21.3. Choose the realism that serves the objective . . . 449
21.3.1. Each kind of fidelity is made a different way . . . 450
21.3.2. What our model needs, and does not . . . 451
21.4. Design the explanation as carefully as the object . . . 452
21.4.1. Predict–observe–explain . . . 452
21.4.2. A worked teaching script . . . 452
21.5. Evaluate learning and maintenance separately . . . 453
21.6. Misleading cues: the real hazard of a teaching model . . . 453
21.7. Engineering bridge: materials and methods for simulation . . . 454
21.8. Worked example: adapting the model to one objective . . . 456
21.9. Workshop exercise . . . 457
21.10. Questions . . . 457
21.11. Answers and reasoning . . . 458
21.12. Interactive tools . . . 459
21.13. Video companion and reading . . . 459
21.14. Exit criteria . . . 459
21.15. What this chapter does not establish . . . 459
Part III: Doing the work . . . 461
22. Building and reviewing your first project . . . 462
22.1. Why this chapter exists . . . 462
22.2. Choose one route . . . 462
22.3. Week 27: confirm the brief . . . 463
22.3.1. The brief you will build against . . . 463
22.3.2. Confirm, in writing, before you cut anything . . . 463

22.3.3. Compare two concepts, and record why . . . 463
22.4. Weeks 28–29: build the first version cheaply . . . 464
22.4.1. Suggested materials . . . 465
22.4.2. Build sequence . . . 465
22.5. Troubleshooting: ask what changed . . . 466
22.6. Run a short bench review . . . 467
22.7. Review the learning claim . . . 467
22.8. Decide when better fabrication is justified . . . 468
22.8.1. Choose a method for each element, not one for the whole . . . 468
22.9. Week 30: review, revise, hand over . . . 469
22.9.1. Assemble the handover . . . 469
22.9.2. The status statement . . . 470
22.9.3. The five-minute explanation . . . 470
22.10. Workshop exercise . . . 471
22.11. Interactive tools . . . 471
22.12. Exit criteria . . . 471
22.13. What this chapter does not establish . . . 472
23. Next steps, specialisms, and finding a place to contribute . . . 473
23.1. Why this chapter exists . . . 473
23.2. Finding a place to contribute . . . 473
23.2.1. Where to look . . . 473
23.2.2. Organisations that exist for this . . . 474
23.2.3. What to actually say . . . 474
23.2.4. Realistic expectations . . . 475
23.3. Choosing a specialism after your first project . . . 476
23.4. Going further: formal study . . . 477
23.5. A compact resource shelf . . . 477
23.6. Exit criteria for this stage . . . 478
23.7. A closing note . . . 478
Appendices . . . 479
Appendix A: Study schedule and assessment method . . . 480
A.1. The overall shape . . . 480
A.2. Week by week . . . 480
A.3. Assessment method . . . 484
A.3.1. Per chapter . . . 484
A.3.2. Where another person is required . . . 484
A.3.3. Final assessment . . . 484
A.4. A repeatable weekly routine . . . 485
A.5. Spaced repetition . . . 485
A.6. Using one main textbook . . . 485
A.7. Cost and access . . . 486

Appendix B: Reusable worksheets . . . 487
B.1. Learning record . . . 487
B.2. One-page project brief . . . 487
B.3. Evidence separation sheet . . . 489
B.4. Task map . . . 489
B.5. Hazard analysis . . . 489
B.6. Requirement and acceptance criterion . . . 490
B.7. Test record . . . 490
B.8. Version record . . . 491
B.9. Learning objective and assessment . . . 492
B.10. Material comparison . . . 493
B.11. Handover pack checklist . . . 493
Appendix C: Physiological reference figures and units . . . 495
C.1. Cardiovascular . . . 495
C.2. Respiratory . . . 496
C.3. Gas partial pressures . . . 496
C.4. General chemistry . . . 496
C.4.1. Constants . . . 496
C.4.2. Useful molar masses . . . 497
C.4.3. Concentration conversions and solutions . . . 497
C.4.4. Standard electrode potentials (orientation) . . . 498
C.5. Chemistry and fluid . . . 498
C.6. Scale reference . . . 498
C.7. Cells and molecular biology . . . 499
C.7.1. Laboratory rules of thumb . . . 500
C.8. Biosignals . . . 500
C.9. Musculoskeletal . . . 501
C.9.1. Approximate joint ranges . . . 501
C.9.2. Approximate material stiffness . . . 501
C.9.3. Two-point discrimination . . . 502
C.10. Special senses . . . 502
C.10.1. Eye . . . 502
C.10.2. Legibility rule of thumb . . . 503
C.10.3. Ear and sound . . . 503
C.10.4. Upper airway, voice, and mouth . . . 503
C.11. Kidneys, gut, and hormones . . . 504
C.11.1. CKD stages by eGFR (mL/min/1.73 m²), after KDIGO . . . 505
C.12. Disease and risk arithmetic . . . 505
C.13. Across the lifespan . . . 506
C.13.1. Vital signs by age, orientation only . . . 506
C.13.2. Body size and composition . . . 506

C.14. Making and manufacture . . . 506
C.14.1. Process accuracy and surface, orientation only . . . 507
C.14.2. Materials and methods notes . . . 507
C.15. Diagnostic imaging . . . 507
C.15.1. Effective dose, orientation only . . . 508
C.15.2. Physical constants and rules . . . 508
C.16. Treatment and life support . . . 508
C.16.1. Rules . . . 509
C.17. Unit conversions . . . 510
Appendix D: Standards and regulatory orientation . . . 511
D.1. Why a maker should know the shape of this . . . 511
D.2. Core process standards . . . 511
D.3. Safety and biological evaluation . . . 512
D.4. Regional regulation . . . 515
D.4.1. Concepts worth knowing by name . . . 515
D.5. Chemical safety . . . 517
D.6. Laboratory biosafety and quality . . . 517
D.7. Health data protection . . . 517
D.8. Accessibility and human factors . . . 518
D.9. Additive manufacturing . . . 518
D.10. Infection prevention . . . 518
D.11. How to read a standard without buying it . . . 519
Appendix E: Final assessment and readiness review . . . 520
E.1. How to use this . . . 520
E.2. Written assessment . . . 520
E.3. Answers and reasoning . . . 521
E.4. Practical assessment . . . 522
E.5. Final readiness review . . . 522
E.6. Completion outcome . . . 523
Appendix F: Working glossary . . . 524
F.1. Anatomy and orientation . . . 524
F.2. General chemistry . . . 524
F.3. Cells, tissues, and regulation . . . 526
F.4. Cell and molecular biology . . . 527
F.5. Skin and contact . . . 529
F.6. Musculoskeletal . . . 530
F.7. Nervous system and sensation . . . 531
F.8. Head, neck, and special senses . . . 531
F.9. Circulation and respiration . . . 533
F.10. Kidneys, gut, and hormones . . . 533
F.11. Disease mechanisms and genetics . . . 534

F.12. Lifespan and the mind . . . 536
F.13. Infection control and materials . . . 537
F.14. Making and manufacture . . . 538
F.15. Design, evidence, and process . . . 539
F.16. Education and simulation . . . 540
F.17. Chemistry and molecular biology . . . 541
F.18. Measurement, imaging, and diagnostics . . . 542
F.19. Diagnostic devices . . . 543
F.20. Treatment and life-support devices . . . 545
F.21. Devices and health informatics . . . 547
F.22. Medical software, health data, and security . . . 548
Appendix G: The resource shelf . . . 551
G.1. Core, free, and sufficient . . . 551
G.2. Optional, when a project justifies it . . . 551
G.3. Where to find people . . . 552
G.4. What not to buy yet . . . 552
G.5. A note on currency . . . 552
Appendix H: Figure credits and licences . . . 554
H.1. Reused figures . . . 554
H.2. What the licences require . . . 555
Appendix I: Device and software index, by maker accessibility . . . 557
I.1. Tier 1 — start here . . . 557
I.2. Tier 2 — with a named clinical partner . . . 560
I.3. Tier 3 — professional territory . . . 563
I.4. Open tools and data worth knowing . . . 565
Bibliography . . . 567

Colophon

Human Biology for the Medical Maker — a self-study book in anatomy, physiology, and collaborative design for an experienced technologist — software, electronics, and a full fabrication workshop — who wants to apply modern technology to healthcare problems and is moving toward biomedical engineering work.

Version 1.0, 2026-10-04. Prepared for Chris.

This book grew out of two shorter documents: a study schedule and resource guide, and a taught handbook of ten lessons. Both are folded into the chapters that follow and substantially expanded.



This is an introductory self-study text. It has not been clinically reviewed, accredited, or endorsed by any of the organisations whose public material it cites.

- Numeric values are **typical adult textbook figures**, given to build intuition and to let you sanity-check an order of magnitude. They are not clinical reference ranges, not thresholds, and not design limits.
- Worked examples, dimensions, test counts, and results are **invented teaching examples**.
- Completing this book prepares you for a **conversation about supervised work**. It does not authorise clinical practice, patient contact, or independent approval of any patient-facing device.
- Cited regulatory material is largely from US agencies (FDA, CDC) because it is freely readable. The requirements that actually apply to you are whatever your host institution and jurisdiction say they are. Establish those with your host.

Preface: what this book is for

You already know how to make things work. You write software and firmware, design and build electronics, draw in CAD, and run a workshop that turns a drawing into an object — machining, turning, welding and fabricating metal, printing, casting and moulding, working glass, cutting and sewing fabric. You know which of those a problem actually needs, and that sometimes the answer is none of them. None of that is in question here.

This book is about putting that whole toolkit to work where it helps people's health: the right technology, chosen for the problem, built to the standard the setting demands, and handed over so someone else can live with it. It is not a catalogue of things to make. It is the knowledge you need to recognise a real problem, understand the body it involves, and work alongside the people who own it.

What this book adds is the other half of the job: enough human biology to follow a clinician's description accurately, enough physiology to understand why a body is not a mechanism with a fixed operating point, and enough of the design and evidence habits of the medical device world to be a **useful collaborator** rather than an enthusiastic liability.

That last distinction is the whole point. A hospital does not need another person who can make a bracket, write an app, or wire a sensor. It has suppliers for all of those. It needs someone who can sit in a conversation with an occupational therapist, hear "the handle is awkward", and understand that **awkward** might mean concentrated pressure, or an unachievable wrist position, or a grip force the person cannot sustain for ninety seconds, or that the thing is impossible to clean and so it never leaves the cupboard. Each of those is a different problem with a different solution — a reshaped grip, a different material, a lighter mechanism, a cleanable redesign, a sensor, a piece of software, or a conversation with the person who orders the equipment — and three of the four are not solved by adding foam.

What you will be able to do

By the end of the book you should be able to:

- Describe a location on the body unambiguously, and read a clinical description without silently mirroring left and right.
- Explain, at the level of a competent amateur, how tissue, bone, muscle, nerve, blood, and air actually behave — including the ways they refuse to behave like engineering materials.
- Explain how the eye, the ear, and the upper airway work well enough to size a label for a viewing distance, judge whether a sound cue will be heard, and know why anything that goes in a mouth is an airway question.
- Look at a proposed object and generate the right **questions**: what does it touch, for how long, who cleans it, who cannot feel it, what happens when it breaks, who reviews it.
- Recognise, for each area of medicine, **what devices and software already exist** — what they measure, what principle they run on, what trains people to use them, and which of them are plausible places for you to contribute.
- Follow a clinical conversation about the common diseases — heart attack, stroke, diabetes,

cancer, sepsis — well enough to see which requirement each one places on a device, and design for children, older adults, and people under cognitive load.

- Explain how ventilators, infusion pumps, defibrillators, dialysis, and radiotherapy work, and why they sit at the top of the risk scale.
- Recognise when software becomes a medical device, handle health data lawfully and securely, and speak the interoperability vocabulary (FHIR, DICOM, SNOMED CT) a hospital uses.
- Explain how X-ray, CT, MRI, ultrasound, and PET each turn a physical interaction into an image, what each is good for, what each costs the patient, and why nothing you build enters an MRI room without approval.
- Read a measurement critically: distinguish accuracy from precision, know why a 99%-accurate test can be mostly wrong, and know which layer of a display is a clinical act rather than an engineering one.
- Turn a vague complaint into a written need statement that a clinician will recognise as their own problem.
- Turn "it seems to work" into a requirement, an acceptance criterion agreed before the test, a recorded result including the failures, and an honest statement of what the result does **not** establish.
- Choose the right technology and fabrication method for a problem — software, electronics, machining, moulding, casting, metal, glass, textiles, additive manufacturing, or none — and explain the choice in terms of cleaning, loads, quantity, maintenance, and risk.
- Build and hand over a small documented prototype that someone else could rebuild and whose limitations are written down.
- Say clearly where your competence ends and whose judgement is needed next.

What this book deliberately does not do

It does not teach you to diagnose, to interpret a patient's measurements, to select a disinfectant, to certify a material, or to approve anything. Several chapters go out of their way to show you the shape of a professional process precisely so that you can recognise when you are standing at its edge.

It also does not try to be a substitute for a degree. A biomedical engineering degree is three or four years of mathematics, signals, mechanics, materials, and physiology. This book is perhaps a hundred and thirty to a hundred and ninety hours. What it buys you is the ability to be **genuinely useful early**, under supervision, while you decide whether to go further — and a vocabulary solid enough that further study lands on prepared ground.

The running project

Throughout Part I and Part II, a single running example appears: a **tabletop elbow teaching model**. Two strips, a pivot, a cord. It demonstrates bending, straightening, and the role of a tendon. It never attaches to a person.

It is deliberately humble. It could be made from card, plywood, acrylic, aluminium, turned brass,

cast silicone, or printed polymer, with or without a sensor and a microcontroller — and choosing between those, for a stated reason, is part of the exercise. It keeps the first project inside familiar workshop territory while the biology, the collaboration, and the evidence habits are the things actually being learned. By Chapter 22 you will have built it, tested it, had it reviewed, found out it teaches one thing wrong, fixed that, and written the handover.

That is the real skill. Making the thing was never the hard part.

The device and software sections

Every chapter in Parts I and II ends with a section called **Devices, software, and training technology**, in the same three-part shape each time: what exists clinically, what software exists, and what is used to **teach** that area of medicine.

They are there for two reasons. The first is orientation — you cannot have a useful conversation with a clinical engineer or a simulation educator without knowing roughly what is in the building. The second is honesty about where you fit: each section names the realistic maker entry point, and Appendix I re-sorts the whole collection by how reachable a project actually is for a supervised learner.

The pattern that emerges is consistent and worth stating early. **Patient-facing clinical devices are mostly not where a supervised newcomer starts.** The realistic places to apply technology early are the ones where the regulatory weight is lighter and the need is real: education and simulation, assistive and daily-living technology designed with the person who uses it, accessibility, workflow and logistics, test and measurement rigs for clinical engineering, and software on synthetic or public data. Across those, the useful technology might be an instrumented trainer, a sewn positioning aid, a machined fixture, a moulded cover, a data pipeline, or an argument for buying something that already exists. Chapter 14 is about making that choice well.

A note on tone

Where this book sounds cautious, it is because the caution is the content. The failure mode for a skilled maker entering healthcare is not incompetence at making. It is confident competence in the wrong frame: a beautifully machined part that concentrates load on fragile skin, an elegant interface that encodes a clinical judgement its author was not qualified to make, a convincing anatomical model that teaches a relationship that does not exist.

Every chapter therefore ends by naming what the chapter does **not** let you conclude. Read those sections as carefully as the rest.

How to use this book

The shape of a chapter

Every chapter in Parts I–III follows the same structure, so you can navigate by habit:

Section	Purpose
Why this chapter exists	The design failure this knowledge prevents. Read it first; it tells you what to look for.
Core material	The actual biology or method, in numbered subsections. Tables and diagrams carry a lot of the load — redraw them rather than only reading them.
Engineering bridge	Where this biology meets numbers, materials, mechanics, or signals. This is the part a biomedical engineering context will expect you to have.
Devices, software, and training technology	What exists clinically in this area, what software exists, what is used to teach it, and the realistic maker entry point. Parts I and II only.
Worked example	A fully reasoned case, usually on the running elbow-model project.
Workshop exercise	Something to make, draw, or measure. Do it before reading the answers.
Questions and **Answers and reasoning**	Self-test. Attempt it closed-book. The answers explain the reasoning, not just the conclusion.
Exit criteria	What "done" means for this chapter. Practical, not just recall.
What this chapter does not establish	The limits. The most important section on the page.

A repeatable weekly routine

Plan for roughly **4–6 hours a week across about 30 weeks**: 130 to 190 hours of core study and bench work. Arranging a clinical partnership often takes longer than the study does, which is why outreach starts in week 1, not week 21.

A week that works:

1. **~1 hour — watch and note.** The companion video for the chapter. Once at normal speed for the shape of it, once slower while drawing.
2. **~1–2 hours — read and draw.** The chapter here, plus the named OpenStax sections [1]. Draw every diagram yourself at least once from memory.
3. **~2 hours — build or apply.** The workshop exercise. Hands on something.
4. **~30 minutes — explain aloud without notes.** To a person, a camera, or an empty room. This is not optional padding; it is where you discover which parts you only recognise rather than know.

For knowledge checks, answer ten questions drawn from your own notes or flashcards without

looking anything up. Aim for at least eight correct, then **correct every error in writing**. Re-attempt the ones you missed a week later.



A quiz score on its own is not completion. Each chapter's exit criteria include practical work, and several require feedback from a clinician or educator. Where a criterion needs another person's review and you have not got that person yet, record it as **pending** — do not quietly mark it done, and do not invent the agreement.

Spaced repetition, concretely

The material in Part I is vocabulary-dense and will fade without deliberate recall. Use a simple schedule per chapter:

When	What	Form
Same day	Explain aloud	No notes, 5 minutes
+1 day	10-question check	Written, closed book
+1 week	Re-attempt only the items missed	Written
+1 month	Re-draw the chapter's main diagram from memory	Paper
On project need	Re-read the relevant section with a specific question	Targeted

Building the book

The source is AsciiDoc. Diagrams are generated at build time from PlantUML, Graphviz, and ditaa source embedded in the chapters, so the figures are editable text rather than images you cannot change.

```
make html     # single-file HTML with left-hand TOC
make pdf      # print-ready PDF
make epub     # EPUB3 for e-readers
make clean
```

Prerequisites are `asciidoctor`, `asciidoctor-diagram`, `asciidoctor-pdf`, `asciidoctor-epub3`, Graphviz (`dot`), and a JRE for PlantUML.

Conventions used

Convention	Meaning
Bold term on first use	Defined there and in the glossary (Appendix F).
`Monospace`	A literal value, filename, or code identifier.
[1]	A bibliography citation. Full list in Bibliography.
NOTE	Useful aside or clarification.
TIP	Something that saves you time at the bench.
WARNING	A place where makers commonly overreach.
CAUTION	A statement that is true but frequently misread as permission.



British and American spellings. Clinical and standards literature mixes them freely — *sterilisation*/*sterilization*, *haemoglobin*/*hemoglobin*, *oesophagus*/*esophagus*, *anaesthesia*/*anesthesia*. This book uses British forms in its own prose and keeps American forms inside quoted titles. Learn to read both; in a literature search you need both spellings or you will miss half the results.

Similarly, *physiotherapist* (UK, and much of the Commonwealth) and *physical therapist* (US) denote closely related professions; *occupational therapist* is used in both.

Part I: Reading the body

Chapter 1. Reading the body and understanding living systems

Weeks 1–2 · approximately 8–12 hours

1.1. Why this chapter exists

An educator points at an elbow and asks for a model that shows how it bends. You go away and build something that bends. It is beautiful. It is also wrong, because it teaches that the cord contracts, and because when you said "put the cable on the inside" you meant one thing and they heard another.

Two failures, two causes. The second is vocabulary: there is a precise, shared coordinate system for describing the body and you were not using it. The first is deeper: you treated a living structure as a mechanism, and a mechanism is a **subset** of what a body is.

This chapter fixes the vocabulary and installs the habit of asking, of any structure, "what does this need in order to stay alive and keep working?"

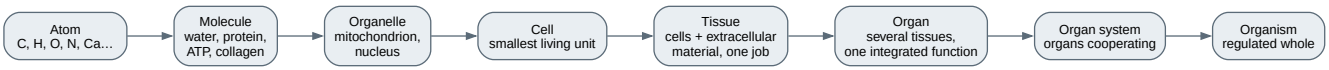
1.2. Anatomy answers *where*; physiology answers *how*

Anatomy is structure: the shapes, positions, and relationships of parts. **Physiology** is function: how those parts work, and how their working is regulated. The elbow's bone contours are anatomy. The fact that muscle activity generates force which a tendon transmits across a joint is physiology.

A model can be anatomically plausible and physiologically misleading. It can also be physiologically honest and anatomically crude — the card-and-cord elbow in Chapter 22 is exactly that, and it is the better teaching aid for its particular objective. Decide which you need. You usually cannot have both at a tabletop budget.

1.2.1. Levels of organisation

Biology is organised hierarchically, and each level introduces properties the level below does not have.



A point that catches engineers: **an organ is never made of one material**. The heart is muscle, but it is also connective tissue, epithelium lining its chambers, nervous tissue carrying its conduction signals, and its own dedicated blood supply. If your mental model of "the heart" is "a pump", you will be surprised by roughly everything that goes wrong with hearts.

In the workshop you might select a metal purely for stiffness and machinability. Living tissue additionally requires:

- a chemical environment within tolerable limits (pH, temperature, osmolarity, ion concentrations),

- for most tissues, a continuing blood supply delivering oxygen and nutrients and removing waste,
- and it **changes with use**: it adapts, it is injured, it repairs, it ages, and it often repairs into something mechanically different from the original.

Cartilage and the inner part of an intervertebral disc have little or no direct blood supply, which is a large part of why they heal poorly. Bone, by contrast, is richly vascularised and remodels continuously. Same skeleton, radically different repair behaviour.

1.3. A shared coordinate system

All anatomical description assumes a standard reference pose, the **anatomical position**: standing upright, facing forward, feet slightly apart, arms at the sides, **palms facing forward**.



Left and right belong to the person being described, not to you and not to the page.

A person facing you has their right side on your left. A radiograph displayed conventionally has the patient's right on the viewer's left. The labels do not change when the person turns around, lies down, or when you rotate the model on the table.

This is the single most common error a newcomer makes in a clinical conversation, and it is the one that causes the most expensive misunderstandings.

The palms-forward detail matters more than it looks. It is why the **radius** is the lateral forearm bone and the thumb is lateral to the little finger. Let the arm hang naturally with the palm facing the thigh and the radius crosses over the ulna — but the anatomical names stay fixed to the reference pose. Describe positions in the reference pose, then separately describe the actual posture.

1.3.1. Directional terms

Pair	Meaning	Example
Anterior (ventral) / posterior (dorsal)	Toward the front / toward the back	The kneecap is anterior to the knee joint. The shoulder blade is posterior.
Superior (cranial) / inferior (caudal)	Toward the head / toward the feet	The chest is superior to the pelvis.
Medial / lateral	Toward / away from the body's midline	In anatomical position the thumb is lateral to the little finger.
Proximal / distal	Nearer / farther from a limb's attachment to the trunk	The elbow is proximal to the wrist; the fingertip is the most distal part of the hand.

Pair	Meaning	Example
Superficial / deep	Nearer / farther from the body surface	Skin is superficial to the underlying muscle.
Ipsilateral / contralateral	Same side / opposite side of the body	A left-sided brain injury commonly affects contralateral limb movement.
Unilateral / bilateral	One side / both sides	A bilateral hand tremor affects both hands.
Palmar (volar) / dorsal (hand)	Palm side / back-of-hand side	A palmar splint sits against the palm.
Plantar / dorsal (foot)	Sole side / top-of-foot side	Plantar surface pressure matters for footwear design.
Prone / supine	Lying face-down / lying face-up	The patient was positioned prone for the procedure.

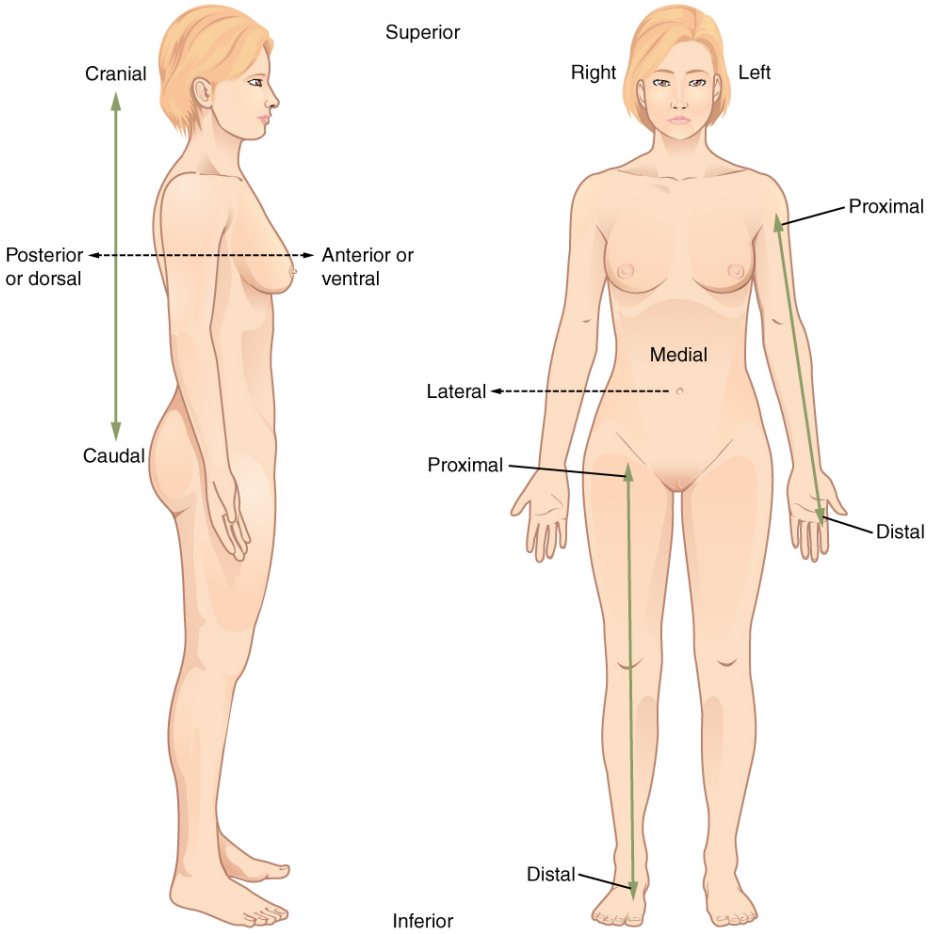
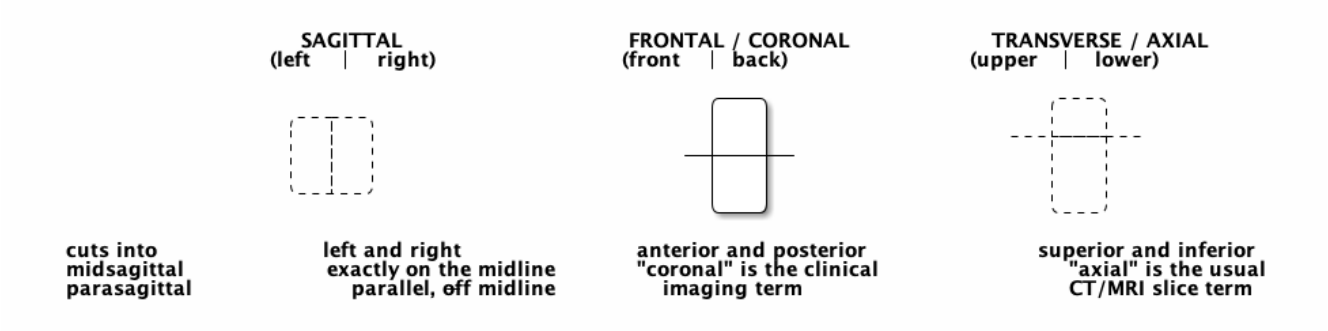


*Figure 1. Directional terms applied to the body in anatomical position. Image: OpenStax College, _Anatomy & Physiology, CC BY 3.0. Credits: Appendix H_*

Proximal and distal are about the limb, not about height. If someone raises their hand above their head, the wrist has not become proximal to the elbow. It is still farther along the limb from the shoulder. This appears in the assessment in Appendix E because it is a reliable trap.

1.3.2. Planes and sections



A **sagittal** plane separates left from right; on the midline exactly it is **midsagittal**, otherwise **parasagittal**. A **frontal** or **coronal** plane separates anterior from posterior. A **transverse** (**axial**, or in imaging, **horizontal**) plane separates superior from inferior. An **oblique** plane is any other angle.

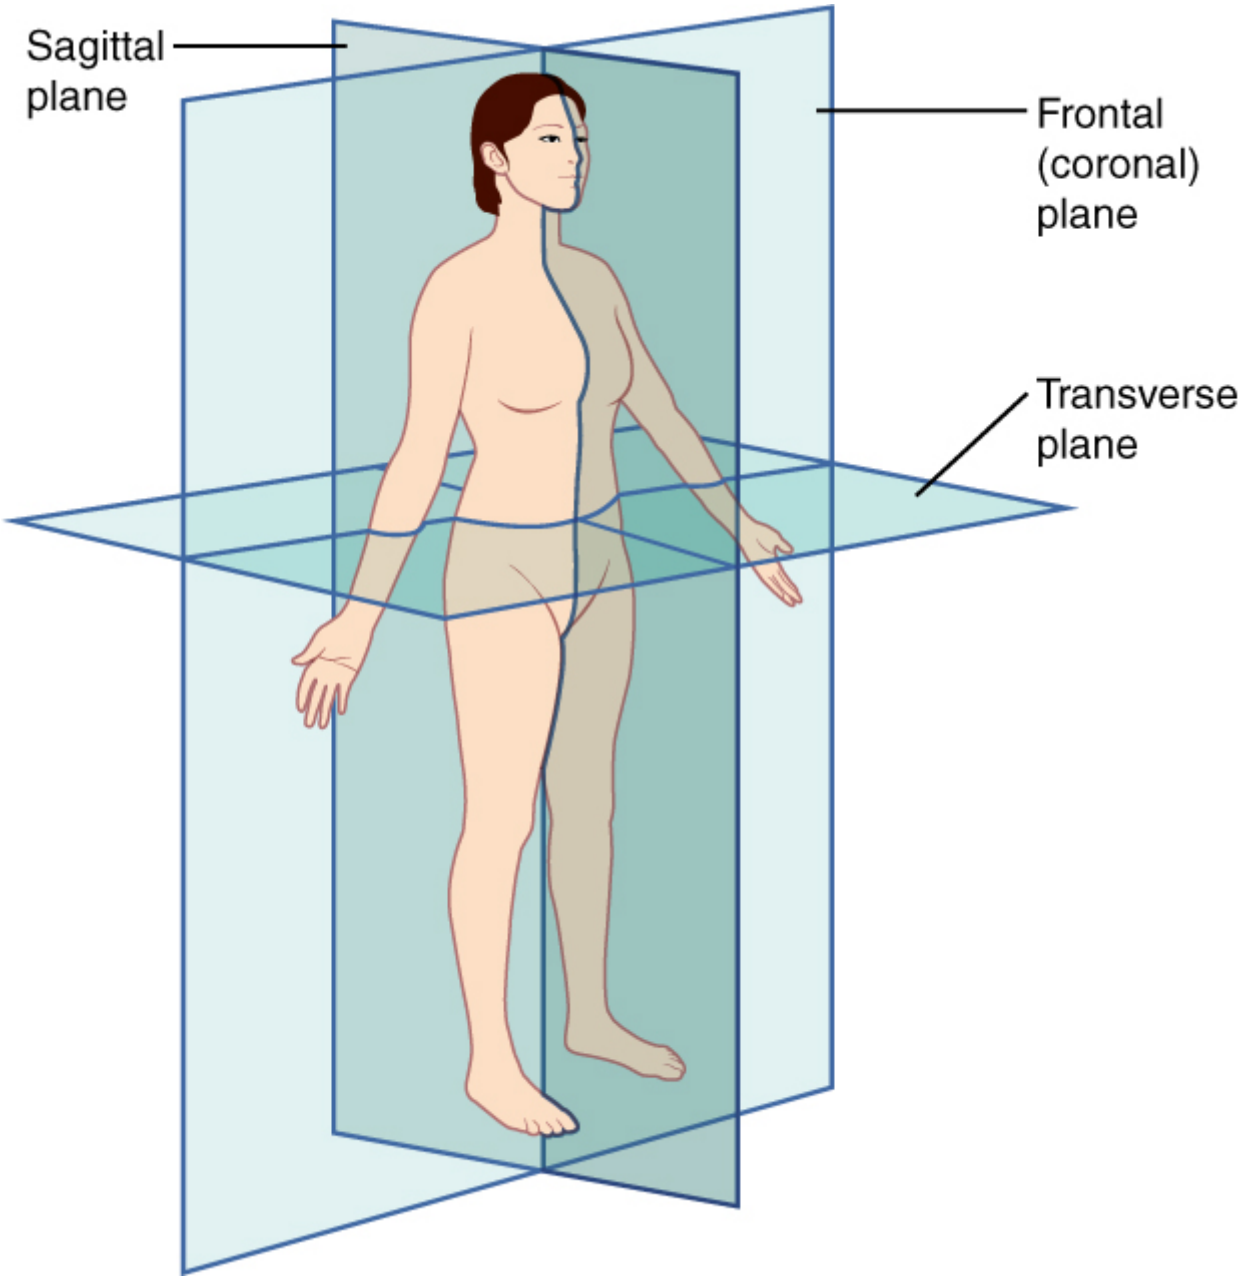


Figure 2. The three principal anatomical planes. Image: Connexions / OpenStax, CC BY 4.0. Credits: Appendix H

If you use CAD, these are section views, and the analogy is genuinely useful. Two places it breaks down:

1. **A flat illustration hides what is behind it.** A textbook diagram of a nerve "running down the arm" has silently chosen a depth. The actual three-dimensional relationship to the artery beside it may be the clinically important fact.
2. **Naming the plane is not reading the image.** A clinical image has orientation markers, a stated patient position, and a convention about which side of the image is the patient's right. Knowing the word "axial" does not qualify you to interpret a CT slice. If a project involves imaging, see the imaging specialism note in Chapter 23 and work with someone who does this professionally.

1.3.3. Body cavities and regions

You need enough of this to follow a conversation, not to pass an exam.

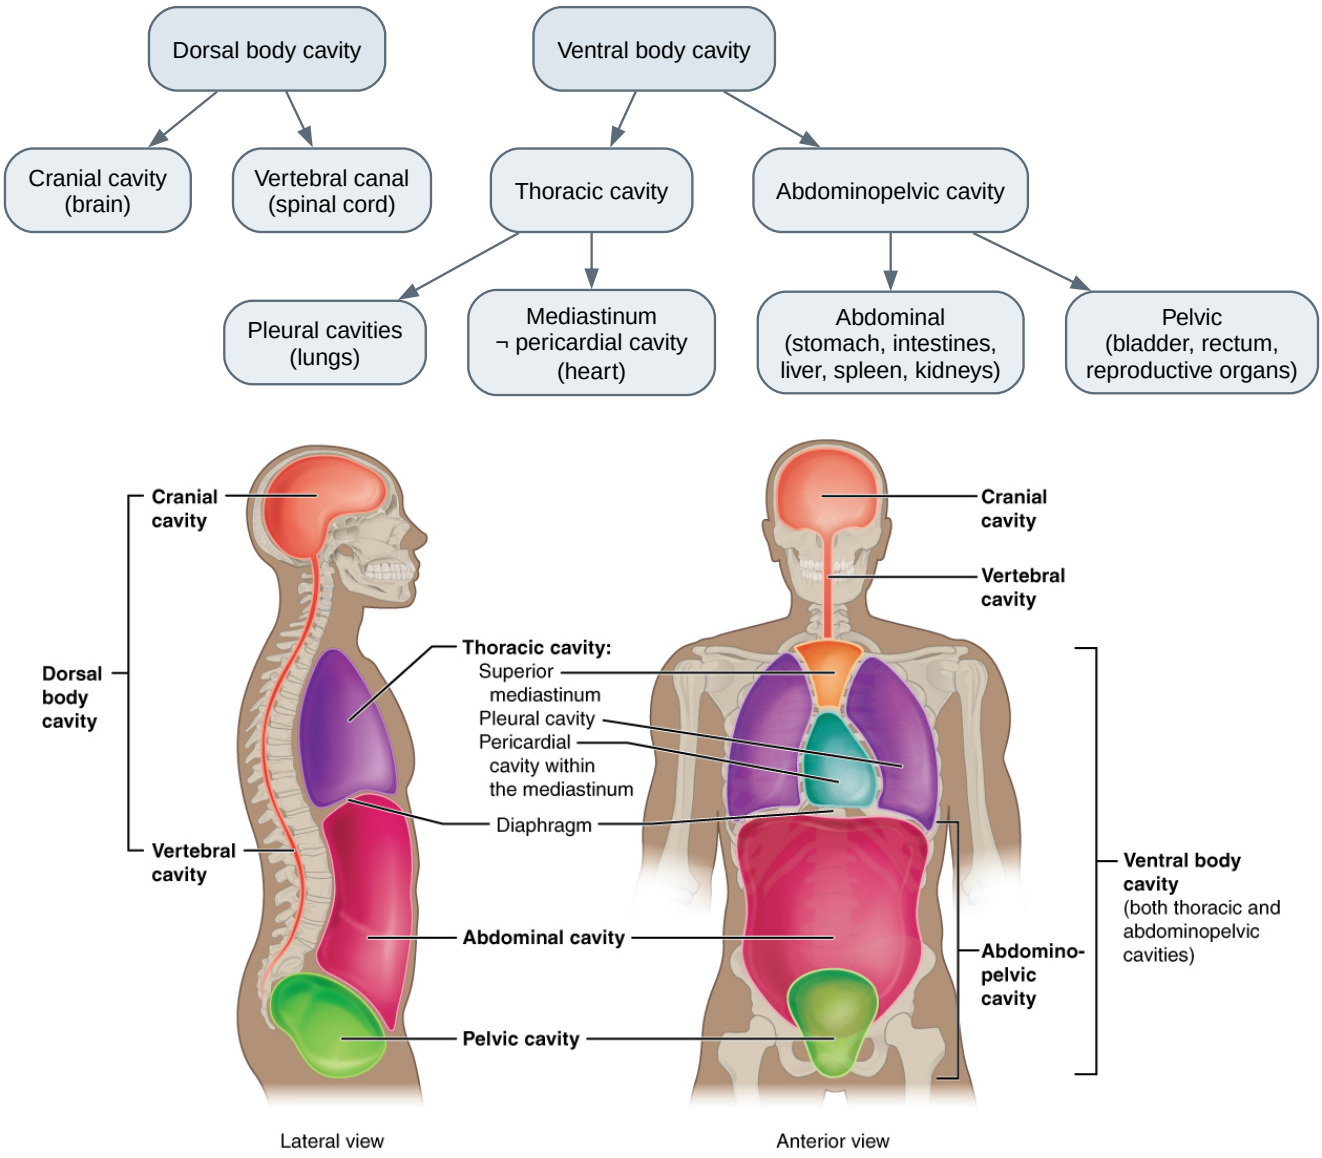


Figure 3. Dorsal and ventral body cavities. Image: Connexions / OpenStax, CC BY 4.0. Credits: Appendix H

The **diaphragm** separates the thoracic from the abdominopelvic cavity, and it is the main muscle of quiet breathing (Chapter 9). Serous membranes — pleura around the lungs, pericardium around the heart, peritoneum in the abdomen — are double-layered sheets with a thin fluid film between

them, which is what lets organs move against each other with very low friction. That is a mechanical fact a maker can appreciate: the body solved its own bearing problem with a fluid film, not a bushing.

For surface description, the nine **abdominal regions** (right/left hypochondriac, epigastric, right/left lumbar, umbilical, right/left iliac/inguinal, hypogastric) or the simpler four **quadrants** (RUQ, LUQ, RLQ, LLQ) are what you will hear. "Pain in the RLQ" is a quadrant description, not a diagnosis.

1.4. Cells maintain an internal environment

This section is about the membrane and the environment it holds. What is inside the cell — the organelles, the machinery that turns genes into proteins, division, and the microbes that share our bodies — is the subject of Chapter 4.

A cell is not a bag. A **plasma membrane** is a phospholipid bilayer studded with proteins, and it is **selectively permeable**: some things cross freely, some need a protein channel or carrier, some are actively pumped against a gradient at the cost of energy, and some do not cross at all.

Mechanism	What moves	Driving force	Needs ATP?
Simple diffusion	Small nonpolar molecules — O_2, CO_2, some lipids	Concentration gradient	No
Osmosis	Water, across a selectively permeable membrane (often via aquaporin channels)	Difference in effective solute concentration	No
Facilitated diffusion	Glucose, ions — via a specific channel or carrier protein	Concentration gradient	No
Primary active transport	Ions against their gradient, e.g. the Na^{+}/K^{+}-ATPase	Direct ATP hydrolysis	Yes
Secondary active transport	A solute coupled to an ion flowing down its own gradient	A gradient previously established using ATP	Indirectly
Bulk transport	Large particles, fluid volumes — endocytosis, exocytosis, phagocytosis	Membrane deformation	Yes

Diffusion is net movement down a concentration gradient — the result of random motion, with no pump required. **Osmosis** is specifically water movement across a selectively permeable membrane in response to a difference in effective solute concentration.

The classic analogy is a smell spreading from one corner of a room. It illustrates the randomness and the net direction. It fails because the room has no selective barrier, no active pumps, and no ongoing energy expenditure to maintain a gradient. **State the limit of your analogy out loud.** A chapter later this becomes a design habit: the model that quietly teaches a false relationship is worse than no model.

1.4.1. Tonicity: why this is not academic

Put a cell in a solution and the water goes where the effective solute concentration tells it to.

Solution	Relative to cell interior	Effect on the cell
Isotonic	Same effective solute concentration	No net water movement; cell volume stable
Hypotonic	Lower solute (more dilute)	Water enters; cell swells and may burst (lyse)
Hypertonic	Higher solute (more concentrated)	Water leaves; cell shrinks (crenates)

This is why intravenous fluids are formulated, not improvised, and why "sterile water" and "sterile saline" are not interchangeable in a clinical context. It is also why wound dressing design is partly a water-management problem. You will not be formulating fluids; you should understand why someone else's choice of fluid was not arbitrary.

1.4.2. Energy, and why oxygen matters

Cells run on chemical energy, overwhelmingly transferred via **ATP** (adenosine triphosphate). Oxygen supports the efficient aerobic pathways that produce most of a cell's ATP; without it, cells fall back on much less productive anaerobic metabolism and accumulate by-products.

You do not need the pathways yet. You need two consequences:

1. **A tissue's function depends on far more than its geometry.** Interrupt the blood supply to a region and its mechanical properties, its sensation, and its ability to repair all change — within minutes for some tissues.
2. **Anything that compresses tissue is potentially restricting its blood supply.** This is the single most important reason that "it fits nicely in CAD" says nothing about whether an object is appropriate against a person for two hours. Chapter 5 develops this.

1.4.3. Enough chemistry to not be lost

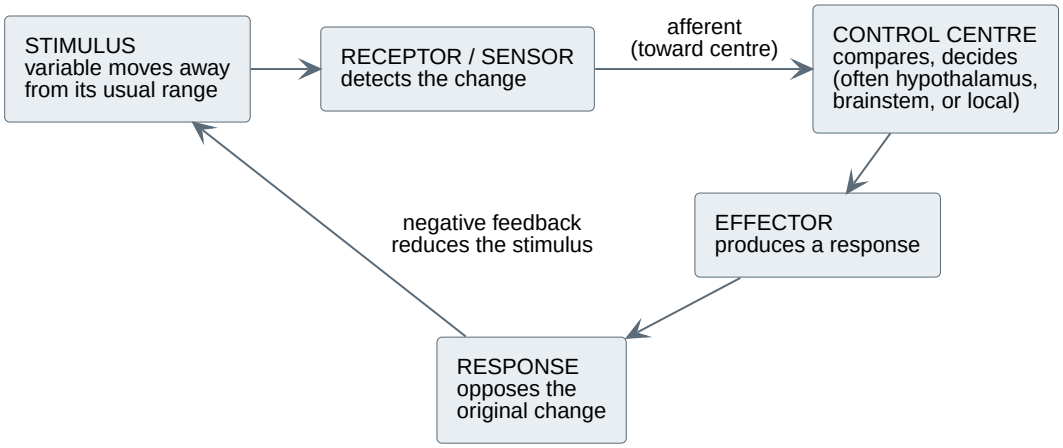
A short orientation, not a chemistry course. Chapter 2 builds general chemistry from the ground up, and Chapter 3 applies it to the molecules of life and the materials you build with.

Idea	What you need to know
Water	The solvent for nearly everything; roughly 50–60% of adult body mass. Its high heat capacity buffers temperature change.
pH and buffers	pH measures hydrogen-ion activity on a log scale. Arterial blood pH is normally held in a narrow band around 7.35–7.45 [1]. Buffer systems (notably bicarbonate) resist change. "Narrow band" is the point: small absolute deviations are physiologically large.
Electrolytes	$Na^{\wedge}$, $K^{\wedge}$, Ca^{2+}, Cl^{-}, and others. Gradients across membranes store energy and underlie nerve and muscle signalling (Chapter 7). Calcium additionally triggers muscle contraction and participates in clotting.

Idea	What you need to know
Macromolecules	Proteins (including structural collagen and elastin, and all enzymes), carbohydrates, lipids, nucleic acids. **Collagen** is the one a mechanical engineer should remember: it is the dominant tensile element in tendon, ligament, skin, and bone's organic matrix.
Enzymes	Protein catalysts, sensitive to temperature and pH. Part of why the body defends those variables so hard.

1.5. Homeostasis is ongoing regulation

Homeostasis is the maintenance of internal conditions within workable ranges. It is not a guarantee, not a fixed set-point, and not a single mechanism. It is many overlapping control loops with finite capacity.



1.5.1. Negative feedback

In **negative feedback** the response opposes the disturbance. "Negative" means *opposing*, not *harmful* — a point worth saying aloud once, because the everyday meaning of the word gets in the way.

Thermoregulation is the standard teaching example. Core temperature rises; thermoreceptors detect it; the hypothalamus acts as the control centre; effectors include sweat glands and dilation of skin blood vessels; the response increases heat loss. On cooling, the loop runs the other way: vessels constrict, shivering generates heat.

Loop component	Temperature rising	Temperature falling
Sensor	Peripheral and central thermoreceptors	Same
Centre	Hypothalamus	Hypothalamus
Effectors	Sweat glands; cutaneous vasodilation	Cutaneous vasoconstriction; skeletal muscle shivering; (in infants) brown fat

Loop component	Temperature rising	Temperature falling
Net effect	Increased heat loss	Reduced heat loss, increased heat production

An electronic thermostat is a reasonable first analogy. Where it fails:

- The body has **many** interacting loops, not one, and they compete. Vasoconstriction to retain heat conflicts with vasodilation to dump the heat of exercise.
- Capacity is finite and **individual**. Age, medication, hydration, acclimatisation, and skin condition all change what the loop can achieve.
- The set-point itself can move — fever is a regulated change in the target, not a failure of regulation.

So: keep the control-systems intuition. Discard the assumption of a single fixed operating point, and discard the assumption that an adequate response always exists.

1.5.2. Positive feedback

Less common, and always a cycle that runs to a defined endpoint rather than a regulator:

- **Haemostasis** — platelet activation recruits more platelets, amplifying clot formation until the breach is sealed.
- **Labour** — oxytocin increases uterine contraction, which increases oxytocin, until delivery.
- **Action potential upstroke** — sodium entry depolarises the membrane, which opens more sodium channels (Chapter 7).

Positive feedback without a terminating event is pathological. A control engineer will recognise the shape immediately.

1.5.3. Variation, and why "normal" is a trap

Physiological variables vary between people, within one person across a day, and with posture, activity, hydration, temperature, emotion, and medication. A published "normal range" is a population-derived interval, often constructed to contain roughly the central 95% of a reference population — which means **by construction** about one in twenty healthy people fall outside it.



Three consequences for anyone who might one day display a number on a screen:

1. A value outside a published range is not automatically abnormal, and a value inside it is not automatically fine.
2. A single measurement is a sample from a distribution, not a property of the person. Trend and context often carry more information than a point value.
3. Deciding what a value **means** for a particular person is a clinical act. Displaying it is not. Keep those separate in your software and in your head (Chapter 9).

1.6. Engineering bridge: scales, units, and language

1.6.1. Orders of magnitude

Intuition for size prevents a class of silly mistakes.

Structure	Typical scale	Maker reference
Protein molecule	~5 nm	—
Cell membrane thickness	~7–8 nm	—
Red blood cell diameter	~7–8 µm	1/10 of a fine human hair
Typical capillary lumen	~5–10 µm	Red cells pass nearly single-file
Skeletal muscle fibre diameter	~10–100 µm	A sheet of printer paper is ~100 µm
Epidermis thickness	~0.05–0.1 mm (much thicker on palm/sole)	Thinner than printer paper
Full skin thickness (epidermis + dermis)	~1–4 mm by site	Thin plywood
Alveolus diameter	~0.2–0.3 mm	Visible as a dot
Nerve fascicle	~0.1–1 mm	—
Typical FDM nozzle	0.4 mm	Larger than an alveolus
Typical CNC end-mill corner radius, laser-cut kerf	~0.1–0.5 mm	Comparable to an alveolus
Sewing thread; fine woven fabric pitch	~0.2–0.5 mm	Comparable to an alveolus

Note the last three lines. The smallest feature that ordinary workshop processes produce — a printed bead, a milling cutter's corner radius, a laser kerf, a stitch — is comfortably bigger than many structures that matter biologically. That is no reason not to make teaching models by any of those methods; it is a reason to be explicit in your labelling about what your model does and does not resolve, whatever it is made from.

1.6.2. Units you will meet

Quantity	Usual clinical unit	Note for an engineer
Pressure (blood)	mmHg	1 mmHg ≈ 133.3 Pa. 120 mmHg ≈ 16 kPa.
Pressure (airway, CVP)	cmH_2O	1 cmH_2O ≈ 98.1 Pa ≈ 0.735 mmHg.
Force	N	Grip strength is often reported in kg-force; 1 kgf ≈ 9.81 N.
Interface pressure	mmHg (wound care) or kPa	Watch for unit switches between papers.

Quantity	Usual clinical unit	Note for an engineer
Temperature	°C	Core temperature is commonly quoted around 36.5–37.5 °C.
Concentration	mmol/L (SI) or mg/dL (US)	Glucose: 5 mmol/L ≈ 90 mg/dL. A real source of error.
Volume (respiratory)	L, mL	Tidal volume ~500 mL in a typical adult at rest.
Potential (biosignal)	mV, µV	ECG ~1 mV; EEG ~10–100 µV. See Chapter 7.



Unit mismatch is a recurring cause of real clinical software incidents — mmol/L versus mg/dL, kg versus lb for weight-based dosing, cm versus inches. If you write software that handles physiological values: carry units in the type, never in a comment; reject rather than guess on ambiguous input; and display the unit next to every number, every time.

1.6.3. Reading medical words

Clinical vocabulary is largely compound Greek and Latin. Learning twenty roots lets you decode hundreds of terms and, more importantly, lets you make a sensible guess and then verify it.

Element	Meaning	Example
`cardi/o-`	heart	cardiology, myocardium, bradycardia
`my/o-`	muscle	myocardium, myopathy, electromyography
`oste/o-`	bone	osteon, osteoporosis, osteotomy
`arthr/o-`	joint	arthritis, arthroplasty, arthroscopy
`derm/o-`, `cutane/o-`	skin	dermis, subcutaneous, dermatology
`neur/o-`	nerve	neuron, neuropathy, neurology
`angi/o-`, `vas/o-`	vessel	angiography, vasodilation
`pulmon/o-`, `pneum/o-`	lung, air	pulmonary, pneumothorax
`nephr/o-`, `ren/o-`	kidney	nephron, renal
`hepat/o-`	liver	hepatic, hepatitis
`haem/o-`, `hem/o-`	blood	haemoglobin, haemostasis, haematoma
`-itis`	inflammation	arthritis, dermatitis, tendinitis
`-ectomy`	surgical removal	appendectomy
`-ostomy`	creating an opening	colostomy, tracheostomy

Element	Meaning	Example
`-otomy`	cutting into	osteotomy, laparotomy
`-plasty`	surgical repair / reshaping	arthroplasty, angioplasty
`-algia`, `-dynia`	pain	neuralgia
`-pathy`	disease	neuropathy, myopathy
`-oma`	tumour, mass	haematoma, lipoma
`-penia`	deficiency	leukopenia
`hyper-` / `hypo-`	above / below normal	hypertension, hypoxia, hypodermic
`brady-` / `tachy-`	slow / fast	bradycardia, tachypnoea
`peri-` / `endo-` / `epi-`	around / within / upon	pericardium, endocardium, epidermis
`-scopy`	viewing	arthroscopy, endoscopy
`-gram` / `-graphy`	record / process of recording	electrocardiogram, radiography

Work through MedlinePlus **Understanding Medical Words** [3] for systematic practice.



Build vocabulary from your actual project, not from a list. Twenty cards drawn from a real conversation with a therapist will stick. Two hundred cards from a textbook index will not.

And **always verify abbreviations.** Clinical abbreviation is dense, local, and ambiguous — the same letters mean different things in different departments and different countries. Asking "sorry, what does that stand for here?" costs you nothing and marks you as careful rather than ignorant.

1.7. Devices, software, and training technology

Each chapter from here on carries a section like this one: what exists clinically, what software exists, what is used to **teach** the subject, and where a maker can realistically contribute. This first one covers the general-purpose equipment that belongs to no single body system.

1.7.1. Clinical devices

Device	Does	Engineering note
Examination couch, hoist, patient transfer aids	Positioning and safe moving of people	Load cases, stability, pinch points, cleanability. Unglamorous and genuinely important
Vital signs monitor (spot-check and continuous)	Heart rate, blood pressure, SpO_2, temperature, respiratory rate	Treated in detail in Chapter 9. Alarm design is a recognised human factors problem

Device	Does	Engineering note
Clinical thermometer — tympanic, temporal, oral, oesophageal	Body temperature at a site	Sites are not interchangeable; each has a different relationship to core temperature
Weighing scales, stadiometer, calipers	Anthropometry, and the basis of weight-based drug dosing	Accuracy here propagates straight into dose calculation
Pressure-redistribution mattresses and cushions	Manage interface pressure over long periods	See Chapter 5. Active (alternating) and reactive (constant low pressure) designs
Patient identification systems — wristbands, barcodes, RFID	Match the right person to the right intervention	Misidentification is a persistent cause of serious incidents. A pure systems-design problem
Call systems and alerting	Summon help	Latency, escalation, and alarm fatigue. Accessible call buttons and switches — big, low-force, mountable — are an electronics and enclosure problem as much as a software one
Positioning aids — wedges, rolls, foam blocks, slide sheets, slings	Support and move people in bed and chair	Foam conversion, upholstery, sewing, and fluid-resistant, wipeable covers. A textiles and cleanability problem more than a shape problem (Chapter 5, Chapter 13)

1.7.2. Software

Software	Does	Note
Electronic health record (EHR/EPR)	The clinical record, orders, results, documentation	The system everything else integrates with. Usability is a major source of clinician burden and of error
HL7 v2, FHIR, DICOM, SNOMED CT, LOINC, ICD	Messaging, API, imaging, and terminology standards	**Learn FHIR if you write clinical software.** It is a modern REST/JSON API and the direction of travel. SNOMED CT and LOINC are the coding systems for findings and observations

Software	Does	Note
Early warning score calculators — NEWS2 and equivalents	Aggregate vital signs into an escalation trigger	A worked example of the inference layer in Chapter 9. A score is not a diagnosis
Clinical decision support	Alerts, order sets, guideline prompts	Alert fatigue is the central design problem
Scheduling, bed management, flow	Operational logistics	Where software skills contribute real value with comparatively low clinical risk
Open EHR platforms — OpenMRS, openEHR, HAPI FHIR	Open-source record systems and toolkits	**Runnable on your own machine with synthetic data.** The realistic way to learn this domain without touching real records
Synthetic data generators — Synthea and similar	Produce realistic but entirely fictional patient records	Exactly what Chapter 19 asks for: develop against plausible data that is not anybody's

1.7.3. Training and simulation

Training technology	Teaches	Form
Anatomical models — skeletons, torsos, organ models	Structure, spatial relationships, vocabulary	The oldest medical training technology. Still everywhere, still breaking, still needing replacement parts — cast, moulded, turned, wired, sewn, or printed according to the part. **A reliable maker entry point**
Virtual dissection tables — Anatomage and similar	Sectional anatomy from real datasets, interactively	Expensive capital equipment; increasingly common in medical schools
Anatomy software — Complete Anatomy, Visible Body, BioDigital	3D interactive anatomy on a tablet or desktop	Commercial, widely licensed by institutions

Training technology	Teaches	Form
Open anatomy resources — Visible Human Project, BodyParts3D, Z-Anatomy	Free anatomical datasets and models	**Z-Anatomy is open source and built on Blender.** Directly usable by a maker
Standardised patients (actors)	Communication, history taking, examination technique	Not a device, but a central simulation modality worth knowing exists
Part-task trainers	One isolated skill — injection, catheterisation, suturing	The category your capstone belongs to (Chapter 21)
Full-body patient simulators — SimMan, HAL and similar	Integrated scenarios with physiological response	High-fidelity manikins with programmable physiology, costing as much as a car. **Consumables and replaceable parts are a genuine maker niche**
Debriefing and video systems	Reviewing performance after a scenario	The debrief, not the manikin, is where most simulation learning happens



Replacement parts and consumables are the realistic way in.

A simulation centre has manikins whose skins tear, whose veins perish, whose jaw hinges break, and whose spare parts cost a great deal and arrive slowly. Replacement parts and reusable consumables are a genuine unmet need, they are low-risk, and they build the relationship that leads to more interesting work. The right method depends on the part: cast silicone for skins and veins, machined or turned metal and acetal for hinges and pivots, sewn covers and straps, re-wired connectors and sensors, and printing where a one-off complex shape is the problem (Chapter 14).

Ask the question from Chapter 18: **"what breaks most often, and what is slowest to replace?"** The answer is usually sitting in a cupboard.

1.8. Worked example: an ambiguous instruction

An educator says: **"Put the cable on the inside of the arm."**

Count the readings:

1. **Medial** — toward the torso, on the inner side of the limb.
2. **Anterior** — on the flexor surface, which casual speech often calls "the inside of the elbow".
3. **Deep** — beneath the skin, internal to the limb.
4. **Within the model's hollow interior** — a statement about your construction, not about anatomy at all.

Four readings, four different builds. The repair is not to demand that the educator speak in Latin. It is to **produce a drawing and ask them to correct it**:

- Name the side explicitly: "the model's left arm".
- Label the segments: "upper arm (humerus region)", "forearm (radius/ulna region)".
- Mark the actual cable route with a line, and dimension it from identifiable features.
- Add a section view if depth matters, and name the plane.
- Write the anatomical term you believe they meant, so they can correct the term as well as the geometry: "cable routed on the **anterior** surface, superficial".

Anatomical vocabulary makes the conversation precise. It does not replace dimensioned, unambiguous views. You need both, and the drawing is what gets corrected.

1.9. Workshop exercise

1. **Body map.** Draw front and side views of a person, large, on paper. Label at least ten regions: cranial, cervical, thoracic, abdominal, pelvic, brachial (arm), antebrachial (forearm), carpal (wrist), femoral (thigh), crural (leg), tarsal (ankle). Add the three principal planes.
2. **Direction drill.** Describe the position of an imaginary forearm-mounted object in anatomical language, from the subject's point of view. Then hand the description to someone else and have them place a sticky note on their own forearm from your words alone. Any ambiguity they hit is a defect in your description, not in their reading.
3. **Vocabulary cards.** Make twenty cards. Front: the term. Back: a plain-English explanation **and** a drawing or concrete example. Draw the content from this chapter's tables plus anything you have actually heard used.
4. **Elbow model sketch.** Sketch the elbow model you will eventually build. Label proximal, distal, anterior, posterior. Mark one section plane and name it. Note which features are anatomical claims and which are construction convenience — you will need that distinction in Chapter 21.

1.10. Questions

1. A person raises their hand above their head. Has their wrist become proximal to their elbow?
2. You are looking at a person facing you. They have a scar on the lateral side of their right forearm. On which side of your field of view do you see it?
3. What is missing from the explanation "the tendon moves because it is stretchy"?
4. In a temperature-control example, what specifically makes the feedback *negative*?
5. Name one way the thermostat analogy for homeostasis misleads.
6. A solution is hypotonic to a cell. Which way does water move, and what happens to the cell?
7. Why does cartilage heal poorly compared with bone?
8. Decode `arthroscopy`, `bradycardia`, and `osteotomy` from their parts.
9. A lab value falls outside the published reference range. Is the person abnormal?

10. Your model's smallest feature — a printed bead, a cutter radius, or a stitch — is about 0.4 mm. Name two anatomical structures it cannot resolve.

1.11. Answers and reasoning

1. **No.** Proximal and distal refer to position along the limb relative to its attachment to the trunk, not to height above the floor. The wrist remains distal to the elbow in every posture.
2. **On your left.** Their right side faces your left. "Lateral" on the right forearm means the thumb side in anatomical position — so you must also establish how the arm is currently rotated before you can point to it on the actual person.
3. It does not identify the **source** of the force or distinguish **generation** from **transmission**. Muscle activity produces force; the tendon transmits it to an attachment. Elastic properties of tendon do matter mechanically, but stretch alone is not the mechanism, and a model that implies the tendon is the motor teaches an error.
4. The response **opposes** the original change, reducing the stimulus that triggered it. "Negative" describes the sign of the loop, not a value judgement.
5. Acceptable answers include: the body runs many competing loops rather than one; the capacity of each loop is finite and varies between individuals; the set-point itself can shift (fever); and a response existing does not mean it is sufficient.
6. Water moves **into** the cell, because the cell interior has the higher effective solute concentration. The cell swells and may lyse.
7. Cartilage has little or no direct blood supply, so the delivery of cells and nutrients required for repair is poor. Bone is richly vascularised and remodels continuously.
8. `arthr/o` joint + `-scopy` viewing → looking inside a joint. `brady-` slow `cardi/o` heart + `-ia` condition → slow heart rate. `oste/o` bone + `-otomy` cutting into → surgically cutting bone.
9. **Not necessarily.** A reference interval is typically constructed to contain about the central 95% of a reference population, so roughly 1 in 20 healthy people fall outside it by definition. Interpretation requires clinical context and is not yours to make.
10. Any two of: capillaries (~5–10 µm), red blood cells (~7–8 µm), individual muscle fibres (~10–100 µm), epidermal layers (~0.05–0.1 mm total), alveoli (~0.2–0.3 mm, marginal). Say so on the label.

1.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Unit-safe clinical arithmetic — pressure units and reading a number with its unit.
- Test accuracy and risk — why one healthy person in twenty falls outside a reference range.
- Unit converter — general unit conversion.

1.13. Video companion and reading

- **Core video:** Crash Course Anatomy & Physiology #1 [7]. Watch once for the shape of it, then again while drawing your body map.
- **Core reading:** OpenStax *Anatomy and Physiology 2e* [1] chapter 1 (especially 1.5 Homeostasis and 1.6 Anatomical Terminology) and chapter 3 (3.1 The Cell Membrane, 3.2 The Cytoplasm and Cellular Organelles). Skim chapter 2 only to clarify unfamiliar chemistry — do not try to learn it.
- **Vocabulary:** MedlinePlus *Understanding Medical Words* [3], completed.
- **Optional structure:** University of Michigan Anatomy Specialization [2] if you want longer lectures; treat it as a **replacement** for some video and reading time, not an addition.

1.14. Exit criteria

- ▢ Correctly describe five locations using anatomical directions, including the subject's left and right, verified by another person following your description.
- ▢ Explain anatomy, physiology, cell, tissue, organ, and system with one example of each, without notes.
- ▢ Use all five core direction pairs correctly in a labelled drawing.
- ▢ Explain diffusion, osmosis, and negative feedback in your own words, **including one stated limitation of each engineering analogy you use**.
- ▢ Decode five unfamiliar medical terms from their roots and verify your guesses.
- ▢ Score at least 8/10 on a terminology check and correct every remaining answer in writing.
- ▢ Re-attempt the missed items one week later.

1.15. What this chapter does not establish

- **You cannot interpret clinical images.** You know the names of the planes. Orientation conventions, windowing, artefacts, and anatomical variation are a specialism.
- **You cannot interpret a measurement.** You know that reference ranges are population-derived intervals. Deciding what a value means for a person is a clinical act.
- **You have not learned metabolism.** Deliberately. Detailed pathways, exhaustive chemical nomenclature, and memorising every anatomical name are all deferred until a project requires them.
- **Vocabulary is not understanding.** Being able to say "the cable runs on the anterior surface" does not mean you know whether it should.

Chapter 2. Foundations of chemistry: atoms to reactions, and how chemists work

Week 3 · approximately 6–8 hours

2.1. Why this chapter exists

Chemistry turns up in almost every conversation this book prepares you for, usually without announcing itself. A nurse says a cleaning solution is "a thousand ppm". A pharmacist says a drug is "the S-enantiomer". A laboratory scientist says a sample was "run on the LC-MS". An engineer says a stainless part "passivated badly" or a lithium cell "went into thermal runaway". A blood gas report gives a partial pressure in kilopascals and a bicarbonate in millimoles per litre. Each of those sentences assumes the same few ideas: atoms and the periodic table, the mole, concentration, reactions and their energy, rates, equilibrium, oxidation and reduction.

If you last did chemistry at school, those ideas are probably still there, but rusty and disconnected from anything you build. This chapter rebuilds them from the ground up, with the engineering arithmetic attached, and then does something a school course rarely does: it shows you **how chemistry is practised** — the branches of the field that touch health, what each of them actually does all day, and the instruments and software they use to do it.

Three practical reasons it matters to someone with a full workshop and an electronics bench:

- **Everything you make is a chemical system.** A machined stainless part, a moulded polypropylene housing, an anodised aluminium plate, a lithium battery, a sewn polyester strap, a glass vessel, and a printed polymer each have a composition, react with what they meet — disinfectants, sweat, saline, oxygen, heat — and change over time. Predicting how is chemistry.
- **Measurement in medicine is overwhelmingly chemical.** Most laboratory tests, many bedside sensors, and a large share of what a clinical engineer maintains measure a concentration. Chapter 15 and Chapter 16 assume you can read one.
- **The workshop is a chemical environment.** Solvents, adhesives, fluxes, resins, etchants, and cleaning agents carry real hazards. Knowing why — and how a chemist assesses them — is part of working safely and of being trusted in someone else's building.



This chapter is general chemistry. Chapter 3 then applies it to the molecules of life — water, pH and buffers, proteins and enzymes, membranes, bioenergetics — and to the materials you build with: polymers, leachables, metals and corrosion, textiles, glass, and adhesion. Chapter 4 takes the molecules into the living cell. Where this chapter touches a topic those chapters treat in depth, it points there rather than repeating it.

2.2. Atoms: what everything is made of

2.2.1. Inside an atom

An atom is a tiny, dense, positively charged **nucleus** surrounded by a cloud of **electrons**.

- The nucleus contains **protons** (positive charge) and **neutrons** (no charge), each with a mass of about 1 atomic mass unit (u, also called the dalton, Da).
- Electrons carry one negative charge each and have about 1/1836 of a proton's mass. Almost all of an atom's mass is in the nucleus; almost all of its volume is electrons.
- The nucleus is roughly 10^{-15} m across; the atom roughly 10^{-10} m. If the nucleus were a pea in the middle of a football pitch, the electrons would be somewhere around the stands.

The **atomic number** (Z) — the number of protons — defines the element. Six protons is carbon, always. The **mass number** is protons plus neutrons. A neutral atom has as many electrons as protons.

Chemistry is the behaviour of the outer electrons. Nuclei do not change in chemical reactions; electrons are shared, transferred, or rearranged between atoms. Nuclear changes — radioactivity, fission, fusion — are physics, and they appear in medicine as imaging and therapy (Chapter 16, Chapter 17).

2.2.2. Electron shells and why atoms react

Electrons occupy **shells** (energy levels) around the nucleus. For the first twenty elements, a useful simplification is that the shells hold 2, 8, 8 electrons in turn. The electrons in the outermost occupied shell — the **valence electrons** — decide how the atom behaves.

Atoms are most stable with a full outer shell, as the **noble gases** (helium, neon, argon) already have. Everything else reacts in ways that get closer to that arrangement:

- **Lose** electrons — sodium (one valence electron) readily gives one up to become Na^+.
- **Gain** electrons — chlorine (seven) readily takes one to become Cl^-.
- **Share** electrons — carbon (four) shares with up to four partners, which is why carbon can build the long chains and rings of organic chemistry.

2.2.3. Ions and electronegativity

An **ion** is an atom or group of atoms with a net charge: a **cation** if positive (Na^+, K^+, Ca^{2+}, Mg^{2+}, H^+), an **anion** if negative (Cl^-, HCO_3^-, HPO_4^{2-}, OH^-). The body's electrolytes are ions, and their gradients across membranes run nerves and muscle (Chapter 7).

Electronegativity is how strongly an atom pulls shared electrons towards itself. On the common Pauling scale: fluorine about 4.0, oxygen 3.4, nitrogen 3.0, carbon 2.6, hydrogen 2.2, sodium 0.9. The **difference** between two bonded atoms predicts the kind of bond:

Electronegativity difference	Bond	Example
Small (< ~0.4)	Non-polar covalent	C–H in fats and polyethylene; the reason they repel water
Moderate	Polar covalent	O–H in water and alcohols; C=O in proteins and polyesters
Large (> ~1.8)	Ionic	Na^+ Cl^- in salt

The bond types themselves — covalent, ionic, hydrogen bonds, van der Waals forces, and the hydrophobic effect — and why water behaves as it does are covered in Chapter 3.

2.2.4. Isotopes

Atoms of one element can have different numbers of neutrons. These are **isotopes**: the same chemistry, different mass, and sometimes an unstable nucleus.

- Carbon-12 (6 protons, 6 neutrons) is stable; carbon-14 (8 neutrons) is radioactive and decays slowly — the basis of radiocarbon dating.
- **Fluorine-18** emits positrons and has a half-life of about 110 minutes; attached to glucose it becomes FDG, the main PET tracer.
- **Technetium-99m** emits a gamma ray with a six-hour half-life; it is the workhorse of nuclear medicine imaging.
- **Iodine-131** emits beta particles that destroy thyroid tissue, which concentrates iodine — a treatment as well as a tracer.

Because isotopes are chemically almost identical, the body handles a radioactive isotope the way it handles the stable one. That is the entire trick of nuclear medicine (Chapter 16).

2.3. The periodic table

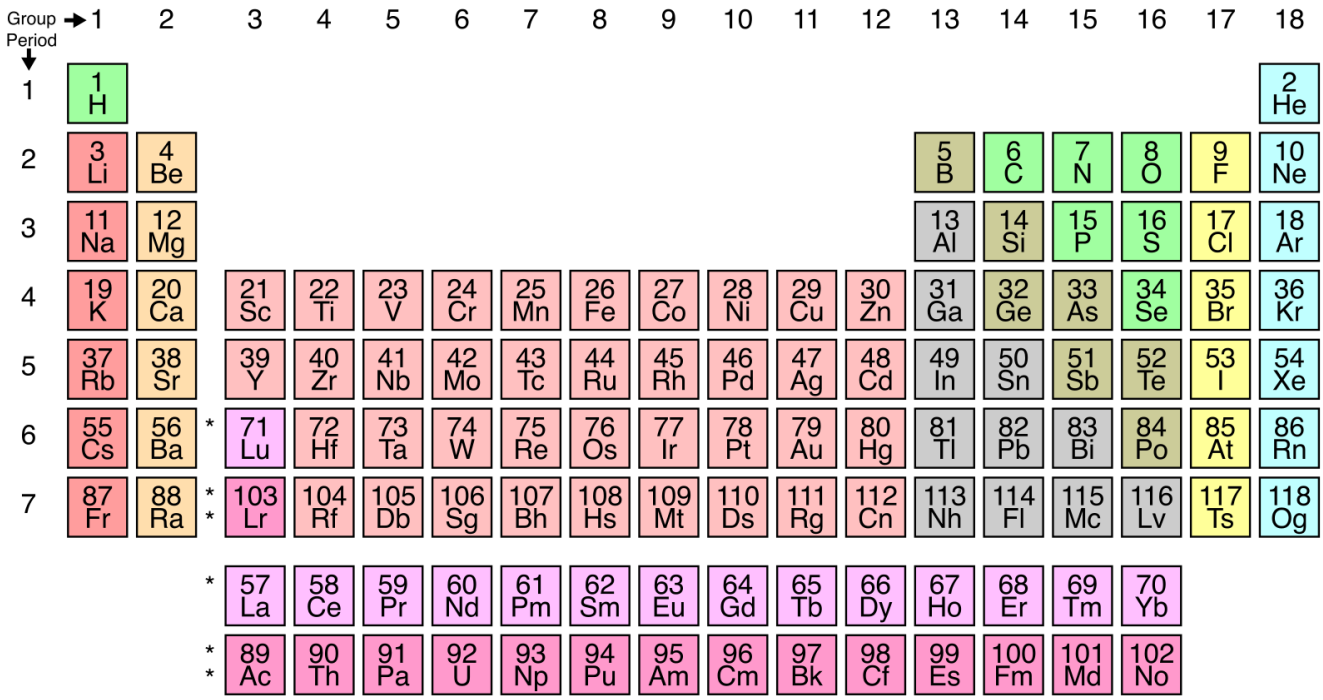


Figure 4. The periodic table, arranged by atomic number. Columns (groups) share outer-electron arrangements and therefore chemistry; rows (periods) add a shell. Image: Offnfopt, public domain. Credits: Appendix H

The table is not a list; it is a map of electron arrangement.

- **Rows (periods)** add a new electron shell. Atoms get bigger going down.
- **Columns (groups)** share the same number of valence electrons, so they behave alike. Group 1 — lithium, sodium, potassium — all form +1 ions and react vigorously with water. Group 17 — fluorine, chlorine, bromine, iodine — all form −1 ions and are strong oxidisers. Group 18 — the noble gases — barely react at all.
- **Metals** sit on the left and in the middle; **non-metals** on the right; a staircase of **metalloids** (boron, silicon, germanium) between them.
- **Trends** run across and down: atoms get smaller and more electronegative left-to-right across a period, and larger and less electronegative going down a group.

2.3.1. The elements of the body

Four elements make up about 96% of body mass. Most of the rest is a handful of minerals, and a long tail of trace elements is essential in tiny amounts.

Element	Approximate share of body mass	Main roles
Oxygen (O)	~65%	Water; every biological macromolecule; respiration
Carbon (C)	~18%	The backbone of all organic molecules
Hydrogen (H)	~10%	Water; organic molecules; pH (H^+)
Nitrogen (N)	~3%	Proteins (amino groups), nucleic acids

Element	Approximate share of body mass	Main roles
Calcium (Ca)	~1.5%	Bone and teeth; muscle contraction; clotting; signalling
Phosphorus (P)	~1%	Bone mineral; ATP; DNA and RNA backbone; membranes
Potassium, sulfur, sodium, chlorine, magnesium	each well under 1%	Electrolytes, membrane potentials, protein structure, enzyme cofactors
Trace: iron, zinc, copper, iodine, selenium, manganese, cobalt, fluorine, molybdenum	tiny	Haemoglobin (Fe), enzymes (Zn, Cu, Se, Mn, Mo), thyroid hormone (I), vitamin B_{12} (Co)

2.3.2. The elements of devices

A second, overlapping set of elements turns up in what you build and what hospitals use:

Element	Where you meet it
Iron, chromium, nickel, molybdenum	Stainless steel; 316L adds molybdenum for chloride resistance (Chapter 3). Nickel is a common skin sensitiser
Titanium	Implants and instruments; a stable oxide film makes it biocompatible
Cobalt and chromium	Wear-resistant alloys for joint replacements and dental frameworks
Aluminium	Housings, frames, heat sinks; anodising thickens its oxide layer
Copper and silver	Wiring and electronics; both are antimicrobial, which is why silver appears in dressings and coatings
Gold and platinum	Corrosion-free contacts and electrodes; platinum–iridium in pacemaker and neurostimulator leads
Silicon	Electronics; silica in glass; the Si–O backbone of silicone elastomers
Lithium	Batteries (Chapter 14); also a psychiatric drug with a narrow therapeutic range
Iodine and barium	X-ray contrast agents, chosen for their high atomic number (Chapter 16)
Gadolinium	MRI contrast agents, chosen for its unpaired electrons
Tungsten and lead	X-ray tube targets and radiation shielding

2.4. States of matter and gases

2.4.1. Solids, liquids, gases, and changes between them

Whether a substance is solid, liquid, or gas depends on how strongly its particles attract one

another compared with their thermal energy. Heating supplies energy; at the **melting point** a solid becomes liquid, at the **boiling point** a liquid becomes gas. Each change absorbs **latent heat** without changing temperature — which is why sweat cools you and why steam at 134 °C delivers so much more energy to a surgical instrument than dry air at the same temperature (Chapter 13).

Vapour pressure is the pressure exerted by the vapour above a liquid. It rises steeply with temperature, and a liquid boils when its vapour pressure equals the surrounding pressure. Volatile liquids — acetone, isopropyl alcohol, ether — have high vapour pressures at room temperature: they evaporate quickly, fill a room with flammable vapour, and cool the skin. An anaesthetic vaporiser is a device engineered around vapour pressure.

2.4.2. The gas laws

Gases are where chemistry meets the physiology of breathing. For an **ideal gas**:

PV = nRT

where **P** is pressure, **V** volume, **n** the amount in moles, **T** absolute temperature (kelvin = °C + 273.15), and **R** the gas constant, 8.314 J/(mol·K). The individual laws you may remember are special cases:

- **Boyle's law** — at constant temperature, P × V is constant. Expand the chest, pressure in the lungs falls, air flows in (Chapter 9).
- **Charles's law** — at constant pressure, V is proportional to T. Gas warms and expands as it enters the airway.
- **Avogadro's law** — at a given temperature and pressure, equal volumes hold equal numbers of molecules. One mole of any ideal gas occupies about 22.4 L at 0 °C and 1 atm, about 24.0 L at 20 °C, and about 25.4 L at body temperature.

Dalton's law of partial pressures says each gas in a mixture contributes pressure in proportion to its share. Dry air at sea level is 760 mmHg in total; oxygen is about 20.9% of it, so its **partial pressure** is about 159 mmHg. By the time air reaches the trachea it has been warmed to 37 °C and saturated with water vapour, whose partial pressure is about 47 mmHg, so the oxygen partial pressure there is (760 − 47) × 0.209 ≈ 149 mmHg. Blood gas reports and the alveolar figures in Appendix C are partial pressures in exactly this sense.

Henry's law says the amount of a gas dissolved in a liquid is proportional to its partial pressure above it. Oxygen is poorly soluble in plasma — about 0.003 mL per 100 mL of blood per mmHg — which is why haemoglobin exists. Raise the pressure and more dissolves: hyperbaric oxygen therapy uses this deliberately, and decompression sickness is what happens when dissolved nitrogen comes back out of solution too fast.

2.5. Amounts and concentrations

2.5.1. The mole

Atoms are too small to count, so chemists count them in batches. One **mole** is 6.022×10^{23} particles — **Avogadro's number**. Its usefulness is that the mass of one mole of a substance, in grams, equals

its molecular mass in daltons:

- Carbon: 12.01 g/mol. Oxygen: 16.00 g/mol. Hydrogen: 1.008 g/mol.
- Water, H_2O: 2 × 1.008 + 16.00 = 18.02 g/mol.
- Sodium chloride, NaCl: 22.99 + 35.45 = 58.44 g/mol.
- Glucose, $C_6H_{12}O_6$: 6 × 12.01 + 12 × 1.008 + 6 × 16.00 = 180.16 g/mol.

Moles are what reactions care about: one molecule of glucose reacts with six of oxygen, not one gram with six grams.

2.5.2. Concentration units, and why medicine uses so many

Unit	Meaning	Where you meet it
mol/L (M), mmol/L, µmol/L	Moles of solute per litre of solution	Electrolytes, glucose (most of the world), creatinine; the SI convention in laboratory medicine
mg/dL, g/L, µg/L	Mass per volume	Glucose and cholesterol in the US; protein; drug levels
% w/v	Grams per 100 mL	Intravenous fluids: 0.9% saline = 0.9 g NaCl per 100 mL = 9 g/L
% v/v	Volume per 100 mL	Alcohol disinfectants: 70% v/v isopropanol or ethanol
ppm, ppb	Parts per million, per billion (in dilute water solutions, 1 ppm ≈ 1 mg/L)	Chlorine disinfectant strength; trace metals; environmental limits
mmol/mol, %	Ratio of one molecule to another	HbA_{1c} (Chapter 10)
mOsm/kg (osmolality), **mOsm/L** (osmolarity)	Moles of dissolved particles per kilogram of water or per litre of solution	Plasma tonicity; IV fluid design; why cells swell or shrink (Chapter 1)
Units (U), international units (IU)	Defined by biological activity, not mass	Insulin, heparin, some vitamins

Two worked conversions show why this matters.

Glucose. A blood glucose of 5.0 mmol/L × 180.16 g/mol = 900 mg/L = 90 mg/dL. The factor of about 18 between the two units is the molar mass divided by ten — and confusing them is a tenfold-plus error with a clinical consequence (Chapter 20).

Saline. 0.9% w/v NaCl is 9 g/L ÷ 58.44 g/mol = 0.154 mol/L = 154 mmol/L of NaCl. Each NaCl gives two particles in solution (Na^+ and Cl^-), so the theoretical osmolarity is about 308 mOsm/L, close to that of plasma — which is the whole reason this concentration was chosen. (The measured osmolality is a little lower, because the ions do not behave entirely independently.)

2.5.3. Dilution

When you dilute a solution, the amount of solute stays the same:

$$C_1 \times V_1 = C_2 \times V_2$$

To make 500 mL of 70% v/v isopropanol from 99% stock: $V_1 = 70 \times 500 / 99 \approx 354$ mL of stock, made up to 500 mL with water. (In practice, alcohol and water mixtures shrink slightly on mixing, which is one reason a laboratory makes up "to volume" in a volumetric flask rather than adding two measured volumes together.)

A **serial dilution** repeats a fixed step: six successive 1-in-10 dilutions give a final concentration of 10^{-6} of the original. Serial dilutions are how microbiologists count bacteria and how calibration standards are made.

2.5.4. Solutions and solubility

A **solution** is a solute dissolved uniformly in a solvent. The rule of thumb is **like dissolves like**: polar and ionic substances dissolve in polar solvents such as water; non-polar substances dissolve in non-polar solvents such as hexane or oils. Alcohols, with a polar end and a non-polar end, sit between — which is why isopropanol is good at removing both skin oils and water-based soils, and why it can craze or dissolve some plastics.

A solution that holds as much solute as it can at a given temperature is **saturated**. Solubility of most solids rises with temperature; solubility of gases falls with it (warm water holds less oxygen). Exceed the limit and the excess **precipitates** — the mechanism behind limescale in steam generators, crystals in a drip line when two incompatible drugs meet, and kidney stones.

2.6. Chemical reactions

2.6.1. Equations and stoichiometry

A chemical equation is a mass balance. Atoms are conserved, so both sides must have the same number of each. The complete oxidation of glucose — what your cells do, in many small steps, to power everything (Chapter 3) — is:

$$C_6H_{12}O_6 + 6\ O_2 \rightarrow 6\ CO_2 + 6\ H_2O$$

Stoichiometry is the arithmetic that follows. One mole of glucose (180 g) consumes six moles of oxygen (6 × 32 = 192 g) and produces six moles of carbon dioxide (6 × 44 = 264 g) and six of water (108 g). Mass in (372 g) equals mass out (372 g).

Reactions are commonly grouped by what happens:

Type	What happens	Example
Synthesis	Two or more reactants combine	Amino acids joining to form a protein; monomers forming a polymer

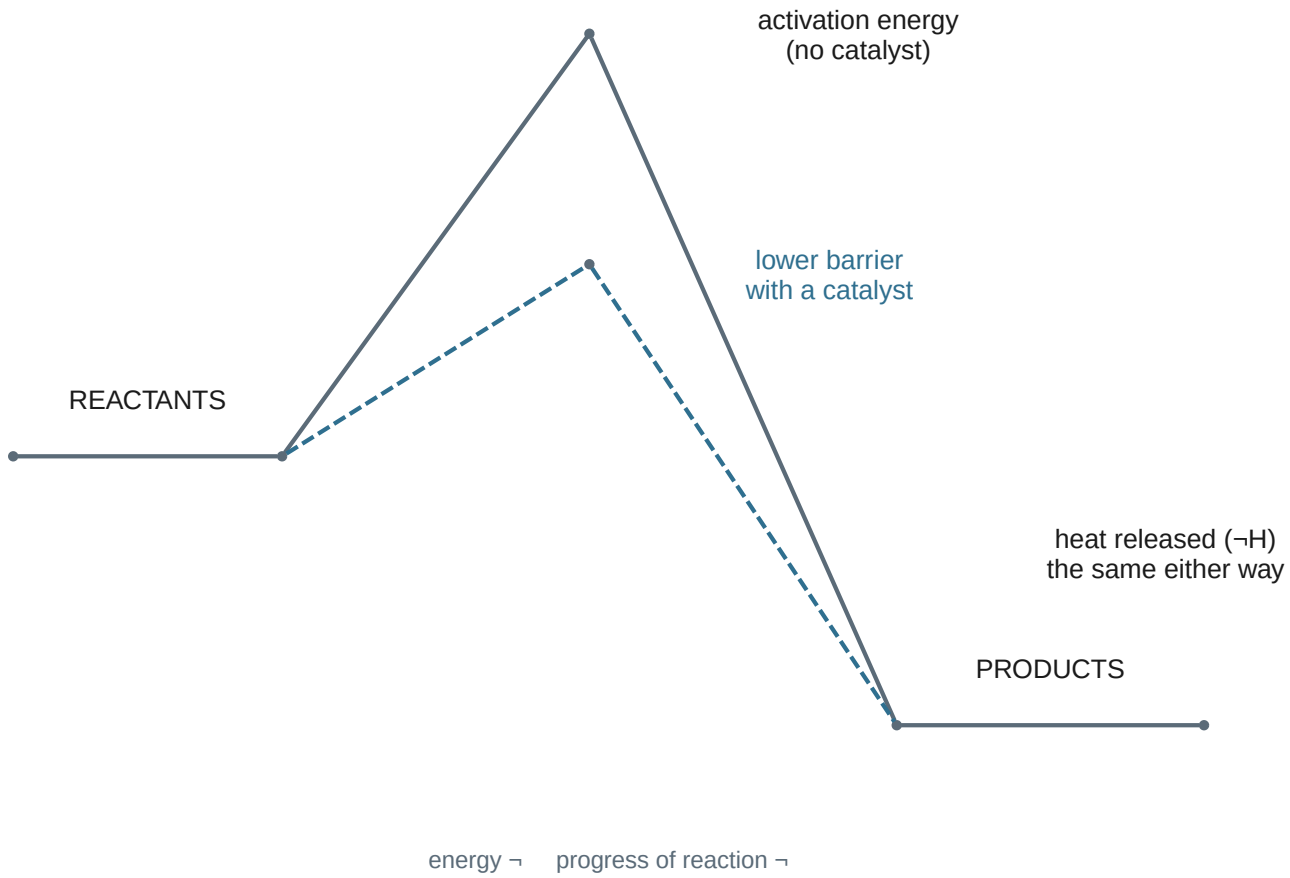
Type	What happens	Example
Decomposition	One substance breaks apart	Hydrogen peroxide → water + oxygen (the fizz on a wound)
Exchange	Partners swap	An acid neutralising a base to form a salt and water
Oxidation–reduction (redox)	Electrons transfer	Rusting; bleach disinfecting; a battery discharging; respiration
Hydrolysis	Water breaks a bond	Digestion of food; degradation of polyesters such as PLA (Chapter 3)
Polymerisation	Small units link into long chains	Curing an epoxy, a silicone, or a dental composite

2.6.2. Energy in reactions

Breaking bonds costs energy; forming bonds releases it. The difference is the reaction's **enthalpy change** (ΔH).

- **Exothermic** reactions release heat (ΔH negative): combustion, the curing of epoxy and polyurethane, dissolving calcium chloride in an instant hot pack, and respiration (about 2,800 kJ per mole of glucose fully oxidised).
- **Endothermic** reactions absorb heat (ΔH positive): dissolving ammonium nitrate in an instant cold pack.

Even an exothermic reaction usually needs a push to start. That push is the **activation energy** — the energy barrier the reactants must climb before they can fall to the products.



A **catalyst** provides an alternative route with a lower barrier, speeding the reaction without being consumed. It changes how **fast** equilibrium is reached, not **where** it lies. Enzymes are biological catalysts (Chapter 3); platinum in a catalytic converter, nickel in hydrogenation, and the tin or platinum compounds that cure silicones are industrial ones.

2.6.3. Rates and kinetics

How fast a reaction goes depends on how often molecules meet with enough energy, in the right orientation. Rate rises with:

- **Concentration** — more molecules, more collisions;
- **Surface area** — powder reacts faster than a lump, a reason fine dusts can explode;
- **Temperature** — a larger fraction of molecules clears the activation barrier;
- **Catalysts** — a lower barrier.

The temperature effect is captured by the **Arrhenius equation**:

$$k = A \times e^{-E_a/RT}$$

where **k** is the rate constant, E_a the activation energy, and **T** the absolute temperature. For many reactions near room temperature, a useful rule of thumb is that the rate roughly **doubles for every 10 °C rise**. That single rule explains why medicines and reagents are refrigerated, why a polymer that is stable for years at room temperature can fail in months in a hot car or a warm cupboard, and why accelerated ageing tests work at all (see the engineering bridge).

Many processes follow **first-order kinetics**: the rate is proportional to how much is left, so a constant **fraction** disappears in each interval. Radioactive decay, the elimination of many drugs (Chapter 17), and the killing of microorganisms by heat or chemicals all behave this way. In sterilisation this is expressed as the **D-value**: the time at a given condition to reduce a microbial population tenfold. Twelve D-values take a million organisms to a one-in-a-million chance of a survivor — the logic behind the sterility assurance level in Chapter 13.

2.6.4. Equilibrium

Many reactions run both ways. When the forward and reverse rates are equal, concentrations stop changing: the system is at **equilibrium**, written with a double arrow (⇌). Equilibrium is dynamic — molecules keep reacting — but the net composition is steady.

Le Chatelier's principle says that if you disturb a system at equilibrium, it shifts to partly counteract the disturbance. Add more of a reactant and more product forms; remove a product and the reaction pulls forward. The body's most important example is the carbon dioxide–bicarbonate system:

$$CO_2 + H_2O \rightleftharpoons H_2CO_3 \rightleftharpoons H^+ + HCO_3^-$$

Breathe out more CO_2 and the equilibrium shifts left, consuming H^+ and raising pH; retain CO_2 and it shifts right, lowering pH. That is how the lungs adjust blood pH within minutes, and why Chapter 3's treatment of buffers starts here.

Another everyday equilibrium is **solubility**: a sparingly soluble salt dissolves until its ions reach a fixed product of concentrations. Push the ion concentrations above it — calcium and oxalate in concentrated urine, calcium carbonate in hard water heated in a boiler — and a solid precipitates.

2.7. Acids, bases, and salts

An **acid** donates hydrogen ions (protons, H^+); a **base** accepts them. **Strong** acids and bases (hydrochloric acid, sodium hydroxide) dissociate completely in water; **weak** ones (acetic acid, ammonia, carbonic acid) only partly, which is what lets them act as buffers. An acid and a base **neutralise** each other to form water and a **salt**.

Acidity is measured on the logarithmic **pH** scale, and the body defends blood pH within a very narrow band using buffers. Both are covered properly in Chapter 3. For a maker, three practical points:

- **Strong acids and alkalis are both corrosive** to tissue, and alkalis are often worse on skin and eyes because they keep penetrating. Oven cleaners, drain cleaners, and some descalers and etchants are in this class.
- **Cleaning chemistry is acid–base chemistry.** Alkaline detergents lift protein and fat; acidic descalers remove mineral deposits; some disinfectants work only within a pH range.
- **Never mix cleaning products.** The classic accidents are chemical: bleach plus an acid releases chlorine gas; bleach plus ammonia releases chloramines.

2.8. Oxidation, reduction, and electrochemistry

2.8.1. Redox

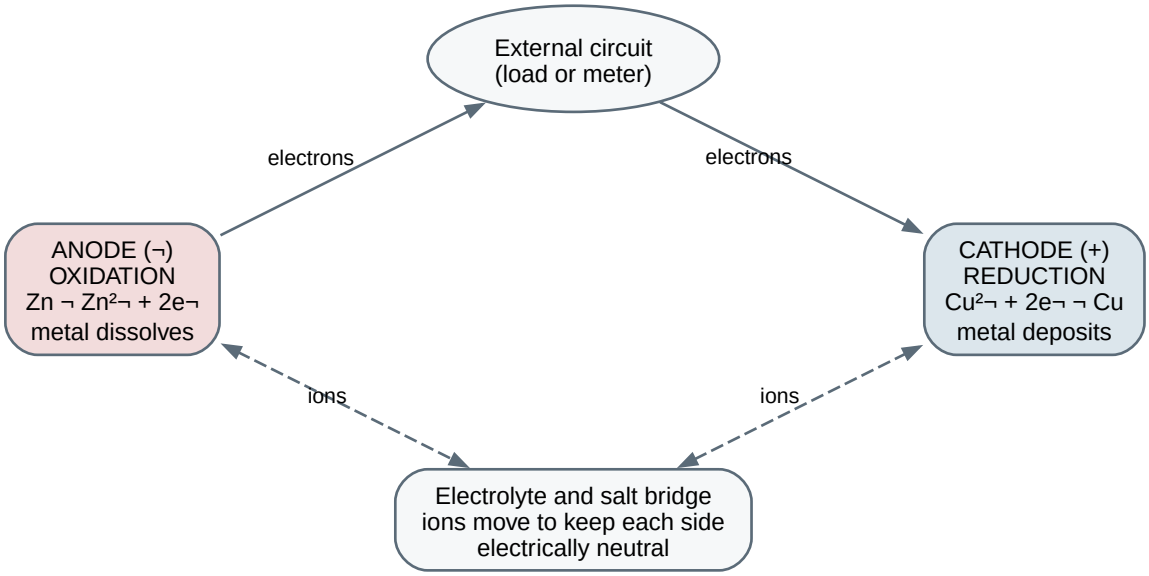
Oxidation is loss of electrons; **reduction** is gain of electrons. ("OIL RIG": oxidation is loss, reduction is gain.) They always happen together, because the electrons must go somewhere. An **oxidising agent** takes electrons and is itself reduced; a **reducing agent** gives them up.

Redox is everywhere in health and engineering:

- **Respiration** oxidises food and reduces oxygen to water (Chapter 3).
- **Disinfection** by sodium hypochlorite (bleach), hydrogen peroxide, peracetic acid, and chlorine dioxide is oxidation of microbial proteins, lipids, and nucleic acids. It is also why these agents corrode metals, bleach fabrics, and degrade some plastics and adhesives.
- **Corrosion** is the oxidation of a metal. Iron rusting, aluminium pitting in chloride, and galvanic attack between dissimilar metals are all redox, treated in Chapter 3.
- **Rechargeable batteries** store energy as a reversible redox reaction.

2.8.2. Electrochemical cells

If you separate the oxidation and the reduction and connect them with a wire, the electron transfer becomes an electric current. That is a **galvanic cell**.



- The **anode** is where oxidation happens; the **cathode** is where reduction happens.
- Electrons flow through the external wire; ions flow through the electrolyte.
- Each half-reaction has a **standard electrode potential**. The cell voltage is the difference: zinc (−0.76 V) against copper (+0.34 V) gives about 1.10 V.

The same physics runs in reverse in **electrolysis**: driving current through a cell forces a reaction that would not happen by itself — electroplating, anodising aluminium, producing chlorine from brine.

Electrochemistry is unusually close to medical devices:

Device or process	Electrochemistry
Glucose meters and CGMs	An enzyme oxidises glucose; the electrode measures the resulting current (amperometry) (Chapter 3)
Blood gas analysers	A glass electrode for pH, a Clark electrode reducing O_2 at a platinum cathode, and a CO_2 electrode based on pH change (Chapter 15)
Ion-selective electrodes	Membrane potentials that depend on one ion: Na^+, K^+, Cl^-, Ca^{2+}
ECG and EEG electrodes	Silver/silver chloride electrodes give a stable, low-noise skin interface (Chapter 7)
Breath alcohol testers	A fuel cell oxidises ethanol and the current is proportional to the amount
Implant batteries	Lithium chemistries chosen for energy density and very low self-discharge over years
Zinc–air hearing aid cells	About 1.4 V, activated by pulling the tab to let oxygen in (Chapter 8)
Corrosion of implants and instruments	Unwanted galvanic cells between dissimilar metals, or at crevices (Chapter 3)

2.8.3. Batteries

A battery is one or more cells packaged together. Its **voltage** comes from the chemistry; its **capacity** (ampere-hours) from how much active material it holds; its **energy** is roughly voltage × capacity.

Chemistry	Nominal cell voltage	Notes
Alkaline (zinc–manganese dioxide)	1.5 V	Primary (single use); voltage sags steadily
Nickel–metal hydride	1.2 V	Rechargeable; robust; self-discharges faster
Lithium primary (e.g. lithium–manganese dioxide coin cells)	~3 V	Long shelf life; coin cells are a serious swallowing hazard (Chapter 8)
Lithium-ion (cobalt, NMC, and similar)	~3.6–3.7 V	High energy density; needs protection circuitry against over-charge, over-discharge, and short circuit
Lithium iron phosphate (LFP)	~3.2 V	Lower energy density; more thermally stable

Chemistry	Nominal cell voltage	Notes
Lead–acid	2.0 V	Heavy; still used in backup supplies and some mobility equipment

Thermal runaway is the failure mode that makes lithium-ion cells a safety topic. Internal damage, over-charging, or overheating starts exothermic decomposition reactions; the heat speeds the reactions (Arrhenius again), and the cell can vent flammable gas and ignite. It is a positive-feedback loop of exactly the kind Chapter 1 warns about. Use cells with protection circuits, charge them with proper chargers, and keep them out of anything that can crush or puncture them (Chapter 14).

2.9. Organic chemistry in one sitting

Organic chemistry is the chemistry of carbon compounds. Carbon makes four bonds, can bond to itself in chains and rings of any length, and forms stable bonds with hydrogen, oxygen, nitrogen, sulfur, and phosphorus. That versatility is why it is the basis of every biological molecule, every drug, and almost every polymer.

2.9.1. Functional groups

The carbon skeleton gives a molecule its shape; **functional groups** — small clusters of atoms attached to it — give it its chemistry. Learn to recognise these and most of biochemistry and polymer science becomes readable.

Group	Structure	Where you meet it
Hydroxyl (alcohol)	–OH	Ethanol and isopropanol disinfectants; sugars; polyvinyl alcohol; makes molecules more water-soluble
Carbonyl — aldehyde	–CHO	Formaldehyde and glutaraldehyde, used as tissue fixatives and (historically) disinfectants; both sensitisers
Carbonyl — ketone	C=O within a chain	Acetone, the workshop solvent; ketone bodies in diabetic ketoacidosis
Carboxyl (acid)	–COOH	Acetic acid; fatty acids; every amino acid; acrylic acid in superabsorbents
Amine	$–NH_2$	Every amino acid; many drugs; epoxy hardeners, which is why they sensitise skin
Ester	–COO–	Fats; aspirin; polyesters such as PET, PLA, and polyester fabric; cyanoacrylate adhesives
Amide	–CONH–	The peptide bond in proteins; nylon; Kevlar
Ether	–O–	Polyethylene glycol; some anaesthetics; epoxy backbones
Thiol	–SH	The amino acid cysteine; disulfide bridges in hair and antibodies
Phosphate	$–PO_4$	ATP; the backbone of DNA and RNA; membrane phospholipids

Group	Structure	Where you meet it
Aromatic ring	Benzene-type ring	Many drugs; polystyrene, polycarbonate, PEEK; phenolic disinfectants
Isocyanate	–N=C=O	Polyurethane foams, coatings, and two-part paints. A potent respiratory sensitiser before it cures

2.9.2. Isomers and chirality

Isomers have the same atoms arranged differently, and can behave completely differently. Ethanol and dimethyl ether are both C_2H_6O; one is a drink, the other a gas used as an aerosol propellant.

A carbon atom bonded to four **different** groups can exist in two mirror-image forms that cannot be superimposed, like left and right hands. These are **enantiomers**, and the molecule is **chiral**. Enantiomers have identical physical properties in a bottle, but the body is built from chiral molecules — its receptors and enzymes are left- or right-handed — so it often treats them differently.

The cautionary example is **thalidomide**, given to pregnant women in the late 1950s and early 1960s: one enantiomer is a sedative, the other is linked to severe birth defects, and the two interconvert in the body, so purifying one would not have made it safe. Today, regulators expect the stereochemistry of a chiral drug to be characterised, and many drugs are sold as a single enantiomer — the "es-" in esomeprazole and the "S-" in S-ketamine mark these.

2.9.3. Polymers, briefly

Polymers are very long molecules built from repeating units (**monomers**): polyethylene from ethylene, nylon from diamines and diacids, silicone from siloxane units, proteins from amino acids, cellulose and cotton from glucose. Their properties depend on chain length, branching, cross-linking, crystallinity, and additives as much as on the monomer. Chapter 3 covers polymer chemistry for makers — glass transition, additives, leachables, degradation — and Chapter 14 covers how each is processed.

2.10. Chemical safety

Everyone who uses chemicals at work in most jurisdictions is covered by a framework that requires hazards to be identified, communicated, and controlled. You should be able to read the documents and symbols of that framework, because a hospital, a university, or a lab will expect it before letting you use their space.

2.10.1. Hazard communication

The **Globally Harmonized System** of classification and labelling (GHS), developed under the United Nations, standardises how chemical hazards are described. It underpins the EU and UK **CLP** regulation and the US OSHA **Hazard Communication Standard**.

GHS pictogram (red diamond)	Meaning, with workshop examples
Flame	Flammable: acetone, isopropanol, ethanol, many aerosols and adhesives
Flame over circle	Oxidiser: hydrogen peroxide (concentrated), some pool chlorine products — feeds a fire
Exploding bomb	Explosive or self-reactive
Gas cylinder	Gas under pressure: oxygen, nitrogen, argon and CO_2 cylinders
Corrosion	Corrosive to skin, eyes, or metals: strong acids and alkalis, some descalers and etchants
Skull and crossbones	Acute toxicity: small amounts can kill
Health hazard (silhouette)	Sensitiser, carcinogen, mutagen, reproductive toxin, organ toxicity: isocyanates, some epoxy components, formaldehyde
Exclamation mark	Irritant, skin sensitiser, harmful, narcotic effects
Environment	Toxic to aquatic life: many solvents, biocides, and metal salts

Every hazardous product comes with a **Safety Data Sheet** (SDS) in a standard sixteen-section format: identity, hazards, composition, first aid, firefighting, spill response, handling and storage, exposure controls and PPE, properties, stability and reactivity, toxicology, ecology, disposal, transport, regulations, and other information. Section 8 (exposure controls) and section 10 (stability and reactivity — what not to mix it with) are the ones people most often skip and most often need.

In the UK the **COSHH** regulations (Control of Substances Hazardous to Health) require a written assessment of each hazardous substance used at work and how exposure is controlled; in the US the OSHA standard requires a written hazard communication programme, labels, SDSs, and training. The principle is the same: identify, assess, control, review.

2.10.2. Controls, in order of preference

The **hierarchy of controls** runs: **eliminate** the substance, **substitute** a safer one, **engineer** the exposure away (enclosure, local exhaust ventilation, a fume hood), apply **administrative** controls (procedures, training, limited quantities), and finally **PPE** (gloves, goggles, respirators). It is the same hierarchy as risk control in Chapter 19, and for the same reason: controls near the top do not depend on someone remembering.

2.10.3. The workshop's own chemistry

Substance	Hazard	Sensible control
Acetone, isopropanol, ethanol	Highly flammable; vapour is heavier than air and travels to ignition sources (acetone's flash point is well below room temperature)	Small quantities, closed containers, ventilation, no ignition sources, metal safety cans for storage

Substance	Hazard	Sensible control
Two-part epoxy	Amine hardeners and uncured resin are skin **sensitisers**: once sensitised, a small exposure provokes dermatitis for life	Nitrile gloves, no skin contact, clean tools away from bare hands
Polyurethane foams, coatings, two-part paints	**Isocyanates** are respiratory sensitisers and a leading cause of occupational asthma	Avoid spraying; local exhaust ventilation; supplied-air protection for spraying is a professional job
Solder flux (rosin)	Fume from heated rosin (colophony) is a respiratory sensitiser	Fume extraction at the iron; lead-free solder still produces flux fume
Cyanoacrylate adhesive	Bonds skin instantly; vapour irritates eyes and airways	Ventilation; acetone-based debonder to hand
Bleach (sodium hypochlorite)	Corrosive; reacts with acids (chlorine gas) and ammonia (chloramines); corrodes metals	Never mix; rinse metal parts after contact
Glass etching and some metal pickling	Hydrofluoric acid penetrates skin painlessly and binds calcium in tissue; small burns can be fatal	Do not use outside a properly equipped laboratory. Use mechanical frosting or safer etching products instead
Resins for printing and casting	Many are skin sensitisers before cure; uncured residue remains on parts	Gloves, complete post-curing, and treating uncured waste as hazardous



Sensitisation is the hazard most often underestimated by experienced makers. An irritant hurts and then stops. A **sensitiser** changes your immune system so that every later exposure — however small — provokes a reaction. It is cumulative, often permanent, and can end your use of an entire class of materials. Epoxies, isocyanates, rosin fume, acrylates, nickel, latex, formaldehyde, and glutaraldehyde are all sensitisers. Treat "I've used it for years without gloves" as a warning, not a reassurance.

2.11. Engineering bridge: chemical arithmetic

2.11.1. Unit discipline with moles

Carry units through every step and let them cancel. To find the mass of potassium chloride needed for 250 mL of a 40 mmol/L solution:

0.250 L × 40 mmol/L = 10 mmol; 10 mmol × 74.55 mg/mmol = 745.5 mg ≈ 0.75 g

The same discipline catches nonsense. If a unit does not cancel to the one you want, the calculation is wrong, however plausible the number looks. That is the chemistry version of the units-in-the-type rule in Appendix C.

2.11.2. Ideal gas arithmetic for cylinders

A small medical oxygen cylinder with an internal volume of 2 L, filled to 200 bar, holds roughly 2 × 200 = 400 L of gas at atmospheric pressure (real gases deviate somewhat at high pressure, so the true figure differs a little, which is why cylinders are labelled with their actual content). At 2 L/min, that is about three hours. The same arithmetic tells you how much stored energy is in a pressurised vessel and why a cylinder with a broken valve becomes a projectile — and why oxygen enrichment turns ordinary materials into fuel.

2.11.3. Accelerated ageing and the Arrhenius rule

Shelf-life claims for packaged devices are often supported by **accelerated ageing**: storing samples at an elevated temperature to simulate a longer time at room temperature. A standard approach for sterile packaging (ASTM F1980) uses an **ageing factor**:

AAF = $\mathbf{Q_{10}}$ raised to the power $\mathbf{(T_{AA} - T_{RT}) / 10}$

where $\mathbf{Q_{10}}$ is the rate increase for a 10 °C rise (2 is a common conservative assumption), $\mathbf{T_{AA}}$ the ageing temperature, and $\mathbf{T_{RT}}$ the real storage temperature. At 55 °C against 25 °C, AAF = 2^3 = 8, so one year of real time is simulated in about 365 ÷ 8 ≈ 46 days.

The limits are chemistry too. The temperature must not cross a transition the material would never meet in service — melting, a glass transition, a change in reaction mechanism — or the test simulates something that will not happen. And accelerated ageing supports a claim; it does not replace real-time ageing data, which is normally collected in parallel.

2.11.4. Faraday's law: how much metal moves

The mass of material deposited or dissolved electrochemically is set by the charge passed:

m = (Q × M) / (n × F)

where **Q** is charge in coulombs, **M** molar mass, **n** electrons per ion, and **F** the Faraday constant, 96,485 C/mol.

- **Plating:** 1 A for 10 minutes is 600 C. For silver (M = 107.9 g/mol, n = 1): 600 × 107.9 ÷ 96,485 ≈ 0.67 g of silver.
- **Corrosion:** a corrosion current of just 1 µA flowing for a year is about 31.6 C. For iron dissolving as Fe^{2+} (M = 55.85, n = 2): 31.6 × 55.85 ÷ (2 × 96,485) ≈ 9 mg. Concentrated at a crevice or a pit rather than spread over a surface, a few milligrams is enough to perforate a thin wall or loosen a fastener — and a small galvanic couple can easily drive more current than that.

2.11.5. Battery energy

Energy (watt-hours) ≈ nominal voltage × capacity (ampere-hours). A 3.7 V, 2,000 mAh lithium-ion cell stores about 7.4 Wh, or about 27 kJ. That is enough, released in seconds during a thermal runaway, to start a fire. Transport regulations place lithium cells and packs in hazard classes by watt-hour rating, which is why a device you build has a shipping and travel consequence as well as a design one (Chapter 14).

2.12. The field in practice: what chemists do, and the tools they use

Chemistry is not one job. Each branch below touches health, and each has its own questions, instruments, and software. The point of this section is that when someone says "send it to the chemistry lab", you know which lab, what it will probably do, and what it will need from you.

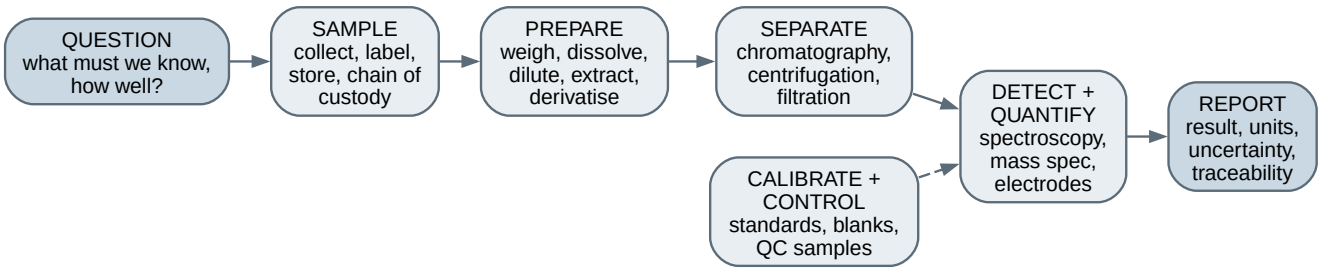
2.12.1. The branches that touch health

Branch	Real-world uses in health	Typical tools
Analytical chemistry	Identifying what something is and measuring how much: an unknown residue on a device, the composition of a polymer, metal ions leaching from an implant, a contaminant in water	Chromatography (HPLC, GC), mass spectrometry, spectroscopy (UV-Vis, FTIR, Raman, NMR), elemental analysis (ICP-MS, XRF), titration
Clinical chemistry and biochemistry	Electrolytes, glucose, kidney and liver markers, hormones, drug levels, blood gases — thousands of patient samples a day (Chapter 15)	Automated analysers, ion-selective electrodes, photometry, immunoassay, LC-MS/MS, LIMS
Pharmaceutical and medicinal chemistry	Designing drug molecules, making them at scale, turning them into a medicine (formulation), and proving every batch is right (quality control)	Synthesis labs, computational chemistry, HPLC and LC-MS for purity, dissolution testers, stability chambers, XRD for crystal form
Materials and polymer chemistry	Choosing and qualifying device materials; extractables and leachables; degradation, coatings, adhesives, failure analysis (Chapter 3)	FTIR, DSC and TGA, GPC (molecular weight), SEM with EDS, contact angle, mechanical testing, accelerated ageing chambers
Electrochemistry	Sensors, batteries, corrosion, electroplating, neural electrodes, implant power	Potentiostats, impedance analysers, battery cyclers, salt-spray cabinets
Environmental and water chemistry	Water for dialysis and sterile services, Legionella control, wastewater from hospitals, air quality in theatres and labs	Conductivity and TOC analysers, chlorine and pH test kits, ICP-MS, particle counters
Toxicology and forensic chemistry	Overdose and poisoning screens; therapeutic drug monitoring; post-mortem analysis	Immunoassay screens confirmed by GC-MS or LC-MS/MS
Food chemistry and nutrition	Nutritional content, enteral feeds, texture-modified diets (Chapter 8), allergens	Proximate analysis, chromatography, rheometers

Branch	Real-world uses in health	Typical tools
Disinfection and sterilisation chemistry	Alcohols, chlorine, peroxide, peracetic acid, aldehydes, ethylene oxide; material compatibility; residues (Chapter 13)	Test strips and titrations for concentration, chemical and biological indicators, residual analysis
Radiochemistry and radiopharmacy	Making and checking radiopharmaceuticals for PET and SPECT within their half-lives (Chapter 16)	Cyclotrons, hot cells, automated synthesis modules, dose calibrators, radio-TLC and radio-HPLC

2.12.2. The analytical workflow

Whatever the branch, an analysis follows the same chain, and errors can enter at every link.



The quality of an analytical result is a property of that whole chain, not of the instrument at the end of it — Chapter 15's point about trustworthy measurement in a single diagram.

2.12.3. Measuring and preparing

Tool	What it does and why it matters
Analytical balance	Weighs to 0.1 mg (some to 0.01 mg) inside a draught shield. Sensitive to vibration, air currents, static, and temperature; checked daily with certified weights
Volumetric glassware	Flasks and pipettes calibrated to contain or deliver one volume precisely (a class A 100 mL flask is good to about ±0.08 mL). Usually borosilicate glass
Graduated cylinders and beakers	Convenient, not accurate. Beaker markings are approximate
Burettes	Deliver measured volumes drop by drop for titration
Micropipettes	Air-displacement pipettes from below 1 µL to 10 mL with disposable tips. Accuracy depends on technique, tip fit, and calibration (ISO 8655); a common source of laboratory error
pH meter	A glass electrode whose voltage changes about 59 mV per pH unit at 25 °C; calibrated against standard buffers before use

Tool	What it does and why it matters
Conductivity meter	Measures how well a solution carries current, a proxy for total dissolved ions. Used for water purity, including water for dialysis and steam generation
Water purification	Deionisation, reverse osmosis, and distillation produce reagent-grade water; its quality is specified and monitored

2.12.4. Separating

Technique	What it does
Filtration	Removes particles; membrane filters at 0.2 µm remove most bacteria, which is how heat-sensitive solutions are sterilised
Centrifugation	Spins samples to separate by density: blood into plasma and cells, cells from culture medium. Force is quoted as relative centrifugal force (RCF, × g), not rpm, because it depends on rotor radius
Distillation	Separates liquids by boiling point; purifies solvents and water
Extraction	Moves a compound from one phase to another — between immiscible solvents, or onto a solid sorbent (solid-phase extraction) — to clean up or concentrate a sample
Thin-layer chromatography (TLC)	Cheap, fast separation on a coated plate; quick purity checks
High-performance liquid chromatography (HPLC)	Pumps a liquid sample at high pressure through a packed column; compounds emerge at different times and are detected, usually by UV absorbance or mass spectrometry. The workhorse of pharmaceutical analysis
Gas chromatography (GC)	Separates volatile compounds in a heated column with a carrier gas. Solvents, alcohols, anaesthetic gases, residual ethylene oxide
Gel electrophoresis	Separates charged molecules — DNA, proteins — by size in an electric field (Chapter 4)

2.12.5. Identifying and quantifying

Instrument	What it measures
UV-visible spectrophotometer	How much light of each wavelength a solution absorbs; concentration from the Beer–Lambert law ($A = \varepsilon c l$). Clinical analysers and pulse oximeters use the same principle
FTIR spectrometer	Infrared absorption by bond vibrations — a fingerprint of functional groups. With an ATR accessory it identifies an unknown plastic, coating, or residue in a minute
Raman spectrometer	Light scattered with a shift in energy, also from bond vibrations; works through glass and packaging and is used to verify incoming raw materials

Instrument	What it measures
NMR spectrometer	Resonance of atomic nuclei in a strong magnetic field reveals how atoms are connected — the definitive tool for molecular structure. The physics is the same as MRI (Chapter 16)
Mass spectrometer (MS)	Ionises molecules and sorts them by mass-to-charge ratio. Paired with chromatography (GC-MS, LC-MS/MS) it identifies and quantifies trace compounds — drugs, toxins, leachables
ICP-MS and ICP-OES	A plasma at several thousand kelvin breaks a sample into atoms; their mass or emitted light identifies elements down to trace levels. Metal leaching from implants and instruments
X-ray fluorescence (XRF)	Elemental composition, non-destructively. Handheld XRF guns verify which grade of stainless steel or alloy a part actually is
X-ray diffraction (XRD)	Crystal structure: which crystal form a drug is in, which phases are present in a metal or ceramic
Titration	Adds a reagent of known concentration until a reaction is complete; still the reference method for many assays and for checking disinfectant strength
Karl Fischer titration	Water content, down to parts per million; critical for moisture-sensitive drugs and polymers
DSC and TGA (thermal analysis)	Differential scanning calorimetry measures heat flow on heating — glass transition, melting, crystallinity. Thermogravimetric analysis measures mass loss — moisture, filler content, decomposition (Chapter 3)
SEM with EDS	Electron microscopy of surfaces with elemental mapping; failure analysis of fractures, corrosion, and contamination
Contact angle goniometer	The angle a droplet makes with a surface, a measure of surface energy; predicts adhesion and wettability

2.12.6. Handling, containment, automation, and data

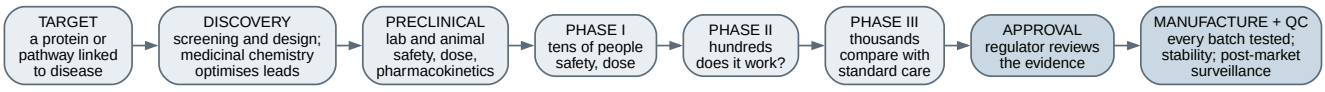
Tool	Purpose
Fume hood	Draws air away from the user to capture vapours and fumes; face velocity commonly around 0.4–0.5 m/s. It protects the user, not the sample, and it is **not** a biological safety cabinet (Chapter 4)
Glovebox	Sealed enclosure with built-in gloves: an inert atmosphere for air-sensitive chemistry, or containment for hazardous material
Schlenk line	Glass manifold for handling air- and moisture-sensitive reagents under vacuum and inert gas
Safety equipment	Eyewash stations, safety showers, spill kits, flammable-storage cabinets, gas detectors
Autosamplers and liquid handlers	Robots that load samples and pipette reagents, allowing hundreds of analyses unattended

Tool	Purpose
LIMS (laboratory information management system)	Tracks samples, tests, results, instruments, and approvals; the audit trail of a regulated lab
ELN (electronic laboratory notebook)	The structured, time-stamped record of what was done. Data-integrity expectations (often summarised as ALCOA+: attributable, legible, contemporaneous, original, accurate, and more) apply (Chapter 20)

2.12.7. Software

Software	What it is for
ChemDraw, Ketcher	Drawing chemical structures and reactions; Ketcher is open source
PubChem, ChemSpider	Free databases of compounds: structures, properties, hazards, literature
NIST Chemistry WebBook	Free reference spectra and thermodynamic data
RDKit, Open Babel	Open-source cheminformatics libraries: reading structures, computing properties, searching for similar molecules. RDKit is the standard Python toolkit
Avogadro, Jmol, PyMOL	Building and viewing molecules in three dimensions
Gaussian, ORCA, and similar	Quantum-chemistry programs that calculate molecular structure and energy from first principles; specialist tools
Instrument software and chromatography data systems	Control instruments, process spectra and chromatograms, and enforce audit trails in regulated labs

2.12.8. From molecule to medicine



Bringing a new drug from target to approval commonly takes ten to fifteen years, and most candidates fail along the way. Chemistry runs through every stage: designing and making the molecule, choosing its salt and crystal form, formulating it as a tablet, injection, or inhaler, scaling the synthesis, and testing every batch for identity, purity, strength, and stability. Medical devices that deliver drugs — inhalers, autoinjectors, pumps — inherit that chemistry (Chapter 17).

2.13. Devices, software, and training technology

2.13.1. Clinical devices built on general chemistry

Device	Chemical principle	Note
Blood gas and electrolyte analysers	Ion-selective electrodes, Clark oxygen electrode, partial pressures, Henry's law	Results feed directly into ventilation and fluid decisions (Chapter 9)
Capnographs	Infrared absorption by CO_2 in exhaled breath	The same physics as FTIR, specialised to one molecule
Oxygen concentrators	Zeolite beads adsorb nitrogen under pressure and release it when depressurised (pressure-swing adsorption), leaving oxygen-enriched air	Oxygen enrichment raises fire risk sharply; smoking and open flames are banned near them
Anaesthetic vaporisers	Vapour pressure of volatile agents, compensated for temperature	Agent-specific; filling the wrong agent is a recognised hazard
Breath alcohol testers	Fuel-cell oxidation of ethanol, or infrared absorption	Calibration and drift are measurement-quality problems (Chapter 15)
Dialysis water treatment	Reverse osmosis, deionisation, conductivity monitoring	Trace contaminants in dialysis water have harmed patients; water quality is specified and monitored (Chapter 10)
Sterilisation chemical indicators	Inks that change colour when exposed to a condition such as steam at temperature	Show that a process condition was met, not that the load is sterile (Chapter 13)
Instant hot and cold packs	Exothermic and endothermic dissolution	Surface temperature limits matter on skin with poor sensation (Chapter 5)
Hyperbaric chambers	Henry's law: more oxygen dissolves at higher pressure	Fire and pressure hazards; every object taken inside is controlled

2.13.2. Software

Software	Does	Note for a maker or developer
PubChem, SDS repositories	Identity, hazards, and properties of any substance you use	The first stop before bringing a chemical into a workshop
RDKit, Open Babel	Programmatic handling of structures and properties	Approachable from Python; useful for anyone building chemistry-aware tools
Unit-aware calculation libraries (e.g. Pint in Python)	Carry units through calculations and refuse mismatches	The software form of the unit discipline in the engineering bridge
LIMS and ELN platforms	Sample tracking and records in regulated labs	Interfaces and integrations are a real need; data integrity rules apply (Chapter 20)

Software	Does	Note for a maker or developer
Virtual laboratory simulations	Practise titrations, dilutions, and instrument operation on screen	Used in teaching before students touch real reagents

2.13.3. Training and simulation

Training technology	Teaches	Form
Laboratory skills training	Pipetting, weighing, dilution, titration, good laboratory practice	Practical classes with water and dyes before real reagents; gravimetric pipette checks
Chemical spill and exposure drills	Spill kits, eyewash and shower use, evacuation	Scenario training in labs, pharmacies, sterile services, and workshops
COSHH and hazard-communication training	Reading labels and SDSs, assessing and controlling exposure	Mandatory in most workplaces that use hazardous substances
Molecular models	Bonding, geometry, chirality	Ball-and-stick kits; physical models made by printing, turning, laser-cutting, or bending wire (Chapter 3)
Instrument teaching rigs	How a spectrophotometer, electrode, or chromatograph works	Simplified, open instruments built for teaching — see below

2.13.4. Where a maker contributes

Contribution	Technologies	Boundary
Teaching instruments: LED colorimeters, simple spectrophotometers, conductivity and pH logging rigs	Electronics, optics, machined or moulded housings, software	Teaching and demonstration only. No patient samples, no clinical claims
Open-source potentiostats and electrochemistry demonstrators	Electronics and firmware; published open designs exist for teaching and research	Research and teaching; never connected to a person

Contribution	Technologies	Boundary
Laboratory fixtures, racks, and holders	Machining, sheet metal, moulding, laser-cut acrylic; materials chosen for the solvents and disinfectants they will meet	Agree chemical compatibility and cleaning with the lab first
Scientific glassware repair and custom apparatus	Scientific glassblowing in borosilicate	A skilled trade; pressure and vacuum glassware has its own safety requirements
Lab automation	Open liquid-handling platforms, custom sample handling, scripts that move data from instruments to LIMS	In a regulated lab, validated under its quality system
Workshop chemical safety	Fume extraction, solvent storage, labelled dispensing, written assessments	Your own workshop first — and a useful demonstration that you take it seriously



An LED colorimeter is the best first build in this chapter. An LED, a photodiode, a cuvette holder machined or moulded to keep the light path fixed, a microcontroller, and some software give you a working Beer–Lambert instrument. Calibrate it with a dilution series of food dye and you have touched concentration, dilution, absorbance, calibration, noise, and uncertainty in one afternoon — the core of laboratory medicine, with no hazardous chemicals and no patients.

2.14. Worked example: "can you make us a dispenser for the chlorine tablets?"

A ward manager asks whether you could make something to help staff make up chlorine disinfectant from tablets. The current method is a bucket, a jug, and a laminated sheet.

The figures below are invented for teaching; real products and local policy differ.

Understand the chemistry first. The tablets release **available chlorine** when dissolved — here, say, 1.5 g per tablet. Local policy (in this invented case) asks for 1,000 ppm available chlorine for routine use. In dilute water solutions 1 ppm ≈ 1 mg/L, so 1,000 ppm is 1,000 mg/L = 1 g/L. One tablet therefore needs:

1.5 g ÷ 1 g/L = 1.5 L of water

Now the questions the chemistry raises.

- **Accuracy.** If staff use a 2 L jug "about two-thirds full" the strength could plausibly be 20% low or high. Too weak may not disinfect; too strong corrodes surfaces, damages equipment, and irritates staff. The useful product is a container with an **unambiguous fill mark** for exactly 1.5 L and an indication of how many tablets — a measurement problem, not a mixing one.
- **Stability.** Hypochlorite solutions decay; local policy typically sets a time after which a made-up

solution must be discarded. The container should show when it was made up — a writable label area or a dated tag — because rate laws apply to buckets too.

- **Materials.** Chlorine solutions corrode many metals and degrade some plastics and rubbers. Polypropylene and HDPE resist them well; mild steel and aluminium do not, and even stainless steel can pit in chloride over time. A machined aluminium fill gauge would be the wrong part, however elegant.
- **Incompatibility.** The container must never be used for acid descalers or ammonia-based cleaners. Distinct colour, shape, and labelling reduce the chance — redundant coding again (Chapter 7).
- **Exposure.** Making up solution releases some chlorine vapour; tablets are an irritant and an oxidiser. The ward's COSHH assessment, PPE, and ventilation apply, and your design must not require anyone to put a hand in the solution.

The candidate solutions span technologies. A moulded polypropylene container with a moulded-in fill line; a laser-etched mark on a commercially available HDPE jug; a sewn or moulded holder that keeps the right jug, tablets, and label together; or — perhaps best — a proposal to switch to a pre-made solution or pre-measured product so that no one mixes anything. That last option belongs to the infection prevention team and procurement, not to you, but noticing it is your job (Chapter 18).

And the measurement. Whatever is built, the result should be checked: chlorine test strips or a titration on a few made-up batches show whether the intended concentration is actually achieved. A design that has never been measured has not been shown to work (Chapter 19).

2.15. Workshop exercise

1. **Build a calibration curve.** Make a serial dilution of food colouring (for example, five steps of 1 in 2). Measure each with a phone camera against a white background or, better, a simple LED-and-photodiode colorimeter. Plot absorbance against concentration and fit a line. Measure an "unknown" made up by someone else and estimate its concentration. Record how you would estimate the uncertainty.
2. **Do the clinical arithmetic.** Calculate, with units shown at every step: the mmol/L of NaCl in 0.9% saline; the mmol/L of glucose in 5% dextrose; the mg/dL equivalent of a glucose of 7.0 mmol/L; the volume of 99% isopropanol needed for 1 L of 70% v/v.
3. **Make a cell and watch corrosion.** Put a strip of zinc (a galvanised washer) and a strip of copper into salt water or a lemon, measure the voltage with a multimeter, and compare it with the standard-potential prediction. Separately, place a stainless screw in contact with an aluminium plate in salt water and another stainless screw alone; inspect both after a week. Wear eye protection; nothing here is hazardous beyond salt water.
4. **Identify plastics without a lab.** Collect offcuts of known plastics. Test which float in water and which sink, then in a saturated salt solution and in diluted alcohol. Compare with published densities and with recycling codes. Write down what this method can and cannot distinguish — and what an FTIR would add.
5. **Audit your own workshop.** Pick five substances you use (a solvent, an adhesive, a flux, a cleaner, a resin or paint). Find each SDS, record the GHS pictograms, the section 8 controls, and

the section 10 incompatibilities, and write a one-page COSHH-style assessment. Change one thing as a result.

6. **Plan an ageing test.** For a packaged item you might make, calculate how long accelerated ageing at 50 °C and at 60 °C would take to simulate two years at 25 °C with Q_{10} = 2. List two material transitions that would make the higher temperature invalid.

2.16. Questions

1. What defines an element, and what changes when an atom becomes an ion? What changes when it becomes a different isotope?
2. Why do sodium and potassium behave alike chemically?
3. Why can fluorine-18 be used to trace glucose metabolism without changing it?
4. Using electronegativity, explain why oil and water do not mix.
5. Convert a blood glucose of 9.0 mmol/L to mg/dL, showing the units.
6. Why is 0.9% the concentration chosen for "normal" saline?
7. Calculate the oxygen partial pressure in dry air at sea level, and explain why it is lower in the trachea.
8. Explain, in terms of activation energy, what a catalyst does and does not change.
9. A reaction rate roughly doubles for each 10 °C rise. Why does this matter for storing a medicine or a polymer part?
10. Explain Le Chatelier's principle using the CO_2–bicarbonate equilibrium.
11. In a corroding galvanic couple, which metal is the anode and what happens to it?
12. Why is a lithium-ion cell a fire hazard, and what feedback loop drives thermal runaway?
13. What is a chiral molecule, and why can two enantiomers of a drug have different effects?
14. Name the instrument you would use to: identify an unknown plastic; check which grade of stainless steel a part is; measure trace nickel leaching into saline; confirm a drug's purity.
15. What is the difference between an irritant and a sensitiser, and why does the difference matter in a workshop?

2.17. Answers and reasoning

1. An element is defined by its **atomic number** — the number of protons. An ion has gained or lost **electrons**, so it carries a net charge but is the same element. A different isotope has a different number of **neutrons**: the same element and essentially the same chemistry, but a different mass and possibly an unstable nucleus.
2. They are in the same group (group 1), so each has one valence electron, which it readily loses to form a +1 ion. Same outer-electron arrangement, similar chemistry.
3. Because isotopes of an element are chemically almost identical, a radioactive atom in a molecule behaves chemically like the stable one. In FDG the fluorine replaces a hydroxyl group rather than being a different isotope of an existing atom, but the molecule is close enough to

glucose to be taken up and trapped like it, while the fluorine-18 emits detectable positrons.

4. Water's O–H bonds are strongly polar, so water molecules attract each other through hydrogen bonds. The C–H bonds of oils are nearly non-polar, so oil cannot join that network. Water molecules maximise their hydrogen bonding by excluding the oil — the hydrophobic effect (Chapter 3).
5. 9.0 mmol/L × 180.16 mg/mmol = 1,621 mg/L; ÷ 10 to convert per litre to per decilitre ≈ 162 mg/dL.
6. 9 g/L ÷ 58.44 g/mol ≈ 154 mmol/L NaCl, giving about 308 mOsm/L of particles — close to plasma osmolality, so cells neither swell nor shrink. (Measured osmolality is slightly lower because the ions are not fully independent.)
7. 0.209 × 760 ≈ 159 mmHg. In the trachea, inspired air is saturated with water vapour at 37 °C (about 47 mmHg), which dilutes the other gases: (760 − 47) × 0.209 ≈ 149 mmHg.
8. A catalyst provides a route with a **lower activation energy**, so more molecules have enough energy to react and the rate rises. It does **not** change the energy difference between reactants and products, the position of equilibrium, or the overall stoichiometry, and it is not consumed.
9. Degradation reactions run faster when warm: a medicine stored 10 °C too warm may lose potency roughly twice as fast, and a polymer may age much faster in a hot vehicle or near a radiator than its room-temperature data suggest. It is also the basis of accelerated ageing tests — and their limits.
10. If a system at equilibrium is disturbed, it shifts to partly counteract the disturbance. Remove CO_2 (faster breathing) and the reactions shift left, consuming H^+ and raising pH; retain CO_2 and they shift right, producing H^+ and lowering pH.
11. The more active (more easily oxidised) metal is the **anode**. It is **oxidised**: it loses electrons and dissolves as ions. The less active metal is the cathode and is protected.
12. The cell stores a lot of energy in a flammable package. Damage, over-charging, or heat starts exothermic decomposition reactions; the heat raises the temperature, which speeds the reactions further (Arrhenius), releasing more heat — a positive-feedback loop that can end in venting and fire.
13. A chiral molecule exists in two non-superimposable mirror-image forms. Receptors and enzymes are themselves chiral, so they can bind one enantiomer and not the other, or respond differently to each — one may be therapeutic and the other inactive or harmful.
14. **Unknown plastic:** FTIR (with ATR). **Stainless grade:** XRF. **Trace nickel in saline:** ICP-MS (or ICP-OES). **Drug purity:** HPLC, typically with UV or mass-spectrometric detection.
15. An **irritant** causes local damage or inflammation that stops when exposure stops. A **sensitiser** primes the immune system so that later exposures — even tiny ones — provoke a reaction, often for life. Workshop staples such as epoxy hardeners, isocyanates, rosin flux fume, acrylates, and nickel are sensitisers, so habitual low-level exposure can end your ability to use them.

2.18. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Solutions — concentration, dilution, osmolarity and chlorine from tablets.
- Chemistry — molar mass, stoichiometry, cell voltages and Faraday mass.
- Spectroscopy — calibration curves and Beer–Lambert.
- Thermochemistry — reaction energy.
- Heart, vessels and oxygen — partial pressures and the gas laws.
- Batch, tolerance and shelf life — accelerated-ageing time (ASTM F1980).

2.19. Video companion and reading

- **Core reading:** OpenStax *Chemistry 2e* [66] — free and complete. Read selectively: chapters 2 (atoms, molecules, and ions), 3 (composition of substances and solutions), 4 (stoichiometry), 9 (gases), 12 (kinetics), 13 (equilibrium), 17 (electrochemistry), and 20 (organic chemistry). Then OpenStax *Anatomy and Physiology* [1] sections 2.1–2.3 to see the same ideas applied to the body.
- **Core video:** the Crash Course **Chemistry** series [67]. Atoms, the periodic table, stoichiometry, gases, equilibrium, and redox map directly onto this chapter.
- **Reference:** PubChem [68] for any substance you meet, and the NIST Chemistry WebBook [69] for spectra and physical data.
- **Safety:** the GHS documentation [70] for pictograms and classifications; then the SDS for every substance in your own workshop.

2.20. Exit criteria

- ☐ Explain atomic number, isotopes, ions, and valence electrons, and use the periodic table to predict which ion sodium, calcium, and chlorine form.
- ☐ Convert between mmol/L and mg/dL, and between % w/v and mmol/L, with units shown at every step.
- ☐ Perform a dilution calculation and a serial dilution, and build a calibration curve from a dilution series.
- ☐ Calculate a partial pressure and explain how the gas laws apply to breathing and to gas cylinders.
- ☐ Explain activation energy, catalysts, the 10 °C rule, and first-order kinetics, and apply them to shelf life and sterilisation.
- ☐ Identify anode and cathode in a galvanic couple, and estimate a corrosion or plating mass with Faraday's law.
- ☐ Recognise the major functional groups and explain chirality with an example.
- ☐ For five substances in your workshop, read the SDS and write an exposure assessment.
- ☐ Name the branch of chemistry and the instrument you would ask for to answer five different practical questions.

2.21. What this chapter does not establish

- **You are not a chemist.** This chapter makes you literate in chemistry's language and arithmetic. Synthesis, method development, and interpretation of analytical results are professional skills.
- **You cannot validate an analytical method.** A home-built colorimeter that produces a straight calibration line is a teaching instrument. Measurement for clinical or regulatory purposes requires validated methods, controls, and accredited laboratories (Chapter 15).
- **You cannot certify a chemical safe to use.** Reading an SDS is necessary; it is not a risk assessment for someone else's workplace. The organisation's safety and infection prevention teams own those.
- **Accelerated ageing is not shelf-life proof.** The arithmetic supports a claim only within the limits of the material's chemistry and alongside real-time data.
- **Nothing here licenses handling hazardous substances you have not been trained and authorised to use** — hydrofluoric acid, concentrated oxidisers, radioactive material, and compressed oxygen above all.

Chapter 3. Chemistry and molecular biology for the maker

Weeks 4–5 · approximately 8–12 hours

3.1. Why this chapter exists

Chapter 1 told you that living tissue "needs a suitable chemical environment" and left it there. This chapter is that sentence unpacked, because the chemistry turns out to be where a surprising amount of device engineering actually lives. It assumes the general chemistry of Chapter 2 — atoms, the periodic table, moles and concentrations, reactions, rates, equilibrium, and redox — and applies it to the molecules of life and the materials you build with.

Three reasons a maker needs it.

First, because materials are chemistry. Whatever you make it from — a moulded or machined polymer, a printed part, cast silicone, stainless steel, glass, a woven strap — it has a composition, a processing history, additives or finishes, and a surface. What comes **out** of it into tissue — **leachables** — is a chemical question with a regulatory standard attached. So is what a disinfectant does **to** it.

Second, because measurement is chemistry. A glucose meter is an enzyme electrode. A blood gas analyser is a set of ion-selective electrodes. A pregnancy test is an immunoassay on nitrocellulose. If you ever want to understand how a diagnostic device works — Chapter 15 — you need the chemistry underneath it.

Third, because the body defends its chemistry hard, and the narrowness of the ranges it defends tells you how consequential small changes are.



This is not a chemistry course — Chapter 2 is the general foundation. This chapter is the subset that keeps recurring in device work, chosen on exactly that criterion. Where something matters enormously in medicine and rarely in making — detailed metabolic pathways, organic synthesis — it is named and skipped.

3.2. Atoms, bonds, and why water behaves as it does

3.2.1. The bonds you will meet

Why atoms bond at all — electron shells, electronegativity, ions — is set out in Chapter 2. Here are the bond types themselves, ranked by what they mean for a device.

Bond	Mechanism	Strength	Where it matters to you
Covalent	Shared electron pair	Strong (~150–500 kJ/mol)	The polymer backbone. Breaking these is degradation

Bond	Mechanism	Strength	Where it matters to you
Ionic	Electrostatic attraction between opposite charges	Strong when dry, much weaker in water	Salts, electrolytes, ion-exchange
Hydrogen bond	Attraction between a polarised H and a nearby electronegative atom	Weak individually (~20 kJ/mol), decisive collectively	Water's behaviour; protein and DNA structure
Van der Waals / London	Transient dipole attraction	Very weak, but scales with contact area	Why gecko adhesion and fine-surface adhesion work
Hydrophobic effect	Not a bond — water excluding nonpolar material to maximise its own hydrogen bonding	Collectively strong	Why lipid membranes self-assemble and proteins fold

The last row is the one people get wrong. The hydrophobic effect is not an attraction between oily things. It is water's own hydrogen-bonding network pushing nonpolar material out of the way. Every cell membrane in your body exists because of it.

3.2.2. Water

Water is roughly 50–60% of adult body mass, and almost every property that matters follows from one fact: the molecule is **polar** and forms extensive hydrogen bonds.

Property	Consequence
Excellent solvent for polar and ionic substances	Blood, cytoplasm, and interstitial fluid are aqueous. Anything implanted or inserted is in a wet, ionic, electrically conductive environment
Poor solvent for nonpolar substances	Lipids separate out. Membranes form. Drug formulation becomes hard for poorly soluble compounds
High specific heat capacity (~4.18 kJ/kg·K)	The body buffers temperature change. Also why a water-filled phantom behaves thermally like tissue
High latent heat of vaporisation	Sweating is an efficient cooling mechanism — and why an occlusive enclosure defeats it (Chapter 5)
Cohesion and surface tension	Capillary action — which is why a press-fit seam wicks fluid and is a cleaning problem (Chapter 13)
Nearly incompressible	Tissue deforms by **moving fluid around**, not by compressing. This is why soft tissue is viscoelastic and time-dependent (Chapter 5)



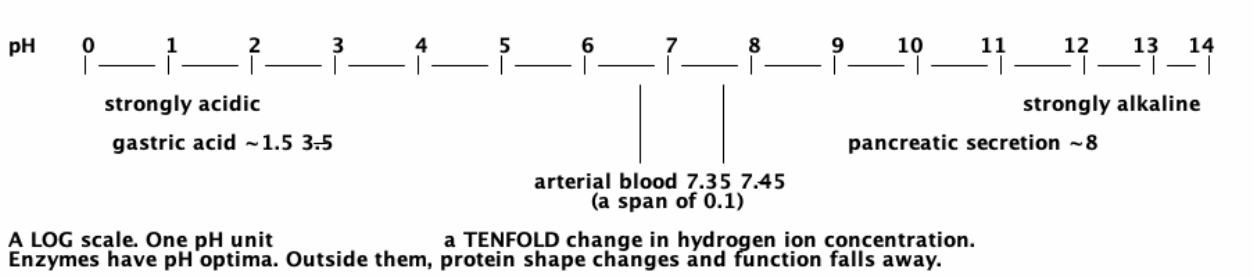
That last row is worth sitting with. When you compress soft tissue, you are not squashing a solid — you are pumping fluid out of a porous matrix. The water has to go somewhere, and it takes time to get there. **That is the physical origin of creep, stress relaxation, and the load–duration trade-off** that Chapter 5 treats

as a design constraint.

3.3. pH, buffers, and why the range is so narrow

pH is the negative base-10 logarithm of hydrogen ion activity. It is a **log scale**: pH 6 has ten times the hydrogen ion concentration of pH 7.

Arterial blood pH is normally held around **7.35–7.45** [1]. That is a span of 0.1 pH units — about a 25% variation in hydrogen ion concentration — and falling outside roughly 6.8–7.8 is incompatible with life.



A **buffer** is a weak acid and its conjugate base together, which resists pH change by mopping up or releasing hydrogen ions. The body's main extracellular buffer is the **bicarbonate** system:

$CO_2 + H_2O \rightleftharpoons H_2CO_3 \rightleftharpoons H^+ + HCO_3^-$

Read that equation left to right and you have the link between **breathing** and **blood chemistry**: blow off more CO_2 and the equilibrium shifts left, consuming H^+ and raising pH. This is why respiratory rate affects acid–base balance, and why Chapter 9's respiratory material and this chapter are the same subject viewed from two ends.

Disturbance	Primary cause	How the body compensates
Respiratory acidosis	CO_2 retention — inadequate ventilation	Kidneys retain bicarbonate (slow, hours to days)
Respiratory alkalosis	Excess CO_2 removal — hyperventilation	Kidneys excrete bicarbonate
Metabolic acidosis	Acid gain or bicarbonate loss	Increased ventilation blows off CO_2 (fast, minutes)
Metabolic alkalosis	Acid loss or bicarbonate gain	Reduced ventilation



Recognising these four names is general knowledge and lets you follow a conversation. **Interpreting a blood gas result is a clinical skill**, and a difficult one. Nothing here moves you toward it.

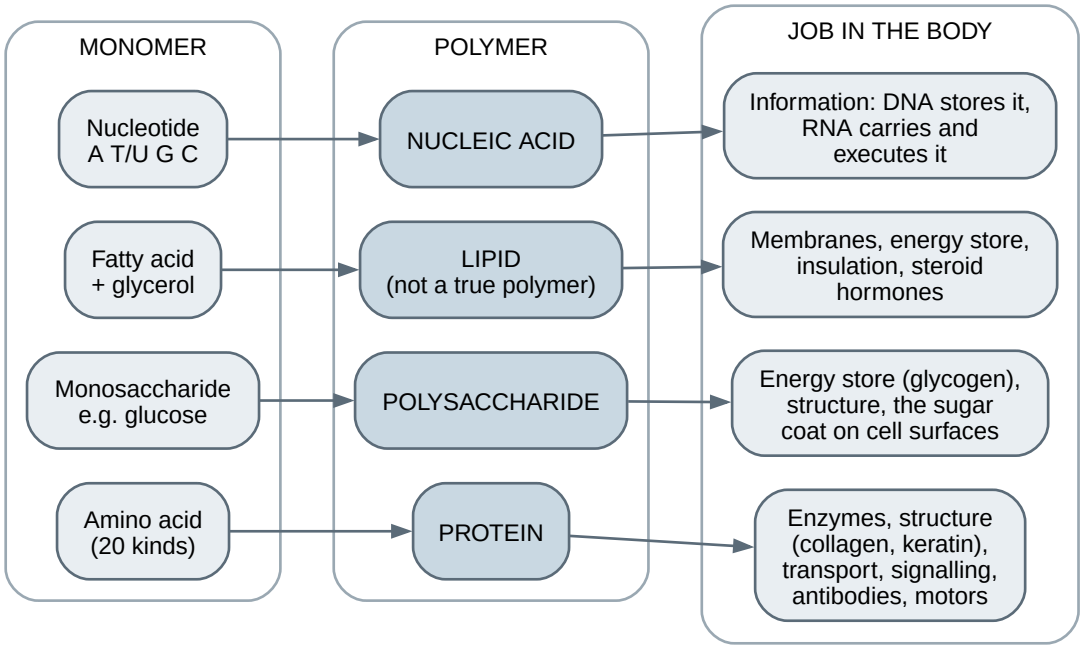
3.3.1. Why pH matters to a device

- **Enzyme function is pH-dependent.** Any device with an enzyme in it — a glucose biosensor,

most notably — has a working pH range.

- **Material degradation is pH-dependent.** Polyesters such as PLA and PLGA hydrolyse, and their own acidic degradation products can lower local pH and accelerate the process.
- **Corrosion is pH-dependent.** Metal implants sit in a warm, salty, variably acidic environment, which is a genuinely aggressive corrosion test.
- **Drug solubility and stability are pH-dependent.** Formulation is largely a pH and excipient problem.

3.4. The four classes of biological molecule



3.4.1. Proteins, in more detail

Proteins do most of the work, and their one organising principle is that **shape determines function**.

Level	What it is	Disrupted by
Primary	The sequence of amino acids, set by the gene	Only by breaking covalent bonds
Secondary	Local regular folds — alpha helix, beta sheet — held by hydrogen bonds	Heat, pH, solvents
Tertiary	The overall 3D fold of one chain	Heat, pH, solvents, some surfaces
Quaternary	Assembly of several chains, e.g. haemoglobin's four	As above

Denaturation is loss of that shape without breaking the sequence. Heat does it; so do extremes of pH, organic solvents, and — relevantly — **adsorption onto a foreign surface**.



Protein adsorption is the first thing that happens when any material meets blood or tissue. Within seconds, a layer of protein coats the surface. Everything the body does next — cell attachment, inflammation, clotting, encapsulation — responds to **that adsorbed layer**, not to your material directly.

This is the single most useful idea in biomaterials for a newcomer. "Is this material biocompatible?" is really "what protein layer forms on it, in what conformation, and what does the body do in response?" It is why surface chemistry and surface topography matter far more than bulk composition, and why a polished and a roughened sample of **the same alloy** can behave differently.

Two protein families worth knowing by name:

- **Collagen** — the dominant tensile protein of tendon, ligament, skin, and bone's organic matrix (Chapter 5). A triple helix, rich in glycine and proline, assembled into fibrils. Roughly a third of total body protein.
- **Haemoglobin** — four subunits, each with an iron-containing **haem** group that binds oxygen. The cooperative binding between subunits is what produces the sigmoid dissociation curve in Chapter 9.

3.4.2. Enzymes

Enzymes are protein catalysts. They lower the activation energy of a reaction without being consumed, and they are **specific** — a given enzyme acts on a given substrate.

Property	Why a device engineer cares
Specificity	It is what makes an enzyme electrode selective. Glucose oxidase responds to glucose and not to the hundreds of other things in blood
Temperature dependence	Rate rises with temperature until the protein denatures. Enzyme-based sensors have a working range and a shelf life
pH dependence	Each has an optimum. Buffering the sensor chemistry is part of the design
Saturation kinetics	At high substrate concentration every active site is busy and the rate plateaus. **This sets the upper end of a biosensor's linear range** — above it, the reading stops tracking
Inhibition	Other substances can block the active site. In a sensor this is **interference**; in pharmacology it is how a large fraction of drugs work



Michaelis–Menten kinetics describes the saturation curve, with $\mathbf{V}_{max}$ the maximum rate and $\mathbf{K}_m$ the substrate concentration at half of it. You do not need to derive it. You do need to know that **every enzyme-based sensor has a concentration above which it stops responding linearly**, and that the manufacturer's stated range is where that was characterised.

3.4.3. Lipids and membranes

A **phospholipid** has a polar phosphate head and two nonpolar fatty acid tails. In water, the hydrophobic effect drives them to self-assemble into a **bilayer**: heads out, tails in.

That bilayer is the basis of every cell membrane, and it has three properties that matter:

1. **It is a selective barrier.** Small nonpolar molecules (O_2, CO_2) cross freely; ions and polar molecules need protein channels or transporters (Chapter 1).
2. **It is fluid.** Components move laterally within it. It is a two-dimensional liquid, not a solid sheet.
3. **It is electrically insulating**, with embedded ion channels. That combination — a capacitor with voltage-gated resistors in parallel — is exactly what makes the action potential possible (Chapter 7).

Also lipid-derived: **cholesterol**, which modulates membrane fluidity and is the precursor of **steroid hormones**; and **triglycerides**, the main energy store in adipose tissue.

3.4.4. Carbohydrates

Beyond energy (**glucose**, stored as **glycogen**), the one to know is the **glycocalyx** — the sugar coating on cell surfaces. It mediates recognition, adhesion, and lubrication, and it is part of why blood vessel interiors are non-thrombogenic. Damage it and clotting behaviour changes.

3.4.5. Nucleic acids, briefly

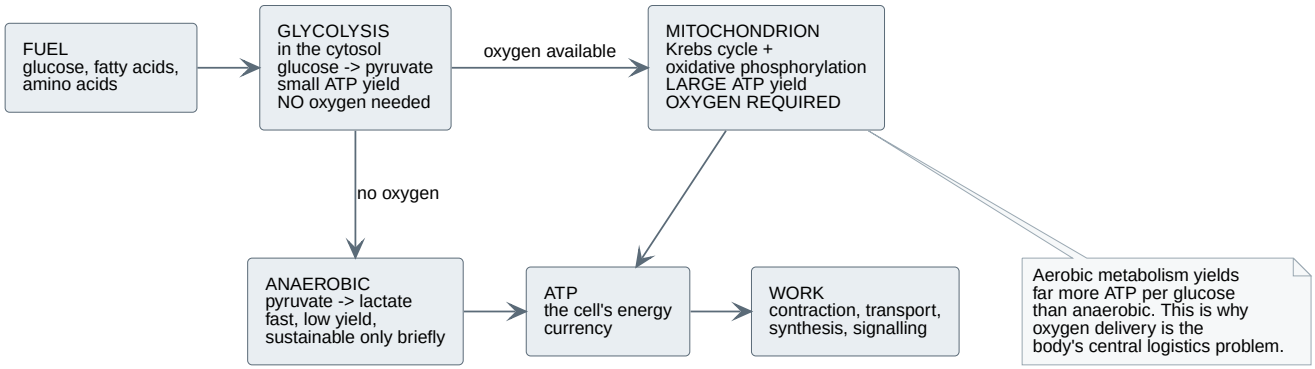
DNA stores the sequence information; RNA carries and executes it.

DNA --(transcription)-→ RNA --(translation)-→ PROTEIN -→ function

This is the **central dogma**. For a maker, three consequences are enough:

- A **mutation** changes a protein's sequence, which can change its shape, which can change or abolish its function. That is the mechanism of genetic disease.
- **PCR** amplifies a specific DNA sequence through repeated thermal cycling — the basis of a huge amount of modern diagnostics (Chapter 15).
- **Sequencing** reads the order of bases. Costs have collapsed, which is why genomic medicine is now a clinical reality rather than a research technique.

3.5. Bioenergetics: where the energy comes from



The practical consequences, which connect directly to earlier chapters:

- **Oxygen delivery is the body's central logistics problem**, and the whole cardiovascular and respiratory apparatus of Chapter 9 exists to solve it.
- **Interrupt perfusion and ATP production falls within minutes.** Ion pumps fail, gradients collapse, cells swell. This is the mechanism behind Chapter 5' insistence that sustained compression is dangerous — you are not bruising tissue, you are starving it.
- **Different tissues tolerate interruption very differently.** Brain and myocardium are intolerant; cartilage and skin tolerate far more.
- **Basal metabolic rate** for a typical adult is roughly 60–75 W — about the same as an old incandescent bulb. Useful when thinking about heat generation inside any enclosure against a person.

3.6. Cell signalling and pharmacology

Cells communicate chemically, and most drugs work by interfering with that.

Mode	Range	Example
Autocrine	The cell signals itself	Some immune regulation
Paracrine	Nearby cells	Local inflammatory mediators
Synaptic	Across a synaptic cleft	Acetylcholine at the neuromuscular junction (Chapter 7)
Endocrine	Via the bloodstream, body-wide	Insulin, thyroid hormone, adrenaline

A **receptor** is a protein that binds a specific signalling molecule and produces a cellular response. An **agonist** binds and activates it; an **antagonist** binds and blocks it. A great many drug names resolve to one of those two words plus a receptor.

Pharmacology term	Meaning
Pharmacokinetics (PK)	What the body does to the drug: absorption, distribution, metabolism, excretion (**ADME**)
Pharmacodynamics (PD)	What the drug does to the body: receptor binding, effect, dose–response

Pharmacology term	Meaning
Bioavailability	Fraction of an administered dose reaching systemic circulation unchanged. Intravenous is 100% by definition
Half-life	Time for plasma concentration to halve. Sets dosing interval
Therapeutic index	Ratio between a toxic dose and an effective dose. A **narrow** index means small errors matter enormously
First-pass metabolism	Drug absorbed from the gut passes through the liver before reaching the circulation, often losing much of its activity



Narrow therapeutic index is the phrase to remember. For some drugs — warfarin, lithium, digoxin, many chemotherapy agents — the gap between an effective and a harmful dose is small.

This is why infusion pumps are among the most heavily regulated devices in existence, why dose-calculation software is a serious safety-critical application, and why a unit error in a medication app is not a cosmetic bug. See Chapter 19.

3.7. Engineering bridge: the chemistry you meet as a maker

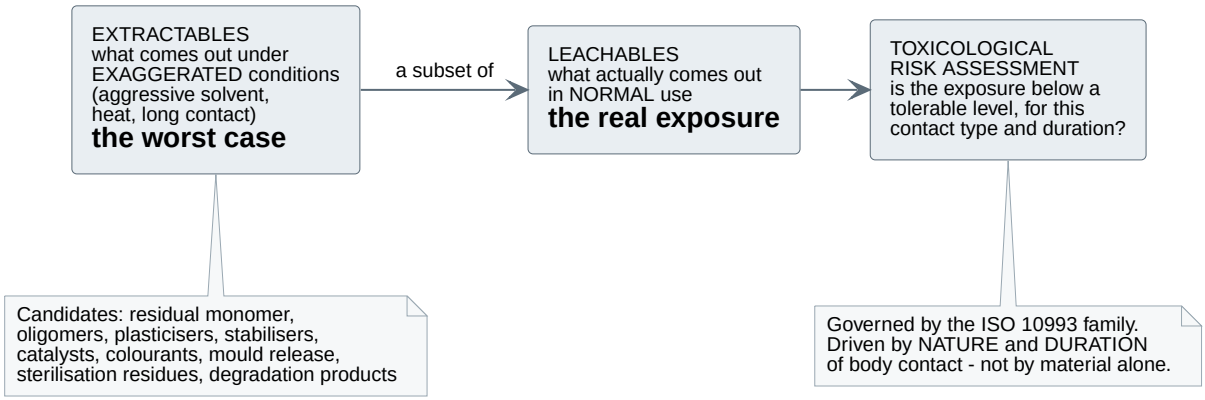
This is the section to come back to. Everything above was background for it.

3.7.1. Polymer chemistry, practically

Concept	Meaning	Why it bites
Thermoplastic vs thermoset	Thermoplastics soften on heating and can be reformed; thermosets cure irreversibly	PP, PC, POM (acetal), PEEK, ABS, PETG, and PLA are thermoplastics — moulded, machined, or printed — and most deform in an autoclave; PEEK and some PP and PC grades are among the exceptions. Epoxies, cured photopolymers, and most silicones are thermosets
Glass transition, T_g	Temperature at which an amorphous polymer goes from glassy to rubbery	PLA's T_g is ~55–65 °C; a car interior in summer exceeds it, and so does a dishwasher (Chapter 13). Polycarbonate's is ~145 °C; PEEK's ~143 °C, with crystallinity keeping it stiff well above that
Crystallinity	Degree of ordered chain packing	Higher crystallinity means more stiffness, better chemical resistance, less transparency
Molecular weight	Average chain length	Falls during processing and degradation. Falling molecular weight is usually the first measurable sign a polymer is breaking down
Additives	Plasticisers, stabilisers, colourants, fillers, flame retardants	**These are the things most likely to leach out.** The datasheet names the polymer, not always the additive package

Concept	Meaning	Why it bites
Hydrolytic degradation	Water cleaves the backbone, mainly in polyesters	PLA, PGA, PLGA, PCL degrade in the body — deliberately, for resorbable devices, and inconveniently otherwise
Oxidative degradation	Oxygen, UV, and reactive species attack chains	Why things yellow, embrittle, and crack with age
Environmental stress cracking	Stress plus a chemical agent cracks a polymer that either alone would not	**Polycarbonate plus certain alcohols and disinfectants.** A real cause of device failure in service
Processing history	The route from feedstock to part changes the part	Injection moulding leaves **residual stress** (which feeds stress cracking) and may leave mould release on the surface. Machined stock is usually stress-relieved but picks up coolant and burrs. Printing leaves layer interfaces, porosity, and anisotropy. The same polymer by three routes is three different parts
Elastomers	Rubbery polymers: silicones, TPU and TPE, natural rubber latex, nitrile	Silicone is the workhorse of skin contact and casting. **Natural rubber latex is a recognised allergen** in healthcare; many organisations avoid it entirely

3.7.2. Leachables, extractables, and biocompatibility



The framing to take away: **biocompatibility is not a property of a material.** It is a conclusion about a **finished device**, in a **specific contact category**, for a **specific duration**. ISO 10993-1 organises the whole evaluation on exactly those axes (Appendix D).

Which is why Chapter 13 insists that a "medical grade" feedstock — filament, moulding pellets, bar stock, silicone kit, or fabric — establishes nothing about your finished part. The grade describes what went in. Your part additionally carries its processing history: a moulding's release agent and residual stress, a machined part's coolant and burrs, a print's layer lines and machine contamination, a fabric's dyes and finishes, and whatever your post-processing left behind.

Source of leachables	Typical culprits
The polymer itself	Residual monomer, low-molecular-weight oligomers
The additive package	Plasticisers, antioxidants, UV stabilisers, colourants
Processing	Catalysts, mould release agents, lubricants, cutting fluids and coolants, machine residues
Post-processing	Solvents (acetone, IPA), uncured resin, adhesive components, polishing compound, passivation acid residues
Textile finishes	Dyes, softeners, water- and stain-repellent finishes, crease-resistant resins, antimicrobial treatments
Sterilisation	Ethylene oxide and its by-products ethylene chlorohydrin and ethylene glycol; radiation-induced free radicals and chain scission
Degradation in service	Hydrolysis and oxidation products, including acidic ones



Several ordinary workshop chemicals are sensitisers. Not merely irritants — **sensitisers**, meaning repeated exposure can produce an allergic response that then persists:

- **uncured photopolymer resin** from resin printing;
- **epoxy** resins and especially their amine hardeners, used in casting, laminating, and bonding;
- **isocyanates** in polyurethane foams, coatings, and some adhesives;
- **rosin (colophony)** in solder flux fume — a recognised cause of occupational asthma;
- **nickel** in many alloys and platings, against the skin.

For anything that will touch a person, the corresponding process steps — full cure, washing, post-curing, ventilation, material selection — must be **specified and verified**, not assumed. That protects the user and it protects you.

3.7.3. Metals and corrosion

Body fluid is warm, oxygenated, around pH 7.4, and contains roughly 0.9% sodium chloride. That is an aggressive environment for metal.

Mechanism	Note
Passivation	Stainless steel and titanium both owe their corrosion resistance to a thin, tenacious, self-repairing oxide film. **Scratch it and it normally re-forms — unless conditions prevent that**
Pitting and crevice corrosion	Localised attack where the passive film cannot re-form — under a washer, in a thread, in a tight joint. Crevices are both a cleaning problem and a corrosion problem

Mechanism	Note
Galvanic corrosion	Two dissimilar metals in electrical contact in an electrolyte: the less noble one corrodes preferentially. **Mixing alloys in one assembly is a design decision, not a convenience**
Fretting corrosion	Micro-motion at an interface continually abrades the passive film. A known issue at modular implant junctions
Ion release	Even slow corrosion releases metal ions. Nickel in particular is a common contact allergen, which is relevant to anything worn

3.7.4. Textiles and glass

Two material families a maker uses constantly and a device engineer is rarely taught.

Textiles — straps, sleeves, covers, slings, padding, garments — are polymer chemistry in fibre form, with a large surface area and close, prolonged skin contact.

Fibre or finish	Practical behaviour
Polyester (PET)	Strong, low moisture uptake, tolerates hot laundering; the default for healthcare textiles
Nylon (polyamide)	Tough and abrasion-resistant; absorbs more water than polyester; webbing and hook-and-loop
Polypropylene	Very low moisture uptake; used in nonwoven disposables, gowns, and wicking layers; poor heat tolerance
Elastane (spandex)	Provides stretch and compression; degraded by heat, chlorine, and some disinfectants, so compression falls with repeated laundering
Cotton	Comfortable and absorbent; holds moisture against skin (Chapter 5' microclimate) and is slow to dry
Coatings and laminates — PU, PVC, silicone	Make a fabric fluid-resistant and wipeable. Coatings crack and delaminate with flexing and disinfectant exposure; once breached, the fabric beneath is uncleanable
Dyes and finishes	The usual source of textile leachables and contact allergy. A fabric's composition label rarely mentions them

Reusable healthcare textiles are commonly laundered with **thermal disinfection** — for example 65 °C held for 10 minutes or 71 °C for 3 minutes in UK guidance — and every component of a sewn item (thread, elastic, hook-and-loop, buckles, padding) must survive that, repeatedly.

Glass is chemically one of the most inert materials you can use, which is why drug vials and laboratory ware are made of it. **Borosilicate** glass (low thermal expansion, about 3.3×10^{-6} per kelvin) tolerates thermal shock and autoclaving; ordinary **soda-lime** glass expands roughly three times as much and cracks if heated unevenly. Pharmaceutical containers are graded by how much alkali leaches from them, and "Type I" borosilicate is the most resistant. Glass's limitations are mechanical, not chemical: it fails suddenly, without yielding, and its fragments are sharp — so in anything handled by a person, how it breaks is a design question.

3.7.5. Adhesion and surface chemistry

Concept	Practical consequence
Surface energy	High-energy surfaces (clean metals, glass) wet easily and bond well. Low-energy ones (PP, PE, PTFE, silicone) resist both. This is why polypropylene is hard to glue and silicone needs a primer
Wettability / contact angle	A cheap, measurable proxy for surface energy. Also predicts how a cleaning solution behaves on your part
Surface treatment	Plasma, corona, flame, and chemical etching raise surface energy — usually temporarily, so there is a working window after treatment
Adhesive chemistry	Cyanoacrylate cures with moisture, is brittle and has poor peel strength. Epoxies are two-part thermosets, strong but with a defined pot life. Silicones cure by condensation or addition; condensation cure releases acetic acid, which can corrode
Roughness	Increases mechanical keying **and** surface area for protein adsorption **and** habitat for biofilm (Chapter 13). The same change helps one objective and harms another

3.8. Devices, software, and training technology

Every chapter from here on ends with this section. The pattern is the same each time: what exists clinically, what software exists, what is used to **teach** the subject, and where a maker can realistically contribute.

3.8.1. Clinical devices built on this chemistry

Device	Chemistry it runs on	Engineering note
Blood glucose meter	Glucose oxidase or glucose dehydrogenase on a disposable electrode; the enzyme reaction produces a measurable current	The classic enzyme electrode. Saturation kinetics set the upper linear range; interferents and haematocrit affect accuracy
Continuous glucose monitor (CGM)	Same enzyme chemistry on a subcutaneous filament, sampling interstitial fluid	Interstitial glucose **lags** blood glucose — a real and clinically important measurement limitation, not a defect
Blood gas analyser	Ion-selective electrodes for pH and ions; amperometric sensor for O_2; Severinghaus electrode for CO_2	Requires frequent calibration and temperature control
Infusion pump	None itself — but it exists because of narrow therapeutic indices	Among the most safety-critical device classes. Dose error reduction software, drug libraries, hard and soft limits
Haemodialysis machine	Diffusion and ultrafiltration across a semipermeable membrane against a dialysate of controlled composition	A mass-transfer device. Membrane area, flow rates, and concentration gradients

Device	Chemistry it runs on	Engineering note
Pulse oximeter	Differential optical absorbance of oxy- and deoxy-haemoglobin	Treated fully in Chapter 9 — including its documented accuracy limitations
Drug-eluting and resorbable implants	Controlled hydrolytic degradation of polyesters (PLGA and relatives) releasing a drug	Degradation rate is tuned by copolymer ratio, molecular weight, and crystallinity
Transdermal patch	Diffusion through the stratum corneum, often with a permeation enhancer	Limited to small, moderately lipophilic molecules — the barrier in Chapter 5 is genuinely good at its job
Pulse-field and electrosurgical devices	Electrochemistry and Joule heating in tissue	High-energy, specialist, heavily regulated
Wound dressings and medical textiles	Hydrocolloids, hydrogels, alginates, foams, and silver or other antimicrobial agents; fibre and coating chemistry	Moisture management is the design variable. Textile and coating chemistry, not shape, decides performance (Chapter 5)
Glass syringes, vials, and ampoules	Type I borosilicate glass chosen for low leaching and chemical durability	Glass is a chemistry decision: alkali leaching, delamination, and breakage are known packaging failure modes

3.8.2. Software

Software	What it does	Relevance
Electronic prescribing with decision support	Checks dose, interaction, allergy, renal adjustment	One of the highest-value and highest-risk clinical software categories
Pharmacokinetic dose calculators	Model drug concentration over time to individualise dosing	Used for narrow-therapeutic-index drugs. Safety-critical
Laboratory information systems (LIS)	Order, track, and report laboratory results	Integration, identity matching, and result-flagging rules
PK/PD modelling and simulation	Research and regulatory submission tools	A quantitative specialism in its own right
Molecular visualisation — PyMOL, UCSF ChimeraX, Jmol, Mol*	Render and manipulate protein and nucleic acid structures from the Protein Data Bank	Free, open, scriptable. **Genuinely usable by a maker today**
Protein structure prediction — AlphaFold and successors	Predict 3D structure from sequence	Has changed structural biology. Database of predicted structures is openly available

Software	What it does	Relevance
Chemistry drawing and modelling — Avogadro, RDKit, Open Babel	Build, convert, and compute on molecular structures	Open source; RDKit is the standard Python cheminformatics toolkit

3.8.3. Training and simulation

Training technology	Teaches	Form
Molecular model kits — physical ball-and-stick	Bonding, geometry, chirality, macromolecular shape	The original task trainer. Still excellent
Physical protein and molecule models	Shape, active sites, scale relationships	**Direct maker opportunity.** Printed from Protein Data Bank meshes; stacked laser-cut acrylic sections; bent-wire backbones in the tradition of the first protein models; turned and drilled ball-and-stick parts for small molecules
Chemistry demonstration apparatus	Diffusion, osmosis, buffering, gas exchange	Glassware, tubing, and inexpensive pH, conductivity, and gas sensors on a microcontroller. Scientific glassblowing and simple electronics make durable, reusable classroom kit
Virtual laboratory software — Labster, LabXchange, PhET	Technique, procedure, and reasoning without reagent cost or risk	Widely used in pre-clinical and undergraduate teaching
Pharmacology simulators	Dose–response, PK curves, receptor behaviour	Often browser-based and free
Medication administration trainers	Drawing up, dilution, labelling, the "five rights"	Manikin arms, simulated vials, barcode-scanning practice systems
Simulated medication and barcode systems	Safe practice of the real scanning workflow without real drugs	A notable niche: simulated-drug labels and packaging are a genuine consumable need
Blood gas and lab-result simulators	Interpretation practice against generated cases	Software; several free web versions exist



The realistic maker entry point in this chapter is the physical molecular model.

Structures come from the Protein Data Bank (free, open). ChimeraX, Mol*, and several dedicated tools export meshes and coordinates. Educators genuinely want models — a haemoglobin model showing the four subunits and the haem pockets explains cooperative binding in a way a diagram struggles to.

Choose the method for what the model must teach. Additive manufacturing is the natural fit for a whole folded surface. Stacked laser-cut acrylic slices show internal cavities and make a clean, cleanable classroom object. A bent-wire backbone shows the fold and is nearly indestructible. Turned and drilled components suit small-molecule kits that students assemble themselves. A chemistry demonstration — a buffer, a diffusion cell, a sensor reading on screen — can teach more than any static model.

It is also low-risk: no patient contact, no infection control beyond ordinary handling, no clinical claim. The accuracy obligation is real but tractable — the structure is published data, and Chapter 21' labelling discipline applies in full. State the resolution, state what the model omits, and state that colour is a convention, not a property of the molecule.

3.9. Worked example: why the handle got sticky

A polymer handle — in this case printed in PLA, though the analysis is the same for a moulded or machined one — has been in service for six months in a warm room. It has become slightly tacky and has fine surface cracks. Nothing was dropped on it. What happened?

Candidate mechanisms, and how you would distinguish them:

Mechanism	Evidence for it	Check
Plasticiser migration	Tackiness, slight dimensional change, the surface feels different from a fresh sample	Compare against a stored control sample of the same batch
Hydrolytic degradation	More likely in a polyester (PLA), accelerated by warmth and humidity; material becomes more brittle	Compare mechanical behaviour with the control; check storage humidity
Oxidative degradation / UV	Yellowing, embrittlement, surface cracking, worse on light-exposed faces	Compare an exposed face with a shaded one
Environmental stress cracking	Fine cracks appearing where the part is stressed — including residual moulding stress — after exposure to a cleaning agent	Which agent is used? Is it compatible with this polymer? Check the cracks' location against the stress pattern
Residual cleaning agent	Tackiness only; no cracking; improves after thorough rinsing	Rinse and re-examine

Note the method, which is the same one Chapter 19 teaches: **keep a control sample**. Without a retained specimen from the same batch, stored in known conditions, you are comparing a six-month-old part against your memory.



Note also that four of those five mechanisms were **predictable from the material choice at the design stage** — and that the cleaning-agent one could only have been caught by asking Chapter 13's question: "what will this actually be cleaned with?"

3.10. Workshop exercise

1. **Build a bilayer.** Using modelling clay or card, make twenty phospholipid molecules with a distinct polar head and two tails. Arrange them into a bilayer. Then show what happens at an edge, and why a sheet spontaneously closes into a vesicle.
2. **pH in practice.** With supermarket indicator paper or a cheap pH pen, measure tap water, a soft drink, soap solution, and vinegar. Plot them on a log scale against the 7.35–7.45 blood range. The point is the **visual compression** of that range.
3. **Buffer demonstration.** Add a few drops of vinegar to (a) plain water and (b) water with a teaspoon of bicarbonate of soda, measuring pH each time. The difference is what a buffer does.
4. **Make a molecule.** Download a structure from the Protein Data Bank — haemoglobin (`1A3N`) or insulin (`4INS`) are good starts. Make a physical model by whichever method suits what you want it to show: a printed surface, stacked laser-cut sections, a bent-wire backbone, or a turned ball-and-stick fragment of the active site. Then write the limitations statement: what resolution, what is omitted, what the colours mean.
5. **Material survey.** For three materials in your workshop — say a polymer, a metal or adhesive, and a fabric or foam — find the actual safety data sheet or technical data. Identify: the polymer, any named additives, the recommended handling, and what it says about skin contact. Note how much it does **not** tell you.
6. **Control samples.** Start the habit now. Make two identical test coupons of anything you are using. Put one in service, store the other in a labelled bag with the date. In six months you will have the comparison you need.

3.11. Questions

1. Why is the hydrophobic effect not an attraction between nonpolar molecules?
2. Blood pH is held between 7.35 and 7.45. In terms of hydrogen ion concentration, how wide is that?
3. Write the bicarbonate buffer equation and explain how hyperventilation raises blood pH.
4. What is the first thing that happens when a material contacts blood, and why does it dominate everything after?
5. Why does every enzyme-based biosensor have an upper limit to its linear range?
6. Why does soft tissue take time to deform under load?
7. What is the difference between an extractable and a leachable?

8. Name three things that could leach from a moulded, machined, or printed part that are not the polymer.
9. Why is mixing two different alloys in one assembly a design decision?
10. What is environmental stress cracking, and why does it make disinfectant choice a materials question?
11. A device uses a drug with a narrow therapeutic index. What does that imply about your software?
12. Why is polypropylene hard to glue?

3.12. Answers and reasoning

1. Because it is driven by **water**, not by the nonpolar molecules. Water maximises its own hydrogen bonding by excluding nonpolar material, which pushes that material together. The association is a consequence of water's behaviour, not of an attractive force between the solutes.
2. About **25%**. pH is a negative log scale, so 0.1 pH units is a factor of $10^{0.1} \approx 1.26$ in hydrogen ion concentration. The narrowness is the point.
3. $CO_2 + H_2O \rightleftharpoons H_2CO_3 \rightleftharpoons H^{\wedge}$ + HCO~3~^−^. Hyperventilation removes CO~2~, shifting the equilibrium to the *left*, which consumes H^, which **raises** pH — respiratory alkalosis.
4. **Protein adsorption.** Within seconds a protein layer coats the surface. All subsequent biological response — cell attachment, inflammation, clotting, encapsulation — is a response to that adsorbed layer and its conformation, not to the bulk material. It is why surface chemistry and topography dominate bulk composition.
5. Because of **saturation kinetics**: above a certain substrate concentration every enzyme active site is occupied, the reaction rate plateaus, and the output stops tracking concentration. The stated measuring range is where the response was characterised as linear.
6. Because tissue is a porous matrix saturated with nearly incompressible water. Deformation requires **fluid to move through the matrix**, which takes time. That is the physical origin of creep, stress relaxation, and the load–duration relationship.
7. **Extractables** are what can be pulled out under deliberately exaggerated conditions — aggressive solvent, elevated temperature, long contact — representing a worst case. **Leachables** are the subset that actually migrates out in normal use, representing the real exposure. Extractables scope the problem; leachables quantify it.
8. Any three of: residual monomer or oligomers; plasticisers; antioxidants or UV stabilisers; colourants; mould release or machine lubricant; solvent from post-processing; cutting fluid; uncured resin; adhesive components; sterilisation residues; degradation products.
9. Because two dissimilar metals in electrical contact within an electrolyte form a galvanic cell, and the less noble metal corrodes preferentially. Body fluid is a good electrolyte. The combination must be chosen deliberately, not by whatever fastener was in the drawer.
10. Mechanical stress plus a chemical agent together cracking a polymer that neither would crack alone — polycarbonate with certain alcohols is the classic case. It makes disinfectant selection a materials-compatibility question, because the host's cleaning agent can destroy a part that

passed every mechanical test.

11. That dose errors have a small margin before harm. Your software needs units carried in the type system, explicit range checking, rejection rather than silent clamping of invalid input, clear display of units, and a tested error path. See Chapter 19. It also means the application is likely a regulated medical device.
12. Because it has **low surface energy**. Adhesives wet it poorly and cannot form good interfacial contact. The fixes are surface treatment — plasma, corona, flame — or a primer, or a mechanical rather than adhesive joint.

3.13. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Solutions — pH, buffers and the bicarbonate system.
- Titration — titration curves.
- Molecules and proteins — molecules, and protein structures from the PDB in 3-D with STL export.

3.14. Video companion and reading

- **Core reading:** OpenStax [1] chapter 2 (The Chemical Level of Organization) — now read it properly rather than skimming, particularly 2.1 Elements and Atoms, 2.2 Chemical Bonds, 2.3 Chemical Reactions, 2.4 Inorganic Compounds (water, pH, buffers), and 2.5 Organic Compounds. Then 3.1 The Cell Membrane and 24.1–24.2 for metabolism overview.
- **Core video:** Crash Course Anatomy & Physiology #2–#3 cover the chemical level; the Crash Course **Chemistry** and **Biology** series are better if you want more depth [6].
- **Structures to handle:** the Protein Data Bank [33] — free, open, and the source for the molecule exercise. Its **Molecule of the Month** series is an outstanding, accessible introduction to structural biology.
- **Materials:** for the device-chemistry material, the ISO 10993 series scope sections (Appendix D) are the practical reference; the FDA additive manufacturing guidance [26] adds the printing-specific questions, and Chapter 14 compares processes.

3.15. Exit criteria

- ☐ Explain the four bond types and the hydrophobic effect, and say why the last is not an attraction between nonpolar molecules.
- ☐ Explain pH as a log scale and state the blood range and what it means in concentration terms.
- ☐ Write the bicarbonate buffer equation and explain the breathing–pH link in both directions.
- ☐ Name the four classes of biological molecule, their monomers, and one job each.
- ☐ Explain why protein shape determines function and what denatures a protein.

- ▢ Explain protein adsorption and why it reframes the biocompatibility question.
- ▢ Explain enzyme saturation and connect it to a biosensor's linear range.
- ▢ Distinguish extractables from leachables, and list five sources of leachables in a made part, naming which depend on the process (moulding, machining, printing, sewing).
- ▢ Name three polymer degradation mechanisms and one corrosion mechanism, with a design consequence for each.
- ▢ Explain why low-surface-energy plastics resist adhesion and name two remedies.
- ▢ Make one molecular model, by a method chosen for what it must show, with a written limitations statement.
- ▢ Start a retained control sample for a material you are using.

3.16. What this chapter does not establish

- **You cannot assess biocompatibility.** You understand why the question is framed around a finished device and a contact category. The evaluation is a specialist activity with a standards family behind it.
- **You cannot select or validate a sterilisation chemistry.** Knowing that ethylene oxide leaves residues is not the same as managing them.
- **You cannot interpret a blood gas, a glucose reading, or any other result.** You know what the chemistry measures. What it means for a person is clinical.
- **You have not learned pharmacology.** You have learned enough vocabulary to follow a conversation and to understand why dose-handling software is safety-critical.
- **A physical molecule model is a schematic.** It renders one deposited structure at one resolution, usually without hydrogens, in conventional colours that are not properties of the molecule. Label it accordingly.

Chapter 4. Foundations of biology: life, cells, and how biologists work

Week 6 · approximately 6–8 hours

4.1. Why this chapter exists

Every chapter after this one describes tissues, organs, and systems. Every one of those is built from cells, run by cells, repaired by cells, and — when something goes wrong — usually broken at the level of cells. Chapter 1 introduced the cell as the unit of life and showed how its membrane holds an internal environment. Chapter 3 introduced the molecules cells are made of. This chapter puts the two together and answers the questions those chapters left open:

- What actually is a cell, what are its parts, and what does each part do?
- How does a sequence of DNA become a working protein, and how is that controlled?
- How do cells divide, specialise, age, and die?
- What are bacteria and viruses, and why are they so hard to get rid of?
- And — the part that matters most for you — **how do biologists actually work?** What equipment do they use, what does each instrument measure, and where do engineering and technology fit into that world?

There are three practical reasons to know this.

First, a great deal of modern medicine is applied cell biology. Insulin is made by bacteria. Cancer is diagnosed by sequencing tumour DNA. Vaccines are designed from viral genomes. Some leukaemias are treated by re-engineering the patient's own immune cells. If you work anywhere near diagnostics, therapeutics, or a hospital laboratory, you are working near cell biology.

Second, the laboratory is a technology-dense workplace. Microscopes, centrifuges, incubators, liquid handlers, sequencers, flow cytometers, bioreactors, and the software that drives and analyses them are all engineered products. Many are expensive, closed, and poorly suited to teaching or low-resource settings. Instrumentation, automation, optics, fixtures, enclosures, and analysis software are all places where a technologist is useful.

Third, cells set the physical limits on what you can build. Diffusion distances, growth rates, sterility, contamination, and the way surfaces become coated in protein and colonised by bacteria are all cell biology with direct design consequences.



This is not a biology degree in one chapter. It is the working core: enough to read a methods section, follow a laboratory conversation, understand what an instrument is for, and recognise when a question needs a biologist rather than a better tool.

4.2. What counts as life

Biologists do not agree on a single crisp definition, but living things share a set of properties. Each has an engineering reading.

Property	What it means	Engineering reading
Organisation	Built from cells, which contain ordered structures	A hierarchy of modules (Chapter 1' levels of organisation)
Metabolism	Takes in energy and matter, transforms them, excretes waste	An open system running far from equilibrium; stop the energy supply and order decays (Chapter 3)
Homeostasis	Holds internal conditions within limits	Feedback control (Chapter 1)
Growth and development	Increases in size and complexity following an internal programme	A build process driven by stored instructions
Reproduction with heredity	Makes copies, passing instructions to offspring	Self-replication with an error rate — and the errors matter
Response to the environment	Senses and reacts	Sensors, actuators, signalling
Evolution	Populations change over generations	An optimiser with no foresight, running continuously

Viruses sit on the boundary. They carry genetic instructions and evolve, but they have no metabolism of their own and can only replicate inside a host cell. Most biologists treat them as non-living infectious particles; for practical purposes in medicine what matters is that they are not cells, so the drugs that kill bacteria do nothing to them.

4.3. The organising ideas

Two ideas make the rest of biology hang together.

4.3.1. Cell theory

1. All living things are made of one or more cells.
2. The cell is the basic unit of structure and function.
3. Cells arise only from existing cells.

The third point sounds obvious and is not: it is why a sterile object stays sterile, why contamination has to come **from** somewhere, and why the chain of infection in Chapter 13 can be broken.

4.3.2. Evolution by natural selection

In every population, individuals vary. Some of that variation is inherited. Some variants survive and reproduce more successfully in their environment, so their variants become more common over generations. Nothing plans this; it is the inevitable result of variation, heredity, and

differential reproduction.

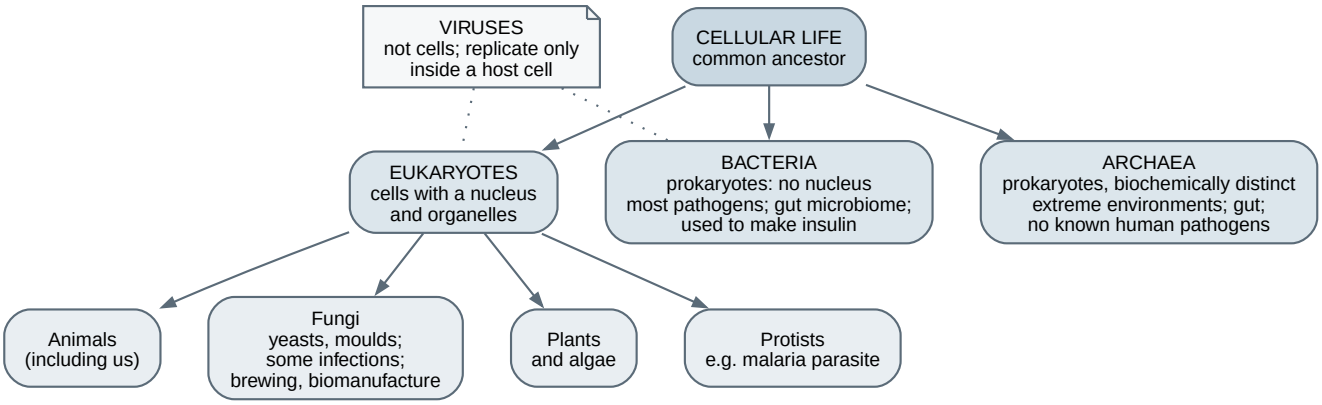
For a technologist, three consequences are worth having.

Bodies are not designed, and it shows. Evolution modifies what already exists; it never starts again. The recurrent laryngeal nerve detours down into the chest and back (Chapter 8). The retina is wired with its nerve fibres in front of the photoreceptors, creating the blind spot. The airway and the food passage cross. These are historical compromises, not optima, and they are the reason "it should work like a well-engineered machine" is so often wrong.

Variation is the normal state. No two people have the same body, the same enzymes, the same drug responses. The variation you design for in Chapter 6 and Chapter 12 is not noise around an ideal; it is the material evolution works with.

Evolution is happening now, fast, in microbes. A bacterial population of a billion cells, dividing every half hour, explores enormous numbers of mutations. Expose it to an antibiotic and the rare resistant cells survive and take over. **Antibiotic resistance** is evolution observed in real time, and it is why antibiotics are prescribed carefully, why courses and doses matter, and why infection prevention (Chapter 13) is cheaper than treatment.

4.4. The tree of life



All cellular life falls into three **domains**. Two of them — **bacteria** and **archaea** — are **prokaryotes**: small cells with no nucleus. The third — **eukaryotes** — have a nucleus and internal membrane-bound compartments. Animals, plants, fungi, and a large assortment of single-celled organisms are all eukaryotes. Viruses infect members of every domain.

4.4.1. Prokaryotic versus eukaryotic cells

Feature	Prokaryote (e.g. E. coli)	Eukaryote (e.g. a human cell)
Typical size	~1–2 µm long	~10–30 µm across
Volume	~1 µm³	~1,000–10,000 µm³
DNA	One circular chromosome in the cytoplasm, plus small **plasmids**	Several linear chromosomes inside a nucleus
Internal membranes	Essentially none	Many: nucleus, ER, Golgi, mitochondria, lysosomes

Feature	Prokaryote (e.g. E. coli)	Eukaryote (e.g. a human cell)
Ribosomes	Smaller (70S)	Larger (80S) in the cytoplasm
Cell wall	Usually; peptidoglycan in bacteria	None in animals; cellulose in plants; chitin in fungi
Division	Binary fission; can be ~20–30 minutes in ideal conditions	Mitosis; ~24 hours is typical for cultured human cells

The differences are not academic: **they are what antibiotics target.** Penicillin and its relatives attack the peptidoglycan cell wall, which human cells do not have. Several antibiotic families block the bacterial 70S ribosome but not the human 80S one. A drug is selective only to the extent that the target differs between pathogen and patient — which is why antibacterial drugs are plentiful, antifungals fewer (fungi are eukaryotes, like us), and antivirals hardest of all (viruses use our own machinery).

4.5. The eukaryotic cell, part by part

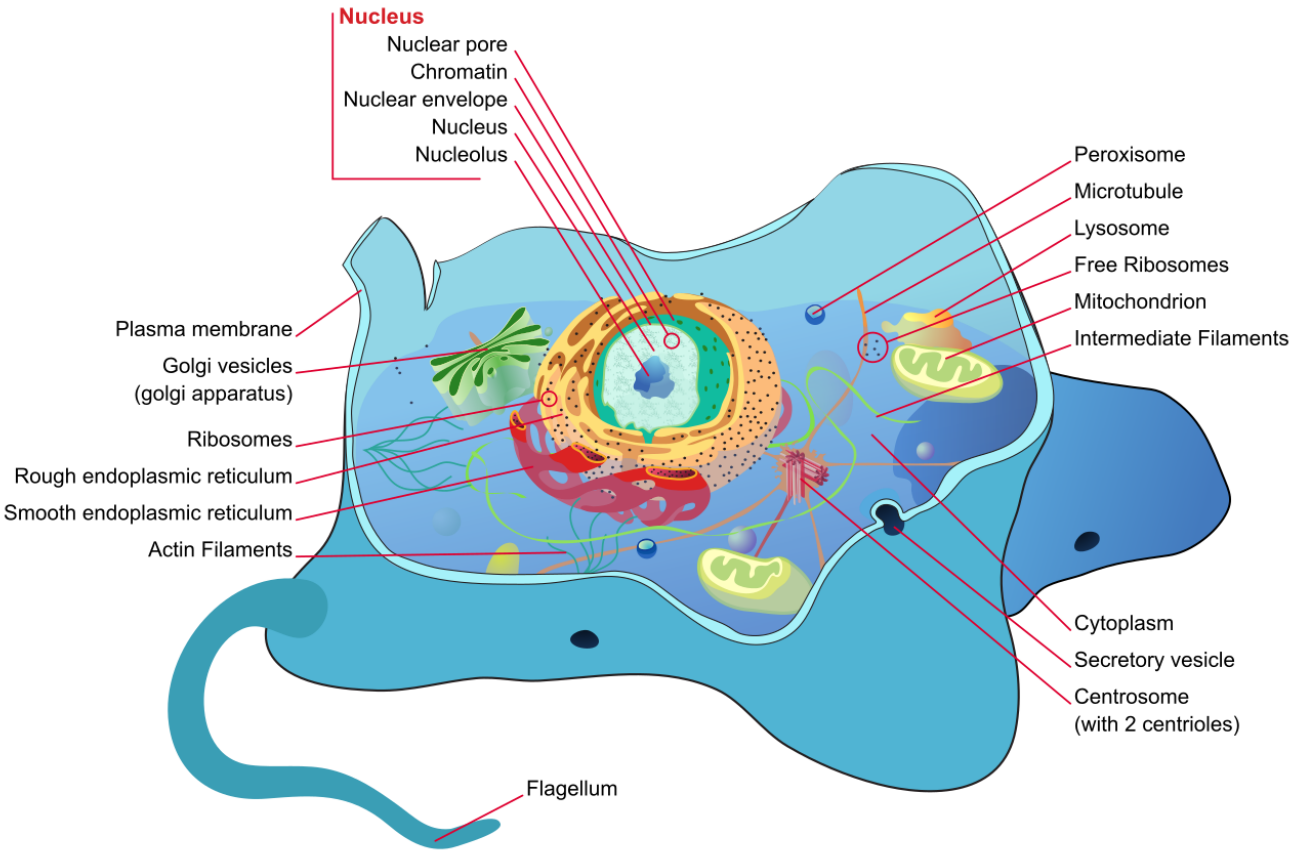


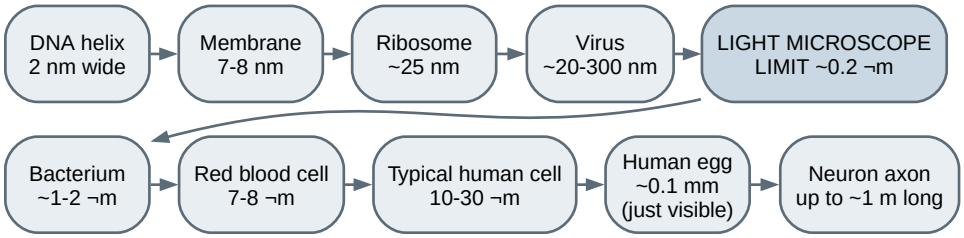
Figure 5. A generalised animal cell. Real cells vary enormously in shape and in which structures are prominent. Image: LadyofHats (Mariana Ruiz), public domain. Credits: Appendix H

Structure	What it does	Why you might care
Plasma membrane	Phospholipid bilayer with embedded proteins; controls what enters and leaves; carries receptors and identity markers (Chapter 3)	Every drug, toxin, and signal acts at or through it. Its proteins are what an immune system or an implant surface "sees"

Structure	What it does	Why you might care
Cytoplasm and cytosol	The gel-like interior and the fluid in it; where much of metabolism happens	Crowded: proteins occupy roughly a fifth to a third of the volume
Nucleus	Holds the chromosomes; site of DNA replication and transcription; separated by a double membrane with pores	About 2 m of DNA packed into a sphere ~5–10 µm across
Nucleolus	Region in the nucleus where ribosomes are assembled	Prominent in cells making lots of protein — a pathologist's clue
Ribosomes	Molecular machines that translate mRNA into protein; free or on the rough ER	Antibiotic targets in bacteria
Rough endoplasmic reticulum (ER)	Membrane network studded with ribosomes; makes proteins destined for membranes or export	The production line for antibodies, hormones, collagen
Smooth ER	Makes lipids; detoxifies drugs (especially in liver); stores calcium (in muscle)	Why the liver handles so many drugs (Chapter 10)
Golgi apparatus	Modifies, sorts, and packages proteins; adds sugar chains	Glycosylation affects how therapeutic proteins behave — a major issue in biomanufacturing
Vesicles	Membrane bubbles carrying cargo between compartments and to the surface	Secretion of hormones, neurotransmitters, enzymes
Lysosomes	Acidic compartments of digestive enzymes; break down worn-out parts and ingested material	Inherited lysosomal enzyme defects cause a family of storage diseases
Peroxisomes	Break down fatty acids and hydrogen peroxide	—
Mitochondria	Produce most of the cell's ATP by aerobic respiration (Chapter 3); have their own small circular DNA, inherited from the mother	Why tissues fail within minutes without oxygen (Chapter 11); the evidence that they were once free-living bacteria
Cytoskeleton	Three filament systems: **actin** (shape, movement, contraction), **microtubules** (tracks for transport, the spindle in division), **intermediate filaments** (mechanical strength)	Cells sense and respond to mechanical load through it — the basis of bone and tissue remodelling (Chapter 6, Chapter 5)
Centrosome	Organises microtubules; duplicates before division	—

Structure	What it does	Why you might care
Cilia and flagella	Motile or sensory projections	Cilia clear mucus from the airway (Chapter 9); sperm swim with a flagellum

4.5.1. Scale



Each step to the right is roughly a factor of three to ten. The light microscope cannot resolve anything left of the highlighted box, which is why viruses and molecular machines were invisible until electron microscopy, and why so much of molecular biology is inferred from indirect measurement rather than seen.

4.6. Membranes and transport, briefly

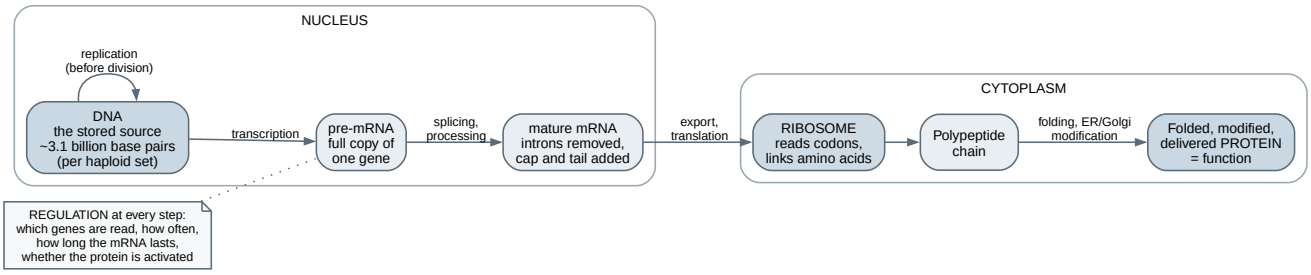
The membrane and the ways things cross it are covered in Chapter 1 (diffusion, osmosis, tonicity, active transport) and Chapter 3 (the bilayer itself). Two points to add here.

Transport proteins are specific and saturable. A channel or carrier handles particular molecules and has a maximum rate. That is why glucose appears in urine once blood glucose exceeds the kidney's reabsorption capacity (Chapter 10), and why a drug that uses a transporter can be blocked by another drug competing for it.

Endocytosis and exocytosis move bulk cargo. Cells engulf particles in membrane pockets (**endocytosis**; **phagocytosis** when the cargo is a bacterium) and release secretions by fusing vesicles with the surface (**exocytosis**). Nanoparticle drug delivery and many viruses enter cells by hijacking endocytosis.

4.7. From DNA to protein: the central dogma in working detail

Chapter 3 gave the one-line version: DNA is transcribed to RNA, which is translated to protein. Chapter 11 covers inheritance and genetic disease. Here is the machinery in between.



4.7.1. DNA and its replication

DNA is a double helix of two strands, each a chain of four nucleotides — **A, T, G, C** — with A pairing to T and G to C across the strands. Because each strand determines its partner, either strand is a template for rebuilding the other. Replication is **semi-conservative**: each new double helix keeps one old strand and gains one new.

The human genome is about **3.1 billion base pairs** per set of chromosomes; most cells carry two sets (46 chromosomes), about 2 m of DNA. It holds roughly **20,000 protein-coding genes**, but those coding sequences are only about 1–2% of the total. The rest includes regulatory sequences, genes for functional RNAs, repeated elements, and a great deal whose function is still debated.

Copying it is a feat of process engineering:

- **DNA polymerase** adds nucleotides at tens of bases per second per fork in human cells, so replication starts simultaneously at tens of thousands of **origins**. S phase still takes several hours.
- Polymerase **proofreads** as it goes, and a separate **mismatch repair** system checks behind it. The overall error rate is roughly **one mistake per billion to ten billion bases** copied.
- Chromosome ends are hard to copy completely, so they carry protective repeated caps, **telomeres**, which shorten with each division in most cells (see "Ageing" below).
- DNA is also damaged constantly — by UV light, reactive oxygen, radiation, and chemicals — and repaired continuously. Defects in repair genes are behind several inherited cancer syndromes (Chapter 11), and radiotherapy works partly by overwhelming repair (Chapter 17).

4.7.2. Transcription and processing

To use a gene, the cell copies it into **messenger RNA** (mRNA). **RNA polymerase** binds at a **promoter**, unwinds the DNA, and builds a complementary RNA strand (with U in place of T).

In eukaryotes the first transcript is processed before it leaves the nucleus:

- **Introns** — non-coding stretches inside the gene — are cut out and the **exons** joined, a process called **splicing**. **Alternative splicing** lets one gene produce several protein variants, one reason 20,000 genes can make a much larger number of distinct proteins.
- A protective cap and a poly-A tail are added, which affect how long the mRNA survives.

4.7.3. Translation and the genetic code

In the cytoplasm, a **ribosome** reads the mRNA three bases at a time. Each triplet — a **codon** — specifies one amino acid, delivered by an adaptor molecule, **transfer RNA** (tRNA).

- There are 4^3 = **64 codons**: 61 specify the 20 amino acids and 3 are **stop** signals. **AUG** both codes for methionine and marks the usual **start**.
- The code is **redundant**: most amino acids have several codons, so many single-base changes ("silent" mutations) do not change the protein.
- The code is **nearly universal** across life. That is what allows a bacterium to make human

insulin from a human gene — the foundation of biotechnology.

The chain then folds (Chapter 3), often with help from **chaperone** proteins, and many proteins are further modified — cut, phosphorylated, glycosylated — before they work. A mis-folded protein may be destroyed, or may aggregate; protein aggregation is central to several neurodegenerative diseases.



Think of it in software terms, then notice where the analogy breaks.

DNA is source code in long-term storage; mRNA is a working copy checked out for execution; the ribosome is an interpreter; the protein is the running process. Splicing is a build step that removes comments. Regulation is configuration.

It breaks in three important places. Nothing in the cell is executed sequentially by a single processor — thousands of copies run in parallel, at rates set by chemistry. "Bugs" are often partial: a mutation may reduce function by 30% rather than crash the program. And the "configuration" — which genes are on in which cell — is as heritable within a body as the code, which is the subject of the next section.

4.7.4. Gene regulation and epigenetics

Every cell in your body carries essentially the same genome. A neuron and a liver cell differ because they **use different genes**. Control happens at every step:

- **Transcription factors** — proteins that bind regulatory DNA sequences — switch genes on or off, often in combination. Hormones such as steroids work largely by activating transcription factors (Chapter 10).
- **Chromatin packing** makes regions of DNA accessible or not.
- **Epigenetic marks** — chemical tags on DNA (methylation) and on the histone proteins it wraps around — lock in patterns of gene use and are copied when the cell divides. They are how a liver cell's daughters stay liver cells. They change with age and environment, and abnormal patterns are a feature of many cancers.
- After transcription, small RNAs and mRNA stability tune how much protein is made; after translation, proteins are switched on and off by modification and destroyed on schedule.

4.8. The cell cycle, division, and death

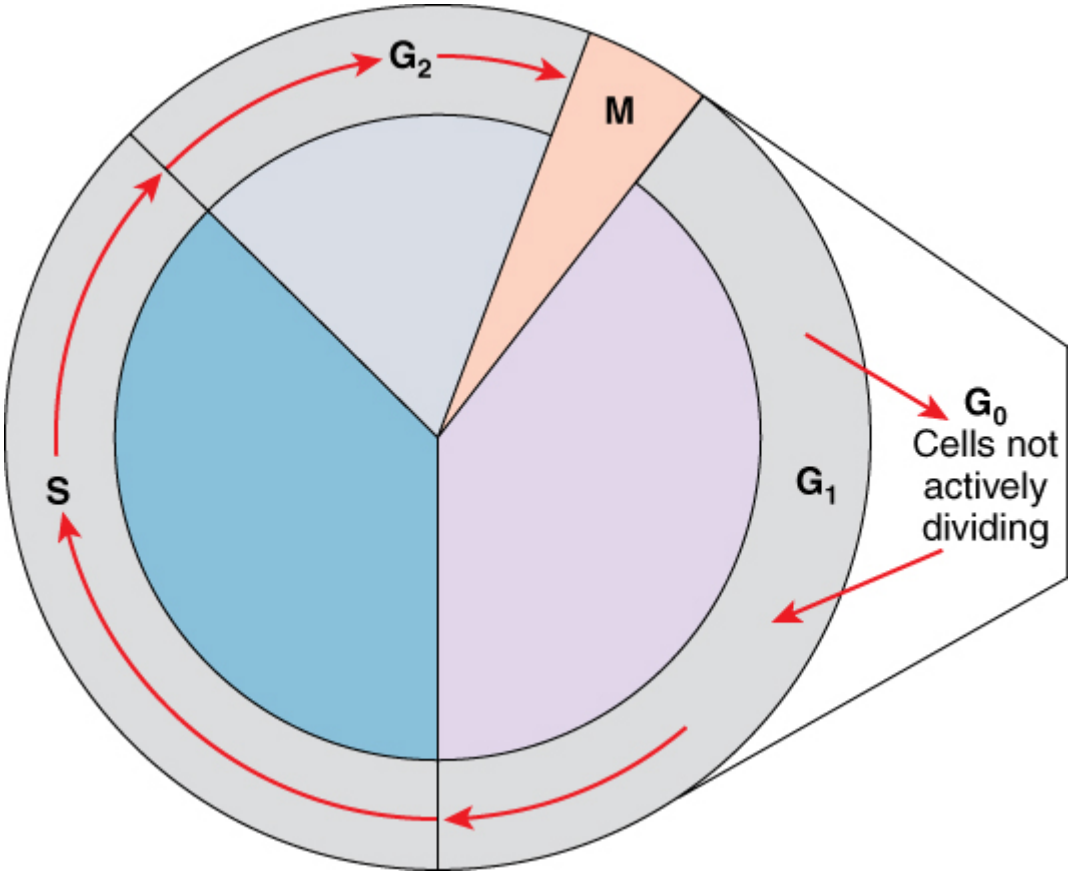


Figure 6. The cell cycle. Cells spend most of their time in interphase (G1, S, G2), growing and copying their DNA, and a short time in mitosis (M). Many cells leave the cycle into G0. Image: OpenStax, CC BY 4.0. Credits: Appendix H

- **G1** — growth and normal work. A checkpoint decides whether to commit to division.
- **S** — DNA is replicated.
- **G2** — preparation; another checkpoint checks the copy.
- **M** — **mitosis**: chromosomes condense, a microtubule **spindle** pulls the copies apart, and the cell splits (**cytokinesis**) into two genetically identical daughters.
- **G0** — a resting state outside the cycle. Most mature neurons and heart muscle cells stay there for life, which is why brain and heart damage is largely permanent.

For a typical human cell in culture the full cycle takes about a day: very roughly G1 ~11 h, S ~8 h, G2 ~4 h, M ~1 h. Rates in the body vary from hours (gut lining, blood-forming cells) to never.

Checkpoints and their genes — notably the tumour suppressor **p53** — stop a cell dividing with damaged DNA. Cancer is, at the cell level, the failure of that control (Chapter 11). It is also why chemotherapy and radiotherapy hurt fast-dividing normal tissues — hair, gut lining, bone marrow — as well as tumours.

Meiosis is the specialised double division that makes eggs and sperm. It halves the chromosome number (46 to 23) and shuffles the parental chromosomes, so every gamete is a new combination. That shuffling, together with mutation, is where inherited variation comes from; errors in it cause conditions such as trisomy 21.

Cell death is part of normal life. **Apoptosis** is tidy, programmed removal; **necrosis** is uncontrolled death from injury. Both are covered in Chapter 11.

4.9. Specialisation and stem cells

A fertilised egg becomes about 30–40 trillion cells of a few hundred recognised types (modern cell atlases distinguish many more). It does so by **differentiation**: progressive, largely stable changes in which genes are used.

Stem cells are cells that can both renew themselves and give rise to specialised cells.

Type	Potential	Example and use
Totipotent	Every cell type, including placenta	The fertilised egg and its first few divisions
Pluripotent	Every cell type of the body	Embryonic stem cells; **induced pluripotent stem cells** (iPSCs) reprogrammed from adult skin or blood cells — used in research, disease modelling, and early clinical trials
Multipotent	Several types within a family	**Haematopoietic stem cells** make all blood cells; bone marrow and stem cell transplants are the longest-established stem cell therapy. Mesenchymal stromal cells (bone, cartilage, fat)
Unipotent / progenitor	One type, or a narrow range	Skin and gut lining renewal



"Stem cell therapy" is also the name of a large unregulated market. Outside a small number of established treatments and properly run clinical trials, clinics offering stem cell injections for arthritis, autism, or ageing are selling unproven and sometimes harmful procedures. If a project brushes against this area, the question to ask is: **is this an approved therapy or a registered trial?**

4.10. Microbes at the cell level

4.10.1. Bacteria

A bacterium is a self-sufficient prokaryotic cell: membrane, cell wall, a single circular chromosome, often plasmids, ribosomes, and sometimes flagella for swimming and **pili** for attachment.

- **Gram staining** — the classic first test in a microbiology lab — splits bacteria by cell wall. **Gram-positive** bacteria have a thick peptidoglycan wall that holds a purple dye; **Gram-negative** bacteria have a thin wall plus an outer membrane, stain pink, and are often harder to treat because that outer membrane keeps drugs out.
- Some bacteria form **spores** — dormant, extremely resistant forms that survive disinfectants and heat that kill active cells (Chapter 13).
- **Plasmids** are small DNA circles that bacteria pass to each other — including across species.

Resistance genes travel on them, which is why resistance can spread far faster than mutation alone would allow.

Growth in a fresh medium follows a characteristic curve:

1. **Lag phase** — adjusting to the new environment; little division.
2. **Exponential (log) phase** — doubling at a constant interval: ~20 minutes for **E. coli** in ideal conditions; 15–20 hours for the tuberculosis bacterium.
3. **Stationary phase** — nutrients run out and waste accumulates; division balances death.
4. **Death phase.**

Biofilms are communities of bacteria attached to a surface and wrapped in a self-made matrix of sugars, proteins, and DNA. A biofilm can tolerate antibiotic concentrations hundreds of times higher than free-floating bacteria and resists cleaning (Chapter 13). Biofilms on catheters, implants, endotracheal tubes, water lines, and the internal channels of endoscopes are a major cause of device-associated infection. **A surface that bacteria can colonise and cleaning cannot reach is the design problem.**

Most bacteria are harmless or helpful. The human gut carries tens of trillions of them, a number of the same order as our own cells, contributing to digestion and immune development (Chapter 10).

4.10.2. Viruses

A virus is genetic material — DNA or RNA — packed in a protein coat, sometimes wrapped in a lipid **envelope** taken from the host cell. Viruses range from about 20 to 300 nm. They:

1. attach to specific receptors on a host cell (which is why each virus infects particular species and tissues);
2. enter, release their genome, and hijack the cell's machinery to copy it and make coat proteins;
3. assemble new particles and leave, often killing the cell.

Retroviruses such as HIV carry RNA and use an enzyme, **reverse transcriptase**, to write it into DNA that integrates into the host genome. Biotechnology borrows both tricks: reverse transcriptase is a standard laboratory tool (it is how RNA viruses are detected by PCR), and engineered viruses are used to deliver genes in gene therapy and CAR-T manufacture.

Enveloped viruses are easier to inactivate with detergents and alcohol than non-enveloped ones, a difference that matters when choosing disinfectants (Chapter 13).

4.11. The immune system at the cell level

The immune system is a distributed sensing and response network built from cells.

Cell	Role	Note
Neutrophils	First responders; engulf and kill bacteria	The most common white cell; pus is largely dead neutrophils

Cell	Role	Note
Macrophages	Engulf debris and microbes; signal to other cells; coordinate repair	Also what surrounds and walls off a foreign body — the **foreign body response** to an implant
Dendritic cells	Capture fragments of invaders and present them to T cells	The link between innate and adaptive immunity; the target of vaccine design
Natural killer (NK) cells	Kill infected or abnormal cells without prior training	—
T cells — helper (CD4) and cytotoxic (CD8)	Helpers coordinate the response; cytotoxic cells kill infected cells	HIV destroys helper T cells. CAR-T therapy re-engineers T cells to attack cancer
B cells and plasma cells	Make **antibodies**, proteins that bind a specific target	The basis of immunoassays (Chapter 15) and monoclonal antibody drugs
Memory cells	Remember a previous encounter for a faster, stronger second response	Why vaccines work

Innate immunity — barriers, inflammation, and the first four rows — responds within hours to anything that looks foreign. **Adaptive** immunity — T and B cells — takes days to learn a specific target but then remembers it. Allergy and autoimmune disease are this system aiming at the wrong target (Chapter 11).

4.12. Metabolism, development, and ageing, in one paragraph each

Metabolism is the network of chemical reactions that extracts energy from food and builds the cell's components. Its energy side — glycolysis, the citric acid cycle, oxidative phosphorylation, and ATP — is in Chapter 3. What to add here is that metabolic differences are diagnostic: cancers' heavy use of glucose is what FDG-PET images (Chapter 16).

Development runs from a single cell through an embryo in which cells divide, move, signal to each other, and differentiate in a coordinated programme. Many birth defects are timing or signalling errors in that programme, and many adult cancers reactivate parts of it.

Ageing at the cell level includes **telomere shortening** (most human cells in culture stop dividing after roughly 40–60 divisions — the **Hayflick limit**), accumulated DNA damage, **senescent** cells that stop dividing and release inflammatory signals, and declining stem-cell function. The consequences at the level of the whole person are in Chapter 12.

4.13. Engineering bridge: cells by the numbers

4.13.1. Why cells are small: surface and volume

A cell imports everything through its surface and uses it throughout its volume. For a sphere of radius **r**:

surface ÷ volume = $4\pi r^2 \div (4/3)\pi r^3 = 3 / r$

Double the radius and the surface available per unit of volume halves. A bacterium ($r \approx 0.5$ μm) has SA:V ≈ 6 μm^{-1}; a 20 μm human cell has 0.3 μm^{-1} — twenty times less. Large cells compensate by being flat (red cells), long and thin (neurons, muscle fibres), or folded (microvilli in the gut, Chapter 10). The same arithmetic governs heat sinks, catalysts, and filters.

4.13.2. Why tissues need blood vessels: diffusion time

The time for a molecule to diffuse a distance **x** scales with the **square** of the distance:

$t \approx x^2 / 2D$

For oxygen in water, $D \approx 2 \times 10^{-9}$ m^2/s:

Distance	Approximate time	Meaning
10 μm (one cell)	~25 ms	Effectively instantaneous
100 μm	~2.5 s	Still adequate
1 mm	~4 min	Too slow for living tissue
1 cm	~7 h	Impossible

Most cells in the body sit within roughly 100–200 μm of a capillary. That one fact explains why tissue-engineered constructs thicker than a few hundred micrometres die in the middle unless they are perfused, why a thick wound dressing or a tight cast can starve tissue (Chapter 5), why cartilage, with no blood supply, heals so poorly (Chapter 6), and why organoids stop growing at a few millimetres. Diffusion is fast over a cell and hopeless over a centimetre; circulation exists to close the gap.

4.13.3. Exponential growth: bacteria, cell culture, and PCR

With doubling time **T**, a population after time **t** is $N = N_0 \times 2^{t/T}$.

- **Bacteria:** one cell doubling every 20 minutes becomes $2^{30} \approx 10^9$ cells in 10 hours. That is why a contaminated surface or solution left warm overnight is a different object in the morning, and why culture-based diagnosis takes "overnight".
- **Cell culture:** human cells doubling every ~24 hours go from a million to a billion in about ten days — the timescale of manufacturing a CAR-T dose.
- **PCR:** each thermal cycle ideally doubles the target DNA, so 30 cycles give $2^{30} \approx 10^9$ copies. Real efficiency is less than perfect: at 90% per cycle the multiplier is $1.9^{30} \approx 2 \times 10^8$ — five times less. Quantitative PCR uses the number of cycles needed to reach a threshold to estimate the starting amount, which is why efficiency must be measured with a standard curve, not assumed.

4.13.4. The genome as information

Four bases carry 2 bits each, so a haploid human genome of 3.1×10^9 bases is about 6.2×10^9 bits ≈ **775 MB** uncompressed — less than a single film download, and mostly shared with everyone else. Two people differ at roughly one base in a thousand.

Copy fidelity: about 10^{-9} to 10^{-10} errors per base per division after repair. That is good for chemistry and poor by storage-engineering standards, where uncorrectable error rates of 10^{-15} per bit are routine. Biology tolerates it — and depends on it, because those errors are the variation evolution selects from. DNA is also the densest durable storage medium known, which is why DNA data storage is an active research field.

4.13.5. Seeing small things: resolution and signal

Light microscopes are limited by diffraction. The smallest resolvable separation is approximately:

$d \approx \lambda / (2\,NA)$

where λ is the wavelength and NA the numerical aperture of the objective (up to ~1.4 for oil immersion). For green light ($\lambda \approx 500$ nm), $d \approx 180$–250 nm. Bacteria are visible; viruses are not. **Super-resolution** techniques beat this limit with clever photophysics (tens of nanometres); **electron microscopes** use electrons with picometre wavelengths and resolve down to fractions of a nanometre; **cryo-electron microscopy** now determines protein structures at near-atomic resolution.

Fluorescence imaging is a photon-counting measurement, so the √N law of Chapter 16 applies: halving the noise needs four times the light. But more excitation light **bleaches** the dye and **damages** living cells. Every live-cell imaging protocol is a negotiation between signal, speed, resolution, and keeping the cells alive — the same trade as dose in radiography.

4.13.6. Spinning things down: relative centrifugal force

Centrifuges are specified in multiples of gravity, not rpm, because the force depends on the radius:

$RCF \approx 1.118 \times 10^{-5} \times r\,(cm) \times rpm^2$

At 3,000 rpm and a 10 cm radius, RCF ≈ 1,000 g — enough to pellet cells. Separating organelles needs 10,000–100,000 g; **ultracentrifuges** reach several hundred thousand g and separate molecules by size and density. A protocol written in rpm is not transferable between rotors; one written in g is. It is a units lesson as clean as any in Appendix C.

4.14. The field in practice: what biologists do, and the tools they use

Biology is not one activity. The table below maps the main branches to what they are **for** in the real world, so that a job advert, a grant, or a hospital department name makes sense.

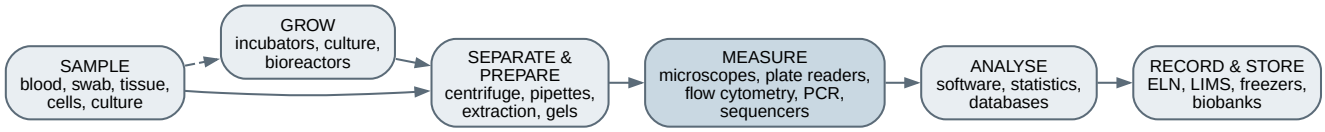
Branch	Real-world use	Where you meet it
Cell biology	How cells work, divide, move, and die; drug effects on cells; toxicity testing	Pharmaceutical research; cancer biology; materials biocompatibility testing (cytotoxicity is ISO 10993-5, Appendix D)

Branch	Real-world use	Where you meet it
Molecular biology	Manipulating and measuring DNA, RNA, and protein	Diagnostics (PCR tests), research, biotechnology — the toolkit underneath most of this table
Genetics and genomics	Inheritance; reading whole genomes	**Newborn bloodspot screening**; diagnosis of rare disease by genome sequencing; **tumour sequencing** to choose targeted cancer drugs; pharmacogenomics (Chapter 11); ancestry and forensics
Microbiology	Identifying organisms and their drug sensitivities	Hospital microbiology labs; infection control (Chapter 13); food, water, and pharmaceutical quality control; sterility testing of devices
Virology	Viruses, their detection, and their control	Respiratory and blood-borne virus testing; vaccine development; public health surveillance including wastewater monitoring
Immunology	The immune response, measured and manipulated	Vaccines; **monoclonal antibody** drugs; immunoassay diagnostics (Chapter 15); transplant matching; allergy testing
Biotechnology and biomanufacturing	Using living cells as factories	**Recombinant insulin** (made in bacteria or yeast since the 1980s); antibodies made in mammalian cell bioreactors; vaccines including **mRNA vaccines**; enzymes for detergents and food
Cell and gene therapy	Cells or genes as the medicine	**CAR-T** therapy for blood cancers (patient's T cells engineered and re-infused); gene therapies for inherited diseases; the first approved **CRISPR**-based therapy (for sickle cell disease and beta thalassaemia, approved 2023)
Synthetic biology	Engineering organisms with designed genetic circuits	Biosensors; microbes producing drug precursors; standardised genetic "parts"
Tissue engineering and regenerative medicine	Building or regrowing tissue	Skin substitutes for burns (Chapter 5); engineered cartilage and bone scaffolds; bioprinting research

Branch	Real-world use	Where you meet it
Organoids and organ-on-chip	Miniature 3D tissues; microfluidic chips lined with living cells under flow	Drug testing and disease models that may reduce animal testing; personalised drug screening on a patient's tumour organoids
Neuroscience	The nervous system from molecules to behaviour	Chapter 7; brain–computer interfaces; neurodegenerative disease research
Bioinformatics and computational biology	Software and statistics for biological data	Genome analysis pipelines; protein structure prediction; image analysis; the most direct entry point for a software-literate engineer
Epidemiology and public health biology	Disease in populations; environmental biology	Outbreak tracking; vaccination programmes; antimicrobial resistance surveillance

4.14.1. A laboratory workflow

Almost every experiment, research or diagnostic, follows the same shape. Tools are easiest to understand by where they sit in it.



4.14.2. Tools for seeing

Tool	What it does	Typical use
Brightfield light microscope	Transmitted light; usually needs staining for contrast	Stained tissue (histopathology), blood films, Gram stains
Phase-contrast and DIC microscopy	Optical tricks that turn tiny refractive-index differences into contrast	Living, unstained cells in culture
Fluorescence microscope	Excites fluorescent dyes or proteins with one wavelength, images the emission at another	Labelling specific proteins, DNA, or cell types; many diagnostic assays
Confocal microscope	Laser scanning with a pinhole that rejects out-of-focus light; optical sectioning	3D images of cells and thick tissue
Light-sheet microscope	Illuminates one thin plane from the side; fast, gentle 3D imaging	Whole embryos, cleared organs, long live imaging
Super-resolution microscopy	Several techniques beating the diffraction limit to tens of nanometres	Molecular organisation inside cells

Tool	What it does	Typical use
Electron microscopy (TEM, SEM)	Electron beam; nanometre resolution; specimens fixed and in vacuum	Viruses, organelles, surfaces of materials and implants
Cryo-electron microscopy	Frozen-hydrated specimens imaged thousands of times and averaged	Protein and virus structures at near-atomic resolution
Live-cell imaging systems	Microscope inside an incubator enclosure, automated over hours or days	Cell division, migration, drug responses over time
Digital slide scanners	Whole microscope slides digitised at high resolution	Digital pathology and AI-assisted diagnosis (Chapter 15)

4.14.3. Tools for growing

Tool	What it does	Note
Biological safety cabinet (BSC)	Enclosed workspace with HEPA-filtered airflow; Class II protects the worker, the sample, and the room	**Not** a fume hood (which protects only the worker from chemical vapours) and not a clean bench (which protects only the sample)
CO_2 incubator	Holds 37 °C, ~5% CO_2 (to buffer the culture medium's pH), and high humidity	Temperature, gas, and humidity control plus contamination control — an embedded-control product
Culture vessels and media	Flasks, plates, and dishes with treated surfaces; nutrient media with buffers, salts, sugars, amino acids, often serum	Surface chemistry decides whether cells attach (Chapter 3)
Shaking incubators and fermenters	Grow bacteria and yeast in liquid with aeration	Recombinant protein production
Bioreactors	Stirred or rocking vessels, single-use bags to thousands of litres, with sensors for temperature, pH, dissolved oxygen, and feeds	Antibody and vaccine manufacture. Oxygen transfer and mixing without shearing the cells is the central engineering problem
Autoclaves	Steam sterilisation of media, glassware, and waste	Chapter 13

4.14.4. Tools for separating and preparing

Tool	What it does	Note
Micropipettes	Measure and transfer microlitre volumes with disposable tips	The most-used instrument in biology; accuracy depends on calibration and technique

Tool	What it does	Note
Liquid-handling robots	Automated pipetting across 96-, 384-, or 1,536-well plates	Throughput and reproducibility; open-source designs exist
Centrifuges and ultracentrifuges	Separate by density and size (see the RCF calculation above)	Blood separation, cell pelleting, DNA preparation
Nucleic acid extraction	Break cells open and purify DNA or RNA, often on silica columns or magnetic beads	The first step in every PCR test and sequencing run
Gel electrophoresis	An electric field pulls DNA or protein through a gel; small fragments travel further	Checking a PCR product's size; separating proteins (SDS-PAGE). Runs at ~100–150 V DC — a real electrical hazard
Protein chromatography	Separates proteins by size, charge, or specific binding	Purifying therapeutic proteins and antibodies
Flow sorting (FACS)	Flow cytometer that physically diverts selected cells into tubes	Isolating specific cell populations, e.g. for CAR-T manufacture

4.14.5. Tools for measuring and counting

Tool	What it measures	Note
Spectrophotometer	Absorbance at chosen wavelengths; DNA concentration at 260 nm, bacterial density as optical density at 600 nm	Beer–Lambert law (Chapter 15)
Plate reader	Absorbance, fluorescence, or luminescence in every well of a microplate	ELISAs, cell viability, enzyme assays, high-throughput screening
Flow cytometer	Thousands of cells per second pass single-file through lasers; scatter and fluorescence classify each cell	Immune cell counts (e.g. CD4 counts in HIV), leukaemia diagnosis, research
Automated cell counters and haematology analysers	Count and size cells by imaging or electrical impedance	The full blood count (Chapter 15)
Western blot	Proteins separated on a gel, transferred to a membrane, detected with antibodies	Confirming a specific protein is present
ELISA and immunoassays	Antibody-based detection with a colour, light, or fluorescence readout	Hormones, infection markers, drug levels; the lateral flow test is a simplified version
Mass spectrometry	Molecules identified by mass-to-charge ratio	Proteomics; bacterial identification in minutes (MALDI-TOF)

Tool	What it measures	Note
Electrophysiology rigs	Measure tiny currents through cell membranes (patch clamp)	Ion channel research; drug cardiac-safety testing (Chapter 7)

4.14.6. Tools for reading and editing genes

Tool	What it does	Note
PCR and qPCR thermocyclers	Amplify a chosen DNA sequence by repeated heating and cooling; qPCR monitors fluorescence in real time	Infection diagnosis (with reverse transcription for RNA viruses), genetic testing, forensics. Thermal uniformity and ramp rate are the engineering problems
Sanger sequencing	Reads one DNA fragment of up to ~1,000 bases very accurately	Confirming specific variants; checking engineered constructs
Short-read (next-generation) sequencing	Billions of short reads (~100–300 bases) in parallel	Whole genomes, exomes, tumour panels, RNA-seq. A human genome now costs hundreds of dollars in reagents at scale
Long-read sequencing	Reads of tens of thousands of bases or more, by single-molecule methods; nanopore devices measure current changes as DNA passes through a pore and can be pocket-sized	Structural variants, complete genomes, rapid field sequencing of pathogens
Microarrays	Hundreds of thousands of DNA probes on a chip	Genotyping, chromosomal microarray testing in developmental disorders
Cloning and plasmids	Cut-and-paste of DNA into small circular carriers that bacteria copy	Making the DNA for recombinant proteins, vaccines, and gene therapy
CRISPR–Cas systems	A guide RNA (~20 bases) directs a nuclease to cut DNA at a chosen sequence; the cell's repair then disables or rewrites the gene. Variants edit single bases without cutting	Research everywhere; approved gene therapy since 2023; also used in rapid diagnostics

4.14.7. Tools for storing, automating, and keeping records

Tool	What it does	Note
−20 °C and −80 °C freezers; liquid nitrogen (−196 °C)	Long-term storage of samples, reagents, and living cells	Freezer failure alarms and temperature logging are classic monitoring problems
Biobanks	Large curated collections of samples with linked data and consent	A governance-heavy data and logistics system (Chapter 20)

Tool	What it does	Note
Laboratory automation	Robotic arms, plate handlers, scheduling software linking instruments	Diagnostic labs process thousands of samples a day this way
LIMS — laboratory information management system	Tracks samples, tests, results, and chain of custody	The lab's equivalent of the hospital record; integrates with the EHR (Chapter 20)
ELN — electronic lab notebook	Records experiments, protocols, and results	Version control for experiments; tamper-evident records

4.14.8. Software

Software	Does	Note
NCBI databases, Ensembl, UniProt	Public reference genomes, genes, variants, and protein records	Free; the starting point for any sequence question
BLAST	Finds similar sequences in databases	"What is this sequence?" in seconds
Biopython, Bioconductor ®	Libraries for sequence handling, statistics, and genomic analysis	The everyday tools of bioinformatics
Galaxy	Web platform for running analysis pipelines without programming	Reproducible workflows; free public servers
Benchling and similar	Commercial cloud platform for sequence design, ELN, and registry	Common in biotech companies
ImageJ / Fiji, CellProfiler, napari	Open-source biological image analysis: measuring, counting, segmenting, tracking	**Directly usable today** on public image sets
FlowJo and open alternatives	Flow cytometry analysis by "gating" populations	Manual gating is subjective — an automation and reproducibility opportunity
AlphaFold and structure databases	Predicted 3D structures of proteins from their sequence	Recognised by the 2024 Nobel Prize in Chemistry; predictions cover essentially all known proteins
PyMOL, ChimeraX	Visualise and analyse molecular structures	Pair with the Protein Data Bank [33]

4.14.9. Biosafety: containment levels

Laboratories that work with living organisms are classified by **biosafety level** according to the risk of the organisms they handle [74] [75]:

Level	Organisms (examples)	Key controls
BSL-1	Not known to cause disease in healthy adults: lab strains of **E. coli**, baker's yeast	Good microbiological practice; open bench; hand washing; decontamination of waste
BSL-2	Moderate risk, treatable: **Staphylococcus aureus**, **Salmonella**, many clinical samples	Restricted access; Class II safety cabinets for aerosol-generating work; PPE; autoclave available; trained staff
BSL-3	Serious, potentially airborne: tuberculosis, some respiratory viruses	Controlled negative-pressure facility; airlocks; respiratory protection; strict procedures
BSL-4	Dangerous and exotic, no treatment: Ebola, Marburg	Full-body positive-pressure suits or Class III cabinets; dedicated, isolated buildings

Hospital diagnostic laboratories work largely at BSL-2, with BSL-3 suites for some organisms. Current international guidance emphasises local risk assessment rather than the level label alone.



Do not culture microorganisms or human cells outside a properly supervised laboratory.

A plate of "harmless" bacteria grown from a swab of a surface is a culture of unknown organisms, potentially including pathogens, at numbers far higher than the original sample. Human cells and blood are potential carriers of blood-borne viruses. Disposal of cultures requires sterilisation, not the bin.

Community biology labs exist that operate at BSL-1 with trained supervision, safety rules, and proper waste handling. That is the right setting for hands-on experience. Equipment you build for a lab — however good — is used under that lab's biosafety rules, not yours.

4.15. Devices, software, and training technology

Biology laboratories are full of instruments that are expensive, proprietary, or poorly suited to teaching and low-resource settings. Open hardware has become a serious movement in response, and it is one of the places a technologist with optics, electronics, machining, and software skills contributes most directly.

4.15.1. Clinical and laboratory devices, with maker entry points

Device area	What exists	Realistic technologist contribution
Microscopy	Research and clinical microscopes; digital slide scanners	**Open-source microscopes** (such as the OpenFlexure design) combine machined or printed flexure stages, off-the-shelf optics, a camera module, and software; used in teaching and in low-resource diagnostic research. Custom stage adapters, phone-camera mounts, illumination modules (LED, Köhler optics) are all well within reach
Temperature and environment control	Incubators, heat blocks, water baths, thermocyclers	Temperature control is a classic embedded-systems problem: sensors, PID control, logging, alarms, and enclosures. Open-source thermocycler and incubator designs exist for teaching
Liquid handling and automation	Commercial robots from bench-top to whole-lab systems	Open-source liquid handlers and protocol software are established; custom labware holders, racks, and adaptors are constant needs — machined, moulded, or printed in a chemically resistant material
Separation	Centrifuges, electrophoresis tanks and power supplies	**Centrifuges store dangerous kinetic energy** and electrophoresis uses lethal-capable voltages in conductive buffer. Build accessories, not the rotor or the power supply
Imaging enclosures and gel documentation	Commercial gel imagers and transilluminators	Light-tight enclosures with blue-LED excitation and amber filters — laser-cut, sheet-metal, or machined — with a camera and software. See the worked example
Bioreactors	Single-use and stainless bioreactors with sensor suites	Teaching bioreactors with instrumented pH, temperature, and dissolved-oxygen sensing are an excellent integrated electronics, fluidics, and control project
Scientific glassware	Flasks, condensers, custom vessels, flow cells	Scientific glassblowing is a recognised specialist trade; custom borosilicate vessels for phantoms and demonstrations (Chapter 14)

Device area	What exists	Realistic technologist contribution
Microfluidics	Chips for organ-on-chip and point-of-care diagnostics	PDMS casting from machined or lithographic moulds, laser-cut laminated chips, and flow control electronics — a fertile research-collaboration area



Open hardware in a clinical laboratory is still a medical device question. A research lab can use a home-built instrument with appropriate checks. A diagnostic laboratory reporting results on patients works under accreditation (for example ISO 15189) and, in many jurisdictions, in vitro diagnostic regulation. An instrument used there must be validated for that purpose (Chapter 15, Chapter 19). The realistic early roles are teaching, research support, and accessories that do not touch the measurement.

4.15.2. Software

Software area	What it does	Technologist entry point
Image analysis	Segment, count, track, and measure cells and tissues	Fiji macros, CellProfiler pipelines, Python with scikit-image or napari on public image sets — immediately accessible
Sequence analysis and pipelines	Alignment, variant calling, expression analysis	Workflow engines, containerised pipelines, and reproducibility tooling are chronic needs
Laboratory data systems	LIMS, ELN, instrument integration	Instrument drivers and data integration; the same interoperability problems as Chapter 20
Instrument control and open protocols	Scripting liquid handlers, microscopes, and plate readers	Python APIs and open protocol formats for automation
Structure and design tools	Protein structure prediction, sequence design	Computational; works on public data with no lab at all

4.15.3. Training and simulation

Training technology	Teaches	Form
Virtual laboratories	Lab technique, safety, and experimental logic	Commercial and free simulations; useful before hands-on time
Physical cell and DNA models	Organelle structure, base pairing, the double helix	Machined, laser-cut, turned, cast, or printed; snap-together base-pair kits

Training technology	Teaches	Form
Microscopy teaching sets	Using a microscope; recognising cell types	Prepared slides plus affordable microscopes — where open-source designs shine
Educational molecular biology kits	DNA extraction, PCR, gel electrophoresis with safe reagents and BSL-1 organisms	Classroom thermocyclers, low-voltage gel systems, blue-light imagers
Pipetting trainers	Accuracy and technique	Gravimetric checking on a balance; dye-based plates read by a camera
Citizen-science games — Foldit, Eterna	Protein and RNA folding	Players have contributed to published structures
Biosafety training	Cabinet use, spill response, waste handling	Courses plus practical assessment



Where this chapter connects you to real work.

- **Image analysis** — counting cells, measuring wound closure in scratch assays, segmenting nuclei — is a software contribution any lab values, and it can be learned on public datasets.
- **Instrument building and adaptation** — open microscopes, environmental control, automation labware, imaging enclosures — combines optics, electronics, machining, and firmware.
- **Teaching equipment** — affordable, robust, safe versions of laboratory instruments — is chronically underserved.

All three start with a conversation with a laboratory scientist about what they actually need (Chapter 18).

4.16. Worked example: twelve gel imagers for a teaching lab

A university teaching lab runs a molecular biology practical: students extract DNA, run PCR, separate the products on an agarose gel, and photograph the bands. The lab has one commercial gel imager for 48 students and asks whether you could build twelve simple ones.

Understand the measurement first. The DNA in the gel is stained with a dye that fluoresces when bound to DNA. Modern "safe" stains are excited by **blue light** (~470 nm) and emit green; older ethidium bromide needs UV. The imager must deliver even excitation light, block that excitation light from the camera with an **amber long-pass filter**, exclude room light, and capture an image a student can save and annotate. The quantity that matters is **sensitivity**: can it show a faint band containing a few nanograms of DNA?

Ask the questions before choosing materials.

1. **Which stain?** Blue-light stains only. A UV transilluminator adds an eye and skin hazard (Chapter 8) the teaching lab does not need.
2. **Where does it sit, and what gets spilled on it?** Electrophoresis buffer and gel fragments. The light surface must be wipeable with the lab's disinfectant and water-resistant (Chapter 13).
3. **What does it connect to?** Students' phones, or a fixed camera on a lab computer? That decides mounting, software, and data handling.
4. **What is it not?** It is not the electrophoresis power supply. The gel is moved to the imager after the run, with the power off. The imager is battery or low-voltage powered and never connected to the gel tank.

Compare methods for the enclosure.

Approach	For	Against
Laser-cut opaque acrylic box, solvent-welded	Fast to make twelve; flat panels suit a box; good finish	Acrylic crazes with some solvents and alcohol wipes — check the lab's cleaning agent (Chapter 14)
Folded and riveted aluminium sheet, powder-coated	Robust, wipeable, survives students	More tooling and finishing time; heavier
Commercial enclosure box, modified	Already sealed and tested	Viewing port and light path harder to integrate
Printed enclosure	Easy to iterate the first one	Twelve large prints are slow, layer lines harbour stain and gel residue, light leaks at seams

The likely answer is a **sheet-metal or acrylic box** (depending on the cleaning agent) with a replaceable glass or polycarbonate stage over an LED array, an off-the-shelf amber filter, a fixed phone or camera cradle, and a simple diffuser. One design, prototyped once, then built as a batch with jigs.

Define acceptance before building. With the lab: "Bands of a standard DNA ladder are visible down to the 5 ng band, with even illumination across a 10 × 10 cm gel, in an image saved to the student's device." Test the first unit against that with a real gel, alongside the commercial imager. Record the failure if there is one — uneven LEDs and filter leakage are typical (Chapter 19).

And the alternative. Before building twelve, ask the price of twelve basic blue-light transilluminators with amber covers and phone stands. If it is close, the right answer may be to buy them and spend your time on the image-analysis worksheet instead.

4.17. Workshop exercise

1. **Look at cells.** With any microscope — a school microscope, a phone with a clip-on lens, or an open-source design you build — look at onion epidermis (stained with a drop of iodine) and a

drop of baker's yeast suspension. Measure cell sizes using a stage micrometer or a printed scale photographed at the same magnification. Compare with the scale figure.

2. **Extract DNA.** Mash strawberries or a banana with a solution of washing-up liquid and salt, filter, and layer ice-cold alcohol on top. Spool the white strands. Write down what each ingredient did (membranes, proteins, solubility). Kitchen materials only; dispose normally.
3. **Measure diffusion.** Set a tray of gelatin or agar, add a drop of food colouring, and photograph it at intervals over a day. Plot spread distance against the square root of time. Estimate D and compare with the oxygen calculation above.
4. **Watch exponential growth safely.** Mix baker's yeast with warm sugar water in several bottles with balloons over the necks, or log the CO_2 with a sensor and a microcontroller. Plot gas production over time and identify lag and exponential phases. Food-grade yeast only, sealed until finished, contents down the drain with plenty of water.
5. **Read a gene.** Find the human insulin gene (**INS**) on NCBI. Use Biopython or an online tool to find the coding sequence, translate it, and compare the protein to the insulin sequence in UniProt. Locate the signal peptide that the ER cleaves off.
6. **Count cells by software.** Download a public fluorescence image of nuclei and count them in Fiji or CellProfiler. Change the threshold by ±10% and record how the count changes. That is Chapter 15's threshold sensitivity, in biology.
7. **Centrifuge arithmetic.** For a bench centrifuge in any datasheet, convert its maximum rpm to RCF at the rotor's maximum radius. Then work out the rpm needed for 500 **g** in a rotor of half that radius.

4.18. Questions

1. List four properties shared by living things. Why are viruses usually treated as non-living, and what follows practically?
2. State the three parts of cell theory, and explain why the third matters for infection control.
3. Why do antibiotics that work against bacteria have no effect on viruses, and why are antifungal drugs harder to make than antibacterial ones?
4. Name five organelles and give one function each.
5. Describe, in order, the steps from a gene to a working protein in a human cell.
6. What does it mean that the genetic code is redundant, and why does its near-universality matter for biotechnology?
7. How can a neuron and a liver cell have the same genome yet be so different?
8. What happens in each phase of the cell cycle, and why do chemotherapy side effects appear in hair, gut, and bone marrow?
9. What is a biofilm, and why is it a device-design problem?
10. Explain why most cells are within about 100–200 µm of a capillary, using a calculation.
11. A PCR runs 35 cycles at 95% efficiency. Roughly how many copies does each starting molecule produce, and why does efficiency matter in qPCR?
12. Why can a light microscope not resolve a virus? What would you use instead?

13. A protocol says "spin at 10,000 rpm". Why is that incomplete, and what should it say?
14. Describe BSL-1 and BSL-2, and explain why you should not grow cultures from swabs at home.

4.19. Answers and reasoning

1. Any four of: organisation into cells, metabolism, homeostasis, growth and development, reproduction with heredity, response to the environment, evolution. Viruses lack their own metabolism and cannot replicate outside a host cell. Practically: they are not cells, so antibiotics do not affect them and they must be controlled by vaccines, antivirals, and infection control.
2. All living things are made of cells; the cell is the basic unit of life; cells arise only from existing cells. The third means contamination must be **introduced** — a sterile item stays sterile unless organisms reach it — which is why breaking a transmission route works.
3. Antibiotics target bacterial structures absent from viruses — cell walls, 70S ribosomes, bacterial enzymes; viruses use the host cell's own machinery. Fungi are eukaryotes like us, so there are fewer fungal targets that human cells lack, and drugs are harder to make selective.
4. Any five, e.g.: nucleus — stores DNA, site of transcription; ribosome — makes proteins; rough ER — makes membrane and secreted proteins; Golgi — modifies and sorts proteins; mitochondria — ATP production; lysosomes — digestion; cytoskeleton — shape, transport, movement.
5. Transcription factors activate the gene; RNA polymerase transcribes it into pre-mRNA; introns are spliced out and a cap and tail added; mRNA is exported; ribosomes translate codons into an amino acid chain; the chain folds and may be modified in the ER and Golgi, then delivered to where it works.
6. Most amino acids are encoded by more than one codon, so many single-base changes do not alter the protein. Because almost all life reads the code the same way, a human gene placed in a bacterium or yeast produces the human protein — the basis of recombinant insulin and much of biomanufacturing.
7. Gene regulation: different transcription factors, chromatin states, and epigenetic marks mean different sets of genes are used in each cell type, and those patterns are inherited when the cell divides.
8. G1: growth and normal function; S: DNA replication; G2: preparation and checking; M: mitosis and cytokinesis; G0: resting outside the cycle. Chemotherapy and radiotherapy hit dividing cells, and hair follicles, gut lining, and bone marrow divide rapidly.
9. A surface-attached community of bacteria in a self-made matrix. It resists antibiotics and disinfectants and forms on catheters, implants, tubing, and inaccessible channels, so design must avoid surfaces and crevices that can be colonised but not cleaned.
10. Diffusion time scales with distance squared: $t \approx x^2/2D$. With $D \approx 2 \times 10^{-9}$ m²/s for oxygen, 100 µm takes ~2.5 s but 1 mm takes ~4 minutes — too slow to supply working cells. So tissue needs capillaries every couple of hundred micrometres.
11. $1.95^{35} \approx 1.4 \times 10^{10}$ (compared with $2^{35} \approx 3.4 \times 10^{10}$ at perfect efficiency). qPCR infers the starting amount from the cycle at which a threshold is crossed; small efficiency differences compound over cycles, so an unmeasured efficiency gives a wrong quantity.
12. Resolution is limited by diffraction to roughly $\lambda/(2NA)$, about 200 nm for visible light, and most viruses are smaller than that. Use electron microscopy (or detect the virus indirectly — by PCR,

antigen tests, or cell effects).

13. Force depends on rotor radius as well as speed. It should state relative centrifugal force in multiples of **g** (and time and temperature), so the protocol transfers between centrifuges.
14. BSL-1: organisms not known to cause disease in healthy adults, handled with good practice on an open bench. BSL-2: moderate-risk organisms and clinical samples, with restricted access, safety cabinets for aerosols, PPE, and autoclaving. A home culture from a swab amplifies unknown organisms, possibly pathogens, by many orders of magnitude, without containment or safe disposal.

4.20. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Biolab — centrifuge RCF, cell counts, PCR amplification and growth.
- Growth curves — logging and fitting growth phases.
- Sequence tools — translating a gene such as INS.
- Sequence alignment — comparing sequences.
- Diffusion, dialysis and clearance — diffusion time, surface to volume, and fitting D from the agar exercise.

4.21. Video companion and reading

- **Core reading:** OpenStax *Anatomy and Physiology* [1] chapter 3, The Cellular Level of Organization — 3.1 to 3.6 (membrane, organelles, nucleus and DNA replication, protein synthesis, cell growth and division, differentiation). For more depth, OpenStax *Biology 2e* [71] units on the cell, genetics, and molecular biology.
- **Core videos:** the Crash Course Biology series, from the series index [6] — the episodes on cells, DNA, the central dogma, and cell division.
- **For an engineer's sense of scale:** Milo and Phillips, *Cell Biology by the Numbers* [73] — free online, and exactly the order-of-magnitude reasoning this chapter uses.
- **Reference:** Alberts et al., *Molecular Biology of the Cell* [72]. Borrow it; use the index.
- **Databases to explore:** NCBI [76], Ensembl, UniProt, and the Protein Data Bank [33].
- **Biosafety:** the WHO *Laboratory Biosafety Manual* [74] — read the introduction and the risk-assessment chapter.

4.22. Exit criteria

- ▢ Sketch a eukaryotic cell from memory, labelling at least eight structures with a function for each.
- ▢ Explain the path from gene to protein, including splicing and the genetic code, to someone who has not studied biology.

- ▢ Give two concrete examples of how evolution explains a feature of the body or a clinical problem.
- ▢ Calculate a diffusion time, a bacterial or PCR amplification, and a centrifuge RCF, with units.
- ▢ For one real-world use of biology (a diagnostic test, a therapy, or a manufacturing process), describe the workflow and name the instrument used at each step.
- ▢ Complete at least two workshop exercises and record what each demonstrated and its limitations.
- ▢ State the four biosafety levels and explain why you will not culture organisms outside a supervised laboratory.

4.23. What this chapter does not establish

- **You are not a laboratory scientist.** Knowing what a flow cytometer measures is not knowing how to run, gate, quality-control, or interpret one. Biomedical scientists and clinical scientists are trained and, in many countries, registered professionals.
- **You may not handle clinical samples or culture organisms** except within a laboratory's training, supervision, and biosafety framework.
- **A home-built instrument is not a validated instrument.** Research use, teaching use, and diagnostic use carry different evidence requirements; only the first two are realistic early on.
- **Cell biology knowledge is not clinical knowledge.** Understanding how a CAR-T product is made says nothing about who should receive it.
- **The field moves quickly.** Sequencing costs, approved gene therapies, and imaging capability change year by year. Treat specific numbers here as orientation and check current sources.

Chapter 5. Tissues, skin, and contact with an object

Week 7 · approximately 4–6 hours

5.1. Why this chapter exists

A good fit in CAD does not establish a suitable interface with a person.

That sentence is the chapter. Everything else is the mechanism behind it: what skin actually is, how it fails, why "it feels comfortable to me" is close to worthless as evidence, and what questions to ask before you let an object rest against somebody.

The failure this prevents is the most common serious harm an amateur-designed contact object can cause, and it is almost always caused by someone competent. A well-machined splint with a crisp edge. A moulded handle with its parting line running across the point of highest load. A neatly sewn strap whose seam allowance sits under a bony prominence. A strap sized to a measurement taken in the morning on a limb that swells by afternoon.

5.2. The four tissue families

Every tissue in the body is classified into one of four families. Learn the families first; the subtypes follow.

Family	Defining job	Where you meet it
Epithelial	Covering and lining; selective barrier; secretion and absorption	Epidermis, gut lining, airway lining, the inner lining of blood vessels (endothelium), gland ducts
Connective	Support, binding, transport, storage; characterised by abundant **extracellular matrix**	Bone, cartilage, tendon, ligament, fat, dermis, blood
Muscle	Generating force by contraction	Skeletal muscle, cardiac muscle (heart wall), smooth muscle (vessel walls, gut, airways, bladder)
Nervous	Rapid communication, integration, control	Brain, spinal cord, peripheral nerves, sensory receptors

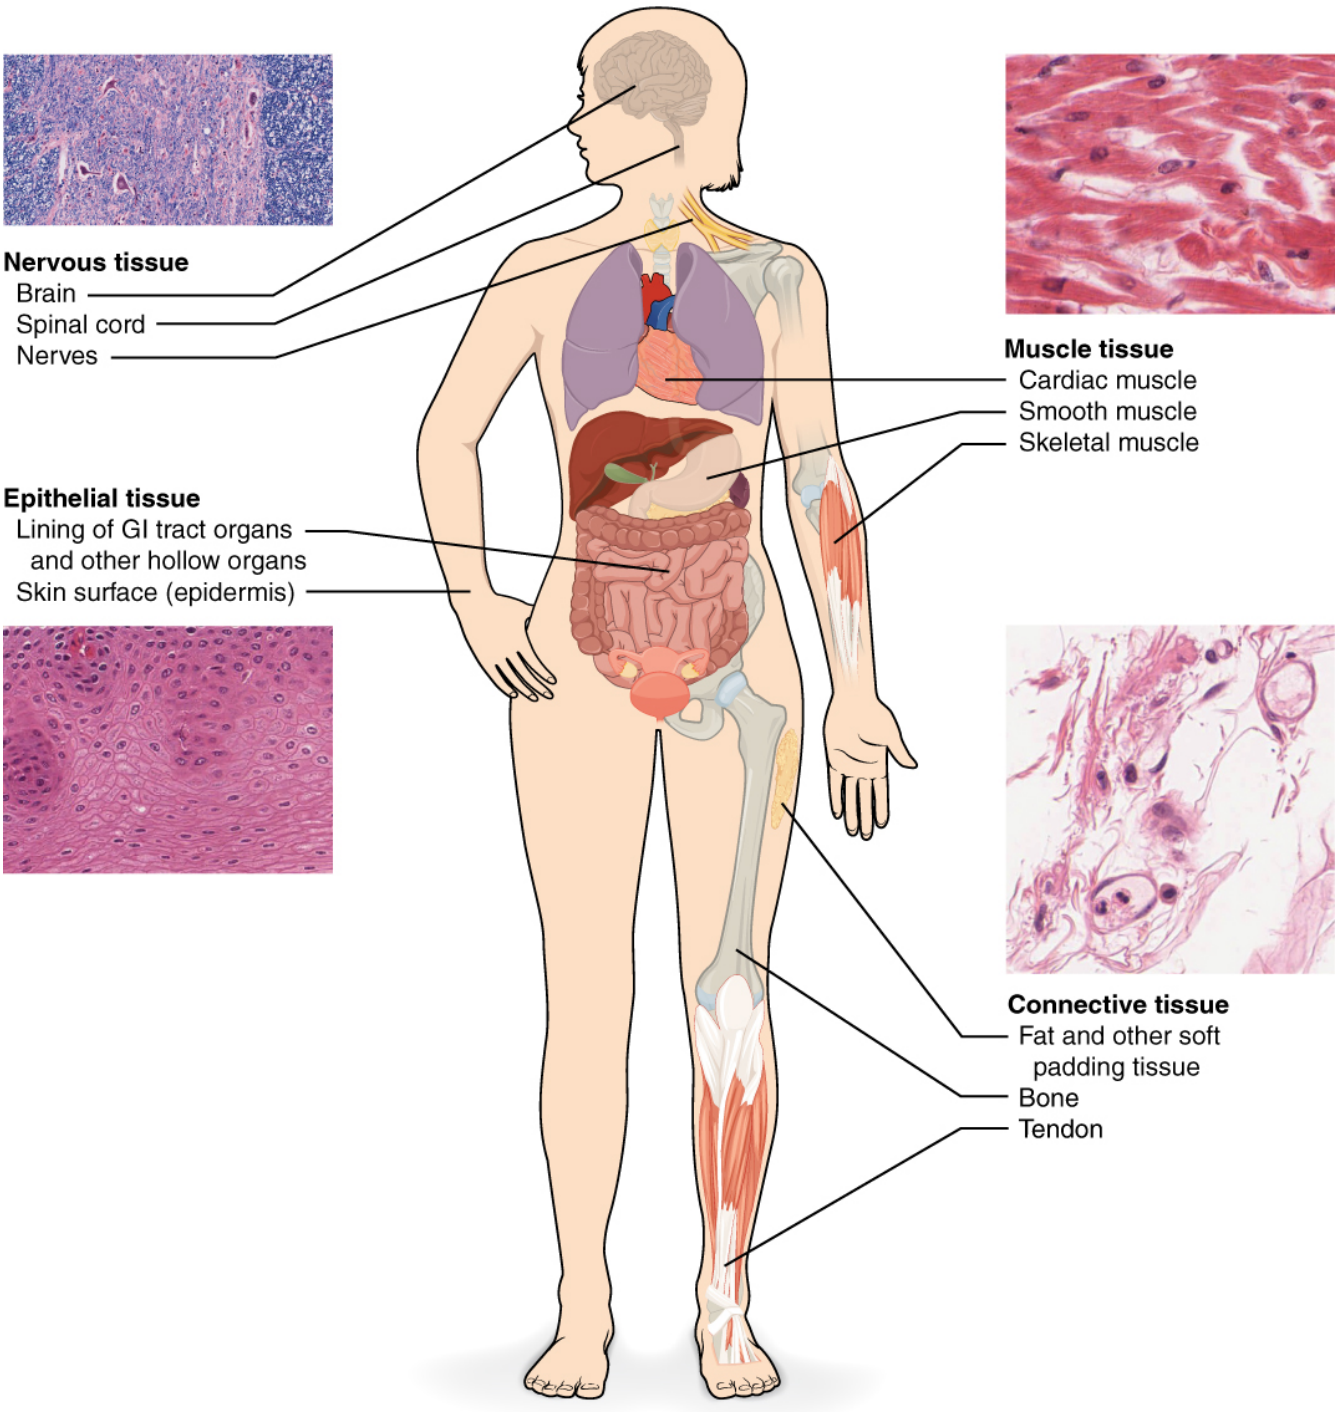


Figure 7. The four tissue families, with representative micrographs. Image: OpenStax College, CC BY 3.0. Credits: Appendix H

5.2.1. Epithelium

Classified by **cell shape** (squamous = flat, cuboidal, columnar) and **number of layers** (simple = one layer, stratified = many). The combination tells you the job.

Type	Location	Why that design
Simple squamous	Alveoli, capillary endothelium	Thin — minimises diffusion distance
Simple cuboidal	Kidney tubules, gland ducts	Secretion and absorption
Simple columnar	Stomach, intestine	Absorption and secretion, often with microvilli

Type	Location	Why that design
Pseudostratified columnar (ciliated)	Trachea, bronchi	Moves mucus upward
Stratified squamous (keratinised)	Epidermis	Abrasion resistance, barrier
Stratified squamous (non-keratinised)	Mouth, oesophagus, vagina	Abrasion resistance, kept moist
Transitional	Bladder	Stretches as volume changes

Two facts with direct design consequences:

1. **Epithelium is avascular.** It has no blood vessels of its own; it is supplied by diffusion from the connective tissue beneath, across a **basement membrane**. This is why the epidermis can be relatively tough and why damage to the layer beneath is what causes real trouble.
2. **A continuous epithelial surface is the body's primary barrier.** Breaching it — anywhere, by any means, including a 3 mm pressure-induced break you did not intend — changes the project from "object near a person" to "object at a wound", which is a categorically different review (Chapter 13).

5.2.2. Connective tissue: the engineer's family

This is where a mechanically-minded reader should spend their attention, because connective tissue is where the body's structural materials live. What distinguishes connective tissue is that its properties come mostly from the **extracellular matrix** the cells secrete, not from the cells themselves.

Matrix component	Mechanical role	Analogy (state its limits)
Collagen fibres (type I predominant in tendon, ligament, skin, bone)	High tensile strength, low extensibility	Like aramid fibre in a composite — but living, continuously turned over, and able to remodel along load direction
Elastin fibres	Recoverable large strain; returns tissue to shape	Like rubber — but interwoven with collagen so stiffness rises sharply at higher strain
Reticular fibres	Fine supporting mesh in soft organs	—

Matrix component	Mechanical role	Analogy (state its limits)
Ground substance (water, proteoglycans, glycosaminoglycans)	Binds water; resists compression; controls diffusion	Like a hydrogel filling a fibre network — and the fluid flow through it is why tissue is time-dependent

Now the specific tissues a maker confuses:

Tissue	Usual connection	Dominant loading	Do not model it as…
Tendon	Muscle → bone	Tension along its axis; transmits muscle force	…the motor. It transmits; it does not generate.
Ligament	Bone → bone across a joint	Tension; constrains and stabilises joint motion	…a tendon. Different job, different place.
Articular cartilage	Covers bone ends inside a synovial joint	Compression, with fluid pressurisation; very low friction surface	…a rubber pad. Its load-bearing depends on fluid trapped in the matrix.
Fibrocartilage	Intervertebral discs, menisci	Compression + tension; shock distribution	…solid. It is partly a fluid-pressure structure.
Adipose (fat)	Beneath skin, around organs	Cushioning, insulation, energy store; varies enormously by site and person	…a uniform layer of known thickness.
Bone	The skeleton	Compression, tension, bending, torsion; see Chapter 6	…inert. It remodels continuously in response to load.



A teaching model that renders tendon, ligament, and cartilage as the same generic "rubbery connector" has quietly taught three errors. If you simplify — and you should — **label the simplification**. "Cord represents a tendon transmitting muscle force; the demonstrator supplies the muscle force" is twelve words and it prevents the single most common misconception about joints.

5.2.3. Muscle and nervous tissue, briefly

Muscle comes in three varieties, and the differences matter when reading clinical language:

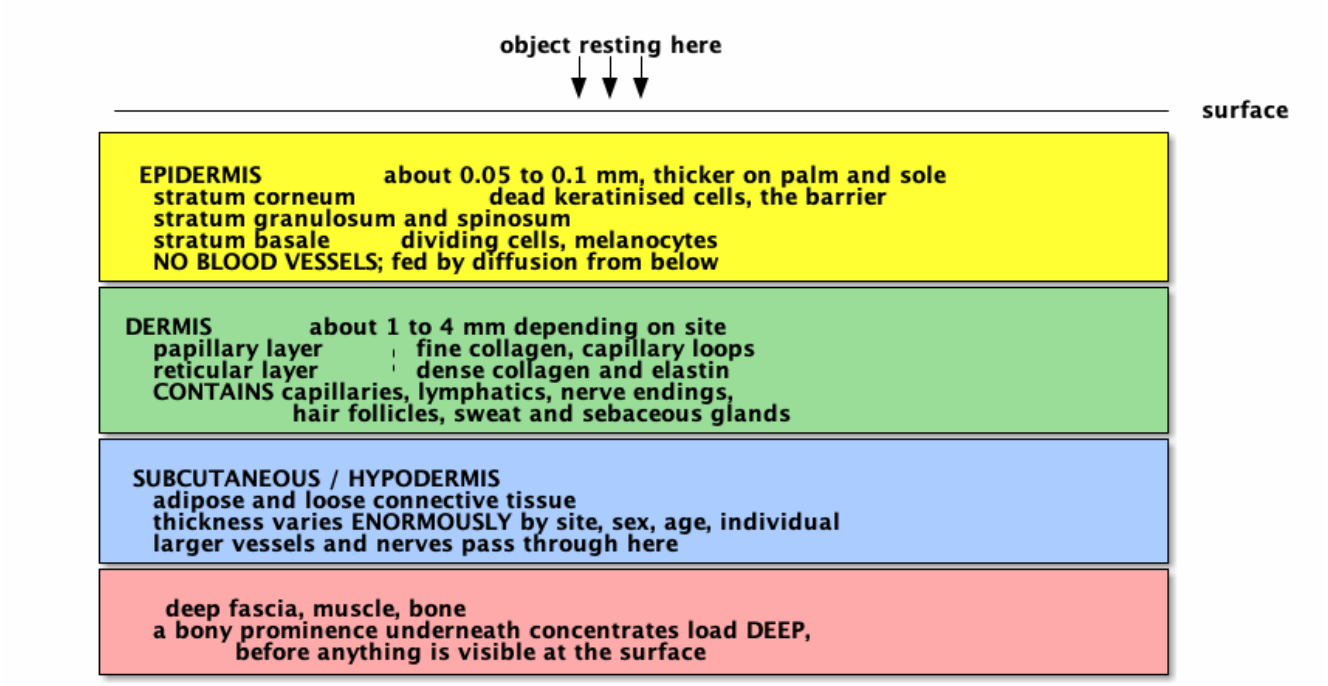
Type	Control	Appearance	Where
Skeletal	Voluntary (plus reflex)	Striated, multinucleate fibres	Attached to skeleton; Chapter 6

Type	Control	Appearance	Where
Cardiac	Involuntary, intrinsically rhythmic	Striated, branched, electrically coupled	Heart wall; Chapter 9
Smooth	Involuntary	Non-striated, spindle cells	Vessel walls, gut, airways, bladder, iris

Nervous tissue is covered in Chapter 7. For now: **neurons** signal, **glial cells** support and insulate them, and the ratio of glia to neurons is high.

5.3. Skin is an organ

Skin is the largest organ by surface area, roughly 1.5–2 m² in an adult, and it is a **layered composite with active biology**, not a cladding material.



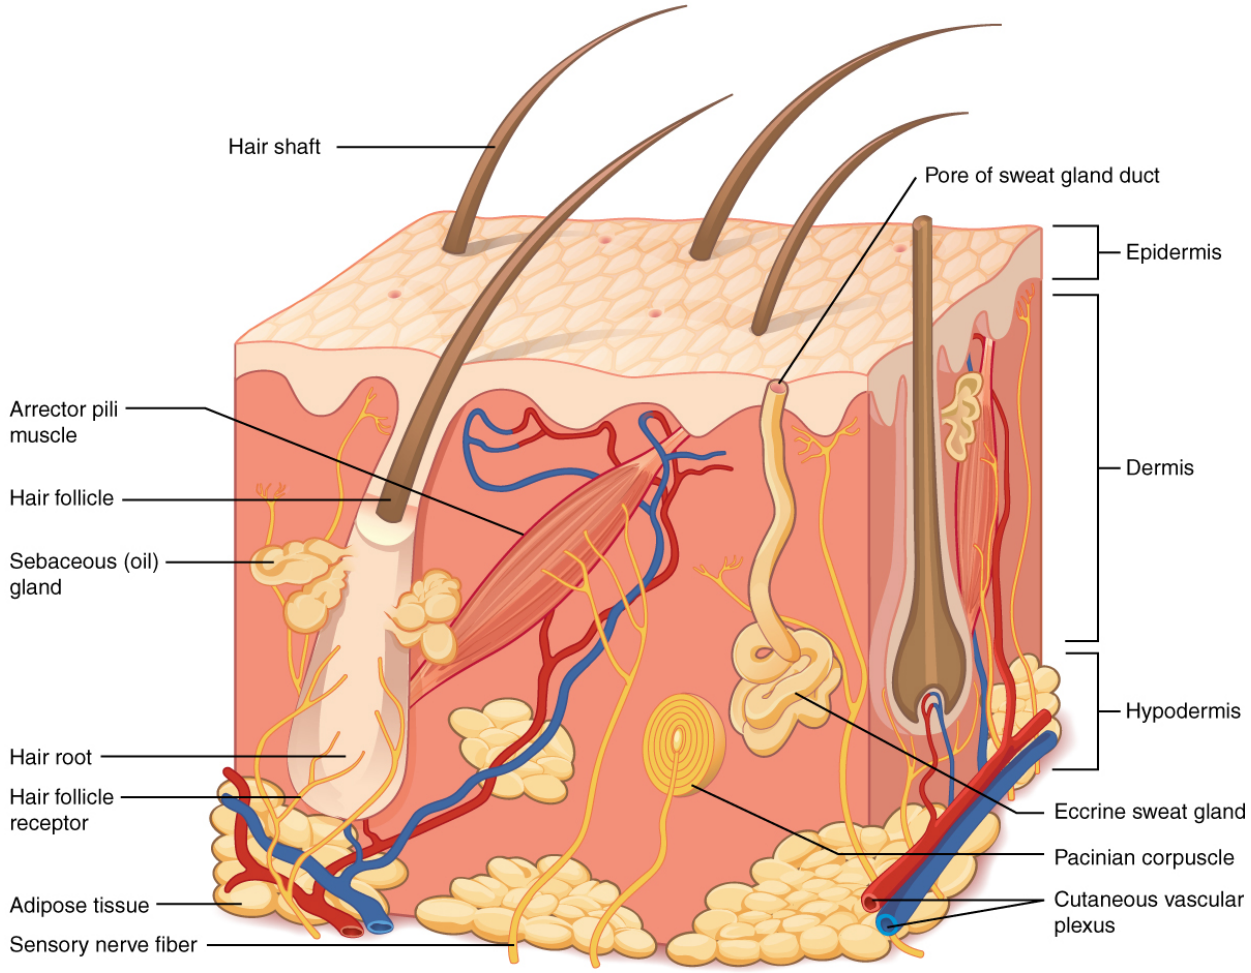


Figure 8. Structure of the skin, showing epidermis, dermis, and hypodermis with their appendages. Image: OpenStax College (Betts, Desaix, Johnson et al.), CC BY 3.0. Credits: Appendix H

5.3.1. Epidermis

Stratified keratinised squamous epithelium, **avascular**, renewed continuously from the **stratum basale** outward. Cells flatten, fill with keratin, die, and are shed from the **stratum corneum**. Full turnover takes roughly a month in healthy adult skin.

Resident cell types worth knowing by name:

- **Keratinocytes** — the bulk; produce keratin, form the barrier.
- **Melanocytes** — produce melanin, which absorbs ultraviolet radiation. Melanocyte **number** is broadly similar across human populations; the amount and type of melanin produced differs. This has a direct engineering consequence for optical sensors (Chapter 9).
- **Langerhans cells** — immune surveillance within the epidermis.
- **Merkel cells** — associated with light-touch sensation (Chapter 7).

The barrier function is largely the stratum corneum plus intercellular lipids. Strip it — with solvent, adhesive, repeated tape removal, or maceration from trapped moisture — and the barrier is compromised before there is any visible wound. **Medical adhesive-related skin injury** is a recognised category of harm, and it is directly caused by design decisions about tapes, dressings, and strap fixation.

5.3.2. Dermis

Dense connective tissue, and the mechanically dominant layer. A **papillary** layer of fine collagen with capillary loops that nourish the epidermis above, and a **reticular** layer of dense collagen bundles plus elastin that provides tensile strength and recoil.

The dermis carries everything functional:

- blood vessels — supplying the epidermis by diffusion, and performing thermoregulation by dilating and constricting,
- lymphatic vessels — fluid return and immune traffic,
- sensory nerve endings — touch, pressure, vibration, temperature, pain,
- hair follicles, sebaceous glands, sweat glands.

5.3.3. Hypodermis (subcutaneous tissue)

Adipose and loose connective tissue beneath the skin proper. Cushioning, insulation, energy storage. **Its thickness varies by site, sex, age, nutrition, and individual by a factor of ten or more.** This is why a design validated on one person's forearm tells you very little about another's, and why the phrase "the standard forearm" should make you suspicious.

5.3.4. What skin does

Function	Consequence for a design
Barrier (mechanical, chemical, microbial, UV)	Any breach changes the project category. Any chemical contact (adhesive, solvent, residual uncured resin, textile dyes and finishes, nickel in a buckle) is a chemical exposure.
Sensation	The person's discomfort report is a safety mechanism. If sensation is reduced, **that mechanism is absent** and the design must not rely on it.
Thermoregulation	Covering skin reduces heat loss and traps moisture. An insulating enclosure is a thermal and a microclimate decision.
Water balance	Occlusion causes maceration — softened, weakened, more easily damaged skin.
Vitamin D synthesis, immune surveillance, excretion	Context; rarely a direct design driver.

5.4. Pressure is only the beginning

Average pressure is force divided by area:

$$\mathbf{p}_{avg} = \mathbf{F} \div \mathbf{A}$$

Take a bench fixture loaded with 20 N over 10 cm²:

$\mathbf{p}_{avg}$ = 20 N ÷ (10 × 10^{-4} m^2) = **20 kPa** ≈ 150 mmHg

Double the contact area at constant force and the average halves to 10 kPa. That is a true and useful calculation. Here is everything it does not tell you.

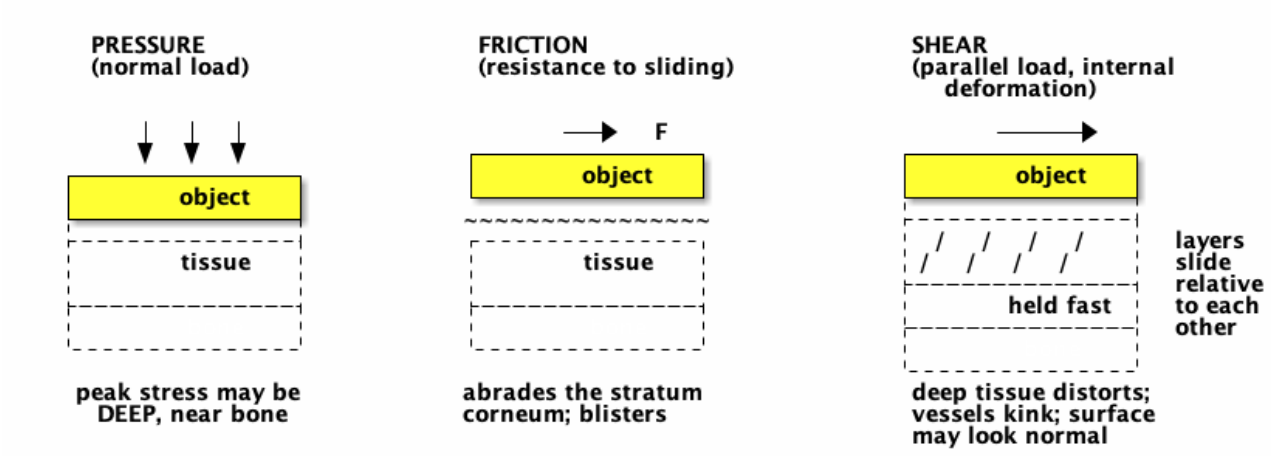
What the pressure calculation cannot establish:

- **Distribution.** Average pressure is an average. A single ridge, seam, layer line, or fastener head can carry a large fraction of the load, and the local peak can be many times the average. The average can look reassuring while a 2 mm edge does the damage.
- **Depth.** Tissue damage from loading often begins **deep**, near a bony prominence, before anything is visible at the surface. The surface pressure you could measure is not the stress in the tissue that fails.
- **Time.** Load and duration trade off non-linearly. A pressure tolerable for two minutes may not be tolerable for two hours, and the relationship is not a simple threshold.
- **The person.** Tissue condition, perfusion, age, nutrition, moisture, medication, sensation, and mobility all change what a given loading does.
- **Movement.** Sliding and repositioning introduce shear and friction, which are different mechanisms (below).

Use $p = F/A$ to generate a **question about force distribution**. Never use it to invent a clinical threshold. Published interface-pressure figures exist in the wound care literature; they are contested, context-specific, and not yours to apply.

5.4.1. Pressure, friction, shear, and microclimate

These are four distinct mechanisms and they are routinely conflated.



- **Pressure** is load normal to the surface. It compresses tissue and can occlude blood vessels, reducing perfusion.
- **Friction** is resistance to sliding at the interface. It abrades the stratum corneum and can cause superficial damage and blistering. Reducing friction is sometimes helpful and sometimes

harmful — low friction means the person slides, which produces shear.

- **Shear** is load acting parallel to the surface, transmitted into the tissue so that layers displace relative to one another. It distorts and can kink the vessels supplying a region. Shear is the mechanism people most often forget, and it is why "I reduced the pressure" can coexist with "the tissue got worse".
- **Microclimate** — local temperature and humidity at the interface. Trapped moisture macerates skin and raises friction; raised temperature increases metabolic demand in tissue whose supply may already be compromised.

A stack of foam sheets with the top sheet pushed sideways while the bottom is held still demonstrates shear usefully. It is a mechanical demonstration. It is not a skin substitute: no circulation, no sensation, no living response, no layered anisotropy.

5.4.2. Soft tissue is viscoelastic, non-linear, and anisotropic

This is the part that most surprises someone used to metals.

Property	What it means	Why you care
Non-linear (J-shaped) stress–strain	Stiffness increases with strain. At low strain, crimped collagen fibres straighten easily; once recruited and aligned, stiffness rises steeply.	There is no single "Young's modulus of skin". A stiffness measured at low strain badly underestimates behaviour at higher strain.
Viscoelastic	Response depends on rate and history. Exhibits **creep** (continued deformation under constant load), **stress relaxation** (falling stress under constant deformation), and **hysteresis** (energy lost per cycle).	A strap tightened to a comfortable tension at 09:00 is at a different tension at 11:00 — in both directions, because the tissue creeps **and** the limb volume changes.
Anisotropic	Properties differ with direction, following collagen fibre orientation (clinically, **Langer's lines**).	Identical loading applied in two directions produces different deformation. A symmetric part is not symmetric in its effect.
Pre-stressed	Skin is under tension in vivo; excised skin retracts.	In-vitro material data does not transfer directly to in-vivo behaviour.
Adaptive	Repeated loading changes the tissue — thickening, remodelling, or injury	The tissue in week six is not the tissue you measured in week one.



Order-of-magnitude stiffness, for intuition only. These are wide literature ranges that depend heavily on site, direction, strain level, strain rate, age, and test method. Treat them as "which decade am I in", never as design values.

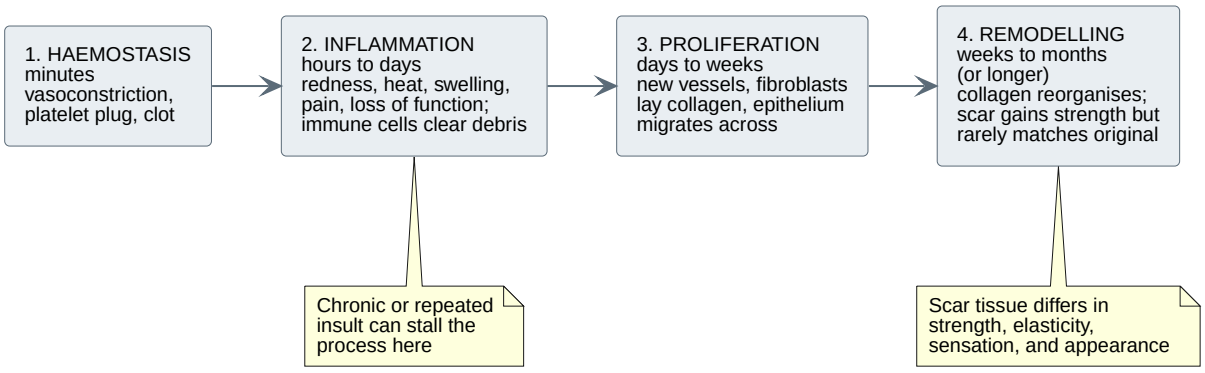
Material	Approximate elastic modulus
Cortical bone	~10–20 GPa

Material	Approximate elastic modulus
Tendon (along fibre, high strain)	~1–2 GPa
Articular cartilage	~0.5–15 MPa (highly rate-dependent)
Skin	~0.1–100 MPa depending on strain level and method
Adipose tissue	~1–10 kPa
Rigid engineering polymers — acetal, ABS, PLA, nylon	~1–4 GPa
Silicone elastomer	~0.1–10 MPa
Polyurethane and closed-cell foams	~0.01–1 MPa (strongly density- and strain-dependent)
Stainless steel 316L	~190–200 GPa
Titanium Ti-6Al-4V	~110 GPa

The useful observation: a rigid plastic part — moulded, machined, or printed — is roughly as stiff as **tendon** and far stiffer than **fat**. Stainless steel is ten times stiffer than cortical bone. Every rigid object you place against a person is, mechanically, much harder than what it touches, and the compliance mismatch is where load concentrates.

5.5. Inflammation and tissue repair

You need the shape of this, because it explains why "it was fine yesterday" is not reassurance and why healing timescales are long.



The classical signs of acute inflammation — redness, heat, swelling, pain, loss of function — are a **normal and necessary** response, not in themselves a sign of mismanagement. But they tell you that the tissue has changed, and therefore that any geometry fitted before the change may no longer fit. Swelling in particular is a direct dimensional problem: a limb's circumference can change substantially over hours.

Repair produces **scar**, not replacement. Scar has different tensile properties, different elasticity, often altered sensation, and does not restore the original architecture. Which is why preventing a 5 mm skin break matters more than being good at responding to one.

5.6. Engineering bridge: the interface as a design problem

5.6.1. Intact skin, compromised skin, and breached skin

These are three different projects.

Contact category	Example	What it implies
No body contact	Tabletop teaching model, bench fixture, software	Workshop review; infection control applies to **handling** (Chapter 13) but not to tissue contact
Intact skin, brief	Handle, control, grip aid used for seconds to minutes	Human factors review; edges, forces, cleaning
Intact skin, prolonged	Splint, strap, seating, prosthetic interface, worn device	Requires clinical assessment. Pressure, shear, microclimate, swelling, sensation, re-check schedule. **Not a maker decision alone.**
Compromised skin	Fragile, oedematous, very thin, scarred, or already reddened skin	Specialist clinical involvement, always
Breached skin, mucosa, or internal	Wound contact, invasive, airway, implant	A different regulatory and clinical universe. Do not approach this as a learning project.



Where your project sits in that table determines who must review it. Work it out **before** you build, and write it in the brief (Chapter 18). A project that moves up a row during development has not become a bigger version of itself; it has become a different project.

5.6.2. A contact design checklist

Ask all of these before you commit geometry. The answers come from the clinical team and the user, not from this book.

1. **Where** exactly does it contact, and over what area? Draw it.
2. **For how long** in one sitting, and how many times a day?
3. **What forces** act, in which directions, and do they change with movement?
4. **Is the skin intact**, and what is its condition? Who assessed that?
5. **Can the user feel** discomfort there, and can they act on it — reposition, remove, call someone?
6. **What is beneath** the contact point — a bony prominence, a superficial nerve, a superficial

vessel?

7. **Does the limb change size** — swelling, time of day, posture, activity, over weeks?
8. **What is the microclimate** — is moisture trapped? Does it need to breathe?
9. **How is it cleaned**, by whom, how often, and does cleaning reach every surface (Chapter 13)?
10. **What are the edges, seams, stitching, fasteners, parting lines, and layer lines** doing at the contact boundary? Where is the highest local stiffness?
11. **What happens when it fails** — does it break into sharp pieces, release suddenly, trap a limb?
12. **Who re-checks it** after a week, a month, a change in the person's condition?
13. **What is the removal and escape path** if something goes wrong?
14. **Is there an existing product** that already does this, and why is it unsuitable?

5.6.3. Design moves, and what each one costs

Padding is not a universal solution, and every modification is a new design with new consequences.

Move	Intended effect	What it also does
Increase contact area	Reduce average pressure	Changes grip diameter, may add weight and bulk, may introduce new edges, may trap more moisture
Add soft padding	Distribute load, improve comfort	Changes diameter and grip; adds seams; absorbs fluid; may be hard to clean or need replacing; may bottom out and conceal a hard edge beneath
Round or chamfer edges	Remove stress concentration	Almost always good; costs a little material and a post-processing step
Reduce stiffness (softer material)	Lower compliance mismatch	May reduce the mechanical function that justified the part
Reduce contact time	Lower the load-duration product	Requires a behavioural or workflow change — needs the user and clinician on board
Add ventilation	Improve microclimate	Creates holes with edges, more surfaces to clean, possible ingress
Make it adjustable	Accommodate swelling and variation	Adds mechanisms that can be set wrongly; needs instructions; new failure modes
Smooth the surface finish	Reduce friction and abrasion; improve cleanability	Extra process step; may reduce grip where grip is needed

Move	Intended effect	What it also does
Change the manufacturing method	Remove a feature the process forced on you — a layer line, a parting line, a seam, a weld	Moving a handle from print to a machined or moulded part, or a strap from a stitched seam to a bonded or knitted one, changes cost, lead time, cleanability, and what you can iterate cheaply (Chapter 14)

The honest statement about any of these is "this changes X, and I have not established Y". Write it that way in your notes.

5.7. Devices, software, and training technology

Skin and soft tissue support a large device industry, and it is the one most accessible to a maker — because much of it is mechanical, much of it is consumable, and the clinical teams who use it are used to improvising.

5.7.1. Clinical devices

Device	Does	Engineering note
Wound dressings — films, foams, hydrocolloids, hydrogels, alginates, antimicrobial	Manage the wound environment: moisture, exudate, protection, sometimes antimicrobial action	A moisture-management problem. Absorbency, vapour transmission rate, adhesion, and atraumatic removal all trade against each other
Negative pressure wound therapy (NPWT)	Applies controlled sub-atmospheric pressure to a sealed wound via a foam dressing	Pump, canister, seal integrity, and alarm handling. A genuine engineering device
Pressure-redistribution surfaces	Mattresses, overlays, and cushions that spread load and manage microclimate	**Reactive** (constant low pressure) and **active** (alternating). The real metric is tissue stress over time, which cannot be measured directly
Interface pressure mapping systems	Sensor mats showing pressure distribution under a seated or lying person	Measures **surface** pressure only. Says nothing about deep tissue stress — the thing that actually fails (Chapter 5)
Compression therapy — bandages, stockings, intermittent pneumatic compression	Manages oedema and venous return	Graduated pressure profiles; application technique dominates the delivered pressure

Device	Does	Engineering note
Orthoses and splints	Support, immobilise, or correct position	Interface pressure, oedema accommodation, and removal path are the design problems. **Clinical involvement from the outset**
Dermatoscope	Magnified, polarised, illuminated view of skin lesions	Cross-polarisation to suppress surface reflection. Now commonly phone-mounted
Skin substitutes and tissue-engineered grafts	Temporary or permanent wound cover	Advanced therapy; a regulated category well beyond this book
Phototherapy and laser systems	Treat psoriasis, neonatal jaundice, vascular lesions, hair removal	Wavelength selectivity against a chromophore. Significant eye and burn hazards
Patient-handling slide sheets and transfer boards	Reduce shear during repositioning	Directly addresses the shear mechanism in Chapter 5. Low-friction woven fabrics on one side, grip on the other — a textile-engineering product
Seating and positioning systems	Cushions, contoured backs, lateral supports, covers	Foam conversion, upholstery, and covers with vapour-permeable, fluid-resistant fabrics. Prescribed and fitted clinically, but much of the craft is textiles and foam
Pressure and microclimate sensing	Instrumented insoles, cushions, and garments in research and some clinical use	Thin-film force sensors, temperature and humidity sensors, and low-power electronics. Sensor drift, hysteresis, and calibration are the hard parts

5.7.2. Software

Software	Does	Note
Wound imaging and measurement apps	Photograph, measure area and depth, track healing over time	Calibration, lighting, and colour constancy are the hard problems. Several use a reference marker in frame
Pressure injury risk assessment tools — Braden, Waterlow, and others	Structured risk scoring to trigger preventive action	A score is a prompt, not a prediction. The inference-layer caution from Chapter 9 applies
Teledermatology platforms	Remote lesion assessment from images	Image quality standards matter enormously; skin tone representation is a known equity issue

Software	Does	Note
AI lesion triage	Flag lesions for review	**Training-set skin tone diversity is a documented weakness.** Same lesson as the pulse oximeter in Chapter 9
Orthotic and prosthetic CAD — Rodin4D, Canfit, and open tools	Design sockets and orthoses from scans	Scan-to-fabrication workflows: CNC-carved positive moulds for thermoformed sockets, and increasingly direct additive manufacture
3D scanning software — Meshroom, photogrammetry pipelines, depth-camera SDKs	Capture body geometry without ionising radiation	**Accessible to a maker today.** Accuracy, repeatability, and soft-tissue deformation during capture are the limitations to state

5.7.3. Training and simulation

Training technology	Teaches	Form
Suture and skin pads	Suturing, knot tying, excision, closure technique	Layered silicone or foam. **Among the easiest things a maker can make well**, and consumed in quantity
Wound care and staging trainers	Recognising wound types and pressure injury categories, dressing selection	Moulded silicone wounds, often wearable over a limb
Moulage	Realistic wound appearance for scenario training	Silicone, wax, pigment. A craft skill with real demand
Injection trainers — intradermal, subcutaneous, intramuscular	Needle angle, depth, site selection	Layered tissue-mimicking pads, sometimes with feedback
Pressure injury staging image sets	Visual recognition and classification	Software and card sets. **Skin tone diversity in these resources is a known gap**
Interface pressure demonstration rigs	Shows how load redistributes with surface and posture	Sensor mat plus display. Teaches the concept without any patient involvement
Shear and friction demonstrators	Why repositioning technique matters	Simple layered-material rigs — the exercise in this chapter, formalised

Training technology	Teaches	Form
Instrumented interface rigs	How load, area, and surface change pressure and shear	Load cells or force-sensitive resistors under a cast silicone "tissue" block, read by a microcontroller and shown live. Electronics, casting, and a machined base
Sewn and upholstered demonstration aids	Strap width, seam placement, and padding effects	A set of straps and covers that differ in one variable each. Textiles make the comparison tangible



Suture pads and moulage are the two highest-demand, lowest-risk items here.

Suture pads are consumed by the hundred in surgical skills teaching, commercial ones are expensive, and a layered silicone pad with a realistic skin–fat–fascia feel is squarely within a maker's competence. The skill is in **mould making** — machined or printed moulds, controlled pour sequence, degassing, and fabric or mesh layers for tear resistance — and in making the hundredth pad match the first. The educator specifies the **feel** they need — the resistance contrast between layers matters more than absolute hardness (Chapter 21) — and you iterate against their judgement.

Neither touches a patient. Both are genuinely wanted. Both build the relationship.

5.8. Worked example: the awkward handle

An occupational therapist mentions that a user finds a handle uncomfortable. Your hands are already reaching for foam.

Stop and ask what "uncomfortable" is. Candidate mechanisms, each with a different fix:

Mechanism	Evidence that would distinguish it	Fix direction
Concentrated force at an edge or seam	Mark shows at a specific line; discomfort localises to a point	Radius the edge, redistribute, move the seam off the load path
Insufficient contact area	Diffuse pressure discomfort across the whole contact	Larger area, contoured shape
Awkward wrist or forearm position	Discomfort is in the wrist or forearm, not at the contact	Change the handle **angle**, not its surface — see Chapter 6
Excessive grip force required	Fatigue builds over seconds to minutes; discomfort in the muscles	Reduce required force, improve friction, change the mechanism
Slipping	User grips harder than necessary to prevent slip	Improve surface friction, add a retaining feature

Mechanism	Evidence that would distinguish it	Fix direction
Thermal or moisture	Sweating, clamminess, discomfort worse in warm conditions	Material and ventilation
Vibration	Only with a powered tool	Isolation — a specialist topic

Note that for three of those seven, adding padding makes things **worse**: it increases diameter (raising required grip force), can increase required grip force to prevent slipping, and can trap moisture.

Then ask the questions padding raises in its own right: does it change the diameter past what the hand can span? Does it create a seam? Does it absorb fluid? Can it be cleaned by the host's actual process? Who replaces it, and how often?

The design move is to find out which mechanism is operating, not to apply the default fix. That requires the therapist and the user, and it is the whole of Chapter 18.

5.9. Workshop exercise

1. **Layered skin model.** Build a skin cross-section from coloured card and foam: epidermis, dermis, subcutaneous tissue, and the muscle/bone beneath. Mark the locations of blood vessels, a nerve ending, and a sweat gland in the correct layer. Label the whole thing **"schematic, not to scale"** — because it is not, and the habit matters. Add a bony prominence underneath one region and mark where you would expect deep stress to concentrate.
2. **Area comparison on paper.** Sketch a narrow and a broad contact surface carrying the same load. Calculate both average pressures. Then mark on each sketch where an edge, a stitched seam, a parting line, or a layer line would concentrate force, and write one sentence on what the average fails to capture.
3. **Shear demonstration.** Stack four foam sheets. Hold the bottom, push the top sideways. Photograph the deformation. Write down three ways this is unlike skin. Repeat with two fabrics laid on top — a smooth woven and a brushed or fleece one — and note how the surface textile changes what reaches the layers below.
4. **Contact map for your own project.** For the elbow teaching model: it contacts hands, a table, and a storage box. Draw every contact. You will reuse this map in Chapter 13.



Do **not** run pressure, constriction, or prolonged-contact experiments on yourself or anyone else. This is not excessive caution: tissue damage from sustained loading develops without warning, can begin deep before it is visible, and cannot be undone by stopping. Pressure and contact limits come from the clinical team and the literature, not from a bench trial you devised.

5.10. Questions

1. Why is a tendon different from a ligament?
2. The epidermis has no blood supply. How is it nourished, and what does that imply about damage to the dermis?

3. You halve the average interface pressure. Have you established that the design is safe on skin?
4. Name the three distinct mechanical mechanisms at a skin interface and the different tissue damage each tends to produce.
5. What does **viscoelastic** mean, and give a concrete consequence for a strap?
6. What is wrong with saying a foam model reproduces living skin?
7. Why can tissue damage be deep before it is visible?
8. A limb is measured at 09:00 and a rigid part is made to fit. Name two reasons it may not fit at 16:00.
9. A user reports no discomfort after an hour of wear. How strong is that evidence, and what changes your answer?
10. List four questions you must ask before designing an object that rests against a person.

5.11. Answers and reasoning

1. Usual connection and function differ. A tendon typically connects muscle to bone and transmits muscle-generated force; a ligament typically connects bone to bone and contributes to joint stability and constraint of motion.
2. By diffusion from capillaries in the papillary dermis, across the basement membrane. It implies that damage to the dermis is more consequential than it looks: it compromises the supply to the epidermis above, as well as destroying vessels, nerve endings, and the tissue's main structural collagen.
3. **No.** You have changed one mechanical average under one assumed load and area. Distribution, peak local stress, depth of stress, duration, the individual's tissue condition, movement, microclimate, and clinical review all remain unresolved.
4. **Pressure** (normal load) — compresses tissue, can occlude perfusion, peak stress often deep near bone. **Friction** — abrades the stratum corneum, superficial damage and blisters. **Shear** — parallel load deforming tissue internally, distorts and kinks vessels, surface may look normal.
5. Its mechanical response depends on rate and loading history, showing creep, stress relaxation, and hysteresis. Consequence: a strap set to a given tension will not be at that tension an hour later, because the tissue creeps and relaxes — and independently, limb volume may change. So a single fitting is not a settled state, and a design that relies on a fixed tension needs a re-check or an accommodating mechanism.
6. The foam lacks living tissue's layered anisotropic structure, circulation, sensation, immune response, repair capacity, and pre-stress. It demonstrates exactly one mechanical idea, and only if you say which one.
7. Loading near a bony prominence produces its highest stress in deep tissue, not at the surface. Deep tissue can be injured while the overlying epidermis — which is tougher and more abrasion-tolerant — still looks intact.
8. Any two of: swelling or oedema; normal diurnal fluid-volume variation; posture or activity changing soft tissue distribution; the tissue creeping under the part's own load; temperature-driven changes in the part's dimensions. The design response is adjustability plus a re-check, not a tighter tolerance.

9. **Weak.** It is a single observation from one person on one occasion. It is **much** weaker if sensation is reduced in that region, because the person's discomfort report — the mechanism you are implicitly relying on — is impaired. It is also weak if the person is unable to reposition or remove the object themselves, or is reluctant to complain. Absence of a report is not evidence of absence of harm.
10. Any four from the checklist above. Strong answers include contact duration, cleaning, whether skin is intact, whether the user can detect and act on discomfort, what lies beneath the contact point, and **who assesses fit and re-checks it**.

5.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Levers, moment arms and contact — contact pressure, and what it cannot tell you.
- Tissues, materials and temperature — tissue and material stiffness on one scale.

5.13. Video companion and reading

- **Core videos:** Crash Course Anatomy & Physiology #2 (tissues) [8] and #6 (the integumentary system) [9]. For more depth, episodes 3–5 and 7 in the series index [6].
- **Core reading:** OpenStax [1] chapter 4 (Tissue Level of Organization, especially 4.1 Types of Tissues and 4.3 Connective Tissue) and chapter 5 (The Integumentary System, especially 5.1 Layers of the Skin and 5.2 Accessory Structures). Use the diagrams and the chapter review questions.
- **Context, not instruction:** if your project involves prolonged skin contact, the international pressure-injury prevention guidance [30] shows you the shape of the professional literature. Read it to understand how complex the question is and who owns it — not to extract a number.

5.14. Exit criteria

- ▢ Identify the four tissue families with an example of each, and correctly distinguish tendon, ligament, and articular cartilage.
- ▢ Label a skin diagram including epidermis, dermis, subcutaneous tissue, and the location of vessels and nerve endings.
- ▢ Explain the difference between pressure, friction, and shear using your own sketch, and name the characteristic damage of each.
- ▢ Complete the bench pressure calculation and state, unprompted, at least four things it cannot establish.
- ▢ Explain viscoelasticity with one concrete consequence for a worn object.
- ▢ List five questions to ask before designing a skin-contact object, including contact duration and cleaning.
- ▢ Explain why "it feels comfortable to me" is insufficient evidence for another user, in terms of

sensation, tissue variation, and duration.

- ▢ Place your own project in the contact-category table and name the reviewer it implies.

5.15. What this chapter does not establish

- **You cannot set a pressure limit.** You now understand why the question is hard. The answer belongs to the clinical team and the specialist literature.
- **You cannot assess anyone's skin.** Recognising that skin condition matters is not the same as assessing it. That is a clinical skill.
- **A bench measurement of interface pressure is not a safety assessment.** Even a real pressure-mapping measurement characterises one loading of one person at one moment, and says nothing about depth, duration, or tissue tolerance.
- **You have not learned biocompatibility.** Material–tissue interaction has its own standards family (Appendix D) and its own testing. "It is skin-safe" is a claim requiring evidence about a **finished** object in a **specific** use.
- **No part of this chapter authorises prolonged-contact or wound-contact work.** Those move you up the contact-category table into territory that requires specialist clinical involvement from the outset.

Chapter 6. Bones, joints, and movement

Weeks 8–9 · approximately 8–12 hours

6.1. Why this chapter exists

This is the chapter where your existing knowledge is most useful and most dangerous.

Most useful, because joints really are pivots, muscles really do apply forces, tendons really are tension members, and the statics you can already do is the right statics. A free-body diagram of a forearm is a genuine free-body diagram.

Most dangerous, because the simplification "joint = hinge, muscle = linear actuator, tendon = cable" is close enough to be convincing and wrong in ways that matter. Real joints have moving centres of rotation. Real muscles have force that depends on their length and velocity. Real tendons are routed over surfaces whose geometry changes the moment arm continuously through the motion. And a person's **comfortable repeated** range is a different thing from their **maximum achievable** range, which is a different thing again from the range a diagnosis might lead you to assume.

6.2. Bone as a material and as a tissue

6.2.1. Structure

Bone is a composite: a **collagen** matrix (tough, tensile) mineralised with **hydroxyapatite**, a calcium phosphate (stiff, compression-bearing, brittle alone). Neither constituent alone would do; the composite is both stiff and reasonably tough.

Two architectural forms:

Form	Structure	Where and why
Cortical (compact)	Dense, organised into **osteons** — concentric lamellae around a central canal carrying a vessel and nerve	Shafts of long bones; outer shell everywhere. Carries bending and torsion
Trabecular (cancellous, spongy)	Open lattice of struts aligned with habitual load paths; marrow in the spaces	Ends of long bones, vertebral bodies, pelvis. High stiffness-to-weight; distributes load into the joint

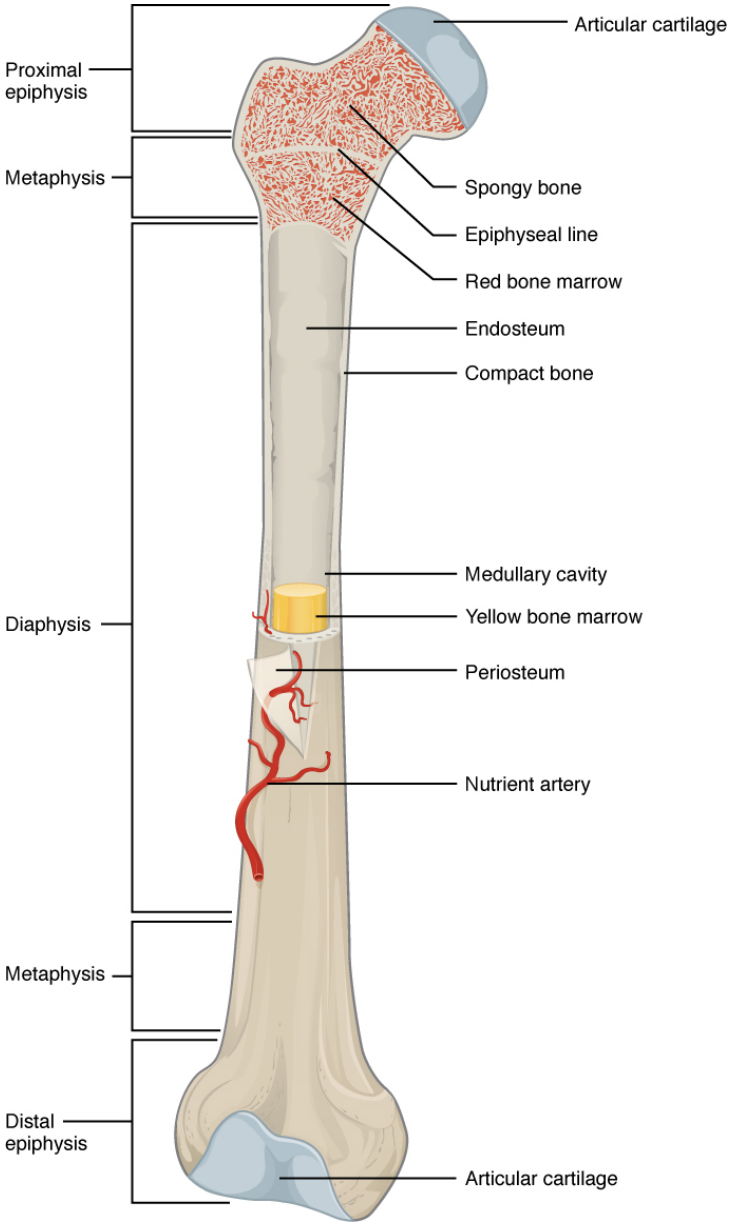


Figure 9. Anatomy of a long bone, showing the dense cortical shell and the trabecular interior. Image: OpenStax Anatomy and Physiology, CC BY 4.0. Credits: Appendix H

If you have ever designed a lattice infill, trabecular bone is the thing you were imitating. The alignment of trabeculae with principal stress directions is a real and observable phenomenon, and the adaptive remodelling of bone to its loading history is loosely called **Wolff's law**. Living cells — **osteoblasts** building, **osteoclasts** resorbing, **osteocytes** sensing within the matrix — continuously rebuild the structure.



The practical consequence of remodelling: **bone responds to how it is loaded.** Unload a bone (bed rest, microgravity, a limb in a cast, or an implant that carries load the bone used to carry) and it loses mass. This is why **stress shielding** is a central concern in orthopaedic implant design, and why "stiffer is better" is false in that context. A stiffer implant takes more load, so the bone takes less, so the bone weakens.

This is a good example of a design principle that is correct in a bridge and wrong in a body.

6.2.2. Mechanical properties, for intuition

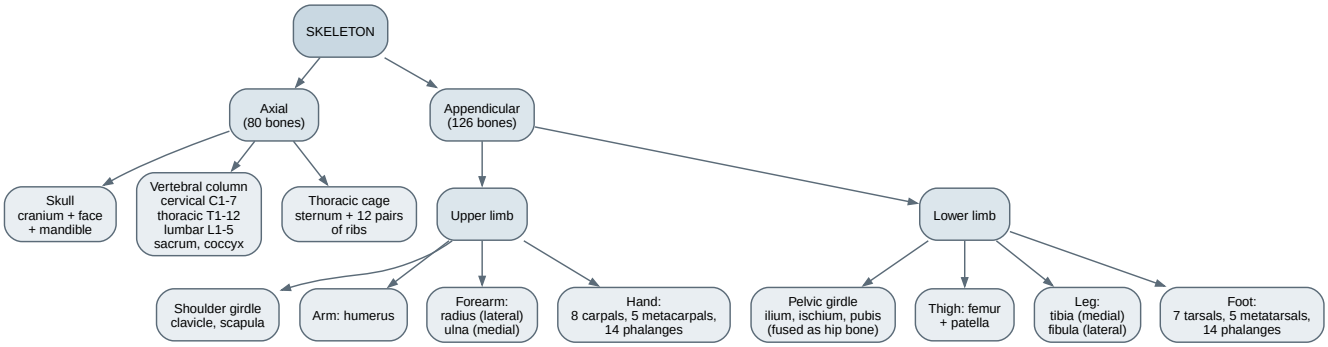
Material	Elastic modulus (approx)	Ultimate strength (approx)	Note
Cortical bone, longitudinal	~15–20 GPa	~100–200 MPa compression; ~50–150 MPa tension	**Anisotropic** — considerably weaker across the grain than along it
Trabecular bone	~0.01–3 GPa	~1–20 MPa	Varies by site and bone density over more than two orders of magnitude
Tendon (along fibres)	~1–2 GPa	~50–100 MPa	Non-linear; low-strain "toe region" is much more compliant
Articular cartilage	~0.5–15 MPa (strongly rate-dependent)	—	Load-bearing depends on pressurised interstitial fluid
Acetal (POM), machined or moulded	~2.5–3 GPa	~60–70 MPa	Low friction, dimensionally stable; a workhorse for bearings, pivots, and handles
PLA, printed	~2–4 GPa	~40–60 MPa (highly anisotropic by layer direction)	Weakest between layers — in Z
Glass-filled nylon, moulded	~5–10 GPa	~100–180 MPa	Stiffer and stronger along the flow direction than across it; weak at weld lines
Aluminium 6061-T6	~69 GPa	~310 MPa	~4× stiffer than cortical bone
Stainless steel 316L	~190–200 GPa	~500–600 MPa	~10× stiffer than cortical bone
Ti-6Al-4V	~110 GPa	~900–1000 MPa	Closest common metal to bone stiffness — part of why it is used in implants



These are literature ranges, not design allowables. Bone properties vary with site, loading direction, strain rate, age, sex, bone mineral density, and donor. They are here so you know whether you are in the right decade, and so you notice the two genuinely useful facts: **bone is anisotropic**, and **so are most things you make**. A printed part is weakest between layers; a fibre-filled moulding is stiffer along the flow than across it and weak where two flow fronts meet; wood follows its grain; a weld or a braze differs from the parent metal; a woven strap stretches more on the bias than along the warp. In every case the orientation is a structural decision, not a default of the process (Chapter 14).

6.3. The skeleton you actually need

Do not memorise 206 bones. Learn the ones in the region your project touches, plus enough of the rest to place a conversation.



6.3.1. The upper limb in detail

Because most assistive and training projects a maker meets involve hands and arms.

Bone	What to know
Clavicle	Collarbone. The only bony strut connecting the arm to the trunk anteriorly. Superficial and commonly fractured.
Scapula	Shoulder blade. Glides on the back of the thorax rather than forming a conventional joint with it — **scapulothoracic** motion contributes a large fraction of what people call "shoulder movement".
Humerus	Upper arm. Head articulates with the scapula’s glenoid; distally it forms the elbow. Its **medial epicondyle** is where the ulnar nerve runs superficially — the "funny bone".
Radius	**Lateral** forearm bone in anatomical position; thumb side. Its head rotates to produce forearm pronation and supination. Its distal end forms most of the wrist joint.
Ulna	**Medial** forearm bone. Its proximal **olecranon** is the point of the elbow. It stays relatively still during pronation/supination while the radius rotates about it.
Carpals	Eight small bones in two rows, with complex intercarpal motion. Wrist movement is distributed across many small articulations, not one hinge.
Metacarpals and phalanges	Palm and fingers. The thumb’s **carpometacarpal** joint is a saddle joint, which is what gives the human thumb its opposition range.

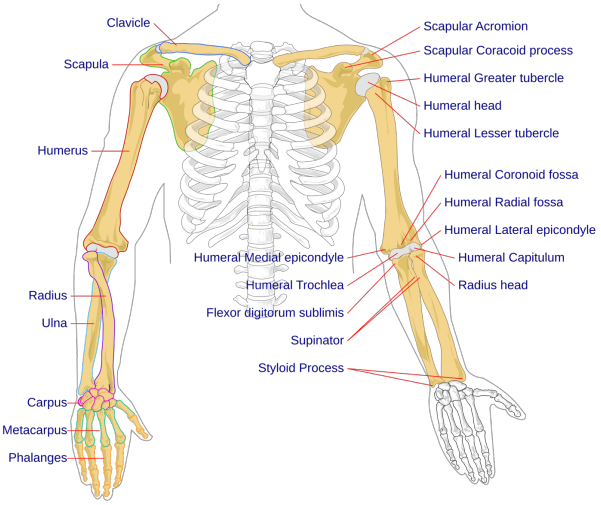


Figure 10. Bones of the upper limb: humerus, radius, and ulna. Image: LadyofHats (Mariana Ruiz Villarreal), public domain. Credits: Appendix H

The elbow is not one joint. The elbow region contains three articulations within one capsule:

- **humeroulnar** — the main hinge for flexion and extension,
- **humeroradial** — also participates in flexion/extension,
- **proximal radioulnar** — rotation of the radius for pronation and supination.

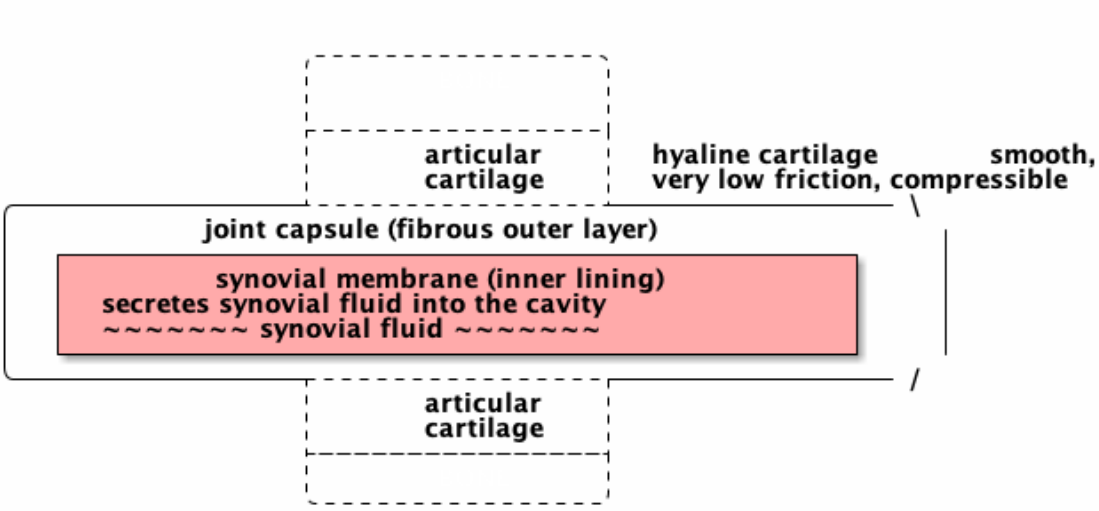
A single-pivot model therefore represents **one** function of the region. That is a legitimate simplification for a teaching aid with a flexion/extension objective (Chapter 21), but the omission must be stated. Saying "this model omits forearm rotation" takes four seconds and prevents a learner forming a wrong mental model.

6.4. Joints

6.4.1. Classification

Structural class	Binding material	Mobility	Example
Fibrous	Dense fibrous connective tissue	Little to none	Skull sutures; the distal tibiofibular joint
Cartilaginous	Cartilage	Limited	Intervertebral discs (fibrocartilage); the pubic symphysis; growth plates in children
Synovial	Joint capsule with a fluid-filled cavity	Free movement	Elbow, knee, hip, shoulder, fingers — almost everything you will model

6.4.2. Anatomy of a synovial joint



PLUS, outside the simple picture

- ligaments (capsular and separate) constraining motion
- muscles and tendons crossing the joint, actively stabilising
- bursae fluid sacs reducing friction where tendons cross bone
- in some joints menisci or labra (fibrocartilage) improving fit
- sensory nerve endings reporting position and pain

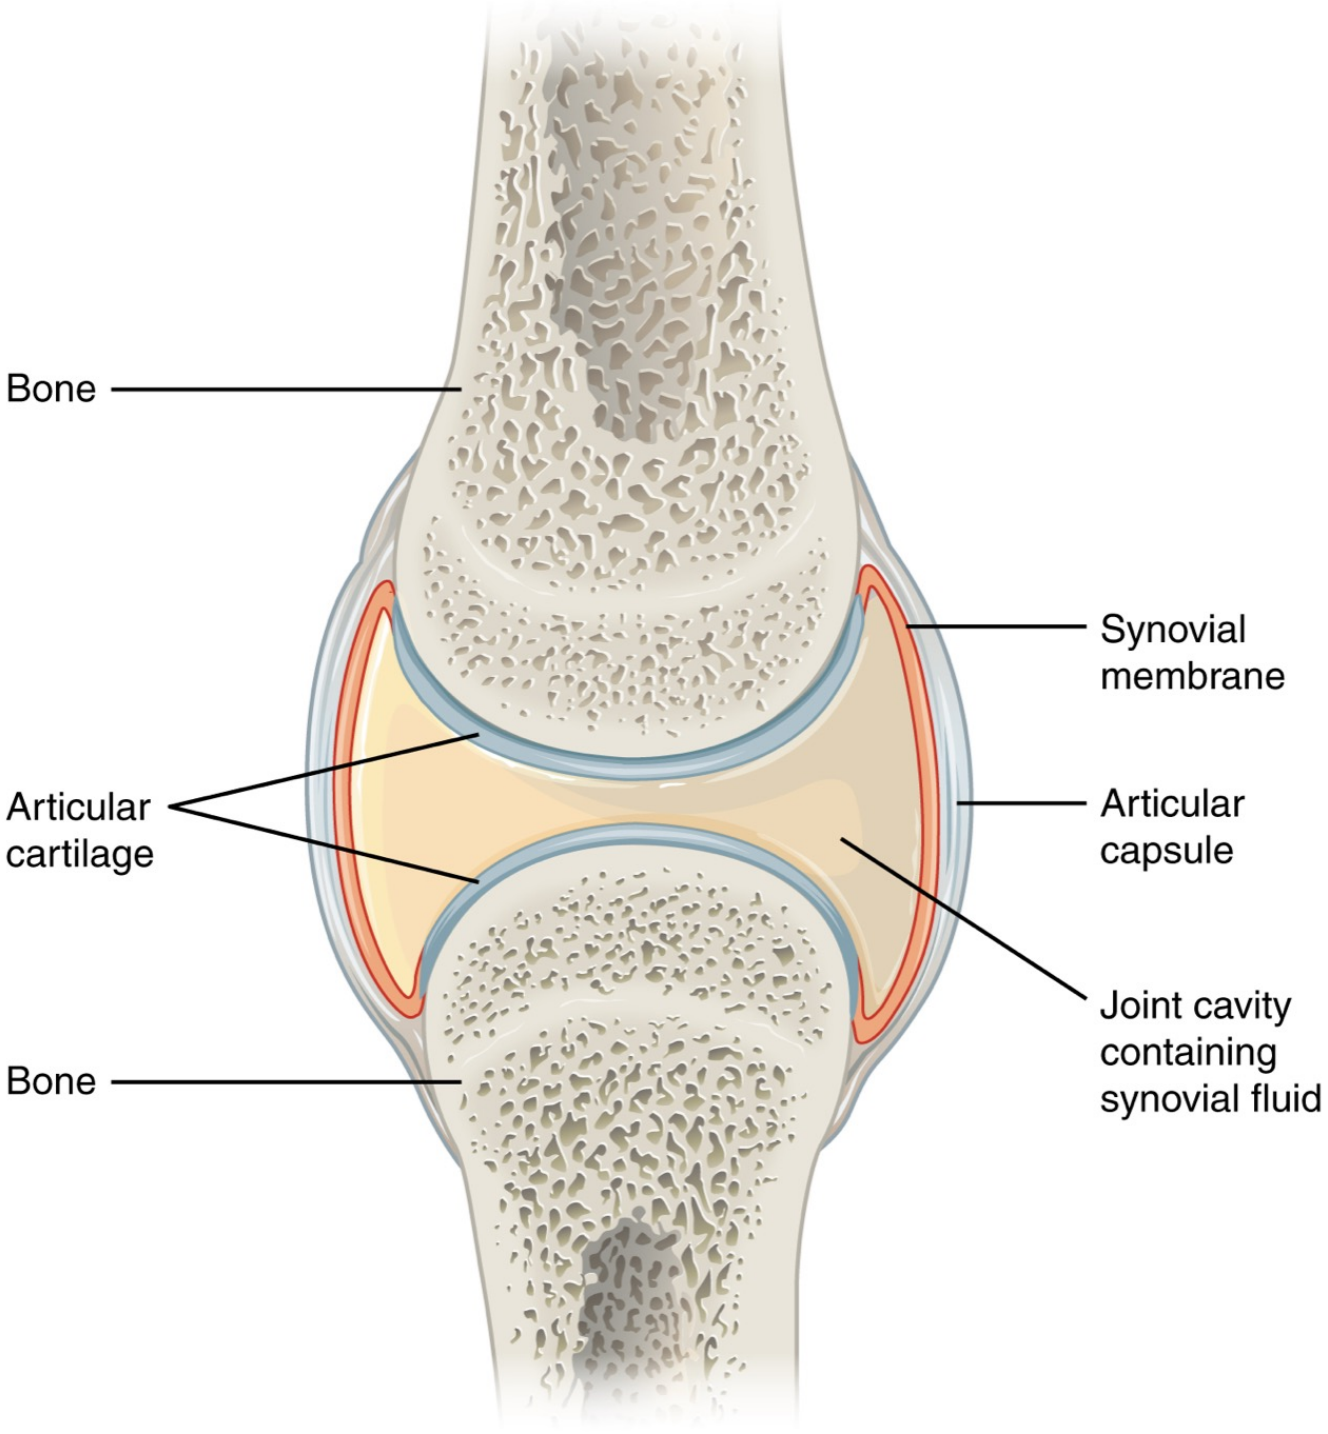


Figure 11. Structure of a synovial joint. Image: OpenStax College, CC BY 3.0. Credits: Appendix H

Coefficients of friction in a healthy synovial joint are extremely low — the literature quotes values in the region of 0.003–0.02 depending on conditions and method, which is better than most engineered dry bearings. The mechanism is not a simple fluid film: it involves fluid pressurised within the cartilage matrix being squeezed out under load (**weeping** or **biphasic** lubrication) alongside boundary lubrication by molecules adsorbed to the surface.

A metal hinge represents the **pivot**. It does not represent cartilage, capsule, fluid, ligamentous constraint, active muscular stabilisation, sensory feedback, or the fact that the centre of rotation moves through the range of motion. Each of those is a thing your model is silently claiming does not exist unless you say otherwise.

6.4.3. The six synovial joint types

Type	Degrees of freedom	Motion	Example
Plane (gliding)	2 (translational)	Sliding	Intercarpal, intertarsal joints
Hinge	1	Flexion/extension	Humeroulnar (elbow), interphalangeal
Pivot (trochoid)	1 (rotational)	Rotation about long axis	Proximal radioulnar; atlantoaxial (head shake)
Condyloid (ellipsoid)	2	Flexion/extension + abduction/adduction	Metacarpophalangeal (knuckle), radiocarpal (wrist)
Saddle	2 (with more range)	As condyloid, greater freedom	Thumb carpometacarpal
Ball-and-socket	3	All, plus rotation	Shoulder, hip

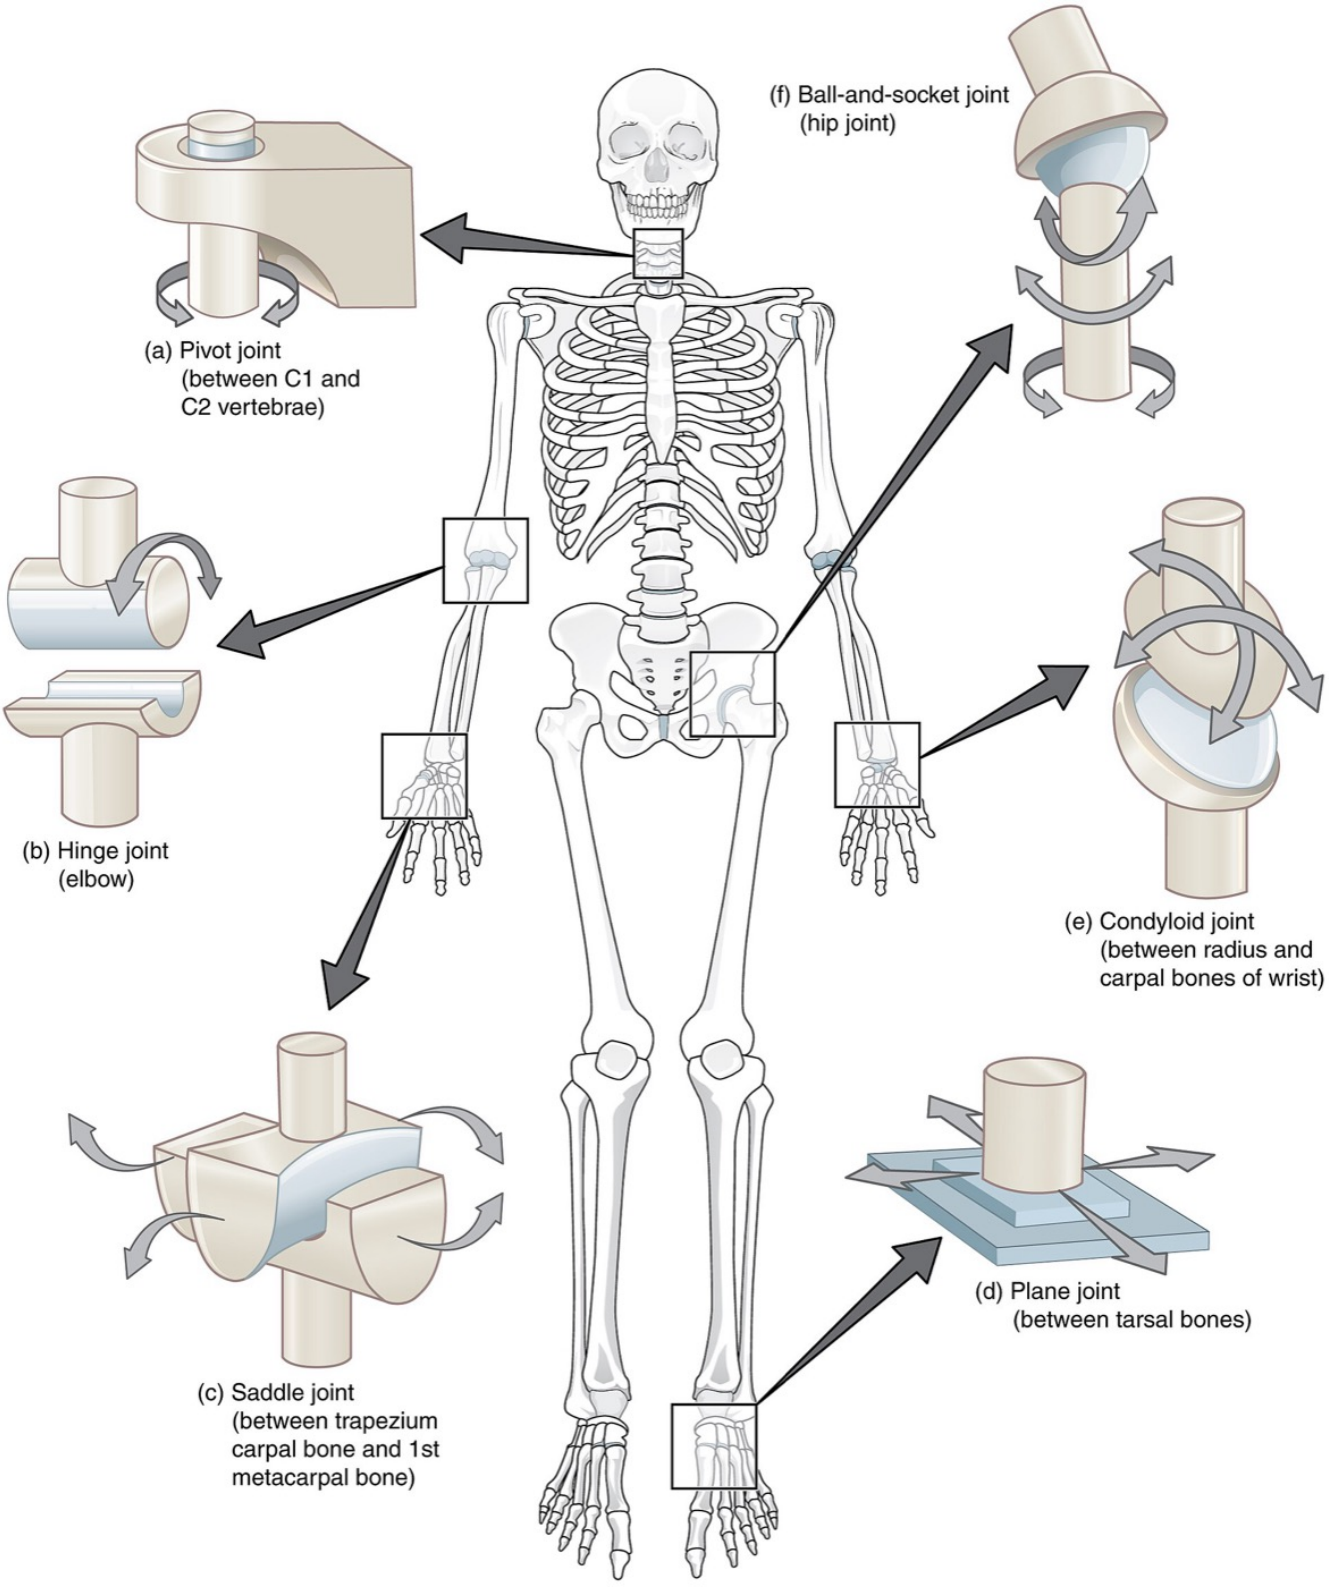


Figure 12. The six types of synovial joint. Image: OpenStax College, CC BY 3.0. Credits: Appendix H

A joint's degrees of freedom are not fixed by geometry alone. The glenohumeral (shoulder) joint is a shallow ball-and-socket with an enormous range, and that range is purchased at the cost of stability — it depends heavily on surrounding muscle, the labrum, and the capsule to stay located. The hip is a deep socket: more stable, less range. This trade-off between mobility and stability is a recurring theme, and it is the same trade-off you make in a mechanism.

6.5. Describing motion precisely

Term	Definition	Example
Flexion	Generally decreases the angle at a joint	Bending the elbow; curling the fingers
Extension	Generally increases the angle; returns from flexion	Straightening the elbow
Hyperextension	Extension beyond the usual anatomical range	Bending the wrist backward past straight
Abduction	Moves a part away from the midline	Raising the arm sideways; spreading the fingers
Adduction	Moves a part toward the midline	Lowering the arm to the side
Circumduction	Conical movement combining flexion, abduction, extension, adduction	Circling the arm at the shoulder
Rotation (medial/internal, lateral/external)	Turning about the long axis	Turning the whole arm inward or outward at the shoulder
Pronation	Forearm rotation bringing the palm to face posteriorly (down, with elbow flexed)	Turning a screwdriver anticlockwise with the right hand
Supination	Forearm rotation bringing the palm to face anteriorly (up, with elbow flexed)	Carrying a bowl of soup
Dorsiflexion / **plantarflexion**	Ankle: toes toward the shin / toes pointing down	Walking heel strike / pushing off
Inversion / **eversion**	Sole turns inward / outward	Rolling an ankle inward
Elevation / **depression**	Moves superiorly / inferiorly	Shrugging the shoulders
Protraction / **retraction**	Moves anteriorly / posteriorly in the horizontal plane	Pushing the jaw forward; pulling the shoulders back
Opposition	Thumb pad toward a fingertip	Picking up a coin — a defining human capability

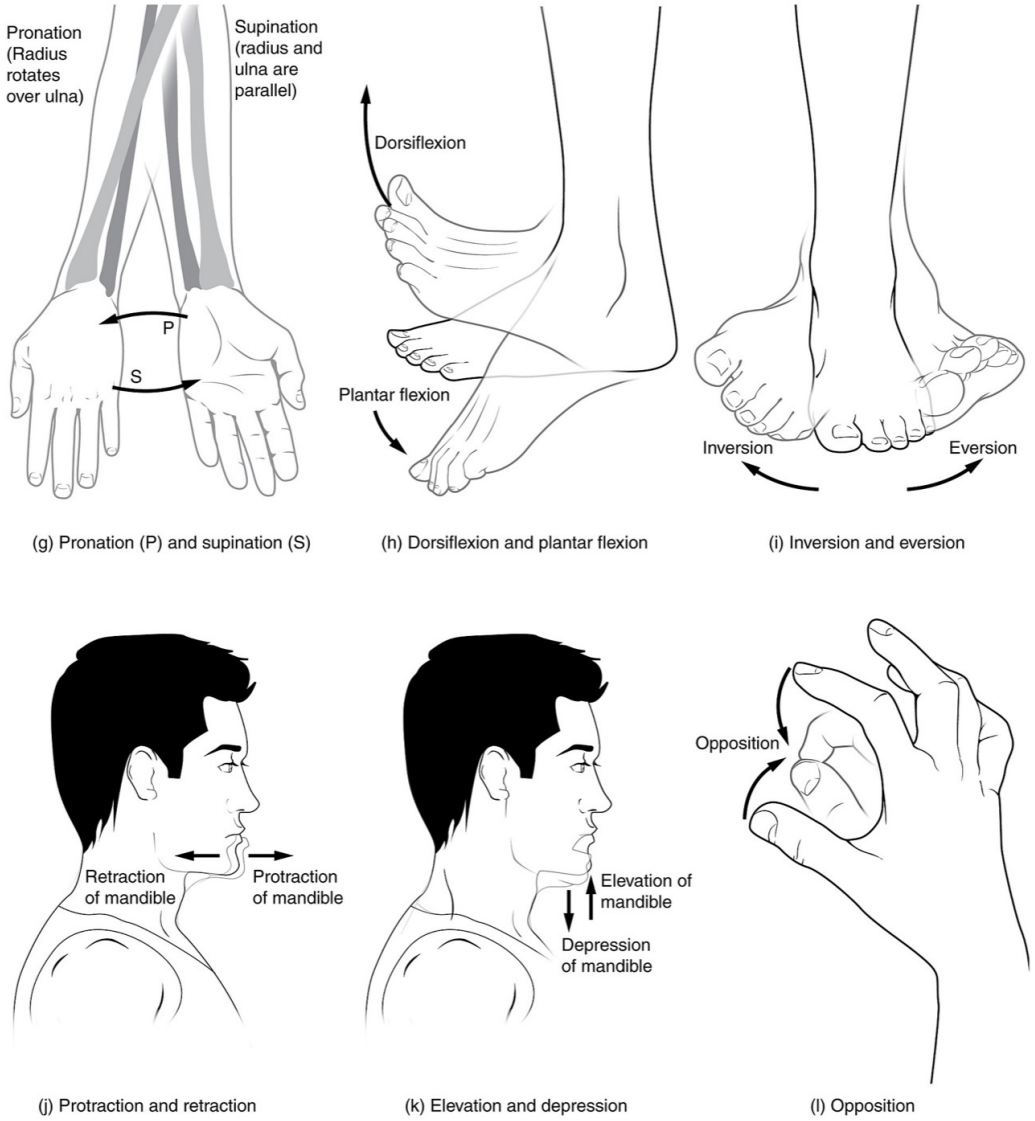


Figure 13. Movements of the body, including rotation, supination and pronation. Image: Tonye Ogele CNX / OpenStax, CC BY 4.0. Credits: Appendix H

Pronation and supination are the two that newcomers get wrong, so anchor them physically. Elbows bent, forearms horizontal: palm up is **supination** (you could carry **soup**). Palm down is **pronation**. The visible result is a hand orientation; the actual movement is the radius rotating about the ulna, and if the elbow is extended, shoulder rotation can disguise or substitute for it. If you ask a person to "turn your palm up" with the arm straight, you may be measuring their shoulder.

6.5.1. Active versus passive movement

Term	Meaning
Active range of motion (AROM)	The range the person can move through using their own muscle action
Passive range of motion (PROM)	The range through which the joint can be moved by an external force
Functional range	The range actually used, comfortably and repeatedly, for a real task

These three are different numbers for the same joint, and the differences are diagnostically meaningful — which is exactly why their measurement is a clinical skill.



Being able to move a model through an angle tells you nothing about a person.

Being able to move **your own** joint through an angle tells you about your own joint on that day. It does not tell you about anyone else's safe or comfortable range, and it certainly does not license you to move someone else's joint. Forcing a joint can injure it.

Range-of-motion assessment, and any device that restricts motion or supports load, needs a physiotherapist's or occupational therapist's functional assessment. Not a diagnosis and not a table — an assessment of **this person**.

6.5.2. Typical ranges, for orientation only

Joint and motion	Approximate adult range	Caveat
Elbow flexion	~0–145°	Varies; many daily tasks use only ~30–130°
Forearm pronation/supination	~80–90° each way	Measured with the elbow flexed
Wrist flexion / extension	~70–80° / ~60–70°	Functional tasks use far less
Wrist radial / ulnar deviation	~20° / ~30°	—
Shoulder flexion	~0–180°	Includes substantial scapulothoracic contribution
Shoulder abduction	~0–180°	As above
Knee flexion	~0–135°	—
Hip flexion	~0–120°	More with the knee flexed
Ankle dorsiflexion / plantarflexion	~20° / ~50°	—



These figures are textbook orientation values for a typical adult [1]. They vary with age, sex, build, activity, prior injury, and measurement method, and published "normal" ranges differ between sources. **Never design to these numbers for a specific person.** Measure — or rather, have a therapist measure — the person. Then design for less than their maximum, because a design that requires someone's end range will be uncomfortable, fatiguing, and abandoned.

6.6. Muscle: how force is actually produced

6.6.1. Structure

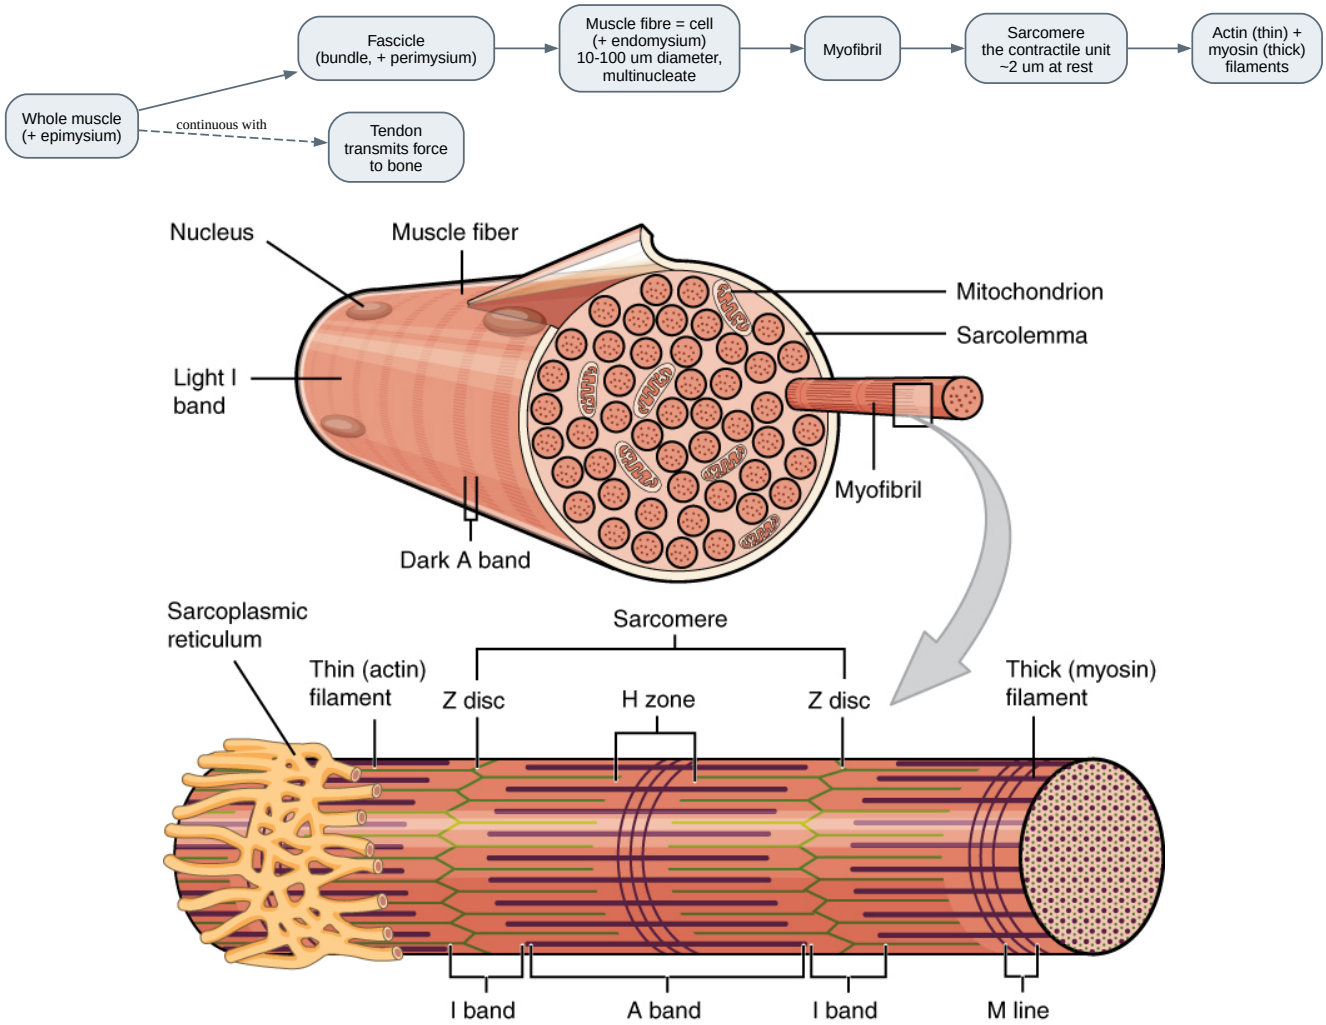


Figure 14. Structure of a skeletal muscle fibre. Image: OpenStax, CC BY 4.0. Credits: Appendix H

Force is generated by myosin heads binding actin and ratcheting — the **sliding filament** mechanism. Calcium released inside the fibre, triggered by a nerve signal at the **neuromuscular junction**, exposes the binding sites. ATP powers the cycle. No calcium, no contraction; no ATP, no release either (which is why rigor mortis happens).

A **motor unit** is one motor neuron plus every muscle fibre it innervates. Force is graded two ways: **recruiting** more motor units, and firing them **faster**. Small precise muscles (extraocular, hand intrinsics) have small motor units — a few fibres per neuron, giving fine control. Large postural muscles have hundreds of fibres per neuron.

> **This is the fact a maker needs most:** the **motor** is the muscle. The **tendon** is a tension member that transmits the force. They are not the same component, and the distinction is the single most common misconception a simplified joint model creates.
>
> The running project in this book exists largely to teach that one distinction well.

6.6.2. Types of contraction

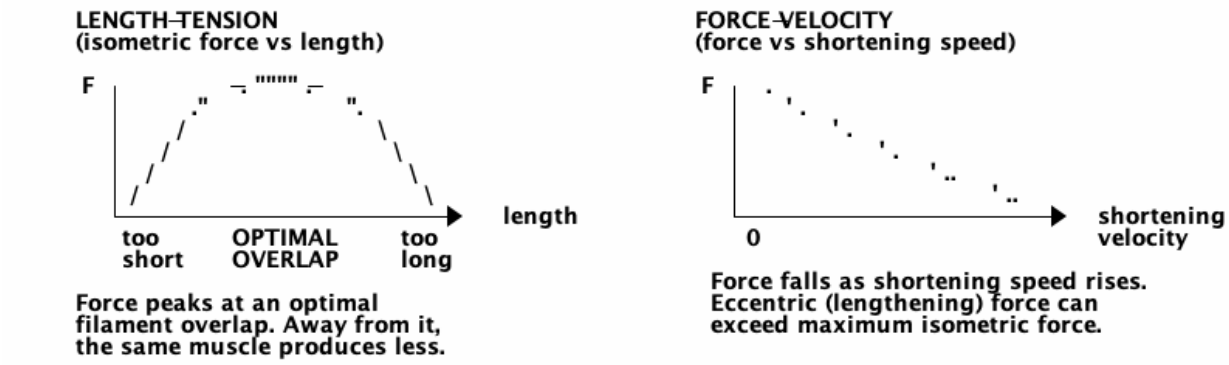
Type	What happens	Everyday example
Isometric	Force produced, muscle length ~constant, no joint movement	Holding a bag still at your side
Concentric	Muscle shortens while producing force; it does positive work	Lifting a bag; the biceps shortening during elbow flexion
Eccentric	Muscle **lengthens** while producing force; it absorbs energy	Lowering a bag slowly; quadriceps on a descending stair

A muscle does not have to shorten in order to produce force. Holding a position is active work. Lowering a load slowly is active work — and eccentric contraction can produce higher forces than concentric.

Two design consequences. First, "the task doesn't involve movement, so it isn't tiring" is wrong: sustained isometric holding is strongly fatiguing, and a design that requires someone to hold a position is making a real demand. Second, if your model represents a muscle as something that only pulls things together, you have omitted half of what muscles do — including the controlled-lowering half that most injuries occur during.

6.6.3. Length–tension and force–velocity

Two relationships that distinguish muscle from a linear actuator:



So the force a muscle can produce depends on the joint angle (which sets muscle length) and on how fast the joint is moving. A "maximum grip strength" figure is therefore a measurement under specific conditions, not a property of the person.

6.6.4. Agonist, antagonist, synergist, stabiliser

Muscles work in groups, not alone.

Role	Function	Elbow flexion example
Agonist (prime mover)	Produces the intended movement	Biceps brachii, brachialis

Role	Function	Elbow flexion example
Antagonist	Opposes it; controls and decelerates the movement	Triceps brachii
Synergist	Assists the agonist, often adding or refining direction	Brachioradialis
Fixator / stabiliser	Holds a proximal segment steady so the movement is effective	Scapular stabilisers holding the shoulder blade

The antagonist is not the enemy. **Co-contraction** — agonist and antagonist firing together — stiffens and stabilises a joint, which is how you keep a hand steady while doing something delicate. It also costs energy, which is why precision tasks are tiring.

6.7. Engineering bridge: levers, moment arms, and statics

6.7.1. Torque

$\boldsymbol{\tau = F \times d_{\perp}}$

Torque is force multiplied by the **perpendicular** distance from the pivot to the force's **line of action**. That perpendicular distance is the **moment arm**.



The moment arm is **not** the distance along the beam to the attachment point. It is the perpendicular distance from the pivot to the line along which the force acts.

This distinction is the single most common arithmetic error in a first biomechanics calculation, and it is also why a cable becomes less effective as a joint changes angle: as the geometry rotates, the perpendicular distance changes even though nothing moved along the beam.

6.7.2. Worked statics: the idealised elbow lever

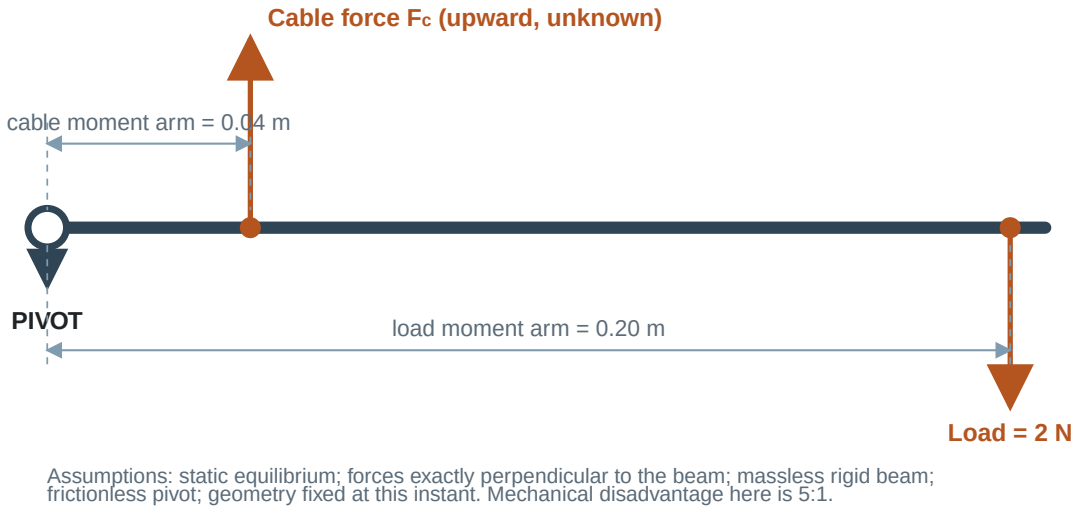


Figure 15. Idealised elbow lever: a pivot, a cable acting at a short moment arm, and a load acting at a long one

Load torque about the pivot:

$\boldsymbol{\tau_L} = \mathbf{F_L} \times \mathbf{d_L} = 2\ \text{N} \times 0.20\ \text{m} = \mathbf{0.40\ N{\cdot}m}$

For static equilibrium the cable must produce an equal and opposite torque:

$\mathbf{F_c} = \boldsymbol{\tau_L} \div \mathbf{d_c} = 0.40\ \text{N{\cdot}m} \div 0.04\ \text{m} = \mathbf{10\ N}$

Halve the moment arm to 0.02 m and the required force doubles to 20 N. The **mechanical disadvantage** here is 5:1 — the cable pulls five times the load it balances.

That ratio is not an artefact of a bad model. It is roughly how your own elbow works. The biceps tendon inserts close to the elbow, so its moment arm is small compared with the distance to your hand, and the muscle must generate several times the force of the load it lifts. The body trades force for **speed and range**: a small muscle shortening produces a large, fast hand movement.

> **Every simplification in that calculation matters.**
>
> What we ignored: the limb's own weight (which is not negligible — a forearm and hand is roughly 2% of body mass, so ~1.4 kg or ~14 N for a 70 kg adult); friction; acceleration (so this applies only to a static hold); the fact that the real moment arm changes continuously with joint angle; the fact that multiple muscles contribute simultaneously and the distribution between them is genuinely indeterminate from statics alone; soft tissue compliance; and joint reaction forces, which are often much larger than either the load or the muscle force alone.
>
> The arithmetic is teachable **because** it is simplified. That same simplification makes it unsuitable as a prescription for a real person. Use it to understand **why attachment geometry matters**, not to specify a force.

6.7.3. Including the limb's own weight

A slightly more honest version. Model the forearm-plus-hand as a rigid segment of mass 1.4 kg (weight ≈ 13.7 N) with its centre of mass 0.14 m from the elbow, holding a 2 N object at 0.33 m, with a muscle moment arm of 0.04 m.

$\mathbf{F_m}$ = [(13.7 × 0.14) + (2 × 0.33)] ÷ 0.04 = (1.92 + 0.66) ÷ 0.04 = **64.4 N**

The segment's own weight dominates: it contributes nearly three times the torque of the held object. This is why "just hold your arm out" is tiring with nothing in your hand, and why a design that requires a sustained unsupported arm position is making a much larger demand than its designer thinks.



Anthropometric segment parameters — segment masses as fractions of body mass, and centre-of-mass locations as fractions of segment length — are tabulated in the biomechanics literature [4]. They are population averages from cadaver and imaging studies. They are good enough for the kind of order-of-magnitude reasoning above and not good enough to characterise an individual.

6.7.4. Lever classes

Class	Arrangement	Property	Body example
First	Fulcrum between effort and load	Can favour either force or speed	Head nodding on the atlanto-occipital joint; triceps extending the elbow
Second	Load between fulcrum and effort	Always mechanical **advantage** (force multiplier)	Standing on tiptoe — rare in the body
Third	Effort between fulcrum and load	Always mechanical **disadvantage**; favours speed and range	Biceps flexing the elbow — the most common arrangement in the body

The body is dominated by third-class levers. It has largely optimised for **range and speed of movement** rather than for force efficiency, and it pays for that with high internal muscle and joint forces. A maker whose instinct is "that's a terrible mechanical arrangement" has correctly identified the trade-off and missed its purpose.

6.7.5. Grip, pinch, and the problem with strength figures

Grip type	Description	Typical use
Power grip (cylindrical)	Whole hand wraps, fingers flexed, thumb opposing	Hammer, handrail, jug handle
Hook grip	Fingers flexed, thumb not involved	Carrying a bag

Grip type	Description	Typical use
Precision / pinch grip	Fingertips, thumb opposing	Holding a pen, a coin, a small latch
Lateral (key) pinch	Thumb pad to the side of the index finger	Turning a key
Tripod (three-jaw) pinch	Thumb to index and middle fingertips	Writing, holding a plug
Spherical / disc grip	Hand spread around a round object	Opening a jar lid

Published adult grip-strength figures span a wide range — the literature reports typical maximal values in the region of a few hundred newtons for healthy young adults, varying hugely with age, sex, hand dominance, posture, wrist angle, dynamometer type, and measurement protocol. Pinch strengths are an order of magnitude lower.



Three reasons not to design to a strength figure:

1. **Maximum is not sustainable.** The force someone can exert once, maximally, in a laboratory is not the force they can produce repeatedly for a minute while doing something else. A widely used ergonomic heuristic is to design for well under a third of maximal voluntary contraction for sustained tasks — but the honest version of that statement is "design for as little force as the function allows, then test with the actual user".
2. **Your users may be nowhere near the population figure.** The whole reason an assistive device exists is often that the person's strength, range, or control differs from the population from which the table was built. Designing to the 5th percentile of a healthy population is still designing for a healthy population.
3. **Strength is only one demand.** Someone may have ample force and insufficient **pinch precision**, or ample precision and insufficient **endurance**, or all three and no way to **see** whether the latch engaged. Chapter 7 decomposes this properly.

6.7.6. Anthropometry and percentiles

Body dimension data exists in standard tables by population, sex, and age [5]. Two things to understand about it:

- **Percentiles are per-dimension, not per-person.** There is no "95th percentile person". Someone at the 95th percentile for stature may be at the 40th for hand breadth. Designing a product by stacking 95th-percentile dimensions produces a shape that fits nobody.
- **Design for a range, not a point.** Adjustability, or a size range, almost always beats a single optimised dimension. The exception is where adjustability introduces a mechanism that can be set wrongly — then you are trading one risk for another, and the trade needs to be explicit.

6.8. Devices, software, and training technology

Musculoskeletal medicine is the most mechanical specialty, and therefore the one where a maker's existing skills — machining, sheet metal and welding, thermoforming, casting, sewing, electronics, and mechanism design — transfer most directly. It is also where the regulatory stakes rise fast — a teaching model and a load-bearing orthosis are very different projects.

6.8.1. Clinical devices

Device	Does	Engineering note
Joint replacement (arthroplasty)	Replaces articular surfaces — hip, knee, shoulder, and others	Bearing couples (metal-on-polyethylene, ceramic-on-ceramic), wear debris, fixation (cemented or bone-ingrowth), and **stress shielding** (Chapter 6). Highly regulated
Fracture fixation — plates, screws, nails, external fixators	Holds bone in position while it heals	Stiffness must be tuned: too flexible and the fracture moves; too stiff and the bone is shielded from load it needs
Orthoses — splints, braces, AFOs, spinal braces	Support, immobilise, offload, or correct alignment	Interface pressure, oedema accommodation, don/doff, and compliance. **Clinical prescription and fitting, not a maker decision**
Prostheses — limb replacement	Replaces a missing limb segment	The **socket interface** is the hard part and the usual cause of abandonment. Laminated carbon fibre, thermoformed sockets, machined titanium and aluminium components, textile liners. Passive, body-powered, and myoelectric control
Mobility aids — sticks, frames, crutches, wheelchairs	Support, offload, and enable movement	Seating and wheelchair prescription is a clinical specialty. Height, grip, and stability are prescribed, not guessed
Assistive and adapted equipment	Grip aids, adapted cutlery, reachers, button hooks, jar openers	**The classic maker-and-OT collaboration.** Individual fit is the whole value proposition. Bent and machined aluminium or stainless, thermoformed or mouldable thermoplastic, turned handles, sewn cuffs and straps, and occasionally a printed one-off

Device	Does	Engineering note
Exoskeletons and powered orthoses	Assist or restore movement	Actuation, control, alignment with the biological joint axis, safety on power loss
Continuous passive motion (CPM) machines	Move a joint through range after surgery	Range limits, force limits, and a reliable stop
Dynamometers and goniometers	Measure strength and joint angle	Protocol dependence is enormous (Chapter 6). Instrumented versions are common
DEXA scanner	Bone mineral density	Dual-energy X-ray absorptiometry. See Chapter 15
Therapeutic ultrasound, TENS, shockwave	Pain and tissue therapy modalities	Dose, duration, and contraindications are clinically specified

6.8.2. Software

Software	Does	Note
OpenSim	Open-source musculoskeletal modelling and simulation — inverse kinematics, inverse dynamics, muscle force estimation	**Free, standard in the research community, and directly usable.** Turns the statics in this chapter into a real tool
Motion capture and gait analysis — Vicon, OptiTrack, and markerless systems	Measures movement in three dimensions	Marker placement error dominates. Markerless pose estimation is improving fast
OpenPose, MediaPipe, and other pose estimators	Markerless body-landmark tracking from ordinary video	**Free and runnable on a laptop.** Accuracy is not research-grade; state that
Surgical planning software	Pre-operative templating and implant sizing from imaging	A regulated category. See the pipeline in Chapter 15
Orthotic and prosthetic CAD	Socket and brace design from scan data	Scan-to-fabrication workflows — CNC-carved moulds for thermoforming and lamination, and increasingly direct additive manufacture — are now mainstream in P&O
Rehabilitation and exercise apps	Prescribe, demonstrate, and track exercise programmes	Adherence is the real clinical problem these address

Software	Does	Note
Finite element analysis — any general FEA package	Stress analysis of bone, implants, and devices	Material models, boundary conditions, and validation are everything. **An unvalidated FEA result is a colourful picture**

6.8.3. Training and simulation

Training technology	Teaches	Form
Articulated skeletons and joint models	Structure, movement terminology, joint mechanics	**Exactly your capstone's category** (Chapter 22). Universally used, frequently broken
Fracture fixation trainers	Drilling, screw placement, plate contouring	Synthetic bone (composite "sawbones" style), often single use. Consumed in quantity
Arthroscopy simulators	Triangulation, scope handling, portal placement	Box trainers with cameras, and VR systems with haptic feedback
Joint injection and aspiration trainers	Landmark identification, needle placement	Often combined with ultrasound guidance phantoms (Chapter 15)
Casting and splinting practice limbs	Application technique, padding, moulding	Simple limb forms; a straightforward making project
Gait and movement demonstration rigs	Lever mechanics, moment arms, muscle action	Cable-and-pulley models. **Your elbow model is the simplest member of this family**
Manual therapy and palpation trainers	Locating structures by touch, joint mobilisation	Requires genuine tactile fidelity (Chapter 21)
Transfer and manual handling training	Safe moving of people; protects staff and patients	Weighted manikins, hoists, slide sheets
Instrumented movement and strength rigs	Measurement technique, protocol dependence, moment arms	Load cells, encoders, and IMUs on a machined or welded frame, logged by a microcontroller. The same electronics underlie clinical dynamometers — a teaching rig makes the protocol effects visible



The line to watch in this chapter is load bearing.

> A model that demonstrates movement is a teaching aid. The moment an object **supports a person's weight**, **restricts a joint's movement**, or **applies force to a limb**, it has become an orthosis or a mobility aid — a prescribed clinical device with assessment, fitting, review, and regulatory obligations behind it.
>
> That transition is easy to make accidentally. "Could you make it a bit stronger so they can lean on it?" is a change of project, not a change of specification. See the contact categories in Chapter 5 and the need discipline in Chapter 18.

6.9. Worked example: moving the cord attachment

Your elbow model has a cord attached 40 mm from the pivot and a guide loop on the upper-arm strip 80 mm above the pivot. You pull. The forearm strip rises.

Now move the attachment to 20 mm from the pivot and pull again. Three things change:

1. **Required force roughly doubles** for the same torque, because the moment arm halved.
2. **The range of movement for a given cord travel increases** — you pull less cord to achieve the same angular sweep. Force traded for range: exactly the body's bargain.
3. **The force direction relative to the beam changes through the motion**, so the effective moment arm varies through the sweep differently than before. The model will feel "strongest" at a different angle.

Now the teaching question. A learner watching this will conclude something about anatomy. What?

If the cord is unlabelled, the likely conclusion is "the cord is the muscle, and it pulls". That is wrong, and your demonstration taught it. The repair is in the labels:

- the free end you pull → **"external force input representing muscle action"**,
- the cord itself → **"tendon representation: transmits force, does not generate it"**,
- the pivot → **"simplified joint: omits cartilage, capsule, ligaments, and the moving centre of rotation"**,
- the guide loop → **"model convenience, not an anatomical tendon route"**.

That last one is important and easy to skip. Your guide loop is where it is because it makes the mechanism work on a tabletop. The real biceps tendon insertion is somewhere else entirely. If the educator wants to teach **actual attachment anatomy**, this design needs substantial revision and anatomical review — it is not a matter of moving the loop.

6.10. Workshop exercise

1. **Build the articulated model.** Two stiff card or lightweight wooden strips, approximately 200 × 35 mm, joined by a freely rotating pivot about 20 mm from one end of each. Round or cover sharp corners.
2. **Predict before you pull.** Attach a non-elastic cord to the moving strip 40 mm from the pivot, passing through a smooth guide on the fixed strip 80 mm from the pivot. **Write down your**

prediction of the direction of movement before pulling. Then pull gently.

3. **Vary the geometry.** Move the attachment to 20 mm and then to 60 mm. Each time: predict, then observe. Record what happens to required force and to the angular sweep per unit of cord travel.
4. **Calculate.** For one configuration, compute the required cord force to balance a 2 N load, stating every assumption. Then estimate the error introduced by ignoring the strips' own weight.
5. **Document the gap.** Write a list titled **"anatomical features this model does not represent"**. Aim for at least eight entries. You will reuse this list in Chapter 21 and Chapter 22.
6. **Functional brief.** Write a one-page functional brief for a hypothetical grip aid that addresses the task, the user, the environment, duration of use, and fatigue — **without** specifying a material or a manufacturing process.



Keep fingers clear of pinch points at the pivot and where the strips approach each other. Keep the model off the body entirely. Do not add significant loads: this exercise is about **motion and geometry**, and large forces add no educational value while adding real risk of a part failing suddenly.

6.11. Questions

1. Why is a single-pivot elbow model incomplete? Name at least four omissions.
2. If the moment arm halves while the required torque stays constant, what happens to the required force?
3. Why can a cable become less effective as a model changes angle, even though nothing moved along the beam?
4. What label prevents a cord from teaching the wrong biological mechanism?
5. Which is the lateral forearm bone in anatomical position, and which rotates during pronation?
6. Distinguish AROM, PROM, and functional range. Which one should a design target?
7. A person holds a 2 N object in their hand with the elbow at 90°. Is the muscle force roughly 2 N, 10 N, or 60 N? Why?
8. Why is an implant that is much stiffer than bone potentially a problem?
9. You measure your own comfortable wrist extension as 65°. What may you conclude about a user’s design requirement?
10. Name a task that is fatiguing despite involving no movement, and say why.
11. Why is there no "95th percentile person"?
12. The body mostly uses third-class levers, which are mechanically disadvantageous. Why?

6.12. Answers and reasoning

1. It omits forearm pronation and supination (the proximal radioulnar articulation); the fact that the region contains three articulations; articular cartilage; the joint capsule and synovial fluid; ligamentous constraint; active muscular stabilisation and co-contraction; the moving centre of

rotation through the range; sensory feedback; and everything living — blood supply, repair, adaptation. Any four of these.

2. It **doubles**. $\tau = F \cdot d\perp$, so at constant torque force is inversely proportional to the moment arm.
3. Because the moment arm is the **perpendicular** distance from the pivot to the force's line of action, and rotating the beam changes that perpendicular distance even with the attachment point fixed on the beam.
4. Something that separates generation from transmission, for example: **"Cord represents a tendon transmitting force; the demonstrator supplies the muscle force."** The key property is that it names what the cord is **not**.
5. The **radius** is lateral (thumb side) in anatomical position, and the radius is the bone that rotates about the ulna during pronation and supination.
6. **AROM** is what the person can achieve with their own muscles; **PROM** is what the joint permits when moved externally; **functional range** is what they actually use comfortably and repeatedly for a real task. Design to the **functional** range, with margin — and have a therapist establish it, not a table.
7. Roughly **60 N** or more. The forearm and hand's own weight acts at a moment arm several times larger than the muscle's, so segment weight dominates. The held object's 2 N contributes only a small fraction. This is the calculation in the Engineering bridge above, and the intuition it buys — **internal forces are much larger than external loads** — is one of the most useful things in this chapter.
8. Because bone remodels in response to the load it carries. A much stiffer implant carries a disproportionate share of the load, the adjacent bone is unloaded, and it resorbs — **stress shielding**. The goal is appropriate load sharing, not maximum stiffness.
9. Almost nothing. It is one measurement, of one person, on one day, with your own measurement technique. The user's range may differ substantially, their **comfortable repeated** range differs from their maximum, and the relevant quantity is what they can sustain while performing the actual task. Have a therapist assess the user.
10. Holding a weight stationary at arm's length; holding a precision tool steady; standing still for a long period. Isometric contraction produces force continuously and consumes energy, and sustained co-contraction for stability is additionally costly. No external work is done, but the muscle is working.
11. Because percentiles are computed independently per dimension, and a person's percentile rank differs across dimensions. Stacking the 95th percentile of every dimension produces a composite that matches no real body. Design for a range.
12. Because the body optimised for **range and speed of movement** rather than force efficiency. A small muscle excursion near the joint produces a large, fast movement at the hand. The cost is high internal muscle and joint forces — a deliberate trade, not a design error.

6.13. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Levers, moment arms and contact — the elbow lever: moment arm, cord force and segment weight through the sweep.
- Tissues, materials and temperature — stiffness, strength and stress shielding.

6.14. Video companion and reading

- **Core videos:** Crash Course Anatomy & Physiology #20 (joints) [11], then episodes 19 and 21–22 on the skeletal system and muscles from the series index [6].
- **Core reading:** OpenStax [1] chapters 6–11, selectively:
  - 6.3 Bone Structure and 6.6 Exercise, Nutrition, Hormones, and Bone Tissue
  - 8.2 Bones of the Upper Limb
  - 9.1 Classification of Joints, 9.4 Synovial Joints, 9.5 Types of Body Movements, 9.6 Anatomy of Selected Synovial Joints
  - 10.3 Muscle Fiber Contraction and Relaxation, 10.4 Nervous System Control of Muscle Tension
  - 11.1 Interactions of Skeletal Muscles, Their Fascicle Arrangement, and Their Lever Systems — this one directly covers the torque material above
- **Optional depth:** the musculoskeletal component of Michigan's Anatomy Specialization [2].
- **For the biomechanics:** Winter, *Biomechanics and Motor Control of Human Movement* [4], is the standard reference for segment parameters and inverse dynamics. Borrow it if the statics interests you; it is not required.

6.15. Exit criteria

- ☐ Identify ten major bones relevant to your chosen limb, including which forearm bone is lateral and which rotates.
- ☐ Demonstrate or diagram eight movement terms accurately, including pronation and supination with the correct anchor.
- ☐ Explain your model's force path and at least **five** ways it differs from living anatomy.
- ☐ Complete the torque calculation, including the version with segment weight, and state every simplifying assumption.
- ☐ Explain the difference between AROM, PROM, and functional range, and say which a design should target and who establishes it.
- ☐ Explain why internal muscle and joint forces exceed external loads, with a number.
- ☐ Write a functional brief for a hypothetical grip aid covering task, user, environment, duration, and fatigue, without specifying a solution.
- ☐ Identify when a proposed device would require therapist assessment — specifically, if it restricts movement, supports a load, is worn for a prolonged period, or depends on the person's range or strength.

6.16. What this chapter does not establish

- **You cannot assess range of motion.** You know the terminology and that AROM, PROM, and functional range differ. Measuring them on a person, and interpreting what the difference means, is a clinical skill.
- **You cannot specify forces for a person.** The statics is correct for a rigid idealised beam. A real limb involves indeterminate muscle distribution, changing geometry, compliant tissue, and large joint reaction forces. Order-of-magnitude reasoning only.
- **You cannot design a load-bearing or motion-restricting device unsupervised.** Splints, braces, supports, and anything that takes body weight all require clinical involvement from the start, not at review.
- **The strength and ROM tables are not design values.** They are population orientation figures, and your users are frequently **not** from that population.
- **Your model is not an anatomy reference.** It is one selected relationship, honestly labelled. Chapter 21 deals with making it teach the right thing.

Chapter 7. Nerves, sensation, and accessible design

Week 10 · approximately 4–6 hours

7.1. Why this chapter exists

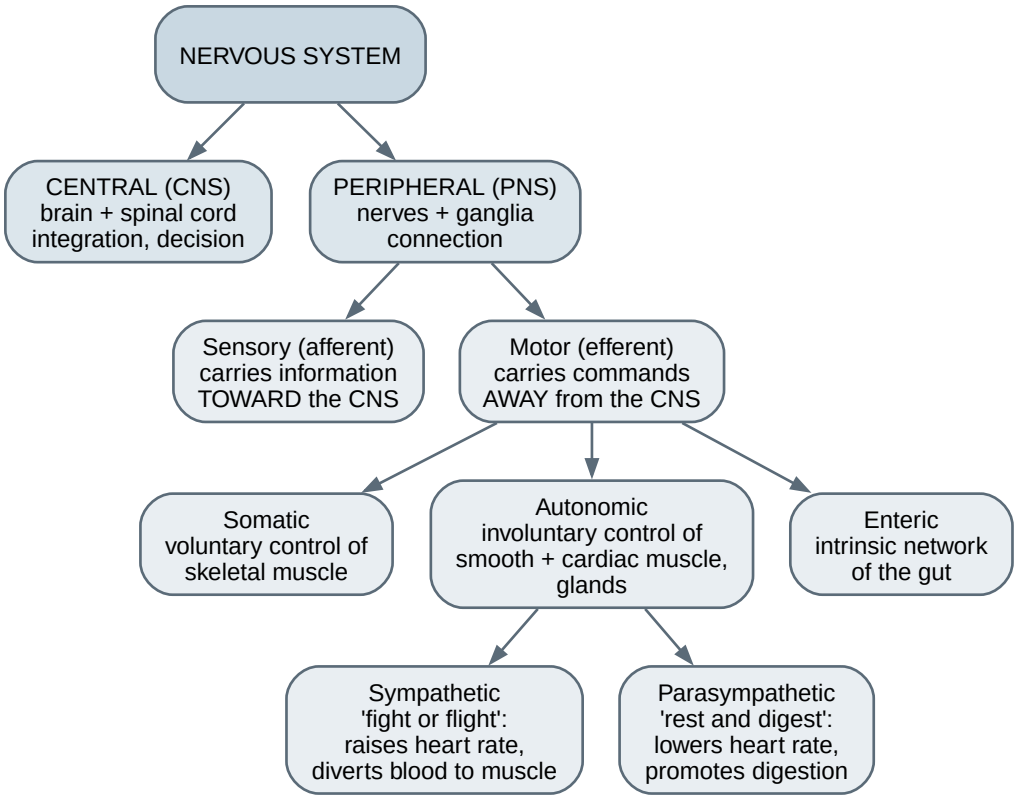
"Can the user operate this control?" is not one question. It is at least four, and they fail independently.

A person may have ample strength and insufficient pinch precision. Another may operate the mechanism perfectly and have no way to see whether it latched. A third may see it, understand it, and be unable to reach it. A fourth may do all of it and be unable to **feel** that the strap beneath is cutting in.

Moving, feeling, seeing, hearing, and interpreting are separate capabilities supported by separate machinery. This chapter gives you enough neurophysiology to know why, and enough design method to act on it.

It is also the chapter closest to classical biomedical engineering, because nerves and muscles are electrical, and electrical signals are measurable. If you ever touch EMG, ECG, EEG, or any biopotential instrument, the physiology here is its foundation.

7.2. The nervous system, organised



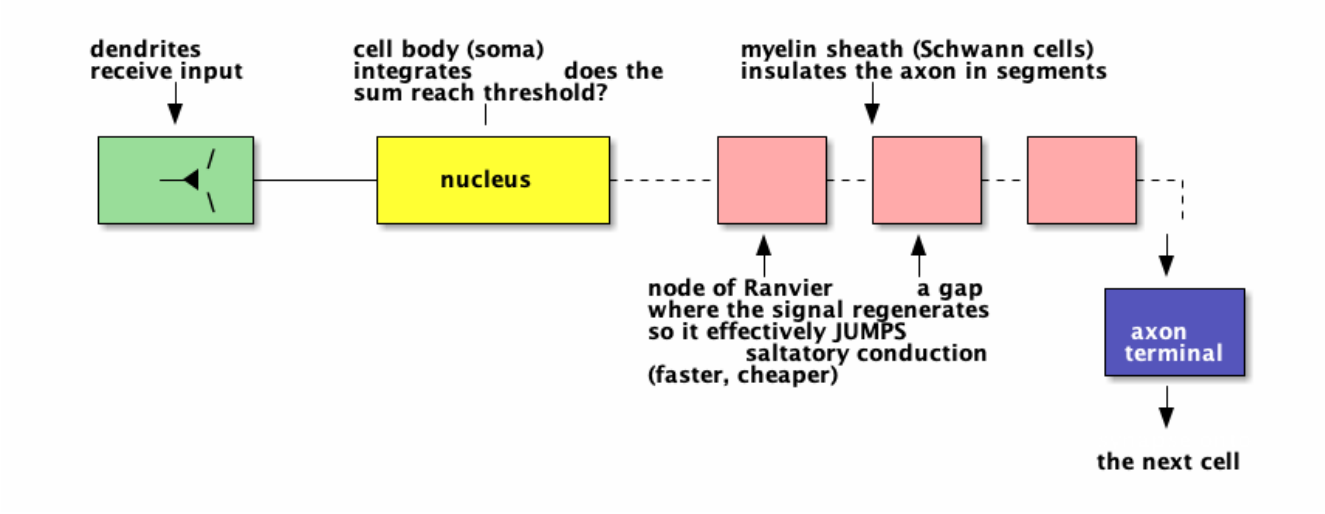
Sensory (afferent) pathways carry information toward the CNS. **Motor (efferent)** pathways carry

commands away from it. The words are worth learning because clinical language uses them constantly, and because the distinction is exactly the one your design review needs: **input to the person** versus **output from the person**.

The **autonomic** division is largely outside conscious control and matters to you mainly as context: it is why heart rate rises with anxiety, why a painful or alarming procedure produces physiological effects, and why "the patient was agitated" is a physiological statement as well as a behavioural one.

7.3. How a neuron works

7.3.1. Structure



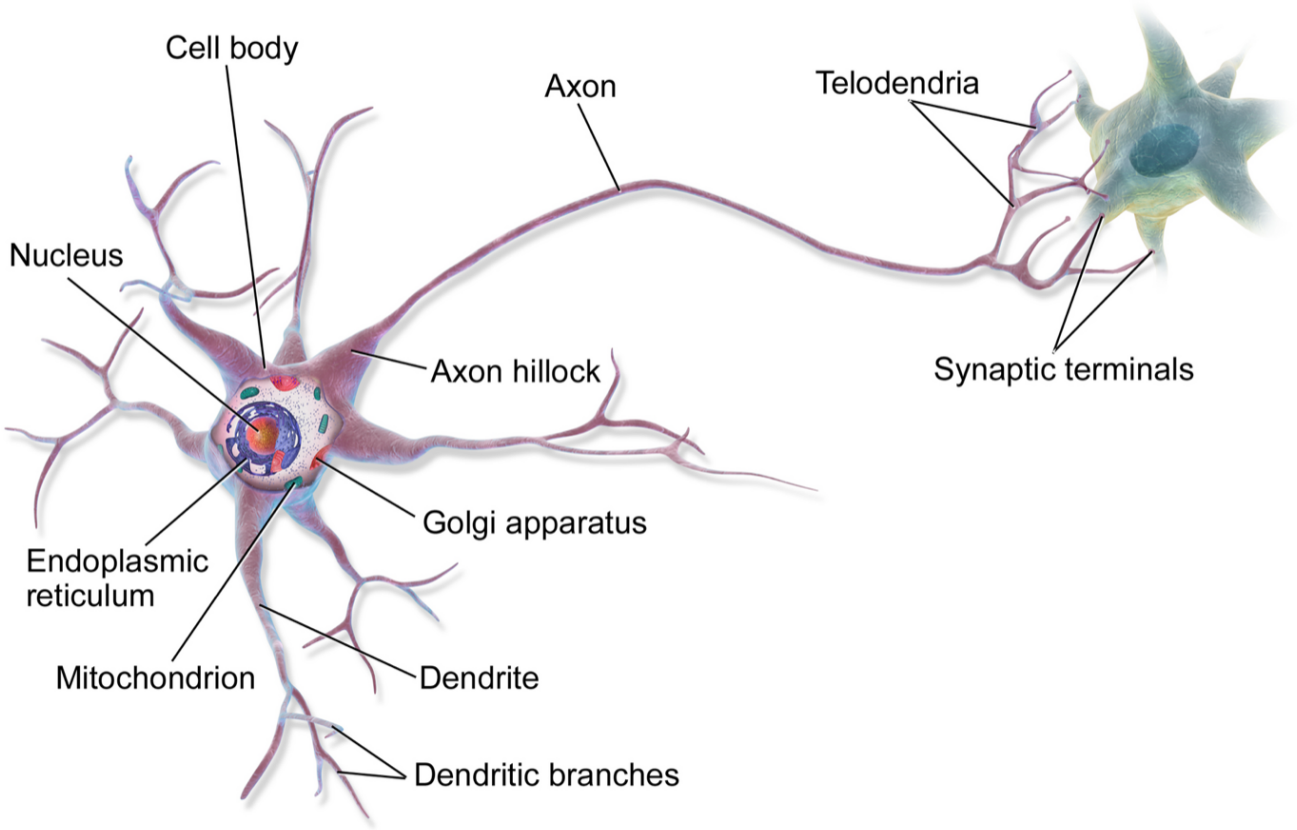


Figure 16. A multipolar neuron. Image: BruceBlaus, CC BY 3.0. Credits: Appendix H

- **Dendrites** receive inputs.
- The **soma** integrates them: does the summed input reach threshold?
- The **axon** conducts the resulting signal, sometimes over a metre.
- **Myelin** — produced by Schwann cells in the periphery, oligodendrocytes in the CNS — insulates the axon in segments. The signal regenerates only at the gaps (**nodes of Ranvier**), so it effectively jumps: **saltatory conduction**, which is far faster and more metabolically efficient than continuous conduction.
- **Axon terminals** transmit to the next cell across a **synapse**.

Conduction velocity depends on axon diameter and myelination, spanning roughly 0.5 m/s for small unmyelinated fibres to around 120 m/s for large myelinated ones. Demyelinating conditions slow or block conduction without destroying the axon, which is why their effects can be profound and variable.

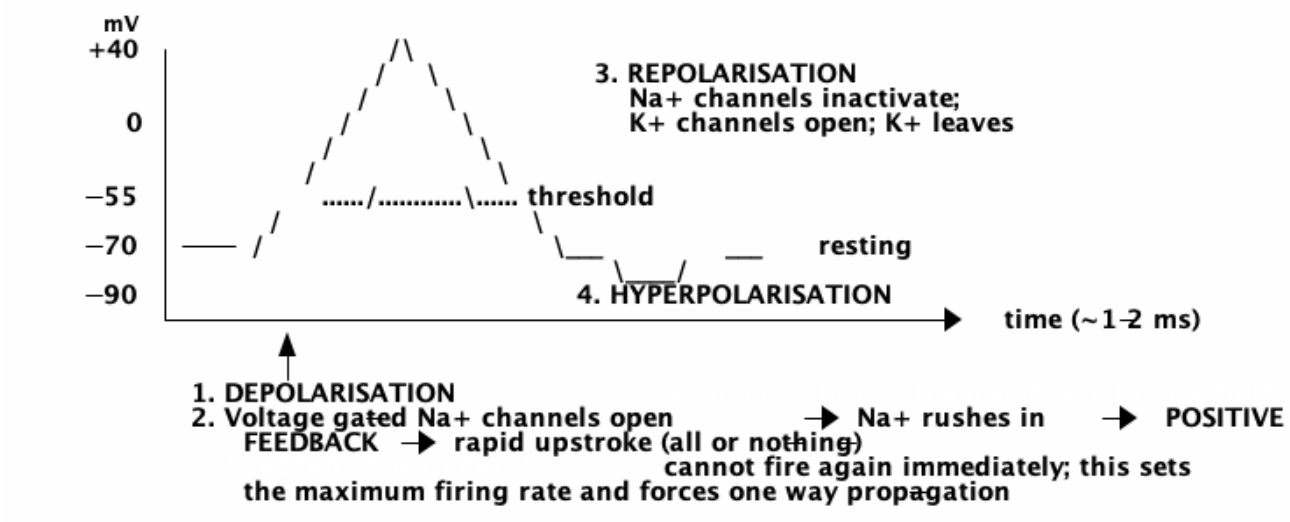
7.3.2. The resting membrane potential

A neuron at rest holds roughly −70 mV inside relative to outside. Three contributors:

1. **Ion gradients** maintained by the Na^{+}/K^{+}-ATPase, which pumps 3 Na^{+} out for every 2 K^{+} in, consuming ATP. High K^{+} inside, high Na^{+} outside.
2. **Selective permeability** — at rest the membrane is far more permeable to K^{+} than to Na^{+}, so the potential sits near the potassium equilibrium potential.
3. **Impermeant anions** inside the cell.

This is a battery maintained by continuous metabolic work. Cut the energy supply and the gradients decay. It is also why electrolyte disturbances — particularly potassium — have immediate effects on nerve and cardiac function.

7.3.3. The action potential



Two properties with engineering consequences:

- **All-or-nothing.** A stimulus either reaches threshold and produces a full action potential, or it does not. The signal is effectively digital in amplitude.
- **Intensity is coded by frequency and by population.** A stronger stimulus does not produce a bigger action potential; it produces **more of them per second**, and recruits **more fibres**. A nervous system is a frequency- and population-coded system, not an amplitude-coded one.

That second point is genuinely useful. If you are ever asked why a nerve-stimulation device controls pulse rate and pulse width rather than "voltage", this is why.

7.3.4. Synapses and the neuromuscular junction

At a chemical synapse, the arriving action potential triggers calcium entry, which causes vesicles to release **neurotransmitter** into a narrow cleft. The transmitter binds receptors on the next cell, which may be excitatory or inhibitory. The receiving cell sums many such inputs in space and time, and fires only if the total reaches threshold.

The **neuromuscular junction** is the specialised synapse between a motor neuron and a skeletal muscle fibre. Its transmitter is **acetylcholine**. This is where Chapter 6' muscle physiology connects: nerve signal → acetylcholine → muscle fibre depolarisation → calcium release inside the fibre → contraction.

Many drugs and toxins act here, which is why the NMJ appears constantly in pharmacology and anaesthesia.

7.4. Sensation

This section covers the **general** senses — touch, pressure, temperature, pain, and body position — whose receptors are spread through the skin, muscles, and joints. The **special** senses — vision, hearing, balance, smell, and taste — have dedicated organs in the head and are the subject of Chapter 8.

7.4.1. The somatosensory receptors

Receptor	Senses	Adaptation	Location
Merkel cell–neurite complex	Sustained light touch, edges, texture, form	Slow	Basal epidermis; dense in fingertips
Meissner (tactile) corpuscle	Light touch, low-frequency flutter, slip detection	Rapid	Papillary dermis of glabrous (hairless) skin
Pacinian (lamellated) corpuscle	High-frequency vibration, deep pressure	Very rapid	Deep dermis and subcutaneous tissue
Ruffini ending	Skin stretch, sustained pressure, finger position	Slow	Dermis and joint capsules

Receptor	Senses	Adaptation	Location
Hair follicle receptor	Hair deflection	Rapid	Around follicles
Free nerve endings (thermoreceptors)	Warm and cold, separately	Varies	Throughout skin
Free nerve endings (nociceptors)	Potentially damaging mechanical, thermal, chemical stimuli	Slow; can **sensitise**	Nearly everywhere except the brain parenchyma
Muscle spindle	Muscle length and rate of change	Slow	Within skeletal muscle
Golgi tendon organ	Muscle tension	Slow	At the muscle–tendon junction
Joint receptors	Joint position and movement	Varies	Capsules and ligaments

Two general lessons:

Rapidly adapting receptors stop reporting a constant stimulus. You stop noticing your clothing within seconds. This is efficient — it frees attention for change — and it is a design hazard: a sustained pressure you **can** feel at minute one you may not notice at minute thirty. Rapid adaptation is part of why pressure injury develops without warning in people whose sensation is otherwise intact.

Receptor density varies enormously by site. Fingertips and lips are densely innervated; the back is sparse. This is measurable as **two-point discrimination** — the minimum separation at which two touches feel like two:

Site	Approximate two-point discrimination threshold
Fingertip	~2–4 mm
Palm	~10 mm
Forearm	~30–40 mm
Back	~40–50 mm or more

Figures are typical orientation values and vary with method, age, and individual.

The design consequence is direct: **tactile markers must be separated by considerably more than the two-point threshold at the site where they will be felt.** Two raised dots 3 mm apart on a handle gripped with the palm are one dot. On a fingertip they are two.

7.4.2. Proprioception

Proprioception is the sense of body position and movement, built mainly from muscle spindles (length), Golgi tendon organs (tension), and joint receptors, integrated centrally.

It is the sense people forget they have, and the one whose loss is most disorienting. Without it you cannot find your own hand in the dark, cannot grade grip force without watching, and cannot maintain a posture without visual attention. Proprioceptive loss makes tasks that look trivial consume the whole of someone's attention.

Design consequence: an interface that can be operated **without looking** — by feel, with distinguishable shapes and clear mechanical detents — places far less load on a person whose vision is occupied, impaired, or needed elsewhere. An interface that **requires** visual confirmation of every action is expensive for everyone and may be unusable for some.

7.4.3. Pain, and why it is not a sensor reading

Nociception is the detection and neural signalling of potentially damaging stimuli. **Pain** is the experience. They are related but not identical, and the gap between them is clinically enormous: pain can occur without ongoing tissue damage, and tissue damage can occur with little pain.

Nociceptors can **sensitise**: after injury or inflammation, the threshold falls and the response grows, so a normally innocuous touch becomes painful. That is why "it only presses lightly" is not reassurance near inflamed tissue.



Absence of a pain report is not evidence of absence of harm.

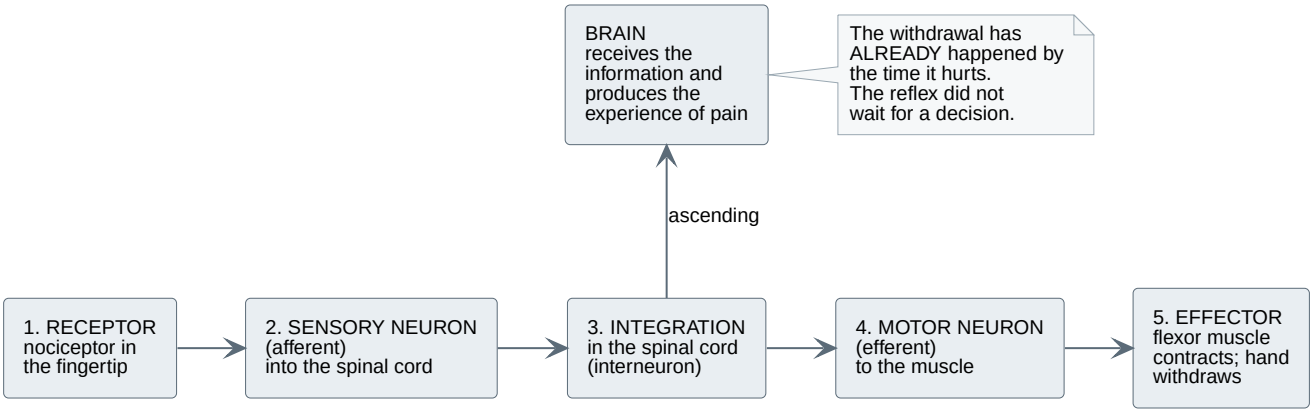
The reasons stack up:

- sensation may be reduced or absent in that region (neuropathy, spinal cord injury, post-surgical, diabetic, after local anaesthesia);
- rapidly adapting receptors stop reporting sustained stimuli;
- damage from sustained loading begins deep, where there is less innervation than at the surface (Chapter 5);
- the person may be unable to communicate, or may be reluctant to complain to someone who has gone to trouble on their behalf;
- attention elsewhere suppresses awareness.

A design must not use the user's discomfort as its safety mechanism. If protection depends on the person noticing a problem, and that person may not notice, the protection does not exist.

7.4.4. Reflexes

A **reflex** is a rapid, involuntary, stereotyped response. The simplest form is a spinal loop that does not wait for a conscious decision.



The withdrawal reflex happens before the pain is experienced. That is a real and instructive ordering: **protective responses can be faster than awareness**.

Which cuts both ways for a designer. A startle or withdrawal response can cause someone to drop, jerk, or pull away — so your mechanism should tolerate a sudden release. And reflex responses are not available for hazards that do not trigger them: sustained pressure produces no withdrawal reflex at all.

7.5. Engineering bridge: biosignals

If you move toward instrumentation, this table is the orientation you need. These are the signals a biomedical instrumentation course begins with.

Signal	Typical amplitude	Typical bandwidth	Source
ECG (electrocardiogram)	~0.5–4 mV at the skin	~0.05–150 Hz (diagnostic)	Cardiac muscle depolarisation
EMG (electromyogram), surface	~0.05–5 mV	~10–500 Hz	Skeletal muscle fibre action potentials
EEG (electroencephalogram)	~10–100 µV	~0.5–100 Hz	Summed cortical activity
EOG (electrooculogram)	~50–3500 µV	~0–30 Hz	Corneo-retinal standing potential as the eye moves
Nerve action potential, surface	~1–20 µV	~100–2000 Hz	Peripheral nerve conduction studies
PPG (photoplethysmogram)	Optical, not electrical	~0.5–20 Hz	Pulsatile blood volume change; basis of pulse oximetry (Chapter 9)



Why this table is the whole of introductory biomedical instrumentation.

Look at the amplitudes. An EEG is tens of microvolts. Mains interference is volts. The signal is five to six orders of magnitude below the ambient noise you are trying to reject, and the "wire" is a gel-coupled electrode on skin whose impedance drifts with moisture, movement, and time.

Everything in the field follows from that: high-input-impedance instrumentation amplifiers, high common-mode rejection ratio, driven-shield and right-leg-drive techniques, careful grounding, motion-artefact rejection, and — crucially — **patient electrical safety**, because you are deliberately making a low-impedance electrical connection to a person.



Never connect anything you have built to a person. Patient-connected electrical equipment is governed by IEC 60601-1 and its particular standards (Appendix D), which specify leakage current limits, isolation, and applied-part classification. Mains-powered bench equipment, development boards, and anything with a USB connection to a laptop are not isolated to those limits.

If a project needs biosignal measurement, use clinically approved instrumentation, operated by someone qualified, or work with recorded or synthetic data. This is an absolute boundary, not a cautious preference.

7.5.1. Psychophysics: the numbers behind "noticeable"

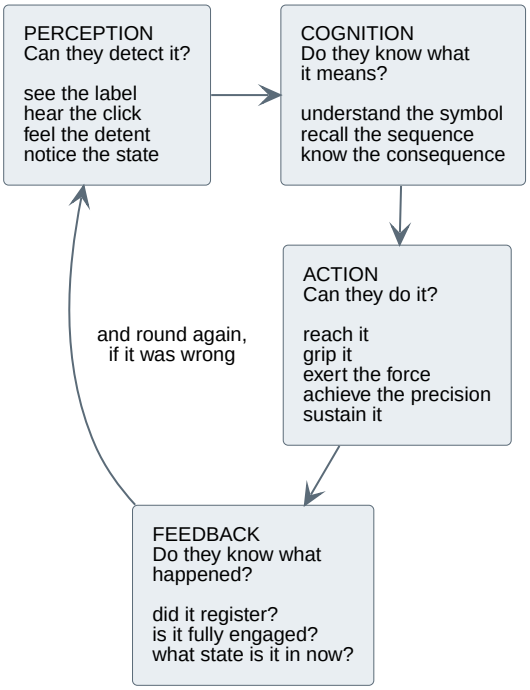
Two ideas worth having:

Just-noticeable difference (JND) — the smallest change in a stimulus that a person reliably detects. For many sensory modalities it is roughly proportional to the baseline magnitude (**Weber's law**): the louder the room, the bigger the change in sound needed to notice it. Consequence: a feedback cue sized for a quiet workshop may be undetectable in a busy clinical area.

Reaction time — a simple visual reaction is typically around 200 ms, auditory slightly faster, and anything requiring a choice or a decision is substantially longer. Fatigue, age, medication, and divided attention all lengthen it. If your interface has a timeout, that timeout is a design parameter with real human limits behind it, and the default in your framework was not chosen with any of them in mind.

7.6. Decomposing a task: perception, cognition, action, feedback

This is the method the chapter exists to deliver. Any interaction with any object decomposes into at least four independent demands.



Work the example. A control requires someone to **see a small symbol**, **pinch a latch**, **hold it open**, and **slide a part into place**. "Can use the control" conceals at least seven demands:

Stage	Demand	Could fail because
Perception	See and resolve a small symbol	Visual acuity, contrast, lighting, size, viewing angle, glare
Cognition	Know what the symbol means	Unfamiliar convention, no label, language, memory load
Cognition	Know the sequence	Order is not discoverable from the object
Action	Reach the latch	Range of motion, posture, obstruction, position of the object
Action	Achieve the pinch	**Pinch precision** and **pinch force** are different capabilities
Action	Sustain the hold	Endurance — an isometric demand (Chapter 6)
Action	Slide the part accurately	Bimanual coordination, tremor, precision
Feedback	Know it is fully engaged	Is there any indication at all? Is it visible, audible, tactile?

Eight failure modes behind four words. Now you can design: each row is a separate intervention, and some are cheap.



The same decomposition works for software. Can the user perceive the control (contrast, size, focus indicator, screen-reader label)? Do they understand it (wording, icon, affordance)? Can they act (target size, pointer precision, keyboard access, timeout)? Do they get feedback (state change, confirmation, error message that says what to do)?

This is not a coincidence. It is why WCAG [22] is organised the way it is, and why its four principles — **perceivable, operable, understandable, robust** — map almost exactly onto this table.

7.7. Designing for variation, not for a diagnosis

A diagnostic label is not a specification.

Two people with the same diagnosis may differ in strength, range, sensation, endurance, tremor, vision, cognition, priorities, living situation, support, and what they actually want to do. Designing from a diagnosis produces a device for a category, and categories do not use devices.

Ask instead:

- What are you trying to do?
- How do you do it now?
- What specifically becomes difficult, and at what point?
- What have you already tried, and why did it not work?
- What would a useful improvement let you do that you cannot do now?

A therapist helps interpret movement, sensation, and functional capability. **The person is the expert on their own experience, priorities, and context.** Both inputs are needed and they are not interchangeable.

7.7.1. Accessible feedback: redundancy, not substitution

Single-channel cue	Fails for	Add
Colour only (red/green)	Colour vision deficiency — around 8% of men and 0.5% of women of northern European descent — plus poor lighting and monochrome printing	Shape, position, text, pattern, icon
Sound only (a click or beep)	Hearing loss; a noisy clinical area; a user wearing ear protection; someone across the room	A visible state change, an alignment mark, a mechanical detent
Text only (small print)	Low vision, dyslexia, unfamiliar language, poor lighting, wet or gloved hands smearing a label	Larger type, high contrast, a tactile feature, a symbol, a shape difference
Tactile only	Reduced sensation; gloved hands; prosthesis	Visible indication, audible confirmation

Single-channel cue	Fails for	Add
Position only (which side the switch is)	Not memorable; symmetrical parts get assembled wrongly	Asymmetric geometry so the wrong assembly is physically impossible

The principle is **redundant coding**: provide the same information through more than one channel, so no single impairment removes it. Note that this is also better for **everyone** — in a noisy room, in poor light, while distracted, while learning.



The strongest control is one where the wrong thing cannot be done. An asymmetric connector that only fits one way beats a label saying "this way up". In risk terms (Chapter 19) this is **design to eliminate** rather than **protective measure** or **information for safety** — and it sits at the top of the control hierarchy for good reason.

7.8. Devices, software, and training technology

Neurology and rehabilitation carry the densest concentration of assistive technology in medicine, and the widest spread of ways to contribute. Much of accessibility work is software, and most of it is not regulated as a medical device. Much of the rest is hardware: switches, interfaces, mounts, and controls that have to be individually fitted, robust, and cleanable — electronics, machining, and textiles as much as code.

7.8.1. Clinical devices

Device	Does	Engineering note
EEG system	Records cortical electrical activity from scalp electrodes	~10–100 µV signals. Electrode impedance, artefact rejection, and the 10–20 placement system. See Chapter 7' amplitude table
EMG and nerve conduction studies	Records muscle electrical activity and measures nerve conduction velocity	Surface and needle electrodes; stimulation and latency measurement
Evoked potential systems	Measures response to a repeated stimulus, averaged out of the noise	Signal averaging is the whole trick: the response is far below the noise floor in a single trial
TENS — transcutaneous electrical nerve stimulation	Pain management by surface electrical stimulation	Pulse rate, width, and amplitude. Widely available, low risk, well understood
Functional electrical stimulation (FES)	Stimulates paralysed muscle to produce useful movement — e.g. foot-drop stimulators	Timing to the gait cycle; electrode placement; fatigue

Device	Does	Engineering note
Deep brain stimulation (DBS)	Implanted electrodes modulating specific brain targets, for Parkinson's and others	Implanted pulse generator, lead migration, programming. Highly specialised
Cochlear implant	Bypasses damaged hair cells, stimulating the auditory nerve directly via an electrode array	Speech processing strategy, electrode array design, and fitting ("mapping")
Hearing aids	Amplify and process sound to match a hearing loss profile	Compression, feedback cancellation, directional microphones. Increasingly over-the-counter
Brain–computer interfaces	Control from neural signals, non-invasive or implanted	Research and early clinical. Signal quality, decoding, and training burden
AAC devices — augmentative and alternative communication	Speech output for people who cannot speak	From simple symbol boards to eye-gaze-driven speech generation. **Access method is the core design problem**
Switch and alternative access hardware	Single switches, sip-and-puff, head arrays, eye gaze, joysticks	**Highly individual.** This is where maker and OT collaboration is most productive: switch electronics and adapters, machined clamps and articulated mounting arms, moulded or turned switch caps, and fabric or padded covers
Environmental control systems	Operate lights, doors, beds, and appliances from an accessible interface	Increasingly built on consumer smart-home hardware
Intrathecal pumps and nerve blocks	Targeted drug delivery for spasticity and pain	Implanted, programmable, narrow therapeutic index (Chapter 3)

7.8.2. Software

Software	Does	Note
Screen readers — NVDA, JAWS, VoiceOver, TalkBack	Non-visual access to computing	**NVDA is free and open source.** Learning to navigate with your screen off for an hour is the single most educational accessibility exercise available
Switch access and scanning software	Operate a whole interface from one or two switches	Scanning rate, selection method, and prediction. Built into both major mobile platforms

Software	Does	Note
Eye-gaze control software	Pointer and keyboard control from gaze	Calibration, dwell time, and fatigue. Hardware is now comparatively affordable
AAC applications — Proloquo2Go, Grid, CoughDrop, open symbol sets	Symbol- and text-based speech generation	**CoughDrop and the Open Symbols set are open.** Vocabulary organisation matters more than the technology
Speech recognition and dictation	Hands-free text entry and control	Accuracy with dysarthric speech is a known gap and an active research area
Accessibility APIs and test tools — axe, WAVE, Lighthouse, Accessibility Inspector	Find and fix interface barriers	**Directly usable by you today.** Automated tools catch perhaps a third of issues; manual testing and real users catch the rest
WCAG and ARIA	The standards your interface should meet	See Chapter 7' decomposition — the four WCAG principles map onto it almost exactly
EEG and signal toolkits — MNE-Python, EEGLAB, OpenBCI software	Analysis of biosignal recordings	**Free, and usable on public datasets.** The safe way to learn biosignal processing
Cognitive and neuro-rehabilitation apps	Structured exercises after stroke or brain injury	Evidence quality varies widely; clinical guidance required

7.8.3. Training and simulation

Training technology	Teaches	Form
Lumbar puncture trainers	Landmark identification, needle angle, the "give" on entering the space	Layered tissue-mimicking torso section with simulated CSF. Tactile fidelity is the whole point (Chapter 21)
Nerve block phantoms	Ultrasound-guided needle placement near simulated nerves	Gel phantoms with embedded structures. **A maker-accessible build** (Chapter 15)
Neurological examination simulators	Reflexes, tone, power, sensation testing	Standardised patients and some instrumented manikins

Training technology	Teaches	Form
Stroke scale training — NIHSS and similar	Structured neurological assessment	Video-based certification courses, widely used
Reflex hammer and sensory testing kits	Clinical examination technique	Monofilaments, tuning forks, two-point discriminators. **Simple, cheap, makeable** — turned and machined metal, calibrated nylon filaments, laser-cut or engraved gauges — and directly connected to this chapter's psychophysics
Disability simulation equipment	Weighted suits, vision-simulating goggles, restricted gloves	**Use with care.** Useful for finding specific interface barriers; widely criticised as a way of "understanding" disability, for the reason in this chapter
EEG and EMG demonstration kits — OpenBCI, Backyard Brains	Biosignal acquisition and processing concepts	Education-grade hardware. **Not clinical**, and never to be presented as such
Tactile and feedback sample boards	How texture, detent, shape, and spacing are perceived	Machined, knurled, engraved, stitched, and moulded samples on one board, plus a few vibration motors and click mechanisms. A direct, hands-on version of redundant coding



The electrical boundary from Chapter 7 applies in full here.

Education-grade biosignal kits exist and are genuinely useful for learning signal processing. They are not clinical instruments, they are not patient-connected equipment under IEC 60601-1, and nothing you build from one may be used to make or inform a clinical judgement about anybody.

Use them on yourself for learning if the manufacturer's instructions permit, use public datasets for anything serious, and keep the two uses entirely separate in your records.



Accessibility is the highest-leverage contribution in this book, in software and in hardware.

On the software side it needs no workshop, no materials, no infection control review, and no patient contact. The tools are free. The standards are public. The need is enormous and chronically under-served, and most clinical and educational software is measurably bad at it. Run `axe` over something. Navigate a site with a screen reader.

> On the hardware side, the work is alternative access: a switch interface that turns a large, low-force button into a keypress; a mounting arm that holds it exactly where one person can reach it; a robust, wipeable enclosure; a cover or strap that keeps it in place. Each is individual, each is made with an OT and the user, and each draws on electronics, machining, and textiles rather than any one process.
>
> Then read Chapter 7' perception–cognition–action–feedback table again and map what you found onto it.

7.9. Worked example: two identical cord ends

Your elbow model has two cord ends. You know which one bends it, because you built it. A learner does not.

First attempt: colour one red and one green. Better than nothing, and insufficient — it fails for colour vision deficiency, in poor light, and in a monochrome photocopy of your instructions.

Second attempt: redundant coding.

- **Text:** one end labelled "PULL TO BEND", the other "SUPPORT — DO NOT PULL".
- **Shape and texture:** different tab geometry on each end — a turned ball on one and a flat knurled or stitched tab on the other — distinguishable by touch alone.
- **Symbol:** a curved arrow showing the direction of the resulting motion.
- **Position:** the active end always on the same side, with a mark on the base.
- **Relationship:** a line drawn on the model from the pull point to the moving segment, so the causal relationship is visible rather than remembered.
- **Colour:** keep it — as a **redundant** cue, not the only one.

Then test it. Hand the model to someone who has not seen it and ask them to explain which cord does what and why, **without hints**. If they cannot, your labelling failed — not their reading. Record whether you prompted them; see Chapter 21.

Second case from the same model: a **click** when a part locks. Audible feedback is useful and cheap. It also fails in a busy room, for a user with hearing loss, and when the educator is facing away. Add a **visible alignment mark** that only lines up when engaged, or a **mechanical stop** the user can feel. Whether either is sufficient depends on the actual task, the actual environment, and the actual users — which means it is a question to test, not to settle by reasoning.

7.10. Workshop exercise

1. **Decompose a task.** Take "press this touchscreen button" and write out every demand under perception, cognition, action, and feedback. Aim for at least eight entries. Then do the same for a physical control on something in your workshop.
2. **Barrier review.** Take a mock software screen or a tabletop control panel — real or drawn. Identify **five** potential barriers, each tied to a specific stage of the decomposition. Redesign it and state, for each change, what it improves **and what it costs**.

3. **Redundant coding.** Apply the two-cord-end treatment to your elbow model. Then test it on someone unfamiliar with it and record the result including any prompts you gave.
4. **Two-point discrimination.** With a willing adult volunteer and two blunt points — a paperclip bent into a U is fine — find the approximate separation at which two touches feel like two on a fingertip, then on the forearm. Light contact only, eyes closed, stop if the person asks. The point is the **ratio** between sites, not an exact number.
5. **Open questions.** Write three open questions for a prospective user. Check each one: does it invite information, or does it fish for approval of a design you have already decided on?



About one-handed trials. Operating your own prototype with one hand, or with heavy gloves, or in dim light, is a legitimate way to **find** problems — it will reveal a two-handed assembly requirement in seconds.

It is not a way to **understand** disability, and it does not establish that your redesign meets anyone's needs. The difference is between generating a hypothesis and confirming one. Keep that distinction explicitly in your notes, because conflating the two is how well-meaning projects go wrong.

7.11. Questions

1. Split "use this touchscreen button" into perception, cognition, action, and feedback.
2. Why is a stronger stimulus not signalled by a bigger action potential, and how **is** it signalled?
3. Give one reason absence of discomfort is weak evidence when sensation is reduced — and one reason it is weak even when sensation is intact.
4. What is proprioception, and name a design consequence of its loss.
5. Why must tactile markers be spaced further apart on a palm than on a fingertip?
6. A control gives feedback by a click. Name two situations where that fails and one alternative for each.
7. Why is colour-only coding insufficient, and what is the general principle that fixes it?
8. What is the difference between nociception and pain?
9. An EEG signal is ~50 µV. Why does that make instrumentation hard, and what must you never do with something you built?
10. A user has the same diagnosis as a previous user you designed for. What may you reuse?
11. Rewrite the question "Would this handle help you?" as something that actually produces information.
12. The withdrawal reflex happens before pain is felt. Give one design implication.

7.12. Answers and reasoning

1. **Perception:** locate the button, resolve it against the background, see its focus or pressed state. **Cognition:** know what it does, know the consequence, know whether it is the right one of several. **Action:** reach it, target it accurately with a finger or pointer, apply it within any time

limit, do it with a tremor or gloved hand. **Feedback:** did it register, what state is the system in now, if it failed does the message say what to do. A strong answer notes that each of these can fail independently.

2. Action potentials are **all-or-nothing** — amplitude is fixed by the ion gradients and channel properties, not by stimulus strength. Intensity is coded by **firing frequency** and by **how many fibres are recruited**.
3. **With reduced sensation:** the person may be entirely unable to detect a harmful contact condition, so the report you are relying on is not available. **Even with intact sensation:** rapidly adapting receptors stop reporting a sustained stimulus; damage from loading begins deep where innervation is sparser; attention elsewhere suppresses awareness; and the person may be reluctant to complain.
4. The sense of body position and movement, from muscle spindles, Golgi tendon organs, and joint receptors. With its loss, a person must **watch** to know where a limb is and to grade force, so every task consumes visual attention. Design consequence: interfaces operable by feel, with distinguishable shapes and clear detents, place far less load on such a user — and on everyone whose eyes are needed elsewhere.
5. Because two-point discrimination threshold is much larger on the palm (~10 mm) than on the fingertip (~2–4 mm) due to lower receptor density and larger receptive fields. Two markers closer than the threshold at the site of contact are perceived as one.
6. Fails in a noisy environment; fails for a user with hearing loss; fails if the operator is facing away or wearing ear protection. Alternatives: a visible alignment mark that lines up only when engaged; a palpable mechanical detent; a state that is visibly different.
7. Because colour vision deficiency is common (roughly 8% of men of northern European descent), and because colour is also degraded by poor lighting, glare, monochrome reproduction, and screen settings. The principle is **redundant coding**: carry the same information on more than one channel — shape, position, text, pattern, texture — so no single impairment removes it.
8. **Nociception** is the neural detection and signalling of potentially damaging stimuli. **Pain** is the conscious experience. They dissociate in both directions: pain without ongoing damage, and damage without pain.
9. It is five to six orders of magnitude smaller than ambient mains interference, acquired through a drifting skin–electrode interface, so it demands high input impedance, high common-mode rejection, careful shielding and grounding, and motion-artefact handling. **You must never connect anything you built to a person** — patient-connected equipment is governed by IEC 60601-1 leakage and isolation requirements that bench and USB-powered equipment do not meet.
10. **Almost nothing about the person.** You may reuse your process, your questions, your documentation templates, and possibly some design concepts as **candidates**. You may not assume the same capabilities, the same priorities, the same environment, or the same solution. The diagnosis is not the specification.
11. Something that asks about the task rather than the artefact: **"Tell me about the last time you tried to do this — what happened?"**, **"Which part of it takes the most effort?"**, **"What would a useful improvement let you do?"** The original question invites politeness, and politeness is not data.

12. Any of: a mechanism must tolerate a sudden jerk or release, because withdrawal happens before deliberate decision; a fast protective response may fire before the user "decides", so a design should not depend on a considered reaction; and conversely, **no** withdrawal reflex exists for sustained pressure, so reflexes provide no protection against the slow hazards.

7.13. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Biosignal analysis — filtering a recording, beat detection, heart-rate variability and EMG onsets.
- FFT and synthetic biosignals — synthetic ECG, EMG and EEG with mains hum, and signal averaging.
- Active filters — notch and band-pass filter design.
- ODE solver — a FitzHugh–Nagumo action potential.
- Labels, alarms and the senses — two-point spacing for tactile markers.

7.14. Video companion and reading

- **Core videos:** Crash Course Anatomy & Physiology #8 (the nervous system) [10], then #12 on the peripheral nervous system from the series index [6]. The episodes on vision, hearing, and balance belong to Chapter 8.
- **Core reading:** OpenStax [1] chapters 12–14 selectively:
  - 12.1 Basic Structure and Function of the Nervous System
  - 12.4 The Action Potential and 12.5 Communication Between Neurons
  - 13.4 Sensory and Motor Pathways
  - 14.2 Central Processing — for the sensory receptors and submodalities
- **Collaborative design:** the Tikkun Olam Makers product development approach [19] — a worked example of designing **with** the person who has the unmet need rather than for a category.
- **Accessibility:** the Web Content Accessibility Guidelines [22]. Read the four principles and the structure even if your project is physical; the thinking transfers.

7.15. Exit criteria

- ▢ Trace a simple sensory input and motor response through a labelled diagram, naming receptor, afferent, integration, efferent, effector.
- ▢ Explain why a person may need protection from a hazard they cannot feel, giving at least three distinct mechanisms.
- ▢ Break one real task into perception, cognition, action, and feedback with at least eight entries.
- ▢ Identify five accessibility questions for your prototype, each tied to a stage of the decomposition, and document two design changes with their costs.

- ▢ Apply redundant coding to at least one cue on your model and test it on someone unfamiliar with it.
- ▢ Separate observed user feedback from your own assumptions in your design notes, explicitly labelled.
- ▢ State the amplitude range of ECG, EMG, and EEG, and explain why nothing you build may be connected to a person.

7.16. What this chapter does not establish

- **You cannot assess anyone's sensation, strength, or cognition.** You know these are distinct capabilities. Measuring them is clinical work.
- **You cannot build patient-connected electronics.** The biosignal table is orientation for a future instrumentation course, not a licence. IEC 60601-1 exists because the failure mode is electrocution.
- **A one-handed or gloved trial is a hypothesis generator, not validation.** It reveals problems; it does not reproduce lived experience or establish that a fix works.
- **Accessibility is not a checklist you can complete alone.** WCAG conformance is testable; whether a thing is usable by a particular person is established **with** that person.
- **You have not learned neurology.** Tremor, spasticity, neuropathy, and the rest appear here as reasons for variation. Their assessment and management are specialist.

Chapter 8. The head and neck: eye, ear, nose, throat, and mouth

Week 11 · approximately 4–6 hours

8.1. Why this chapter exists

Chapter 7 treated seeing and hearing as **capabilities** — things a user may or may not have, which a design must not depend on alone. It never opened the organs that provide them. This chapter does.

The head and neck are the densest region of the body for a designer. In a volume smaller than a shoebox sit four of the five special senses, the only route for air into the lungs, the only route for food into the gut, the crossing point where those two routes share a passage, the voice, the teeth, and the blood supply to the brain. Almost every structure is close to the surface. Almost nothing is padded.

You will meet this region in three ways:

- **As a place devices sit.** Glasses, goggles, headsets, hearing aids, oxygen tubing, masks, straps, lanyards. Each presses on thin tissue over bone or cartilage, and each can cause the device-related pressure injury of Chapter 5 in a place nobody thought to look.
- **As the reason an interface works or fails.** Every visual and auditory cue in Chapter 7' redundant-coding table depends on an eye or an ear with known, measurable limits. A label is an optics problem. A beep is an acoustics problem.
- **As simulator and training territory.** Airway heads, ear and eye examination trainers, and surgical airway trainers are among the most consumed — and most maker-accessible — task trainers in medical education.

And one rule that belongs at the start rather than the end: **the airway and the food passage cross.** Anything that goes into the mouth is, until shown otherwise, something that can go into the lungs.

8.2. The special senses, and the nerves that carry them

The **general senses** of Chapter 7 — touch, pressure, temperature, pain, position — use receptors distributed through the body. The **special senses** use dedicated organs, all in the head, wired to the brain directly by **cranial nerves** rather than through the spinal cord.

Sense	Organ	Receptor and stimulus	Cranial nerve
Vision	Eye	Photoreceptors (rods and cones) — light, ~400–700 nm	II (optic)
Hearing	Cochlea (inner ear)	Hair cells — fluid vibration driven by sound	VIII (vestibulocochlear)

Sense	Organ	Receptor and stimulus	Cranial nerve
Balance	Vestibular apparatus (inner ear)	Hair cells — head rotation, linear acceleration, gravity	VIII
Smell (olfaction)	Olfactory epithelium (roof of nasal cavity)	Chemoreceptor neurons — airborne molecules	I (olfactory)
Taste (gustation)	Taste buds (tongue, palate, pharynx)	Chemoreceptor cells — dissolved molecules	VII, IX, X

You do not need all twelve cranial nerves by heart. You do need to know that they exist, that they leave the skull through small openings, and roughly what each does — because clinical conversation about this region is conducted in them.

Nerve	Name	What it does that a maker might care about
I	Olfactory	Smell. Commonly reduced with age and after viral illness
II	Optic	Vision
III, IV, VI	Oculomotor, trochlear, abducens	The six muscles that move each eye; lid elevation; pupil constriction
V	Trigeminal	Sensation of the face, cornea, nasal and oral lining; the chewing muscles
VII	Facial	Facial expression, **eye closure**, taste from the front of the tongue, tears and saliva
VIII	Vestibulocochlear	Hearing and balance
IX	Glossopharyngeal	Taste from the back of the tongue; sensation of the pharynx; the gag reflex (afferent)
X	Vagus	Larynx and pharynx muscles — voice and swallowing; autonomic control of heart and gut
XI	Accessory	Neck and shoulder muscles (sternocleidomastoid, trapezius)
XII	Hypoglossal	The tongue's muscles



The corneal blink reflex is a two-nerve loop worth remembering. Touch the cornea and both eyes close. Sensation in on V, command out on VII. A person with facial nerve weakness (a Bell's palsy, for example) may be unable to close one eye — so the cornea dries and is damaged **without any loss of sensation**. Same lesson as Chapter 7: protection depends on the whole loop, and any link can fail.

8.3. The eye

8.3.1. Structure

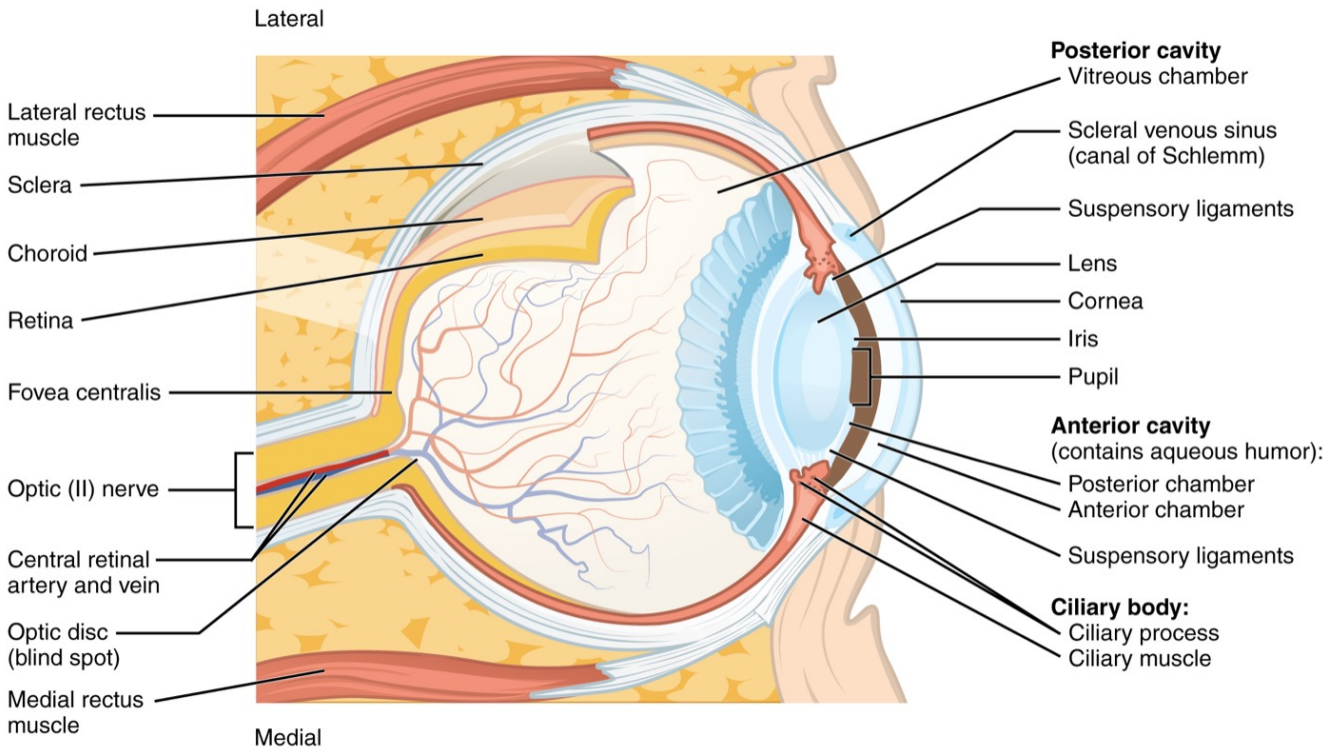


Figure 17. The eye in horizontal section. Image: OpenStax, CC BY 3.0. Credits: Appendix H

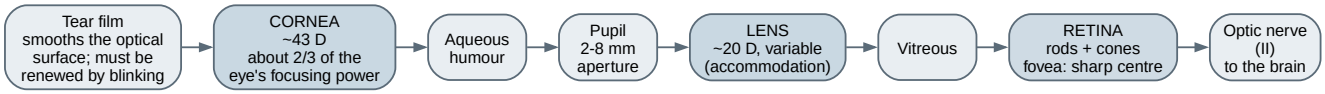
The eyeball is roughly a sphere about 24 mm in axial length, built in three layers:

- **Fibrous layer** — the tough white **sclera**, continuous at the front with the clear, dome-shaped **cornea**. The cornea is about 0.5 mm thick at its centre and 11–12 mm across. It has **no blood vessels** — it is fed by the tear film and the aqueous humour — and it is one of the most densely innervated tissues in the body.
- **Vascular layer (uvea)** — the **choroid**, which supplies the outer retina; the **ciliary body**, which makes aqueous humour and holds the lens on suspensory ligaments; and the **iris**, whose muscles set the **pupil** between roughly 2 and 8 mm.
- **Neural layer** — the **retina**, a sheet of photoreceptors and processing neurons lining the back of the eye, leaving through the **optic disc** as the optic nerve.

Two fluid compartments keep the shape: **aqueous humour** in front of the lens, constantly produced and drained through a meshwork at the angle between iris and cornea; and the gel-like **vitreous** behind it. The balance of aqueous production and drainage sets the **intraocular pressure**, typically around 10–21 mmHg. When drainage fails and pressure rises, the optic nerve is slowly damaged — **glaucoma** — usually with no symptoms until the loss is substantial.

Around the eyeball are the **accessory structures**: the eyelids, the **conjunctiva** (a thin mucous membrane over the white of the eye and inside the lids), the **lacrimal apparatus** that makes and drains tears, and six **extraocular muscles** per eye that aim it.

8.3.2. The light path



Two things about this path surprise most engineers.

The cornea does most of the focusing, not the lens. Refraction is strongest where the change in refractive index is largest, and that is the air–cornea interface. The cornea contributes roughly 43 of the eye's ~60 **dioptres**; the lens contributes about 20 and is the **adjustable** part. This is why reshaping the cornea (laser refractive surgery) changes vision so much, why a scratched or dry cornea blurs everything, and why vision under water is so poor — water removes most of the cornea's power.

Accommodation is a muscle relaxing a spring. The lens is elastic. To focus near, the ciliary muscle contracts, the ligaments slacken, and the lens rounds up to more power. The lens stiffens steadily with age, so the **range** of accommodation shrinks: roughly 10 D at 20, around 4–6 D at 40, and about 1 D by 60. That is **presbyopia**, and it is universal. The near point recedes from about 10 cm to about a metre, which is why reading glasses appear in the mid-forties and why small print on a device defeats a large share of its users.

8.3.3. Refractive error

Condition	What is wrong	Corrected with
Myopia (short sight)	Eye too long for its power; distant objects focus in front of the retina	Negative (diverging) lens
Hyperopia (long sight)	Eye too short; near objects especially are focused behind the retina, and the lens must work constantly to compensate	Positive (converging) lens
Astigmatism	Cornea curved differently in different meridians; lines at some angles blur	Cylindrical lens component
Presbyopia	Loss of accommodation with age	Reading addition, bifocals, varifocals

Uncorrected refractive error is the largest single cause of vision impairment worldwide [41] — and the cheapest to fix. A pair of spectacles is a medical device in most jurisdictions, and a remarkably effective one.

8.3.4. The retina

The retina has two receptor families:

- **Rods** — roughly 90–120 million. Very sensitive, so they work in dim light; one kind only, so they give no colour; densest away from the centre.
- **Cones** — roughly 6 million. Three types, sensitive to short, medium, and long wavelengths; **colour** is the brain's comparison between them. Concentrated in the **fovea**, a pit about 1.5 mm across at the centre of the **macula**.

Only the central couple of degrees of the visual field is seen sharply. Everything else is coarse, motion-sensitive, and colour-poor. You do not notice because the eyes move several times a second and the brain stitches the result. The design consequence: **something the user is not looking at is something they may not see**, however large the screen.

The optic disc has no receptors at all. Each eye has a **blind spot** about 15° to the side of the point of fixation, which you never notice because the other eye covers it and the brain fills in. You can find yours in the workshop exercise.

Adaptation runs at very different speeds in the two directions. Moving into bright light recovers in seconds. Moving into the dark takes 20–30 minutes for full rod sensitivity — a fact relevant to anyone designing displays for dimmed wards, darkened imaging rooms, or night work.

8.3.5. Where vision fails, and why the pattern matters

Low vision is not one condition, and the **pattern** of loss determines what helps.

Cause	Typical pattern	Design implication
Cataract — clouding of the lens	Blur, **glare**, reduced contrast, colours dulled and yellowed	Avoid glossy surfaces and glare; high contrast; larger text
Glaucoma — optic nerve damage	Loss **from the periphery inwards**; central vision may be preserved until late	Do not place critical information at the edges; do not rely on noticing an indicator "out of the corner of the eye"
Age-related macular degeneration	Loss **of the centre**; peripheral vision preserved	Magnification helps only partly; large, simple layouts; non-visual alternatives
Diabetic retinopathy	Patchy, variable loss; may fluctuate	Redundant coding (Chapter 7); never assume a stable capability
Colour vision deficiency	Reduced discrimination, usually red–green	Chapter 7: colour must never be the only channel



Contrast sensitivity falls with age even when acuity is "normal". A 75-year-old who reads the bottom of the chart may still struggle with grey text on a grey background in a dim room. Acuity charts measure high-contrast detail. Interfaces mostly fail on low contrast.

8.3.6. How the eye protects itself

The eye is defended by the bony orbit, the lids, the blink, and the tear film. A resting blink rate is roughly 15–20 per minute, and it falls sharply during concentrated visual work — reading, screen use, driving a simulator — which is why prolonged screen tasks and headsets produce dry, sore eyes. The cornea's dense innervation makes even a small foreign body intensely painful: an alarm system, not a design margin.

Light itself is a hazard. Intense or prolonged exposure to ultraviolet, blue light, or a laser can

damage the cornea, lens, or retina, and the retina has no pain receptors to warn you. This is why the UV-A hand hygiene box of Chapter 13 must shield the source from the eyes, and why lamps and lasers have their own photobiological safety standards (Appendix D).

8.4. The ear

8.4.1. Three parts, three media

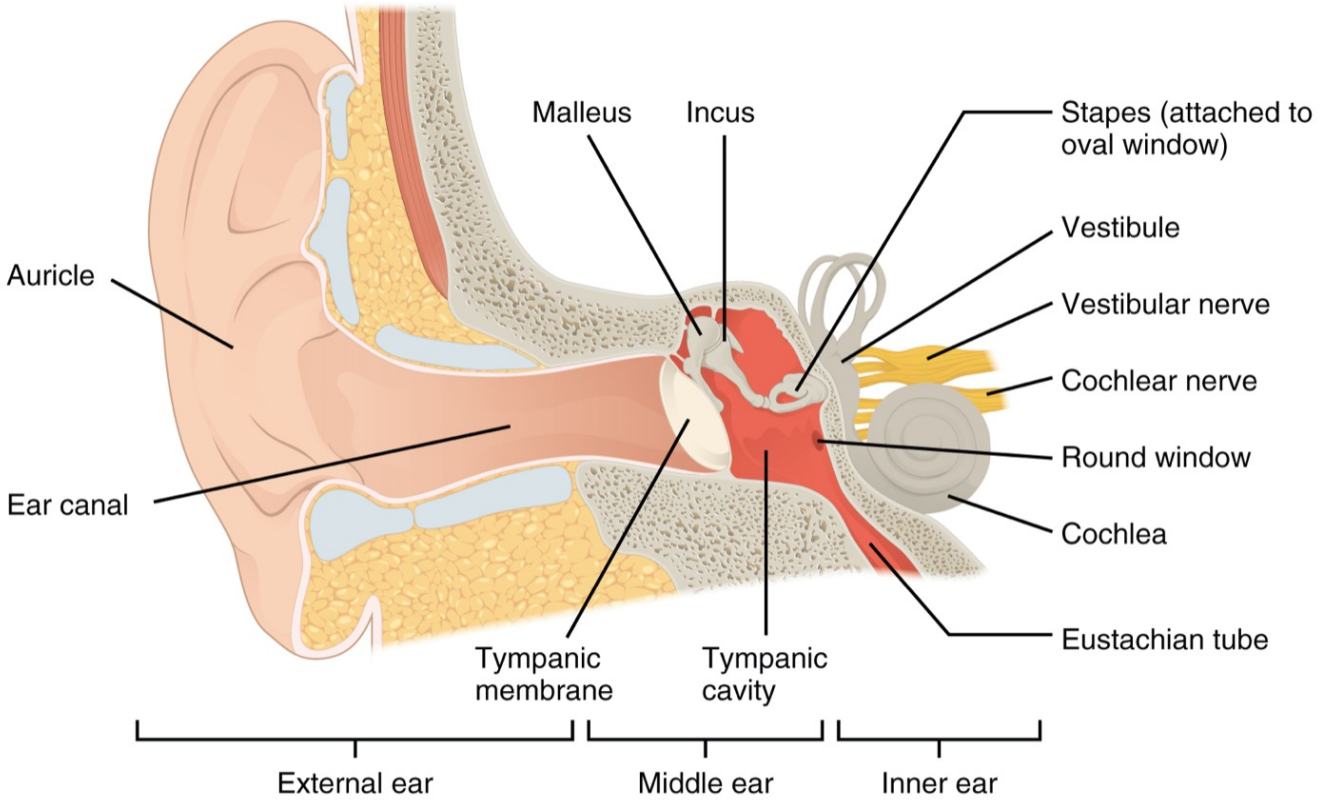
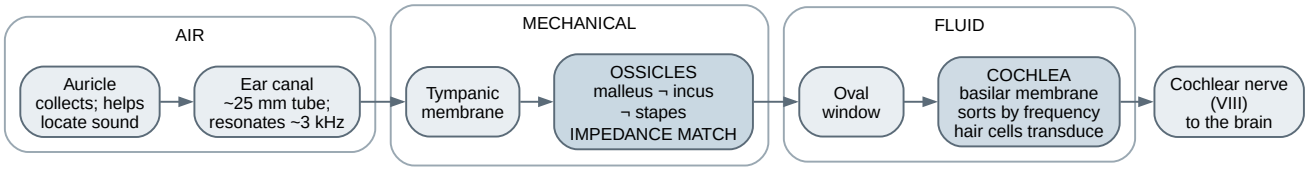


Figure 18. The external, middle, and inner ear. Image: OpenStax, CC BY 4.0. Credits: Appendix H

- **External ear** — the **auricle** (pinna), shaped cartilage covered with thin skin, and the **ear canal**, an S-shaped tube about 25 mm long and 7 mm across, ending at the **tympanic membrane** (eardrum) — about 8–10 mm across and 0.1 mm thick. The outer part of the canal is cartilage; the inner part is bone, covered with skin so thin that touching it is painful and bleeding is easy. This is why an otoscope speculum or an earmould is shaped and inserted with care.
- **Middle ear** — an **air-filled** cavity in the temporal bone containing the three smallest bones in the body, the **ossicles**: malleus, incus, and stapes. The stapes footplate sits in the **oval window** of the inner ear. The **Eustachian (auditory) tube** connects the middle ear to the back of the nose to equalise pressure — the "pop" when you swallow on a descending aircraft. In young children the tube is shorter and more horizontal, which is part of why middle ear infection is so common in them.
- **Inner ear** — a **fluid-filled** labyrinth of bone containing the **cochlea** (hearing) and the **vestibular apparatus** (balance).

8.4.2. The sound path



8.4.3. The cochlea

The cochlea is a fluid-filled tube coiled about two and a half turns, divided along its length by the **basilar membrane** — roughly 35 mm long if uncoiled. When the stapes pushes the oval window, a travelling wave runs along the membrane. The membrane is narrow and stiff at the base and wide and floppy at the apex, so each frequency peaks at a different place: **high frequencies near the base, low frequencies near the apex.** This **tonotopic** map is a mechanical spectrum analyser, and it is the reason a cochlear implant's electrode array stimulates different frequency bands at different depths (Chapter 7).

Along the membrane sit about 3,500 **inner hair cells**, which do most of the signalling to the brain, and about 12,000 **outer hair cells**, which act as amplifiers — they change length in time with the sound, sharpening tuning and boosting quiet sounds. They also emit faint sounds of their own, **otoacoustic emissions**, which a microphone in the ear canal can detect. That is the basis of newborn hearing screening.

> **Human hair cells do not regenerate.** Noise, some drugs, and age destroy them permanently. Noise-induced and age-related hearing loss both affect **high frequencies** first — which is where the consonants that distinguish words live. The typical complaint is not "I cannot hear" but "I can hear you, but I cannot understand you, especially in a noisy room."
>
> A device beep at 3–4 kHz may therefore be exactly the sound an older user cannot hear, while a lower-pitched tone of the same level is perfectly audible.

8.4.4. Two kinds of hearing loss

- **Conductive** — sound does not reach the cochlea efficiently: wax, fluid behind the eardrum, a perforation, fixed ossicles. Often treatable, and often helped by amplification or by **bone conduction**, which vibrates the skull and bypasses the outer and middle ear.
- **Sensorineural** — the cochlea or nerve is damaged. Amplification helps audibility but not clarity; distortion and poor performance in noise persist. Severe cases may be candidates for a cochlear implant.

8.4.5. Balance

The **vestibular apparatus** has two parts:

- Three **semicircular canals**, roughly at right angles, detect **rotation** of the head — angular acceleration — by the lag of their fluid as the head turns.
- Two **otolith organs**, the utricle and saccule, detect **linear acceleration and gravity** using hair cells loaded with tiny crystals. They tell you which way is down.

Together they form an inertial measurement unit — gyroscopes and accelerometers — and the brain fuses their output with vision and proprioception (Chapter 7). The fusion has a reflex attached: the **vestibulo-ocular reflex** counter-rotates the eyes when the head turns, so the image stays still on the retina. You can read this page while shaking your head; you cannot read it while shaking the page.

When the sources disagree — the eyes report motion that the inner ear does not, or the reverse — the result is **motion sickness**. A virtual-reality headset is a device for creating exactly that disagreement. Latency between head movement and display update, low frame rates, and artificial locomotion are the usual culprits, and some users, including many older adults and people with vestibular disorders, are badly affected. If you build simulation in VR, sickness is a design requirement, not a user weakness.

8.5. The nose and sinuses

The **nasal cavity** is divided by the **septum** (cartilage at the front, bone behind) into two passages. From each side wall project three scroll-shaped **turbinates** (conchae) covered with a richly vascular mucous membrane. They multiply surface area and turn the airflow turbulent, so that inspired air is **warmed, humidified, and filtered** in a few centimetres. By the time it reaches the lungs, air is close to body temperature and nearly saturated with water vapour; the nose does the larger share of that work.

The nose is also the largest single resistance in the airway — roughly half of the total during quiet breathing. That is why people mouth-breathe during exercise and with a blocked nose, and why the nasal passages are the narrowest point a nasal cannula, nasal swab, or nasogastric tube must negotiate.

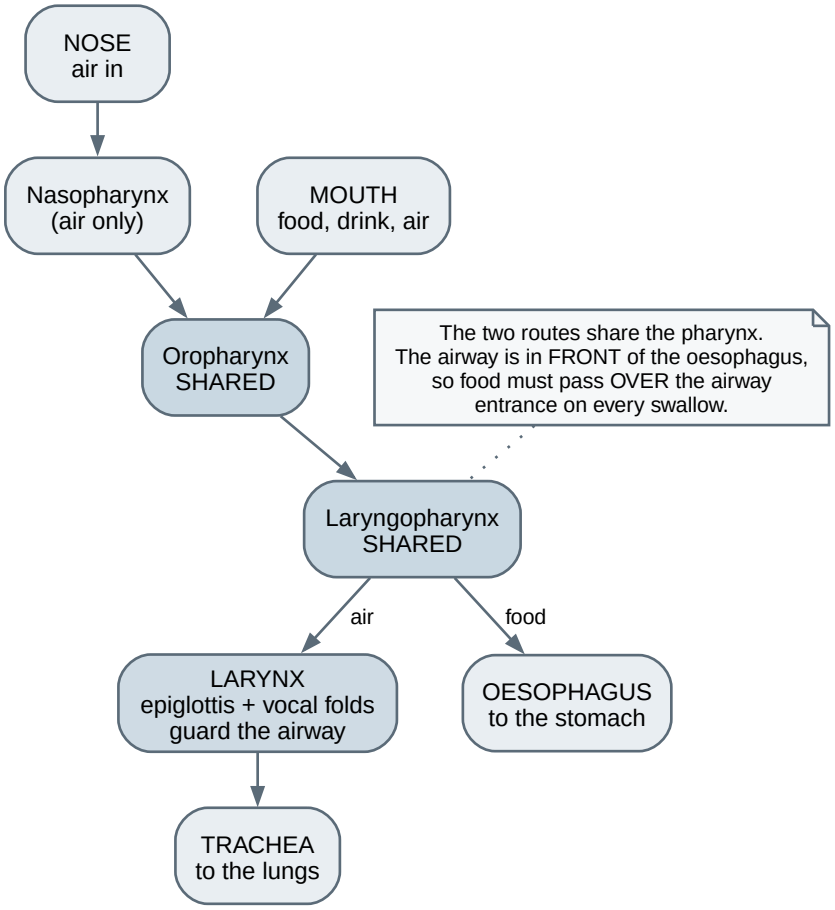
Around the nasal cavity are four pairs of air-filled **paranasal sinuses** — frontal, maxillary, ethmoid, and sphenoid — which drain into it through small openings. When those openings block, the sinuses fill and ache.

The **olfactory epithelium** is a small patch, a few square centimetres, at the very roof of the nasal cavity. Its receptor neurons — of roughly 400 functional types — send fibres up through perforations in the skull base (the **cribriform plate**) directly into the brain. Smell provides most of what we call **flavour**: taste on its own reports only sweet, sour, salty, bitter, and umami. Loss of smell is common with age and after viral infections, and it removes a warning channel — for smoke, gas, spoiled food — that people rarely realise they were relying on.

The front of the septum carries a dense network of small vessels (**Little's area**), the source of most nosebleeds. Pressure from a cannula prong, a nasal tube, or its securing tape on the septum or the rim of the nostril can cause pressure injury in a few days — on a **mucous membrane**, which the skin staging systems of Chapter 5 do not describe.

8.6. The throat: pharynx and larynx

8.6.1. The crossing



The **pharynx** is a muscular tube from the back of the nose to the top of the oesophagus, in three parts: the **nasopharynx** (air only, behind the nose, where the Eustachian tubes and adenoids are), the **oropharynx** (behind the mouth, where the tonsils are), and the **laryngopharynx** (behind the larynx). The lower two carry both air and food.

In the airway diagram of Chapter 9 the pharynx is one box between the nose and the larynx. That box is where the two systems cross, and the crossing is managed by a reflex.

8.6.2. The larynx and the voice

The **larynx** is a box of cartilages at the top of the trachea:

- the **thyroid cartilage**, the prominence at the front of the neck;
- the **cricoid cartilage** below it — the only **complete** ring of cartilage in the airway;
- the **epiglottis**, a leaf-shaped flap that tilts over the airway entrance during swallowing;
- the **vocal folds**, two bands of muscle and ligament that can close the airway completely, open it wide for breathing, or be held together and set vibrating by air from the lungs.

The gap between the vocal folds is the **glottis**, the narrowest part of the adult airway. Everything an airway manager does — intubating, sizing a tube, choosing a supraglottic device — is organised around getting past this point. The space between the cricoid and the thyroid cartilage, the **cricothyroid membrane**, is the site of the emergency surgical airway; it is a few millimetres under the skin of the front of the neck.

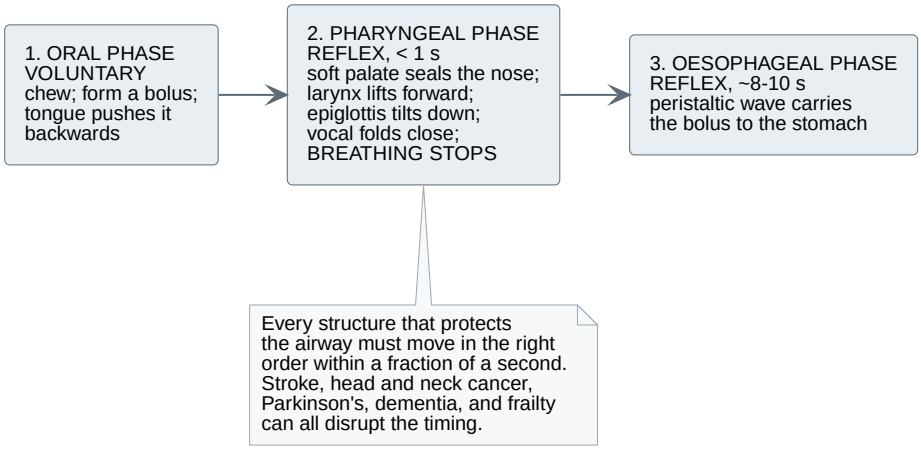
The voice is a **source–filter** system, which will be familiar if you have done any audio work. The

vibrating vocal folds are the source: a buzz with a fundamental frequency typically around 85–155 Hz in adult men and 165–255 Hz in adult women, rich in harmonics. The throat, mouth, and nose are the filter: their changing shape creates resonances, the **formants**, that make vowels distinguishable. Speech technology — synthesis, recognition, the electrolarynx used after laryngectomy — is built on this model.

The nerves to the larynx are branches of the vagus (X). The left **recurrent laryngeal nerve** loops down into the chest under the aorta before returning to the larynx, which is why surgery in the neck or chest can leave a hoarse voice.

8.6.3. Swallowing

A swallow happens several hundred times a day and is one of the most complex reflex sequences in the body.



Dysphagia is difficulty swallowing. Its central hazard is **aspiration**: food, drink, or saliva passing below the vocal folds into the airway, which can cause choking or pneumonia. Normally aspiration triggers a forceful cough. But **silent aspiration** — aspiration with no cough at all — is common in people with neurological conditions. It is Chapter 7' rule again: **absence of a visible reaction is not evidence of absence of harm.**

Management is clinical and individual — speech and language therapists assess swallowing, often with X-ray video (videofluoroscopy) or a nasal endoscope. A common intervention is modifying the **texture** of food and the **thickness** of drinks, now described internationally by the IDDSI framework [42]: eight levels from thin liquids (0) to regular food (7), each defined by a simple test — how fast a drink flows through a syringe, whether food holds its shape on a fork. Note the shape of that standard. It is a measurement method a carer can perform in a kitchen. That is good test design, and worth studying for its own sake.

> **If it can go in a mouth, assume it can go in an airway.**
>
> Small parts, detachable caps, pieces that can be bitten off a soft simulator, beads, button cells, magnets: in a child, in a person with dementia, or in anyone with a swallowing disorder, these are choking and aspiration hazards. Toy regulations use a test cylinder about 31.7 mm across for children under three; anything that fits inside is treated as a small part. That number is a regulatory convention for one population, not a safety limit for all of them — but it is a useful reminder of

> how large "small" is.
>
> Button cells and magnets are worse than inert parts: a swallowed button cell can burn through tissue in hours. Anything you build that contains one needs a battery compartment that cannot be opened without a tool.

8.6.4. When the upper airway is bypassed

Two situations change the anatomy in ways you must understand before designing anything used near the neck:

- A **tracheostomy** is an opening in the front of the trachea, below the larynx, holding a tube. The person breathes through the tube, so inspired air is no longer warmed and humidified by the nose — hence the **heat and moisture exchanger** (HME) worn over it. Air does not pass the vocal folds unless a one-way **speaking valve** redirects it, so speech may be absent. Smell is reduced because air no longer flows through the nose.
- After a **total laryngectomy** the larynx is removed and the trachea is brought to a permanent opening in the neck (a **stoma**). There is then **no connection at all** between the mouth and nose and the lungs. Oxygen given by face mask reaches nothing. People who breathe this way are called **neck breathers**, and some carry cards saying so.

For a maker this is not airway management — that is Tier 3 (Appendix I). It is awareness: a shower guard, a scarf clip, a stoma cover, or a training manikin for these users must respect an open airway at the front of the neck.

8.7. The mouth

The **oral cavity** is bounded by the lips, cheeks, hard palate (bone) and soft palate (muscle, which lifts to seal off the nose during swallowing), and the floor of the mouth, and contains the tongue and teeth.

- **Teeth** — 20 primary, 32 permanent. **Enamel** is the hardest material in the body, with an elastic modulus around 80 GPa — comparable to aluminium. Beneath it, **dentine** is closer to cortical bone (Chapter 6). Neither regenerates once lost.
- **The tongue** is a **muscular hydrostat**: muscle with no bone, changing shape at constant volume, like an octopus arm. It carries most of the **taste buds** — a few thousand in all — and is essential to chewing, swallowing, and speech.
- **Salivary glands** — parotid, submandibular, and sublingual — produce roughly 0.5–1.5 L of saliva a day. Saliva lubricates, starts starch digestion, buffers acid, and protects the teeth. **Dry mouth** is a common side effect of medication and of radiotherapy to the head and neck, and it makes eating, speaking, and wearing a denture hard.
- **The jaw** hinges and slides at the **temporomandibular joint**. A typical adult can open to about 40–55 mm between the incisors; maximal biting force in the molar region is commonly several hundred newtons. Both numbers vary widely and both matter to mouthpieces, bite blocks, and oral simulators.

The oral lining is mostly **non-keratinised** mucous membrane: it heals faster than skin and absorbs

some drugs directly (the reason for medicines dissolved under the tongue). It is also heavily colonised by bacteria. Oral hygiene in ventilated patients is an infection control measure (Chapter 13), because organisms from the mouth are a route to pneumonia.

8.8. The face and neck as a place devices sit

Site	Why it is vulnerable	Devices that injure it
Bridge of the nose	Thin skin directly over bone and cartilage; no padding	Masks (CPAP, oxygen, respirators), goggles, glasses, headsets
Behind and on top of the ears	Thin skin over poorly perfused cartilage	Oxygen tubing looped over the ears, mask straps, glasses arms, headset bands
Nostrils and septum	Mucosa and thin skin; small area takes the whole load	Nasal cannula prongs, nasogastric tubes and their tape
Lips, corners of the mouth, tongue	Mucosa; moisture; constant movement	Endotracheal tube holders, bite blocks, mouthpieces
Cheekbones and chin	Bony prominences under thin tissue	Mask cushions and harnesses
Front and sides of the neck	Airway, major vessels, and nerves close to the surface	Collars, straps, lanyards, tracheostomy ties

These are classic **medical device–related pressure injuries**, and they follow every rule of Chapter 5: small area, sustained load, thin tissue over bone, moisture, shear from movement, and a user who may not be able to report discomfort. Health workers wearing tight respirators for long shifts developed the same injuries on their own faces.

The neck adds two further hazards. On each side, where the carotid artery divides, the **carotid sinus** contains pressure sensors; external pressure there can slow the heart and drop the blood pressure. And anything that encircles the neck — a lanyard, a cord, a strap — is a strangulation risk, which is why clinical lanyards have breakaway clasps.

For eyewear and headsets the critical dimension is the **interpupillary distance** — the distance between pupil centres — typically about 54–74 mm in adults. An optical device that cannot be adjusted across that range will be wrong for a large share of its users.

8.9. Engineering bridge: optics, acoustics, and airflow

8.9.1. Dioptres

Optical power in dioptres is the reciprocal of focal length in metres: $P = 1 / f$. A +2 D lens focuses parallel light at 0.5 m. Powers of thin lenses in contact add, which is why dioptres are the working unit of eye care.

Use it to reason about everyday cases:

- A myope with a −2.0 D prescription sees clearly without glasses only out to 1 / 2.0 = **0.5 m** — the **far point**.
- As a rough rule, each millimetre of extra eye length adds about 2.5–3 D of myopia. A 24 mm eye that is 1 mm too long is a strongly myopic eye.
- A 60-year-old with about 1 D of accommodation who wants to read at 40 cm needs 1 / 0.40 = 2.5 D of focusing. The eye supplies 1; reading glasses supply the other **+1.5 D**.

8.9.2. Visual angle and legibility

The legibility of anything depends not on its size but on the **angle it subtends at the eye**. For small angles:

θ (radians) ≈ height ÷ distance, and 1 arcminute ≈ 0.000291 rad.

Standard acuity — "20/20" or "6/6" — is the ability to read a letter that subtends **5 arcminutes** overall, built from strokes of 1 arcminute. So the smallest letter a person with standard acuity can just read at distance **d** is:

h ≈ d × 5 × 0.000291 ≈ d × 0.00145

Viewing distance	Threshold letter height, 6/6	Threshold, 6/12	Threshold, 6/18 (moderate impairment begins below this)
0.4 m (reading)	0.6 mm	1.2 mm	1.7 mm
1 m	1.5 mm	2.9 mm	4.4 mm
3 m	4.4 mm	8.7 mm	13 mm
6 m	8.7 mm	17 mm	26 mm

These are **thresholds**: letters a person can just make out under test conditions with high contrast. Comfortable, fast, error-free reading needs a reserve of **two to three times** threshold, more in poor light, at low contrast, or for a user who is moving, stressed, or older. That margin is a design decision, and it should be written down.

8.9.3. Decibels

Sound pressure level is logarithmic:

$L = 20 \log_{10}(p / p_0)$ dB SPL, with p_0 = 20 μPa, roughly the quietest sound a young ear can detect at 1 kHz.

Level, dB(A)	Typical source (orientation only)
0	Threshold of hearing, young adult
30	Whisper; quiet bedroom at night
45–55	Quiet office; many wards during the day

Level, dB(A)	Typical source (orientation only)
60	Conversation at 1 m
70–80	Busy street; vacuum cleaner; some alarms at 1 m
85	Occupational action level for an 8-hour day
100	Power tools near the ear
120–140	Threshold of pain

Rules of thumb worth carrying:

- **+6 dB doubles sound pressure**; **+10 dB** is perceived as roughly **twice as loud**.
- **Two equal, independent sources add 3 dB**, not double the decibels: 60 dB + 60 dB = 63 dB.
- In open space, level drops about **6 dB for each doubling of distance**. Rooms reflect sound, so indoors the fall-off is less — but a quiet cue still fades fast.
- **Exposure trades level against time.** Under the widely used 85 dB(A), 3-dB exchange criterion, each 3 dB increase halves the allowable daily duration: 8 hours at 85, 4 at 88, 2 at 91, 1 at 94, about 15 minutes at 100. Workshop machinery reaches those levels easily. Wear hearing protection; hair cells do not come back.
- **dB(A)** weights the measurement toward the frequencies the ear hears best. **dB HL** ("hearing level"), used on an audiogram, is referenced to normal hearing at each frequency, so 0 dB HL means "typical" at every frequency. They are different scales; do not mix them.

8.9.4. Two pieces of ear engineering

The ear canal is a quarter-wave resonator. A tube closed at one end resonates where its length is a quarter of a wavelength: f = c / 4L = 343 / (4 × 0.025) ≈ **3.4 kHz**. The ear canal boosts sound in the 2–4 kHz region by around 10–15 dB, which is part of why we are most sensitive there — and why blocking the canal with an earbud, an earmould, or a hearing aid **removes** that gain unless the device restores it. The same blockage traps bone-conducted sound, making the wearer's own voice boom: the **occlusion effect**, a classic hearing-aid complaint and a vent-design problem.

The middle ear is an impedance-matching transformer. Sound moving from air into fluid would be almost entirely reflected — about 99.9% of the energy — because the acoustic impedances differ by a factor of several thousand. The middle ear compensates in two ways. The effective area of the eardrum (~55 mm²) is about 17 times that of the stapes footplate (~3.2 mm²), and the ossicles add a lever advantage of about 1.3. Together:

pressure gain ≈ 17 × 1.3 ≈ 22, or 20 $\log_{10}$ 22 ≈ 27 dB.

That is roughly the hearing loss a person has when the ossicular chain is disrupted — which is what makes conductive loss of this size **predictable from the mechanics**. An engineer would recognise the whole arrangement as a step-down transformer between a low-impedance source and a high-impedance load.

8.9.5. Airflow in the upper airway

A resting adult's **peak** inspiratory flow is roughly 25–30 L/min, and much higher with exertion or breathlessness. A nasal cannula supplies oxygen at perhaps 1–6 L/min. So with every breath the person draws in a mixture of cannula oxygen and room air, in a proportion that depends on how they are breathing. The oxygen fraction actually inhaled is therefore **not known** from the flow setting — it is estimated, and the estimate is rough. Keep that in mind if you ever build a simulator or training software that displays one: the honest output is a range, not a number (Chapter 9, Chapter 15).

Higher flows dry and irritate the nasal lining, because the nose cannot humidify that much air. That is why high-flow nasal oxygen systems heat and humidify the gas — and why a person breathing through a tracheostomy, with the nose bypassed entirely, needs an HME.

8.10. Devices, software, and training technology

ENT, ophthalmology, audiology, speech and language therapy, and dentistry are five separate professions, each with its own equipment and standards. The head and neck also contain the most heavily used task trainers in emergency medicine.

8.10.1. Clinical devices

Device	Does	Engineering note
Spectacles and contact lenses	Correct refractive error	Contact lenses are regulated medical devices on the cornea: oxygen transmissibility, fit, and cleaning regime are the safety-critical variables
Intraocular lenses	Replace the clouded lens in cataract surgery — one of the most common operations performed	Implant. Optical design, material, and power calculation from eye measurements
Tonometers	Measure intraocular pressure	Contact (applanation) and air-puff types. Indirect measurement, affected by corneal thickness — a sensor error model in a single instrument
Ophthalmoscope, slit lamp, fundus camera	View and photograph the inside of the eye	Illumination and viewing optics sharing one small pupil; light safety limits apply
Optical coherence tomography (OCT)	Cross-sectional imaging of the retina at micrometre resolution	Low-coherence interferometry. Retinal layers a few micrometres thick are resolved without touching the eye
Low-vision aids	Magnifiers, video magnifiers, telescopic aids	Magnification, field of view, and working distance trade against each other — the optics of this chapter's bridge in a box

Device	Does	Engineering note
Otoscope	Views the ear canal and eardrum	Speculum sizing; the bony canal is easily injured
Audiometer	Measures hearing thresholds frequency by frequency	Calibrated transducers; output in dB HL; must be tested in a quiet booth
Tympanometer	Measures how easily the eardrum moves as ear-canal pressure is varied	A **compliance** measurement — conceptually the same quantity as the test lung in Chapter 9
Otoacoustic emission screeners	Newborn hearing screening	Detects sounds emitted by the outer hair cells — an active cochlea proves itself
Hearing aids, bone-conduction devices, cochlear implants	Restore audibility	See Chapter 7. Earmould and shell fit are made from ear-canal impressions or scans — custom-fitted manufacturing on a large scale
Ear tubes (tympanostomy tubes)	Ventilate the middle ear in recurrent fluid or infection	Tiny implant through the eardrum
Nasal cannulas, high-flow nasal oxygen, CPAP masks	Deliver oxygen or pressure through the nose	Interface design is a **pressure-injury** problem (this chapter, Chapter 5) as much as a gas-delivery one
Laryngoscopes and videolaryngoscopes	Expose the glottis for tube placement	A camera on a blade: optics, lighting, fogging, and cleaning (Chapter 13)
Tracheostomy tubes, speaking valves, HMEs	Maintain and condition an airway that bypasses the upper airway	Life-sustaining. Tier 3 (Appendix I)
Electrolarynx and voice prostheses	Restore voice after laryngectomy	The source–filter model made literal: an external buzzer replaces the vocal folds
Dysphagia assessment — videofluoroscopy and endoscopic evaluation	Show what a swallow actually does	X-ray video and flexible endoscopy. Clinical procedures
Dental devices — intraoral scanners, aligners, crowns, surgical guides	Restore and move teeth	**One of the largest industrial uses of 3D printing in medicine.** Scan-to-print workflows are mature here years ahead of most other specialties

8.10.2. Software

Software	Does	Note
Screen magnification — built-in OS magnifiers and zoom	Enlarges any part of the screen	Test your interfaces at 200% and 400%. Layouts that break are a barrier for low-vision users
Colour-vision and low-vision simulators — Color Oracle, browser DevTools vision emulation	Preview an interface as it appears with colour deficiency or blur	**Free and immediate.** A finding tool, not a substitute for testing with users (Chapter 7)
Live captioning and transcription	Real-time text from speech	Built into the major mobile and desktop platforms. Accuracy drops with noise, accent, and medical vocabulary
Hearing-aid and sound-amplification apps	Fitting adjustment, environmental programs, personal amplification	Increasingly over-the-counter; the user is now often the person doing the fitting
Hearing-screening apps	Rough threshold estimates with consumer headphones	Screening only. Uncalibrated transducers in an uncontrolled room are not an audiometer
Speech analysis — Praat	Pitch, formants, spectrograms, voice measurements	**Free and widely used in speech science.** The source–filter model on your laptop
Text-to-speech and speech recognition	Alternative output and input	See Chapter 7' AAC and dictation entries
Eye-tracking software	Gaze-based input and research measurement	Calibration and the VOR; relates directly to the EOG row of Chapter 7' biosignal table
Public retinal and ENT image datasets	Fundus photographs and OCT scans for research and education	The safe route into ophthalmic image analysis — the same reasoning as Chapter 15's open imaging data

8.10.3. Training and simulation

Training technology	Teaches	Form
Airway management heads	Bag–mask ventilation, supraglottic devices, intubation	Realistic upper airway in a head and neck. **The replaceable parts — tongues, airways, teeth — wear out constantly** (Appendix I)
Cricothyroidotomy trainers	The emergency surgical airway	A neck with a replaceable skin layer and membrane. **Consumed on every attempt; a maker-accessible build**
Tracheostomy care trainers	Suctioning, tube change, stoma care	Neck and upper chest with a stoma
Otoscopy trainers	Ear examination and recognition of eardrum appearances	Model ears with interchangeable eardrum inserts. **Small, precise, makeable**
Ear-wax removal trainers	Irrigation and microsuction	Model canal with simulated wax; tactile fidelity of the canal wall matters
Fundoscopy trainers	Examining the retina through the pupil	Model eyes with interchangeable retinal images and adjustable pupil size — an optics project with a real teaching purpose
Ophthalmic surgical simulators	Cataract and retinal surgery	Mostly virtual reality with instrumented hand-pieces. Commercial and specialised
Nasal swab and nasogastric tube trainers	Correct angle and depth through the nose	Sagittal-section models. Widely needed after respiratory pandemics
Dental phantom heads and typodonts	Every practical skill in dental education	Replaceable teeth — printed, cast, or moulded from scanned originals
Swallow and voice teaching models	Anatomy of the larynx and the swallow sequence	Articulated models showing epiglottis and laryngeal elevation. Close cousins of your capstone project



In this region, three ordinary maker materials and devices become hazards.

- **Light sources.** UV, high-intensity blue LEDs, and lasers can damage the eye without pain. Any light source in something you build that a person might look into must be assessed against the photobiological and laser safety standards in Appendix D, or physically shielded.
- **Sound.** An alarm or buzzer placed close to the ear can exceed safe levels at the ear, even if it seems moderate at arm's length — level rises 6 dB with every halving of distance.

- **Small parts.** As above: anything that can be detached, bitten off, or swallowed.

And nothing you make goes in an eye, an ear, a nose, a mouth, or an airway on a person. Those are Tier 3 in Appendix I without exception.

Good entry points in this chapter span the whole workshop.

- **Cast and moulded consumables** — replaceable necks, cricothyroid membranes, tongues, and ear canals for airway and otoscopy trainers — are constantly needed: silicone casting into machined or printed moulds, with layered skin and fabric reinforcement.
- **Optics and glass** — fundoscopy trainer eyes need a real lens, a pupil aperture, and an interchangeable retinal image; an honest optics project rather than a shell.
- **Electronics and sound** — calibrated sound sources for audibility demonstrations, light-level and contrast test rigs, and switch or caption hardware for people with sensory loss.
- **Calculation before construction** — the **visual-angle and decibel calculations** of the engineering bridge let you check any label, screen, or alarm you design against human limits before you build it — and then test it.

8.11. Worked example: the label at the back of the room

Chapter 7 gave your elbow model redundant coding: a "PULL TO BEND" label, distinct tab shapes, an arrow, and a click when a part locks. The educator now plans to use the model in front of a class of twenty, with the furthest learner about 4 m away. Is the coding still working?

The label. It was printed with letters 3 mm tall — comfortable in the hand. At 4 m:

- standard acuity threshold: 4 m × 0.00145 ≈ **5.8 mm**;
- a learner with 6/12 vision — common, and often uncorrected for distance — ≈ **11.6 mm**;
- with a reading reserve of two to three times threshold: ≈ **23–35 mm**.

A 3 mm label is below threshold for **everyone** in the room beyond about 2 m. It was designed for a viewing distance that is no longer the use case.

The click. Suppose — an invented figure — the click measures about 50 dB(A) at 0.5 m. At 4 m, three doublings of distance away, open-space fall-off would bring it to about 50 − 18 = **32 dB(A)**. A teaching room with twenty people in it is commonly 45–55 dB(A). The click is inaudible to most of the class, and even more so for anyone with high-frequency hearing loss, since a click is mostly high-frequency energy.

The redesign. Not a bigger model — the model's labels are correct for its original use, being handled. Instead:

- add a **large demonstration card**, with letters sized to the room (tens of millimetres) and the same arrow symbol as the model, so the two are visibly linked;
- make the locked state **visible at distance** — a large, high-contrast flag that only appears when the part is engaged;
- or put the model under a document camera and project it, which solves size, but not the click.

Then test it. Stand at the back of the actual room, in its actual lighting, with someone who did not make the model. Record what they could and could not read, and whether you prompted them.

The general lesson is in the first sentence of the redesign: **a cue that works is a cue that works at a stated distance, level, and lighting, for a stated user.** Change any of those and it must be checked again. The calculations here took two minutes and caught the problem before anybody stood in the room.

8.12. Workshop exercise

1. **Find your blind spot.** On a card, draw a small cross and, 8 cm to its right, a dot. Close your left eye, look at the cross with your right, and move the card slowly toward and away from you until the dot vanishes. Measure the eye-to-card distance **D** and calculate the angle: arctan(8 / D). Compare with the ~15° in the text.
2. **Measure a visual angle.** Print the same word at five letter heights, from 2 mm to 20 mm. Find the distance at which you can just read each one, and calculate the angle subtended in arcminutes. This is not a clinical acuity test — the point is to see that the **angle** is roughly constant while the size and distance change.
3. **Map a room's sound.** Use a sound-level meter app on a phone. Record the level in your workshop with and without a machine running, in a quiet room, and in a busy public space. Note that the phone is **uncalibrated**: record the readings as relative, not absolute. Then measure your elbow model's click at 0.5 m, 1 m, and 2 m.
4. **Feel a swallow.** Rest your fingertips lightly on the front of your neck over the thyroid cartilage and swallow. Feel the larynx rise and move forward. Try to breathe in at the moment of swallowing; you will find you cannot.
5. **Simulate and redesign.** Run a colour-vision and blur simulation over a photograph of your model's labels and over one screen of any software you use. List what fails, then apply this chapter's legibility calculation to fix one item.
6. **Map pressure on your own face.** After an hour wearing glasses, safety goggles, or a respirator, look for red marks. Record their positions on a sketch of the face and name the bony prominence or cartilage under each one.

8.13. Questions

1. Name the five special senses, their organs, and the cranial nerve for each.
2. Which part of the eye provides most of its focusing power, and why there rather than the lens?
3. What is presbyopia, and what does it imply for a label on a device used by adults of all ages?
4. Glaucoma and macular degeneration both cause low vision. Why might a design that helps one

fail the other?

5. Explain how the middle ear achieves its pressure gain, and estimate the gain in decibels.
6. Why does age-related hearing loss make speech harder to **understand** rather than only quieter, and what does this imply for an alarm tone?
7. A buzzer measures 70 dB(A) at 1 m. Roughly what is it at 4 m in open space, and at 0.25 m?
8. What causes motion sickness in a VR simulator, and what follows for its design?
9. What is silent aspiration, and which principle from Chapter 7 does it illustrate?
10. Why is face-mask oxygen ineffective for a person after total laryngectomy?
11. List four sites on the face and neck where devices commonly cause pressure injury, and why each is vulnerable.
12. What letter height does a person with 6/12 vision need to just read a label at 3 m, and what would you actually specify?

8.14. Answers and reasoning

1. **Vision** — eye — II (optic). **Hearing** — cochlea — VIII. **Balance** — vestibular apparatus — VIII. **Smell** — olfactory epithelium — I. **Taste** — taste buds — VII, IX, X. A strong answer adds that special senses go directly to the brain by cranial nerves, not through the spinal cord.
2. The **cornea**, about 43 of roughly 60 dioptres. Refraction is greatest where the change in refractive index is largest, which is the air–cornea interface; inside the eye the lens sits between fluids of similar refractive index, so its contribution is smaller. The lens matters because its power is **adjustable**.
3. Age-related loss of the lens's ability to change shape, so near focus is lost — the near point moves from about 10 cm in youth to about a metre by 60. For a device: small print read at close range will be unreadable to many older users without reading glasses, and a user may not have their glasses with them. Size text for the stated viewing distance with a reserve, and do not hide essential information in fine print.
4. **Glaucoma** removes peripheral vision first; central vision may be intact. **Macular degeneration** removes central vision; the periphery is preserved. A design that places key information in the periphery, or depends on noticing a peripheral indicator, fails the first. A design that relies on magnifying fine central detail helps the first and fails the second. The pattern of loss determines what helps, so the diagnosis alone is not the specification.
5. By **area ratio** — the effective area of the eardrum is about 17 times that of the stapes footplate, concentrating force — and by the **ossicular lever**, about 1.3. Pressure gain ≈ 17 × 1.3 ≈ 22, which is 20 $\log_{10}$ 22 ≈ 27 dB. Without it, most of the sound energy would be reflected at the air–fluid boundary.
6. It destroys hair cells at the cochlear base first, which serve **high frequencies** — where the consonant sounds that distinguish words are carried. Vowels, which are loud and lower in frequency, are still heard, so speech is audible but unclear, especially in noise. A high-pitched alarm (3–4 kHz) may be inaudible to exactly the users most likely to have this loss; a lower-frequency or multi-frequency tone, plus a visible cue, is more robust.
7. 4 m is two doublings of 1 m, so about 70 − 12 = **58 dB(A)**. 0.25 m is two halvings, so about 70 + 12

= **82 dB(A)** — close to the occupational action level, delivered to an ear very close to the device. Indoors the fall-off with distance will be smaller than this.

8. **Sensory conflict**: vision reports self-motion that the vestibular system does not detect, or the reverse — worsened by display latency, low frame rate, and artificial locomotion. Design consequences: keep motion-to-display latency low, avoid artificial movement the user did not cause, provide stable visual reference, keep sessions short, and treat sickness as a requirement to test for, not a user failing.
9. Aspiration — material passing below the vocal folds — **without a cough** or any other visible reaction. It illustrates that **absence of a reaction is not evidence of absence of harm**: a protective response that depends on the person noticing or reacting cannot be assumed.
10. After total laryngectomy the trachea opens directly at a stoma in the neck and there is no connection between the mouth and nose and the lungs. Oxygen applied to the face reaches nothing; it must be applied at the stoma.
11. Any four of: **bridge of the nose** (thin skin over bone and cartilage); **ears** (thin skin over poorly perfused cartilage, loaded by tubing and straps); **nostrils and septum** (mucosa, small loaded area, prongs and tubes); **lips and mouth** (mucosa, moisture, movement); **cheekbones and chin** (bony prominences); **neck** (superficial airway and vessels; strangulation risk). All follow Chapter 5' pattern of sustained load on thin tissue over a hard structure.
12. Threshold for 6/12 at 3 m: 3 × 0.00145 × 2 ≈ **8.7 mm**. That is the height at which the label is **just** readable under good conditions. A specification would add a reserve of two to three times — roughly 17–26 mm — state the assumed lighting and contrast, and then be tested in the real room.

8.15. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Labels, alarms and the senses — letter height for a reading distance, alarm audibility, noise exposure, the ear canal, and colour-vision simulation.
- Physics — the eye in dioptres: far point, accommodation and reading add.
- Acoustics — adding decibels and inverse square.
- Audio scope — spectrograms of the voice.

8.16. Video companion and reading

- **Core videos:** the Crash Course Anatomy & Physiology episodes on **taste and smell**, **vision**, and **hearing and balance**, from the series index [6].
- **Core reading:** OpenStax [1] selectively:
  - 14.1 Sensory Perception — taste, smell, audition, equilibrium, and vision
  - 13.4 The Peripheral Nervous System — the cranial nerves
  - 22.1 Organs and Structures of the Respiratory System — the nose, pharynx, and larynx

    - 23.3 The Mouth, Pharynx, and Esophagus — including deglutition (swallowing)
- **Context:** the WHO **World report on hearing** [40] and **World report on vision** [41]. Read the summaries; they explain how common these losses are and how much of the burden is preventable or cheaply correctable.
- **Test design worth studying:** the IDDSI framework [42] — a texture-measurement standard designed so that a carer in a kitchen can perform every test.

8.17. Exit criteria

- ▢ Trace the light path through the eye and the sound path through the ear from memory, naming each structure and its job.
- ▢ Explain the middle ear's impedance matching and calculate its approximate gain.
- ▢ Draw the airway–food crossing and explain, step by step, how a swallow protects the airway.
- ▢ Calculate threshold and specified letter heights for at least one real label at its real viewing distance, and record the assumptions.
- ▢ Measure (relatively) a sound cue at three distances and compare with the ambient level of the room where it will be used.
- ▢ Identify every point where a head-worn or face-worn item you own presses on thin tissue over bone or cartilage.
- ▢ State why nothing you make goes into an eye, ear, nose, mouth, or airway on a person.

8.18. What this chapter does not establish

- **You cannot test anyone's vision or hearing.** A letter chart and a phone app show you the principles. Acuity, visual field, and audiometric testing are clinical measurements made with calibrated equipment by qualified people.
- **You cannot assess a swallow or recommend a diet texture.** Dysphagia assessment and management are the work of speech and language therapists and the clinical team. The IDDSI levels appear here as an example of good measurement design, not as something to apply.
- **You have not learned airway management.** The anatomy here lets you follow a conversation about it and understand the trainers. Managing a real airway is clinical practice under training and supervision.
- **Light and sound safety are not settled by intuition.** If a design has a light source someone might look into, or a sound source near an ear, the applicable standard and a measurement decide whether it is safe — not whether it seemed fine on the bench.
- **ENT, ophthalmology, audiology, speech and language therapy, and dentistry are separate specialties.** This chapter is a map of the territory. Each of them is a field you would enter through a named partner, as Chapter 18 describes.

Chapter 9. Circulation, breathing, and the wider body

Week 12 · approximately 4–6 hours

9.1. Why this chapter exists

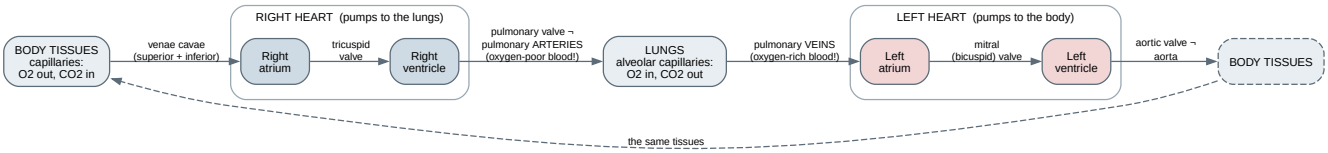
Two reasons, and they pull in different directions.

The first is comprehension. You need enough physiology to follow a clinical conversation without stopping it every third sentence, and to recognise when a seemingly small design change touches an essential body function. "Can you make the strap a bit tighter?" is a different question if the strap crosses a superficial artery.

The second is restraint. This is the chapter where someone fluent in software and electronics is most likely to build something dangerous without touching a single sharp edge. A monitor display is easy to make. It is also a device that can look authoritative while showing a stale value in the wrong units, and the gap between **displaying a number** and **supporting a clinical decision** is the most important boundary in this book.

9.2. The heart and the two circuits

9.2.1. Chambers and valves



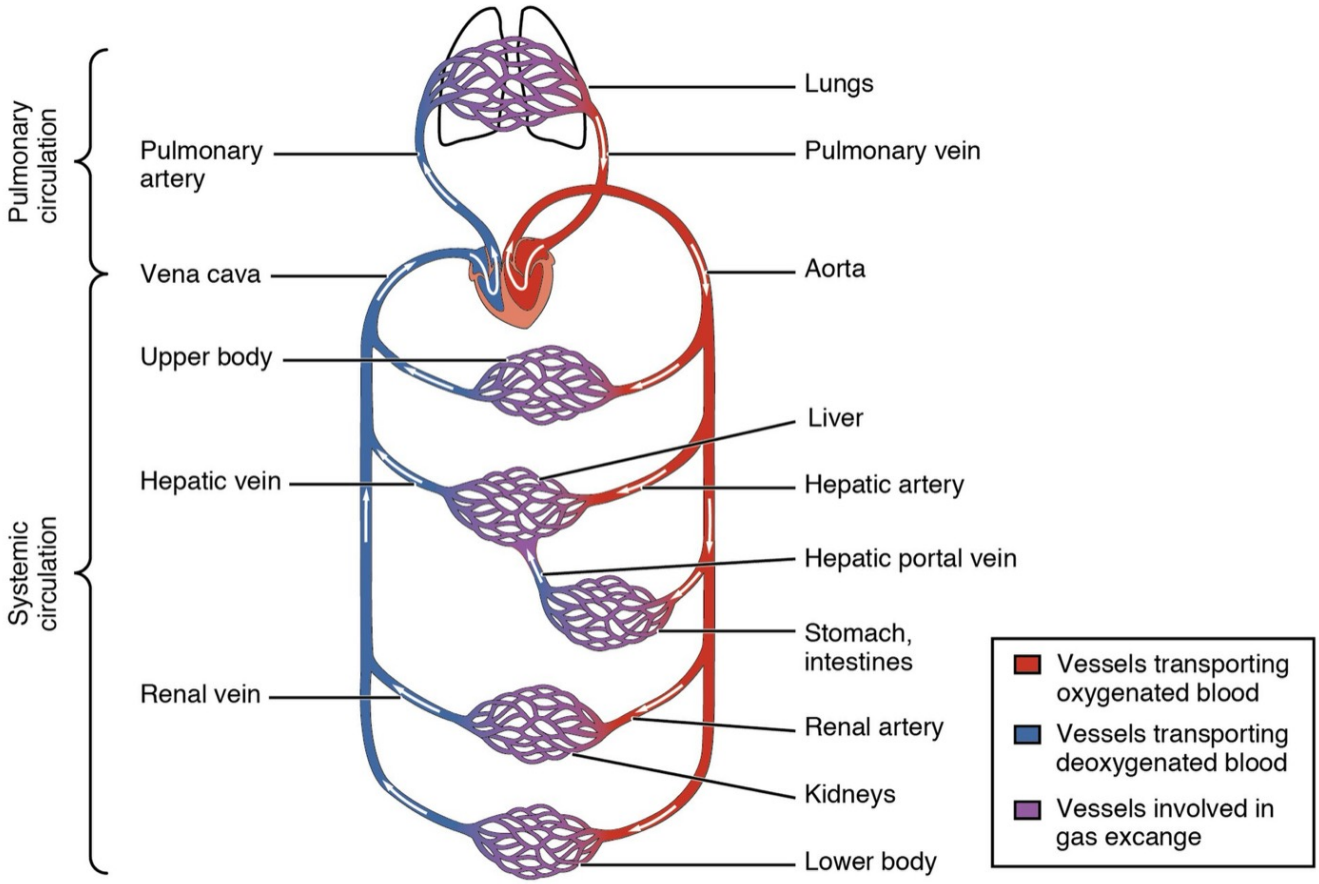


Figure 19. Blood flow through the heart. Image: OpenStax College, CC BY 3.0. Credits: Appendix H

Memorise this sequence; you will use it constantly:

> body → **superior and inferior venae cavae** → **right atrium** → tricuspid valve → **right ventricle** → pulmonary valve → **pulmonary arteries** → **lungs** → **pulmonary veins** → **left atrium** → mitral valve → **left ventricle** → aortic valve → **aorta** → body

Arteries and veins are defined by direction of flow, not by oxygen content.

An **artery** carries blood **away from** the heart. A **vein** carries blood **toward** the heart. That is the whole definition.

The **pulmonary arteries** carry relatively oxygen-**poor** blood from the right ventricle to the lungs. The **pulmonary veins** carry oxygen-**rich** blood from the lungs to the left atrium. The shortcut "arteries are red, veins are blue" works for the systemic circuit and is wrong for the pulmonary circuit, which is exactly why it is the standard exam question.

9.2.2. Vessel types

Vessel	Wall	Function	Mechanical note
Artery	Thick, elastic and muscular layers	Carries blood from the heart at high pressure	Elastic recoil in the aorta smooths pulsatile flow — a **Windkessel** (air-chamber) effect, a genuine engineering analogy
Arteriole	Mostly smooth muscle	Regulates flow into capillary beds	The principal site of resistance control; small radius changes have large effects
Capillary	A single layer of endothelium	**Exchange**: gases, nutrients, waste, fluid	~5–10 μm lumen; thin wall minimises diffusion distance
Venule / vein	Thin, distensible	Returns blood at low pressure; **reservoir** holding most of the blood volume	Valves prevent backflow in limbs; skeletal muscle contraction acts as a pump
Lymphatic	Thin, valved	Returns interstitial fluid and carries immune traffic	Its failure causes lymphoedema — relevant to any compressive design

9.2.3. The cardiac cycle

Phase	Ventricles	What you would hear or see
Atrial systole	Filling completes; atria contract to top up ventricles	End of diastole
Ventricular systole — isovolumetric contraction	Pressure rises, all valves shut	First heart sound (S1, "lub") is AV valves closing at its start
Ventricular systole — ejection	Pressure exceeds arterial pressure; semilunar valves open; blood ejected	The arterial pulse you can palpate
Ventricular diastole — isovolumetric relaxation	Pressure falls, semilunar valves shut	Second heart sound (S2, "dub") is semilunar valves closing
Ventricular diastole — filling	AV valves open; passive filling, then atrial systole	The longest phase at rest

Some numbers for intuition in a typical resting adult [1]:

Quantity	Typical resting value	Relationship
Heart rate (HR)	~60–80 beats/min	—

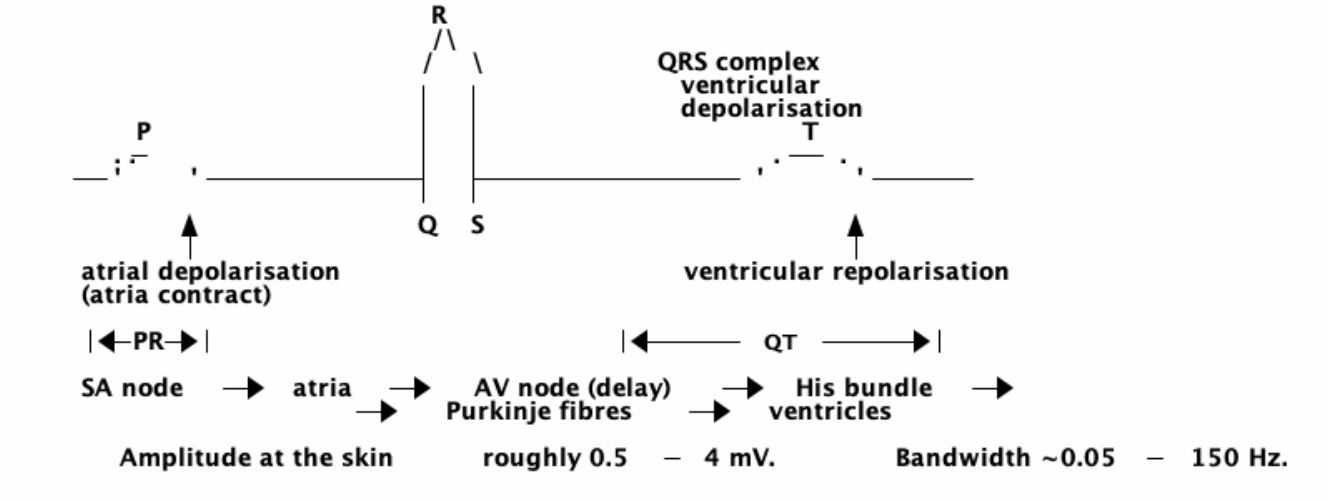
Quantity	Typical resting value	Relationship
Stroke volume (SV)	~70 mL	Volume ejected per beat
Cardiac output (CO)	~5 L/min	CO = HR × SV
End-diastolic volume	~120 mL	Volume before ejection
End-systolic volume	~50 mL	Volume remaining after ejection
Ejection fraction	~55–70%	SV / EDV
Total blood volume	~5 L	So the whole volume circulates roughly once a minute



CO = HR × SV is the single most useful equation in cardiovascular physiology, and it explains a great deal of clinical reasoning. Output can be raised by beating faster or by ejecting more per beat. Both have limits, and at very high rates filling time shortens enough that stroke volume falls — so output does not rise indefinitely with rate.

9.2.4. The conduction system and the ECG

The heart generates its own rhythm. The **sinoatrial (SA) node** depolarises spontaneously, acting as the pacemaker; the signal spreads across the atria, is briefly delayed at the **atrioventricular (AV) node** (allowing the ventricles to fill), then travels rapidly down the **bundle of His**, the left and right **bundle branches**, and the **Purkinje fibres** to depolarise the ventricles from the apex upward.



The ECG is the **electrical** activity recorded at the body surface. It is not the mechanical pump action, and the two can dissociate — electrical activity can be present without effective mechanical output. Knowing the letters P, QRS, and T, and what each corresponds to, is the right depth for this book.



Reading an ECG is a clinical skill that takes years to do well and is routinely got wrong by people who have had some training. Knowing what P, QRS, and T **represent** is general knowledge. Interpreting a trace, recognising arrhythmias, or

setting alarm thresholds is not, and no part of this chapter moves you toward it.

And as in Chapter 7: do not connect anything you built to a person.

9.3. Engineering bridge: haemodynamics

The circulation is a pressurised flow network, and the engineering analogies are genuinely productive — provided you know where they break.

9.3.1. The Ohm's law analogue

ΔP = **Q** × **R** — pressure drop = flow × resistance

Flow is driven by a pressure difference against a resistance. In the circulation, mean arterial pressure minus central venous pressure drives cardiac output against total **systemic vascular resistance**.

9.3.2. Poiseuille's law, and why arterioles matter so much

For steady laminar flow of a Newtonian fluid through a rigid cylindrical tube:

$$\mathbf{Q} = (\pi \cdot \mathbf{\Delta P} \cdot \mathbf{r}^4) \div (8 \cdot \boldsymbol{\eta} \cdot \mathbf{L})$$

The exponent on radius is **four**. Halve a vessel's radius and flow falls to one sixteenth of its previous value, all else equal. This is why the arterioles — small vessels wrapped in smooth muscle — are the body's flow-control valves, and why **vasoconstriction** and **vasodilation** are such powerful physiological levers.

It is also, directly, why compression matters. Anything that reduces a vessel's radius reduces flow steeply. An object pressing on tissue is a flow-control device whether you intended it to be or not.



Where Poiseuille's law fails in the body, and each failure matters:

- **Blood is not Newtonian.** It is a suspension of cells; its apparent viscosity depends on shear rate, haematocrit, and vessel diameter.
- **Vessels are not rigid.** They are compliant and they taper and branch. Compliance is what stores energy in systole and releases it in diastole.
- **Flow is not steady.** It is pulsatile, with wave reflections from branch points.
- **Flow is not always laminar.** It becomes disturbed at branches, stenoses, and valves — which is what produces audible murmurs and bruits.
- **The network is actively regulated.** Radius is under continuous neural, hormonal, and local metabolic control. The "resistor" changes value in response to the flow through it.

Use the r^4 relationship to understand **sensitivity to radius**. Do not use the equation to compute a flow in a person.

9.3.3. Pressure, and what a blood pressure reading is

Arterial pressure is conventionally reported as systolic over diastolic in mmHg — for example 120/80. **Systolic** is the peak during ventricular ejection; **diastolic** is the trough between beats. **Mean arterial pressure** is often estimated as

$$\mathbf{MAP} \approx \mathbf{P}_{dia} + \tfrac{1}{3}(\mathbf{P}_{sys} - \mathbf{P}_{dia})$$

because at rest the heart spends roughly twice as long in diastole as in systole. **Pulse pressure** is the difference, systolic minus diastolic.

Typical pressures around the circuit, for orientation only:

Location	Approximate pressure (mmHg)
Aorta / large systemic arteries	~120/80
Arterioles	~60 falling to ~30
Capillaries	~35 at the arterial end to ~15 at the venous end
Venae cavae / right atrium	~0–5
Pulmonary artery	~25/10



The pulmonary circuit runs at roughly one fifth the pressure of the systemic circuit. Both receive the same cardiac output. The lungs are a low-resistance, low-pressure, high-compliance bed, and that is why the right ventricle is much thinner-walled than the left.

That is a nice piece of engineering reasoning available from first principles: same flow, much lower pressure, therefore much lower resistance, therefore a very different pump requirement.

Note also the capillary pressures. They are small, and they are the pressures at which exchange happens. An external interface pressure of 20 kPa — which is 150 mmHg, the bench example from Chapter 5 — is several times capillary pressure. That is the mechanism by which sustained external loading reduces tissue perfusion, and it is worth sitting with for a moment: the numbers a maker generates casually are large compared with the numbers the microcirculation works at.

9.3.4. Blood itself

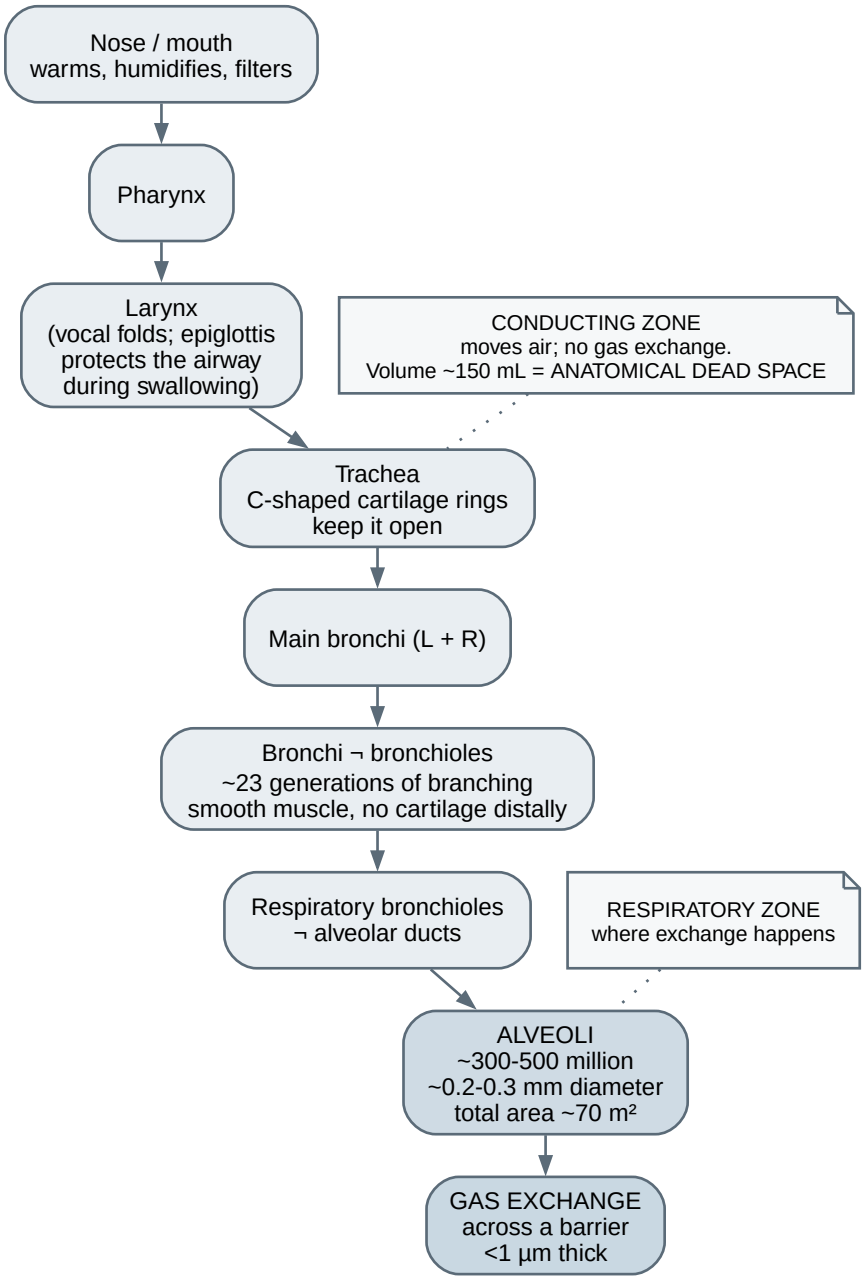
Component	Share	Function
Plasma	~55% of volume	Water, proteins (albumin, clotting factors, antibodies), electrolytes, transport
Erythrocytes (red cells)	~45% (the **haematocrit**)	Oxygen and CO_2 transport via haemoglobin

Component	Share	Function
Leukocytes (white cells)	<1%	Immune defence
Platelets	<1%	Haemostasis

Haemostasis — the stopping of bleeding — proceeds by vascular constriction, a platelet plug, then a fibrin clot produced by an amplifying cascade. It is the positive-feedback example from Chapter 1, and it is why any object that could abrade or puncture skin in someone on anticoagulant medication is a larger concern than it looks.

9.4. Breathing

9.4.1. Anatomy



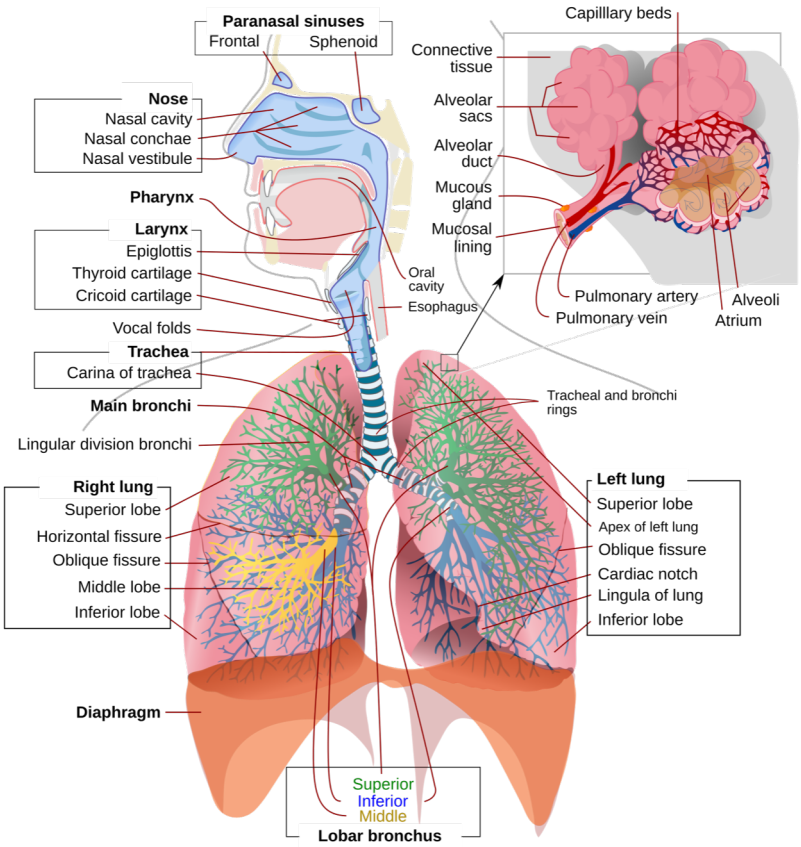
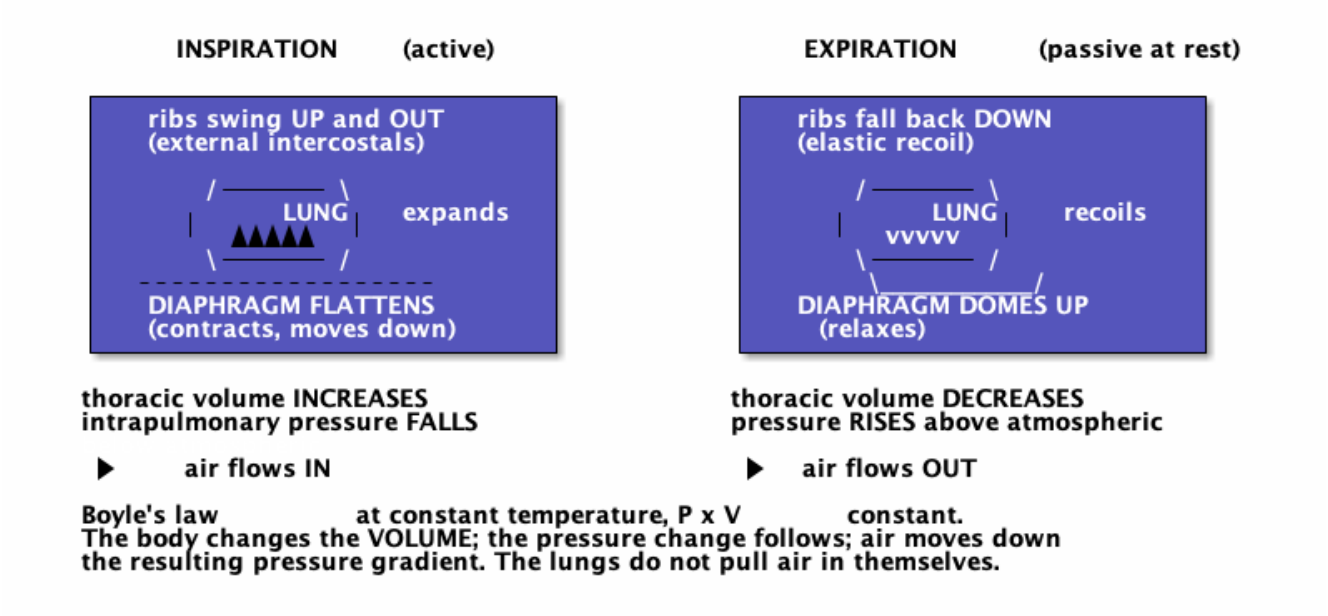


Figure 20. The respiratory system. Image: LadyofHats and Jmarchn, public domain. Credits: Appendix H

The upper part of this path — nose, pharynx, and larynx — is also where the airway crosses the food passage, and it is covered in detail in Chapter 8. This chapter starts at the larynx and works down.

The **conducting zone** — nose to bronchioles — moves air and conditions it but performs no gas exchange. Its volume, around 150 mL in an adult, is **anatomical dead space**. That single fact explains why shallow rapid breathing is inefficient: if tidal volume is 500 mL, 150 mL of each breath never reaches an exchange surface, so 350 mL is useful. Halve the tidal volume to 250 mL and double the rate, and the minute volume is unchanged but useful **alveolar** ventilation collapses from 350 to 100 mL per breath.

9.4.2. Mechanics of ventilation



Quiet inspiration is **active**: the diaphragm contracts and flattens, the external intercostals lift and expand the ribcage, thoracic volume increases, intrapulmonary pressure falls below atmospheric, and air flows in. Quiet expiration is largely **passive**: the muscles relax and the elastic recoil of lung and chest wall does the work. Forced expiration recruits abdominal and internal intercostal muscles.

> **The lungs are not muscles and do not inflate themselves.** They are elastic bags that follow the thoracic cavity because the pleural space couples them to the chest wall. Lose that coupling — air entering the pleural space, a **pneumothorax** — and the lung collapses away from the chest wall even though nothing is wrong with the lung tissue.
>
> This is also why a model that shows a balloon inflating inside a jar is a genuinely good demonstration of **ventilation mechanics** and says nothing at all about gas exchange.

Three mechanical properties that govern breathing, and which an engineer will recognise:

Property	Definition	Clinical relevance
Compliance	Change in volume per unit change in pressure, $\Delta V/\Delta P$	Low compliance (stiff lungs — fibrosis, oedema) means more pressure for the same volume. High compliance with lost recoil (emphysema) makes expiration hard.
Airway resistance	Pressure drop per unit flow	Dominated by the radius of the conducting airways. Bronchoconstriction (asthma) raises it steeply — r^4 again.

Property	Definition	Clinical relevance
Surface tension / surfactant	Fluid lining the alveoli would tend to collapse them (Laplace: $P = 2\gamma/r$)	**Surfactant** reduces surface tension and, crucially, does so more in smaller alveoli, stabilising them against collapse. Its absence in premature infants is respiratory distress syndrome.

That surfactant mechanism is worth a moment. By Laplace's law, a smaller bubble has a higher internal pressure for the same surface tension, so small alveoli should empty into large ones and collapse. Surfactant's concentration-dependent effect lowers tension more in smaller alveoli, equalising pressures and keeping the population stable. It is an elegant solution to a real instability, and it is a good example of a biological mechanism with a clean physical explanation.

9.4.3. Lung volumes

Volume	Typical adult	Definition
Tidal volume (V_T)	~500 mL	Air moved in one quiet breath
Anatomical dead space	~150 mL	Conducting-zone volume; no exchange
Alveolar ventilation per breath	~350 mL	V_T − dead space
Respiratory rate	~12–20 breaths/min	At rest
Minute ventilation	~6 L/min	V_T × rate
Inspiratory reserve volume	~3000 mL	Extra you can inhale beyond tidal
Expiratory reserve volume	~1100 mL	Extra you can exhale beyond tidal
Residual volume	~1200 mL	Cannot be exhaled; keeps alveoli open
Vital capacity	~4600 mL	Maximum exhalable after maximum inhalation
Total lung capacity	~5800 mL	Vital capacity + residual volume

Values are typical young adult male figures from [1] and vary substantially with sex, height, age, and fitness. Spirometry results are reported against predicted values derived from those variables, which is itself an interesting measurement-design problem.

9.4.4. Gas exchange

Ventilation is air movement. **Diffusion** across the alveolar–capillary barrier is gas exchange. **Perfusion** is blood flow through the tissue. All three must be adequate and **matched**.

Term	What it is	Failure mode
Ventilation (V)	Fresh air delivered to alveoli	Obstruction, weak respiratory muscles, poor compliance

Term	What it is	Failure mode
Diffusion	O_2 and CO_2 crossing the barrier down partial-pressure gradients	Thickened barrier (fibrosis, oedema), reduced surface area (emphysema)
Perfusion (Q)	Blood flow past the alveoli	Pulmonary embolism, low cardiac output
V/Q matching	Ventilation and perfusion in the same places	Shunt (perfused, not ventilated) or dead space (ventilated, not perfused)

Approximate partial pressures, for orientation:

Location	P_{O_2} (mmHg)	P_{CO_2} (mmHg)
Dry atmospheric air at sea level	~160	~0.3
Alveolar air	~104	~40
Arterial blood	~95–100	~40
Venous blood (mixed)	~40	~45

Oxygen diffuses from alveolus (104) into blood arriving at 40; CO_2 diffuses from blood (45) into alveolar air (40). Note how small the CO_2 gradient is — CO_2 is far more soluble, so a small gradient suffices.

Most oxygen is carried bound to **haemoglobin**, not dissolved. The **oxygen–haemoglobin dissociation curve** relating saturation to partial pressure is **sigmoid**, which has two useful consequences: in the lungs, the flat upper portion means saturation stays high even if alveolar P_{O_2} drops somewhat; in the tissues, the steep middle portion means a modest fall in P_{O_2} unloads a lot of oxygen. The curve shifts right (unloading more readily) with increased temperature, increased CO_2, and lower pH — exactly the local conditions in actively metabolising tissue.

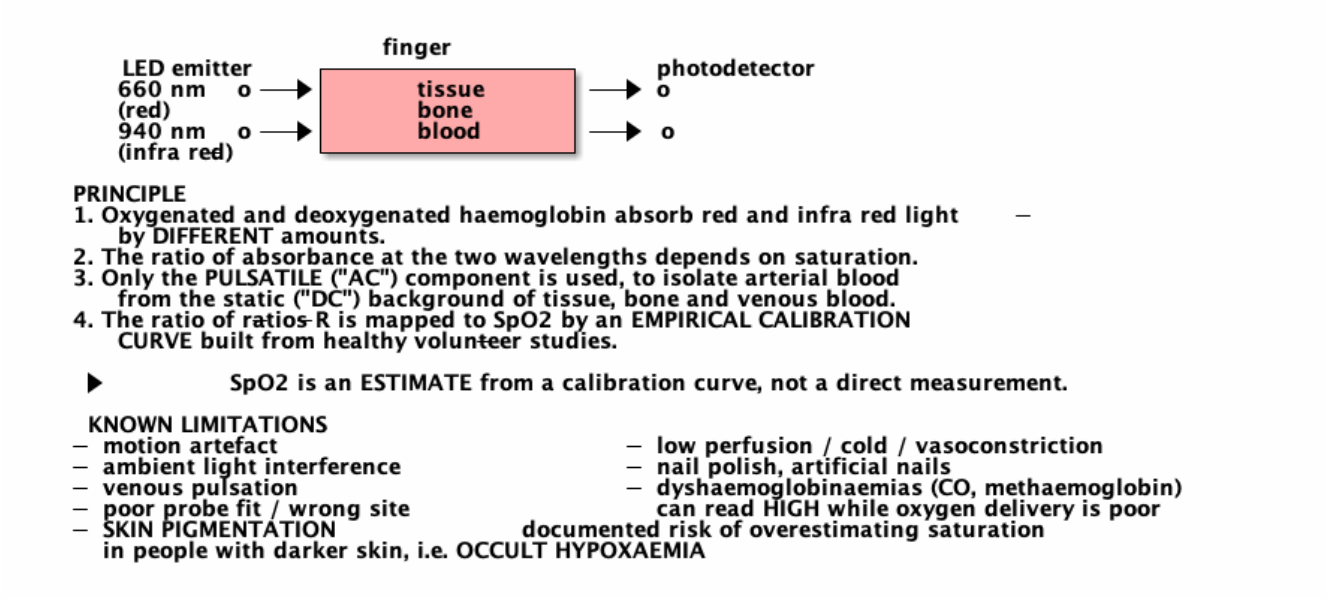
9.5. What the vital signs actually measure

Sign	What it is	How it is obtained	What it is not
Heart rate	Beats per minute	Palpation, ECG, or inferred from a photoplethysmogram	Not a measure of cardiac output, blood pressure, or adequacy of perfusion
Blood pressure	Pressure within the arterial circulation, reported systolic/diastolic	Oscillometric cuff, auscultation, or an invasive arterial line	Not a measure of flow. Pressure can be maintained while flow is inadequate
Respiratory rate	Breaths per minute	Observation, impedance, capnography-derived	Not a measure of tidal volume, alveolar ventilation, or gas exchange

Sign	What it is	How it is obtained	What it is not
Oxygen saturation (SpO□)	**Estimated** proportion of haemoglobin binding sites occupied by oxygen	Pulse oximeter: optical absorbance at two wavelengths	Not a measure of ventilation, not of oxygen **delivery**, not of CO_2, not of arterial oxygen **content**
Temperature	Body temperature at the measured site	Oral, tympanic, axillary, rectal, skin, oesophageal — all different	Not interchangeable between sites; site must be stated
Pain score	A self-reported rating	Numeric or visual scale	Not a physical measurement at all; it is a report, and it is the person's

9.5.1. How a pulse oximeter works, and where it fails

Worth doing properly, because it is the single most instructive measurement device for a maker: simple principle, enormous clinical reliance, well-documented limitations.



The pigmentation issue is the most important thing in this chapter for anyone who writes software.

Pulse oximeters have been shown to **overestimate** arterial oxygen saturation more often in people with darker skin pigmentation, meaning true hypoxaemia can be missed. The FDA has issued safety communications on pulse oximeter accuracy and limitations [27], and this is an active area of regulatory and research attention.

Three lessons generalise far beyond oximetry:

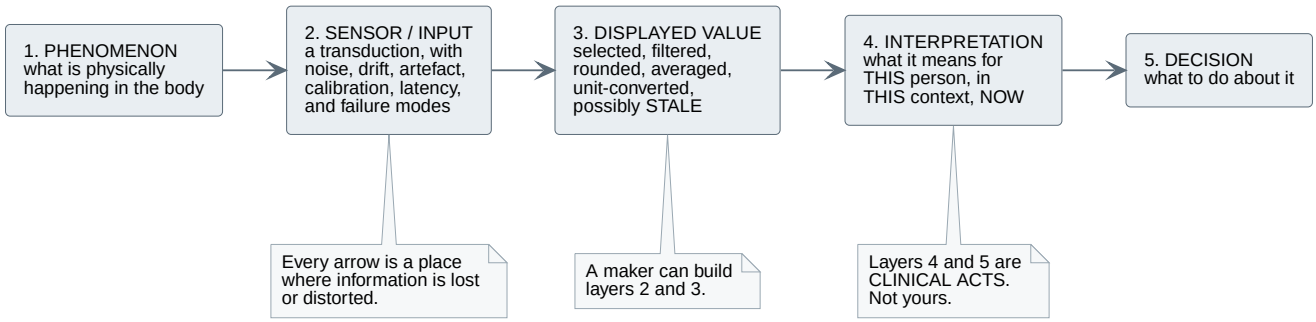
1. **A device calibrated on a non-representative population performs worse on the people it was not calibrated on.** This is the same failure as a machine-learning model trained on a skewed dataset. It is not a new problem

introduced by AI; it is a measurement problem that AI made more visible.

2. **An estimate presented as a measurement conceals its own uncertainty.** "SpO□ 95%" looks like a measurement. It is a calibration-curve output with a confidence interval that the display does not show.
3. **If you build an interface, you inherit responsibility for how its uncertainty is communicated.** Choosing to display a bare number is a design decision with clinical consequences.

9.5.2. The four layers: phenomenon, sensor, display, interpretation

This is the mental model to take away.



Software can look authoritative while its inputs or assumptions are wrong.

A plausible display conceals: a stale value whose input stopped minutes ago; the wrong units; rounding that hides a trend; an average over a window that smooths away the event that mattered; a default that was never configured; a sensor that has detached; and an inference presented as a measurement.

Concrete rules for any physiological display you ever write, even a classroom mock-up:

- Make **simulation status** unmissable, not a footnote.
- Show **units** next to every number, always.
- Show the **age** of the value, and change its appearance when it goes stale. Never repeat the last value as though it were current.
- Distinguish **no data** from **zero**. They are completely different states and conflating them has killed people.
- Make **invalid or out-of-range input** explicit rather than clamping it silently.
- Do not label a value "normal" because it fell inside a range you chose. You are not the person who decides what is normal for this person.
- Make **paused** visually distinct from **live**.
- Handle the **interruption** case — connection lost, app backgrounded, device asleep.

9.6. The other systems: a map, not a textbook

Your goal here is to know where to look, and to recognise when a project has quietly become a project about one of these. The table is the map; the short sections after it give each system enough substance to follow a conversation, and name what a maker must not assume. The kidneys, gut, and hormones are covered properly in Chapter 10.

System	Core function	Key structures	A maker might meet it as
Renal / urinary	Regulates fluid volume, electrolytes, and pH; excretes waste as urine	Kidneys (~1 million **nephrons** each), ureters, bladder, urethra	Catheter management, fluid measurement, dialysis ancillaries, continence aids — all highly specialised
Digestive	Mechanical and chemical breakdown, absorption, elimination	Mouth, pharynx, oesophagus, stomach, small and large intestine, liver, pancreas, gallbladder	Feeding aids, adapted cutlery, enteral feeding ancillaries, stoma appliances
Endocrine	Hormonal regulation over minutes to years	Hypothalamus, pituitary, thyroid, parathyroids, adrenals, pancreatic islets, gonads	Insulin delivery, glucose monitoring — a heavily regulated area
Immune / lymphatic	Defence; interstitial fluid return	Lymph nodes and vessels, spleen, thymus, tonsils, bone marrow, white cells	Infection control (Chapter 13); lymphoedema garments
Reproductive	Reproduction; hormone production	Female: ovaries, uterine tubes, uterus, cervix, vagina, breasts. Male: testes, epididymis, vas deferens, prostate, seminal vesicles, penis	Obstetric and neonatal simulation; pelvic and breast examination trainers; positioning aids
Integumentary	Barrier, sensation, thermoregulation	Skin and appendages	Chapter 5 — the system you meet most often
Nervous	Communication and control	Chapter 7	Accessibility, assistive technology, interfaces
Special senses; head and neck	Vision, hearing, balance, smell, taste; the upper airway, voice, and swallowing	Eye, ear, nose and sinuses, pharynx, larynx, mouth and teeth	Chapter 8 — eyewear, hearing technology, face-worn interfaces, airway trainers

9.6.1. Renal and urinary

Each kidney filters the whole blood plasma many times a day — on the order of 180 L of filtrate daily, of which all but roughly 1–2 L is reabsorbed. The **nephron** is the working unit: a filter (the glomerulus) followed by a long tubule that selectively takes back water, salts, glucose, and

bicarbonate. The kidneys therefore set the body's water balance, its sodium and potassium, and — with the lungs — its pH (Chapter 3). They also make hormones that control red cell production and help regulate blood pressure.

For a maker: **urine output is a vital measurement**, often recorded hourly in acutely ill people, so anything that touches collection or measurement — bags, meters, holders — sits inside a clinical data chain. Urinary catheters are the commonest source of healthcare-associated infection in many settings (Chapter 13). Kidney failure is treated by **dialysis**, which is mass transfer across a membrane — a chemical engineering problem with a person attached, and firmly Tier 3.

9.6.2. Digestive

The gut is a tube from mouth to anus — several metres long — with glands opening into it. The **mouth and pharynx** (Chapter 8) start the process and manage the swallow. The **stomach** stores and acidifies (pH ~1.5–3.5). The **small intestine**, with its folded, villus-covered lining, does nearly all of the absorption; its surface area is enormous for its length. The **large intestine** recovers water and salts. The **liver** processes absorbed nutrients and most drugs, makes bile and clotting proteins, and clears toxins; the **pancreas** supplies digestive enzymes and bicarbonate as well as insulin.

For a maker: feeding is a daily-living task with real demand for adaptive design — cutlery, plate guards, cups that do not require neck extension — and it is also where swallowing safety (Chapter 8) and texture-modified diets apply. **Enteral feeding** through a tube uses connectors deliberately designed so they **cannot** be mated with intravenous lines — the ISO 80369 series of small-bore connectors exists because misconnections killed people. That is design-to-eliminate (Chapter 7) written into a standard. **Stoma appliances** are a skin-adhesion problem (Chapter 5) on a site that changes shape.

9.6.3. Endocrine

Hormones are the slow, broadcast channel of control: released into the blood, acting on any cell with the right receptor (Chapter 3), over minutes to years. The hypothalamus and pituitary at the base of the brain direct much of it. The thyroid sets metabolic rate; the adrenals produce cortisol and adrenaline; the pancreatic islets produce **insulin** and **glucagon**, the opposed pair that holds blood glucose in range — negative feedback in its clearest form (Chapter 1).

For a maker: **diabetes technology** — glucose meters, continuous glucose monitors, insulin pumps, and the software that links them — is the largest and most visible consumer-facing medical-device market, with a large and expert open-source user community. It is also a domain where a unit error (mmol/L versus mg/dL, Appendix C) or a dosing bug can kill. Treat it as Tier 3 for anything that computes or delivers a dose.

9.6.4. Immune and lymphatic

The **lymphatic system** returns the ~2–3 L per day of fluid that leaks out of capillaries and is not reabsorbed, passing it through **lymph nodes** — filters where immune cells sample what the fluid carries. When lymphatic drainage is damaged (after cancer surgery or radiotherapy, for example) fluid accumulates: **lymphoedema**, managed with compression garments and bandaging, which are a pressure-and-fit problem of exactly the kind in Chapter 5.

The **immune system** has two layers. **Innate** immunity — barriers, inflammation, and phagocytic cells — responds within minutes to hours to anything foreign. **Adaptive** immunity — lymphocytes and antibodies — learns a specific target over days and remembers it. The materials questions of Chapter 3 and Chapter 13 are immune questions underneath: a surface the body sees as foreign provokes inflammation, and a surface bacteria can colonise is a surface the immune system will struggle to clear.

9.6.5. Reproductive

The reproductive organs produce gametes and sex hormones; in women they also support pregnancy and birth. You are unlikely to design for this system directly, but it is prominent in **simulation**: obstetric and birthing simulators, neonatal manikins, pelvic, breast, and prostate examination trainers, and cervical screening trainers are core teaching tools, and their replaceable parts need making. They also raise particular requirements for **dignity, consent, and inclusive design** — in who the trainer represents, and in how real examinations are taught — that the educator, not the maker, sets.

Pregnancy changes almost every system in this chapter: blood volume rises by around 40–50%, cardiac output rises, breathing changes as the diaphragm is displaced, and posture and balance shift. A device or positioning aid validated on non-pregnant adults has not been validated on pregnant ones.

Two extra concepts worth having, because they cut across systems:

Fluid compartments. Body water (~50–60% of mass) divides into **intracellular** (~2/3) and **extracellular** (~1/3), and the extracellular into **interstitial** fluid and **plasma**. Oedema is excess interstitial fluid, and it is a dimensional problem for anything worn.

Thermoregulation. Heat is produced by metabolism and lost by radiation, convection, conduction, and evaporation. Skin blood flow and sweating are the main effectors (Chapter 1). Anything that covers skin reduces evaporative and convective loss and traps moisture — an enclosure is a thermal decision.

9.7. Devices, software, and training technology

This is the densest device chapter in medicine. Cardiology, respiratory medicine, anaesthesia, and critical care between them account for a large share of all medical equipment — and almost all of it is life-supporting, patient-connected, or both.

Read this section as a **map**, not an invitation. Almost nothing here is a maker project. The training column is where you can contribute.

9.7.1. Clinical devices

Device	Does	Engineering note
Stethoscope	Acoustic auscultation of heart, lungs, and bowel	Still the emblem of clinical medicine. Electronic versions add amplification, filtering, and recording

Device	Does	Engineering note
ECG machine — 12-lead and telemetry	Records cardiac electrical activity	~0.5–4 mV, 0.05–150 Hz. Electrode placement errors are a common and consequential fault
Pulse oximeter	Estimates arterial oxygen saturation optically	Covered in full in this chapter, including its documented accuracy limitations
Blood pressure monitor — oscillometric, auscultatory, invasive arterial line	Measures arterial pressure	Cuff size error is a large and under-appreciated source of bias
Defibrillator / AED	Delivers a controlled shock to terminate certain arrhythmias	Energy delivery, impedance compensation, rhythm analysis. AEDs are designed for untrained users — a human factors triumph worth studying
Pacemaker and ICD	Implanted rhythm management	Lead design, battery life, telemetry, MRI conditionality
Mechanical ventilator	Supports or replaces breathing	Modes, pressure and volume targets, triggering, alarms, and patient–ventilator synchrony. Among the most safety-critical devices in existence
CPAP and NIV	Non-invasive pressure support, for sleep apnoea and respiratory failure	**Mask interface fit is the usual failure point** — and it is a Chapter 5 contact-design problem
Oxygen delivery — nasal cannulae, masks, high-flow systems	Delivers supplemental oxygen at a controlled concentration	Flow, FiO_2, humidification. High-flow nasal oxygen changed respiratory care
Capnograph	Measures exhaled CO_2 continuously	Infrared absorption. The waveform carries more information than the number
Spirometer	Measures lung volumes and flows	Results are reported against predicted values derived from age, sex, height, and ethnicity — the ethnicity correction is actively contested
Nebuliser and inhaler	Delivers drug to the airway	Particle size determines deposition site. Inhaler technique is a major real-world problem, and a training opportunity

Device	Does	Engineering note
Infusion pump and syringe driver	Delivers fluid or drug at a controlled rate	Dose error reduction software, drug libraries, free-flow protection (Chapter 3)
ECMO and cardiopulmonary bypass	Takes over gas exchange and circulation	Pumps, oxygenators, anticoagulation, and heat exchange. Specialist perfusion practice
Suction, chest drains, airway devices	Clear secretions, drain air or fluid, maintain an airway	Simple in principle, unforgiving in use

9.7.2. Software

Software	Does	Note
Patient monitoring and central telemetry	Displays and distributes continuous physiological data	**Alarm design and alarm fatigue are the dominant human factors problem** in critical care. IEC 60601-1-8 governs alarm signals
ECG interpretation software	Automated rhythm and morphology analysis	Always reported as **provisional**, pending clinician over-read. A good example of inference-not-decision
Early warning scores — NEWS2 and equivalents	Aggregate vital signs into an escalation trigger	The inference layer, formalised. The score prompts a human; it does not decide
Ventilator waveform analysis	Detects asynchrony and abnormal mechanics	Teaching tools exist and are genuinely useful
Cardiac imaging analysis	Echo quantification, cardiac MRI analysis	See Chapter 15
Sleep study analysis	Scores respiratory events and sleep stages from polysomnography	Partially automated, clinician-reviewed
Open physiological datasets and toolkits — PhysioNet, WFDB, NeuroKit2	Published ECG, EEG, PPG, and respiratory recordings, with analysis libraries	**Exactly where to learn this material if you write signal-processing software.** Real signals, no patient contact, no device to build
Resuscitation and ALS decision aids	Algorithm prompts during a cardiac arrest	Cognitive aids under extreme time pressure — a human factors specialism

9.7.3. Training and simulation

Training technology	Teaches	Form
CPR manikins with feedback	Compression depth, rate, recoil, ventilation volume	Sensors plus real-time feedback measurably improve performance. **Consumables — airways, lungs, face shields — are constantly needed**
AED trainers	Pad placement and the shock sequence, with no actual energy delivery	Deliberately inert. A good example of designed-in safety
Full-body patient simulators — SimMan, HAL, and similar	Integrated scenarios: airway, breathing, circulation, drugs, team working	Programmable physiology, palpable pulses, chest rise. Very expensive; **spare parts and consumables are a genuine maker niche**
Airway management trainers	Bag-mask ventilation, supraglottic airways, intubation	Head-and-torso trainers; difficult-airway variants
Auscultation trainers	Recognising heart and lung sounds	Manikins with speakers, plus phone apps and recorded sound libraries. **Free sound libraries exist**
ECG rhythm simulators	Rhythm recognition and response	Software and hardware generators feeding a real monitor
Ventilator simulators and test lungs	Modes, settings, waveform interpretation, troubleshooting	Mechanical test lungs with settable compliance and resistance. **A tractable making project** — a bellows with interchangeable springs and orifices demonstrates compliance and resistance honestly
Central line and chest drain trainers	Insertion technique, often ultrasound-guided	Layered torso sections; consumable inserts
Inhaler technique trainers	Correct inhaler use — a large real-world problem	Whistle-feedback devices and placebo inhalers
Perfusion and bypass simulators	Circuit management for perfusionists	Highly specialised
Moulage and scenario props	Realism and context for team training	**Low-risk, high-demand making work** — silicone and wax casting, fabrics and dressings, and simple electronics for bleeding pumps or sound cues



This chapter's devices are where a maker must be most careful.

Ventilators, defibrillators, infusion pumps, oxygen delivery, and monitoring are **life-supporting or life-sustaining**. They are heavily regulated for good reason. The improvised-ventilator projects of 2020 are instructive: enormous effort, very few devices that were ever clinically usable, and a great deal of expert time spent explaining why.

The useful contributions here are **training aids, consumables, spare parts, props, and software that touches recorded or synthetic data**. That list is not a consolation prize — simulation centres need all of it, constantly, and cannot easily get it.



The mechanical test lung is an unusually good project.

It is a bellows with a settable spring (compliance) and an interchangeable orifice (resistance). It connects a ventilator to something that behaves plausibly, it teaches the two mechanical properties in this chapter honestly, it touches no patient, and it needs no electronics. It is also a genuinely mixed-workshop build: a sewn or moulded bellows, turned or drilled orifice plates of known diameter, sheet-metal or machined end plates, and springs selected by measured rate. Add a pressure sensor and a microcontroller later, if a display is wanted, and the electronics stays outside the gas path.

It also makes the physics tangible in a way the equations do not: change the spring and the pressure needed for the same volume changes; narrow the orifice and the flow resistance rises steeply, exactly as the r^4 relationship predicts.

9.8. Worked example: the classroom mock monitor

An educator wants a fake patient monitor for a teaching scenario. You can build this in an afternoon. Here is how to build it **responsibly**.

Establish first, before writing code:

- What is the **learning objective**? (The educator defines this, not you — Chapter 21.)
- Which values must be shown, and must they **change**, or is a static display enough?
- Does the scenario require the numbers to behave in a clinically meaningful way? If so, **the educator specifies the behaviour**. You are not qualified to decide what a deteriorating patient's numbers do.
- Where will it run, who can see it, and **could it ever be mistaken for a real monitor**?

Build it with:

- A permanent, high-contrast **"SIMULATION ONLY — NOT A MEDICAL DEVICE"** banner that cannot be dismissed, hidden, or scrolled off. Not a footer. Not a splash screen.
- Obviously fictional patient data. No real name, no real identifier, nothing derived from a real

record.

- Units beside every value.
- An explicit visual difference between **live**, **paused**, and **stale**.
- No "normal" or "abnormal" indicator unless the educator defined the criteria and takes responsibility for them.
- A visible version number, so a result can be traced to the build that produced it (Chapter 19).

Do not:

- connect it to any real device, sensor, or clinical system;
- use real patient data, even de-identified, without the host's formal information governance approval;
- implement alarms, thresholds, or trend interpretation of your own devising;
- deploy it anywhere a passer-by could read it as real;
- let it run on a device that also displays real clinical information.

Note how little of that is programming. Almost all the engineering here is in the **framing**.

9.9. Workshop exercise

1. **Draw both circuits from memory.** The full loop: body → venae cavae → right atrium → right ventricle → pulmonary arteries → lungs → pulmonary veins → left atrium → left ventricle → aorta → body. Mark where exchange happens. Check you have not connected the aorta to the right ventricle.
2. **Draw the air path.** Nose to alveolus, marking the boundary between conducting and respiratory zones. Then calculate alveolar ventilation per breath for V_T = 500 mL and again for V_T = 250 mL at double the rate, and write one sentence on what the comparison shows.
3. **Calculate.** MAP for 118/76. Cardiac output for HR 72 and SV 68 mL. The flow reduction if a vessel's radius falls by 20%, by Poiseuille.
4. **Mock monitor, if software is your thing.** Build the display above. Deliberately implement **one** misleading behaviour — a value that persists after its input stops. Then fix it, and write down what the fix changes about what a viewer can conclude.
5. **Trace a strap.** Take any worn object you own. Identify where it crosses a point at which an artery or nerve runs superficially — the medial elbow, the fibular neck at the outer knee, the wrist. Then write the question you would ask a clinician about it. Do not experiment.

9.10. Questions

1. Why is "arteries contain oxygenated blood" an unreliable definition?
2. Trace blood from the right ventricle to the aorta.
3. Can an inflating lung model prove that gas exchange is adequate?
4. A display repeats the last reading after its input stops. What information is missing, and why

does it matter?

5. Cardiac output is 5 L/min at HR 70. What is stroke volume?
6. A vessel's radius halves. By what factor does flow change, all else equal? Why does that matter for a strap?
7. Why is shallow rapid breathing inefficient?
8. What does SpO□ actually estimate, and name three things it does not tell you?
9. Why does a pulse oximeter use only the pulsatile component of the optical signal?
10. Why might a pulse oximeter overestimate saturation, and what is the general lesson?
11. The pulmonary circuit carries the same cardiac output as the systemic circuit at about a fifth the pressure. What does that tell you about pulmonary resistance, and about the right ventricle?
12. A patient's blood pressure is maintained. Is their tissue perfusion adequate?

9.11. Answers and reasoning

1. Arteries are defined by carrying blood **away from the heart**, not by oxygen content. The pulmonary arteries carry relatively oxygen-poor blood toward the lungs, and the pulmonary veins carry oxygen-rich blood back — the standard counterexample.
2. Right ventricle → pulmonary valve → pulmonary arteries → lungs (gas exchange in the alveolar capillaries) → pulmonary veins → left atrium → mitral valve → left ventricle → aortic valve → aorta.
3. **No.** Inflation demonstrates **ventilation mechanics** — volume change driving airflow down a pressure gradient. Gas exchange additionally requires diffusion across an intact alveolar–capillary barrier and adequate **perfusion**, matched to ventilation. A model can inflate beautifully with no exchange modelled at all.
4. The **age** of the value and the **status of the input**. It matters because a retained value is visually indistinguishable from a current one, so a viewer will reasonably conclude the patient's state is being observed when it is not. "No data" and "unchanged" are entirely different clinical situations.
5. SV = CO / HR = 5000 mL/min ÷ 70 ≈ **71 mL per beat**.
6. Flow falls to **one sixteenth** (r^4). For a strap: a modest reduction in vessel radius produces a steeply disproportionate reduction in flow, so compression is a far more powerful intervention than its gentleness suggests. "It's only a little tight" is not a reassuring statement about flow.
7. Because each breath wastes the ~150 mL anatomical dead space. At V_T 500 mL, 350 mL per breath reaches alveoli; at V_T 250 mL, only 100 mL does. Doubling the rate keeps minute ventilation the same, but useful alveolar ventilation falls from 350 to 100 mL per breath — at twice the rate, 200 mL in the time one deeper breath delivered 350 mL. That is a large loss for identical measured effort.
8. It estimates the **proportion of haemoglobin binding sites occupied by oxygen in arterial blood**, via a two-wavelength absorbance ratio mapped through an empirical calibration curve. It does **not** tell you: ventilation adequacy; CO_2 level; oxygen **delivery** (which also needs haemoglobin concentration and cardiac output); total arterial oxygen content; or whether the haemoglobin present is capable of carrying oxygen at all (carbon monoxide poisoning can

produce a reassuring reading).

9. To isolate the **arterial** contribution. The static component — tissue, bone, venous blood, pigment — does not pulse; the arterial volume does. Using only the pulsatile component removes most of the static background. It is also why low perfusion and motion both degrade the measurement: the first shrinks the signal, the second counterfeits it.
10. Because the calibration curve was derived from study populations, and absorbance is affected by skin pigmentation in ways that calibration on a non-representative sample does not capture — producing a documented risk of **occult hypoxaemia** in people with darker skin [27]. The general lesson: **a device calibrated on a non-representative population performs worse on everyone it was not calibrated on**, and an estimate displayed as a bare number hides its own uncertainty.
11. Since $\Delta P = Q \cdot R$ and flow is the same while pressure drop is about five times lower, **pulmonary vascular resistance is roughly five times lower** than systemic. Consequently the right ventricle needs to generate far less pressure, which is why its wall is much thinner than the left ventricle's.
12. **Unknown from that fact alone.** Pressure and flow are different quantities. Perfusion can be inadequate at a maintained pressure — for example if resistance is high, output is low, or distribution is abnormal. This is also the general form of the trap: a reassuring number on one variable is not reassurance about a different variable. And the judgement is clinical, not yours.

9.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Heart, vessels and oxygen — cardiac output, MAP, resistance, Poiseuille radius, gases and the O curve.
- Test lung and breathing mechanics — lung volumes, dead space, airway radius and the equation of motion.
- Fluids — laminar flow and its r sensitivity.
- Biosignal analysis — the pulse-oximeter ratio of ratios, shown as a range.

9.13. Video companion and reading

- **Core videos:** Crash Course Anatomy & Physiology #25 (the heart) [12] and #31 (the respiratory system) [13]. Episodes 27 and 32 add useful detail.
- **Core reading:** OpenStax [1]:
  - 19.1 Heart Anatomy and 19.3 Cardiac Cycle
  - 20.1 Structure and Function of Blood Vessels and 20.2 Blood Flow, Blood Pressure, and Resistance
  - 22.1 Organs and Structures of the Respiratory System, 22.3 The Process of Breathing, 22.4 Gas Exchange
  - Skim chapter summaries for 17 (endocrine), 21 (immune/lymphatic), 23 (digestive), and 25

(urinary) **for orientation only**. Return to any of them in depth only when a project requires it.

- **Measurement:** FDA information on pulse oximeter accuracy and limitations [27]. Read this one properly — it is a short, concrete, real example of everything this chapter is about.

9.14. Exit criteria

- ▢ Trace circulation through the heart and lungs from memory without reversing any major connection.
- ▢ Explain why arteries are not defined by oxygen content, with the pulmonary counterexample.
- ▢ Distinguish ventilation, diffusion, and perfusion, and explain what V/Q mismatch means.
- ▢ Calculate MAP, cardiac output, and a Poiseuille radius effect, and state the assumptions each calculation makes.
- ▢ Explain dead space and why shallow rapid breathing is inefficient, with numbers.
- ▢ Describe what each of the five vital signs represents **and at least two things each does not**, without claiming to diagnose from any of them.
- ▢ Explain the pulse oximeter principle and name five of its limitations, including the pigmentation issue and its general lesson.
- ▢ Identify a misleading state in a fictional monitor and propose clearer feedback.
- ▢ Name the four layers — phenomenon, sensor, display, interpretation — and state which two are clinical acts.

9.15. What this chapter does not establish

- **You cannot interpret any measurement.** You know what the numbers represent. What they mean for a person is a clinical act, and so is deciding what to do about it.
- **You cannot set alarm limits or thresholds.** These are clinical and regulatory decisions with a substantial human factors literature behind them, including alarm fatigue as a recognised hazard.
- **You cannot build a monitoring device.** Breathing circuits, medication delivery, infusion, airway devices, and clinical monitoring are specialised regulated projects requiring expert oversight from the outset. They are not scaled-up versions of a classroom mock-up.
- **A model that demonstrates mechanics does not demonstrate physiology.** Inflation shows ventilation. Exchange is a separate process and needs separate representation and separate labelling.
- **The physiological numbers here are orientation, not reference ranges.** They let you check an order of magnitude. They do not let you assess anybody.

Chapter 10. Kidneys, gut, and hormones: the internal regulators

Week 13 · approximately 4–6 hours

10.1. Why this chapter exists

Chapter 9 ended with a one-page map of the systems this book had not yet opened: kidneys, gut, hormones, immunity, reproduction. A map is enough to know where you are. It is not enough to follow a ward round, to understand why a drainage bag must hang where it hangs, or to see why a single unit error in diabetes software can kill.

This chapter opens three of those systems — the **kidneys and urinary tract**, the **digestive system**, and the **endocrine (hormone) system** — because they have something in common that matters to an engineer. They are the body's **slow regulators**. The nervous system of Chapter 7 acts in milliseconds. These systems act over minutes, hours, and days, keeping the internal environment of Chapter 1 inside the narrow ranges that cells tolerate: water, salt, potassium, acid, glucose, calcium.

You will meet them in three ways:

- **As a source of data.** Urine output, fluid balance, blood glucose, kidney function, and hormone levels are some of the most frequently recorded numbers in medicine. Every one of them comes with units that differ between countries and laboratories.
- **As territory full of everyday devices.** Urinary catheters and drainage bags, continence products, feeding tubes and their connectors, stoma appliances, glucose meters, continuous glucose monitors, and insulin pumps are used by millions of people at home, not just in hospital.
- **As control systems.** Hormonal regulation is feedback control with gain, delay, and a set point. Automated insulin delivery is literally a control engineer's problem, and a community of patients and developers built working closed-loop systems in open source before the first commercial one was approved.

10.2. The kidneys and urinary tract

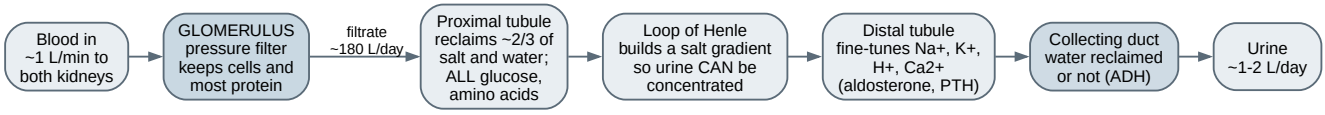
10.2.1. Anatomy

The two **kidneys** lie against the back wall of the abdomen, roughly at the level of the lowest ribs, each about 11–12 cm long and weighing around 150 g. Together they receive roughly **20–25% of the cardiac output** — about a litre of blood a minute — far more than their own metabolism needs. That flow is there to be **processed**.

From each kidney a **ureter** — a muscular tube about 25–30 cm long — carries urine down to the **bladder**, using waves of peristalsis rather than gravity alone. The bladder is a stretchy muscular bag (the **detrusor** muscle) that stores urine until a convenient moment. The **urethra** carries it out: about 4 cm long in women and about 20 cm in men, a difference that shapes infection risk, catheter design, and continence products.

10.2.2. The nephron: a filter followed by a refinery

Each kidney contains around a million **nephrons**. Each nephron does the same job in two stages.



1. **Filtration.** At the **glomerulus**, a knot of capillaries inside a cup-shaped capsule, blood pressure pushes water and small solutes through a three-layer filter. Cells and most proteins stay behind. The rate of filtration from both kidneys together is the **glomerular filtration rate** (GFR) — about 120 mL/min in a healthy young adult, roughly **180 litres a day**. Your entire plasma volume is filtered about sixty times a day.
2. **Reabsorption and secretion.** That filtrate contains everything the body needs — water, sodium, glucose, amino acids, bicarbonate — so the tubule spends most of its energy taking it back. More than 99% of the filtered water is reclaimed. Meanwhile the tubule **secretes** some substances into the urine: hydrogen ions, potassium, and many drugs.

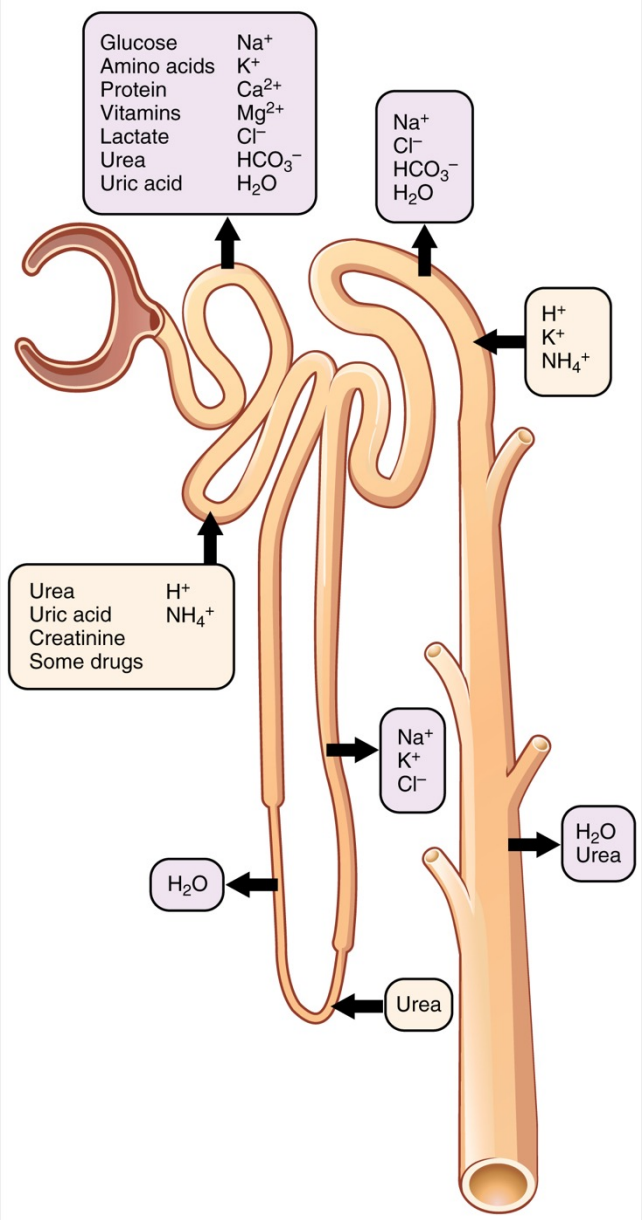


Figure 21. What the nephron reabsorbs (arrows outward) and secretes (arrows inward) along its length. Image: OpenStax, CC BY 3.0. Credits: Appendix H

The design looks wasteful — filter everything, then laboriously take most of it back. It is in fact robust: anything the body has **not** specifically decided to keep is removed by default, including toxins and drug metabolites the kidney has never encountered. A whitelist, not a blacklist.

10.2.3. What the kidneys regulate

Quantity	How	Main controls
Water	Varying how much water the collecting duct reclaims. Urine can range from very dilute (~50 mOsm/kg) to about four times as concentrated as plasma (~1200 mOsm/kg)	**Antidiuretic hormone** (ADH, vasopressin) from the posterior pituitary; thirst

Quantity	How	Main controls
Sodium and blood volume	Sodium reabsorption; water follows sodium	**Renin–angiotensin–aldosterone system**: low pressure at the kidney releases renin, which leads to angiotensin II (constricts vessels) and aldosterone (retains sodium)
Potassium	Secretion in the distal nephron	Aldosterone; plasma potassium itself. Both too high and too low disturb heart rhythm (Chapter 7' resting potential)
Acid–base	Reclaiming filtered bicarbonate; excreting roughly 50–100 mmol of acid a day	With the lungs, which remove CO_2 (Chapter 3)
Calcium and phosphate	Reabsorption; activating vitamin D	Parathyroid hormone (below)
Red cell production	Releasing **erythropoietin** when oxygen delivery to the kidney falls	The reason kidney failure causes anaemia
Blood pressure	Volume, plus renin–angiotensin	Long-term blood pressure is largely a kidney decision

The kidneys must excrete a minimum of roughly 0.4–0.5 L of urine a day to clear the day's waste solutes at maximum concentration. That is why **urine output** is watched so closely in unwell people: falling output can be the first sign that the kidneys are short of blood flow, and it is often recorded hourly from a catheter with a measuring chamber.

10.2.4. Measuring kidney function

GFR cannot be measured directly at the bedside, so it is **estimated** — usually from the blood level of **creatinine**, a waste product of muscle that is filtered and barely reabsorbed. If filtration halves, creatinine roughly doubles. Laboratories report an **estimated GFR** (eGFR) calculated from creatinine, age, and sex with a published equation, and chronic kidney disease is staged by it:

Stage	eGFR (mL/min/1.73 m²)	Description
G1	≥ 90	Normal filtration (other evidence of kidney damage needed to call it disease)
G2	60–89	Mildly reduced
G3a / G3b	45–59 / 30–44	Moderately reduced
G4	15–29	Severely reduced
G5	< 15	Kidney failure; dialysis or transplant may be needed

Staging after KDIGO [52]. Note the strange unit: GFR is normalised to a body surface area of 1.73 m²,

which is itself an estimate. Creatinine also depends on muscle mass, so the estimate is least reliable at the extremes — in frail older people, amputees, body builders, and children (Chapter 12). eGFR is also used to **adjust drug doses**, which is why a kidney-function number feeds directly into prescribing software (Chapter 20).

10.2.5. The bladder and continence

The bladder fills quietly; stretch receptors report its volume. A first urge typically comes at around 150–250 mL, and a typical adult capacity is roughly 400–600 mL. Emptying is a reflex — the detrusor contracts while the sphincters relax — coordinated by a centre in the brainstem and **permitted or postponed** by the cortex. Two sphincters guard the outlet: an internal one of smooth muscle, not under voluntary control, and an external one of skeletal muscle that is. The **pelvic floor** muscles support the whole arrangement.

Continence therefore depends on muscles, nerves, the brain, and — often forgotten — the ability to **get to a toilet in time and manage clothing**. The common types of incontinence map onto those components:

Type	Mechanism	Where design can help
Stress	Weak sphincter or pelvic floor; leakage on coughing, lifting, laughing	Pelvic floor training aids; pads; supportive devices
Urge	Overactive detrusor; sudden urgency	Toilet access, signage, timing aids, reminders
Overflow	Bladder cannot empty (obstruction, weak detrusor, nerve damage); constant dribbling	Catheterisation, often intermittent
Functional	The urinary system works, but mobility, dexterity, vision, or cognition prevent getting there in time	**This is a design problem.** Clothing fastenings, toilet height and rails, way-finding, bedside urinals usable with one hand (Chapter 7, Chapter 6)

Incontinence is common, under-reported because of embarrassment, and a major cause of skin damage — moisture plus friction (Chapter 5) — and of falls, as people hurry to the toilet at night.

10.2.6. Catheters

A **urinary catheter** drains the bladder through a tube.

- An **indwelling urethral catheter** (often called a Foley) is held in the bladder by a small water-filled balloon and drains continuously into a bag.
- **Intermittent self-catheterisation** uses a fresh catheter several times a day and removes it each time — a task requiring dexterity, vision, and privacy, and a frequent subject of adaptive aids.
- A **suprapubic catheter** enters the bladder through the lower abdominal wall.
- **External catheters** — sheaths for men, wicking devices for women — collect urine without entering the body.

Indwelling catheters are among the commonest causes of healthcare-associated infection: the tube is a highway for bacteria, and organisms form biofilms on its surfaces (Chapter 13). The design rules for the drainage system follow directly from that: keep it **closed**, keep the bag **below the bladder** so urine cannot flow back, keep it **off the floor**, keep the tubing **free of kinks**, and empty the bag without touching its outlet to anything.

10.3. The digestive system

10.3.1. The tube and its glands

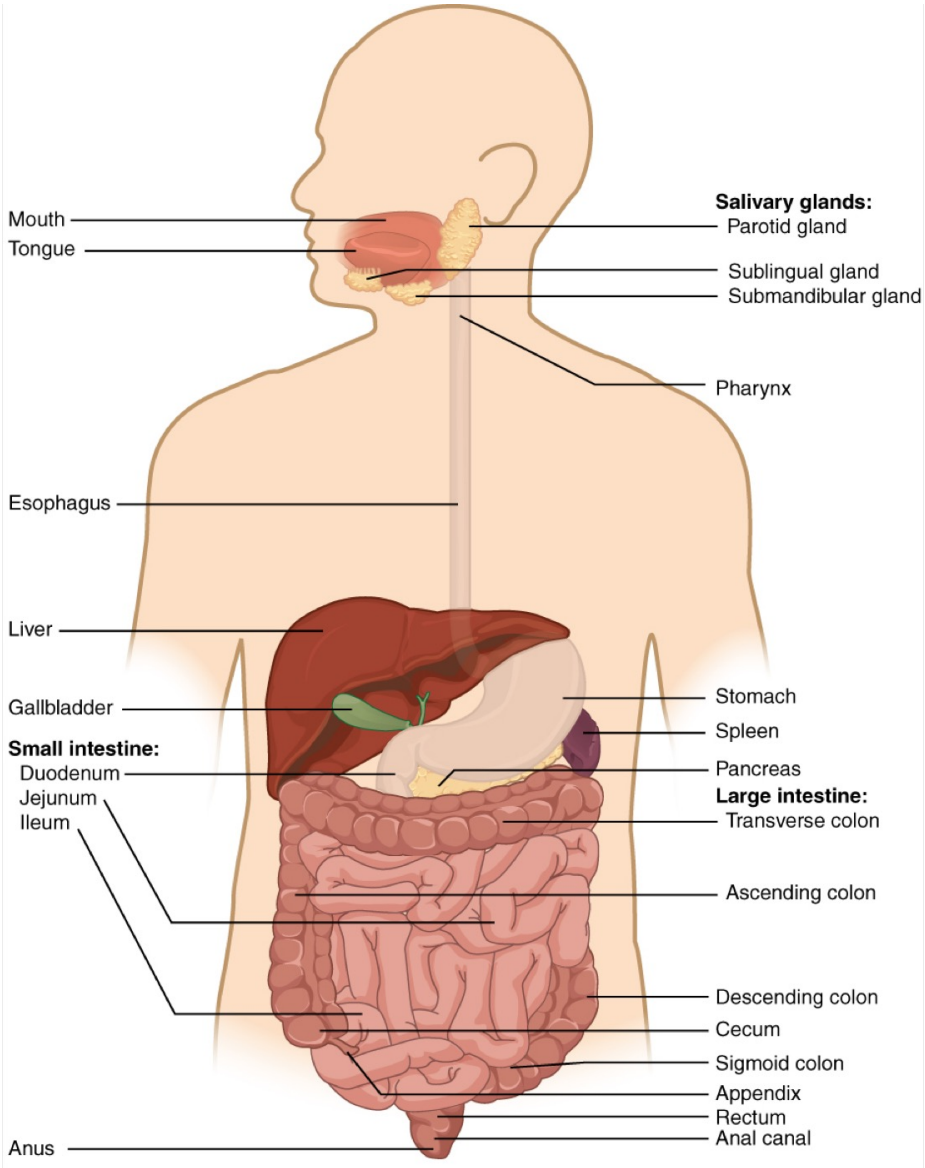
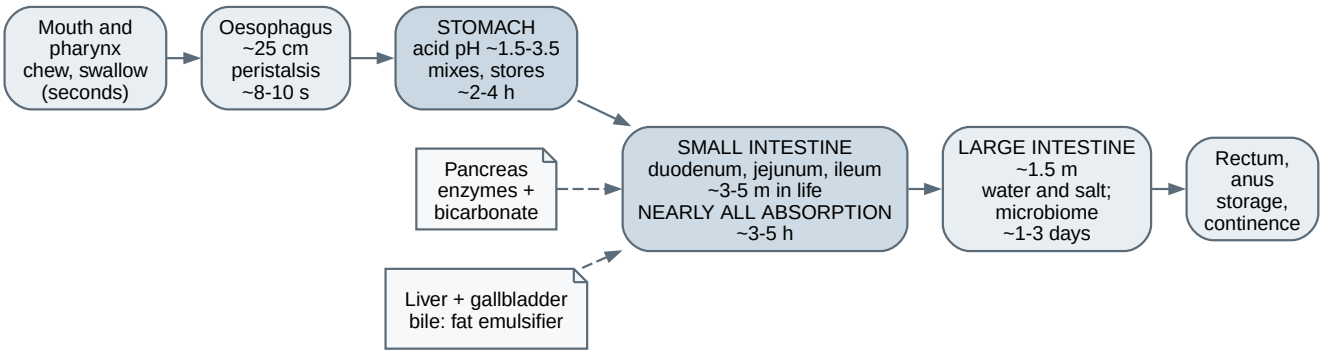


Figure 22. The digestive system. Image: OpenStax, CC BY 3.0. Credits: Appendix H

The digestive system is a tube — mouth, pharynx, oesophagus, stomach, small intestine, large intestine, rectum, anus — plus glands that empty into it: the salivary glands, liver and gallbladder, and pancreas. Food inside the tube is, strictly, still **outside** the body. Only what crosses the lining is absorbed.



Transit times vary widely between people and from day to day; mouth-to-anus commonly takes one to three days. The figures above are orientation only.

10.3.2. Stomach

The stomach is a muscular reservoir that can expand from about 50 mL empty to more than a litre after a meal. Its lining secretes hydrochloric acid (pH roughly 1.5–3.5), which kills many swallowed microbes and starts protein digestion, together with enzymes and a protective layer of mucus that keeps it from digesting itself. It churns food into a semi-liquid **chyme** and releases it into the small intestine a little at a time through the **pyloric sphincter**, over a few hours.

For a device designer the stomach matters as the destination of feeding tubes, as an acid environment that degrades many materials and drugs, and as the reason that **the position of a nasogastric tube must be confirmed before anything is put down it**: a tube that has slipped into the airway instead (Chapter 8) delivers feed into the lungs.

10.3.3. Small intestine

Nearly all digestion and absorption happens here. The first part, the **duodenum**, receives bile from the liver and gallbladder and enzyme-rich, bicarbonate-rich juice from the pancreas, which neutralises the stomach acid. The **jejunum** and **ileum** absorb sugars, amino acids, fats, vitamins, minerals, and most of the water. The final part of the ileum is the only place vitamin B_{12} and bile salts are absorbed — so removing it surgically has specific, predictable consequences.

The lining is folded, the folds are covered in finger-like **villi**, and the cells on the villi carry **microvilli**: three levels of folding to increase surface area. How much area is a good lesson in numbers that become "facts". Older textbooks quoted 200–300 m² — "a tennis court". A careful re-measurement published in 2014 put the whole digestive tract's absorptive area at roughly 30 m². Both figures appear in good sources. When a number is famous, check where it came from.

10.3.4. Large intestine

The colon absorbs most of the remaining water and salt, turning several litres of liquid a day into a much smaller volume of faeces. Fluid accounting is striking: around 2 L a day is drunk and eaten, and roughly 7 L more is **secreted** into the gut as saliva, gastric juice, bile, pancreatic juice, and intestinal fluid. Normally all but 100–200 mL is reabsorbed. If that reabsorption fails — in severe diarrhoea, or from a high-output stoma — a person can lose litres of fluid and salt a day and become dangerously dehydrated within hours.

The large intestine also houses most of the **gut microbiome**: tens of trillions of bacteria, roughly as many cells as the human body has. They ferment fibre the body cannot digest, produce some vitamins, train the immune system, and resist colonisation by harmful organisms. Antibiotics disturb them, which is one reason antibiotic use has side effects and why **Clostridioides difficile** infection follows antibiotic treatment (Chapter 13).

The gut also has its own nervous system — the **enteric** division of Chapter 7, hundreds of millions of neurons — which coordinates peristalsis largely independently of the brain.

10.3.5. Liver, gallbladder, and pancreas

The **liver**, at about 1.5 kg, is the largest internal organ. It receives about a quarter of the cardiac output, most of it through the **portal vein**, which carries blood straight from the gut. Everything absorbed from the intestine therefore passes through the liver before it reaches the rest of the body. Among its many jobs:

- storing glucose as glycogen and releasing it between meals (with the pancreas, below);
- making plasma proteins, including albumin and most **clotting factors**;
- breaking down and clearing drugs, alcohol, hormones, and toxins — including the **first-pass metabolism** that removes much of some oral drugs before they reach the circulation (Chapter 3);
- making **bile** (roughly 0.5–1 L a day), which emulsifies fats so enzymes can reach them;
- processing haemoglobin breakdown products. When this fails, a yellow pigment, bilirubin, accumulates: **jaundice**.

The **gallbladder** stores and concentrates bile and squeezes it into the duodenum after a fatty meal. Stones can form in it and block its outflow.

The **pancreas** is two organs in one. Most of it is **exocrine**, producing digestive enzymes and bicarbonate. Scattered through it, making up only 1–2% of its mass, are the **islets of Langerhans** — the endocrine pancreas — which make insulin and glucagon.

10.3.6. When the tube cannot be used normally

Situation	What is done	Design-relevant point
Cannot swallow safely (Chapter 8)	**Enteral feeding** through a tube: nasogastric (through the nose), or **gastrostomy** (through the abdominal wall into the stomach, e.g. PEG) for longer term	Tube position verification; securing without pressure injury; feed pumps; connectors
Gut cannot absorb	**Parenteral nutrition**: nutrients given intravenously	Central venous access; infection risk; a pharmacy-compounded product

Situation	What is done	Design-relevant point
Part of the bowel removed or rested	A **stoma**: bowel brought out through the abdominal wall. **Colostomy** (colon, firmer output), **ileostomy** (small bowel, liquid and corrosive output), **urostomy** (urine)	Appliances are adhesive, changed daily or every few days, on skin that moves and changes shape — Chapter 5 in its hardest form



The connector that cannot be connected wrongly.

Feeding tubes once used the same Luer connectors as intravenous lines. People died when liquid feed, or drugs meant for the gut, were connected into a vein. The international response was the **ISO 80369** series of small-bore connector standards, which gives each application — enteral, respiratory, neuraxial, intravascular, and others — a geometrically **incompatible** connector. The enteral design is widely known as **ENFit** [54].

This is the top of the risk-control hierarchy in Chapter 19: not a warning label, not a colour, but a shape that makes the error physically impossible. Anything you design that sits in a feeding line must use the correct connector and must not include adaptors that defeat it.

10.4. The endocrine system

10.4.1. Hormones as a broadcast channel

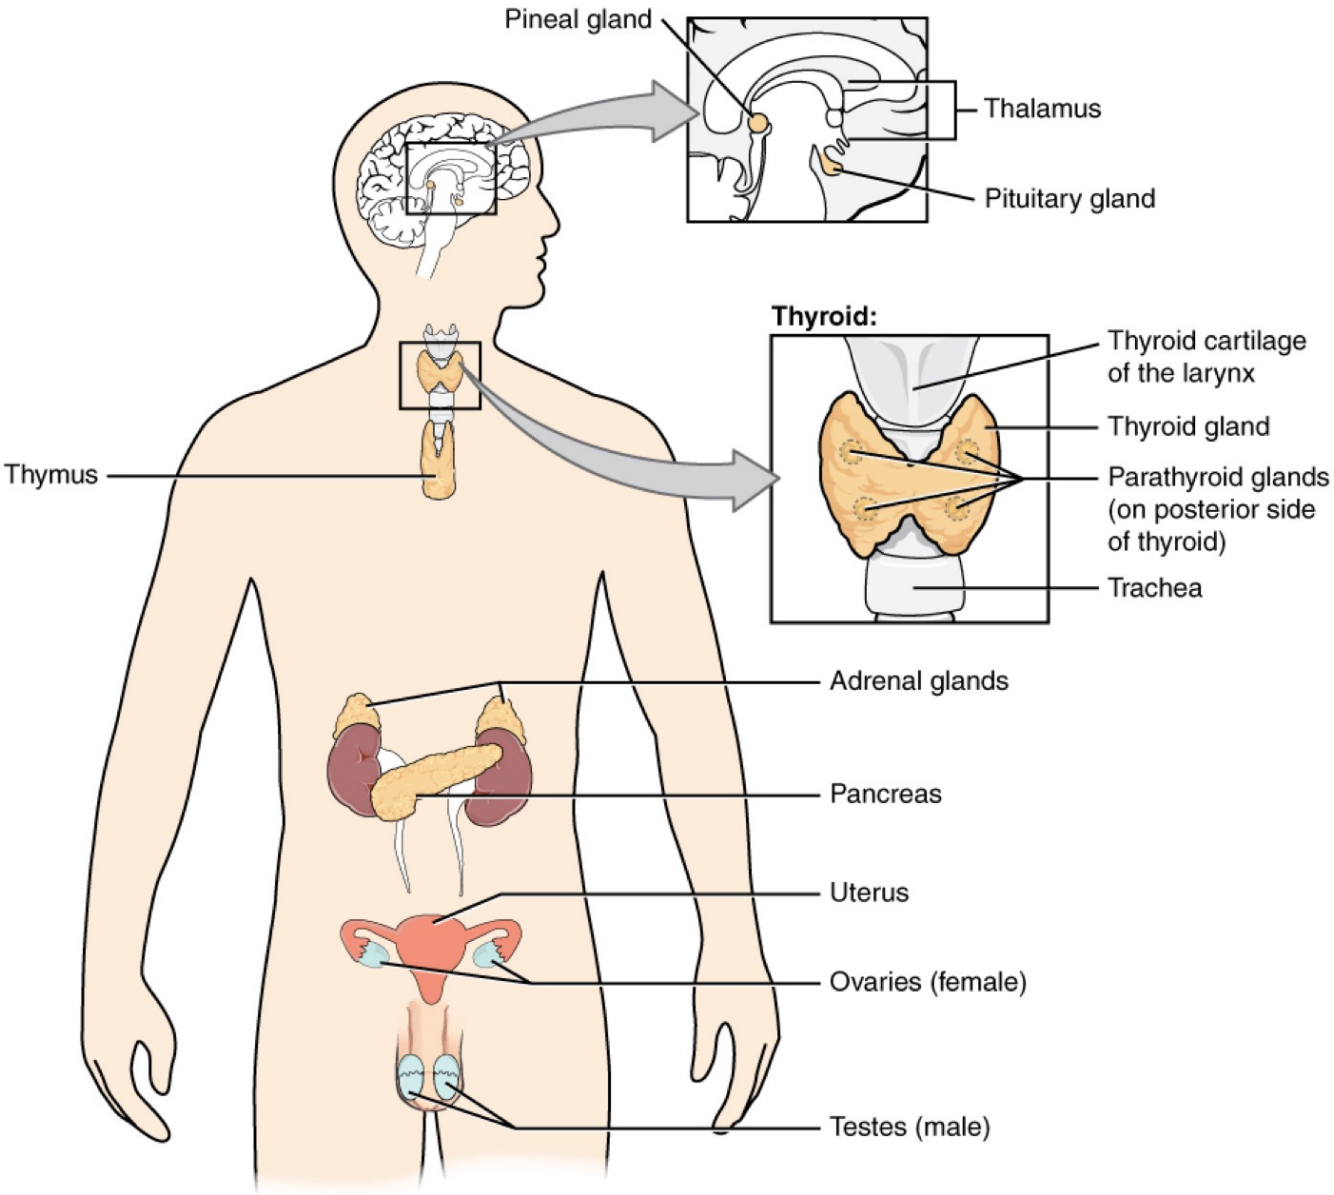


Figure 23. The major endocrine glands. Image: OpenStax, CC BY 3.0. Credits: Appendix H

A **hormone** is a chemical messenger released into the blood that acts on any cell carrying its receptor (Chapter 3' cell signalling). Where the nervous system is a wired network with point-to-point messages, the endocrine system is a **broadcast**: slow to start, slow to stop, reaching everywhere at once. Hormones differ enormously in how long they last in the blood:

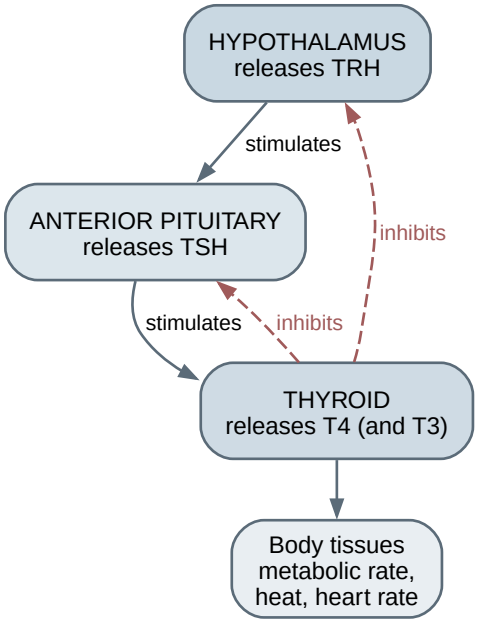
Hormone	Approximate half-life in blood	Consequence
Adrenaline	~2–3 min	Fast on, fast off: the "fight or flight" surge
Insulin	~4–6 min	Can be adjusted minute by minute
Cortisol	~1–1.5 h	Effects last hours
Thyroxine (T4)	~7 days	Thyroid changes take weeks to settle after a dose change

A control engineer would read that table as a list of **time constants**. They decide how fast each loop can respond, and therefore how often it makes sense to measure and adjust.

10.4.2. The hypothalamus and pituitary: the control hierarchy

The **hypothalamus**, at the base of the brain, links the nervous and endocrine systems. It signals the **pituitary gland** hanging beneath it, which in turn signals glands throughout the body. The **posterior pituitary** releases ADH (water balance) and oxytocin (childbirth, breast-feeding). The **anterior pituitary** releases growth hormone, prolactin, and four hormones that control other glands: TSH (thyroid), ACTH (adrenal cortex), and FSH and LH (ovaries and testes).

These form **axes** — three-level cascades with negative feedback at each level:



The thyroid axis shown is the cleanest example. Low thyroid hormone raises TSH; high thyroid hormone suppresses it. That is why **TSH** is the usual first blood test of thyroid function: it is the controller's output, and it amplifies small changes in the thing being controlled. The **adrenal axis** (CRH → ACTH → cortisol) has the same shape, with a strong daily rhythm — cortisol peaks in the early morning — which is why the time of a blood sample must be recorded.

10.4.3. The main glands

Gland	Main hormones	Too little / too much, briefly
Thyroid (front of the neck, Chapter 8)	T4 and T3, which set metabolic rate; needs dietary iodine	Hypothyroidism: tiredness, cold, weight gain, slowness. Hyperthyroidism: weight loss, heat intolerance, tremor, fast heart
Parathyroids (four, behind the thyroid)	Parathyroid hormone (PTH): raises blood calcium	Calcium disorders affect nerves, muscle, heart rhythm, bone
Adrenal cortex	Cortisol (stress, glucose, immunity); aldosterone (sodium, potassium)	Too little cortisol can be life-threatening during illness; long-term steroid medicines suppress the body's own production

Gland	Main hormones	Too little / too much, briefly
Adrenal medulla	Adrenaline and noradrenaline	The sympathetic surge of Chapter 7, delivered by blood
Pancreatic islets	Insulin (beta cells) lowers glucose; glucagon (alpha cells) raises it	Diabetes, below
Ovaries and testes	Oestrogen, progesterone, testosterone	Puberty, fertility, menopause, bone density (Chapter 12)
Kidneys, heart, gut, fat	Erythropoietin, natriuretic peptides, gut hormones, leptin	Many organs are part-time glands

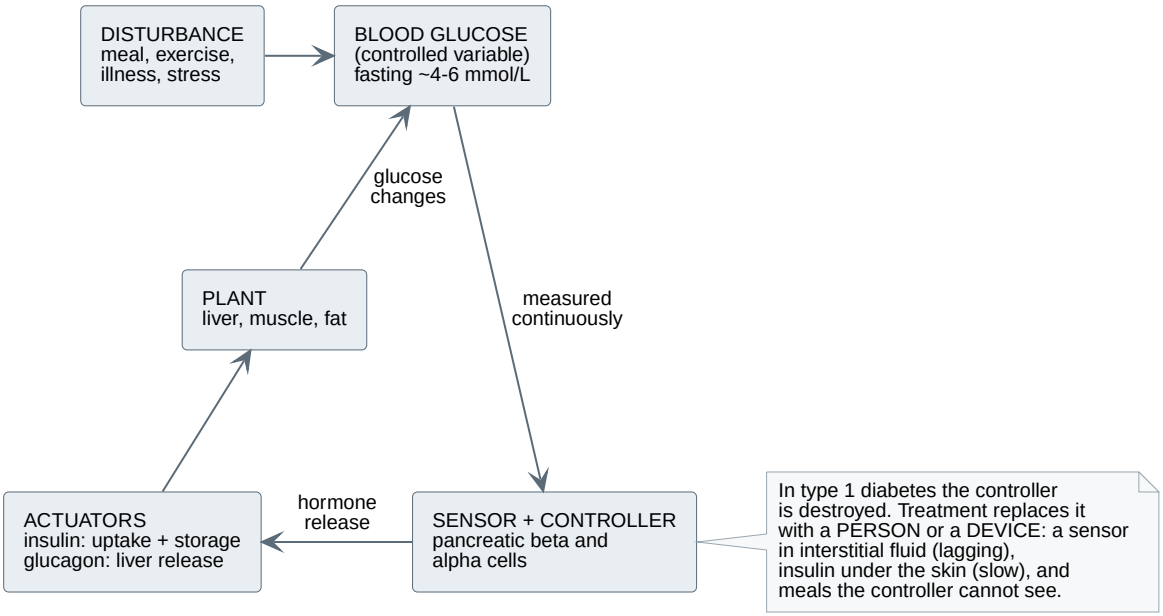
10.4.4. Calcium

Blood calcium is held within a remarkably narrow band — roughly 2.2–2.6 mmol/L — because nerve and muscle excitability depend on it. When it falls, the parathyroids release PTH, which draws calcium out of bone, reduces its loss in urine, and activates **vitamin D** in the kidney so that more is absorbed from the gut. The skeleton of Chapter 6 is therefore not only a structure but a **calcium reservoir** under hormonal control — one reason bone density changes with age, hormones, and kidney disease.

10.4.5. Glucose: the most engineered loop in the body

The brain alone uses roughly 120 g of glucose a day and has little capacity to store it, so blood glucose must be kept up between meals and kept down after them. Two opposed hormones from the pancreatic islets do most of the work:

- **Insulin**, released when glucose rises, moves glucose into muscle and fat cells, tells the liver to store it as glycogen, and suppresses the liver's own glucose output.
- **Glucagon**, released when glucose falls, tells the liver to release glucose.



Fasting glucose is normally about 4–6 mmol/L (Appendix C). Below about 4 mmol/L is

hypoglycaemia: sweating, shaking, confusion, and — if severe — seizures and loss of consciousness. Prolonged **high** glucose damages blood vessels, nerves, eyes, and kidneys over years. The **HbA1c** blood test, which measures how much haemoglobin has become glycated, reflects average glucose over the previous two to three months.

Diabetes is failure of this loop:

- **Type 1** — the immune system destroys the beta cells. No insulin is produced; it must be replaced, for life, from diagnosis. Often begins in childhood.
- **Type 2** — the tissues respond poorly to insulin (**insulin resistance**) and the beta cells cannot keep up. Far more common; associated with age, weight, and genetics; managed with lifestyle, tablets, other injectables, and sometimes insulin.
- **Gestational** — arising in pregnancy.

Diabetes is common, chronic, and largely **self-managed** at home — which is why its technology is consumer technology, and why the person using it is usually also the person making decisions from it.

10.5. Engineering bridge: filters, transit, loops, and units

10.5.1. Clearance: the kidney as a mass-transfer device

Clearance is the volume of plasma completely cleared of a substance per minute:

C = (U × V) / P

where **U** is the substance's concentration in urine, **V** the urine flow rate, and **P** its concentration in plasma.

Work an example with creatinine. Urine creatinine 10 mmol/L, urine flow 1.2 mL/min, plasma creatinine 0.1 mmol/L (100 µmol/L):

C = (10 × 1.2) / 0.1 = 120 mL/min

— a typical adult GFR, as an order of magnitude. Because creatinine is filtered and barely reabsorbed or secreted, its clearance approximates GFR. The **same equation** describes an artificial kidney: a dialyser's clearance is the blood flow rate multiplied by the fraction of the solute removed on one pass.

10.5.2. Dialysis

When kidneys fail, **dialysis** replaces their filtering — not their hormones.

- **Haemodialysis** pumps blood, typically at 300–400 mL/min, through a cartridge of thousands of hollow fibres made of a semi-permeable membrane. **Dialysate** flows the opposite way on the outside — **countercurrent**, to keep a concentration gradient along the whole length. Small solutes move by **diffusion**; excess water is removed by applying a pressure difference,

ultrafiltration. Sessions are usually about four hours, three times a week, through a surgically created **arteriovenous fistula** in the arm.

- **Peritoneal dialysis** uses the person's own peritoneum — the lining of the abdominal cavity — as the membrane. About 2 L of dialysate is run into the abdomen through a permanent catheter, left to equilibrate, and drained, several times a day or overnight by a machine.

Compare the totals. Healthy kidneys filtering at 120 mL/min, continuously, process about **1,200 litres of plasma a week**. A dialyser clearing urea at, say, 200 mL/min for twelve hours a week clears about **144 litres**. Dialysis works — people live on it for many years — but it is an order of magnitude less filtration, delivered in pulses, with none of the kidney's hormonal functions. That arithmetic explains fluid and diet restrictions, the fatigue, and why transplantation is preferred where possible. Dialysis equipment is life-sustaining and firmly Tier 3 (Appendix I); the arithmetic is yours to understand.

10.5.3. Transit, surface, and residence time

The gut is a **plug-flow reactor** with walls that absorb. What is absorbed depends on surface area, residence time, and the concentration gradient — the same variables as any exchanger. Speed up transit (diarrhoea, a shortened bowel) and absorption falls; remove the terminal ileum and specific substances are no longer absorbed at all. A drug designed to dissolve in a particular segment, or a "modified-release" tablet, is a design for a particular position and time in this reactor — which is why crushing tablets to put them down a feeding tube can be dangerous, and is a pharmacist's decision.

10.5.4. Hormonal loops as control systems

Every hormonal loop has the parts of a feedback controller (Chapter 1):

- a **set point** (not fixed — it moves with time of day, age, illness, pregnancy);
- a **sensor** (cells that detect glucose, calcium, osmolality, thyroid hormone);
- a **controller** with a **gain** (how much hormone per unit error);
- an **actuator** (the hormone and its target tissue); and
- **delays** — secretion, transport, receptor action, and the hormone's half-life.

Delay is the enemy of control. A loop with a long delay and a high gain overshoots and oscillates. The body's loops are tuned for their own delays. Artificial loops have **new** delays that the body never had to manage.

Automated insulin delivery shows this vividly. Compare the natural and artificial loops:

Element	Natural (pancreas)	Artificial (automated insulin delivery)
Sensor	Beta cells sample **blood** directly, continuously	CGM measures **interstitial fluid**; lags blood by roughly 5–15 minutes; noise and drift

Element	Natural (pancreas)	Artificial (automated insulin delivery)
Actuator	Insulin released into the **portal vein**, reaching the liver first, acting within minutes	Rapid-acting insulin injected **under the skin**: begins working in ~10–20 min, peaks at ~1–2 h, lasts ~3–5 h
Feedforward	Gut hormones warn of an incoming meal	The user must usually announce meals and estimate carbohydrate
Opposing actuator	Glucagon raises glucose	Most systems have **insulin only**: they can stop insulin but cannot push glucose up

The algorithms in commercial systems use variants of PID and **model-predictive control**. They are deliberately conservative, because insulin already given cannot be taken back. A community of people with diabetes and their families, unwilling to wait for commercial products, reverse-engineered pump and CGM communications and built open-source closed-loop systems — OpenAPS, AndroidAPS, and Loop among them [53]. Their story is instructive in both directions: it showed what users could achieve and how strong the unmet need was, and it raised exactly the questions of verification, cybersecurity, and responsibility that Chapter 19 and Chapter 20 address. Regulators have since created pathways for **interoperable** sensors, pumps, and controllers.

10.5.5. Units: where this chapter's numbers bite

Quantity	Unit A	Unit B	Conversion and hazard
Glucose	mmol/L	mg/dL	× 18. A reading of "6" is an everyday value in mmol/L and dangerously low in mg/dL; "108" is an everyday value in mg/dL and impossible in mmol/L
Creatinine	µmol/L	mg/dL	× 88.4. 1.0 mg/dL ≈ 88 µmol/L
HbA1c	mmol/mol (IFCC)	% (DCCT)	Not a simple ratio: an offset as well as a scale, (% − 2.15) × 10.929; 48 mmol/mol ≈ 6.5%. Both are in clinical use
Urea / BUN	Urea, mmol/L	Blood urea nitrogen, mg/dL	Different substances counted differently; BUN mg/dL × 0.357 ≈ urea mmol/L
Insulin	Units per mL	—	Standard U-100, but concentrated U-200, U-300, and U-500 insulins exist. Volume is **not** dose unless concentration is known



In this domain, a unit or concentration error is a dosing error, and a dosing error can kill within hours.

If you write software that touches any of these values: store the unit with the value

(Appendix C), reject ambiguous input, never infer the unit from the magnitude silently, display the unit every time, and treat any code that **calculates or recommends** a dose as a regulated medical device. That last point is not a technicality. It is Tier 3 (Appendix I) for good reason.

10.6. Devices, software, and training technology

10.6.1. Clinical devices

Device	Does	Engineering note
Urinary catheters and drainage systems	Drain the bladder; bags with measuring chambers for hourly output	Balloon retention; material (latex, silicone, coatings); closed drainage; anti-reflux valves. Infection is the dominant failure mode
Intermittent catheters	Single-use catheters for self-catheterisation	Hydrophilic coatings for low friction; packaging that can be opened with limited dexterity
Continence products	Pads, sheaths, external wicking devices, bedside urinals	Absorption, skin microclimate (Chapter 5), discretion, dignity
Bladder scanners	Estimate urine volume in the bladder by ultrasound	Avoids unnecessary catheterisation; an **estimate** with known error (Chapter 16)
Haemodialysis machines and dialysers	Remove solutes and fluid from blood	Blood pump, countercurrent hollow-fibre membrane, dialysate preparation from purified water, air detection, pressure and conductivity monitoring. Life-sustaining
Peritoneal dialysis cyclers	Automate overnight dialysate exchanges at home	Fluid volume accounting; infection at the catheter exit site
Enteral feeding tubes, pumps, and ENFit connectors	Deliver liquid nutrition into the stomach or small bowel	ISO 80369-3 connectors; pump accuracy; tube securement without pressure injury
Stoma appliances	Collect output from a colostomy, ileostomy, or urostomy	Adhesive chemistry, skin protection, filters for gas, fit to a changing shape
Endoscopes and capsule cameras	See inside the gut	Chapter 16; reprocessing (Chapter 13)
Blood glucose meters	Measure capillary glucose from a finger-prick	Amperometric enzyme strips (Chapter 3); accuracy requirements in ISO 15197; haematocrit and temperature interference

Device	Does	Engineering note
Continuous glucose monitors (CGM)	Measure interstitial glucose every few minutes for days to weeks	Enzyme sensor on a fine filament under the skin; lag, drift, adhesive wear, alarms that must wake a sleeping child's parents
Insulin pens and pumps	Deliver insulin under the skin; pumps give a programmable background (basal) rate plus boluses	Dose accuracy at microlitre volumes; occlusion detection; usability; cybersecurity
Automated insulin delivery (closed loop)	Adjusts insulin automatically from CGM readings	A medical control system with a human in the loop. See the bridge above

10.6.2. Software

Software	Does	Note
Diabetes data platforms	Collect and display meter, CGM, and pump data for users and clinics	Time-in-range and glucose pattern reports; data from many vendors; privacy (Chapter 20)
Open-source looping projects — OpenAPS, AndroidAPS, Loop	User-built automated insulin delivery	**Read the documentation, not the code, first.** The safety reasoning, testing advice, and warnings are an education in themselves. Not something to build for anyone else
Nightscout	Open-source remote CGM monitoring, e.g. for parents	A real example of patient-led software meeting an unmet need
Fluid balance charts and electronic observations	Record intake and output	Often still on paper; totals, time windows, and units are frequent error sources
eGFR calculators and drug-dosing references	Estimate kidney function and adjust doses	Calculation is easy; **validation and responsibility** are not. Regulated when it drives dosing
Physiological simulators — open glucose–insulin models (e.g. the UVA/Padova simulator family used in research), teaching models of renal function	Test algorithms **in silico** before any human use	A simulator is a model of a population, not of a person (Chapter 21)

10.6.3. Training and simulation

Training technology	Teaches	Form
Catheterisation trainers	Urethral catheter insertion with aseptic technique; male and female anatomy	Pelvic models with replaceable urethral inserts and fluid reservoirs. **Consumable parts wear out**
Nasogastric tube trainers	Insertion, securing, and position checking	Head and torso sections (Chapter 8)
Stoma care trainers	Changing appliances, skin care, measuring the stoma	Abdominal pads with stoma models of various shapes. **A maker-accessible silicone casting project**
Injection pads and insulin pump demonstrators	Injection technique, site rotation, pump operation	Practice pads; "trainer" pumps that deliver saline or nothing
Dialysis teaching rigs	Diffusion, countercurrent exchange, ultrafiltration	Dialysis tubing, salt solutions, and a conductivity meter — a classic bench demonstration
Abdominal examination simulators	Palpation of liver, spleen, and masses; listening to bowel sounds	Instrumented manikins and standardised patients
Continence and dignity training	Respectful continence care; product selection	Scenario-based teaching with real products. The skills are communication as much as technique



Where a maker fits in this chapter. Physical: catheterisation and stoma trainer consumables (silicone casting, fabric-backed abdominal pads), adaptive aids for self-catheterisation and continence (with an OT and the user — often sewn, bent wire or sheet metal, or moulded rather than printed), and bench demonstrations of diffusion and dialysis using laboratory glassware and tubing. Software: data visualisation and education on **synthetic or public** glucose data, unit-safe data handling, and accessibility of the diabetes apps millions rely on. Not: anything that calculates a dose, connects to a pump, or sits in a feeding line.

10.7. Worked example: the drainage-bag hanger

A ward sister asks you for something simple. Catheter drainage bags hang from a hook on the bed frame, but when patients sit out in a chair the bag ends up on the floor, or draped over the chair arm, or — worse — lifted onto the patient's lap. Could you make a clip-on hanger for the chairs?

It sounds like Tier 1 furniture. Work through what the physiology and infection control of this chapter say first.

Functional requirements, each with its reason:

1. **Below the bladder at all times**, including when the patient leans or the chair reclines — so urine cannot flow back from the bag into the bladder. Specify a minimum drop from seat height, and check it with the chair's full range of adjustment.
2. **Clear of the floor** — floors are heavily contaminated, and a bag outlet touching the floor is a route for organisms into a closed system (Chapter 13).
3. **Tubing without kinks or loops below the bag inlet** — a kink stops drainage; a dependent loop lets urine pool and run backwards.
4. **Measuring scale readable in place** — hourly output is clinical data. If staff must unhook the bag to read it, they will read it less often, or less accurately (Chapter 15).
5. **Outlet accessible for emptying** without the outlet touching the hanger or the floor.
6. **Cleanable** between patients, with no crevices (Chapter 13), and **no sharp edges** near legs that may have fragile, oedematous skin (Chapter 5).
7. **Does not interfere with transfers** — a hoist sling, a standing aid, or a person swinging their legs out must not catch on it.
8. **Dignity**: a cover option, so the bag's contents are not on display to visitors.

What you do not get to decide: whether a given patient's catheter should be on a leg bag instead (a clinical choice), the cleaning agent (Chapter 13), and whether the design is adopted (the ward and infection prevention team).

Then test it on the actual chairs, in the actual ward, with the staff who empty the bags — and record what they had to do differently.

The general lesson: a passive bracket on a chair is in a **closed fluid system** that has a direction of flow, an infection risk, and a measurement role. None of that is visible from the bracket. All of it comes from the physiology.

10.8. Workshop exercise

1. **Clearance arithmetic.** Calculate creatinine clearance for three invented cases: (a) urine creatinine 8 mmol/L, flow 1.5 mL/min, plasma 0.09 mmol/L; (b) the same urine values with plasma 0.3 mmol/L; (c) case (a) with the plasma value mistakenly entered in µmol/L. Comment on (c).
2. **Bench dialysis.** Fill a length of dialysis (visking) tubing with salt solution and a little starch, tie it off, and place it in a beaker of tap water. Measure the water's conductivity (a cheap TDS meter works) every five minutes for an hour, and test the water for starch with iodine at the end. Then repeat with the water stirred, and with the water replaced every ten minutes. Explain the differences using concentration gradient and countercurrent flow.
3. **A toy glucose loop.** In a spreadsheet or Python, model glucose as a tank with a meal input, a constant drain, and an insulin "valve" controlled proportionally to the error from a set point. Add a 15-minute sensor delay and a 60-minute actuator delay. Vary the gain and watch for overshoot. **Label every output "toy model — not physiology".** The point is the shape of the problem, not the numbers.
4. **A tube to scale.** Using string or tape, lay out the gut at 1:10 scale with the approximate lengths

in this chapter. Mark where absorption of glucose, B_{12}, and water mainly happens.

5. **Units in the type.** Write a small data type or class for a glucose reading that cannot be created without a unit, converts explicitly, and refuses to compare readings in different units. Write three tests that would have caught a mmol/L–mg/dL confusion.
6. **Connector awareness.** Find published images or specifications of Luer and ENFit connectors and sketch both. Write down, in one paragraph, why making them incompatible is a stronger control than colour-coding them.

10.9. Questions

1. What are the two stages of urine formation, and roughly how much fluid is filtered each day?
2. Why is filtering everything and reabsorbing most of it a robust design?
3. Name four things the kidneys regulate besides water.
4. Why is eGFR least reliable in a frail older person with low muscle mass?
5. List the four common types of incontinence and say which is most directly a design problem.
6. Give three rules for a urinary drainage system and the reason for each.
7. Why must the position of a nasogastric tube be confirmed before use?
8. What problem does the ISO 80369 series solve, and where does it sit in the risk-control hierarchy?
9. Explain the thyroid axis and why TSH is the usual first thyroid test.
10. Compare the natural glucose loop with an automated insulin delivery system, naming two delays and one missing actuator in the artificial version.
11. A glucose reading of 5.5 is entered into software that assumes mg/dL. What happens, and how should the software have been designed?
12. Roughly how does weekly small-solute clearance by haemodialysis compare with healthy kidneys, and what does dialysis not replace?

10.10. Answers and reasoning

1. **Filtration** at the glomerulus, then **reabsorption and secretion** along the tubule. About 180 L a day of filtrate (GFR ~120 mL/min), of which more than 99% of the water is reclaimed, leaving ~1–2 L of urine.
2. Anything the body has not specifically arranged to keep is excreted by default, including unfamiliar toxins and drug metabolites. A whitelist is safer than a blacklist because it does not need to anticipate every harmful substance.
3. Any four of: sodium and blood volume; potassium; acid–base balance; calcium and phosphate (via vitamin D activation); red cell production (erythropoietin); long-term blood pressure.
4. eGFR is calculated from blood creatinine, which comes from muscle. Low muscle mass produces less creatinine, so blood creatinine looks normal even when filtration is reduced — the estimate overstates kidney function, with consequences for drug dosing.

5. **Stress, urge, overflow, functional.** Functional incontinence — where the urinary system works but mobility, dexterity, vision, or cognition prevent timely toileting — is most directly a design problem: clothing, toilet access, rails, signage, and aids.
6. Any three of: **closed system**, so organisms cannot enter; **bag below the bladder**, so urine cannot flow back; **off the floor**, because floors are contaminated; **no kinks or dependent loops**, so drainage continues and urine does not pool; **outlet kept clean** when emptying.
7. Because a tube passed through the nose can enter the airway instead of the oesophagus, and feed or drugs delivered into the lungs can be fatal (Chapter 8).
8. Misconnection between small-bore connectors of different applications — feed or drugs into a vein, for example. Making the connectors physically incompatible is **design to eliminate** the hazard: the top of the hierarchy, above protective measures and information for safety.
9. The hypothalamus releases TRH, which makes the anterior pituitary release TSH, which makes the thyroid release T4/T3; thyroid hormone feeds back to inhibit both. TSH is the controller output and responds strongly to small changes in thyroid hormone, so it is a sensitive early indicator.
10. Natural: beta cells sense **blood** glucose and release insulin into the **portal vein**, acting in minutes; glucagon opposes it; gut hormones give advance notice of meals. Artificial: the CGM **lags** blood glucose by ~5–15 min; subcutaneous insulin takes ~10–20 min to start and peaks at ~1–2 h; most systems lack **glucagon** (cannot raise glucose) and need meals announced.
11. 5.5 mg/dL is about 0.3 mmol/L — profound hypoglycaemia — so the software may raise a false emergency, or (if the assumption runs the other way in a dosing context) recommend a dangerous action. The value should have carried its unit; ambiguous input should have been rejected, not guessed; and the unit should have been displayed beside every value.
12. Healthy kidneys filter roughly 1,200 L of plasma a week; haemodialysis at ~200 mL/min for ~12 h a week clears roughly 140–150 L — an order of magnitude less, delivered intermittently. Dialysis does not replace the kidney's hormonal roles: erythropoietin, vitamin D activation, and much of blood-pressure regulation.

10.11. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Unit-safe clinical arithmetic — glucose, creatinine, HbA1c and clearance with units that cannot be confused.
- Diffusion, dialysis and clearance — clearance and the artificial kidney.
- Half-lives, decay and steady state — hormone half-lives as time constants.
- PID controller — a toy glucose loop with sensor and actuator delays.
- ODE solver — a minimal glucose model.

10.12. Video companion and reading

- **Core videos:** the Crash Course Anatomy & Physiology episodes on the **urinary system**, the

digestive system, and the **endocrine system**, from the series index [6].

- **Core reading:** OpenStax [1] selectively:
  - 25.3–25.6 — the kidney's anatomy, urine formation, and tubular reabsorption
  - 25.8–25.9 — hormonal regulation of kidney function and fluid volume
  - 26.1–26.4 — body fluids, electrolytes, and acid–base balance (skim)
  - 23.1, 23.4–23.6 — overview of the digestive system; stomach; small and large intestines; accessory organs
  - 17.1–17.6 and 17.9 — hormones, the pituitary and hypothalamus, thyroid, parathyroid, adrenals, and the endocrine pancreas
- **Kidney disease classification:** the KDIGO chronic kidney disease guideline [52]. Read the staging tables; skip the treatment.
- **Patient-led engineering:** the OpenAPS reference design and documentation [53]. Read it as a case study in requirements and risk written by users.
- **Connectors:** GEDSA's material on ENFit and ISO 80369 [54].

10.13. Exit criteria

- ☐ Draw a nephron from memory, labelling filtration, reabsorption, and where ADH and aldosterone act.
- ☐ Calculate a clearance from urine and plasma values, and show what a unit error does to the result.
- ☐ Trace the digestive tube from mouth to anus with approximate lengths and transit times, and place the liver, gallbladder, and pancreas on it.
- ☐ Explain the thyroid axis and the glucose–insulin loop in control-system terms: sensor, controller, actuator, set point, delay.
- ☐ Write down the unit pairs for glucose, creatinine, and HbA1c, with conversions, and implement a unit-carrying glucose type with tests.
- ☐ State the drainage-system rules for a urinary catheter and the reason for each.
- ☐ Explain why ENFit exists and which level of the risk-control hierarchy it represents.

10.14. What this chapter does not establish

- **You cannot interpret kidney, liver, or hormone results.** You know what the tests measure. What a result means for a person depends on context only clinicians are positioned to weigh.
- **You cannot manage catheters, feeding tubes, or stomas.** Their insertion, care, and troubleshooting are clinical and nursing skills, learned under supervision.
- **You cannot design dosing software or insulin-delivery algorithms for anyone's use.** The bridge explains why the problem is hard. Building something people rely on to dose insulin is regulated medical-device work with a clinical and regulatory team — the open-source projects are explicit that their users take responsibility for themselves alone.

- **The toy glucose model is not physiology.** It shows the shape of a control problem. Real glucose–insulin models are validated research tools, and even they model populations, not individuals.
- **You have not learned nephrology, gastroenterology, or endocrinology.** This chapter is enough physiology to follow a conversation in each and to recognise when a project has become one of them.

Chapter 11. How disease works: from mechanism to symptom

Week 14 · approximately 6–8 hours

11.1. Why this chapter exists

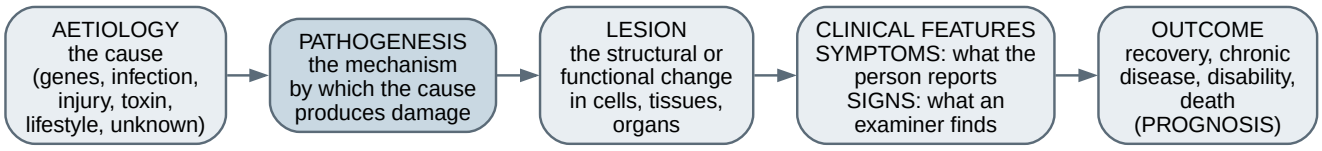
Every chapter so far has described the body working **normally**. But nobody commissions a device for a body that is working normally. Glucose monitors exist because of diabetes, compression stockings because of thrombosis, inhalers because of asthma, offloading boots because feet ulcerate, and most of the training aids in Chapter 21 exist because clinicians must learn to recognise and manage disease.

That makes disease the hidden source of most requirements you will ever be given. When an occupational therapist says "she has diabetic neuropathy", they are not giving you a label — they are giving you a specification in compressed form: **this person cannot feel a pressure injury forming, may not see well enough to inspect for one, and heals slowly when one occurs.** If you cannot decompress the label, you will design for the wrong person.

This chapter is **pathophysiology for an engineer**: not a catalogue of diseases, but the small number of mechanisms from which most of them are built — cell injury, inflammation, infection, clotting, blocked vessels, failing pumps, uncontrolled growth, mistaken immunity, wear, and the genetic code that underlies them all. It ends with the arithmetic of risk, because every clinical claim you read about a disease or a device is expressed in it, and it is routinely presented in the most flattering form available.

11.2. The language of disease

Clinical descriptions follow a fixed sequence, and the vocabulary maps onto it.



Term	Meaning
Aetiology	The cause. Often multiple, often partly unknown
Pathogenesis	The mechanism — the sequence of events from cause to damage
Symptom	Something the person experiences and reports: pain, breathlessness, tiredness
Sign	Something an examiner observes or measures: a heart murmur, swelling, a raised temperature
Acute / chronic	Sudden onset and short course / long-lasting, often permanent. Not a statement of severity

Term	Meaning
Prognosis	The expected course and outcome
Risk factor	Something associated with a higher probability of disease — not necessarily its cause
Comorbidity	Two or more conditions present together, which is the norm in older patients
Incidence	New cases arising in a population over a period
Prevalence	Total existing cases at a point in time

Two of these pairs repay attention.

Signs versus symptoms is the distinction of Chapter 9's four-layer model in clinical dress: a symptom is reported, a sign is measured. Both are **data**, and neither is the disease. A device that reports a sign — a temperature, a saturation — is reporting evidence, not a diagnosis.

Incidence versus prevalence matters whenever someone tells you how big a need is. Roughly, **prevalence ≈ incidence × average duration**. A condition that is common but quickly resolved (a sprained ankle) has high incidence and low prevalence. A condition that is uncommon but lifelong (type 1 diabetes) can have a modest incidence and a large prevalence. A device for the first serves a constant stream of new users for a short time; for the second, a stable population for decades. Those are different businesses and different designs.

11.3. Cell injury and cell death

11.3.1. Why oxygen supply is the commonest failure

Most cells survive on aerobic respiration (Chapter 3). Cut off their oxygen — **hypoxia** — or cut off their whole blood supply — **ischaemia**, which also removes glucose and waste clearance — and the energy supply fails within minutes. Ion pumps stop, the cell swells with water, membranes fail, and enzymes released from damaged structures begin to digest the cell from within.

Tissues differ enormously in how long they tolerate this:

Tissue	Tolerance of complete ischaemia (orientation)	Consequence
Brain	Minutes	The brain is ~2% of body mass but uses ~20% of resting oxygen, and has almost no reserve
Heart muscle	Irreversible injury begins after roughly 20–40 minutes	The basis of "time is muscle" in heart attack treatment
Kidney	Tens of minutes to hours	Relevant to transplantation and surgery
Skeletal muscle	A few hours	Tourniquet times are limited for this reason

Tissue	Tolerance of complete ischaemia (orientation)	Consequence
Skin and connective tissue	Hours	But sustained pressure keeps them ischaemic for exactly that long (Chapter 5)

These figures are rough orientation; tolerance depends on temperature, partial versus complete loss of flow, and the individual.

For untreated large-vessel stroke, a widely quoted estimate is the loss of about **1.9 million neurons per minute**. Every system of rapid stroke assessment, scanning (Chapter 16), and treatment is built around that number.

11.3.2. Two ways to die

- **Necrosis** is uncontrolled death from injury. Cells swell and burst, spill their contents, and provoke **inflammation**. It is messy, and it is what happens in a heart attack, a pressure injury, or a burn.
- **Apoptosis** is **programmed** cell death: an orderly self-dismantling in which the cell shrinks, packages its contents, and is quietly eaten by neighbours without inflammation. The body uses it constantly — to remove damaged cells, shape organs during development, and end immune responses. **Failure** of apoptosis is one of the hallmarks of cancer.

A cell that is stressed but not killed can **adapt**: grow larger (**hypertrophy**, as in a heart muscle working against high blood pressure), multiply (**hyperplasia**), shrink (**atrophy**, as in an unused limb in a cast), or change type (**metaplasia**, as in the airway lining of a smoker). Adaptations are reversible up to a point and then become disease in their own right.

11.4. Inflammation

Inflammation is the body's general-purpose response to injury and infection. Chapter 5 described it as the first phase of healing; here it is as a disease mechanism, because it is the common pathway of a remarkable fraction of illness.

Acute inflammation develops over minutes to hours. Small vessels dilate and become leaky; fluid and proteins leak into the tissue; white cells, mainly **neutrophils**, migrate out of the blood to attack the cause. The classical signs, described two thousand years ago, follow directly from that vascular change:

- **redness** and **heat** — more blood flow;
- **swelling** — leaked fluid (oedema);
- **pain** — chemical mediators and pressure sensitising nociceptors (Chapter 7);
- **loss of function** — all of the above.

Chronic inflammation persists for weeks to years, with different cells (macrophages and lymphocytes) and a tendency to **scar** — fibrosis — and to destroy the tissue it is meant to protect. Chronic inflammation drives rheumatoid arthritis, inflammatory bowel disease, chronic lung

disease, and atherosclerosis, and contributes to diabetes and many cancers.

For a designer, two consequences recur. Inflamed tissue is **swollen**, so its dimensions change — the strap that fitted at nine in the morning does not fit at five. And inflamed tissue is **sensitised**, so a contact that is harmless on healthy skin is painful or damaging here.

11.5. Infection, sepsis, and shock

11.5.1. Infection and the host response

An **infection** is the invasion and multiplication of micro-organisms — bacteria, viruses, fungi, parasites — in the body. Whether it causes disease depends on both sides: the organism's **virulence** and the person's **defences**. Chapter 13 covers how organisms travel between people and objects; here the question is what happens once they arrive.

Most of the harm in many infections comes not from the organism directly but from the **host response**: inflammation, fever, and the immune attack. That is usually proportionate and protective. Sometimes it is not.

11.5.2. Sepsis

Sepsis is defined as life-threatening organ dysfunction caused by a **dysregulated host response** to infection. The infection may begin anywhere — lungs, urinary tract, abdomen, skin, a line or catheter — but the response becomes body-wide: vessels dilate and leak everywhere, blood pressure falls, clotting is activated in small vessels, and organs including the kidneys, lungs, brain, and liver begin to fail. Global estimates are in the order of tens of millions of cases and around ten million deaths a year.

Sepsis is a time-critical condition, and its early features are non-specific — fever or low temperature, fast heart rate, fast breathing, confusion, feeling "very unwell". That combination of urgency and vagueness is why hospitals use **early warning scores** built from routine observations, and why automated sepsis alerts are a large, and contested, software category (see below).

11.5.3. Shock

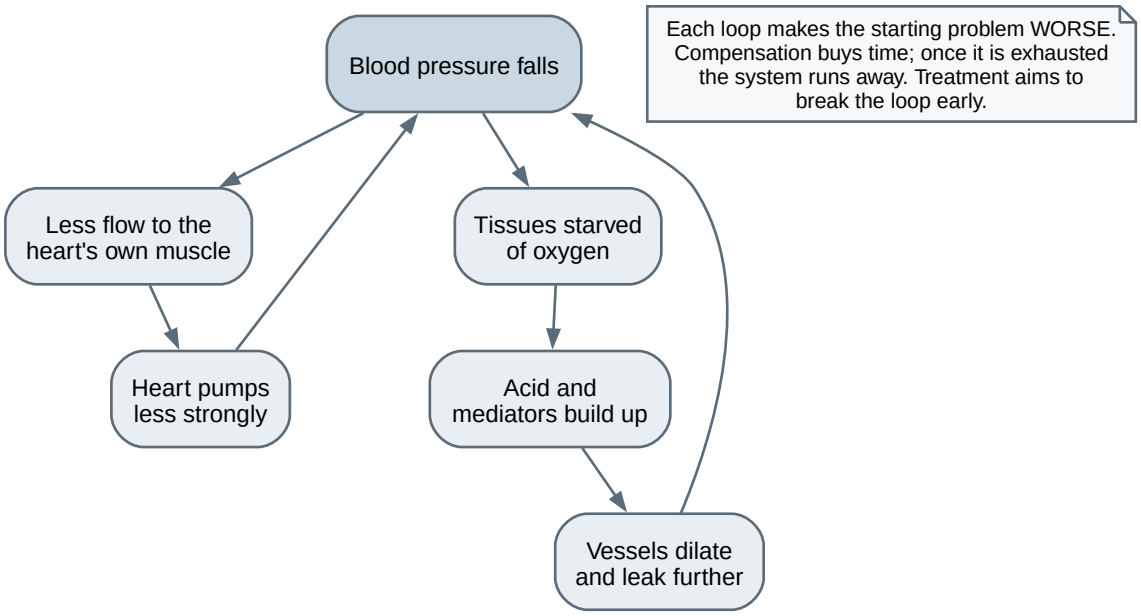
Shock is circulatory failure: the tissues are not receiving enough oxygen to meet their needs, whatever the blood pressure reads. Using Chapter 9's haemodynamics, there are four mechanisms — each a failure of a different term in the circuit:

Type	What fails	Examples
Hypovolaemic	Not enough volume in the circuit	Bleeding, severe dehydration, burns
Cardiogenic	The pump	Major heart attack, severe heart failure, arrhythmia
Distributive	The pipes dilate, so pressure collapses although volume is normal	**Septic**, anaphylactic, neurogenic (spinal injury)

Type	What fails	Examples
Obstructive	Flow is physically blocked	Large pulmonary embolism, fluid compressing the heart, tension pneumothorax

In every type, the body compensates first: heart rate rises, vessels in skin and gut constrict, breathing quickens. Blood pressure may be **maintained** until compensation is exhausted — and then fall abruptly. That is a crucial lesson about measurement: **a normal blood pressure does not exclude shock**, and a falling one is a late sign, not an early one.

11.5.4. When negative feedback turns positive

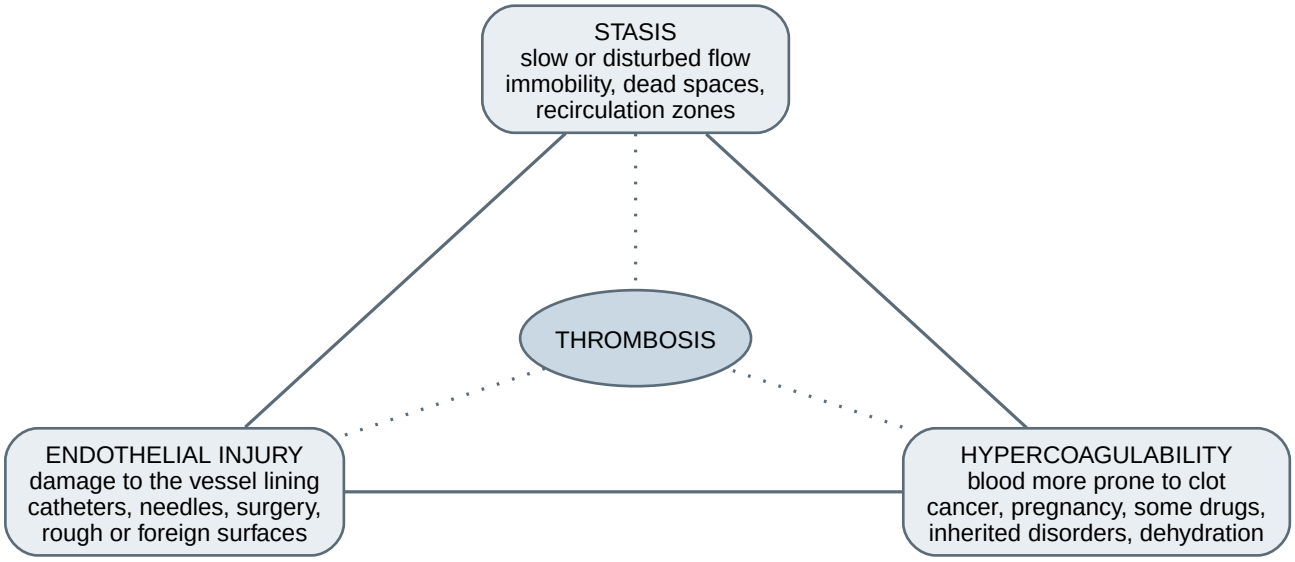


Chapter 1 introduced **negative feedback** as the basis of homeostasis and **positive feedback** as a rare, self-limiting exception. Much of severe acute illness is negative feedback **overwhelmed**, so that loops which normally restore stability begin to amplify the disturbance. Shock is the clearest example; uncontrolled bleeding with consumption of clotting factors, and some arrhythmias, are others. An engineer will recognise a control system driven out of its stable operating region. The clinical consequence is the same as the engineering one: **intervene while the system is still recoverable**, because past a certain point it is not.

11.6. Thrombosis and embolism

Blood must stay liquid inside vessels and clot instantly outside them. **Haemostasis** manages that balance through platelets and a cascade of clotting proteins. A **thrombus** is a clot formed inside a vessel when it should not be. An **embolus** is anything — usually a fragment of thrombus, but also air, fat, or debris — carried by the blood until it lodges and blocks a smaller vessel.

In the nineteenth century, Rudolf Virchow described three conditions that favour thrombosis. They remain the framework, and they are also a design checklist.



- **Deep vein thrombosis (DVT)** forms in the leg veins, typically during immobility — after surgery, in hospital, on long journeys. It is why patients wear **compression stockings**, use **intermittent pneumatic compression** devices, and are mobilised early. Stasis is the target of all of them.
- **Pulmonary embolism (PE)** is a fragment of DVT carried through the right heart into the lung arteries — obstructive shock if large enough.
- **Arterial thrombosis** forms on a damaged artery wall, usually on atherosclerotic plaque, and causes heart attacks and most strokes.
- **Atrial fibrillation** — an irregular, ineffective quivering of the upper chambers — leaves blood stagnant in the left atrium, where clots form and may travel to the brain. It is a major cause of stroke and the reason consumer devices that screen for it exist.



Virchow's triad is a checklist for anything in contact with flowing blood.

Does the design create **stasis** — dead-end volumes, crevices, recirculation behind a step, slow-flow corners? Does it **injure** the vessel wall or present a rough, foreign surface to blood? Does the patient population have **increased clotting tendency**? Catheters, cannulas, dialysis circuits, heart valves, stents, and blood pumps are all designed against these three questions, and anticoagulation is the price paid when the design cannot eliminate them. That is Tier 3 work (Appendix I), but the reasoning is a good habit for any fluid path — including the "dead legs" that harbour biofilm in Chapter 13's water systems.

11.7. Blocked pipes: atherosclerosis, heart attack, and stroke

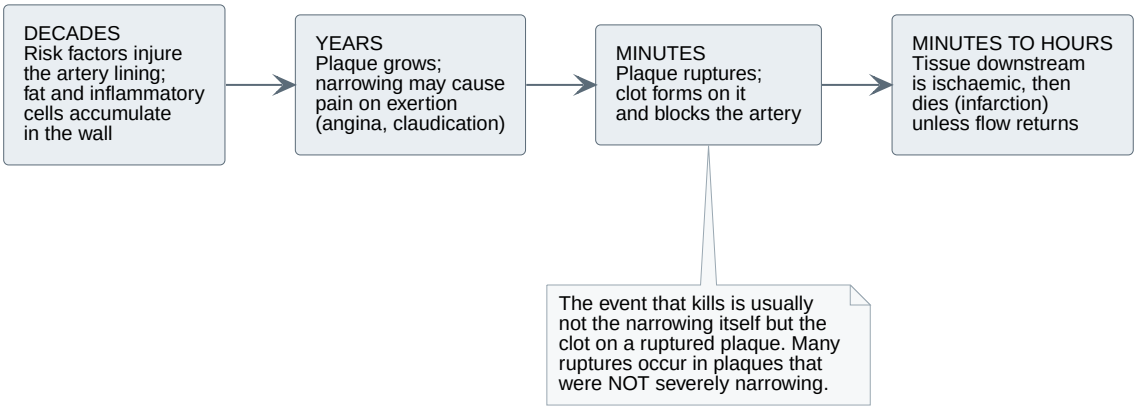
11.7.1. Atherosclerosis

Atherosclerosis is the slow build-up of **plaques** — fatty deposits, inflammatory cells, and fibrous tissue — inside the walls of arteries. It begins in early adult life and progresses over decades. Risk factors are familiar: smoking, high blood pressure, raised cholesterol, diabetes, age, male sex,

family history, inactivity, and obesity.

Plaques cause harm in two ways:

- **Slowly**, by narrowing the artery (**stenosis**). Recall Poiseuille's law from Chapter 9: resistance rises with the inverse fourth power of radius, so flow is preserved surprisingly well until the narrowing is severe and then falls steeply. Reduced flow causes pain on exertion — **angina** in the heart, **claudication** (calf pain on walking) in the legs — when demand outstrips supply.
- **Suddenly**, when a plaque **ruptures** or erodes. Its contents are intensely thrombogenic; a clot forms on it within minutes and may block the artery completely.



11.7.2. Heart attack

A **myocardial infarction** is death of heart muscle from loss of its blood supply, usually by thrombosis on a ruptured coronary plaque. The muscle downstream becomes ischaemic, then begins to die from the inside of the wall outwards over the following hours. Treatment — reopening the artery with a balloon and stent, or with clot-dissolving drugs — is a race against that clock. Damaged muscle releases **troponin** into the blood, which is why that blood test, with the ECG, is the first investigation (Chapter 16).

Dead heart muscle is replaced by scar, which does not contract. The consequences may include **heart failure** and dangerous **arrhythmias**, because scar disrupts the conduction system of Chapter 9.

11.7.3. Stroke

A **stroke** is the sudden loss of brain function from a vascular cause. Two mechanisms, with opposite treatments:

- **Ischaemic stroke** (roughly 85% in many populations) — a vessel is blocked, by a thrombus formed locally or an embolus from the heart or a neck artery. Treatment aims to **restore flow**: clot-dissolving drugs or mechanical removal of the clot through a catheter.
- **Haemorrhagic stroke** — a vessel **bursts** and bleeds into or around the brain. Clot-dissolving drugs would be disastrous.

That is why the first test is a CT scan to exclude bleeding (Chapter 16). The symptoms of both — sudden weakness of one side, facial droop, speech disturbance, loss of vision — depend on **which**

part of the brain has lost its supply, not on the mechanism. A stroke survivor's needs are therefore highly individual: weakness, spasticity, sensory loss, visual field loss, language difficulty, swallowing difficulty (Chapter 8), and cognitive change can occur in any combination — the clearest possible case for Chapter 7' rule that a diagnosis is not a specification.

11.8. Failing pumps and airways

11.8.1. Heart failure

Heart failure means the heart cannot pump enough blood to meet the body's needs at normal filling pressures. It is a syndrome — the end-point of many causes: heart attack, long-standing high blood pressure, valve disease, cardiomyopathies, arrhythmias. Clinicians often divide it by **ejection fraction** (Appendix C): **reduced** (the muscle contracts weakly; typically EF below about 40%) or **preserved** (the muscle is stiff and fills poorly, though EF looks normal).

Its features follow from the plumbing. Pressure backs up behind the failing pump: fluid accumulates in the lungs (breathlessness, especially lying flat) and in the legs and abdomen (**oedema**). Forward flow is reduced: fatigue, exercise intolerance, and reduced kidney perfusion, which in turn causes the kidneys to retain salt and water — a maladaptive feedback loop that worsens the congestion.

For a designer, heart failure means a user who tires quickly, may be breathless when lying flat, has legs that change size through the day, and is often managed with daily weighing, because a rapid weight gain is fluid. Home monitoring and telehealth for heart failure is a substantial device and software category.

11.8.2. Asthma and COPD

Both are **obstructive** lung diseases: the airways resist airflow, especially on breathing out, and spirometry shows a reduced FEV_1/FVC ratio (Chapter 16).

- **Asthma** is chronic airway inflammation with **variable, reversible** narrowing — airway smooth muscle constricts in response to triggers (allergens, cold air, exercise, infection). It often begins in childhood, and between attacks breathing may be normal.
- **COPD** (chronic obstructive pulmonary disease) is **persistent, largely irreversible** obstruction from chronic inflammation, narrowed and mucus-filled small airways, and destruction of alveolar walls (**emphysema**), which removes gas-exchange surface and the elastic recoil that drives exhalation. Smoking is the main cause worldwide; indoor smoke from cooking fuels is another.

The cornerstone of treatment for both is **inhaled** medicine — and inhaler technique is notoriously poor. A large proportion of people use their inhalers incorrectly, which is why inhaler **training devices**, spacers, and dose counters exist. That is a human factors problem as much as a pharmacological one.

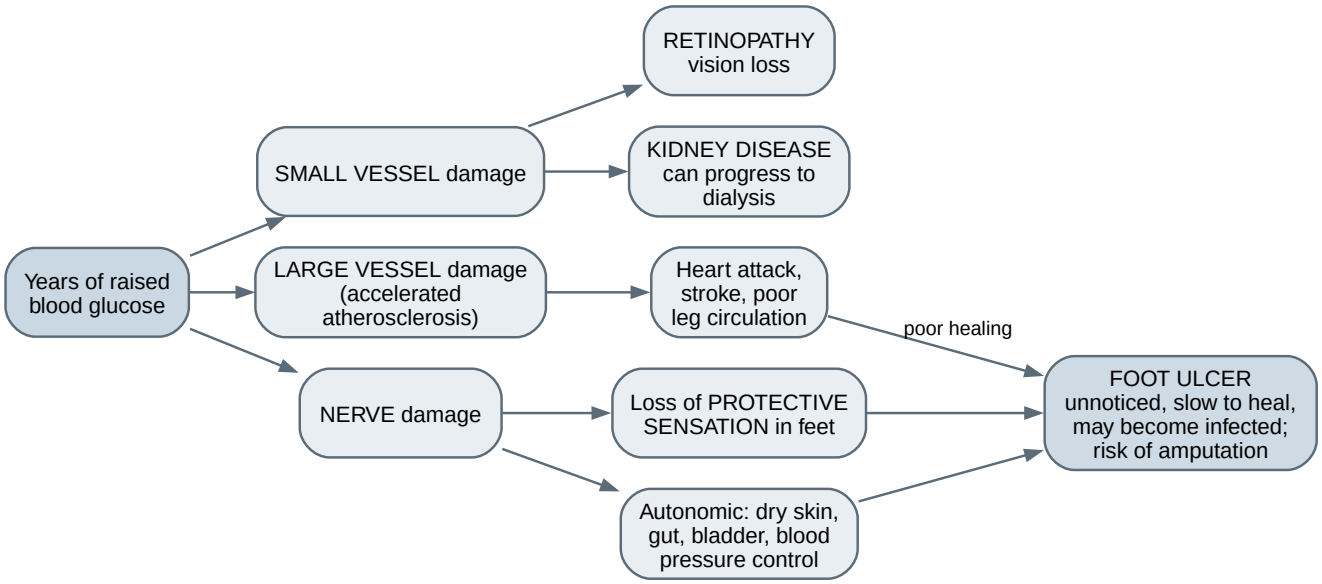
11.9. Diabetes and its complications

Diabetes mellitus is chronically raised blood glucose from failure of the insulin system (Chapter 10). Over half a billion adults live with it worldwide.

- **Type 1 diabetes** is **autoimmune** destruction of the insulin-producing beta cells of the pancreas. It usually begins in childhood or young adulthood, and insulin must be replaced for life.
- **Type 2 diabetes**, around 90% of cases, combines **insulin resistance** — tissues respond poorly to insulin — with a gradual failure of insulin secretion. It is strongly associated with excess weight, inactivity, age, and family history, and is managed with lifestyle change, tablets, injectables, and eventually often insulin.

The **HbA_{1c}** blood test measures the fraction of haemoglobin carrying attached glucose. Because red cells live about three months, it reflects average glucose over roughly the previous two to three months — an integrator, where a finger-prick meter is a point sample and a continuous glucose monitor is a time series (Chapter 16).

The danger of diabetes is mostly in its **complications**, which come from glucose damaging blood vessels and nerves over years.



The **diabetic foot** is where every mechanism in this book converges. **Neuropathy** removes the pain that would warn of a stone in the shoe or a blister. **Autonomic neuropathy** dries and cracks the skin. **Vascular disease** slows healing. **Retinopathy** may prevent the person inspecting their own feet. Raised glucose impairs the immune response, so infection takes hold. The lifetime risk of a foot ulcer in people with diabetes has been estimated at roughly one in five to one in three, and diabetes is a leading cause of lower-limb amputation.

Loss of protective sensation is screened with a **10 g monofilament** pressed against the sole: if the person cannot feel it, they cannot be relied upon to feel a developing injury. That is Chapter 7' **absence of a pain report is not evidence of absence of harm**, made into a routine clinical test.

11.10. Cancer

11.10.1. What makes a cell cancerous

Cancer is uncontrolled growth of cells that have accumulated genetic changes allowing them to escape the normal controls on division, location, and death. The changes are mostly **acquired** during life — from copying errors, chemicals such as those in tobacco smoke, radiation, some viruses — and a minority are **inherited**.

A widely used framework lists the **hallmarks** cancer cells acquire: they keep signalling themselves to grow; ignore signals to stop; resist apoptosis; divide indefinitely; recruit their own blood supply; invade and spread; reprogram their metabolism (the reason FDG-PET works, Chapter 16); and evade the immune system. Genetic instability and inflammation help them get there.

- A **benign** tumour grows locally, stays encapsulated, and does not spread. It may still cause harm by pressure — a benign brain tumour in a rigid skull can be lethal.
- A **malignant** tumour **invades** surrounding tissue and can **metastasise**: cells enter blood or lymph vessels and establish secondary tumours elsewhere — commonly lymph nodes, liver, lungs, bone, and brain. Metastasis causes most cancer deaths.

11.10.2. Staging and treatment

Staging describes how far a cancer has spread, usually by the **TNM** system: the size and extent of the primary *T*umour, involvement of lymph *N*odes, and distant *M*etastasis. Stage, with the cancer's type and molecular features, drives both prognosis and treatment. Staging uses imaging (CT, MRI, PET-CT) and tissue diagnosis (biopsy), as Chapter 16 describes.

Treatment combines several modalities, covered further in Chapter 17:

Modality	Principle
Surgery	Remove the tumour, often with surrounding tissue and nearby lymph nodes
Radiotherapy	Ionising radiation damages DNA; tumour cells repair it less well than normal cells. Delivered precisely, in fractions, from outside or from sources placed inside the body
Chemotherapy	Drugs that kill rapidly dividing cells — which is why they also affect hair, gut lining, and bone marrow
Hormone therapy	Blocks the hormones some cancers (many breast and prostate cancers) depend on
Targeted therapy	Drugs aimed at a specific molecular change in the tumour, identified by testing
Immunotherapy	Releases the brakes on the patient's immune system, or engineers immune cells, to attack the tumour

For a maker, cancer care is surprisingly rich in physical needs: radiotherapy immobilisation masks and positioning aids, prosthetics after surgery, lymphoedema garments (Chapter 9), wigs and head

coverings, central venous access ports, and a great many training models for examination and screening. Notice how few of those are rigid parts: they are thermoplastic mesh, compression textiles, sewn and fitted soft goods, cast silicone, and implantable devices made under full medical manufacturing controls.

11.11. Immunity turned inward, and wear

11.11.1. Autoimmunity

In **autoimmune disease** the adaptive immune system attacks the body's own tissues. The target determines the disease:

Condition	Target
Type 1 diabetes	Insulin-producing cells of the pancreas
Rheumatoid arthritis	Joint linings (synovium); also other organs
Multiple sclerosis	Myelin in the brain and spinal cord (Chapter 7)
Coeliac disease	The small-intestine lining, triggered by gluten
Systemic lupus erythematosus	Many tissues: skin, joints, kidneys, blood
Graves' disease, Hashimoto's thyroiditis	The thyroid

Autoimmune diseases are typically **chronic and relapsing**: periods of activity and remission. A device designed on a good day may be unusable during a flare. Many are treated with drugs that suppress the immune system — which makes the user more vulnerable to infection, and makes Chapter 13's cleanability rules more important, not less.

11.11.2. Degenerative disease

- **Osteoarthritis** is the commonest joint disease: progressive loss of articular cartilage (Chapter 6), remodelling of the underlying bone, bony spurs, and pain and stiffness. It is not simply "wear and tear" — it involves the whole joint and low-grade inflammation — but load, injury, age, and obesity are major contributors. The hands, knees, hips, and spine are most affected, and the hand changes are directly relevant to grip and pinch design.
- **Osteoporosis** is loss of bone mass and microarchitecture, so that bones fracture from minor falls — typically the wrist, spine, and hip. It is assessed by DEXA (Chapter 16), defined by a T-score of −2.5 or below, and is common in older women in particular. A hip fracture in an older person is a life-changing and often life-shortening event — the reason falls prevention is a major design and research area.
- **Neurodegenerative diseases** progressively destroy particular populations of neurons. **Alzheimer's disease**, the commonest cause of dementia, involves accumulation of abnormal proteins (amyloid plaques and tau tangles) and loss of neurons, beginning with memory. In **Parkinson's disease**, dopamine-producing neurons in the midbrain are lost; symptoms — tremor, slowness, stiffness, balance problems — typically appear only after a large proportion of them has gone. Both are covered further, as users, in Chapter 12.

11.12. Genes, inheritance, and genetic disease

11.12.1. From code to protein, briefly

Chapter 3 introduced DNA and proteins. The human genome is about 3 billion base pairs, in 23 pairs of chromosomes — one of each pair from each parent — and contains roughly 20,000 protein-coding genes. A gene is **transcribed** into messenger RNA and **translated** into a protein, three bases specifying each amino acid.

A **mutation** is a change in the DNA sequence. Most have no effect. Some change a protein's structure — and since structure is function (Chapter 3), its behaviour. Mutations can be single-base changes, insertions and deletions, or large rearrangements of chromosomes. They may be **inherited** (present in every cell, from conception) or **acquired** (arising in one cell line during life, as in most cancers).

11.12.2. Patterns of inheritance

Pattern	How it works	Examples
Autosomal dominant	One altered copy is enough. An affected parent passes it to each child with probability 1/2	Huntington's disease, many inherited cancer predispositions (e.g. **BRCA1/2**), familial hypercholesterolaemia
Autosomal recessive	Two altered copies are needed. Two unaffected **carrier** parents have a 1/4 chance with each child of an affected child, and 1/2 of a carrier child	Cystic fibrosis, sickle cell disease, spinal muscular atrophy
X-linked	The gene is on the X chromosome. Males, with one X, are affected far more often; carrier females may be mildly affected	Haemophilia A and B, Duchenne muscular dystrophy, common red–green colour vision deficiency (Chapter 7)
Mitochondrial	The small genome inside mitochondria is inherited from the mother only	Some rare disorders of energy metabolism
Multifactorial	Many genes, each with a small effect, plus environment	Most common diseases: type 2 diabetes, heart disease, most cancers, many psychiatric conditions

The last row matters most for population health. "It runs in the family" usually means multifactorial **risk**, not destiny.

11.12.3. Genetic testing and pharmacogenomics

Genetic tests range from a single known variant, through panels of genes and whole-exome sequencing, to whole-genome sequencing (Chapter 15' molecular diagnostics). Their results are **probabilistic** more often than people expect: many variants are classified as of **uncertain significance**, and the same variant may confer very different risks in different people. Results also affect relatives who never consented to be tested, which is why genetic testing is accompanied by

genetic counselling and why genomic data are among the most sensitive data there are (Chapter 20).

Pharmacogenomics uses genetic variation to predict drug response. Variation in liver enzymes such as **CYP2D6** changes how fast some drugs are activated or cleared; for example, people who convert codeine to morphine unusually fast can be harmed by ordinary doses. Testing for a particular immune-system variant before prescribing the HIV drug abacavir prevents a serious hypersensitivity reaction. It is the clearest current example of **personalised medicine** — and, like Chapter 3's narrow therapeutic index, a reminder that "the standard dose" is a population average.

11.13. Engineering bridge: failure, growth, and risk

11.13.1. Disease as failure mode — and where the analogy breaks

Engineers already think in failure modes, and the analogy is useful:

Engineering failure	Disease mechanism
Supply interruption	Ischaemia — heart attack, stroke, pressure injury
Blocked or narrowed pipe	Atherosclerosis, thrombosis, embolism, airway obstruction
Pump failure	Heart failure, cardiogenic shock
Control loop failure	Diabetes (glucose control), runaway positive feedback in shock
Sensor failure	Neuropathy — and every downstream injury it allows
Fatigue and wear	Osteoarthritis, some valve disease
Corrosion	Chronic inflammation slowly destroying the tissue it protects
Firmware corruption	Mutation
Uncontrolled replication	Cancer
Security breach	Infection
Intrusion detection attacking the system it protects	Autoimmunity

The analogy breaks in three places, and they matter. Living tissue **repairs and adapts**, so damage is not simply cumulative. Causes are **probabilistic and multifactorial**, so most individual cases cannot be attributed to one cause the way a failure analysis would like. And the "system" is the person, who has their own priorities about which risks are worth taking. Use the analogy to ask good questions; do not use it to predict outcomes.

11.13.2. Exponential growth

Bacteria in ideal conditions can double every twenty minutes or so. Cancer cells double far more slowly — weeks to months — but the arithmetic is the same:

$$N = N_0 \times 2^{t/T_d}$$

A single cell, about 10 µm across, needs roughly 30 doublings to reach a tumour of about 1 cm and a

billion cells — around the smallest size many imaging methods reliably detect. Ten more doublings reach about a trillion cells and a kilogram. So by the time a tumour is detectable, most of its doublings have already happened, and the remaining ones are few. That is the arithmetic behind screening programmes, and behind Chapter 15's warning that detection thresholds, not just test accuracy, determine what screening can achieve.

The same arithmetic, run backwards, explains why disinfection is specified in **log reductions** (Chapter 13): when populations double, killing 99% leaves a population that regrows to its original size in under seven doublings.

11.13.3. Relative and absolute risk

Every claim about a disease, a risk factor, a treatment, or a device is ultimately a statement about risk, and the same result can be presented in very different-looking ways.

Suppose a trial finds that, over five years, a heart attack occurs in **2%** of people taking a placebo and **1%** of people taking a drug.

- **Relative risk** = 1% ÷ 2% = 0.5.
- **Relative risk reduction** = 50%. ("Halves your risk!")
- **Absolute risk reduction** = 2% − 1% = **1 percentage point**.
- **Number needed to treat** (NNT) = 1 ÷ 0.01 = **100**. A hundred people must take the drug for five years for one to avoid a heart attack. Of the other ninety-nine, ninety-eight would not have had one anyway, and one has one despite the drug.

That last sentence is worth reading twice, and it is why NNT is the most honest of the four. Applied to a population at **lower** baseline risk — 0.2% instead of 2% — the same 50% relative reduction is an absolute reduction of 0.1 percentage points and an NNT of 1000.



Relative risk without baseline risk is uninterpretable. "Doubles your risk" of something with a one-in-a-million baseline is a different statement from "doubles your risk" of something with a one-in-five baseline.

When you write about a device's benefit — in a brief, a report, a funding application — state the baseline, the absolute difference, and the number needed. If you are given only a relative figure, ask for the rest. This is Chapter 15's prevalence lesson in another form.

11.14. Devices, software, and training technology

Almost every device in this book is a response to a disease. This section groups the ones whose purpose is to **monitor, manage, or prevent** the mechanisms above, rather than to diagnose them (Chapter 16) or treat them acutely (Chapter 17).

11.14.1. Clinical devices

Device	Does	Engineering note
Continuous glucose monitors and insulin pumps	Measure interstitial glucose; deliver insulin; increasingly linked as automated "closed-loop" systems	A control loop around a failed biological control loop. Sensor lag, calibration, occlusion detection, and dose safety (Chapter 10)
Home blood pressure monitors	Track hypertension, the leading modifiable risk factor for stroke and heart disease	Oscillometric (Chapter 16). Validation and cuff size are the main error sources
Inhalers, spacers, and nebulisers	Deliver drugs to the airways in asthma and COPD	Particle size, inspiratory flow, and **technique**. Dose counters and smart inhaler add-ons record use
Peak flow meters and home spirometers	Track airway obstruction	Effort-dependent measurement; coaching is part of the device
Compression stockings and intermittent pneumatic compression	Prevent DVT by reducing venous stasis	Graduated pressure profiles; the same pressure-on-skin risks as Chapter 5
Anticoagulation self-testing meters	Measure clotting time (INR) for people on warfarin	Point-of-care testing (Chapter 15) at home
Wearable rhythm monitoring	Detects atrial fibrillation from pulse or single-lead ECG	Screening populations raises Chapter 15's positive predictive value problem directly
Offloading footwear and total contact casts	Redistribute pressure from a diabetic foot ulcer so it can heal	Pressure redistribution for a user who cannot feel the pressure. Prescribed and fitted clinically — Tier 3
Heart failure telemonitoring	Daily weight, blood pressure, symptoms; some implanted pressure sensors	Trend detection rather than thresholds; alert fatigue
DNA sequencers and genotyping platforms	Read genetic variation for diagnosis, cancer profiling, and pharmacogenomics	Massive data volumes; interpretation, not sequencing, is now the bottleneck

11.14.2. Software

Software	Does	Note
Early warning scores — e.g. NEWS2	Combine routine observations into a score that triggers escalation	Simple, transparent, and widely validated. A useful model of what good clinical software looks like: explainable and tied to a defined action
Automated sepsis alerts	Predict sepsis from electronic record data	**A cautionary category.** At least one widely deployed proprietary model performed far worse on independent validation than claimed, missing many cases and alerting on many non-cases. Validate on **your** population (Chapter 15)
Cardiovascular risk calculators	Estimate ten-year risk of heart attack or stroke from risk factors	Output is **absolute** risk — exactly the quantity this chapter argues for
Diabetes data platforms and open-source closed-loop systems	Aggregate meter, CGM, and pump data; some community-built systems automate insulin dosing	A remarkable patient-led engineering community — and a vivid example of the regulatory and safety questions in Chapter 20
Genomic databases — ClinVar, gnomAD, Ensembl	Variant interpretation, population frequencies, gene annotation	**Free and public.** Look up a well-known variant and read how its significance is classified
Synthetic patient generators — Synthea [39]	Realistic fictional patient records including disease progression	The safe way to build and test disease-related software
Epidemiological tools	Risk, rate, and survival calculations	R, Python, and spreadsheet implementations abound; the arithmetic in the bridge is the core

11.14.3. Training and simulation

Training technology	Teaches	Form
Pathology specimens and virtual pathology	What diseased organs and tissues actually look like	Museum specimens; whole-slide images (Chapter 15); cast or printed replicas of diseased anatomy from imaging
Deteriorating-patient and sepsis scenarios	Recognition and escalation	High-fidelity manikins whose vital signs follow scripted trajectories — Chapter 9's mock monitor with a disease model behind it

Training technology	Teaches	Form
Diabetic foot examination models	Monofilament testing, ulcer recognition, footwear assessment	Foot models with interchangeable lesions; moulage. **Makeable with silicone casting**
Inhaler technique trainers	Correct inhalation timing and flow	Placebo devices; flow-indicating whistles. Simple physical feedback devices
Breast, testicular, and skin examination models	Finding lumps and suspicious lesions	Silicone models with embedded masses of graded size and stiffness — Chapter 5' material knowledge applied
Stroke recognition and rehabilitation training	Rapid assessment; rehabilitation technique	Video, standardised patients, and rehab equipment practice
Pedigree and genetic counselling exercises	Inheritance reasoning, communicating risk	Paper and software exercises; role play



The maker's way into this chapter is through the training models.

Examination models with **embedded findings** — breast lumps of graded size and depth, diabetic foot lesions, skin lesions, swollen joints — are wanted, consumed, and directly built on the silicone and casting skills of Chapter 5, with machined or printed moulds, fabric reinforcement, and inserts of graded stiffness. Electronics can add what silicone cannot — a palpable pulse, a temperature difference, a sensor that records how firmly the learner pressed. The disease mechanisms in this chapter tell you **what** the finding should be: how large, how stiff, how deep, and how it differs from the normal tissue around it. Agree each with an educator, exactly as Chapter 21 requires.

11.15. Worked example: a foot-inspection aid

A diabetes specialist podiatrist asks whether you could make a simple aid to help patients inspect the soles of their own feet each day. Daily inspection is a standard recommendation, because an ulcer caught early heals far more readily. The aid would not touch any wound. It is an inspection tool, not a treatment.

The naive brief: "a mirror on a stick". It would take an afternoon.

Decompress the diagnosis. Work through each mechanism in this chapter and ask what it implies for the **user**:

Mechanism	What it means for the user	Requirement it generates
Peripheral neuropathy (feet)	Cannot feel a sharp edge, a hot surface, or a pressure point on the foot	No sharp edges or corners anywhere the aid might contact the foot; nothing that heats
Peripheral neuropathy (hands) — common too	Reduced grip feedback and dexterity	Large, high-friction grip; operable without fine pinch (Chapter 6, Chapter 7)
Retinopathy, cataract	May not see a small red mark in a mirror at arm's length	Magnification and good lighting; or a camera with a large, high-contrast display; legibility calculated, not guessed (Chapter 8)
Obesity, osteoarthritis, back pain	Cannot bend to bring the foot close to the eyes	Works with the foot on the floor; long reach without bending
Balance problems, orthostatic hypotension (autonomic neuropathy)	Risk of falling when lifting a foot or bending	Usable seated; no need to stand on one leg
Immune impairment	Infection risk	Cleanable; no crevices (Chapter 13)
Cognitive change, low mood — both more common with long-standing diabetes	May not remember or be motivated to inspect daily	Simple, quick, part of an existing routine; possibly a reminder — a question for the user, not an assumption

The "mirror on a stick" fails at least four of those rows. A better concept — perhaps a floor-standing mirror with integral lighting and magnification that the seated user slides their foot over, or a phone holder at floor level — emerges **from the mechanisms**, not from a sketch.

Then stop and check scope. The moment someone suggests "and it could photograph the foot and flag ulcers automatically", the aid becomes software making a clinical claim (Chapter 20), with Chapter 15's full burden of sensitivity, specificity, and prevalence. That may be worth doing, but it is a different project with a different review route. Record the suggestion; do not quietly absorb it.

And involve the users. The table is a list of hypotheses about what people with diabetes may need. Which of them apply to **these** patients, and which matter most, is established with the podiatrist and with the patients themselves (Chapter 18).

11.16. Workshop exercise

1. **Decompress a diagnosis.** Choose one condition from this chapter — stroke, heart failure, COPD, rheumatoid arthritis, or Parkinson's disease. List every mechanism involved and, for

each, what it implies for someone using a simple everyday object such as a kettle, a door handle, or a phone. Aim for at least eight implications. Then mark which ones you would need to verify with a clinician and a user.

2. **Draw the loop.** Sketch the feedback loops in heart failure: reduced output → kidneys retain salt and water → more fluid → more congestion. Mark which loops are negative feedback and which have become self-reinforcing.
3. **Virchow on paper.** Take a drawing of any fluid path — a tube junction, a three-way tap, a pump housing (real or from an online datasheet). Mark every place that could create stasis, every surface that is rough or abrupt, and write one sentence on how each could be improved. This is a reasoning exercise; nothing goes near blood.
4. **Exponential arithmetic.** Calculate how many doublings take one 10 µm cell to a 1 cm tumour (assume spheres and constant cell size). Then calculate how long that takes at doubling times of 60 and 200 days. Write two sentences on what this means for screening.
5. **Rewrite a headline.** Find a news story or press release claiming that something "cuts the risk" of a disease by a percentage. Find the baseline risk in the underlying study, and rewrite the claim as an absolute risk reduction and an NNT. If the baseline is not available, say so.
6. **Draw a pedigree.** Invent a three-generation family with an autosomal recessive condition affecting one grandchild. Mark the carriers who must exist, and calculate the probability that an unaffected sibling of the affected child is a carrier. (Hint: it is not 1/2.)

11.17. Questions

1. Distinguish aetiology from pathogenesis, and a sign from a symptom, with an example of each.
2. Why is the brain so much less tolerant of ischaemia than skin?
3. Contrast necrosis and apoptosis. Which provokes inflammation?
4. A strap fits at 9 a.m. and is too tight at 5 p.m. on a person with an inflammatory joint disease. Give two mechanisms.
5. Name the four types of shock and the part of the circulation each represents. Why does a normal blood pressure not exclude shock?
6. State Virchow’s triad, and apply each element to the design of something in contact with flowing blood.
7. Why can a heart attack occur in an artery that was not severely narrowed?
8. Why must bleeding be excluded before treating a stroke, and how is that done?
9. List four complications of diabetes and explain why they converge on the foot.
10. What distinguishes a malignant tumour from a benign one? What does TNM describe?
11. Two unaffected parents have a child with cystic fibrosis. What is the chance their next child is affected? Is a carrier?
12. A device "reduces falls by 40%". The baseline fall rate in the trial population was 10% per year. Calculate the absolute risk reduction and the NNT, and say what else you would want to know.

11.18. Answers and reasoning

1. **Aetiology** is the cause (e.g. smoking); **pathogenesis** the mechanism by which it produces damage (chronic airway inflammation and alveolar destruction leading to COPD). A **symptom** is reported by the person (breathlessness); a **sign** is found by an examiner or a measurement (a low oxygen saturation, wheeze heard with a stethoscope).
2. The brain has a very high, continuous oxygen demand — around a fifth of resting consumption for about 2% of body mass — and almost no energy reserve, so its cells fail within minutes. Skin has a low metabolic rate and can survive hours of reduced flow, which is also why pressure injury can develop without immediate symptoms.
3. **Necrosis** is uncontrolled death from injury: cells swell and rupture, releasing contents that **provoke inflammation**. **Apoptosis** is programmed, orderly self-destruction without inflammation, used routinely to remove unwanted cells. Failure of apoptosis contributes to cancer.
4. Any two of: inflammation causes **swelling** (leaky vessels, fluid in tissue) that may increase with activity; **dependent oedema** accumulates in limbs through the day with gravity and reduced movement; **sensitisation** makes the same pressure more uncomfortable. Design for dimensional change, adjustability, and a reviewed fit.
5. **Hypovolaemic** — volume; **cardiogenic** — the pump; **distributive** — the vessels (resistance); **obstructive** — a physical blockage to flow. Compensation (raised heart rate, vasoconstriction) can **maintain** blood pressure while tissue perfusion is already inadequate; a falling blood pressure is a late sign.
6. **Stasis**, **endothelial injury**, **hypercoagulability**. A design should avoid dead spaces and recirculation zones (stasis); avoid damaging the vessel wall and present smooth, compatible surfaces (injury); and recognise that some users' blood is more prone to clot, which no design can remove (hypercoagulability) — hence anticoagulation when the first two cannot be eliminated.
7. Because most heart attacks are caused by **rupture** of a plaque and a clot forming on it, not by gradual narrowing. Plaques that rupture are often not severely stenosing, so a test that measures narrowing does not capture all the risk.
8. Ischaemic and haemorrhagic stroke cause similar symptoms but need opposite treatment; clot-dissolving drugs given into a bleed would make it worse. A non-contrast CT scan of the head shows fresh blood, so it is done first.
9. Any four of: peripheral **neuropathy** (loss of protective sensation); **autonomic neuropathy** (dry, cracking skin); **peripheral arterial disease** (poor healing); **retinopathy** (cannot inspect the foot); impaired **immune response** (infection); **kidney disease**. On the foot, unnoticed injury, poor healing, and infection combine — leading to ulceration and, in the worst cases, amputation.
10. A malignant tumour **invades** surrounding tissue and can **metastasise** to distant sites through blood or lymph; a benign tumour grows locally without invading or spreading. TNM describes the extent of the primary *T*umour, lymph *N*ode involvement, and distant *M*etastasis.
11. Both parents must be carriers. Each child has a **1/4** chance of being affected, **1/2** of being a carrier, and 1/4 of inheriting neither altered copy. (For an **unaffected** sibling, the chance of being a carrier is 2/3.)

12. Absolute risk reduction = 10% × 0.40 = **4 percentage points** (from 10% to 6%). NNT = 1 ÷ 0.04 = **25** people using the device for a year to prevent one fall. You would also want: the population and setting of the trial (is it like your users?), how falls were counted, the duration, adverse effects, whether injuries — not just falls — were reduced, and the confidence interval.

11.19. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Test accuracy and risk — relative and absolute risk, NNT, and incidence against prevalence.
- Genetics — inheritance and pedigree risk.
- Biolab — exponential growth and doubling times.
- Heart, vessels and oxygen — the oxygen dissociation curve and its shifts.

11.20. Video companion and reading

- **Core videos:** Crash Course Anatomy & Physiology episodes on the immune system and on the cardiovascular system, from the series index [6], watched now with mechanisms of failure in mind.
- **Core reading:** OpenStax [1], selectively — the "Disorders of…" and "Homeostatic Imbalances" feature boxes throughout the systems chapters, together with:
  - 3.6 Cellular Differentiation and the cell-cycle material of chapter 3, for cancer
  - 18.1–18.5 for blood and clotting, in chapter 18 The Cardiovascular System: Blood
  - 21.1–21.3 for inflammation and immunity, in chapter 21 The Lymphatic and Immune System
  - 28.7 Patterns of Inheritance
- **Context:** a trusted public-facing medical encyclopaedia, such as MedlinePlus [3], for any condition you meet in a project. Read the patient description first; it tells you what the person experiences.
- **Risk literacy:** any short guide to absolute versus relative risk and number needed to treat — then practise on real headlines, as in the exercise.

11.21. Exit criteria

- ▢ Explain the path from aetiology to outcome, and the difference between signs and symptoms, using one condition you know.
- ▢ Describe ischaemia, necrosis, apoptosis, inflammation, thrombosis, and embolism in one or two sentences each.
- ▢ Name the four types of shock and explain why blood pressure is a late indicator.
- ▢ Apply Virchow’s triad to a drawing of a fluid path and record your findings.
- ▢ Draw the diabetic-foot convergence diagram from memory and derive at least five design requirements from it.

- ▢ Decompress one diagnosis into at least eight user implications, marked as hypotheses to verify.
- ▢ Convert a relative risk claim into an absolute risk reduction and an NNT, stating the baseline.
- ▢ Work an autosomal recessive inheritance problem correctly.

11.22. What this chapter does not establish

- **You cannot diagnose anyone.** Recognising the mechanism of a disease is not recognising the disease in a person. Diagnosis combines history, examination, and investigation, and is clinical work.
- **You cannot advise anyone about their condition, their risk, or their treatment.** Not even informally, and not even when asked by someone you are designing for. Refer the question to their clinical team.
- **A diagnosis decompressed by you is a list of hypotheses.** Which mechanisms affect a particular person, and how much, is established with that person and their clinicians.
- **You cannot interpret genetic test results.** Variant interpretation and genetic counselling are specialist work, and results carry implications for relatives.
- **The failure-mode analogies are teaching tools.** Biology repairs, adapts, and varies in ways that machines do not; the analogy suggests questions, not answers.
- **You have not learned pathology.** Each disease named here is a field. This chapter is a map of the mechanisms they share.

Chapter 12. Across the lifespan, and the mind

Week 15 · approximately 4–6 hours

12.1. Why this chapter exists

Almost every number in this book so far has quietly described **one person**: a healthy adult of working age, somewhere near the middle of the distribution. The reference tables say "typical adult". The joint ranges are adult ranges. The vital signs are adult vital signs.

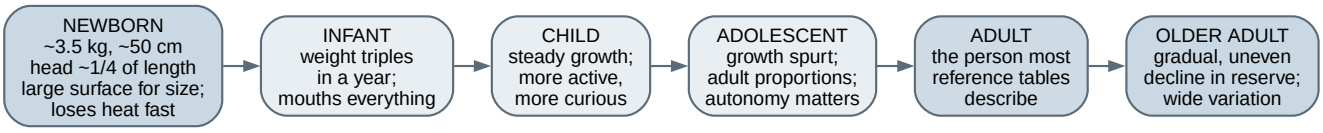
Most of the people who use healthcare are not that person. Hospitals are full of the very young and the very old, and of people whose thinking, attention, or mood is affected by illness, medication, stress, or a long-standing condition. A design that works for the person who built it may fail a three-year-old, an eighty-five-year-old, or anyone on the worst day of their life — and it will usually fail them in ways you did not think to test.

This chapter covers three things:

- **Children** — who are not small adults. Their proportions, physiology, and behaviour differ in ways that change what is safe and what works.
- **Older adults** — whose bodies change system by system, gradually and very unevenly, and who make up a large and growing share of healthcare's users.
- **The mind** — attention, memory, cognitive load, dementia and delirium, mental health, neurodiversity, and stress. Every interface is operated by a mind, and every safety measure in a mental-health setting is shaped by one particular risk.

The common thread is a design principle you have met before in Chapter 7: **design for the range of real users, not the middle of it**. Here that range is wider than anywhere else in the book.

12.2. Children are not small adults



12.2.1. Size and proportion

A newborn weighs about 3.5 kg and is about 50 cm long. Weight roughly doubles by five months and triples by the first birthday. But growth is not uniform enlargement: the **proportions** change.

- **The head is large.** At birth it is about a quarter of body length; in an adult, about an eighth. A baby's centre of mass is high, and a large occiput (the back of the skull) flexes the neck forward when the baby lies on its back — which matters to anything that positions an infant's head and airway.

- **The limbs are short** relative to the trunk, and lengthen through childhood. The adolescent growth spurt changes proportions again, quickly.
- **Surface area is large relative to mass.** A newborn has roughly two to three times as much skin surface per kilogram as an adult. Heat, water, and anything absorbed through the skin move proportionally faster.

The practical consequence is that a paediatric device cannot be designed by scaling an adult one. A mask, a cuff, a strap, a seat, or a handle that is "70% of adult size" fits no actual child, because no child is a 70% adult in every dimension.

12.2.2. The airway

The paediatric airway differs from the adult airway of Chapter 9 and Chapter 8 in ways that make it both easier to obstruct and harder to manage:

- the tongue is large relative to the mouth;
- the larynx sits higher in the neck;
- young infants breathe predominantly through the nose, so a blocked nose — from secretions or from a badly placed device — is a much bigger problem than in an adult;
- above all, the airway is **narrow**.

That last point has a sharp engineering consequence. From Poiseuille's law (Chapter 9), resistance to laminar flow rises with the inverse **fourth power** of the radius. Take an infant airway about 4 mm across — radius 2 mm. One millimetre of swelling all round reduces the radius to 1 mm:

$$\mathbf{R_{new} / R_{old} = (2 / 1)^4 = 16}$$

The same millimetre of swelling in an adult airway 16 mm across changes the radius from 8 mm to 7 mm, a resistance increase of $(8/7)^4 \approx 1.7$. The same insult that is trivial in an adult can be life-threatening in an infant. Anything near a child's face, nose, or neck deserves that arithmetic.

12.2.3. Physiology by age

Age	Heart rate (beats/min)	Breathing rate (breaths/min)	Note
Newborn	~110–160	~30–60	Fast rates are normal, not alarming
Infant (to 1 year)	~100–160	~30–40	
Toddler (1–2 years)	~95–150	~25–35	
Pre-school (2–5 years)	~95–140	~25–30	
School age (5–12 years)	~80–120	~20–25	
Adolescent and adult	~60–100	~12–20	Chapter 9

These are typical orientation ranges; published paediatric ranges differ between sources, and the differences are exactly why clinical tools state which reference they use. Blood pressure is lower in children and rises through childhood.

The design point is that **an alarm limit is a function of age**. A heart rate of 150 is unremarkable in a newborn and alarming in an adult. Software that displays, flags, or simulates vital signs (Chapter 9) without an age-appropriate reference is wrong for every child it meets.

Other differences matter just as much:

- **Body water** is about 75–80% of body mass at birth, against 50–60% in adults, and young infants turn over fluid quickly. Dehydration develops fast.
- **Blood volume** is around 80–90 mL per kilogram in a newborn — so a 3.5 kg baby has only about 300 mL of blood in total. Repeated blood samples are a real proportion of it, which is why neonatal sampling devices are designed to minimise volume.
- **Thermoregulation** is limited. Newborns cannot shiver effectively; they generate heat from **brown fat** and lose it rapidly across their large surface. Incubators and radiant warmers exist for this reason, and an exposed baby on a cold surface cools within minutes.
- **Skin** in newborns, and especially in babies born early, is thinner and less of a barrier than adult skin (Chapter 5). It absorbs chemicals more readily and is easily injured by adhesives: removing a dressing or a sensor's tape can strip the outer layer. Medical adhesive–related skin injury is a recognised neonatal problem, and a design concern for every sensor that sticks on.

12.2.4. Development and behaviour

Children change in what they **do**, not only in size. Roughly:

- sitting unsupported around six months, crawling and pulling to stand later, walking around the first birthday;
- a fine **pincer grasp** — picking up small things between thumb and finger — from around nine to ten months;
- **mouthing** everything within reach, most intensely in the first two years;
- climbing, opening, unscrewing, and testing everything as soon as they can.

Milestones vary widely and are not targets. The design point is about **capability racing ahead of judgement**. A small part that a six-month-old could not pick up, a two-year-old can pick up and swallow (Chapter 8). A cord that an infant cannot reach, a toddler can reach and wrap around a neck. Child-resistant closures, small-parts limits, cord-length limits, and finger-entrapment rules in toy and nursery standards all come from this race.

12.2.5. Consent, assent, and the family

A child cannot usually give legal consent, so a parent or guardian does. But good practice also seeks the child's own **assent** — their agreement, in language they understand — and older children may be able to consent for themselves, depending on their understanding and the jurisdiction. For a maker, this shapes any user research with children (Chapter 18): you are designing with **two** users, the child and the carer, whose needs may differ, and any contact with children runs through the

host's safeguarding process without exception.

12.2.6. Why paediatric devices are a field of their own

- **Weight-based everything.** Drug doses, fluid volumes, ventilator volumes, and defibrillation energies are calculated per kilogram. A decimal-point slip gives a **tenfold** error, and tenfold dosing errors are a recognised paediatric hazard. Software that calculates doses must make that error hard to commit and easy to see (Chapter 19).
- **Many sizes.** A paediatric unit stocks airways, masks, cuffs, and cannulas in many sizes, and choosing the right one under pressure is itself a task. Colour-coded, length-based systems — a tape laid alongside the child that reads off an estimated weight and the matching equipment sizes — exist to make that choice fast and reliable.
- **Small markets.** Children are a small market for any one product, so many devices are adult devices used "off-label" or adapted. That is a known gap, and one where careful, well-documented adaptation by makers working with paediatric teams has real value.

12.3. Pregnancy, briefly

Chapter 10 and Chapter 9 noted that pregnancy changes almost every system. A few changes have direct design consequences:

- **Blood volume** rises by roughly 40–50% and cardiac output by a similar order. Resting heart rate rises.
- **Lying flat on the back** in later pregnancy can let the uterus compress the large vessels returning blood to the heart, lowering cardiac output — which is why pregnant people are positioned tilted to one side, and why a couch, table, or positioning aid that assumes a supine adult is not neutral.
- **Ligaments soften, the centre of mass moves forward**, and balance changes, so trips and falls become more likely and more consequential.
- **Breathing** changes as the diaphragm is pushed upward; breathlessness on exertion is common.
- Imaging that uses ionising radiation is weighed especially carefully (Chapter 16).

Two people are affected by any intervention in pregnancy, and most devices and drugs are validated on neither.

12.4. Older adults

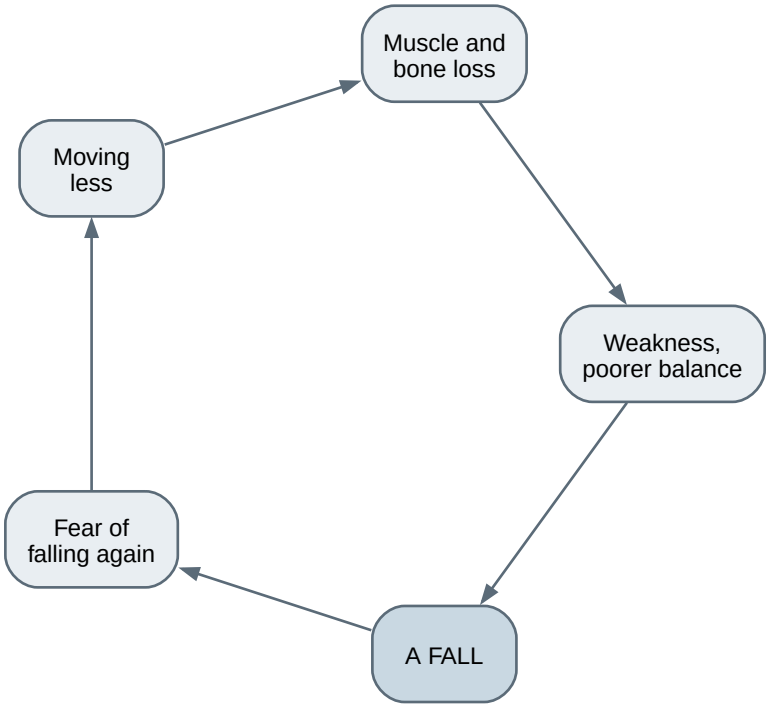
12.4.1. Ageing is gradual, uneven, and individual

Ageing is not a disease, and older adults are not a uniform group. Two eighty-year-olds can differ more than two thirty-year-olds: one runs marathons; the other has multiple chronic conditions and needs help to stand. What ageing **does** reliably reduce is **reserve** — the margin each system has between its everyday load and its maximum. With reduced reserve, a small additional stress (an infection, a new medicine, an unfamiliar ward, a night without sleep) can tip a person from coping to not coping.

System	Typical change with age	Design consequence
Skin (Chapter 5)	Thinner dermis; the junction between epidermis and dermis flattens, so layers separate more easily; slower healing; reduced sensation	Skin tears from adhesives, straps, and edges; pressure injury develops faster
Muscle	Gradual loss of mass and strength — **sarcopenia** — accelerating in later life; power (force × speed) falls faster than strength	Lower grip and pinch forces; slower rise from a chair; fatigue
Bone (Chapter 6)	Bone mass falls; especially rapid in women after menopause. **Osteoporosis** makes fracture from a minor fall likely	A fall that would bruise a young adult can break a hip or wrist
Joints	Osteoarthritis is common; range of motion reduced; stiffness	Reaching, twisting, and fine manipulation become harder
Vision and hearing (Chapter 8)	Presbyopia universal; cataract, macular degeneration, glaucoma more common; high-frequency hearing loss	Label size and contrast; alarm pitch; redundant coding (Chapter 7)
Balance and gait	Slower walking, shorter steps, slower reactions, reduced proprioception	Falls; trip hazards; time limits that assume young walking speeds
Kidneys (Chapter 10)	Filtration rate falls on average with age, very variably	Drugs cleared by the kidneys accumulate
Thirst and temperature	Blunted thirst; reduced ability to regulate temperature	Dehydration and heat illness occur more readily
Cognition	Processing speed slows; most knowledge and skills are retained; dementia becomes more common but is **not** normal ageing	Fewer steps, more time, clearer feedback

12.4.2. Frailty

Frailty is the state of reduced reserve across several systems at once, so that a person is vulnerable to a sudden, disproportionate decline after a minor event. One widely used definition identifies it from five features — unintentional weight loss, exhaustion, weakness (low grip strength), slow walking speed, and low physical activity — with three or more indicating frailty [49]. Note that two of the five are **measurements** — grip force and gait speed — which is why simple, reliable measuring devices matter in this field.



The cycle above is why falls matter beyond the injury itself, and why devices that help people stay **active and confident** — not merely devices that detect a fall — are valuable.

12.4.3. Falls

About a third of people over 65 living in the community fall each year, and the proportion rises with age. Falls are a leading cause of injury, of hip fracture, and of loss of independence. Common contributors are muscle weakness, balance and gait problems, poor vision, some medicines (particularly sedatives and those that lower blood pressure), needing the toilet urgently at night, unfamiliar surroundings, and hazards: loose cables, poor lighting, slippery floors, beds and chairs at the wrong height.

Most of those contributors are **environmental or design** factors. That is the maker's opening: grab rails, stable seating, lighting, contrasting edges on steps, raised toilet seats, bed and chair heights, and devices that do not create trip hazards of their own. A charging cable trailing across the floor from a bedside device is a fall waiting to happen.

12.4.4. Multimorbidity and polypharmacy

Older adults often live with several long-term conditions at once — **multimorbidity** — and take several medicines; five or more is commonly called **polypharmacy**. Each medicine may be reasonable on its own; together they interact, add side effects, and become hard to manage. Ageing also changes how drugs are handled (Chapter 3): kidney clearance falls, the proportion of body fat rises and of water falls (changing how drugs distribute), and the brain becomes more sensitive to sedatives.

Two design consequences:

- **Medication management is a daily task with real failure modes** — pill organisers, dispensers, reminder systems, and packaging that can be opened by arthritic hands but not by children (a genuine tension, not an easy specification).

- **A guideline written for one condition may not fit a person with five.** Software and devices built around a single disease can conflict with each other in the same person.

12.4.5. Designing for older users

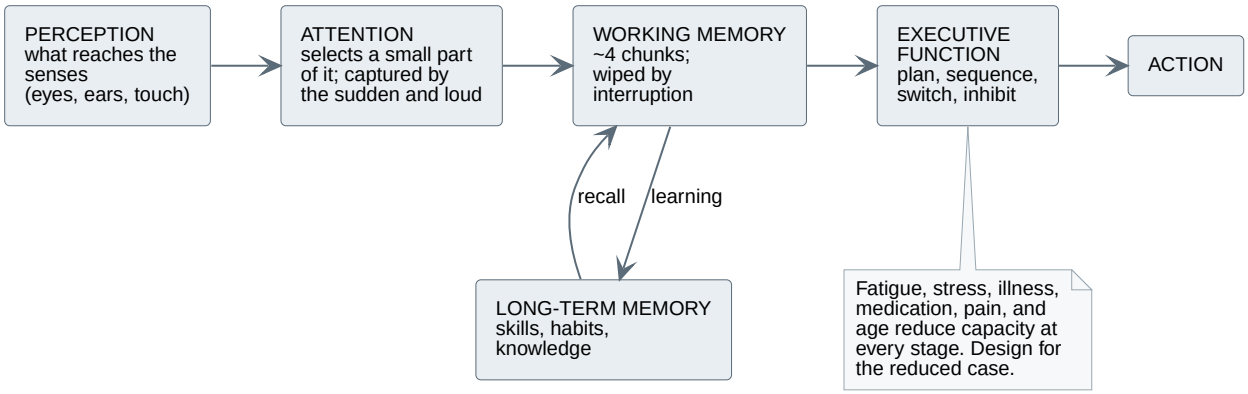
- **Text and contrast** sized for the actual viewing distance with a reserve (Chapter 8).
- **Sound cues at lower frequencies**, always duplicated visually.
- **Forces**: low operating forces, large levers, no fine pinch, no simultaneous two-handed actions.
- **Time**: no short timeouts; walking-speed assumptions that include slow walkers. (Many pedestrian crossings assume about 1.2 m/s; a substantial share of older people walk more slowly than that.)
- **Steps**: fewer, each with a clear result; recoverable errors.
- **Respect**: older adults are adults. Designs that look childish or stigmatising are rejected — often rightly — and unused devices protect no one.

12.5. The mind

12.5.1. Cognition: the resources every interface spends

Three components of cognition matter most to design:

- **Attention** — what is being processed now. It is limited and easily captured by anything sudden, loud, or moving. Sustained attention to a dull task declines within tens of minutes: the **vigilance decrement**, which is one reason monitoring by eye alone fails.
- **Working memory** — the mental scratchpad holding what you are currently using. It is small: classic estimates of "seven plus or minus two" items have been revised downward, and a capacity of **around four meaningful chunks** is now widely cited [50]. It is also fragile: an interruption can wipe it.
- **Executive function** — planning, sequencing, switching between tasks, inhibiting the wrong action. It is the first thing to fail under fatigue, stress, illness, and many neurological conditions.



This is the **cognition** stage of Chapter 7' perception–cognition–action–feedback model, opened up. **Cognitive load** is the demand a task places on these limited resources. An interface that requires a user to remember a code from one screen while entering it on another, to hold a sequence of steps in mind, or to translate units in their head, spends working memory that a stressed, tired, ill, or

older user may not have.

The cheapest cognitive-load reductions are well known: put information where it is needed (**knowledge in the world**, not in the head); show the current state; use one action per step; make the next step obvious; use consistent conventions; prevent errors rather than report them; and allow recovery.

12.5.2. Dementia and delirium

These two are often confused and the difference matters.

	Dementia	**Delirium**
Onset	Gradual, over months to years	Sudden, over hours to days
Course	Progressive	Fluctuates; often worse at night; usually reversible if the cause is treated
Attention	Often preserved early	Impaired — the core feature
Causes	Alzheimer's disease (most common), vascular disease, and others	Infection, medication, dehydration, pain, surgery, constipation, sleep loss, unfamiliar surroundings — often several at once
Urgency	Long-term care and support	**A medical emergency** that is frequently missed

Delirium is common in older people in hospital and is often the **quiet** kind — the person becomes withdrawn and drowsy rather than agitated, and is easily overlooked. People with dementia are at much higher risk of it. Some of what helps is environmental and low-tech: visible clocks and calendars, good lighting and darkness at night, glasses and hearing aids actually in use, familiar objects, minimal moves between wards. That makes it a design topic, not only a medical one.

Dementia-friendly design principles overlap heavily with good design for everyone: clear contrast between floor, walls, and fittings; avoiding patterned or shiny floors that can be misread as steps or water; signs with pictures and words at eye level; doors and controls that look like what they are; and objects that are familiar rather than novel.

12.5.3. Mental health

Mental health conditions are common — surveys in many countries find that around one adult in five experiences one in a given year — and they coexist with physical illness far more often than chance would predict. Briefly:

- **Depression** — persistent low mood, loss of interest, low energy, disturbed sleep and appetite, sometimes thoughts of suicide.
- **Anxiety disorders** — excessive worry or fear that interferes with daily life, including panic attacks and phobias.
- **Psychosis** — experiences that are not shared by others, such as hearing voices or holding

beliefs firmly that others do not share, as in some forms of schizophrenia or bipolar disorder.

- **Eating disorders, personality disorders, substance use, and post-traumatic stress** — each with its own specialist services.

Language matters. Say "a person with schizophrenia", not "a schizophrenic"; "died by suicide", not "committed suicide"; "a person living with dementia", not "a dementia sufferer". Clinical teams will notice, and so will the people you design with.

12.5.4. Learning disability, autism, and neurodiversity

A **learning disability** (the term used in UK healthcare; "intellectual disability" elsewhere) is a significantly reduced ability to understand new or complex information and to learn new skills, starting in childhood, with lasting effects on everyday life. People with learning disabilities experience well-documented health inequalities, including from information and devices they cannot use.

Autism is a lifelong difference in how people communicate, interact, and experience the world. Many autistic people — many prefer that identity-first phrasing — experience **sensory** differences: heightened or reduced sensitivity to noise, light, touch, smell, or textures. **Neurodiversity** is the broader idea that variation in how brains work, including autism, ADHD, and dyslexia, is part of human diversity rather than a set of defects.

Design consequences are concrete: **easy-read** information combining short sentences and pictures; predictable sequences with no surprises; the option to preview what will happen; control over sound and light; avoiding sudden alarms where possible; materials and textures that do not distress. A noisy, bright clinical environment is itself a barrier.

12.5.5. Stress and performance

Clinical staff and patients are often operating under stress. Moderate arousal can sharpen performance on simple tasks, but high stress narrows attention ("tunnel vision"), shrinks effective working memory, and pushes people toward habitual actions — the right thing if the habit is right, the wrong thing if the device works differently from the one they trained on. The rough inverted-U relationship between arousal and performance is often attributed to Yerkes and Dodson; it is a useful picture rather than a precise law.

Design implications for anything used in an emergency:

- **Make the critical action the most obvious one** and the dangerous action hard to do by accident.
- **Keep conventions consistent** with what users already know; novelty costs most when stress is highest.
- **Reduce alarms to those that need action** — Chapter 16 and Appendix D' alarm standard address **alarm fatigue** for this reason.
- **Assume interruptions** and make it easy to see where a task was left.

12.5.6. Ligature risk: a design requirement from mental-health settings

In inpatient mental-health settings, a leading method of suicide is hanging, and the points used are often ordinary fixtures: door handles, hinges, curtain and shower rails, coat hooks, taps, pipework, window fittings, and cords. A **ligature point** is anything a cord or fabric could be tied or looped around and that could bear a person's weight; a **ligature** is anything that could be used as the cord.

Mental-health buildings are therefore designed and furnished to **anti-ligature** principles [51]: rounded, sloped, and flush surfaces with nothing to loop over; fittings designed to **collapse or release** under a low load; tamper-resistant fixings; and, in some areas, door-top alarms that detect weight on a door. Wards carry out regular **ligature risk assessments** of every fixture and object.

For a maker this is not abstract. **Any object you make that will go onto a mental-health ward** — a training model, a sensory aid, a piece of furniture, a device with a cable or a strap — must be assessed against that environment before it is introduced. Cords, cables, lanyards, straps, and plastic bags are ligature risks; rigid parts may be ligature points or may be broken into sharp fragments. The ward team and their risk assessment decide; a breakaway design, as for lanyards (Chapter 8), is often part of the answer but rarely all of it.

12.6. Engineering bridge: scaling, distributions, and limits

12.6.1. Scaling with body size: weight versus surface area

Many quantities in physiology do not scale in proportion to body mass. Metabolic rate, for example, scales roughly with mass to the power three-quarters (**Kleiber's law**), so small bodies burn more energy **per kilogram** than large ones. Heat loss scales with surface area. That is why some drug doses and physiological indices are normalised to **body surface area** (BSA) rather than weight.

A commonly used estimate is the **Mosteller formula**:

BSA (m²) = √(height (cm) × weight (kg) / 3600)

Work two examples:

- Adult, 170 cm, 70 kg: √(170 × 70 / 3600) = √3.31 ≈ **1.82 m²**, or 0.026 m² per kg.
- Newborn, 50 cm, 3.5 kg: √(50 × 3.5 / 3600) = √0.049 ≈ **0.22 m²**, or 0.063 m² per kg.

The newborn has about **2.4 times** the surface area per kilogram. Heat, water, and anything applied to the skin are lost or absorbed correspondingly faster — which is the arithmetic behind incubators and behind caution about skin products in neonates. Which normalisation a given dose uses is a clinical and regulatory decision; the point for you is that **"per kilogram" is a modelling choice, not a law of nature**, and software must state which model it implements.

12.6.2. Growth charts are distributions

A child's height and weight are plotted on **growth charts** [47] showing **percentile** (centile) curves: a child on the 25th centile for weight is heavier than a quarter of children of the same age and sex

in the reference population.

Two lessons from statistics carry straight over from Chapter 1 and Chapter 15:

- **Being on a low or high centile is not by itself abnormal.** By construction, 2% of healthy children lie below the 2nd centile.
- **The trajectory matters more than the point.** A child steadily following the 9th centile may be fine; a child who drops from the 75th to the 9th over a few months needs attention. A single measurement cannot show a trend — and a measurement error (wrong scale, clothes, nappy, a misread decimal) can create a false one.

12.6.3. Designing across the whole range

Anthropometric design traditionally targets the **5th to 95th percentile** of the intended user population (Chapter 6). Two warnings when the population spans ages:

- **The range widens enormously.** Hand size, reach, grip force, and height across children, adults, and older adults vary by factors, not percentages. A single fixed design covering all of them is usually impossible; adjustability, multiple sizes, or a deliberately narrowed user population — **stated in the requirements** — are the honest options.
- **Percentiles do not combine.** A person at the 50th percentile for height is not at the 50th for arm length, grip strength, and hand width simultaneously. Designing for a "50th-percentile person" designs for nobody.

And remember that excluding the 5% at each end of one dimension excludes 10% of people on that dimension alone; across several independent dimensions, a "5th–95th" design can exclude far more than 10% of users. The people excluded are disproportionately children, older adults, and disabled people.

12.6.4. Choice time and working memory

Hick's law describes how decision time grows with the number of equally likely options:

$$T = a + b \times \log_2(n + 1)$$

where **a** and **b** are constants for a person and task. The logarithm is good news — doubling the options adds a fixed increment, not double the time — but the constants are not fixed across people: **b** grows with age, fatigue, stress, and cognitive impairment. Combined with a working memory of about four chunks, it gives practical rules:

- a single screen of eight clear options may be faster than three nested menus — **if** the options are well labelled and the user does not need to remember anything between them;
- any task that requires holding more than about four new items in mind will fail for some users some of the time;
- a sequence the user must **recall** fails more often than one they can **recognise** on the screen.

12.6.5. The physics of a fall

Falling converts potential energy into kinetic energy, and the body must absorb it on impact:

$E = m \times g \times h$

For a 70 kg person whose hip drops about 0.9 m in a sideways fall from standing:

$E \approx 70 \times 9.81 \times 0.9 \approx 620$ J

How much **force** that produces depends on how far and how quickly the impact is stopped — the same energy absorbed over a shorter distance gives a larger force. On a hard floor the stopping distance is a few centimetres of soft tissue over the hip, and the peak force can reach several kilonewtons — enough to exceed the strength of an osteoporotic femur. The engineering responses all work on the variables in that statement:

- **reduce h** — low beds that sit close to the floor; a fall from a 20 cm bed carries well under a quarter of the energy of a fall from a 90 cm one;
- **increase stopping distance** — hip protectors with energy-absorbing or force-spreading pads, impact-absorbing flooring;
- **prevent the fall** — rails, lighting, footwear, strength and balance exercise, and removing trip hazards;
- **shorten the time on the floor** — fall detection and alerting, because lying on the floor for hours after a fall causes its own harm (Chapter 5).

12.7. Devices, software, and training technology

12.7.1. Clinical devices

Device	Does	Engineering note
Incubators and radiant warmers	Keep newborns warm and humidified	Closed-loop temperature control from a skin probe; humidity; access without heat loss. Life-supporting — Tier 3 (Appendix I)
Neonatal phototherapy	Blue light treats jaundice by converting bilirubin in the skin	Wavelength, irradiance, and eye protection for the baby
Paediatric sizes of everything	Masks, cuffs, airways, cannulas, sensors	Size selection under pressure is a task; colour coding and length-based systems help
Length-based weight-estimation tapes	Estimate weight and equipment sizes from a child's length in an emergency	A measurement system turned into a lookup table — simple, robust, and a lesson in designing for stress
Syringe pumps with paediatric drug libraries	Deliver small, precise volumes with dose limits	Weight-based limits; tenfold-error prevention (Chapter 17)

Device	Does	Engineering note
Hip protectors	Reduce the force on the hip in a fall	Energy absorption or force shunting; only effective if worn — comfort is a safety feature
Low beds, bed and chair sensors	Reduce fall height; alert staff when a person at risk gets up	Pressure mats and sensors trade missed falls against alarm fatigue
Personal alarms and fall detectors	Summon help after a fall	Accelerometer-based detection; false alarms versus missed falls is an ROC trade (Chapter 15)
Walking aids, grab rails, raised seats	Support mobility and transfers	Load-bearing — prescribed by therapists (Appendix I)
Medication dispensers and organisers	Support taking the right medicine at the right time	Opening force versus child resistance; reminders; carer visibility
Anti-ligature fittings	Remove ligature points in mental-health settings	Collapsing rails, sloped surfaces, load-release hooks, tamper-resistant fixings
Brain stimulation therapies — ECT, transcranial magnetic stimulation	Treat severe depression and some other conditions	Specialist, regulated, and delivered under strict protocols
Light therapy devices	Bright-light exposure for seasonal depression and circadian problems	Illuminance at the eye; glare; photobiological safety (Chapter 8)
Sensory rooms and sensory equipment	Calming or stimulating environments for autistic people and people with learning disabilities or dementia	User control over light and sound; durability; cleanability (Chapter 13)

12.7.2. Software

Software	Does	Note
Growth-chart software — e.g. the WHO Anthro tools	Plot and calculate centiles and z-scores	**Free.** Reference population and units must be explicit
Paediatric dose calculators	Weight- or BSA-based dosing	Medication-handling software is high risk (Appendix I); tenfold errors are the canonical failure
Early warning scores	Combine vital signs into a score that triggers review	Adult and paediatric versions use **different** thresholds — the age table above in code

Software	Does	Note
Cognitive screening apps	Brief tests of memory and attention	Screening only; results depend heavily on language, education, hearing, and vision
Digital mental-health apps	Self-help, therapy programmes, mood tracking	Evidence quality varies enormously; crisis handling is a safety requirement, not a feature
Easy-read and plain-language tools	Readability checking; symbol libraries; simplified layouts	**Directly usable by you.** Test with the intended readers, not with a formula alone
Fall-detection algorithms on wearables	Detect falls from motion data	Public datasets exist; sensitivity and false-alarm rates must be reported honestly
Reminder and prompting systems	Support memory and routine for people with dementia or brain injury	Simplicity and reliability matter more than features

12.7.3. Training and simulation

Training technology	Teaches	Form
Neonatal, infant, and child manikins	Resuscitation, airway management, cannulation, recognition of the sick child	Sized to age; some with programmable vital signs. **Paediatric replacement parts are scarce and expensive**
Preterm and birthing simulators	Delivery, neonatal resuscitation, emergencies in childbirth	From simple pelvis models to high-fidelity birthing manikins
Paediatric task trainers	Infant lumbar puncture, intraosseous access, heel-prick sampling	Small, precise, and makeable with the right materials (Chapter 21)
Ageing-simulation suits	Reduced vision, hearing, joint range, and grip	**Use with care** — as Chapter 7 warns, useful for finding barriers, not for "understanding" ageing
Dementia awareness simulations	What disorientation and sensory overload may feel like	Widely used; contested for the same reason. Pair with people who have lived experience
Simulated patients (actors)	Communication, consent, de-escalation, mental-health assessment	The highest-fidelity "device" for the mind is a trained person
De-escalation and restraint-reduction training	Recognising and calming distress; safe holding as a last resort	Scenario-based; heavily governed

Training technology	Teaches	Form
Ligature-cutter training	Rapid, safe removal of a ligature in an emergency	Specialised cutters and practice rigs — a training object with a precise purpose



Three populations, three firm boundaries.

- **Children**: anything a child could reach must be assessed for small parts, cords, entrapment, button cells, and magnets — and any contact with children goes through safeguarding. Nursery and toy standards exist for good reasons and are not optional because a hospital is the setting.
- **Older and frail adults**: anything that bears weight or helps someone stand or walk is a prescribed device (Appendix I). A rail or seat that fails under an unexpected load causes exactly the fall it was meant to prevent.
- **Mental-health settings**: nothing enters a ward without the ward's ligature risk assessment. The most innocent cord is the most dangerous object.



The best entry points here are informational and environmental. Easy-read versions of patient information, better labelling, clocks and signage for wards, accessible controls, and replacement parts for paediatric manikins are all in demand, need no patient contact, and draw on the whole toolkit — writing and graphic design, signage and sheet materials, fabrics, electronics, and software — provided they are made **with** the people who will use them.

12.8. Worked example: the elbow model goes to two new places

Your elbow model has done well. Two requests arrive in the same week:

1. A **play specialist on a children's ward** would like to use it to explain to children what will happen before an operation on the arm.
2. An **occupational therapist running a falls-prevention group** for older adults would like to use it to show how muscles and tendons move a joint, as motivation for strength exercises.

Same object, two populations. Work through each.

The children's ward.

- **Small parts.** The removable cord tabs and the pivot pin: do they fit the small-parts cylinder of Chapter 8? If so, they must be captive or redesigned. Any coin cell from a later version is excluded outright.
- **Cords.** The cord is a strangulation risk. Shorten it, make the loop impossible to put over a head, or use a breakaway attachment, and never leave the model unattended with a child.
- **Pinch and entrapment.** The hinge closes like a pair of scissors. Gaps that a small finger can

enter but not easily escape — especially at the moving joint — must be removed or guarded.

- **Cleanability.** Children's wards are infection-control priorities (Chapter 13); card is out, smooth wipeable surfaces are in.
- **Language and labels.** "PULL TO BEND" is fine for adults; for a young child, a picture of the arm bending and a coloured handle may say more. The play specialist decides; you supply options.

The falls-prevention group.

- **Labels.** The group sits round a table at 1–2 m; work out the letter height with Chapter 8's rule for a 6/12 user with reserve — roughly 6–9 mm at 1 m, more at 2 m — and make the contrast high.
- **Grip and force.** Members may have arthritic hands. The cord tab must be large enough to hold with the whole hand, and the force to bend the model low.
- **Cognitive load.** The instructions go from four steps to one action with a visible result: "Pull the red handle — watch the forearm rise."
- **Respect.** The model must look like a piece of anatomical equipment, not a toy. The same bright colours that welcomed the children may put the adults off.

What carries across. Both versions came from the same question — **who exactly will handle this, where, and with whom present?** — asked before changing a single part. Record both populations in the requirements, with their own acceptance tests (Chapter 19), and mark the play specialist's and therapist's approval as **pending** until you have it.

And one more question you should ask of both: **might this model ever be taken onto a mental-health ward?** If the answer is yes or unknown, the cord needs the ward's risk assessment before it goes anywhere near one.

12.9. Workshop exercise

1. **Surface area.** Calculate your own BSA with the Mosteller formula, then a newborn's (50 cm, 3.5 kg) and a ten-year-old's (about 140 cm, 32 kg). Express each per kilogram and plot the three points. Explain in one paragraph what the curve means for heat loss.
2. **The fourth power.** Plot airway resistance against radius for radii from 1 mm to 10 mm, normalised to the value at 10 mm. Mark where an infant and an adult airway sit, and the effect of 1 mm of swelling at each.
3. **Fall energy.** Measure the height of a bed, a chair seat, and your own hip when standing. Calculate the potential energy for a 70 kg person falling from each. Then look around one room in your home and list every trip hazard and every object at the wrong height.
4. **Cognitive-load audit.** Take any multi-step task on a device or app you use. Count how many things the user must hold in mind between steps. Redesign it so that the number is never more than two, and state what each change costs.
5. **Easy-read rewrite.** Take a paragraph of instructions — from this book, from a device leaflet, or from a hospital website — and rewrite it in easy-read style: one idea per sentence, short common words, a picture for each step. Then ask someone unfamiliar with it to follow it, and record where they hesitated.

6. **Ligature walk.** In your own home, identify every fixture that could bear a load and that something could be looped around. You do not need to change anything. The point is to see how ordinary such points are — and why anti-ligature design is a specialised discipline.

12.10. Questions

1. Give three ways a newborn differs from an adult that matter to a device touching the skin.
2. Why can one millimetre of swelling be trivial in an adult airway and life-threatening in an infant's? Show the calculation.
3. A heart rate of 150 beats per minute is displayed. What else must you know before deciding whether it is unusual?
4. What does "weight-based dosing" mean for software, and what is the characteristic error?
5. Name four changes of ageing and one design consequence of each.
6. What is frailty, and which two of its common defining features are measurements?
7. State three contributors to falls that are environmental or design factors.
8. Distinguish delirium from dementia in three ways, and explain why delirium is treated as urgent.
9. What is the approximate capacity of working memory, and give two interface features that spend it unnecessarily.
10. What is a ligature point, and why must an object for a mental-health ward be assessed before it is introduced?
11. Use the Mosteller formula to compare BSA per kilogram in a newborn and an adult. What follows?
12. Why does designing for "the 5th to 95th percentile" on several dimensions exclude more than 10% of users, and who is disproportionately excluded?

12.11. Answers and reasoning

1. Any three of: **thinner, more permeable skin** (more absorption of chemicals, more injury from adhesives); **larger surface area per kilogram** (faster heat and water loss); **immature thermoregulation** (cannot shiver effectively); **small blood volume** (sampling matters); **fragile skin layers** that separate when tape is removed.
2. Resistance scales with the inverse fourth power of radius. An infant airway of radius 2 mm narrowed to 1 mm has $(2/1)^4 = 16$ times the resistance. An adult airway of radius 8 mm narrowed to 7 mm has $(8/7)^4 \approx 1.7$ times the resistance.
3. The person's **age** above all — 150 is normal in a newborn and abnormal in a resting adult — plus context: activity, fever, distress, crying, medication, and how the measurement was taken. An alarm limit is a function of age.
4. Doses and volumes are calculated per kilogram (or per m^2 of BSA), so the software must capture weight reliably, with its unit, and compute and display the dose with its basis visible. The characteristic error is a **tenfold** error from a misplaced decimal point or a unit confusion (g versus kg, mg versus micrograms).

5. Any four, for example: **thinner skin** — skin tears from adhesives and edges; **sarcopenia** — lower operating forces needed; **osteoporosis** — falls cause fractures, so prevent falls and reduce fall height; **presbyopia and cataract** — larger, higher-contrast text, low glare; **high-frequency hearing loss** — lower-pitched and visual alarms; **reduced kidney clearance** — drugs accumulate; **slower gait** — no short time limits.
6. A state of reduced reserve across several systems, making a person vulnerable to disproportionate decline after a minor stress. Common definitions include **grip strength** and **walking speed**, both measured.
7. Any three of: poor lighting; loose cables and rugs; slippery floors; beds, chairs, or toilets at the wrong height; lack of grab rails; unfamiliar layouts; devices that create trip hazards.
8. **Onset**: delirium sudden, dementia gradual. **Course**: delirium fluctuates and is often reversible; dementia progresses. **Attention**: impaired in delirium, often preserved early in dementia. Delirium signals an acute underlying problem — infection, medication, dehydration — and is associated with worse outcomes if missed, so it is urgent.
9. Around four meaningful chunks. Features that spend it: requiring a code or value to be remembered from one screen to another; long sequences the user must recall rather than recognise; mental unit conversions; unclear system state, so the user must remember what they have already done.
10. A fixture or object around which a cord or fabric could be looped and which could bear a person's weight. Hanging is a leading method of suicide in inpatient mental-health settings, so every object and fixture — including cords and straps, which are ligatures themselves — must be assessed against that risk by the ward team.
11. Newborn: $\sqrt{(50 \times 3.5 / 3600)} \approx 0.22$ m², about 0.063 m²/kg. Adult: $\sqrt{(170 \times 70 / 3600)} \approx 1.82$ m², about 0.026 m²/kg — roughly 2.4 times less per kilogram. The newborn loses heat and water and absorbs skin-applied substances proportionally faster.
12. Each dimension excludes 10% of people, and the excluded people differ between dimensions, so the combined exclusion grows with the number of dimensions. Children, older adults, and disabled people are disproportionately at the extremes and so are excluded first.

12.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Unit-safe clinical arithmetic — body surface area, and per-kg against per-m².
- Levers, moment arms and contact — fall energy and stopping distance.
- Test lung and breathing mechanics — airway radius in infants and adults.
- Labels, alarms and the senses — what changes with age in seeing and hearing.
- Test accuracy and risk — percentiles and growth-chart exclusion.

12.13. Video companion and reading

- **Core videos:** the Crash Course Anatomy & Physiology episodes on **the brain** and **memory**, from

the series index [6], for the cognition section.

- **Core reading:** OpenStax [1] selectively:
  - 28.5 Maternal Changes During Pregnancy, Labor, and Birth
  - 28.6 Adjustments of the Infant at Birth and Postnatal Stages
  - 6.6 Exercise, Nutrition, Hormones, and Bone Tissue — for osteoporosis
  - the "Aging and the …" feature boxes found in many chapters, which summarise how each system changes with age
- **Growth and percentiles:** the WHO Child Growth Standards [47]. Look at one chart and find the median, the 3rd, and the 97th centiles.
- **Ageing:** the WHO fact sheet on ageing and health [48].
- **Frailty:** Fried et al. [49] — read the abstract and the five criteria.
- **Working memory:** Cowan [50] — the abstract is enough.
- **Mental-health environments:** the UK Health Building Note on adult mental-health units [51] — skim the sections on ligature risk to see how a design requirement is written for this setting.

12.14. Exit criteria

- ▢ Explain, with numbers, three ways a newborn's physiology differs from an adult's and one design consequence of each.
- ▢ Reproduce the airway fourth-power calculation and the Mosteller BSA comparison without notes.
- ▢ List the main age-related changes in at least six systems, each with a design consequence.
- ▢ Distinguish dementia from delirium, and name three environmental measures that help people with either.
- ▢ Complete a cognitive-load audit of a real task and reduce the number of things held in mind to two or fewer.
- ▢ Produce an easy-read version of one set of instructions and test it on one person, recording hesitations.
- ▢ For one object you have made, write down whether it could reach a child, a frail older adult, or a mental-health ward, and what would have to change — or be assessed — first.

12.15. What this chapter does not establish

- **You cannot assess a child's development, an older person's frailty, or anyone's cognition or mental health.** The concepts here let you follow the conversation and design more inclusively. Assessment is clinical work by trained professionals.
- **You cannot calculate doses.** Weight and surface-area scaling appear here to show why software must be careful, not to teach dosing.
- **You cannot make anything ligature-safe by your own judgement.** Anti-ligature design is a specialist field with its own guidance, products, and testing, and every ward makes its own risk

assessment.

- **Simulation suits and awareness exercises do not give you lived experience.** As in Chapter 7, they help find barriers; they do not tell you what it is like to be old, disabled, or unwell, and they do not validate a design.
- **Working with children and with people who may lack capacity has legal and safeguarding requirements** that differ by jurisdiction. Your host establishes them; you follow them without exception.

Part II: Putting an object into a clinical world

Chapter 13. Infection control and choosing materials

Week 16 · approximately 4–6 hours

13.1. Why this chapter exists

A part coming off the machine — out of the printer, off the lathe, out of the mould, off the sewing machine, out of the annealing oven — is the **beginning** of the object's life, not the end of its design.

The object will be handled, carried, set down, used, set down again, cleaned by somebody who is not you with a process you did not choose, stored in a cupboard, used by the next person, dropped, repaired, and eventually thrown away. Every one of those stages is a design input, and most of them are decided at the design stage — in CAD, on the pattern, at the bench — by whether there is a cavity you cannot reach, a seam that traps fluid, a fastener that loosens, a surface the approved wipe cannot wet for long enough.

This chapter is also where a maker most often makes a claim they cannot support. "It's PLA, that's food safe", "it's surgical stainless", "it's medical-grade silicone", and "this material is used in medical devices" are all sentences that sound like evidence and are not.

The chapter's second half is about **materials and finished objects**, whatever made them. Choosing the technology and the fabrication method in the first place — software, electronics, printing, machining, moulding, metalwork, glass, textiles — is the subject of Chapter 14.

13.2. Microbiology, in the amount you need

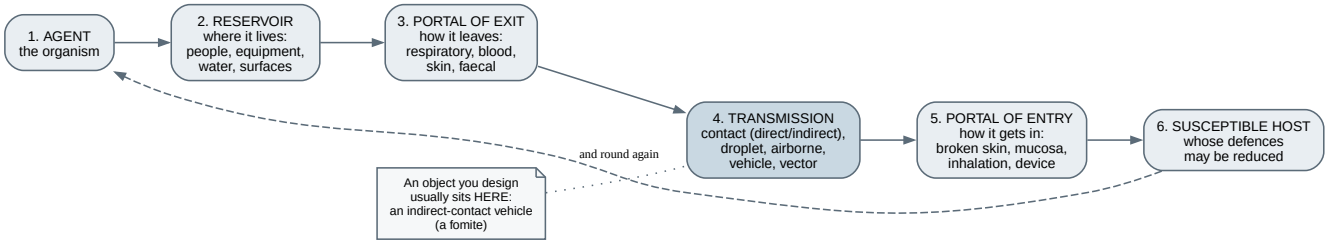
Agent	What it is	Why it matters to an object
Bacteria	Single-celled organisms; many harmless, some pathogenic	Multiply on surfaces given moisture and nutrients; some form **biofilm**
Bacterial spores	Dormant, highly resistant forms produced by some bacteria (e.g. *Clostridioides difficile*)	Survive many disinfectants. Sterilisation is defined partly by its ability to kill them
Viruses	Genetic material in a protein coat; replicate only inside host cells	**Enveloped** viruses (influenza, SARS-CoV-2) are generally easier to inactivate; **non-enveloped** (norovirus) are tougher
Fungi	Yeasts and moulds	Grow in damp conditions; relevant to porous and absorbent materials
Prions	Misfolded proteins	Extraordinarily resistant; standard sterilisation is not assumed sufficient. A specialist domain

Biofilm deserves its own paragraph. Bacteria attached to a surface can secrete a protective matrix,

forming a community that is dramatically more resistant to disinfectants and to mechanical removal than the same organisms free-floating. Biofilm forms preferentially in exactly the places a maker creates without thinking: crevices, threads, lumens, scratches, porous surfaces, and the gaps between press-fitted parts.

This is the microbiological reason behind what otherwise sounds like fussiness about surface finish. A scratched or porous surface is not merely harder to clean; it provides protected habitat.

13.2.1. The chain of infection



Infection control works by breaking the chain at the cheapest available link. Hand hygiene is the single most effective measure, and it breaks link 4. Your object sits at link 4 as well: a **fomite**, an inanimate vehicle for indirect contact transmission.

> **An object can participate in transmission even if it never touches a patient.**
>
> The shared teaching model in your running project is handled by eight learners, set down on a table, put in a box, carried between rooms, and handled again next week. Nobody's skin was broken, no patient touched it. It is still a surface in a healthcare environment, and the host's infection prevention team has a view about surfaces in healthcare environments.
>
> So the question is never "does it touch a patient?" It is "**where does it go, who touches it, and what else do they touch next?**"

13.2.2. Routes of transmission

Route	Mechanism	Design relevance
Direct contact	Person to person	Not usually your object
Indirect contact	Via a contaminated intermediate object — a **fomite**	**This is your object.** Hands, surfaces, shared equipment
Droplet	Larger respiratory particles over short distances	Relevant to anything near a face; shields, masks
Airborne	Smaller particles that remain suspended	Specialist; ventilation engineering
Common vehicle	Shared food, water, medication, solution	Relevant to anything holding fluid
Vector-borne	Via an organism such as an insect	Rarely relevant in this context

13.3. Cleaning, disinfection, and sterilisation are three different things

These words are not interchangeable, and treating them as synonyms is the most common error in this area.

Process	Definition	Typical methods
Cleaning	Physical **removal** of soil, organic material, and contaminants. It does not necessarily kill anything, and it is the **prerequisite** for everything else — residual soil physically shields organisms from a disinfectant or sterilant	Detergent and water, mechanical action, ultrasonic bath, automated washer-disinfector
Disinfection	A process that **inactivates many disease-causing microorganisms**. Effectiveness depends on the agent, concentration, contact time, temperature, and the organism. It is explicitly **not** equivalent to sterilisation — notably, many disinfectants do not reliably kill bacterial spores	Alcohol, chlorine-releasing agents, quaternary ammonium compounds, hydrogen peroxide, thermal disinfection
Sterilisation	A process **intended to eliminate all forms of microbial life**, including spores. Defined probabilistically — a sterility assurance level, conventionally 10^{-6} — and validated, monitored, and documented	Steam (autoclave), ethylene oxide, vaporised hydrogen peroxide, gamma or electron-beam irradiation



Cleaning must come first, always. Residual blood, protein, or other soil physically protects organisms and can inactivate the chemical agent. A sterilisation cycle run on a dirty instrument is not a sterilisation cycle. This is why automated reprocessing workflows are sequential and why "I gave it a wipe" does not substitute for any step in them.

13.3.1. Spaulding classification

The standard framework for deciding **what level** of processing an item needs, based on **how it is used**:

Category	Contact	Required processing	Examples
Critical	Enters sterile tissue or the vascular system	**Sterilisation**	Surgical instruments, implants, needles
Semi-critical	Contacts mucous membranes or non-intact skin	At minimum **high-level disinfection**	Endoscopes, respiratory therapy equipment, some probes

Category	Contact	Required processing	Examples
Non-critical	Contacts intact skin only	**Low- or intermediate-level disinfection**	Blood pressure cuffs, stethoscopes, bed rails, most equipment surfaces



The category is set by the use, not by the object. The same piece of plastic is non-critical on a desk and critical inside a body. This is why "what is this made of?" is the wrong first question and "where does it go and what does it touch?" is the right one.

Your tabletop teaching model is **non-critical at most**, and arguably outside the scheme entirely as general equipment. That is a good place to be for a first project, and you should state it explicitly in your brief rather than leaving it implied.

13.3.2. Steam sterilisation, since it is the common one

Autoclaving is the default where material permits. Typical cycle conditions in healthcare practice:

Type	Temperature	Hold time	Note
Gravity displacement	~121 °C	~15–30 min	At ~103 kPa (15 psi) above atmospheric
Pre-vacuum	~132–135 °C	~3–4 min	Vacuum removes air so steam reaches surfaces

Figures are illustrative of commonly cited cycles; actual validated parameters are set by the institution and the equipment, per standards such as ISO 17665 and AAMI ST79 (Appendix D).



Autoclaving will destroy much of what a general workshop produces.

PLA softens well below 100 °C — its glass transition is around 55–65 °C. It will deform in a dishwasher, let alone an autoclave. PETG is better but still marginal. Acrylic, many hot-mel and cyanoacrylate bonds, most elastic and hook-and-loop fabrics, and almost all electronics are also casualties. The things that routinely survive — 316 stainless steel, borosilicate glass, many silicone grades, and engineering polymers such as PEEK, PPSU, and some polypropylenes — are survivors because they were **chosen** for it. Even materials that survive thermally may change dimension, become brittle, discolour, or degrade over repeated cycles.

If a project ever needs steam sterilisation, that constraint drives material and process selection **from the beginning**. It is not something to discover at the end.

13.4. The finished object is the design



A material claim is not device evidence.

"Food safe", "medical grade", "biocompatible", and "used in medical devices" are

claims about **a material as supplied**, usually under specified conditions. None of them is a claim about **your finished object** in **your intended use**.

Between the raw material and your part sit the process and everything it touches. Every method has its own list:

Process	What sits between the datasheet and the part
Additive (printing)	Thermal history; additives and colourants; the hot end, nozzle, and residue of previous materials; build-plate adhesion aids; layer voids; support marks and residue; for resins, incomplete cure
Machining (CNC, lathe, mill)	Cutting fluids and coolants; tool marks and burrs; embedded swarf; machine-shop contamination from other materials; heat-affected surfaces on plastics
Injection moulding	Mould-release agents; colourant masterbatch; regrind content; flash and parting lines; sink marks and internal voids; moulded-in stress that later cracks under disinfectants
Casting (silicone, polyurethane, resin)	Mould release; cure inhibition from contaminated moulds or gloves; incomplete cure and post-cure; trapped bubbles; pigments and fillers
Sheet metal and welding	Heat tint and oxide on welds; flux residues; free iron from tools, needing passivation of stainless; crevices at lap joints and spot welds; galvanic couples between dissimilar metals
Glasswork	Residual stress if annealing is inadequate; chips and sharp edges; seals to other materials; surface treatments
Textiles and soft goods	Finishes, dyes, and flame retardants; shrinkage and change after laundering; stitch holes through coatings; elastic and hook-and-loop that degrade and trap soil
Electronics	Flux residues; conformal coatings; battery chemistry; adhesives and potting compounds; ingress paths around buttons, connectors, and seams

Plus, for all of them: post-processing chemicals and abrasives, adhesives, fasteners, coatings, and handling.

A part made by any process is not the material on the datasheet. The FDA's framing, written for additively manufactured devices but true of every process, is instructive: materials are evaluated **in the context of a finished device and its intended use**, not approved in the abstract [25].

13.4.1. What actually varies: additive manufacturing

Every process has variables that change the finished part from run to run. Additive manufacturing is shown in detail because its variables are the least visible from the outside; the next table gives the equivalent questions for other methods.

Factor	Why it matters
Build orientation	Determines where layer boundaries fall relative to load. Printed parts are markedly weaker between layers. Also determines where supports contact and scar the surface
Layer height and nozzle size	Sets surface roughness, which drives cleanability and biofilm habitat
Infill and wall count	An infilled part contains enclosed voids. If the shell is ever breached, those voids are inaccessible and uncleanable
Temperature and flow	Affects layer adhesion and porosity. The same file on the same printer a month later is not guaranteed identical
Support removal	Leaves marks, residue, and sometimes embedded fragments
Post-processing	Sanding introduces particles and changes roughness; vapour smoothing introduces solvent; painting and coating add a new material with its own questions; acetone, isopropanol, and resin washes leave residues
Resin printing specifically	**Uncured resin is an irritant and a sensitiser.** Adequate washing and post-curing are essential and are process steps that must be specified and verified, not assumed
Machine history	A shared hobby printer has an unknown contamination history. Colourants, previous materials, and lubricants are all present
Traceability	Which file, which slicer profile, which machine, which spool, which date? Without this you cannot reproduce a part, and you cannot connect a test result to the thing tested (Chapter 19)

13.4.2. What actually varies: other methods

Method	What varies	Why it matters for cleaning and contamination
CNC machining	Tool wear, feeds and speeds, coolant, deburring, fixturing	Surface roughness and burrs set cleanability; coolant and swarf residues must be removed; internal threads and blind holes are as hard to clean as any printed cavity
Injection moulding	Resin lot, regrind, moulding temperature and pressure, release agent, tool wear	Usually excellent, repeatable surfaces — once the tool is right. Release-agent residue and moulded-in stress are the hidden variables; a change of resin lot or moulder is a change of product
Silicone and polyurethane casting	Mix ratio, degassing, cure temperature and time, mould condition	Under-cured material is tacky and leaches; bubbles at the surface become pits that hold soil; cure inhibition can be invisible

Method	What varies	Why it matters for cleaning and contamination
Sheet metal and welding	Weld quality, heat input, finishing, passivation	Stitch welds and lap joints make crevices; continuous welds ground smooth do not. Unpassivated stainless can rust where tools left free iron
Glasswork	Annealing schedule, wall thickness, joint design	Inadequately annealed glass can fail later under thermal or mechanical shock — during cleaning, for instance — and broken glass is a sharps injury
Textiles and soft goods	Fabric lot and finish, seam construction, thread, laundering history	Fabric is cleaned by **laundering**, not wiping; coatings crack and seams wick; every wash changes dimension and strength
Electronics assembly	Solder process, flux and cleaning, enclosure sealing, firmware version	Enclosures must resist the approved wipe and fluid ingress; the firmware version is part of the configuration (Chapter 20)

Each of these methods has its own traceability record: drawing revision and material certificate for a machined part, resin lot and moulder for a moulding, mix and cure log for a casting, fabric batch and pattern revision for a sewn item, board revision and firmware build for electronics.

13.4.3. Design for cleanability

Rules of thumb that cost nothing at the CAD stage and are expensive later:

Principle	In practice
Make every surface reachable	If a wipe on a gloved finger cannot reach it, it is not being cleaned. Avoid deep narrow slots, blind holes, internal lattices, and re-entrant features
Eliminate crevices, or make them openable	A press-fit seam is a capillary gap. Either seal it properly or make it a deliberate, tool-free, repeatable disassembly
Radius internal corners	Sharp internal corners trap soil and resist wiping. Generous radii clean better
Minimise fasteners, and avoid exposed threads	Threads are hard to clean and loosen with handling. Prefer captive or smooth features
Smooth, non-porous, non-absorbent surfaces	Which rules out unsealed wood, open-cell foam, uncoated fabric, as-cast or bead-blasted textures, and unsmoothed prints for anything that must be **wiped** clean to a standard
Make metal joints continuous	Continuous welds ground flush, not stitch welds or lap joints; no blind tubes; passivate stainless after fabrication

Principle	In practice
Choose the cleaning route for soft goods	Fabric items are laundered, wiped (if coated and sealed), or single-patient use. Decide which at the start; hook-and-loop, exposed foam, and stitched-through coatings defeat all three
Protect the electronics	Sealed enclosures with gaskets or overmoulding, no gaps around buttons, membrane switches rather than open mechanical ones, connectors that can be capped
No fluid traps	Orient so liquid drains. Ask where water sits when the object is set down the way people actually set it down
Design for inspection	Someone must be able to **see** whether it is clean and whether it is damaged. A closed shell cannot be inspected
Label durably	Version and warning labels must survive the actual cleaning process. A paper label is not a cleanable surface, and a label that washes off takes your traceability with it
Plan for disposal	What happens at end of life? Mixed materials and embedded electronics complicate it



Wood and fabric. Both are legitimate workshop materials, and they behave very differently in a clinical setting.

Wood is porous and absorbent and is usually inappropriate for anything shared that must be cleaned. "Sealed wood" raises further questions: which sealant, is it intact, what happens when it is scratched, is it compatible with the disinfectant, how is damage detected? Use it where cleaning requirements permit — a workshop demonstration, a bench fixture, a model that stays in your own space — and be explicit in your notes that this is why.

Fabric is everywhere in healthcare — slings, straps, garments, compression sleeves, positioning cushions, curtains — because it has a recognised cleaning route of its own: **laundering** with thermal disinfection (UK guidance, for example, specifies holding at 65 °C for at least 10 minutes or 71 °C for at least 3 minutes), or a wipeable coated fabric with sealed seams, or single-patient use. A fabric design is cleanable if it survives its route: the fabric, thread, elastic, fastenings, and labels all go through the wash or the wipe, and the item still fits and functions afterwards. Agree the route with the host first.

13.5. Engineering bridge: comparing candidate materials

Here is a comparison of the sort you should produce — and note how much of it is **unknowns**, honestly marked.

Material	Typical process	Thermal limit	Cleanability	Autoclavable?	Unresolved
PLA	FDM printing	Softens ~55–65 °C (T_g)	Poor as printed: layer lines, porosity	**No**	Additives and colourants unknown; machine contamination history; layer porosity
PETG	FDM printing	T_g ~80 °C	Better than PLA, still layered	Generally no	As above; chemical compatibility with specific disinfectants
ABS / ASA	FDM printing	T_g ~105 °C	Can be vapour-smoothed, which helps	Marginal, and dimensional change likely	Smoothing solvent residue; styrene emissions during printing
Polycarbonate	FDM, machining	High	Good when smooth	Sometimes, material-dependent	Susceptible to stress cracking with some chemicals — compatibility must be checked
Polypropylene	Injection moulding, some FDM	~100 °C+	Good; widely used in medical disposables	Often yes, grade-dependent	Grade and additives; mould release; difficult to print, with adhesion and warping
Acrylic (PMMA)	Laser cutting, machining, casting	T_g ~105 °C	Good when smooth and edges polished	No	**Crazes and cracks with alcohol-based wipes**, especially at laser-cut edges and drilled holes — check before choosing it for anything wiped
Acetal (POM)	Machining; injection moulding	Continuous use ~90–100 °C	Good; low friction, dimensionally stable	Limited — repeated steam cycles degrade it, homopolymer worst	Grade; compatibility with chlorine-based agents
UHMWPE	Machining	Softens ~80 °C+	Good; very low friction, tough	Generally no	Creep under load; bonding is difficult
PEEK	Machining; injection moulding	T_g ~143 °C; continuous use ~250 °C	Excellent	Yes, repeatedly	Cost; machining coolant residues; grade (implant grades are a separate supply chain)

Material	Typical process	Thermal limit	Cleanability	Autoclavable?	Unresolved
Photopolymer resin	SLA / DLP / MSLA	Varies widely	Can be very smooth	Usually no	**Cure state.** Uncured resin is an irritant and sensitiser. Wash and cure process must be specified and verified
Stainless steel 316L	Machining, sheet metal and welding, SLM printing	Very high	Excellent when polished and joints are continuous	Yes	Weld heat tint and passivation; printed porosity and residual powder; surface finish; crevices at joints
Aluminium (anodised)	Machining, sheet metal	High	Good	Anodising may be attacked by some agents	Chemical compatibility of the anodised layer
Silicone elastomer	Casting, moulding	High for many grades	Good; non-porous	Many grades yes	Grade matters enormously; cure system; post-cure; tears trap soil
Borosilicate glass	Scientific glassblowing, lampworking	Very high; resists thermal shock better than soda-lime glass	Excellent	Yes	Annealing; breakage and sharps risk; joints to metal or polymer; chips harbour soil
Coated fabric (PU- or PVC-coated polyester)	Cutting and sewing, welded seams	Low to moderate	Wipeable while the coating and seams are intact	No	Coating cracking and delamination over time; disinfectant compatibility; stitch holes unless seams are welded or taped
Polyester webbing and woven textiles	Cutting and sewing	Launderable at thermal-disinfection temperatures, grade-dependent	Laundered, not wiped	Generally no	Shrinkage and strength after repeated washes; hook-and-loop and buckles trap soil; who launders it, and where
Hardwood, sealed	Machining	Low	**Poor** — porous, absorbent	No	Sealant identity and integrity; damage detection. Generally unsuitable for shared clinical use



This table is a worked example of how to structure a comparison, not a materials specification. Every "unresolved" column entry is the honest part. A

comparison whose unknowns column is empty is a comparison that has not been done properly.

Do not give any material a blanket "hospital safe" rating. The host's infection prevention team specifies the process; the process constrains the material; the intended use determines the process. That is the direction of the logic, and it runs opposite to the way a maker naturally wants to work.

13.6. Devices, software, and training technology

Infection prevention is a systems discipline more than a device one, but the equipment that supports it is substantial — and the training side is unusually accessible to a maker.

13.6.1. Clinical devices and equipment

Device	Does	Engineering note
Steam steriliser (autoclave)	Saturated steam under pressure; the default where materials permit	Air removal is the critical step — trapped air prevents steam contact. Pre-vacuum cycles exist for exactly that reason
Washer-disinfector	Automated cleaning plus thermal disinfection	Load design, spray coverage, and the A_0 concept — time at temperature
Low-temperature sterilisation — ethylene oxide, vaporised hydrogen peroxide	Sterilises heat-sensitive items	EtO is effective and leaves residues requiring aeration (Chapter 3). VH_2O_2 has material compatibility limits
Endoscope reprocessor	Cleans and high-level disinfects long narrow lumens	**One of the hardest reprocessing problems in medicine.** Channel patency, drying, and traceability. Associated with real outbreaks when it fails
Ultrasonic cleaner	Cavitation removes soil from complex geometry	Effective in crevices a brush cannot reach — which is also where biofilm lives
UV-C disinfection systems	Room and surface decontamination between occupancies	Line of sight only; shadowed surfaces are untreated. **Serious eye and skin hazard**
HEPA filtration and ventilation	Controls airborne transmission; the basis of isolation rooms and theatres	Air changes per hour, pressure differentials, directional flow

Device	Does	Engineering note
PPE — gloves, masks, respirators, gowns, eye protection	Barrier protection	Fit testing for respirators is not optional; a poorly fitting FFP3 is a comfort item. Donning and doffing **sequence** matters as much as the equipment
Sterile barrier systems — pouches, wraps, containers	Maintains sterility after processing until point of use	Covered by ISO 11607. The package is part of the device
Hand hygiene dispensers and monitoring	The single most effective intervention	Placement drives use. Electronic monitoring systems exist and are contested
Chemical and biological indicators	Verify that a sterilisation cycle met its conditions	Biological indicators use resistant spores — a direct test of the thing that defines sterilisation
Sharps and waste management systems	Safe disposal; segregation by waste stream	Sharps injury prevention is a well-studied design problem

13.6.2. Software

Software	Does	Note
Infection surveillance systems	Detect and track healthcare-associated infections and outbreaks	Data linkage and case definition are the hard parts
Instrument tracking and traceability	Links each tray through cleaning, sterilisation, and use to a patient	Barcode or RFID. **A recall problem: if a cycle fails, which patients were affected?**
Antimicrobial stewardship tools	Guide appropriate antibiotic prescribing	Addresses resistance — one of the major long-term threats in medicine
Hand hygiene compliance monitoring	Electronic or observational auditing	The Hawthorne effect is a genuine confounder in observational auditing
Outbreak modelling	Transmission modelling and scenario analysis	Epidemiological modelling; largely open-source toolchains
Sterilisation cycle monitoring and documentation	Records and validates each load	Part of the quality system; required for traceability

Software	Does	Note
Reprocessing instruction management (IFU)	Manufacturers' instructions for reprocessing each device	Enormous, fragmented, and a real practical burden on sterile services

13.6.3. Training and simulation

Training technology	Teaches	Form
Fluorescent marker systems — Glo Germ and equivalents	Makes contamination transfer **visible** under UV	Lotion or powder plus a UV lamp. **Cheap, dramatic, and genuinely effective.** The best single teaching tool in infection control
Hand hygiene training boxes	Shows missed areas after handwashing	A UV box; **a straightforward making project**
Aseptic technique trainers	Sterile field setup, non-touch technique	Practice kits with simulated consumables
Donning and doffing practice	PPE sequence — doffing is where contamination happens	Practice PPE plus observer checklists and fluorescent markers
Simulated contamination scenarios	Transmission routes through a whole workflow	Marker-based tracing across a simulated ward
Surgical scrub trainers	Technique and timing	Often combined with fluorescent marker verification
Cleanability demonstration coupons	Why surface finish and joints matter	Matched samples — printed, machined, polished, bead-blasted, cast, welded, sewn, coated fabric — soiled with fluorescent marker and wiped under the approved method. **A maker-accessible teaching set that shows this chapter's argument in one glance**
Catheter and line care trainers	Insertion and maintenance bundles that prevent infection	Task trainers plus checklist practice
Sterile services training rigs	Instrument inspection, set assembly, packaging	Real trays and instruments; inspection under magnification
Outbreak tabletop exercises	Decision making and communication under uncertainty	Facilitated scenarios; no equipment needed



The UV hand hygiene box is the best first build in this chapter.

> A light-tight enclosure, a UV-A source, a viewing aperture, and a hand opening. Pair it with commercial fluorescent lotion. It is a genuinely mixed build: folded and powder-coated sheet metal or a sealed laminate for the enclosure, UV-A LEDs with a driver, a door or hand-opening interlock and a timer in the electronics, and a wipeable gasket around the opening. It costs very little, it makes an invisible problem visible, and infection prevention teams use them constantly for teaching and for audit.
>
> Note what the design requires you to think about, all of it from this book: the enclosure must be **cleanable** (this chapter); the UV source must be **shielded from the eyes** — a hazard the user cannot perceive, which is Chapter 7' point exactly; and the learning objective — "identify which areas you missed" — must be stated before you build, which is Chapter 21.

13.7. Worked example: the pivot cover

Your elbow model's pivot could be covered two ways.

Option A — a closed one-piece shell. One part, no fasteners, nothing to lose, quick to make, nothing rattles.

Option B — two removable halves. Interior accessible, inspectable, cleanable, repairable.

Neither is correct in the abstract. The comparison looks like this:

Criterion	Closed shell	Removable halves
Part count	1	2 + fasteners
Assembly error risk	None	Real — can be reassembled wrongly or left loose
Interior access	None	Yes
Inspection	Cannot see inside	Can
Cleaning	Exterior only; interior is a sealed unknown	Interior reachable
Capillary gaps	At the shell-to-body join only	At the join, plus the fastener seats and the split line
Loss of parts	None	Fasteners can be lost
Repair	Replace whole	Replace the broken half
Cost of a wrong choice	Have to remake it	Someone reassembles it incorrectly

The question that resolves it is not in this table. It is: does the intended setting require interior access, what maintenance is expected, who performs it, and **is a cover needed at all?** Removing the cover entirely may be the best answer, and it only appears once you ask what the cover is for.

Then, and only then, the method. Suppose the educator wants a cover, wiped with an alcohol-

based wipe between sessions. The candidate ways of making it each carry their own cleaning profile:

Method	Strengths for this job	Cleaning questions
Printed shell (PETG)	Fast to iterate; one-off cost low	Layer lines and porosity; seal or smooth the surface; enclosed infill if breached
Machined acetal halves	Smooth, tough, repeatable; low friction at the pivot	Captive fasteners; deburred edges; coolant removed
Folded 316 stainless sheet	Durable, wipeable, survives anything	Bend radii and edges; sharp corners; weight on a light model
Cast silicone boot	No seams, no fasteners; flexes with the joint	Grade and cure; tears trap soil; stretches over time
Laser-cut acrylic	Transparent, so the mechanism stays visible	**Crazes with alcohol wipes** — likely to fail this cleaning route

None is "the answer" — that depends on quantity, the cleaning route, durability, and what the cover is for. The point is that the method is a design decision with cleaning consequences, made **after** the use is understood (Chapter 14).

Now a separate point about evidence. You run the model through fifty handling cycles and observe that a fastener works loose by cycle thirty-two. That is a real and useful bench result about **mechanical durability under your handling conditions**. It establishes nothing whatsoever about microbial cleanliness. Match the evidence to the claim.

13.8. Workshop exercise

1. **Contact map.** Draw every surface touched during one complete use and reset of your model. Include the hidden steps people forget: carrying it, opening the box, moving the labels, handling the cord, setting it on the table, putting it away. Mark each contact with **who** touches it and **what they touch next**.
2. **Material comparison.** Compare three candidates for a training fixture made by different methods — for example a printed or moulded polymer, a machined or sheet-metal construction, and a sewn or cast soft version. Record for each: intended setting, fabrication method, surface characteristics, cleaning questions, maintenance, and **explicitly, what evidence is missing**. Do not rank them overall.
3. **Design for cleanability review.** Take your current model and find three features that would be difficult to clean. Redesign each, and note what the redesign costs.
4. **Find the question owner.** Identify, by role, who at a potential host institution would have to review cleaning and any patient-contact requirements. Write down the question you would ask them. If you have a host, ask it.
5. **Traceability record.** Write down everything needed to reproduce one part of your model, using the record that fits its method — for a printed part: file and revision, slicer profile, machine, nozzle, material and spool, non-default settings, post-processing, date; for a machined

part: drawing revision, material and certificate, machine, tooling, coolant, finishing; for a cast part: material lot, mix ratio, cure schedule, mould revision; for a sewn item: fabric and thread batch, pattern revision, seam types; for electronics: board revision, components, firmware build. Then see whether someone else could reproduce it from your record alone.

6. **Cleanability coupons.** Make small matched samples by at least four methods available to you — for instance printed, machined, cast silicone, and coated fabric. Soil them with fluorescent lotion, wipe each once with the same technique, and photograph them under UV. Record what the surface did, not what you expected.

13.9. Questions

1. Explain cleaning, disinfection, and sterilisation, and why the order matters.
2. Why can an object participate in transmission without ever touching a patient?
3. What determines a device's Spaulding category?
4. Why are bacterial spores significant in defining sterilisation?
5. What is biofilm and why does surface finish matter to it?
6. "The filament is medical grade." What does this establish about your printed part?
7. Why is an infilled printed part a cleaning concern if its shell is breached?
8. Give three ways the same STL produces a different physical object.
9. You run fifty handling cycles with no failures. What have you established?
10. Name five features in a typical workshop assembly — printed, machined, welded, or sewn — that are hard to clean.
11. Why is PLA unsuitable for anything that must be autoclaved?
12. "It's injection moulded in medical-grade polypropylene." What could still be wrong with the finished part?
13. A strap for a positioning aid is made of polyester webbing with a hook-and-loop closure. How is it cleaned, and what will defeat that?
14. Why might a laser-cut acrylic cover fail in service when it is wiped with an alcohol-based wipe?
15. Your model is used only on a classroom table and never touches a patient. Does infection control apply?

13.10. Answers and reasoning

1. **Cleaning** physically removes soil and contaminants; **disinfection** inactivates many disease-causing microorganisms, with effectiveness depending on agent, concentration, contact time, temperature, and organism; **sterilisation** is a validated process intended to eliminate all microbial life including spores. Cleaning must come first because residual soil physically shields organisms and can inactivate the chemical agent.
2. Because indirect contact transmission via a **fomite** is a recognised route: hands touch the object, then touch something or someone else. The object is a link in the chain of infection regardless of whether a patient contacted it.

3. **Its intended use** — specifically what tissue it contacts. Critical items enter sterile tissue or the vascular system; semi-critical contact mucous membranes or non-intact skin; non-critical contact intact skin only. The material is irrelevant to the categorisation.
4. Because spores are highly resistant dormant forms that survive many disinfectants. Sterilisation is distinguished from disinfection precisely by being a process intended to eliminate them, defined to a stated sterility assurance level.
5. Biofilm is a surface-attached microbial community protected by a self-produced matrix, markedly more resistant to disinfectants and mechanical removal than free-floating organisms. Rough, scratched, porous, or crevice-rich surfaces provide protected habitat and shield organisms from both the chemical agent and the wiping action.
6. **Almost nothing about your part.** It is a claim about the material as supplied under stated conditions. Your part additionally involves thermal history, additives and colourants, machine contamination, layer porosity, surface roughness, support marks, residues, and post-processing. Evaluation is of a **finished device in its intended use**.
7. Because infill creates enclosed voids. While the shell is intact they are sealed and irrelevant; once breached — by a crack, a scratch through the wall, a drilled hole — they become an interior volume that cannot be reached, inspected, or cleaned.
8. Any three of: build orientation; layer height and nozzle; infill and wall count; temperature and flow rate; material batch or spool; machine calibration and wear; support placement and removal; post-processing; ambient conditions.
9. That the assembly survived fifty cycles under **your** handling conditions, with **your** applied forces, on **that** version. It does not establish long-term service life, does not establish behaviour under other people's handling, and establishes **nothing at all** about cleanliness, material safety, or clinical suitability.
10. Any five of: exposed threads; press-fit seams and split lines; blind holes and blind tubes; internal lattice or infill voids; sharp internal corners; support-scarred surfaces; layer lines; stitch welds and lap joints; burrs; knurling and texture; uncoated fabric, exposed foam, or hook-and-loop; stitch holes through a coating; paper labels; recessed fastener heads; gaps around buttons and moving parts.
11. Because PLA's glass transition is around 55–65 °C, far below steam sterilisation temperatures of ~121–135 °C. It will soften and deform. It is also unsuitable for thermal washer-disinfectors and often even for a dishwasher.
12. The material claim is about the resin as supplied. The part may carry mould-release residue, colourant or regrind that changes its properties, flash and parting-line crevices, sink voids, or moulded-in stress that cracks under disinfectants. A change of resin lot or moulder can change all of these. The claim is about a material; the evaluation is of the finished part in its intended use.
13. By **laundering** at a thermal-disinfection temperature (or by single-patient use) — not by wiping. Hook-and-loop traps lint, hair, and soil and degrades with washing; the webbing may shrink or lose strength; and the cleaning route must be agreed with whoever actually launders it.
14. Acrylic is prone to **crazing** — networks of fine cracks — with alcohols, especially where the material is stressed, such as laser-cut edges and drilled holes. The cracks spread, trap soil, and eventually the part breaks. A material can be fine in the workshop and fail on the first week of

the real cleaning route.

15. **Yes.** It is a shared object in a healthcare environment, handled by multiple people who then touch other things. It is non-critical in the Spaulding sense, which determines the **level** of processing — it does not exempt it. The host's infection prevention team decides what applies, and the cleaning and storage arrangement must be agreed before the object is accepted into a shared space.

13.11. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Tissues, materials and temperature — polymer temperatures against laundry and steam sterilisation, with evidence gaps.
- Biolab — log reduction, D-value and sterility assurance level.
- Solutions — available chlorine from tablets.

13.12. Video companion and reading

- **Core videos:** CDC *Inside Infection Control* [23]. Watch episodes on infection-control goals, surface transmission, cleaning versus disinfection, why cleaning matters, and contact time. These sit within a COVID-era series — take the foundational concepts and follow your host's **current** procedures for anything actual.
- **Core reading:** FDA *3D Printing of Medical Devices* [25]. Read it as a case study in how a regulator thinks about **any** manufactured object: the framing of finished-device-in-intended-use applies equally to a machined, moulded, cast, welded, or sewn part, and it is the idea to take away.
- **Reference:** CDC *Guideline for Disinfection and Sterilization in Healthcare Facilities* [24] for the definitions and the Spaulding classification. Read the definitions section; do not attempt to select a process from it.
- **Optional technical depth:** FDA *Technical Considerations for Additive Manufactured Medical Devices* [26], especially its cleaning and manufacturing sections.



These are US sources, chosen because they are freely readable and clearly written. They are not the requirements that apply to you. Your host institution and your jurisdiction determine that — in the UK, for example, decontamination in healthcare is governed by national guidance and local policy, and in the EU by the Medical Device Regulation and harmonised standards. Establish the applicable framework **with your host**, early.

13.13. Exit criteria

- ☐ Explain cleaning, disinfection, and sterilisation accurately, without treating them as interchangeable, and explain why order matters.

- ▢ Classify an item by Spaulding category and explain that the category follows from use, not material.
- ▢ Complete a contact map for your model, including the handling steps people forget.
- ▢ Identify five potential cleaning or contamination problems in a prototype, and propose a redesign for three with their costs.
- ▢ Produce a three-material comparison that clearly marks unknowns and does not rank overall.
- ▢ Explain what a material claim establishes and what it does not, in terms of finished device and intended use.
- ▢ Produce a traceability record from which someone else could reproduce your part, in the form that fits its fabrication method.
- ▢ For one object, name its cleaning route — wiped, laundered, reprocessed, or single use — and show that every material, joint, and fastening in it survives that route.
- ▢ Identify by role who at a host institution must review cleaning and patient-contact requirements.
- ▢ Explain why a bench wipe test does not establish validated clinical reprocessing.

13.14. What this chapter does not establish

- **You cannot select a disinfectant or a sterilisation cycle.** These are institutional decisions, validated against specific loads and equipment, with monitoring and documentation requirements.
- **You cannot declare any material safe.** Not for skin, not for a clinical environment, not for repeated processing. "Safe" is a claim about a finished object in a use, with evidence behind it.
- **You cannot validate a cleaning process.** Validation involves defined soil challenges, recovery methods, and acceptance criteria. Wiping something and looking at it is not that.
- **A durability result is not a cleanliness result.** Keep bench mechanical evidence strictly separate from any claim about contamination.
- **This chapter does not cover biocompatibility testing.** Material–tissue interaction has its own standards family (Appendix D). If a project involves anything beyond brief intact-skin contact, that becomes relevant and it is specialist territory.

Chapter 14. Making it: choosing the technology and the method

Week 17 · approximately 6–8 hours

14.1. Why this chapter exists

Every earlier chapter has asked **what the body needs**. This one asks **how you make something that meets that need** — and the honest first answer is that the method should be the last decision, not the first.

A workshop shapes how its owner sees problems. Someone with a printer sees printable problems; someone with a lathe sees turned parts; someone who writes software sees an app. A technologist who is genuinely useful in a clinical setting has more than one of those tools, and — more importantly — the judgement to pick the one the problem calls for, including the frequent case where the right answer is **not to make anything**.

You already have a wide toolset: software and electronics, additive manufacturing, CNC and other subtractive work, moulding and casting, metalwork, glasswork, and textiles. This chapter does not teach you to use them. It does three other things:

1. It gives you a **decision framework** that starts from the need and its constraints — quantity, contact, cleaning, load, accuracy, maintenance, risk tier — and only then reaches a method.
2. It goes through each technology family as a clinical setting sees it: what it is good at, where it fails, and the specific traps — cleanability, material certification, traceability, process variability — that matter more in healthcare than in a workshop.
3. It explains what changes as an object moves from a one-off to a batch to a product: why a hospital may prefer a commercial item, what **process validation** means, and what a maker must hand over.

Chapter 13 already covered cleaning and sterilisation, the gap between a material claim and a finished object, the variables of additive manufacturing, design for cleanability, and a material comparison. This chapter builds on all of it.

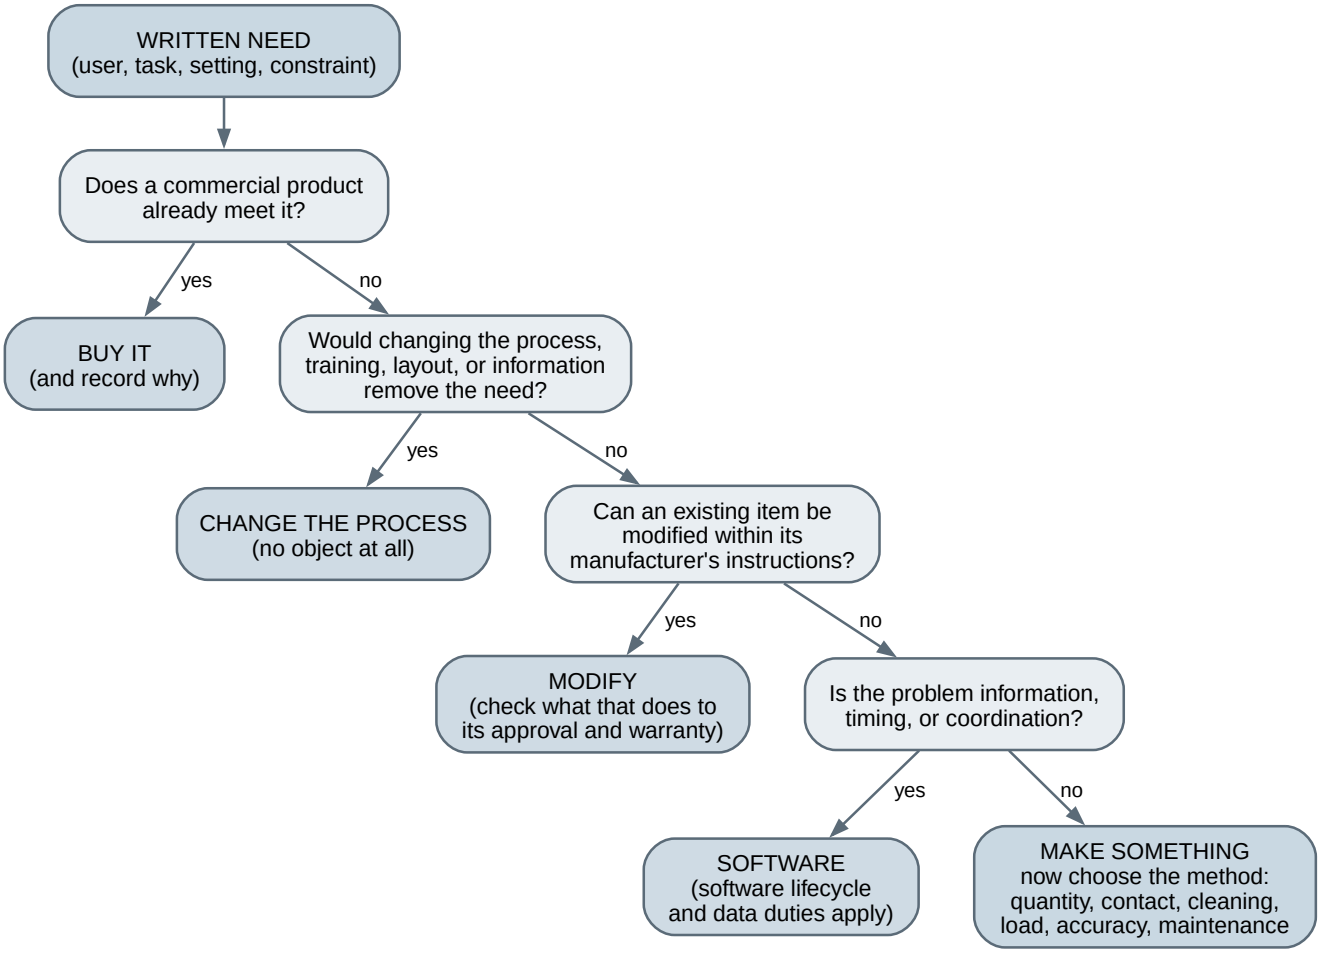
14.2. Start from the problem, not the tool

14.2.1. The questions that choose the method

Before any method is considered, the need should be written down (Chapter 18) and the following answered. Each answer removes options.

Question	Why it matters	What it rules in or out
How many, and how often?	One-off, a batch of tens, hundreds, or continuous supply	Tooling-based processes (moulding, pressing, die-cutting) only pay at volume; one-offs favour printing, machining, hand work, and software
What does it touch, and for how long?	Intact skin, broken skin, mucosa, nothing (Chapter 5)	Contact category drives material evidence and often removes the method entirely (Appendix I Tier 3)
How is it cleaned, by whom, how often, with what?	Wipe, washer-disinfector, steam, or never (Chapter 13)	Porous materials, crevices, exposed electronics, and heat-sensitive polymers fail here first
What loads, and what happens when it fails?	Static, cyclic, impact; failure mode slow or sudden (Chapter 6)	Layered and bonded structures are weakest across their joins; ductile metals warn before they break, many plastics and glass do not
How accurate must it be, and what does it mate with?	Fit with a commercial device, a body part, or another made part	Sets tolerance, which sets process and inspection method
Who maintains it, and for how long?	A ward, a technician, a family, nobody	Favours standard fasteners, available spares, simple materials, and documentation
Does it need power, sensing, or logic?	Indication, measurement, control, communication	Brings in electronics, firmware, batteries, EMC, and software lifecycle obligations
Where is it in the risk tiers?	Appendix I	A Tier 3 object is not a fabrication question for you at all

14.2.2. The options before fabrication



The diagram is deliberately pessimistic about making things. That is not modesty; it is how a clinical engineering department thinks. A commercial product arrives with evidence, a manufacturer who answers for it, spares, a cleaning instruction validated by someone, and a recall route. A made object arrives with **you**, and you will not be on the ward at three in the morning.

The four alternatives deserve their place:

- **Buy it.** The search is part of the work. A good brief often finds an existing product nobody on the team knew about — and that is a successful outcome, worth recording.
- **Change the process.** The "lost item" that needs a holder may be better solved by moving where it is stored. The "awkward" sequence may be better solved by a checklist.
- **Modify.** A small adaptation to a commercial item can be the most useful contribution of all — and it can also void the manufacturer's evidence for that item. Any modification of a medical device is a question for the clinical engineering team, not a workshop decision.
- **Software.** If the difficulty is **knowing** — what is where, what happens next, what the state is — a physical object may be the wrong medium.

Only after those does the method question begin.

14.3. The technology families

Each section follows the same pattern: what the family is good at, what it is bad at, orientation figures, and the traps that matter most in clinical use. The figures are typical shop-floor orders of

magnitude, not specifications; machines, materials, and operators vary widely.

14.3.1. Software

Software is the cheapest thing to change and the hardest thing to see. It belongs at the front of the list because many needs that arrive as "can you make a …?" are information problems: scheduling, tracking, reminding, explaining, displaying.

- **Good at:** information, coordination, teaching, simulation, data handling, and anything that must change often.
- **Bad at:** anything requiring physical force, physical protection, or reliability when the power, the network, or the login fails.
- **Clinical traps:** intended use can make it a medical device; health data brings legal obligations; it must be maintained for as long as it is used (Chapter 20).

When software beats hardware: the need is for information at the point of use; the "physical" fix would need frequent revision; or the same problem exists in many places and copying a file is cheaper than making objects. **When hardware beats software:** the user cannot or will not use a screen; the environment is wet, gloved, or hurried; or failure must be visible and mechanical rather than silent.

14.3.2. Electronics and embedded systems

Sensors, microcontrollers, displays, motors, LEDs, wireless links, and batteries let an object **sense, decide, and signal**. They are also the point at which a project acquires the most new obligations.

Concern	What it means in practice
Patient connection	Anything electrically connected to a person — electrodes, conductive contacts, sensors in body contact powered from mains — is IEC 60601-1 territory: leakage currents, isolation, applied parts. This is the absolute boundary of Chapter 7
Power	Low-voltage, battery-powered, self-contained designs are the realistic space for a maker. Mains-powered designs bring insulation, creepage, and earthing requirements
Batteries	Lithium cells store a great deal of energy in a small volume. Use protected cells or packs with protection circuits, certified chargers, and enclosures that tolerate swelling. Cells and packs have their own safety and transport standards (Appendix D)
EMC	A device must neither be disrupted by nor disrupt nearby equipment. Hospitals are dense with radio transmitters and sensitive monitors (IEC 60601-1-2 for medical equipment)
Enclosure and ingress	IP ratings (IEC 60529) describe protection against solids and water: IPX4 resists splashing, IP67 survives brief immersion. Wipe-cleaning in practice means fluid running into every seam

Concern	What it means in practice
Cleaning chemicals	Alcohols, quaternary ammonium compounds, chlorine releasers, and peroxides attack membrane keypads, label adhesives, some polycarbonate and acrylic housings (environmental stress cracking), and exposed connectors. Check the chemical, not the word "wipeable"
Firmware	Firmware is software. If the device's function matters clinically, so does its version control, testing, and update process (Chapter 20)
Failure	Electronics fail silently. Ask what the user sees when a sensor drifts, a battery is low, or the microcontroller resets — and design that indication deliberately (Chapter 7' redundant coding)

- **Good at:** indication (lights, sounds, haptics), measurement of the environment or an object, timing, logging, automation, accessibility (switch interfaces, environmental control).
- **Bad at:** tolerating fluids, heat, impact, and cleaning; working for years without attention; failing obviously.
- **Realistic medical examples:** switch-adapted toys and controls for children with disabilities; timers and prompts; training aids that light or sound when a step is done correctly; test fixtures and signal simulators for bench checking equipment; data loggers for temperature or usage of equipment (not people).

14.3.3. Additive manufacturing

Additive manufacturing builds a part layer by layer. It is one family among many, with real strengths and clear limits; Chapter 13 covered the process variables that matter for cleanability.

Process	Typical accuracy	Typical surface (Ra, as built)	Character
Material extrusion (FDM/FFF)	~±0.2–0.5 mm	~10–25 µm or worse; visible layers	Cheap, fast to iterate, wide polymer choice. Weak between layers; porous; infill voids
Vat photopolymerisation (SLA, DLP)	~±0.05–0.15 mm	~1–5 µm	Fine detail and smooth surfaces. Resin chemistry; washing and post-cure are critical; many resins are brittle and degrade in UV
Powder bed fusion, polymer (SLS, MJF)	~±0.2–0.3 mm	~5–15 µm; slightly porous	Strong nylon parts without supports; complex internal geometry. Powder trapped in cavities; surface porosity
Powder bed fusion, metal (LPBF/SLM)	~±0.05–0.2 mm	~5–15 µm as built	Titanium, cobalt–chrome, stainless; porous implant surfaces and lattices. Industrial equipment, heavy post-processing, powder hazards

- **Genuine strengths:** complex one-offs; geometry derived from imaging (anatomical models, patient-matched guides — a regulated activity when used on a patient); internal channels; rapid

iteration; jigs and fixtures for other processes; short runs of parts no one else makes.

- **Wrong choice when:** the part must be wiped, washed, or sterilised repeatedly and the process leaves layers and voids; the load runs across layers; a cheap standard item exists; the quantity justifies a tool; or tolerance and surface must be guaranteed part to part.
- **Clinical traps:** everything in Chapter 13' "what actually varies" table, plus the temptation to treat a successful print as a finished device.

Mature medical uses exist, and they are instructive precisely because of the discipline around them: dental models and aligner moulds, surgical planning models, patient-matched cutting guides, and porous titanium implant surfaces — all made under quality systems with validated processes [26].

14.3.4. CNC machining and other subtractive methods

Milling, turning, routing, drilling, grinding, and profile cutting by laser, waterjet, or plasma start from solid stock and remove material.

Process	Typical accuracy	Typical surface (Ra)	Character
CNC milling and turning	~±0.02–0.1 mm routinely; tighter with care	~0.8–3.2 µm as machined; < 0.4 µm polished	Solid, dense, isotropic material with certified properties. Geometry limited by tool access
CNC routing (sheet and plastics)	~±0.1–0.2 mm	Depends on material and cutter	Fast for flat parts in HDPE, acetal, acrylic, plywood, composites
Laser cutting	~±0.1 mm; kerf ~0.1–0.3 mm	Edge quality varies by material	Very fast for sheet; heat-affected edges. Some plastics release hazardous fumes (PVC releases chlorine compounds — never laser-cut it)
Waterjet	~±0.1–0.25 mm	Slightly rough, tapered edge	Cuts metal, glass, stone, and thick plastics without heat

- **Strengths in clinical use:** the part is the **certified stock material**, fully dense, with no internal voids — exactly what Chapter 13 asks for. Surfaces can be finished smooth enough to clean well. Material certificates trace the stock to its melt. Repeatability is excellent once the program is proven.
- **Weaknesses:** tool access limits internal geometry; material waste; setup and programming time per design; sharp edges and burrs unless deliberately removed.
- **Clinical traps:** burrs and sharp edges injure skin and tear gloves; machining fluids and coolants must be removed; deep tapped holes and internal corners trap soil; some plastics (acrylic, polycarbonate) craze when cleaned with alcohol after machining stresses.
- **Medical examples:** surgical instruments, implant components (under quality systems), equipment brackets and mounts, test fixtures, phantom housings, bench jigs, adaptive equipment parts in acetal or HDPE.

14.3.5. Moulding, casting, and forming

These methods shape material against a **tool** — a mould, a die, a former — which costs time and money once and then produces many identical parts.

Method	Good at	Watch for
Injection moulding	Thousands to millions of identical, accurate, smooth parts at low unit cost. The dominant process for medical disposables	Tooling from a few thousand (aluminium, prototype) to tens or hundreds of thousands (multi-cavity steel) and weeks of lead time. Design rules: uniform walls (~1–3 mm), draft (~0.5–2°), no undercuts without slides. Medical-grade resins with change notification; cleanroom moulding for sterile products
Overmoulding and insert moulding	Soft grips on rigid cores; embedded metal inserts; sealed electronics	Bond strength between materials; crevices at the interface
Silicone casting (platinum- or tin-cure)	Soft, skin-like, durable parts in small batches: simulators, grips, seals, pads. Hardness from very soft gels to Shore A ~80	Cure inhibition (platinum cure is poisoned by sulphur, some printed masters, latex); bubbles; tear propagation; release agents left on surfaces
Polyurethane and resin casting	Rigid or elastomeric short runs from silicone moulds	Isocyanate handling safety; moisture sensitivity; shrinkage
Vacuum forming and thermoforming	Shells, trays, covers, splints, orthoses, positioning aids from sheet (PETG, polypropylene, polyethylene, ABS, low-temperature thermoplastics)	Thinning at corners and deep draws; trimmed edges; a single-sided tool controls only one surface
Dip moulding and rotational moulding	Hollow, seamless parts; covers and bladders	Wall-thickness variation

The key engineering decision is the **cost model**, worked in the bridge section below: tooling is a fixed cost, so these methods win only beyond a break-even quantity. For a maker, the most accessible members of the family — silicone casting, vacuum forming, and low-pressure resin casting — are also the most medically useful, because **soft, cleanable, skin-compatible-looking** surfaces are exactly what simulators, grips, and pads need.

Orthotics and prosthetics workshops are thermoforming shops: polypropylene and polyethylene heated and drawn over a plaster or digitally milled positive. Many of the skills transfer directly; the clinical judgement does not (Chapter 6).

14.3.6. Metalwork

Sheet metal, fabrication, welding, brazing, forging, and finishing remain central to healthcare

equipment: trolleys, frames, rails, brackets, instruments, and fittings.

Material	Strengths	Clinical traps
Stainless steel 304 (A2)	General-purpose, formable, weldable, widely available	Pits in chloride environments — chlorine-releasing disinfectants and saline
Stainless steel 316 / 316L (A4)	Molybdenum improves chloride resistance; "L" (low carbon) resists corrosion near welds. The usual choice for equipment that is wiped and washed	Still corrodes if contaminated with carbon-steel particles; needs passivation after fabrication
Martensitic stainless (e.g. 410, 420)	Hardenable; used for cutting instruments	Lower corrosion resistance; needs correct heat treatment and passivation
Titanium and its alloys	Light, strong, corrosion-resistant, biocompatible; implant alloys such as Ti-6Al-4V	Expensive; galling; difficult to weld outside inert conditions
Aluminium alloys	Light, easy to machine and form; anodising gives colour and hardness	Anodising is attacked by alkaline detergents used in washer-disinfectors; galvanic corrosion against stainless (Chapter 3)
Mild steel	Cheap, strong, easy to weld	Rusts; painted or plated finishes chip, exposing it. Rarely appropriate in clinical areas

- **Passivation** — a chemical treatment that removes free iron and restores the protective chromium oxide film on stainless steel — is a routine finishing step after machining or fabrication, not an optional extra.
- **Welds and brazed joints** create crevices, heat tint, and rough beads. Continuous, ground, and passivated welds clean well; stitch welds and overlapping joints trap soil.
- **Fasteners** should match the base metal (stainless with stainless) and avoid exposed threads in cleaned areas.
- **Medical examples:** drip stands, trolleys, frames for patient lifting (Tier 3 for load), instrument trays, bed accessories, wall-mounted brackets, and the stainless benches of sterile services.

14.3.7. Glasswork

Glass is chemically inert, transparent, heat-resistant in its borosilicate form, and supremely cleanable — and it breaks into sharps. Healthcare uses it carefully.

- **Borosilicate glass** (expansion around 3.3×10^{-6} per K) tolerates heat and thermal shock far better than ordinary soda-lime glass (around 9×10^{-6} per K). It is the material of laboratory glassware, many vials, and scientific glassblowing.
- **Scientific glassblowing** produces bespoke laboratory apparatus, flow cells, condensers, and

repair of existing glassware. Universities and research laboratories still keep glassblowers for exactly the one-off work a printer cannot do in an inert, heat-proof material.

- **Optics** — lenses, prisms, windows, mirrors — underpin endoscopes, microscopes, ophthalmic instruments, and the eye models of Chapter 8. Optical work is a precision craft of its own.
- **In phantoms and teaching models,** glass tubes and vessels give clear, stable, rigid channels — for flow phantoms, for instance — where plastics would scratch, cloud, or absorb dyes.
- **Clinical traps:** breakage creates sharps and, in patient areas, risk of injury or self-harm (Chapter 12); thermal shock fractures soda-lime glass; glass drug ampoules can shed fragments when opened, which is why filter needles exist. Clinical areas therefore replace glass with polymers wherever transparency does not require glass.

14.3.8. Textiles and soft goods

Straps, slings, garments, covers, pads, positioning aids, harnesses, and bags are made from fabrics — and much of daily clinical life rests on them. Textiles are the technology closest to the body's own soft tissue (Chapter 5).

Consideration	In practice
Fibre and construction	Woven fabrics resist stretch; knits stretch and conform; nonwovens are cheap and often disposable. Polyester and nylon dominate where strength and laundering matter
Technical fabrics	Spacer fabrics (breathable three-dimensional knits for pressure and moisture management), laminated or coated fabrics (impermeable, wipeable covers for mattresses and cushions), elastics with specified stretch for compression
Seams	Every stitch hole is a path for fluid. Covers that must be fluid-proof use welded (ultrasonic or radio-frequency) or taped seams; a stitched cover over foam lets fluid into the foam
Laundering	Healthcare linen is commonly thermally disinfected — for example at 65 °C for 10 minutes or 71 °C for 3 minutes — and tumble dried. Elastics, foams, adhesives, prints, and hook-and-loop all age under that regime. Design for the actual laundry process, not a domestic wash
Flammability	Textiles and upholstered items in healthcare are subject to fire-performance requirements; check what applies before choosing fabric or foam
Hook-and-loop fasteners	Convenient, adjustable, and accessible for weak hands — and they collect lint, hair, and soil, lose grip with laundering, and are hard to clean. Close them before washing; consider buckles or magnetic closures where cleaning matters
Fit and pressure	A strap is a pressure source along its edge. Width, edge finish, and padding determine whether it is comfortable or injurious (Chapter 5, Chapter 8)

Consideration	In practice
"Antimicrobial" claims	Treated fabrics may reduce microbial growth on the fabric under test conditions; they do not replace cleaning, and the claim itself may be regulated
E-textiles	Conductive yarns, printed circuits, and sensors in fabric. Washing, flexing, and connection points are the usual failure modes

- **Medical examples:** positioning wedges and covers, slings for limbs, wheelchair and seating accessories, compression garments (Tier 3 when therapeutic), bags and pouches for equipment, weighted or sensory items for children, carrying solutions for oxygen or feeding pumps, adaptive clothing with accessible closures.
- **Pattern-making** is CAD in another medium: a flat pattern is the development of a three-dimensional surface, and grading a pattern across sizes is the anthropometry of Chapter 6.

14.3.9. Wood and other workshop materials

Wood, cork, leather, card, and foam board are honest prototyping and workshop materials. They are excellent for:

- concept models and full-size mock-ups for discussion;
- jigs, templates, and bench fixtures;
- objects for home, community, or outdoor use where clinical cleaning does not apply — play equipment, garden aids, adapted furniture for a person's own home;
- the first cardboard version of almost anything (Chapter 22).

They are generally **excluded from clinical areas** because they are porous and absorbent, so they cannot be cleaned to a standard and they harbour microorganisms (Chapter 13). Where wood does appear in healthcare equipment — seating bases in wheelchairs, for example — it is enclosed, sealed, and covered, and the cover is the cleanable surface.

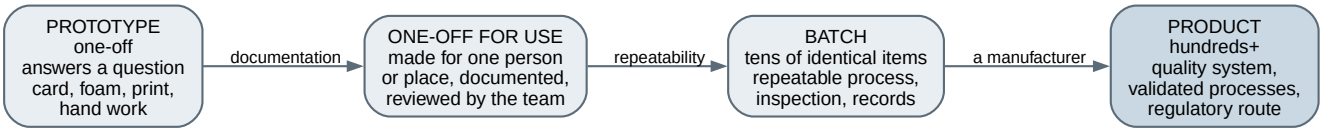
14.3.10. Combining methods

Most good objects are **hybrid**: a machined aluminium frame with a moulded grip; a printed housing around a commercial sensor; a laser-cut acrylic panel with a silicone gasket; a textile cover over a thermoformed shell. Each interface between methods is also an interface between materials, with its own crevice, bond, galvanic couple, or thermal mismatch.

The discipline that keeps hybrids sane is **design for manufacture and assembly** (DFMA):

- fewer parts, each with an obvious orientation;
- standard fasteners and standard stock wherever possible;
- features that make wrong assembly impossible (Chapter 7' design to eliminate);
- parts that can be inspected, cleaned, and replaced individually;
- every part traceable to a drawing, a file, a material, and a process.

14.4. From prototype to product



Each step to the right adds obligations that have nothing to do with the workshop:

- **Prototype.** Its job is to answer a question. It can be crude; it must be labelled as a prototype and kept away from use.
- **One-off for use.** Made for a specific person or setting under the host's review process. It needs drawings, materials, cleaning instructions, a version, and a named owner (Chapter 19).
- **Batch.** Thirty identical teaching models or fifty brackets are no longer a craft object. Every unit must match the reviewed design, which means a written process, jigs, inspection criteria, and records of which materials went into which units.
- **Product.** A medical device placed on the market requires a manufacturer with a **quality management system** (ISO 13485; in the US the Quality Management System Regulation, which incorporates it by reference), design controls, risk management, and a regulatory route. That is a company's job, and it is the point at which a maker's design is usually transferred to one.

14.4.1. Process validation

Some processes produce outputs that cannot be fully checked afterwards. You cannot inspect every weld internally, test every moulded part to failure, or prove every sealed pouch is sterile. For these, the **process itself** is validated, so that its output can be trusted without destructive testing:

- **IQ — installation qualification:** is the equipment installed and set up as specified?
- **OQ — operational qualification:** across the allowed range of settings, does the process produce acceptable output? Where are its limits?
- **PQ — performance qualification:** under normal production conditions, does it consistently produce acceptable output?

Moulding, welding, heat sealing, sterilisation, adhesive bonding, and many additive processes are typical candidates. You do not need to run a validation; you do need to recognise when "we tested one and it was fine" is not an answer.

14.4.2. Supplier controls and material certificates

A certificate of conformity or a mill certificate says what a batch of material **is**: alloy, composition, properties, heat number. Medical-grade polymers come with biocompatibility test data and, crucially, a **change notification** commitment — the supplier will tell you before changing the formulation. A hardware-store sheet of plastic comes with neither. When a project moves from one-off to batch, **where the material came from** becomes part of the design.

14.4.3. Why the hospital may prefer to buy

It is common, and correct, for a clinical engineering team to choose a commercial item over a better-fitting made one. The commercial item has a manufacturer responsible for it, evidence, validated cleaning instructions, spare parts, service, and a recall route. Your made item has your time. That is not a judgement on your skill; it is a judgement on **continuity**. The way to compete is to make continuity part of what you deliver.

14.4.4. What a maker hands over

Item	Contents
Drawings and models	Dimensioned drawings with tolerances; native CAD and neutral formats (STEP for solids, DXF for profiles, Gerbers for circuit boards); pattern files for textiles
Bill of materials	Every part, material, grade, supplier, and part number, including fasteners, threads, adhesives, and finishes
Process notes	Machine settings, CAM programs, slicer profiles, mould and cure parameters, sewing specifications, firmware build instructions and versions
Inspection and test	What is checked on each unit, how, and the acceptance criteria (Chapter 19)
Use and care	Intended use, cleaning method and agents, inspection for damage, end-of-life criteria
Spares and repair	Which parts wear, how to replace them, and where to get them
Limitations and status	What has and has not been tested, reviewed, or approved — the honest statement of Chapter 22

14.5. Engineering bridge: costs, tolerances, surfaces, and energy

14.5.1. The cost model

Every method has a **fixed cost** (tooling, programming, setup, fixtures) and a **variable cost** per unit (material, machine time, labour). Total cost for **n** units:

$$C(n) = F + v \times n$$

Two methods break even where their totals are equal:

$$n_{BE} = (F_A - F_B) / (v_B - v_A)$$

Work an invented example — a small cleanable clip:

	Printed and finished by hand	Injection moulded
Fixed cost F	~£0 (design already done)	~£8,000 (aluminium tool)
Variable cost v	~£6.00 per part (material, machine time, finishing labour)	~£0.80 per part

n_{BE} = (8000 − 0) / (6.00 − 0.80) ≈ 1,540 parts.

Below about 1,500 parts, printing is cheaper; above it, moulding is. But the comparison is incomplete without the **non-financial** columns: the moulded part is dense, smooth, repeatable, and made from certified resin; the printed part is none of those. For a cleaned, repeated-use item, the break-even on **suitability** may arrive long before the break-even on cost. Many methods sit between these extremes — CNC, casting in silicone moulds, thermoforming — with modest fixed costs and moderate unit costs.

14.5.2. Tolerance stack-up

When several parts sit in a chain, their tolerances add. For four parts each ±0.1 mm:

- **Worst case:** ±(0.1 + 0.1 + 0.1 + 0.1) = **±0.4 mm**.
- **Statistical (root sum of squares):** ±√(4 × 0.1²) = **±0.2 mm** — the likely spread if errors are independent and centred.

Design to the worst case where the consequence of misfit is serious; use the statistical estimate where it is not and the volume is high. Either way, **the method's typical tolerance must be smaller than the chain allows**, which is often the decisive reason to choose machining over printing for a part that mates with a commercial device.

14.5.3. Surface roughness and cleanability

Surface roughness is commonly reported as **Ra**, the arithmetic mean deviation of the surface profile, in micrometres. Hygienic-design guidance in food and pharmaceutical equipment commonly targets around **0.8 µm Ra or smoother** for product-contact surfaces; surgical instruments are typically finer still. Compare:

Surface	Typical Ra (orientation)
Polished stainless, electropolished	~0.1–0.4 µm
Injection moulded, polished tool	~0.05–0.4 µm
Machined (milled or turned)	~0.8–3.2 µm
Resin print (SLA/DLP)	~1–5 µm
Powder bed print, polymer or metal, as built	~5–15 µm
FDM print, as built	~10–25 µm or worse
Bacterium, for scale	~1 µm

Roughness is not the whole story — crevices, porosity, and geometry matter more — but the table shows at a glance why a bacterium-sized valley is a feature of some processes and not others.

14.5.4. Battery energy

A typical lithium-ion cylindrical cell (the common 18650 size) stores roughly 9–13 Wh at a nominal 3.6–3.7 V. A small four-cell pack therefore holds around 40–50 Wh — about 150–180 kJ — in something the size of a phone. If a cell is short-circuited, crushed, overcharged, or manufactured with a defect, it can enter **thermal runaway**: internal heating that releases flammable electrolyte and oxygen, spreading cell to cell and very hard to extinguish.

Design consequences: protected cells or packs from reputable sources; a protection circuit for overcharge, over-discharge, over-current, and short circuit; a charger designed for the chemistry; mechanical protection from crushing and puncture; venting space; and no improvised cells in anything that will sit on a ward, near oxygen, or on a person. Primary (non-rechargeable) cells and coin cells have their own hazard, especially ingestion (Chapter 8).

14.5.5. Process capability, in brief

For a batch, the question is not "is this part in tolerance?" but "will the next hundred be?". **Process capability** compares the spread of a process with the tolerance:

$$C_p = (USL - LSL) / 6\sigma, \text{ and } C_{pk} = \min(USL - \mu, \mu - LSL) / 3\sigma$$

where USL and LSL are the upper and lower specification limits, σ the process standard deviation, and μ the process mean. A C_{pk} of 1.0 means the process just fits; 1.33 or more is a common target. Measure ten consecutive parts from your printer, your mill, or your mould; calculate C_{pk} for one critical dimension; and you have learned more about your process than a single good part could tell you.

14.6. Devices, software, and training technology

14.6.1. Where making happens in healthcare

Setting	What is made or maintained there	How a maker fits
Clinical engineering (medical physics and engineering) workshops	Repairs, modifications, mounts, fixtures, test equipment, bespoke one-offs under the hospital's procedures	The natural host for supervised work; they own the review process and the risk
Hospital and university maker spaces	Prototypes, training aids, research apparatus	Often the entry point; check what governance applies to anything leaving the room
Orthotics and prosthetics laboratories	Thermoformed orthoses, laminated sockets, footwear modifications, scanning and milling of positives	Thermoforming, lamination, CAD/CAM. Clinical prescription and fitting stay with the orthotist or prosthetist

Setting	What is made or maintained there	How a maker fits
Rehabilitation engineering and assistive technology services	Seating, switch access, environmental controls, communication-aid mounting	Electronics, metalwork, and textiles together; strongly user-led
Dental laboratories	Crowns, bridges, dentures, aligners, guards — milled, printed, cast, and pressed	One of the most mature digital manufacturing workflows in medicine
Sterile services	Not a making department, but the one that decides whether your object survives reprocessing	Ask them early (Chapter 13)
Medical illustration and simulation technicians	Moulage, task-trainer repair, props, casting	Silicone, textiles, paint, and craft skills in daily use

14.6.2. Software and design tools

Tool family	Examples	Note
Mechanical CAD	FreeCAD, Onshape, Fusion, SolidWorks, OpenSCAD	Parametric models make variants and size ranges cheap; keep native files and neutral exports
CAM	CAM in FreeCAD or Fusion; slicers for additive work	The toolpath or slicer profile is part of the design record
Electronics CAD	KiCad and commercial equivalents	Schematics, layouts, bills of materials, and Gerbers; version them like code
Firmware	Embedded toolchains and frameworks (Arduino, PlatformIO, Zephyr and others)	Firmware is software; see Chapter 20 for lifecycle and security
Simulation	FEA (FreeCAD FEM, CalculiX, commercial packages), circuit simulation	Useful for comparison and insight; validated against a physical test before trusted
Pattern and textile CAD	Seamly2D and commercial pattern software	Parametric patterns graded to anthropometric sizes
Laser and profile cutting	Vector drawing tools (Inkscape and others) and machine software	Engrave labels and version numbers directly into parts — durable traceability
Document and record control	Version control (Git) for CAD, firmware, and documents alike	The single most useful habit for traceability (Chapter 19)

14.6.3. Training and simulation made by each method

Method	Training technology	Note
Silicone casting	Suture pads, skin and wound models, injection pads, examination models with embedded findings	Consumables — constantly needed (Chapter 5, Chapter 11)
Machining and metalwork	Rigid frames for task trainers, articulated skeletal models, mounting hardware, probe holders	Durable; cleanable; repairable
Thermoforming	Shells for manikin parts, trays for kits, covers	Cheap, light, wipeable
Additive manufacturing	Anatomical models from imaging, airway trees, complex joints, bespoke connectors	Strong where geometry is complex and quantities small (Chapter 15)
Electronics	Feedback for CPR depth, needle position, pressure, or correct sequence; simulated monitors on fictional data	Never patient-connected; label simulated displays unmistakably (Chapter 9)
Glass and optics	Eye models, fundoscopy trainers, flow phantoms	Precision and clarity; breakage risk managed by enclosure
Textiles	Wearable simulators, moulage garments, positioning practice aids, manikin clothing	Fit to a manikin is anthropometry again
Software	Scenario engines, virtual microscopy, image libraries, quizzes	Often the cheapest to update and distribute (Chapter 21)



Range is the asset. The person who can say "this should be a thermoformed shell with a machined hinge and a cast silicone grip, and the counter is a twenty-line microcontroller program — or actually, the ward needs a different storage location and no object at all" is rare and genuinely valuable. Present your skills as a **range of options weighed against a need**, not as access to a particular machine.

14.7. Worked example: thirty elbow models for a teaching programme

Your tabletop elbow model — two segments, a pivot, a cord standing in for a tendon — has been reviewed and approved as a single teaching aid. The educator now asks for **thirty**, to be used by groups of students across a year, wiped between sessions with the department's standard detergent-and-alcohol wipes, and stored in a cupboard. Labelled "PULL TO BEND" as in Chapter 7.

Step 1 — the requirements that choose the method.

- Quantity: 30, plus spares; possibly another 30 next year.

- Cleaning: frequent wipe-down with detergent and alcohol. No immersion, no heat.
- Load: hand forces only; dropped onto a hard floor occasionally.
- Accuracy: the pivot must be smooth and repeatable; parts must interchange so a broken segment can be replaced from spares.
- Labels and coding: must survive a year of wiping.
- Maintenance: a technician, not you, will replace cords and pivots.

Step 2 — candidate methods.

Candidate	Strengths	Problems for this need
FDM printing, as the original	No tooling; same files	~120 parts with layer lines that collect soil; variability between prints; printed or sticker labels wear off; slow; inconsistent pivot fit across thirty
Laser-cut acrylic	Fast, cheap, clear, crisp engraved labels	**Acrylic crazes and cracks with alcohol wipes**, especially at laser-cut edges with residual stress — a classic environmental stress cracking failure (Chapter 3)
CNC-routed acetal or HDPE sheet, stainless pivot	Chemically robust to the wipes; dense, smooth, cleanable; engraved labels; parts identical and interchangeable	Some programming and fixturing time; acetal does not glue easily (design for mechanical fastening)
Silicone or urethane casting	Good for soft parts	Unnecessary for rigid segments; slower per unit
Injection moulding	Ideal parts	Tooling cost far beyond the break-even at 30–60 units

Step 3 — the choice, and why. CNC-routed acetal or HDPE segments, with a stainless 316 shoulder bolt as the pivot and a nylon washer; a replaceable braided polyester cord with a crimped end; labels and arrows **engraved** into the surface and colour-filled, so the redundant coding of Chapter 7 survives cleaning. A routing jig makes every segment identical; an assembly jig sets the cord length. The original printed model remains the reference prototype.

Step 4 — what to check before making thirty. Confirm the exact wipe product with infection prevention and control and check its compatibility with acetal or HDPE and the label fill. Make **three** first, wipe them a few hundred times on the bench with the actual product, drop them, and inspect. Measure the pivot bore on all three and calculate the spread. Only then cut the rest.

Step 5 — the handover. Drawings and CAM files; bill of materials with material grades and suppliers; jig drawings; assembly and inspection instructions; cleaning instructions naming the agent tested; spare-parts list (cords, bolts, washers, one spare of each segment); end-of-life criteria (cracks, illegible labels, a pivot that binds); and the status statement: approved design, batch made on date **X**, three units wipe-tested, no clinical use.

The lesson. The method that made the best prototype was not the method that made the best batch. And the decisive fact — that the department's wipes would crack acrylic — came from asking how the object would be **cleaned**, not how it would be **made**.

14.8. Workshop exercise

1. **Decision framework.** Take three real or invented needs — a holder for a feeding pump on a wheelchair, a training model of the knee for physiotherapy students, and a reminder system for a care-home medication round. Run each through the decision diagram and the question table. Record where you stopped and why. At least one should end without fabrication.
2. **Method comparison.** For one object you have made recently, write the comparison table of the worked example for at least four methods from different families. Include the cleaning agent and the quantity as explicit columns.
3. **Break-even.** Get real or realistic quotes for a small part by printing, CNC, and moulding (online quoting services make this quick). Plot cost against quantity from 1 to 10,000 on a log scale and mark the break-even points.
4. **Wipe test.** Take offcuts of five materials you use — for example acrylic, polycarbonate, PETG, acetal, and anodised aluminium. Wipe each 200 times with an alcohol wipe, flexing the acrylic and polycarbonate slightly while wet. Inspect under a magnifier for crazing, softening, colour change, or label lift.
5. **Capability.** Make ten copies of one simple part by a method you use. Measure one critical dimension on each, calculate the mean, standard deviation, C_p, and C_{pk} against a tolerance you choose, and decide whether your process could supply a batch.
6. **Handover pack.** Write the handover pack of the table above for one object you have made, then give it to someone else and ask them to identify a part, a material, and a cleaning agent without asking you.

14.9. Questions

1. Name four alternatives to fabrication that should be considered first, and give an example of each.
2. Which questions about a need most strongly constrain the choice of manufacturing method?
3. Why is a battery-powered, self-contained electronic aid a realistic maker project when a mains-powered, body-contacting one is not?
4. List three ways cleaning chemicals damage electronic devices.
5. Give two genuine strengths of additive manufacturing in medicine and two situations where it is the wrong choice.
6. Why does a machined part from certified stock often clean and perform better than a printed part of the same geometry?
7. What is the break-even quantity between two methods, and what does a cost-only comparison leave out?
8. Why should 316L rather than 304 stainless usually be chosen for equipment cleaned with

chlorine-releasing agents, and what finishing step should follow fabrication?

9. Why do clinical areas avoid glass, and where is it still the right material?
10. Why does a stitched cover over foam fail as a fluid barrier, and what alternatives exist?
11. Give two problems with hook-and-loop fasteners in clinical use.
12. What is process validation, and for which kinds of process is it necessary?
13. Four parts in a chain each have a tolerance of ±0.05 mm. What are the worst-case and statistical stack-ups?
14. List five items a maker should hand over with a batch of made objects.

14.10. Answers and reasoning

1. **Buy** an existing product (a commercial holder exists); **change the process** (move the storage location rather than make a holder); **modify** an existing item within its manufacturer's instructions (an approved accessory rail); **software** (a reminder or tracking app instead of a physical tag). Each removes the need to make and maintain a new object.
2. Quantity, contact with the body, the cleaning or sterilisation regime, loads and failure consequences, required accuracy and mating parts, who maintains it, whether it needs power or logic, and its risk tier. Quantity and cleaning usually eliminate the most options fastest.
3. Self-contained low-voltage battery power with no electrical connection to the person avoids the leakage-current, isolation, and applied-part requirements of IEC 60601-1 that apply to patient-connected and mains equipment — the failure mode of which is electrocution. The realistic space is indication, timing, logging, and accessibility.
4. Any three of: alcohols and quaternary ammonium compounds cause environmental stress cracking in some polycarbonate and acrylic housings; fluids wick into seams, connectors, and membrane keypads; chlorine releasers corrode contacts and metal parts; solvents lift label adhesives and printed legends.
5. **Strengths:** complex one-off or patient-matched geometry (planning models, guides); internal channels and lattices; rapid iteration; jigs for other processes. **Wrong choice:** repeatedly cleaned or sterilised items where layers and voids trap soil; loads across layers; quantities beyond a tooling break-even; parts needing guaranteed tolerance and surface part to part; cases where a cheap standard item exists.
6. It is fully dense and isotropic with no internal voids or layer boundaries, its material is traceable by certificate, its surface can be finished smooth, and the process is highly repeatable — all properties Chapter 13 asks of a cleanable object.
7. The quantity at which the two methods' total costs (fixed plus variable) are equal: $n_{BE} = (F_A - F_B)/(v_B - v_A)$. A cost-only comparison ignores suitability — cleanability, material certification, consistency, strength, lead time, and maintenance.
8. 316L contains molybdenum, which improves resistance to chloride pitting, and low carbon, which reduces corrosion susceptibility at welds. After machining or fabrication it should be **passivated** to remove free iron and restore the protective oxide film.
9. Glass breaks into sharps, creating injury and self-harm risks, and soda-lime glass fractures with thermal shock. It remains right where inertness, heat resistance, transparency, or optical

quality are essential: laboratory glassware, optics, vials, and some phantoms — usually in controlled settings or enclosed.

10. Every stitch hole is a path for fluid into the foam, which then cannot be cleaned. Fluid- proof covers use welded or taped seams and laminated or coated fabrics.
11. Any two of: they collect lint, hair, and soil; they are hard to clean; they lose grip after repeated laundering; open hook surfaces snag other fabrics and skin.
12. Establishing by objective evidence that a process consistently produces output meeting its requirements — through installation, operational, and performance qualification — so that each output need not be fully tested. It is necessary where output cannot be fully verified by inspection: moulding, welding, sealing, bonding, sterilisation, and many additive processes.
13. Worst case $\pm(4 \times 0.05) = \mathbf{\pm 0.20\ mm}$. Statistical $\pm\sqrt{(4 \times 0.05^2)} = \pm 0.05 \times 2 = \mathbf{\pm 0.10\ mm}$.
14. Any five of: dimensioned drawings and native and neutral CAD files; bill of materials with grades and suppliers; process notes and machine programs or profiles; inspection and test criteria; use, cleaning, and inspection instructions naming the agents tested; spare parts and repair instructions; end-of-life criteria; a status statement of what has and has not been tested or approved.

14.11. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Batch, tolerance and shelf life — break-even, tolerance stack-up, Cp and Cpk, surface roughness and accelerated ageing.
- Tissues, materials and temperature — choosing materials against process temperatures.
- Print tolerance — fits for printed parts.
- Box maker — enclosures.
- OpenSCAD — parametric models.
- STL viewer — checking meshes.
- Battery pack — pack energy.

14.12. Video companion and reading

- **Additive manufacturing in medicine:** the FDA's overview [25] and technical considerations [26] — read them for the **discipline** around the process, which transfers to every other method.
- **Cleaning and materials:** revisit Chapter 13 and the CDC guideline [24] with your own materials list in hand.
- **Design for manufacture:** any standard text on design for manufacture and assembly, and the design guides published by moulding, machining, and sheet-metal service bureaus, which are free and practical.
- **Volunteer engineering in practice:** REMAP [31] and Tikkun Olam Makers [18] — look at the range of methods in their published projects: woodwork, metalwork, electronics, sewing, and

printing side by side.

- **Standards to know by name:** ISO 13485, process validation guidance, and the battery, enclosure, and EMC standards in Appendix D.

14.13. Exit criteria

- ▢ Apply the decision diagram to three needs, with at least one ending without fabrication, and record the reasoning.
- ▢ Produce a method comparison across at least four technology families for one real object, with cleaning agent and quantity as explicit columns.
- ▢ Calculate a break-even quantity from realistic costs and state what the cost comparison omits.
- ▢ Complete a wipe test on at least five materials and record the results with photos.
- ▢ Calculate C_p and C_{pk} for one dimension from ten parts of your own process.
- ▢ Write a handover pack for one made object and have someone else use it without your help.
- ▢ Explain, for each technology family in this chapter, one clinical trap specific to it.

14.14. What this chapter does not establish

- **You are not a manufacturer.** Making a batch does not make you a medical device manufacturer, and placing a product on the market requires a quality system, a regulatory route, and a legal entity that answers for it.
- **A process you have run is not a validated process.** Validation is a documented, statistical activity under a quality system.
- **Material compatibility is a test result, not a datasheet line.** What a cleaning agent does to your material, in your geometry, with your residual stresses, is established by testing with the actual agent — and confirmed by the people who own the cleaning process.
- **Electrical safety is not established by a small voltage alone.** Batteries store real energy, and anything connected to a person is outside your scope.
- **Range of skill does not replace the team.** Orthotists, prosthetists, clinical engineers, dental technicians, and sterile services staff hold the clinical and process judgement for their fields. Your methods are offered **to** them, not substituted for them.

Chapter 15. Measurement, imaging, and the diagnostic laboratory

Week 18 · approximately 6–8 hours

15.1. Why this chapter exists

Most of what a hospital **does** before it treats anybody is **measure**. Imaging, laboratory medicine, physiological measurement, and point-of-care testing together account for an enormous share of clinical activity, capital equipment, software, and training.

A technologist meets this world sooner than expected. Cast ultrasound and needle-guidance phantoms, machined calibration jigs and test fixtures, electronic signal simulators for checking monitors, specimen-handling and transport aids, training datasets and the software that serves them, and physical anatomical models derived from imaging are all realistic first projects — and every one of them requires you to understand what the underlying measurement actually is before you can represent it honestly.

It is also where Chapter 9's four-layer model — phenomenon, sensor, display, interpretation — gets its most thorough workout.

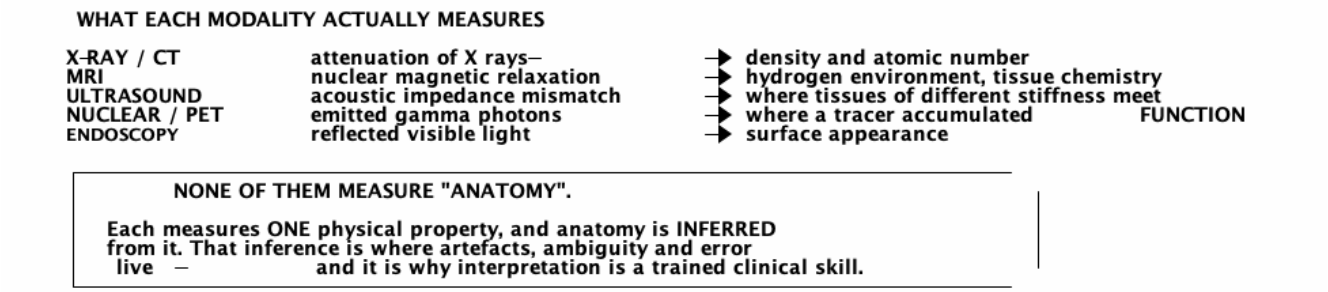
15.2. Medical imaging

This section is the overview: what each modality measures, and the quality and dose trades common to all of them. **How** each machine works — the physics, the reconstruction, the safety environment, and what each is used for — is the subject of Chapter 16.

15.2.1. The modalities

Modality	Physical principle	Good at	Costs and limits
Radiography (plain X-ray)	Differential attenuation of X-rays by tissue density and atomic number; projection onto a detector	Bone, chest, foreign bodies. Fast, cheap, everywhere	**Ionising radiation.** A projection — 3D collapsed into 2D, so structures superimpose
Computed tomography (CT)	Many X-ray projections around the body, reconstructed into cross-sectional slices	Excellent spatial resolution, bone and acute bleeding, fast	**Higher radiation dose.** Soft-tissue contrast poorer than MRI
Magnetic resonance imaging (MRI)	Hydrogen nuclei aligned in a strong magnetic field, perturbed by radiofrequency pulses; relaxation signals are spatially encoded by gradient coils	Superb soft tissue contrast, no ionising radiation, many contrast mechanisms (T1, T2, diffusion, functional)	Slow, expensive, noisy, claustrophobic. **Strong static field — a serious projectile hazard.** Contraindicated with some implants

Modality	Physical principle	Good at	Costs and limits
Ultrasound	Pulses of high-frequency sound; echoes from interfaces between tissues of differing acoustic impedance are timed and reconstructed	Real-time, portable, cheap, no ionising radiation. Excellent for guidance, obstetrics, vessels, heart	Operator-dependent; poor through bone and gas; limited depth at high frequency
Nuclear medicine / PET / SPECT	An injected radiopharmaceutical concentrates by **function**; its emissions are detected externally	Shows physiology and metabolism rather than structure	Ionising, low spatial resolution, requires radiopharmacy
Fluoroscopy	Continuous X-ray imaging in real time	Guiding catheters, implants, swallow studies	Cumulative dose to both patient and operator
Endoscopy	Direct optical visualisation through a flexible or rigid scope	Direct mucosal view; allows biopsy and intervention	Invasive; reprocessing is a major infection-control challenge (Chapter 13)
Optical coherence tomography (OCT)	Interferometry with near-infrared light	Micrometre-resolution cross-sections of retina and vessel wall	Very shallow penetration



15.2.2. Image quality, and the trade you cannot escape

Parameter	What it means
Spatial resolution	Smallest distinguishable detail
Contrast resolution	Smallest distinguishable difference between tissues
Temporal resolution	How fast you can capture a changing scene
Signal-to-noise ratio	How far the signal rises above the noise
Field of view and **slice thickness**	How much is covered, and how thick each slice is
Partial volume effect	One voxel spanning two tissues reports their average — a real source of error in segmentation

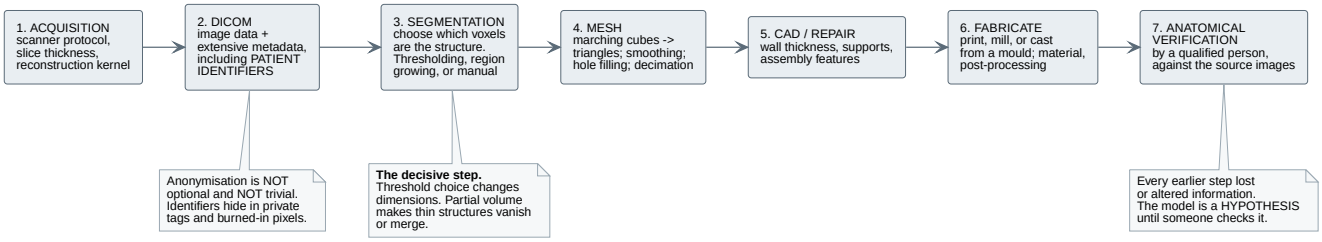
In every modality these trade against one another, and in the ionising ones they trade against **dose**. More signal means more radiation, or more time, or less coverage. The clinical choice of protocol is a deliberate compromise made by people trained to make it.

ALARA — As Low As Reasonably Achievable — is the governing principle for ionising radiation. Every X-ray and CT is a justified decision that the information is worth the dose.

A maker occasionally encounters a request to "just scan it and make a model". Where the scan is of a **person**, that request has a dose, a justification, and a referral pathway behind it. Scanning an **object** on clinical equipment is a different matter and still runs through the department's process.

15.2.3. DICOM, and the path from scan to physical model

The pipeline below ends in a physical object, but the fabrication step is the least interesting part of it. The output geometry may be printed, milled from foam or plastic, turned into a mould for casting silicone or gel, used to cut a template, or never made at all — a well-segmented 3D model on a screen answers many planning and teaching questions on its own. Choosing among those is a fabrication decision (Chapter 14). The steps that decide whether the result is **true** all happen before it.



An attractive reconstruction is not an accurate one.

Each step above is lossy or interpretive. Threshold choice alone can shift a bone surface by a millimetre or more. Partial volume effects make thin walls disappear. Smoothing removes real detail along with noise. Decimation changes geometry. None of this is visible in the finished object — printed, milled, or cast — which looks equally convincing either way.

Do not treat a reconstruction as clinically accurate without validation by someone qualified to do it. And imaging data is patient data from the first file — information governance applies before you open it (Chapter 18).

15.3. Laboratory medicine

Roughly two-thirds of clinical decisions are said to involve laboratory data. The discipline splits into several areas:

Discipline	Measures	Typical methods
Clinical biochemistry	Electrolytes, glucose, renal and liver markers, enzymes, hormones, drug levels	Spectrophotometry, ion-selective electrodes, immunoassay, mass spectrometry
Haematology	Cell counts, cell morphology, clotting	Impedance and optical flow cytometry, microscopy, coagulometry
Microbiology	Identification of organisms and their antibiotic susceptibility	Culture, Gram stain, MALDI-TOF mass spectrometry, molecular methods
Histopathology and cytology	Tissue and cell structure and disease	Fixation, embedding, sectioning, staining, microscopy, immunohistochemistry
Immunology	Antibodies, antigens, immune function	ELISA and related immunoassays, flow cytometry
Molecular diagnostics / genomics	Nucleic acid sequences	PCR, qPCR, sequencing
Transfusion science	Blood groups and compatibility	Agglutination, antibody screening

15.3.1. The measurement principles worth knowing

Principle	How it works
Spectrophotometry	Light at a chosen wavelength passes through a sample; absorbance relates to concentration via the **Beer–Lambert** law. The workhorse of clinical chemistry — and the same principle as the pulse oximeter in Chapter 9
Ion-selective electrode	A membrane selectively permeable to one ion develops a potential that varies logarithmically with its activity. How sodium, potassium, chloride, and pH are measured
Amperometry	Current from an electrochemical reaction at an electrode is proportional to analyte concentration. The glucose meter in Chapter 3
Immunoassay	An antibody binds its specific antigen; the bound complex is detected by an enzyme, fluorescent, or gold-particle label. ELISA in the lab; **lateral flow** on a strip — the pregnancy test and the COVID rapid test are the same architecture
Flow cytometry	Cells pass single-file through a laser; scattered light and fluorescence classify and count them. Also the basis of the automated full blood count
PCR	Thermal cycling with a DNA polymerase doubles a target sequence each cycle; **qPCR** monitors amplification in real time. Exponential amplification is what makes tiny quantities detectable — and also what makes contamination catastrophic

Principle	How it works
Mass spectrometry	Ionise, separate by mass-to-charge ratio, detect. Extremely specific. MALDI-TOF identifies bacteria in minutes; LC-MS/MS is the reference method for many analytes
Microscopy	Still central to haematology, histopathology, and microbiology. Increasingly digital

15.3.2. Point-of-care testing

Testing performed near the patient rather than in a central laboratory: glucose meters, blood gas analysers, lateral flow tests, portable coagulation monitors, haemoglobin meters.

The trade is consistent and worth understanding: **speed and accessibility against analytical performance, quality control, and traceability**. A central lab has calibration regimes, external quality assurance, trained staff, and an audit trail. A device on a ward has none of those by default, which is why point-of-care programmes have their own governance.

15.4. Engineering bridge: what makes a measurement trustworthy

This vocabulary is the heart of the chapter, and it applies to anything you ever build that produces a number.

Term	Meaning
Accuracy	Closeness to the true value. Freedom from systematic error
Precision	Closeness of repeated measurements to each other. Freedom from random error
Trueness and **bias**	Trueness is agreement of the **average** with the true value; bias is the difference
Repeatability	Precision under the same conditions, same operator, short interval
Reproducibility	Precision across different operators, instruments, days, sites
Resolution	Smallest change the instrument can display
Sensitivity (analytical)	Smallest amount reliably detectable — the limit of detection
Specificity (analytical)	Freedom from response to other substances. The opposite of **interference**
Linearity and range	Over what span the output tracks the input proportionally
Calibration	Establishing the relationship between instrument output and known standards
Traceability	An unbroken chain of calibrations linking your result to a reference standard

Term	Meaning
Drift	Slow change in output over time with no change in input
Uncertainty	The quantified doubt about a result. **A result without it is incomplete**

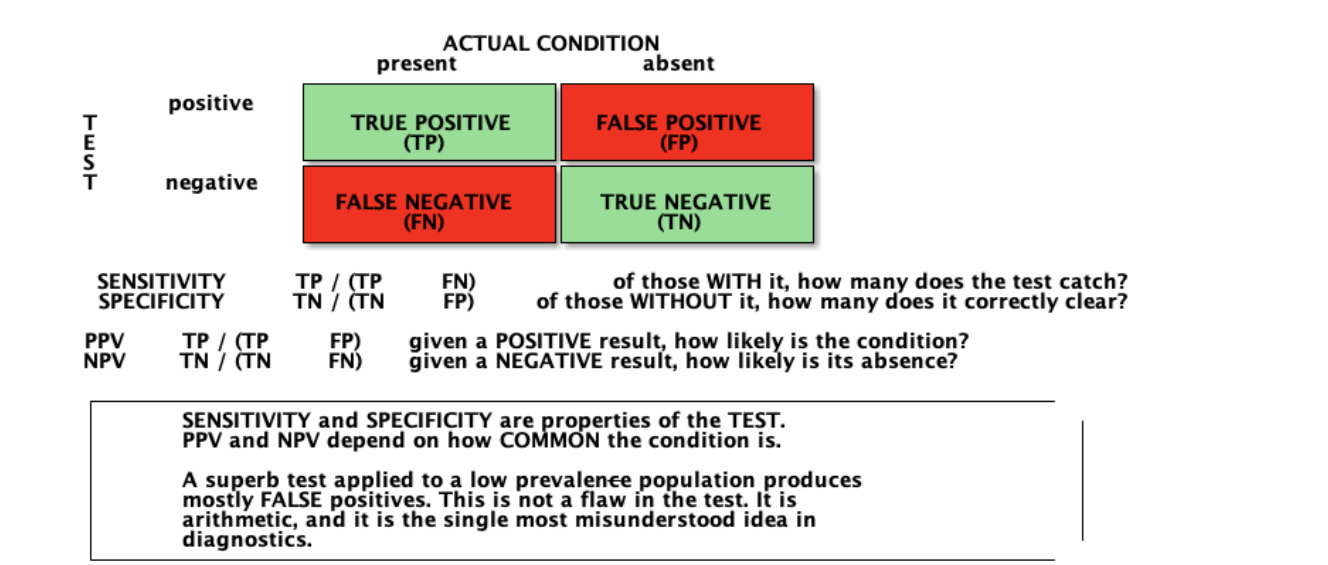


Accuracy and precision are independent. An instrument can be precise and consistently wrong — that is bias, and it is the more dangerous failure because the repeatability looks reassuring.

This is exactly the pulse oximeter problem in Chapter 9: a device can give highly repeatable readings that are systematically biased in one population. Precision concealed the bias.

15.4.1. Diagnostic performance: the other sensitivity and specificity

Confusingly, **sensitivity** and **specificity** mean something different in diagnostics than in instrumentation. Both meanings are in common use.



Work the arithmetic once, because it is genuinely counter-intuitive.

A test with **99% sensitivity** and **99% specificity**, applied to 10,000 people in a population where the condition affects **1 in 1,000**:

	Condition present	Condition absent	Total
Test positive	10	100	110
Test negative	0	9,890	9,890
Total	10	9,990	10,000

Of 110 positive results, only 10 are correct. **The positive predictive value is about 9%** — a positive result from a 99%-accurate test means the condition is probably **absent**.



This is why screening programmes are designed by epidemiologists rather than enthusiasts, why "our model is 99% accurate" is close to meaningless without prevalence, and why any software you write that classifies anything clinical must report more than an accuracy figure.

It is also the reason confirmatory testing exists. A positive screen is a **reason to look further**, not a diagnosis.

15.4.2. ROC curves and the threshold decision

Most tests produce a continuous value, and someone must choose the cut-off. Move it toward higher sensitivity and you catch more cases at the cost of more false positives; move it the other way and the reverse.

The **receiver operating characteristic** curve plots sensitivity against (1 − specificity) across all possible thresholds, and the area under it summarises discrimination independent of any particular cut-off.

Where the threshold is actually set is a clinical and ethical decision, not a mathematical one. It depends on the cost of a missed case versus the cost of a false alarm, and those costs are not commensurable by any formula you can write.

15.5. Devices, software, and training technology

15.5.1. Clinical devices

Device	Does	Engineering note
Digital radiography system	Flat-panel detector, generator, positioning equipment	Detector quantum efficiency; dose-area product monitoring
CT scanner	Rotating source–detector, helical acquisition, reconstruction	Iterative and deep-learning reconstruction now reduce dose substantially
MRI scanner	Superconducting magnet, gradient and RF coils	Cryogenics; acoustic noise; **ferromagnetic projectile hazard**; implant compatibility
Ultrasound system	Piezoelectric transducer array, beamforming, Doppler	Handheld and phone-connected probes have changed access dramatically
Automated lab analyser	Sample handling, pipetting, reaction, detection, often thousands of tests per hour	Robotics, fluidics, optics, and scheduling software
Haematology analyser	Full blood count by impedance and optical methods	High throughput, automated flagging for manual review

Device	Does	Engineering note
PCR thermocycler / qPCR instrument	Thermal cycling with optical detection	Thermal uniformity and ramp rate are the core engineering problem
Mass spectrometer	Analyte identification by mass-to-charge ratio	MALDI-TOF for microbiology; LC-MS/MS for reference chemistry
Digital slide scanner	Whole-slide imaging at microscope resolution	Enables remote reporting and computational pathology. Files are gigabytes
Point-of-care analysers	Blood gas, glucose, coagulation, lateral flow readers	Cartridge-based microfluidics; the design challenge is self-contained reagent handling

15.5.2. Software

Software	Does	Note for a maker
PACS — picture archiving and communication system	Stores, retrieves, and distributes medical images	The hospital's image backbone. DICOM is its interchange format
RIS / LIS	Radiology and laboratory information systems: ordering, workflow, reporting	Integration via HL7 and increasingly FHIR
3D Slicer	Free, open-source medical image computing: segmentation, registration, 3D modelling	**The standard free route from DICOM to a surface mesh** (STL or OBJ) for any fabrication route, or to a 3D view that needs no fabrication at all. Well documented, large community
ITK-SNAP, 3D Slicer, InVesalius, Horos	Segmentation and visualisation	All free; InVesalius and Horos are approachable starting points
OsiriX, Materialise Mimics	Clinical-grade image processing and segmentation	Commercial; Mimics Medical is regulatory-cleared for some planning uses
ITK / VTK / SimpleITK / pydicom	Libraries for image processing, visualisation, and DICOM handling	**The toolkit if you write code.** `pydicom` + `SimpleITK` in Python gets you a long way
AI triage and detection tools	Flag suspected findings — intracranial haemorrhage, pulmonary nodules, fractures	A fast-growing regulated category. Performance claims need the prevalence caveat above

Software	Does	Note for a maker
Computational pathology	Analysis of whole-slide images	Enormous data volumes; active research and early clinical deployment
OHIF Viewer, dcm4che, Orthanc	Open-source DICOM viewer, toolkit, and server	**Lets you build a complete imaging pipeline on your own machine, with public datasets**

15.5.3. Training and simulation

Training technology	Teaches	Form
Ultrasound phantoms	Probe handling, orientation, needle guidance	Gelatin, agar, ballistic gel, or commercial polymer, with embedded targets. **A classic maker project**
Vascular access trainers	Cannulation, central line placement	Simulated vessels in tissue-mimicking material, often with artificial blood
Radiology teaching files and public datasets	Interpretation, pattern recognition	Radiopaedia, TCIA, and similar collections of anonymised cases
Virtual microscopy	Histology and pathology without physical slides	Whole-slide images served in a browser; widely used in medical schools
CT and MRI simulators	Protocol selection, artefact recognition, dose awareness	Software; some vendor-provided
Phlebotomy and specimen-handling trainers	Venepuncture, labelling, correct tube order, transport	Manikin arms plus consumables; labelling practice is a known weak point
Laboratory skills simulators	Pipetting, dilution, assay setup	Virtual labs plus physical pipetting practice rigs
Anatomical models derived from imaging	Spatial relationships, surgical planning, patient explanation	Printed, milled, or cast from a mould made from the segmented geometry; often combined with silicone soft tissue. Widely requested — and the one with the clearest accuracy obligation
Instrument and analyser test fixtures	Checking that equipment reads what it should	Machined or sheet-metal holders, electronic signal and pressure simulators, reference samples. **Clinical engineering and laboratories value these, and they are rarely glamorous enough for anyone else to make**

Training technology	Teaches	Form
Radiation safety trainers	Scatter, inverse square, shielding, dosimetry	Simulation software plus survey-meter practice with check sources



The ultrasound phantom is the best first project in this chapter.

It is cheap, it uses mould-making and casting skills a workshop maker already has, it needs no patient contact, and educators want them because commercial phantoms are expensive and degrade. The mould can be anything you can make watertight — a fabricated acrylic box, a turned or milled former, a sheet-metal tray, a printed shell — and the material choice (gelatin, agar, PVA cryogel, silicone) is where the real decisions are. Gelatin or agar with scatterers and embedded targets gets you a usable needle-guidance phantom for the cost of a few pounds of materials.

The honest limitations are also clear and easy to state: a simple phantom's acoustic properties — speed of sound, attenuation, backscatter — are not those of tissue; it has a short shelf life; it grows mould. None of which stops it teaching probe handling and needle alignment, if you say so.

Agree the target with the educator first: **which** skill, at **which** depth, with **what** needle. The phantom follows from the objective, exactly as Chapter 21 requires.

15.6. Worked example: the request for a model from a scan

A surgeon asks for a physical model of a patient's pelvis from a CT scan, to plan an operation and to explain it to the patient.

What you must establish before touching a file:

Question	Why
Is there **information governance approval**?	This is patient data. Approval precedes access, not the other way round (Chapter 18)
Who **anonymises** it, and how is that verified?	Identifiers live in private DICOM tags and sometimes burned into pixel data. Removing the obvious fields is not sufficient
What is the model **for**?	Planning, rehearsal, patient explanation, and implant sizing are different requirements with different accuracy needs
Does it need to be **physical** at all?	A rotatable on-screen model or an existing anatomical model may answer the question. If a physical one is needed, the purpose decides the method: a rigid printed or milled bone, a cast silicone model with soft tissue, or a cutting template

Question	Why
What **accuracy** is needed, and who says so?	"Looks right" is not a specification. If sizing decisions depend on it, the tolerance must be stated
Who performs **anatomical verification**?	A qualified person, against the source images. Not you
Does it enter the **sterile field**?	If so this is a different project entirely — materials, sterilisation, and regulatory status all change
What is the **scan protocol**?	Slice thickness and reconstruction kernel set the achievable fidelity. A 5 mm slice cannot support a 0.5 mm claim
Who **owns** it afterwards, and how is it disposed of?	It is derived from patient data

What you then do, and record: the DICOM series used; the segmentation method and every threshold; smoothing, hole-filling, and decimation settings; scale verification against a known dimension; the fabrication method, material, machine, and settings (print orientation, toolpath and cutter, or mould and casting compound, as applicable); measured dimensional checks against the source images; and the verification sign-off.

What you write on the model and in the handover:

> "Anatomical model derived from CT series **[identifier]**, slice thickness **[x]** mm. Segmentation by **[method]** at threshold **[value]**. Smoothed and decimated — fine detail and thin structures may be lost or merged. Dimensional check against source images: **[result]**. Anatomically verified by **[name, role, date]**. **Not sterile. Not for intraoperative use.** Scale 1:1 verified by **[method]**."

That paragraph is the deliverable as much as the object is. Without it the model is a convincing-looking artefact whose provenance nobody can reconstruct — which is exactly the failure mode Chapter 21 calls a misleading cue, in physical form.

15.7. Workshop exercise

1. **Build the pipeline.** Install 3D Slicer. Download a **public, already anonymised** dataset — the Cancer Imaging Archive and 3D Slicer's own sample data are both suitable. Segment a bone and export a surface mesh. Then make it physical by whichever route your workshop supports best — print it, mill it from foam or tooling board, or make a two-part mould and cast it — at a stated scale. Record why you chose that route.
2. **Document what you lost.** For that model, write down every threshold, smoothing, and decimation setting, and measure one dimension on the model against the same dimension in the source images. Report the difference.
3. **Threshold sensitivity.** Re-segment the same structure at three different thresholds. Measure

the same dimension on each. The spread is your answer to "how accurate is a model made from a scan?"

4. **Make an ultrasound phantom.** Gelatin with a scatterer (flour, psyllium husk, or cornflour) and embedded targets — a balloon of water, a length of tubing, an olive. If you can borrow a probe, practise needle alignment. Write its limitations.
5. **Do the prevalence arithmetic.** Take a test with 95% sensitivity and 90% specificity. Compute PPV at 1% prevalence and at 40% prevalence. Write one sentence about what that means for a screening app.
6. **Read a DICOM file.** With `pydicom`, open a public DICOM file and list its tags. Find the pixel data, the slice thickness, and at least three tags that **would** carry identifying information in a clinical file.

15.8. Questions

1. What physical property does each of X-ray, MRI, and ultrasound actually measure?
2. Why does none of the imaging modalities "measure anatomy"?
3. What is the partial volume effect and why does it matter for segmentation?
4. Name the three most information-lossy steps between a CT scan and a physical model.
5. Why is an attractive reconstruction not necessarily an accurate one?
6. Distinguish accuracy from precision, and say which failure is more dangerous.
7. What is the difference between repeatability and reproducibility?
8. A test is 99% sensitive and 99% specific. Prevalence is 1 in 1,000. What is the approximate PPV and what does it mean?
9. Why are sensitivity and specificity properties of the test but PPV and NPV not?
10. Why does qPCR's exponential amplification make contamination so serious?
11. What is the trade made by point-of-care testing?
12. Who sets the threshold on a continuous diagnostic test, and on what basis?

15.9. Answers and reasoning

1. **X-ray/CT**: attenuation of X-rays, which depends on density and atomic number. **MRI**: nuclear magnetic relaxation of hydrogen nuclei, which depends on their chemical environment. **Ultrasound**: reflection of sound at interfaces between tissues of differing **acoustic impedance**.
2. Because each measures **one physical property**, and anatomy is **inferred** from the spatial distribution of that property. The inference step is where artefacts, ambiguity, and error live, and it is why interpretation is a trained skill rather than a reading-off.
3. One voxel spanning two tissues reports their **average** value. Thin structures can therefore disappear below a threshold, and adjacent structures can merge. It directly limits the smallest feature a segmentation can faithfully recover.
4. Any three of: **segmentation threshold choice** (the most decisive); **smoothing**, which removes

real detail with noise; **decimation**, which changes geometry; the original **slice thickness and reconstruction kernel**, which cap what was ever available; and the **fabrication tolerance and post-processing** of whichever method makes the object.

5. Because every step from acquisition to fabrication is lossy or interpretive, and none of that loss is visible in the finished object. A model built on a wrong threshold looks exactly as convincing as one built on a right threshold.
6. **Accuracy** is closeness to the true value — freedom from systematic error. **Precision** is closeness of repeated measurements to one another — freedom from random error. **Precise but inaccurate is more dangerous**, because the consistency looks like reliability while the bias goes undetected.
7. **Repeatability** is precision under the same conditions — same operator, instrument, and short interval. **Reproducibility** is precision across different operators, instruments, days, and sites. Reproducibility is the harder and more clinically relevant one.
8. Of 10,000 people, 10 have the condition and the test finds all 10; of the 9,990 without it, 1% — about 100 — test positive anyway. So 110 positives, 10 of them true: **PPV ≈ 9%**. A positive result from this excellent test still means the condition is probably absent, and a confirmatory test is required.
9. Because sensitivity and specificity are conditional on the **true state** — they describe how the test behaves given disease or given health, which does not depend on how common disease is. PPV and NPV are conditional on the **test result**, so they depend on the mix of diseased and healthy people being tested, which is prevalence.
10. Because amplification is exponential: a few stray template molecules carried over from a previous reaction or from the environment are amplified with the same efficiency as the real sample, producing a convincing false positive. It is why PCR facilities use physically separated pre- and post-amplification areas and unidirectional workflow.
11. **Speed and accessibility** against **analytical performance, quality control, and traceability**. A central lab has calibration regimes, external quality assurance, trained operators, and an audit trail; a device at the bedside has none of those by default, which is why point-of-care testing needs its own governance programme.
12. A clinician, with epidemiological input — on the basis of the relative **costs** of a missed case and a false alarm in that specific clinical context. The ROC curve describes the available trade-offs; it does not choose between them, and no formula can, because the two costs are not commensurable.

15.10. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Test accuracy and risk — sensitivity, specificity, PPV against prevalence, and ROC thresholds.
- CT and MRI lab — how a segmentation threshold moves a measured dimension.
- DICOM viewer and de-identification check — reading image headers, Hounsfield units, and finding identifiers.

- Statistics — measurement statistics.

15.11. Video companion and reading

- **Imaging physics:** Radiopaedia [38] is free, well written, and the standard first reference for both physics and interpretation. Start with the modality overview articles, not the case library.
- **Hands-on:** the 3D Slicer documentation and tutorials [34]. Working through the segmentation tutorial with sample data is the single most useful afternoon in this chapter.
- **Data:** The Cancer Imaging Archive [35] provides large, properly anonymised, openly available imaging datasets for exactly this kind of practice.
- **Diagnostics:** for the prevalence arithmetic, any introductory clinical epidemiology text; the key skill is doing the 2×2 table yourself until it stops being surprising.
- **Code:** `pydicom`, `SimpleITK`, and `VTK` documentation if you want to build a pipeline rather than drive a GUI.

15.12. Exit criteria

- ☐ State what physical property each major imaging modality measures, and why none of them measures anatomy directly.
- ☐ Explain the partial volume effect and its consequence for segmentation.
- ☐ Complete the DICOM → segmentation → mesh → physical model pipeline on a public dataset, by a fabrication route you chose and justified, with every setting recorded.
- ☐ Quantify how much a dimension changes across three segmentation thresholds.
- ☐ Write a limitations statement for a model derived from imaging.
- ☐ Define accuracy, precision, trueness, bias, repeatability, reproducibility, and uncertainty, and explain why precise-but-biased is the dangerous case.
- ☐ Compute PPV from sensitivity, specificity, and prevalence, and explain the result in plain language.
- ☐ Explain why a positive screening result is not a diagnosis.
- ☐ Build one ultrasound phantom and state at least four of its limitations.
- ☐ Name the approvals required before touching clinical imaging data.

15.13. What this chapter does not establish

- **You cannot interpret any image or any result.** You know what the measurement is. What it means is a trained clinical skill, and for imaging a long and specialised one.
- **You cannot handle clinical imaging data without approval.** Information governance comes first, always, and anonymisation is a verified process, not a best effort.
- **You cannot certify a model's anatomical accuracy.** A qualified person checks it against the source images. Every time.
- **Nothing here is for the sterile field or intraoperative use.** Those raise materials, sterilisation,

and regulatory questions this book does not cover.

- **A phantom is not tissue.** Its acoustic properties differ, it degrades, and it teaches handling rather than pathology. Say so.
- **A reported accuracy figure is not performance.** Without prevalence, a classification accuracy number says very little, and any software you write should report more than it.

Chapter 16. Diagnostic devices: how they work and what they are for

Week 19 · approximately 6–8 hours

16.1. Why this chapter exists

Chapter 15 gave you the **framework**: every modality measures one physical property and anatomy is inferred from it; image quality trades against dose and time; a test's usefulness depends on prevalence as much as on its accuracy. It summarised the imaging modalities in one table and moved on.

This chapter opens the boxes. For each of the common diagnostic devices — X-ray, CT, MRI, ultrasound, nuclear medicine and PET, endoscopy, and the physiological tests of heart, lungs, sleep, and nerves — it answers three questions:

1. **What physical interaction does it exploit?** What goes into the body, and what comes out?
2. **How is that turned into an image or a number?** Detector, reconstruction, display.
3. **What is it good for, what is it bad at, and what does it cost the patient?**

You need this for practical reasons. A phantom must reproduce the property the device measures — X-ray attenuation for CT, relaxation times for MRI, acoustic impedance for ultrasound — or it teaches nothing. Software that handles these images must respect what the pixel values mean. Anything you take into a scanner room must be compatible with the physics in it. And you cannot hold a useful conversation with a radiographer, sonographer, physicist, or physiologist unless you know what their machine does.

There is a quieter reason too. These machines are some of the most elegant engineering in existence — a hospital basement full of superconducting magnets, particle accelerators, phased arrays, and inverse problems. You will enjoy them more for understanding them.

16.2. The map: which question, which device

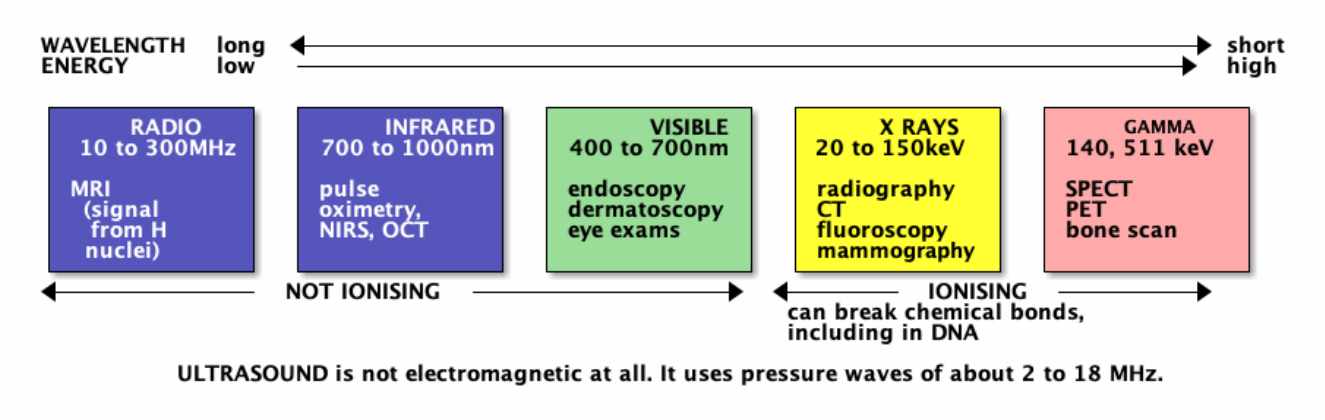
Diagnosis is a sequence of questions, and each family of devices answers a different kind.

The question	Device family	Examples
What does it look like inside? (structure)	Anatomical imaging	X-ray, CT, MRI, ultrasound
What is it doing? (function, metabolism)	Functional imaging; physiological measurement	PET, SPECT, functional MRI, Doppler, ECG, spirometry, EEG

The question	Device family	Examples
What can be seen directly? (surface)	Endoscopy and optical examination	Gastroscopy, colonoscopy, bronchoscopy, dermatoscopy, ophthalmoscopy (Chapter 8)
What is it made of, chemically?	Laboratory medicine	Blood tests, urine tests, microbiology (Chapter 15)
What are the cells? (tissue diagnosis)	Biopsy and histopathology	Needle biopsy, excision biopsy, cytology

Most diagnoses use several of these in sequence: a history and examination, then a cheap, widely available test, then a more specific one only if the first leaves a question open. Keep that sequencing in mind; it explains why the expensive machines are not used first.

16.3. Radiation, ionising and not



The single most important divide is whether a modality uses **ionising** radiation — photons energetic enough to knock electrons out of atoms and break chemical bonds, including those in DNA. X-ray imaging, CT, fluoroscopy, and nuclear medicine do. MRI, ultrasound, and optical methods do not.

Ionising dose is expressed as **effective dose** in millisieverts (mSv), a weighted estimate of whole-body risk. It is a population-level quantity for comparing procedures, not a measurement of harm to an individual. Some orientation figures:

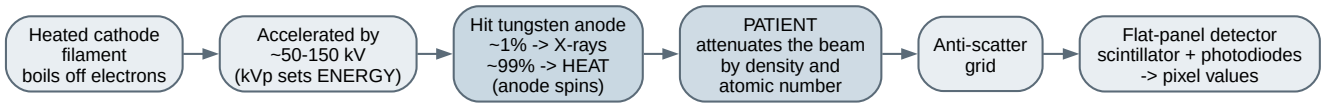
Source	Typical effective dose	Equivalent natural background
Dental bitewing X-ray	~0.005 mSv	About a day
Chest X-ray (single view)	~0.02–0.1 mSv	A few days to ~10 days
Long-haul flight	~0.03–0.08 mSv	About a week
Bone density (DEXA)	~0.001 mSv	Hours
Mammogram (screening)	~0.4 mSv	~7 weeks
CT head	~2 mSv	~8 months

Source	Typical effective dose	Equivalent natural background
CT chest	~6 mSv	~2 years
CT abdomen and pelvis	~8–10 mSv	~3 years
Bone scan (nuclear medicine)	~6 mSv	~2 years
PET-CT	~10–25 mSv	~3–8 years
Natural background, per year	~2–3 mSv	—

Figures are typical published adult values for comparison; actual doses vary widely with protocol, equipment, and body size. The risk from a single diagnostic dose is small and statistical — a slight increase in lifetime cancer probability, larger for children, whose tissues are more sensitive and who have longer to live. That is why **justification** and **ALARA** (Chapter 15) govern every exposure, and why paediatric protocols are separately optimised.

16.4. Projection radiography: the plain X-ray

16.4.1. How X-rays are made



An X-ray tube is a vacuum diode. A heated filament releases electrons; a high voltage accelerates them across the tube into a tungsten target. As they decelerate in the target they emit X-rays — mostly **bremsstrahlung** ("braking radiation"), a continuous spectrum up to the peak voltage. About 1% of the electrical energy becomes X-rays; the rest becomes heat, which is why the anode is a spinning disc in an oil-cooled housing, and why the tube's thermal capacity limits how fast a CT scanner can work.

The operator sets two things:

- **kVp** (peak kilovoltage) sets the **energy** of the photons — their penetrating power. Low kVp (25–35 kV for mammography) gives high contrast between soft tissues; higher kVp (100–150 kV for a chest) penetrates more and reduces dose for a given image.
- **mAs** (tube current × time) sets the **number** of photons — the image's noise level and, proportionally, the dose.

16.4.2. How the body makes the image

As the beam passes through tissue it is attenuated exponentially — the Beer–Lambert law of Chapter 15 again:

$$I = I_0 e^{-\mu x}$$

where μ is the linear attenuation coefficient of the material and x its thickness. Two physical interactions dominate:

- **Photoelectric absorption** — the photon is absorbed outright. Its probability rises roughly with the **cube of atomic number** and falls with the cube of photon energy. Calcium in bone (Z = 20) absorbs far more than the carbon, oxygen, and hydrogen of soft tissue. This is where **contrast** comes from.
- **Compton scatter** — the photon bounces off an electron and changes direction. It depends mainly on **density**, not atomic number, and scattered photons reaching the detector add a fog that reduces contrast. An **anti-scatter grid** of thin lead strips aligned with the beam absorbs most of them.

The result is a **shadow**: a projection of everything between the tube and the detector, collapsed into two dimensions. A radiograph shows five basic densities, from dark to light: **air, fat, soft tissue and fluid, bone, metal**. Everything a radiologist reads is boundaries between those five.

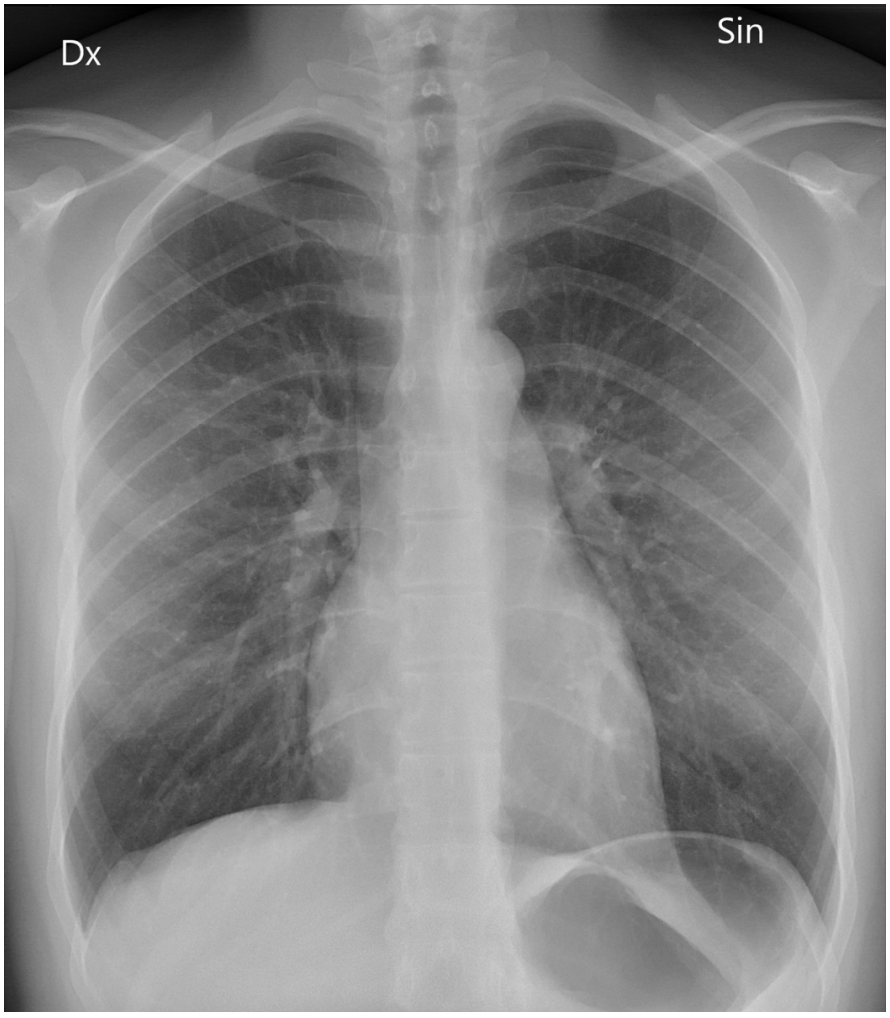


Figure 24. A normal posteroanterior chest radiograph. Air in the lungs is dark; the heart, diaphragm, and soft tissues are grey; the ribs, spine, and clavicles are white. Image: Mikael Häggström, CC0. Credits: Appendix H

Contrast agents exploit the photoelectric effect deliberately. Iodine (Z = 53) and barium (Z = 56) have absorption edges — sharp jumps in photoelectric absorption — at about 33 and 37 keV, right in the diagnostic energy range. Iodine injected into a vessel makes blood visible; barium swallowed coats the gut.

16.4.3. The radiographic family

Variant	How it differs	Used for
Plain radiography	Single projection, typically two views at right angles	Fractures, chest (pneumonia, fluid, collapsed lung), foreign bodies, bowel obstruction, dental
Mammography	Low kVp, molybdenum, rhodium or tungsten targets, firm compression to spread and thin the breast. **Tomosynthesis** takes several angled views and reconstructs thin slices	Breast cancer screening and diagnosis; detecting tiny calcifications
DEXA (dual-energy X-ray absorptiometry)	Two photon energies; because bone and soft tissue attenuate them differently, the two images can be solved for bone mineral content alone	Bone density; osteoporosis, reported as a **T-score** relative to young adults. Tiny dose
Fluoroscopy	Continuous low-dose X-ray at video rate	Swallow studies (Chapter 8), joint injections, placing lines, catheters, and implants
Angiography and digital subtraction angiography (DSA)	Iodine injected into vessels during fluoroscopy; a pre-contrast "mask" image is subtracted, leaving only the vessels	Coronary angiography, stroke clot retrieval, vascular stenting
Dental panoramic and cone-beam CT	Rotating source around the head; cone-beam CT gives 3D at modest dose	Dental implants, orthodontics, maxillofacial surgery — feeding a mature scan-to-manufacture chain of milled crowns, printed guides, and cast and machined prostheses

The fundamental limitation of all projection imaging is **superimposition**: a lesion in front of a rib, behind the heart, or inside the spine's shadow can be invisible. The solution to that is the next machine.

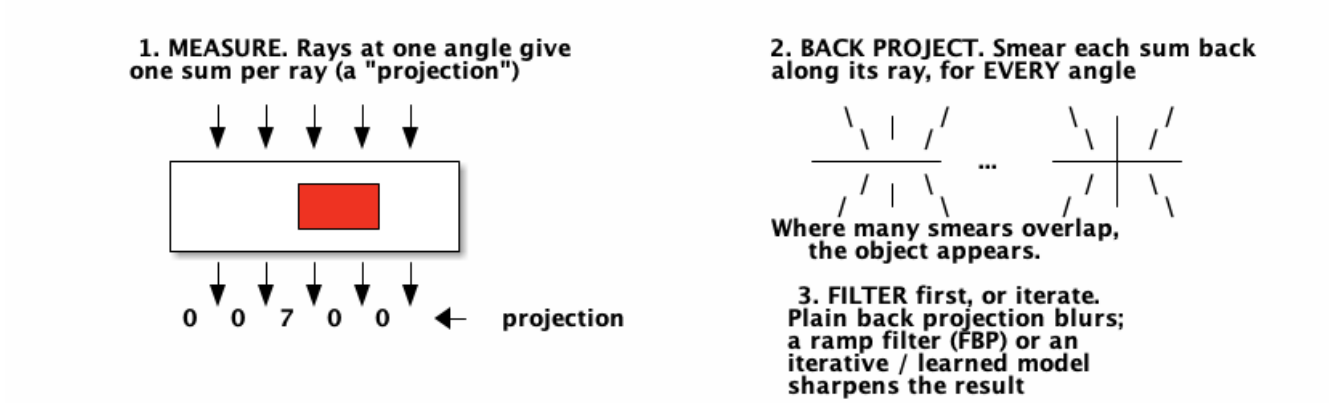
16.5. Computed tomography (CT)

16.5.1. Many shadows, solved backwards

CT takes the X-ray idea and solves its central problem. An X-ray tube and an opposed arc of detectors — 64 to 320 rows of them in a modern scanner — rotate around the patient, typically two to four times per second, while the table moves through the gantry. The helical path collects hundreds of projections from every angle.

Each detector reading is a **line integral**: the total attenuation along one ray through the body. The reconstruction problem is to find the attenuation of every small volume element (**voxel**) from hundreds of thousands of such sums. It is an **inverse problem**, and its mathematics — the Radon

transform, published in 1917, decades before anyone could use it — is the foundation of all tomography.



The classical method is **filtered back-projection** (FBP): high-pass filter each projection, then smear it back across the image along the path it came from. Modern scanners use **iterative reconstruction** — guess an image, simulate the projections it would produce, compare with what was measured, correct, repeat — and increasingly **deep-learning reconstruction**. Both allow lower dose for the same image quality, because they model the noise rather than just inverting it.

16.5.2. Hounsfield units

A CT pixel is not a brightness. It is a **calibrated measurement** of attenuation, in **Hounsfield units** (HU), defined so that water is 0 and air is −1000:

$$\mathbf{HU = 1000 \times (\mu - \mu_{water}) / (\mu_{water} - \mu_{air})}$$

Tissue	Approximate HU
Air	−1000
Lung	−900 to −500
Fat	−100 to −50
Water	0
Soft tissue, muscle, unclotted blood	+20 to +60
Fresh clotted blood	+60 to +90
Iodine-enhanced vessel	+150 to +400 or more
Cancellous bone	+300 to +800
Cortical bone	+1000 to +2000
Metal	above +3000 (often saturates)

A CT scanner produces 12-bit values spanning thousands of HU; a screen shows perhaps 256 grey levels and the eye distinguishes fewer. So every CT image is displayed through a **window**: a centre (level) and a width that map a chosen band of HU onto black-to-white. Everything below the band is black, everything above is white.

Window	Level (HU)	Width (HU)	Shows
Brain	~40	~80	Grey versus white matter; fresh blood stands out
Soft tissue	~40	~400	Abdominal organs, muscle
Lung	~−600	~1500	Airways and lung tissue
Bone	~400	~2000	Cortex and trabeculae

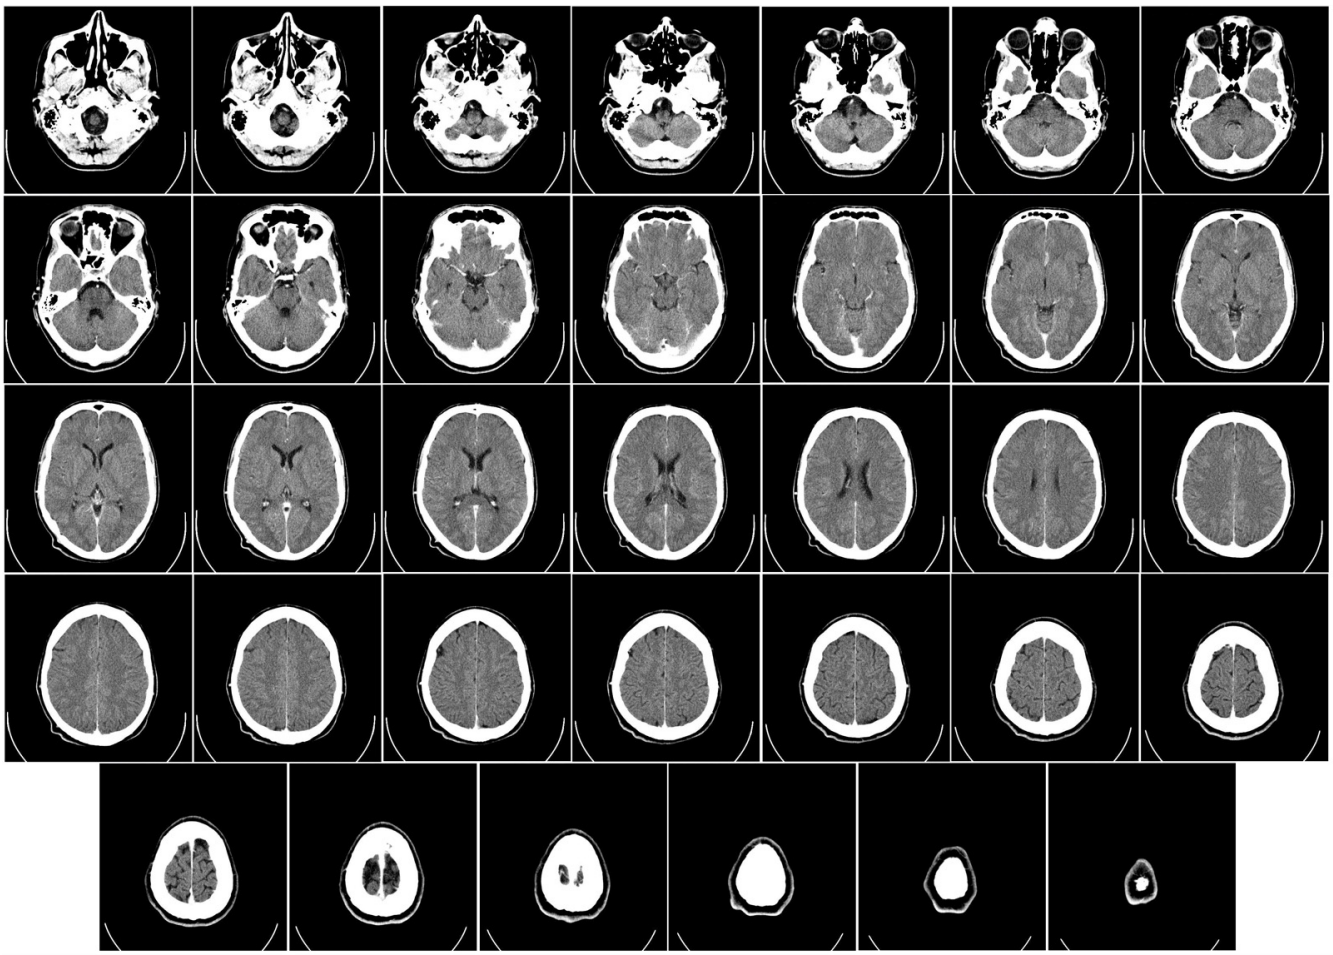


Figure 25. A contrast-enhanced CT of the brain, from skull base (top left) to vertex, displayed in a brain window. Bone saturates to white; air is black. Image: Department of Radiology, Uppsala University Hospital, CC0. Credits: Appendix H

For anyone writing imaging software: the window is a display transform, not the data.

Segmentation thresholds, mesh exports for fabrication, and measurements must operate on the **stored HU values** (after applying the DICOM rescale slope and intercept), not on the windowed screen image. A bone model thresholded at a screen grey level rather than at a HU value is a model of the display settings, not of the bone — the same trap as Chapter 15's scan-to-model pipeline.

16.5.3. What CT is for, and what it costs

CT is fast — a whole chest in a few seconds — available around the clock, and excellent for bone,

lung, acute bleeding, and anything where speed matters. It is the workhorse of emergency medicine:

- **Suspected stroke:** a non-contrast head CT in minutes excludes bleeding before clot-dissolving drugs are given; CT angiography finds a blocked vessel for clot retrieval.
- **Major trauma:** whole-body CT finds injuries in one pass.
- **Pulmonary embolism:** CT pulmonary angiography shows clots in the lung arteries.
- **Abdominal pain, cancer staging, surgical planning:** contrast-enhanced CT, often timed to catch arterial or venous phases.
- **Screening:** low-dose CT for lung cancer in high-risk smokers; CT colonography.

Its costs are **dose** — the highest routine diagnostic dose in medicine — and the risks of **iodinated contrast**: allergic-like reactions (usually mild) and a possible strain on already impaired kidneys. Soft-tissue contrast is poorer than MRI's.

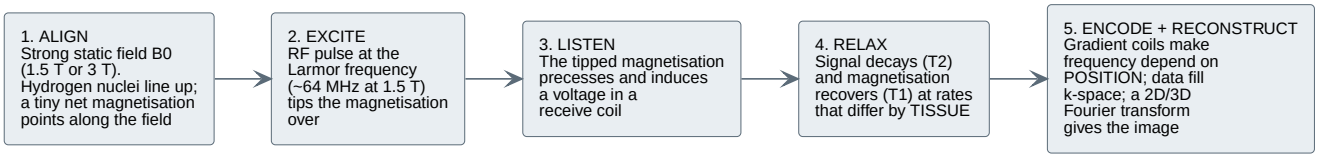
Artefacts you should recognise: **streaks** from metal implants (the beam is almost fully absorbed, so the reconstruction lacks data), **beam hardening** (low-energy photons are filtered out passing through bone, so the beam's average energy rises and attenuation is under-estimated), **motion**, and **partial volume** (Chapter 15).

Two newer developments: **dual-energy CT** acquires at two energies and separates materials — iodine from calcium, for instance — and **photon-counting CT**, clinically available since the early 2020s, measures the energy of each individual photon rather than integrating total charge, giving sharper images at lower dose.

16.6. Magnetic resonance imaging (MRI)

MRI is the device engineers most often want explained, and it repays the effort. It uses no ionising radiation at all — only a strong magnetic field and radio waves — and its contrast comes from the **chemical environment** of water.

16.6.1. The physics in five steps



Step 1 — align. Hydrogen nuclei (protons), abundant in water and fat, behave like tiny spinning magnets. In a strong field $\mathbf{B_0}$ — clinically 1.5 or 3 tesla, some 30,000 to 60,000 times the Earth's field — slightly more of them align with the field than against it. The excess is only a few parts per million, but there are so many protons in a body that the net **magnetisation** is measurable. That tiny excess is also why MRI is signal-starved, and why stronger magnets give better images.

Step 2 — excite. The protons **precess** about the field like tilted spinning tops, at the **Larmor frequency**:

$\mathbf{f} = \bar{\gamma} \times \mathbf{B_0}$, with $\bar{\gamma}$ = 42.58 MHz per tesla for hydrogen.

At 1.5 T that is about 64 MHz; at 3 T, about 128 MHz — the same band as FM radio, which is why the scanner room is a shielded **Faraday cage**. A radiofrequency pulse at exactly that frequency is absorbed — resonance — and tips the net magnetisation away from the field.

Step 3 — listen. The tipped magnetisation keeps precessing, a rotating magnet inside a coil of wire. It induces a small oscillating voltage — the **MR signal** — in a receive coil placed close to the body part. Coil design (head coils, knee coils, flexible surface arrays) matters enormously to image quality.

Step 4 — relax. The signal does not last. Two independent processes return the system to equilibrium, at rates that depend on the molecular environment of the water:

- **T1 (longitudinal) relaxation** — the magnetisation **regrows** along the main field as the protons give energy to their surroundings. Fast in fat, slow in free water.
- **T2 (transverse) relaxation** — the precessing protons **lose phase** with each other and the signal decays. Slow in free water, faster in tissue.

Tissue (at 1.5 T, approximate)	T1 (ms)	T2 (ms)
Fat	~250	~80
White matter	~600–800	~80
Grey matter	~900–1000	~100
Cerebrospinal fluid	~4000	~2000

Step 5 — encode and reconstruct. If every proton precessed at the same frequency, the coil would receive one undifferentiated signal. **Gradient coils** superimpose small, precisely controlled field slopes on B_0 so that the Larmor frequency — and the phase — varies with position. The received signal is then a sum of frequencies, each from a different place. The raw data are samples of the image's **spatial-frequency spectrum**, called **k-space**, and a two- or three-dimensional **Fourier transform** turns them into the image.

That last point should feel familiar to anyone who has done signal or image processing. The centre of k-space holds the low spatial frequencies — overall contrast and brightness; the edges hold the high frequencies — fine detail and edges. Much of modern MRI speed-up (parallel imaging, compressed sensing, deep-learning reconstruction) is about acquiring **less** of k-space and reconstructing the rest intelligently.

The rapidly switched gradient coils, vibrating in a strong field, are also why the scanner is **loud** — often above 100 dB(A). Hearing protection is mandatory for every patient (Chapter 8).

16.6.2. Choosing the contrast: weighting and sequences

The operator controls contrast by choosing **when** to excite and **when** to listen: the repetition time (**TR**) and echo time (**TE**). Short TR and short TE emphasise T1 differences; long TR and long TE emphasise T2. The same anatomy can look completely different.

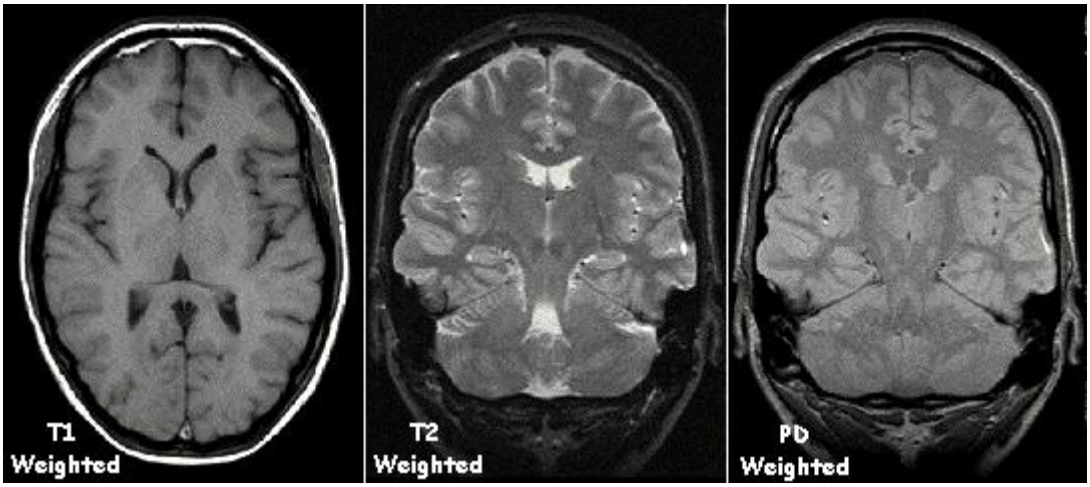


Figure 26. The same brain imaged with T1, T2, and proton-density weighting. On T1, fluid is dark and fat bright; on T2, fluid is bright. Image: KieranMaher, public domain. Credits: Appendix H

Sequence	What it shows	Typical use
T1-weighted	Fat bright, fluid dark; crisp anatomy	Anatomy; with gadolinium contrast, areas of leaky vessels (tumours, inflammation)
T2-weighted	Fluid bright	Most disease raises water content: oedema, inflammation, many tumours, cysts
FLAIR	T2-weighted with the signal from free fluid suppressed	Brain lesions next to the fluid-filled ventricles — multiple sclerosis, for example
Fat-suppressed (e.g. STIR)	Fat signal removed	Bone marrow oedema, subtle fractures, inflammation
Diffusion-weighted (DWI)	Signal depends on how freely water molecules move	**Acute stroke** — visible within minutes, long before CT changes; tumour cellularity
MR angiography	Flowing blood made bright, with or without contrast	Vessels of the brain and neck
Functional MRI (BOLD)	Small signal changes from blood oxygenation, because deoxyhaemoglobin is magnetic	Mapping brain activity; pre-surgical planning; research
Cardiac MRI	Fast, ECG-synchronised imaging of the beating heart	Heart muscle function, scarring, inflammation
MR spectroscopy	The frequencies of protons in different molecules	Chemical composition of a region, e.g. in brain tumours

16.6.3. What MRI is for, and what it costs

MRI has the best soft-tissue contrast of any modality and no ionising dose. It is the first choice for the brain and spinal cord, joints (ligaments, cartilage, menisci), soft-tissue tumours, the pelvis, the liver and biliary tree, and the heart.

Its costs: it is slow (typically 20–60 minutes, lying still), expensive, loud, and enclosed — claustrophobia stops a meaningful fraction of scans. **Gadolinium** contrast agents, which shorten T1, are generally very safe but carry a rare risk of a fibrosing condition in severe kidney failure (mainly with older agents), and small amounts are retained in the body. And the magnet creates a safety environment unlike anywhere else in the hospital.

16.6.4. MR safety: the field is always on



A clinical MRI magnet is superconducting and is never switched off. It is at full field when the room is empty, at night, and during a fire.

The hazards, in order of how often they cause harm:

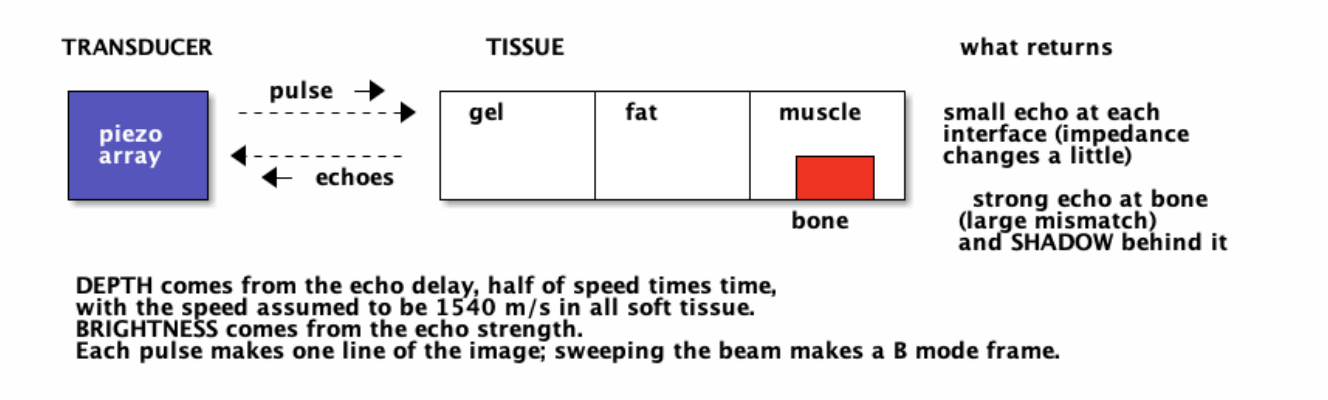
- **Projectiles.** Ferromagnetic objects — oxygen cylinders, scissors, pens, hairpins, tools, wheelchairs, floor polishers — are pulled into the bore with enough force to kill. Force rises steeply near the bore, so an object that seems only mildly attracted at the door can become uncontrollable a metre closer.
- **Implants and devices.** Pacemakers, neurostimulators, cochlear implants, pumps, and metallic fragments can move, heat, or malfunction. Many modern implants are **MR Conditional** — safe only under stated conditions of field strength, gradient, RF power, and position.
- **RF heating.** The RF pulses deposit energy, limited by the **specific absorption rate** (SAR, W/kg). Conductive loops — ECG leads, cables, even skin-to-skin contact between thighs — can concentrate it and cause burns.
- **Peripheral nerve stimulation** from fast gradient switching: tingling or twitching.
- **Acoustic noise:** above.
- **Quench:** an emergency or fault can make the magnet lose superconductivity and boil off its liquid helium through a vent; a failed vent can displace oxygen in the room.

Access is therefore controlled by **zones** — conventionally I (public) to IV (the magnet room) — with screening of every person and object before entry, supervised by trained MR staff. Items are labelled by the ASTM F2503 convention: **MR Safe** (non-conducting, non-metallic, non-magnetic), **MR Conditional** (safe under specified conditions), or **MR Unsafe**.

Nothing you build may enter Zone IV without the MR safety lead's explicit approval, however harmless it looks. A hand magnet test is not a safety assessment: some stainless steels are weakly magnetic in a hand test and strongly attracted at 3 T, and non-magnetic metals still heat and distort images.

16.7. Ultrasound

16.7.1. Pulse and echo



An ultrasound **transducer** is an array of **piezoelectric** elements — ceramic crystals that deform when a voltage is applied and generate a voltage when deformed. Each element is pulsed briefly, sending a short burst of sound at 2–18 MHz into the body, then listens.

Every boundary between tissues of different **acoustic impedance** (Z = density × speed of sound) reflects part of the pulse. The machine assumes the speed of sound in soft tissue is **1540 m/s** and converts each echo's arrival time into depth:

depth = **c** × **t** / **2**

and its strength into brightness. One pulse gives one line of the image; firing the elements in sequence, with electronic **beamforming** to steer and focus, builds a 2D frame — **B-mode** — many times a second.

The fraction of energy reflected at a boundary is:

$$\mathbf{R = Z_2 - Z_1) / (Z_2 + Z_1^{\,2}}$$

Between fat and muscle, about 1% is reflected — enough to see, little enough to continue. Between soft tissue and bone, roughly 40%: a bright echo and a **shadow** behind. Between soft tissue and air, **more than 99.9%**: nothing gets through. That is why **coupling gel** is essential, why ultrasound cannot see through gas-filled bowel or lung, and why bone blocks the view.

16.7.2. Frequency is a trade between detail and depth

Higher frequency means shorter wavelength and finer resolution — at 5 MHz the wavelength in tissue is 1540 / 5,000,000 ≈ 0.3 mm. But attenuation rises with frequency, roughly **0.5 dB per centimetre per MHz**, counted both ways. With a usable dynamic range of about 60 dB:

- at 3 MHz: 60 / (2 × 0.5 × 3) ≈ **20 cm** of useful depth — abdomen, heart, pregnancy;
- at 10 MHz: 60 / (2 × 0.5 × 10) ≈ **6 cm** — tendons, thyroid, superficial vessels, needle guidance.

That is why there are different probes for different jobs, and why the same device cannot image a deep liver and a finger tendon equally well.

Frame rate is limited by physics too. Each line must wait for the deepest echo to return: for 15 cm, 2 × 0.15 / 1540 ≈ 195 µs. With 128 lines per frame that is 25 ms, or about 40 frames per second. Deeper

imaging or more lines means fewer frames.

16.7.3. Doppler: measuring movement

Sound reflected from moving blood returns with its frequency shifted:

$$\Delta f = 2 f_0 v \cos\theta / c$$

where **v** is blood speed and **θ** the angle between the beam and the flow. The shifts fall in the **audible** range — which is why Doppler machines have loudspeakers and why the fetal heartbeat "sounds" the way it does.



Figure 27. Colour and spectral Doppler of a carotid artery. Colour shows flow direction and speed in the vessel; the trace on the left shows velocity over time. Note the probe frequency (3.8 MHz), the angle correction (60°), and the safety indices MI and TIS displayed at the top. Image: Etan J. Tal, CC BY 3.0. Credits: Appendix H

Work the example in the figure: 3.8 MHz, angle 60°, peak velocity 0.88 m/s.

Δf = 2 × 3,800,000 × 0.88 × cos 60° / 1540 ≈ 2.2 kHz — squarely audible.

Now see the measurement problem. The machine computes velocity from the frequency shift **using the angle the operator enters**. At 60°, an error of 5° in that angle changes cos θ by about 15%, and the reported velocity by roughly 13–18%: dividing by a too-small cosine inflates it more than a too-large one shrinks it. Velocity thresholds are used to grade the narrowing of arteries. This is Chapter 15's lesson in a single knob: **a measurement is only as good as its least-controlled input**, and here that input is a person's hand.

Modes you will meet: **colour Doppler** (flow direction and speed overlaid on the image), **spectral Doppler** (a velocity–time trace from one spot), **power Doppler** (sensitive to slow flow, no direction), **M-mode** (one line plotted against time, for heart valves), and **elastography** (tissue stiffness, estimated from the speed of shear waves the machine generates — used to assess liver fibrosis without a biopsy).

16.7.4. What ultrasound is for, and what it costs

Ultrasound is real-time, portable, cheap, and free of ionising radiation. It is the first test for pregnancy, the gallbladder, kidneys, thyroid, testes, deep vein thrombosis, and many tendon and joint problems. **Echocardiography** — ultrasound of the heart, from the chest or from a probe in the oesophagus — is the workhorse of cardiology. **Point-of-care ultrasound** at the bedside answers yes-or-no questions in minutes (fluid around the heart? free fluid in the abdomen after trauma?), and ultrasound **guides needles** into vessels, nerves, joints, and lesions.

Its limits are physics and people. It cannot see through bone or gas. And it is the most **operator-dependent** modality: the image exists only while someone holds the probe in the right place at the right angle, and a missed view leaves no record that it was missed.

Safety: diagnostic ultrasound has an excellent record, but it deposits energy. Machines display a **thermal index** (TI, a heating estimate) and **mechanical index** (MI, a cavitation-risk estimate), and sonographers keep both as low as reasonably achievable, especially in early pregnancy and over the eye.

16.8. Nuclear medicine: SPECT and PET

Every modality so far shows **structure**. Nuclear medicine shows **function**: where a substance goes, how fast, and how much accumulates. A **radiopharmaceutical** — a biologically meaningful molecule labelled with a radioactive atom — is injected, inhaled, or swallowed, it distributes according to physiology, and a camera detects the gamma rays that come out.

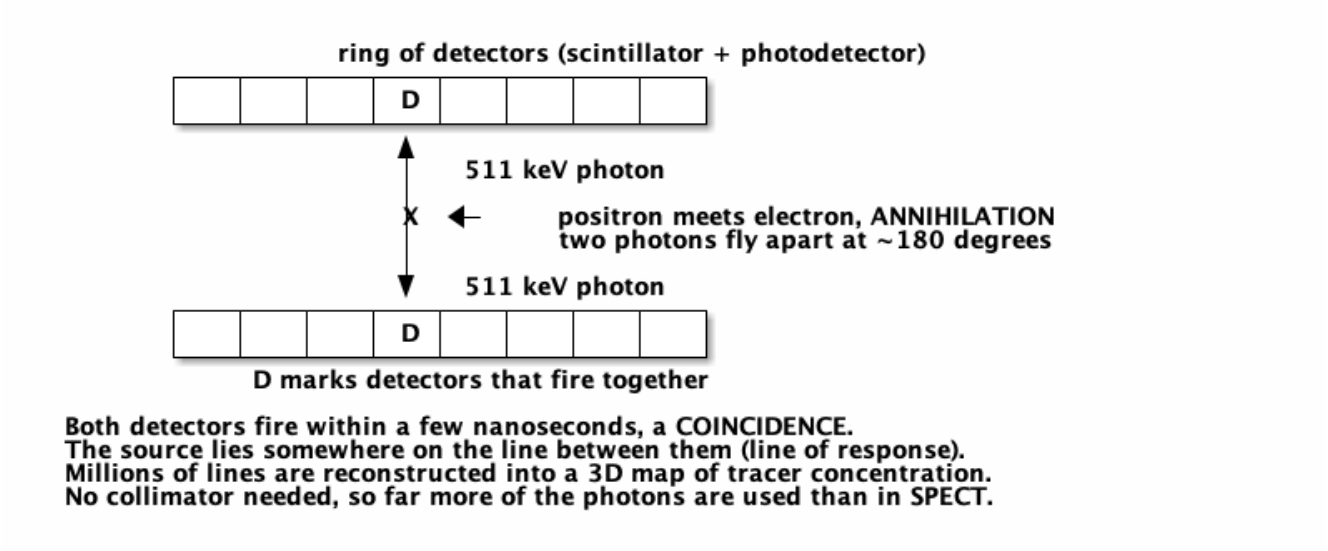
16.8.1. Gamma camera and SPECT

The workhorse isotope is **technetium-99m**: it emits a single 140 keV gamma ray, has a 6-hour half-life (long enough to image, short enough to fade), and is produced on site from a **generator** containing its parent, molybdenum-99. Attached to different carrier molecules it goes to bone, heart muscle, lungs, kidneys, or thyroid.

A **gamma camera** detects the photons with a scintillation crystal that flashes when struck, and photodetectors that locate each flash. To know **where** a photon came from, a thick lead **collimator**

with thousands of parallel holes admits only photons travelling straight in — and rejects all the others, roughly 9,999 in 10,000. That waste is the reason nuclear images are noisy and coarse (resolution around a centimetre). Rotating the camera around the patient gives a 3D reconstruction: **SPECT** (single-photon emission computed tomography).

16.8.2. PET



Positron emission tomography uses isotopes that emit **positrons** — anti-electrons. A positron travels a millimetre or so, meets an electron, and the pair annihilate into **two 511 keV photons travelling in opposite directions**. A ring of detectors looks for pairs arriving within a few nanoseconds of each other. Each such **coincidence** defines a line on which the event happened, so no collimator is needed: PET is far more sensitive than SPECT and sharper, around 4–5 mm clinically. **Time-of-flight** PET goes further, using the difference in arrival times — tens to hundreds of picoseconds — to locate the event along the line.

The dominant tracer is **¹⁸F-fluorodeoxyglucose (FDG)**: glucose with one oxygen replaced by radioactive fluorine. Cells take it up as if it were glucose and phosphorylate it, but cannot metabolise it further, so it is **trapped** in proportion to glucose use. Many cancers are avid glucose consumers, so they light up — but so do the brain, the heart, active muscle, healing tissue, infection, and inflammation, and FDG is excreted through the kidneys into the bladder. Patients fast beforehand and rest quietly during the uptake period, so that muscle and insulin do not distort the map.

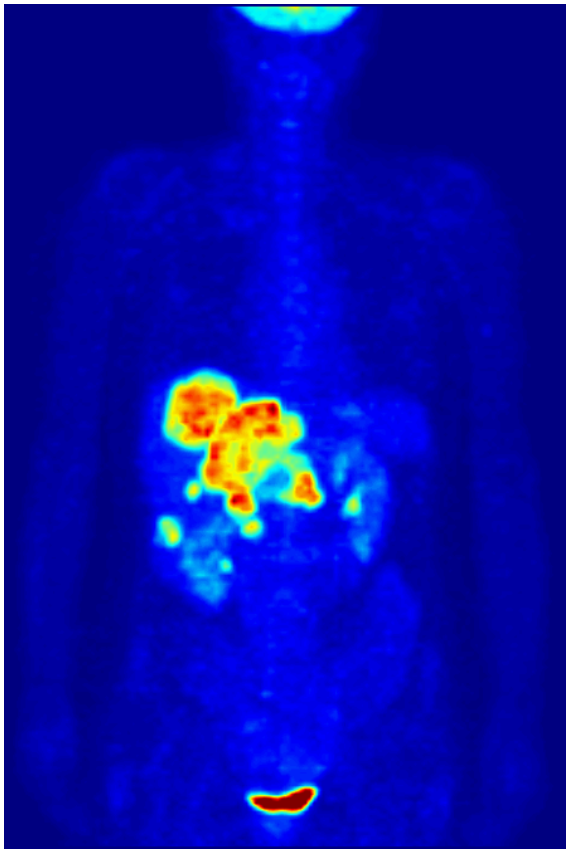


Figure 28. Whole-body FDG PET (maximum-intensity projection) of an adult woman, one hour after injection. Normal uptake appears in the brain, heart, kidneys, and bladder; the bright cluster in the upper abdomen is metastatic cancer in the liver. Image: Jens Maus, public domain. Credits: Appendix H

Fluorine-18 has a half-life of about 110 minutes, so it is made in a **cyclotron** — a particle accelerator — and delivered by road the same morning. Every hour of delay costs about 30% of the activity. Other PET tracers target specific biology: **PSMA** tracers for prostate cancer, **amyloid** and **tau** tracers for dementia, **rubidium-82** or **ammonia** for heart perfusion.

PET is almost always combined with CT (**PET-CT**) or MRI (**PET-MRI**) in one machine. The anatomical scan locates the functional signal — "the hot spot is **that** lymph node" — and supplies the attenuation map needed to correct the PET data for photons absorbed in the body.

16.8.3. The standardised uptake value

PET is quantitative, and its headline number is the **SUV**:

SUV = tissue activity concentration ÷ (injected activity ÷ body weight)

An SUV of 1 means the tissue holds exactly the average concentration; a tumour might be 5 or 15. It is used to compare scans before and after treatment.

It is also a textbook example of a measurement with many confounders: blood glucose, time from injection to scan, body composition (fat takes up little FDG, so heavy patients have inflated SUVs if weight is used), scanner calibration, reconstruction settings, the region drawn, and **decay correction**. In the figure's example — 371 MBq injected, scanned 60 minutes later — the activity remaining at scan time is $371 \times 2^{-60/110} \approx 254$ MBq. Forget the decay correction and every SUV is wrong by a factor of about 1.46. The number looks like a lab value; it behaves like a protocol.

16.8.4. What nuclear medicine is for, and what it costs

Study	Answers
FDG PET-CT	Cancer staging and response to treatment; infection and inflammation of unknown source
PSMA PET	Where prostate cancer has spread
Bone scan	Bone turnover — metastases, stress fractures, infection; very sensitive, not very specific
Myocardial perfusion (SPECT or PET)	Which parts of the heart muscle lose blood supply under stress
V/Q scan	Pulmonary embolism, especially in pregnancy or when CT contrast is unsuitable
Thyroid scan	Overactive nodules
Renal scans	Function of each kidney separately; drainage
Sentinel node mapping	The first lymph node draining a tumour, for the surgeon to remove
Amyloid / tau PET	Supporting diagnosis of Alzheimer's disease

Nuclear medicine also **treats**: radioiodine for an overactive or cancerous thyroid, and the newer **theranostics**, in which the same targeting molecule carries an imaging isotope to find disease and a therapeutic isotope (such as lutetium-177) to irradiate it.

Costs: ionising dose (to the patient, who remains mildly radioactive for hours, and to staff), cost, limited availability, and a logistics chain set by half-lives. Spatial resolution is poor; specificity often is too. The images are maps of physiology, not pictures, and they need anatomy beside them.

16.9. Endoscopy and direct visualisation

An **endoscope** puts a light and a camera inside the body. Modern flexible scopes carry a chip camera at the tip, LED or fibre-optic illumination, a steerable bending section controlled by wires from the handle, and **working channels** for air, water, suction, and instruments — biopsy forceps, snares to remove polyps, injection needles.

Procedure	Looks at
Gastroscopy (upper GI endoscopy)	Oesophagus, stomach, duodenum
Colonoscopy	The whole large bowel; the standard for finding and removing polyps
Bronchoscopy	Airways; biopsies of lung lesions and lymph nodes
Cystoscopy	Bladder and urethra
Nasendoscopy	Nose, pharynx, larynx (Chapter 8)

Procedure	Looks at
Laparoscopy, arthroscopy	Abdomen, joints — rigid scopes through small incisions; usually surgical
ERCP	Bile and pancreatic ducts, combining a side-viewing scope with fluoroscopy
Endoscopic ultrasound	An ultrasound probe on the scope tip, imaging through the gut wall
Capsule endoscopy	The small bowel, by a swallowed camera about 11 × 26 mm that takes several images a second for 8–12 hours

Endoscopy's great advantage is that it **sees the surface directly and can take tissue**. Its costs are invasiveness, sedation, a small risk of perforation or bleeding, and **reprocessing**. A flexible endoscope is a long, narrow, multi-channel, heat-sensitive instrument that cannot be steam sterilised. Outbreaks of drug-resistant infection have been traced to duodenoscopes whose complex tip mechanisms could not be reliably cleaned — the most expensive lesson in Chapter 13's argument that cleanability is a design requirement.

Related optical devices: the **dermatoscope** (magnified, polarised view of skin lesions), **colposcope** (cervix), **otoscope** and **ophthalmoscope** (Chapter 8), and **OCT** (Chapter 15).

16.10. Physiological measurement

Many of the most common diagnostic tests produce a **signal or a number**, not an image. Most were introduced earlier; here they are gathered by what they test.

16.10.1. The heart

- **12-lead ECG.** Ten electrodes give twelve "views" of the heart's electrical activity (Chapter 9, Chapter 7' biosignal table). Rhythm, conduction, heart attack, strain, electrolyte effects. Cheap, fast, and the first test for chest pain or palpitations.
- **Ambulatory ECG.** A **Holter** monitor records for 24–48 hours; adhesive patch recorders for up to two weeks; an **implantable loop recorder** under the skin for around three years — because symptoms that come once a month are rarely present during a ten-second ECG. Consumer watches now record a single-lead ECG and screen for atrial fibrillation.
- **Exercise stress testing.** ECG (sometimes with imaging) during graded exercise, to provoke ischaemia that is absent at rest.
- **Echocardiography.** Above, under ultrasound.
- **Cardiac catheterisation.** Catheters threaded from the wrist or groin measure pressures inside the heart and, with iodine contrast and fluoroscopy, show the coronary arteries. A **pressure wire** across a narrowing gives the **fractional flow reserve** (FFR) — the pressure ratio across it during maximal flow — which decides whether it is worth stenting.
- **Blood tests** — troponin for heart muscle damage, natriuretic peptides for heart failure (Chapter 15).

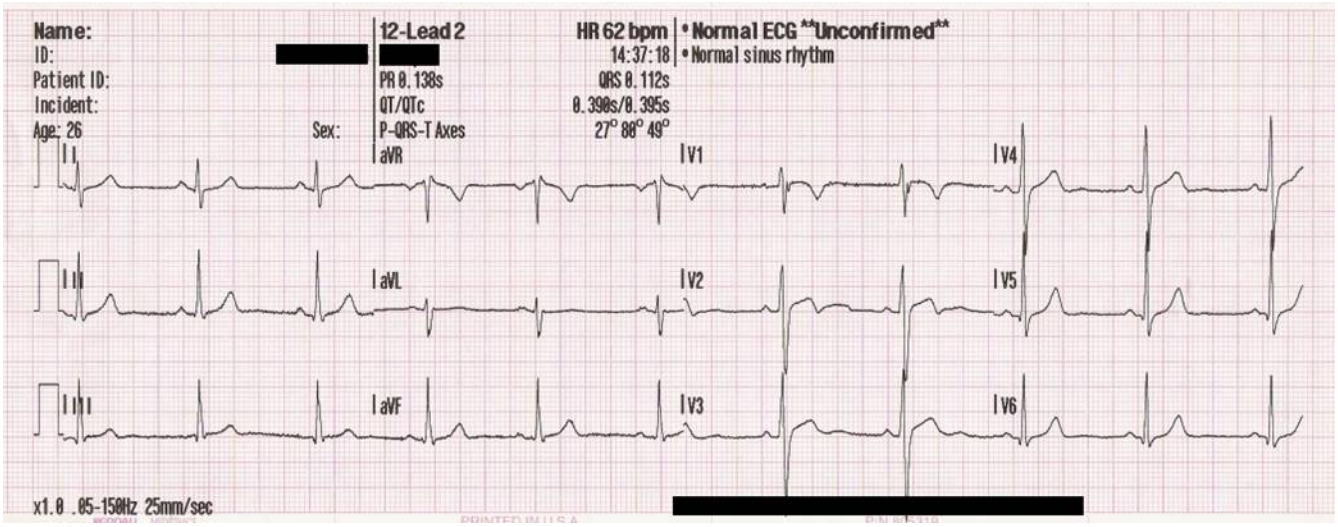


Figure 29. A normal 12-lead ECG. Each row shows the same heartbeats from different electrode combinations; the bottom strip is a continuous rhythm lead. Image: MoodyGroove, public domain. Credits: Appendix H

16.10.2. Blood pressure

Most automated cuffs are **oscillometric**. The cuff inflates above systolic pressure and deflates in steps; the arterial pulsations transmitted to the cuff grow and then shrink. The pressure at the **largest** oscillation corresponds closely to **mean** arterial pressure. Systolic and diastolic values are then **estimated** from the oscillation envelope using manufacturer-specific ratios — which is why automated devices must be validated against reference methods, and why two validated brands can disagree in an individual.

Cuff size is a classic measurement error: a cuff too small for the arm over-reads, sometimes by more than 10 mmHg. **Ambulatory monitoring** over 24 hours avoids the "white coat" rise seen in clinics. In intensive care, an **arterial line** measures pressure directly and continuously through a cannula and a pressure transducer — a sensor, a fluid column, and a dynamic response that must be checked (Chapter 9).

16.10.3. The lungs

- **Spirometry.** The patient breathes in fully and blows out as hard and fast as possible. The key outputs are **FEV_1** (volume exhaled in the first second) and **FVC** (total volume). A reduced FEV_1/FVC ratio — conventionally below about 0.7 — indicates **obstruction** (asthma, COPD); a proportional reduction of both suggests **restriction** (stiff lungs, weak muscles, chest wall problems). The flow–volume loop shows the shape. Spirometry is effort-dependent: technique and coaching are part of the measurement.
- **Peak flow meter.** A cheap handheld device for tracking asthma at home.
- **Full lung function tests — body plethysmography** (lung volumes, measured in a sealed box using Boyle's law) and **gas transfer** (DLCO, how well a tracer gas crosses into the blood).
- **Pulse oximetry and blood gases** — Chapter 9.

16.10.4. Sleep

Polysomnography records, overnight, the EEG, eye movements (EOG), chin and leg muscle activity (EMG), airflow at the nose and mouth, chest and abdominal breathing effort, oxygen saturation, ECG, and body position — a dozen or more channels, synchronised. Sleep stages come from the EEG/EOG/EMG; breathing events from the airflow, effort, and saturation channels. The headline number for **obstructive sleep apnoea** is the **apnoea–hypopnoea index** (AHI): breathing pauses and reductions per hour of sleep, with conventional cut-offs around 5, 15, and 30 per hour for mild, moderate, and severe. Simplified **home sleep apnoea tests** record the breathing channels alone.

16.10.5. Nerves, muscle, and brain

EEG (epilepsy, coma, sleep), **EMG and nerve conduction studies** (nerve and muscle disease), and **evoked potentials** are described in Chapter 7. **Audiometry**, **tympanometry**, and **otoacoustic emissions** are in Chapter 8.

16.10.6. Metabolism and temperature

- **Glucose meters and continuous glucose monitors** (Chapter 3). A CGM measures glucose in **interstitial fluid**, not blood, so it lags behind blood glucose by roughly 5–15 minutes — a lag that matters most when glucose is changing fastest.
- **Thermometers.** Core temperature is estimated from a site — oral, ear, forehead, armpit, rectal, bladder, oesophageal — each with its own offset and error. Infrared ear and forehead thermometers are convenient and sensitive to technique and environment. **A temperature without its site is an incomplete measurement** (Appendix C).
- **Urine dipsticks** — colour-change chemistry pads for glucose, protein, blood, white cells, and nitrites.

16.11. Tissue diagnosis: the biopsy

For many diseases — cancer above all — imaging raises a suspicion and **histopathology** settles it. A **biopsy** removes tissue for examination under the microscope (Chapter 15):

- **Fine-needle aspiration** — a thin needle draws out cells; quick, minimally invasive, cells only.
- **Core needle biopsy** — a spring-loaded needle takes a cylinder of tissue, preserving its architecture. Usually **image-guided** by ultrasound or CT, which is where many needle-guidance phantoms earn their keep.
- **Endoscopic biopsy** — forceps through a scope channel.
- **Punch, shave, and excision biopsies** — of skin and accessible lesions.
- **Surgical biopsy** — removing a lymph node or a whole lesion.

The tissue is then fixed, processed, cut into sections a few micrometres thick, stained, and examined — increasingly also tested for molecular markers that guide treatment. **Liquid biopsy**, detecting tumour DNA circulating in the blood, is an emerging complement.

16.12. Choosing a test

The table below shows **typical** first-line choices for common questions, so that you can see the logic. It is orientation only: real pathways depend on the patient, local guidelines, and availability, and change over time.

Clinical question	Typical first test(s)	Why that one
Is the bone broken?	Plain X-ray	Fast, cheap, excellent for bone; CT or MRI if X-ray is normal and suspicion remains
Is this a stroke, and is it bleeding?	Non-contrast CT, CT angiography; MRI with DWI	Speed; bleeding must be excluded before clot-dissolving treatment
Is this chest pain a heart attack?	ECG and troponin blood test	Cheap, quick, and directly tied to treatment decisions
Is there a clot in the lungs?	CT pulmonary angiography; V/Q scan in some cases	Direct view of the clot; V/Q avoids iodine and lowers breast dose
Is there a clot in the leg?	Compression ultrasound	No radiation; veins that do not compress contain clot
Are there gallstones?	Ultrasound	Stones are highly reflective; no radiation
Is it appendicitis?	Ultrasound in children and pregnancy; CT in many adults	Dose trade against diagnostic accuracy
Is the knee ligament torn?	MRI	Soft-tissue contrast
How far has the cancer spread?	CT; PET-CT for many cancers	Anatomy plus function
Are the bones thinning?	DEXA	Specific, quantitative, tiny dose
Breast cancer screening	Mammography (± tomosynthesis)	Population evidence; detects calcifications
Is the heart pumping well?	Echocardiogram; natriuretic peptide blood test	Real-time function, no radiation
Is this breathlessness asthma or COPD?	Spirometry	Distinguishes obstruction from restriction
Does this person stop breathing in sleep?	Home sleep test or polysomnography	Only an overnight recording shows it
Are these seizures epilepsy?	EEG; brain MRI	Electrical activity plus structure

Notice how often the first choice is the **cheapest adequate** test, and how often radiation dose shapes the decision for children and in pregnancy.

16.12.1. The modalities side by side

Modality	Ionising?	Typical resolution	Typical time	Main constraint
X-ray	Yes (low)	~0.1–0.2 mm	Milliseconds	Superimposition
CT	Yes (higher)	~0.3–0.6 mm	Seconds	Dose; contrast risks
MRI	No	~0.5–1 mm	20–60 min	Magnet safety; time; cost; claustrophobia
Ultrasound	No	~0.1–1 mm, frequency-dependent	Real time	Bone and gas; operator dependence
SPECT	Yes	~10 mm	15–45 min, after uptake delay	Resolution; logistics
PET	Yes	~4–5 mm	~15–30 min, after ~60 min uptake	Cost; cyclotron supply; dose
Endoscopy	No	Surface, microscopic with magnification	10–60 min	Invasive; reprocessing
OCT	No	~5 µm	Seconds	Penetration of 1–2 mm

16.13. Engineering bridge: the mathematics underneath

Look past the hardware and the same handful of ideas recur. Each will be familiar from somewhere in engineering or software.

Idea	Where it appears	Familiar as
Exponential attenuation ($I = I_0 e^{-\mu x}$)	X-ray, CT, PET attenuation correction, ultrasound depth loss, pulse oximetry	Absorbance, signal loss in a cable
Inverse problems	CT, SPECT, PET reconstruction; EEG source localisation	Solving Ax = b when A is huge and ill-conditioned
Fourier transforms	MRI k-space; ultrasound and Doppler spectra; ECG filtering	FFTs, image compression, audio
Time of flight	Ultrasound depth; time-of-flight PET	Sonar, lidar, network latency
Resonance	MRI (Larmor frequency); piezoelectric transducers	Tuned circuits, radio

Idea	Where it appears	Familiar as
Counting statistics	Every photon-counting modality: CT, SPECT, PET, X-ray	Poisson noise; shot noise

The last one deserves a line of arithmetic, because it explains more than any other single fact about diagnostic imaging. When you count **N** random events, the noise is **√N**, so the signal-to-noise ratio is:

SNR = N / √N = √N

To **halve** the noise you need **four times** the counts — four times the dose, or four times the scan time, or a detector four times as sensitive. That square-root law is why dose reduction is hard, why nuclear images are grainy, why better reconstruction (which uses the counts more cleverly) is worth so much engineering, and why there is no free lunch in Chapter 15's image-quality trade.

A few more calculations you should be able to do in your head:

- **Larmor frequency:** 42.58 MHz/T × 3 T ≈ 128 MHz.
- **Echo depth:** an echo returning after 65 μs came from 1540 × 0.000065 / 2 ≈ 5 cm.
- **Decay:** activity after time **t** = $A_0 \times 2^{-t/T½}$. Fluorine-18 two hours after production: $2^{-120/110} \approx 0.47$ of the original.
- **Hounsfield units:** a voxel reading +50 is soft tissue; −80 is fat; +1200 is cortical bone.

16.14. Devices, software, and training technology

This whole chapter is a device survey. This section concentrates on where someone with engineering skills — software, electronics, machining, moulding and casting, metal, glass, or textiles — can actually contribute. Very little of it is the scanner itself; most of it is everything that surrounds the scanner.

16.14.1. Clinical devices and the maker's entry points

Area	What surrounds the scanner	Maker-relevant work
CT and radiography	Positioning aids, immobilisation devices, lead shielding, quality-control phantoms	**CT phantoms**: machined blocks of tissue-equivalent plastics, cast resins and gels, or printed parts whose infill density is tuned to target HU values; QC test objects. **Positioning aids** — foam wedges with wipeable, sewn or welded covers, machined acrylic supports — which must themselves be radiolucent and cleanable

Area	What surrounds the scanner	Maker-relevant work
MRI	Coils, positioning pads, MR-conditional monitoring and infusion equipment	**MRI phantoms** — water or agarose doped with paramagnetic salts to set T1 and T2, in glass or machined acrylic containers. Non-metallic positioning pads and straps (foam, fabric, hook-and-loop). Anything for Zone IV must be designed MR Safe from the outset and approved
Ultrasound	Probes, gel, covers, needle guides	Cast **phantoms** (Chapter 15), probe holders, needle-guide fixtures for teaching; electronic and acoustic test rigs for probe checks
Nuclear medicine	Shielding, syringe shields, dose calibrators, fillable phantoms	Fillable QC phantoms; **nothing** that handles radioactive material without the radiation protection adviser
Endoscopy	Scope reprocessing, storage, and handling	Storage and drying cabinets (sheet metal, airflow, electronics for monitoring), transport trays and covers — Chapter 13's cleanability rules apply in full
Physiological measurement	Electrodes, cuffs, mouthpieces, sensors	Machined test fixtures, pressure and flow rigs, and electronic signal simulators for **bench** checking of equipment — never connected to a person (Chapter 7); training arms and cuffs that involve sewing, casting, and pneumatics

16.14.2. Software

Software	Does	Note
ASTRA Toolbox, TIGRE	Open-source tomographic reconstruction (CT, cone-beam CT)	**Write your own CT reconstruction** from simulated or public projection data
BART (Berkeley Advanced Reconstruction Toolbox)	Open-source MRI reconstruction, including parallel imaging and compressed sensing	Pairs with public raw k-space data
fastMRI dataset	Large public set of raw MRI k-space data with reconstructions	The standard benchmark for learned MRI reconstruction
Field II, k-Wave	Ultrasound field and wave propagation simulation	Model a transducer, a phantom, and the image it would make

Software	Does	Note
GATE / Geant4, STIR	Monte Carlo simulation of nuclear imaging; open-source PET/SPECT reconstruction	How PET physics is designed and taught
3D Slicer, OHIF, pydicom, SimpleITK	Viewing, windowing, segmentation, HU measurement	Chapter 15. Apply rescale slope and intercept; threshold on HU, not grey levels
WFDB, NeuroKit2, PhysioNet	ECG and physiological signal processing, with open recordings	Build an ECG beat detector or an apnoea-event detector on public data
Bloch-equation simulators	Simulate magnetisation, relaxation, and pulse sequences	The most direct way to understand T1, T2, TR, and TE — several free and browser-based ones exist

16.14.3. Training and simulation

Training technology	Teaches	Form
Anthropomorphic radiography phantoms	Positioning, exposure selection, image critique	Human-shaped bodies with real or tissue-equivalent bone. Expensive; partial models are makeable
CT dose phantoms	Dose measurement and protocol optimisation	Standard acrylic cylinders (16 and 32 cm) with holes for an ionisation chamber
MRI safety training	Zone control, screening, projectile awareness	Courses and drills; physical props for screening practice — **props that are MR Safe**
Ultrasound simulators	Probe handling and image interpretation	Phantoms, and instrumented mock probes driving recorded or synthetic images
Echocardiography simulators	Cardiac views, transthoracic and transoesophageal	Manikin with a tracked probe and synthetic image generation
Endoscopy simulators	Scope handling, navigation, polyp removal	Virtual reality with haptic feedback; physical colon and airway models
Laparoscopic box trainers	Camera handling, instrument coordination, suturing	A box, a camera, ports, and tasks. **A classic and well-documented maker build**
Bronchoscopy airway models	Navigating the branching airways	Built from CT data of an airway tree — printed directly, or printed as a core and cast in silicone for realistic flexibility; the scan-to-model pipeline in action

Training technology	Teaches	Form
ECG and patient monitor simulators	Rhythm recognition; checking that monitors respond correctly	Signal generators that output ECG waveforms into equipment, not into people. **An electronics project with real use in teaching and equipment checking**
Spirometry and blood-pressure training	Coaching technique; cuff selection and placement	Training arms with adjustable pulse; calibration syringes



Three of the best projects in the book are in this chapter, and they use three different skill sets.

- **A laparoscopic box trainer** — sheet material, a camera, lighting, and well-designed tasks. Physical, cheap, wanted, and squarely Tier 1.
- **An ECG or monitor waveform simulator** — a microcontroller, a DAC, a resistor network to produce realistic lead amplitudes, and an enclosure. For checking and teaching equipment, never for a person.
- **A CT reconstruction from scratch** — take a phantom image, simulate projections at many angles, add Poisson noise, then reconstruct with simple back-projection, filtered back-projection, and an iterative method. A weekend in Python teaches more about CT than a month of reading, and every image you produce illustrates the square-root law.

16.15. Worked example: the positioning model that went to the MRI room

The radiography educator who reviewed your elbow model has a new request. Students practising MRI knee positioning use a volunteer; could you make a knee model, roughly life-size, that they can strap into a knee coil on the actual scanner table during quiet sessions?

Your first prototype uses what your workshop does well:

- femur and tibia turned and milled from acetal, joined with M4 stainless steel screws;
- a cast silicone soft-tissue sleeve;
- a small neodymium magnet as a catch for the removable patella;
- a coin-cell LED that lights when the leg is correctly rotated;
- a loop of copper wire as a "nerve" for teaching anatomy;
- a sewn webbing strap with a steel buckle to hold it in the coil.

Walk through the physics.

1. **The magnet** is the worst item. Near the bore of a 1.5 T or 3 T scanner it will be pulled violently,

and it can damage the coil, the patient-contact surfaces, or someone's hand. **MR Unsafe.**

2. **The screws.** "Stainless" covers many alloys; some are ferromagnetic, and even austenitic grades can become weakly magnetic after machining. Even fully non-magnetic metal distorts the field and creates image artefacts. Replace with **nylon or PEEK fasteners**.
3. **The coin cell and LED circuit** contain ferromagnetic parts and conductive loops. A cell can be attracted, and the circuit can have currents induced by the switched gradients and RF. Remove it. Indicate correct rotation by **geometry** — an alignment mark that lines up only in the correct position (Chapter 7' redundant coding again).
4. **The copper loop** is non-magnetic, but a conductive loop in an RF field can heat. On a model that will be scanned, it is also an artefact source. Replace with a **channel moulded into the silicone** or a coloured cord.
5. **The strap.** The webbing is fine; the steel buckle is not. Use a plastic side-release buckle or hook-and-loop fastening — and check any thread, rivets, or eyelets too.
6. **The acetal and silicone** are fine magnetically. Silicone gives an MR signal of its own, which can be useful (it looks like tissue) or confusing; dense plastics show little or nothing. Worth telling the educator.

Then follow the process, not your reasoning. Even with every metal part removed, the model is not yours to declare MR Safe. Take the design, a materials list, and a sample to the department's **MR safety lead**. Ask how they assess and label items for Zone IV, and whether the model must be marked. Record the decision and who made it.

And the lesson about the brief. "Make a knee model" was a training-aid request. "That goes into a 3 T magnet room" turned it into a project with a physical-safety requirement that dominates every material choice. Ask where the object will be used **before** you choose how to fasten it.

16.16. Workshop exercise

1. **Back-projection by hand.** On squared paper, draw a 3 × 3 grid and fill it with numbers (0–9) — your "object". Write down the row sums and column sums: two "projections". Give only the sums to someone else and ask them to reconstruct the grid. Then add the two diagonal sums and try again. Record how many solutions each set of projections allows. That is why CT needs hundreds of angles.
2. **Time-of-flight ranging.** With an ultrasonic distance sensor (the common 40 kHz type) and a microcontroller, measure distances to a wall by timing echoes, using 343 m/s for sound in air. Then move a cardboard sheet partly into the beam and watch the echo. This is B-mode in one dimension — **in air, on objects only**.
3. **Window a real CT.** Download a public CT scan (TCIA, or the sample data in 3D Slicer). Measure the HU of air, fat, muscle, and bone with a region tool. Display the same slice in brain, soft-tissue, lung, and bone windows and note what each hides.
4. **k-space in ten lines.** In Python, load any greyscale image, take its 2D FFT, and reconstruct after (a) keeping only the central 10% of k-space (10% of the width in each direction, so 1% of the samples), and (b) removing it and keeping only the edges. Describe what each looks like and connect it to MRI.

5. **Decay planning.** An FDG dose is calibrated at 07:30. The patient is injected at 09:00 and scanned at 10:00. Calculate the fraction of the calibrated activity remaining at injection and at scan time. Then calculate the SUV error if a scan delayed to 10:30 is processed as if it happened at 10:00.
6. **Square-root law.** Simulate photon counting: generate Poisson random numbers with means of 10, 100, 1000, and 10,000 and plot the ratio of mean to standard deviation. Confirm it follows √N.

16.17. Questions

1. Which diagnostic modalities use ionising radiation, and which do not?
2. Why does bone appear white on an X-ray, and why does iodine make a good contrast agent?
3. What problem of projection radiography does CT solve, and how?
4. What is a Hounsfield unit, and why must segmentation operate on HU rather than on the displayed grey levels?
5. Explain, in four or five sentences, how MRI produces a signal and how that signal is located in space.
6. On which MRI weighting is fluid bright, and why is that useful for finding disease?
7. Name three MRI hazards and explain why "the scanner is switched off at night" is false.
8. Why can ultrasound not image through the lung or bowel gas, and why is gel needed?
9. Why does a higher-frequency ultrasound probe give a sharper but shallower image?
10. In a Doppler measurement at an insonation angle of 60°, why does a small error in the entered angle matter so much?
11. How does PET locate an event without a collimator, and why does FDG show cancer — and what else does it show?
12. Give two reasons a PET SUV can differ between two scans of an unchanged tumour.
13. Why would a 12-lead ECG at the clinic miss a cause of occasional palpitations, and what devices address that?
14. Choose a first-line test for: a suspected broken wrist; suspected gallstones; suspected torn knee ligament. Justify each in one sentence.

16.18. Answers and reasoning

1. **Ionising:** plain X-ray, mammography, DEXA, fluoroscopy, angiography, CT, SPECT, PET, bone scans. **Non-ionising:** MRI, ultrasound, endoscopy, OCT, and the physiological measurements (ECG, spirometry, EEG, and so on).
2. Bone contains calcium, whose higher atomic number makes photoelectric absorption much more likely than in soft tissue, so fewer photons reach the detector behind it. Iodine has a high atomic number and an absorption edge at about 33 keV, in the diagnostic energy range, so even modest concentrations absorb strongly.
3. **Superimposition**: a radiograph collapses three dimensions into two. CT acquires projections from hundreds of angles and **reconstructs** the attenuation of each voxel by filtered back-projection or iterative methods, producing true cross-sections.

4. A calibrated scale of X-ray attenuation with water = 0 and air = −1000. The displayed image is a **window** — an arbitrary mapping of a band of HU to grey levels — so a threshold on screen intensity depends on the display settings rather than on the tissue. Thresholds and exports must use HU (after applying rescale slope and intercept).
5. A strong field aligns a small excess of hydrogen nuclei, giving a net magnetisation that precesses at the Larmor frequency. An RF pulse at that frequency tips it over; as it precesses it induces a signal in a receive coil, which decays and recovers at tissue-specific rates (T2, T1). Gradient coils make the frequency and phase depend on position, so the signal samples the image's spatial-frequency content (k-space). A Fourier transform turns those samples into an image.
6. **T2-weighted** (and FLAIR, with free fluid suppressed). Most disease processes — inflammation, oedema, many tumours — increase tissue water, so they stand out as bright.
7. Any three of: projectile attraction of ferromagnetic objects; implant movement, heating, or malfunction; RF heating and burns from conductive loops; peripheral nerve stimulation; acoustic noise; helium release during a quench. The magnet is **superconducting** and stays at full field permanently; only a deliberate or emergency quench removes it.
8. Soft tissue and air differ so greatly in acoustic impedance that more than 99.9% of the energy is reflected at the boundary, so nothing passes beyond it. Gel removes the air layer between the probe and the skin for the same reason.
9. Higher frequency means shorter wavelength and therefore finer resolution, but attenuation rises roughly in proportion to frequency (~0.5 dB/cm/MHz, both ways), so the echo from depth falls below the noise sooner.
10. Velocity is computed from the frequency shift divided by $\cos\theta$. Near 60°, $\cos\theta$ changes steeply: an error of 5° changes the computed velocity by roughly 13–18%. Velocity thresholds grade arterial narrowing, so an operator's angle placement can move a patient between categories.
11. Positron annihilation produces two 511 keV photons travelling in opposite directions; detectors on opposite sides of the ring registering within nanoseconds of each other define a line containing the event. FDG is taken up like glucose and trapped, so it shows tissues with high glucose use — many cancers, but also brain, heart, active muscle, infection, inflammation, and healing tissue, with excretion into the urinary tract.
12. Any two of: different time between injection and scanning; different blood glucose; weight change or use of body weight versus lean mass; incorrect decay correction; different scanner, calibration, or reconstruction; differently drawn region of interest.
13. The arrhythmia may simply not be happening during a ten-second recording. Holter monitors (24–48 h), patch recorders (up to ~2 weeks), event recorders, implantable loop recorders (~3 years), and consumer single-lead ECG devices extend the observation window.
14. **Wrist:** plain X-ray — fast, cheap, excellent for bone. **Gallstones:** ultrasound — stones are highly reflective and there is no radiation. **Knee ligament:** MRI — the soft-tissue contrast shows ligaments, menisci, and cartilage that X-ray and CT do not.

16.19. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- CT and MRI lab — projection puzzle, sinograms, back-projection, HU windowing, k-space and MRI weighting.
- Ultrasound, Doppler and dose numbers — pulse-echo depth, reflection, Doppler angle error, air ranging and dose comparisons.
- DICOM viewer and de-identification check — windowing a real CT from a public archive.
- Half-lives, decay and steady state — FDG decay and SUV correction.
- Biosignal analysis — ECG and PPG analysis.
- Probability — counting statistics.

16.20. Video companion and reading

- **Patient-level explanations of every test:** RadiologyInfo [43], written by radiologists for the public. Read the pages for X-ray, CT, MRI, ultrasound, and PET-CT — they are short and accurate, and they describe the patient's experience, which is what your designs will meet.
- **MRI physics, one question at a time:** Questions and Answers in MRI [44]. Free, thorough, and organised by the questions engineers actually ask.
- **Radiation dose in context:** Image Wisely [45].
- **Cases and images:** Radiopaedia [38] — browse normal anatomy on each modality before looking at disease.
- **Optional textbook:** Bushberg et al., *The Essential Physics of Medical Imaging* [46] — the standard reference if you want the full mathematics. Borrow it.
- **Core reading:** OpenStax [1] section 1.7 Medical Imaging gives a short overview; it does not treat the physics in depth, which is what this chapter is for.

16.21. Exit criteria

- ☐ For X-ray, CT, MRI, ultrasound, and PET, state in two sentences each what physical interaction is used, what is detected, and how the image is formed.
- ☐ Sort a list of tests into ionising and non-ionising, and place five of them on the dose table in the right order of magnitude.
- ☐ Measure HU of four tissues in a public CT dataset and explain the difference between stored values and a display window.
- ☐ Reproduce the Doppler equation, work one example, and state the effect of a 5° angle error at 60°.
- ☐ Complete the k-space exercise and explain what the centre and edges of k-space carry.
- ☐ List every material and component in a hypothetical object for an MRI room and mark each as acceptable, questionable, or excluded, with reasons — and name who would make the final decision.
- ☐ Choose and justify a first-line test for five common clinical questions, labelled as orientation, not advice.

16.22. What this chapter does not establish

- **You cannot interpret a diagnostic image or trace.** Reading images and signals is a specialist clinical skill learned over years. Recognising that a CT is in a lung window, or that an ECG has twelve leads, is not reading it.
- **You cannot operate diagnostic equipment on people.** Radiographers, sonographers, physiologists, nuclear medicine technologists, and endoscopists are trained and, in most places, registered professionals. Exposing a person to ionising radiation requires a legal justification and an authorised operator.
- **You cannot declare anything MR Safe, radiolucent, or radiation-safe.** Those are determinations made under the department's procedures by named people — the MR safety lead, the medical physicist, the radiation protection adviser.
- **The test-choice table is not clinical guidance.** It shows reasoning patterns. Real pathways are set by guidelines, the individual patient, and local services.
- **Physics knowledge is not imaging-physics qualification.** Medical physicists and clinical scientists hold formal training and registration. If a project needs a dose calculation, an acceptance test, or a safety assessment, it needs one of them.

Chapter 17. Treatment and life-support devices: how they work and what they are for

Week 20 · approximately 6–8 hours

17.1. Why this chapter exists

Chapter 16 opened the machines that **find out** what is wrong. This chapter opens the machines that **do something about it**: the ventilator that breathes for a person, the pump that delivers a drug drop by drop, the defibrillator that restarts a heart, the dialysis machine that stands in for failed kidneys, the electrosurgical unit that cuts and seals, the linear accelerator that kills a tumour, and the implants that stay inside a body for decades.

Almost all of them appear in Appendix I under **Tier 3 — professional territory**. That is not a reason to skip them. It is a reason to understand them, for three practical reasons:

- **They are the context of most clinical work you will be near.** A training manikin is ventilated, a simulation scenario uses an infusion pump, a positioning aid sits in a radiotherapy beam. You cannot build the accessories, trainers, and software around these devices without knowing what they do and how they fail.
- **They are where the engineering lessons were learned the hard way.** Alarm fatigue, unit errors in drug calculations, a radiotherapy machine whose software overdosed patients, burns from an electrosurgical return pad, surgical fires: the safety culture of Chapter 19 was written in these incidents.
- **They are the devices you will be asked to make.** In 2020, makers worldwide were asked to build ventilators. A few understood what a ventilator has to do. This chapter is so that you would be one of them — and would know, from the physics, why the honest answer was usually to build something else.

For each device family the chapter asks the same questions as Chapter 16: **What physical principle does it use? How does it work? What is it for? How does it fail? Who operates it?**

17.2. Life support, and why it is different

A **life-supporting** or **life-sustaining** device is one whose failure, or whose incorrect use, can kill within minutes. Ventilators, infusion pumps delivering critical drugs, defibrillators, pacemakers, dialysis and bypass circuits all qualify. They are placed in the highest regulatory risk classes, and three engineering consequences follow:

- **Single-fault safety.** The design must stay safe — or fail to a safe state, with an alarm — after any single component failure: a sensor, a valve, a power supply, a line of code. "It works when everything works" is not a specification.
- **Alarms are part of the therapy.** The device cannot fix every problem, so it must tell a person,

reliably and in time. That makes alarm design a safety function — and alarm **overload** a hazard in its own right.

- **The operator is part of the system.** Most harm from these devices involves a person setting, connecting, or interpreting them, which is why Chapter 19' usability engineering is not optional here.



Alarm fatigue is a recognised patient-safety hazard.

A typical intensive care bed may generate hundreds of alarms a day, and studies repeatedly find that a large majority of them are false or need no action. Staff learn — correctly, most of the time — that an alarm probably means nothing. Then one does. Alarm-related deaths led accreditation bodies to make alarm safety a formal hospital priority [62].

Every alarm you add to anything competes for the same attention. Design alarms that are rare, specific, actionable, and distinct (Chapter 8' acoustics, Chapter 7' redundant coding), and treat "add an alarm" as a decision with a cost, not a free safety improvement.

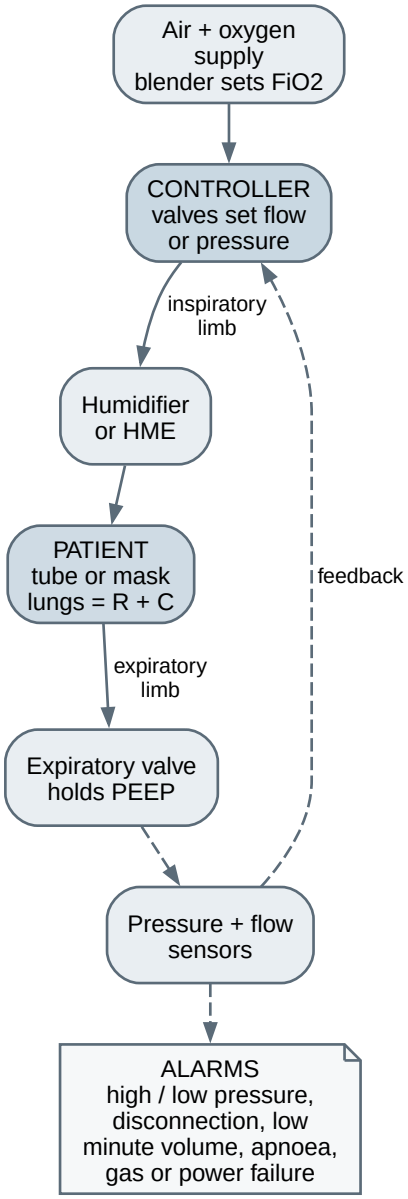
17.3. Mechanical ventilation

17.3.1. Breathing in reverse

In normal breathing (Chapter 9) the diaphragm contracts, chest pressure falls **below** atmospheric, and air is **drawn** in. A ventilator does the opposite: it **pushes** gas in under **positive pressure**, then lets the lungs' elastic recoil push it out.

That reversal has consequences beyond the lungs. Positive pressure in the chest squeezes the great veins and reduces blood returning to the heart, so starting ventilation can drop blood pressure. Too much pressure or volume injures the lung itself — **ventilator-induced lung injury** — which is why modern practice uses relatively small breaths, around 6 mL per kilogram of **predicted** body weight in the most severe lung disease, and limits the pressure at the end of inspiration. A tube through the vocal folds (Chapter 8) bypasses the nose, so the gas must be warmed and humidified by machine, and it gives bacteria a route to the lungs (Chapter 13).

17.3.2. The circuit



A ventilator is a closed-loop flow and pressure controller. A blender mixes air and oxygen to the set **fraction of inspired oxygen** (**FiO_2**, from 0.21 to 1.0). Valves deliver each breath; sensors measure pressure and flow continuously; an expiratory valve closes before the lungs empty completely, holding **positive end-expiratory pressure** (**PEEP**) — typically at least 5 cmH_2O — to keep small airways and alveoli from collapsing.

17.3.3. Modes, in one distinction

There are dozens of named modes, with names that differ between manufacturers. Nearly all are variations on one choice: **what does the machine control, and what does it let vary?**

	Volume control	**Pressure control**
Operator sets	Tidal volume and flow	Inspiratory pressure and time
Machine delivers	The set volume, whatever pressure that takes	The set pressure, whatever volume that produces

	Volume control	Pressure control
If the lungs stiffen or the tube kinks	**Pressure rises** — risk of injury; high-pressure alarm	**Volume falls** — risk of under-ventilation; low-volume alarm
Monitor closely	Airway pressures	Delivered tidal and minute volume

Breaths may be **mandatory** (machine-timed), **assisted** (triggered by the patient's own effort, detected as a small pressure or flow change), or **spontaneous with support**. Weaning a person off a ventilator is a gradual handover of the work of breathing back to them.

17.3.4. Without a tube: non-invasive support and oxygen

- **CPAP** (continuous positive airway pressure) applies one constant pressure through a mask. It splints the upper airway open in **obstructive sleep apnoea** (Chapter 16) and helps fluid-filled lungs.
- **Bilevel** non-invasive ventilation (**NIV**, often called BiPAP) alternates a higher inspiratory and lower expiratory pressure, assisting each breath — used, for example, when carbon dioxide builds up in severe COPD.
- **High-flow nasal oxygen** delivers heated, humidified gas at up to about 60 L/min through wide nasal prongs — enough to exceed the patient's own inspiratory flow, so the delivered oxygen fraction is far more predictable than with an ordinary cannula (Chapter 8).
- **Oxygen supply** comes from hospital piped systems, high-pressure **cylinders** (typically around 140–230 bar when full), or **concentrators**, which use pressure-swing adsorption on a zeolite to strip nitrogen from room air, typically yielding about 90–95% oxygen at a few litres per minute.

For a maker, two of these matter immediately. **Masks** are a pressure-injury problem on the bridge of the nose (Chapter 8, Chapter 5), and **oxygen is a fire accelerant**: materials that are barely flammable in air burn fiercely in enriched oxygen. Nothing with a spark, a hot spot, or oil or grease goes near an oxygen supply.

17.3.5. Who operates it

Intensive care doctors and anaesthetists prescribe ventilation; nurses, respiratory therapists (in some countries) or physiotherapists manage it at the bedside; clinical engineering maintains the machines. Home ventilation and CPAP are increasingly managed by patients and families — which makes the user interface a home-use design problem.

17.4. Infusion pumps and syringe drivers

17.4.1. Two mechanisms

	Volumetric (large-volume) pump	Syringe driver
Mechanism	Peristaltic fingers or a cassette squeeze a section of a dedicated administration set	A lead screw drives the plunger of a syringe at a controlled speed

	Volumetric (large-volume) pump	Syringe driver
Typical use	Fluids, antibiotics, feeds; rates of several to hundreds of mL/h	Small volumes of concentrated, potent drugs at low rates — down to fractions of a mL/h; neonates; pain relief; end-of-life care
Typical accuracy	Around ±5% over time, dependent on the set	Good once running; start-up can be slow at low rates

17.4.2. How they fail

Each failure mode has a design response, and every response has a limitation:

- **Free flow.** If the set is removed from the pump without being clamped, gravity can empty the bag into the patient in minutes. Response: **anti-free-flow** clamps that close automatically when the set is removed.
- **Occlusion.** A kinked line or closed tap stops flow. The pump detects rising pressure — but at low rates pressure builds slowly, so an occlusion can go **unalarmed for a long time**, and when it is released the stored volume arrives as a **bolus**.
- **Start-up delay and siphoning in syringe drivers.** Mechanical slack and the compliance of the syringe and tubing mean that at very low rates actual delivery can lag the set rate for many minutes. A syringe mounted well above the patient and not properly secured can **siphon**; raising or lowering the pump can push or pull a small bolus. For a potent drug given at 0.5 mL/h, a "small" bolus is a large dose.
- **Air in the line.** Detected ultrasonically; the pump stops and alarms.
- **Programming error.** The dominant failure mode. A rate entered as 50 instead of 5.0, a dose in mg where mcg was intended, weight in pounds treated as kilograms, a rate (mL/h) confused with a dose (mg/h), the wrong drug selected from a list. Ten-fold and thousand-fold errors are documented repeatedly by medication-safety organisations [63].

17.4.3. Smart pumps and their limits

Dose error reduction software — a **drug library** loaded onto "smart" pumps — holds the concentrations and usual dose ranges for each drug in each clinical area. A **soft limit** asks "are you sure?"; a **hard limit** refuses. These catch many gross errors. They cannot catch an error within the allowed range, a wrong drug selected, a library bypassed by programming the pump in "basic" mode, or a correct number for the wrong patient. **Interoperability** — the electronic prescription sending its numbers to the pump directly — removes the manual entry step, and with it a whole class of keying errors, at the cost of new software failure modes (Chapter 20).

Insulin pumps (Chapter 10) are the same idea, worn on the body: a small syringe driver delivering a continuous **basal** rate plus **boluses** for meals through a fine subcutaneous cannula. Paired with a continuous glucose monitor and an algorithm, they form **automated insulin delivery** ("closed-loop" or "hybrid closed-loop") systems, in which the algorithm adjusts the basal rate and the person still announces meals.

17.5. Defibrillators, pacemakers, and implantable defibrillators

17.5.1. Defibrillation

In **ventricular fibrillation** the heart's electrical activity has disintegrated into chaotic, uncoordinated wavelets (Chapter 9). The muscle quivers and pumps nothing. A **defibrillator** passes a large, brief current through the chest that depolarises a critical mass of heart muscle at once, so the chaotic wavelets are extinguished and the heart's own pacemaker can take over.

The machine charges a **capacitor** and discharges it through two pads or paddles. Energy is quoted in joules — commonly 120–200 J for modern **biphasic** external devices (current flows one way, then reverses), which are more effective at lower energy than older **monophasic** waveforms (typically up to 360 J). The chest's resistance — **transthoracic impedance**, often in the region of 50–150 Ω — determines the current that actually flows, so modern devices measure impedance through the pads and adjust the waveform to compensate.

A related procedure, **cardioversion**, delivers a smaller shock **synchronised** to the ECG's R wave, to reset organised but abnormal rhythms such as atrial fibrillation without triggering fibrillation by shocking during the vulnerable phase.

17.5.2. The AED: usability as a life-saving feature



Figure 30. An automated external defibrillator. The lid pictograms show pad placement and chest compressions; voice prompts lead an untrained rescuer through each step. Image: VoLid27, CC0. Credits: Appendix H

An **automated external defibrillator** is designed to be used by someone who has never seen one, under extreme stress, in a public place. It turns on when the lid opens. It talks. The pads show

where they go. It analyses the rhythm itself and will only shock a rhythm its algorithm classifies as shockable, so a bystander cannot shock a person who does not need it. It prompts chest compressions and times them.

The stakes justify the effort: for a cardiac arrest from ventricular fibrillation, the chance of survival falls substantially with every minute without defibrillation. The AED is one of the best examples in medicine of Chapter 7' perception–cognition–action–feedback decomposition applied to a device — every stage of the task is designed for, for a user who brings no knowledge at all — and of a high-risk device made safe for lay use **by** its design rather than by training.

17.5.3. Pacemakers and ICDs

A **pacemaker** is a battery-powered pulse generator, usually implanted under the skin below the collarbone, with one to three leads passed through veins into the heart (or, in **leadless** designs, a capsule placed inside the heart itself). It **senses** the heart's own activity and **paces** — delivers a brief pulse of a few volts lasting a fraction of a millisecond — only when the natural rhythm is too slow or a beat fails to arrive. Batteries typically last many years, and a replacement means a further operation.

Pacing modes are summarised in a short letter code: the first letter is the chamber paced, the second the chamber sensed, the third the response to sensing. **VVI** paces and senses the ventricle and inhibits pacing when a natural beat is sensed; **DDD** does both chambers and keeps atria and ventricles coordinated. The code is worth recognising; the choice is a cardiologist's.

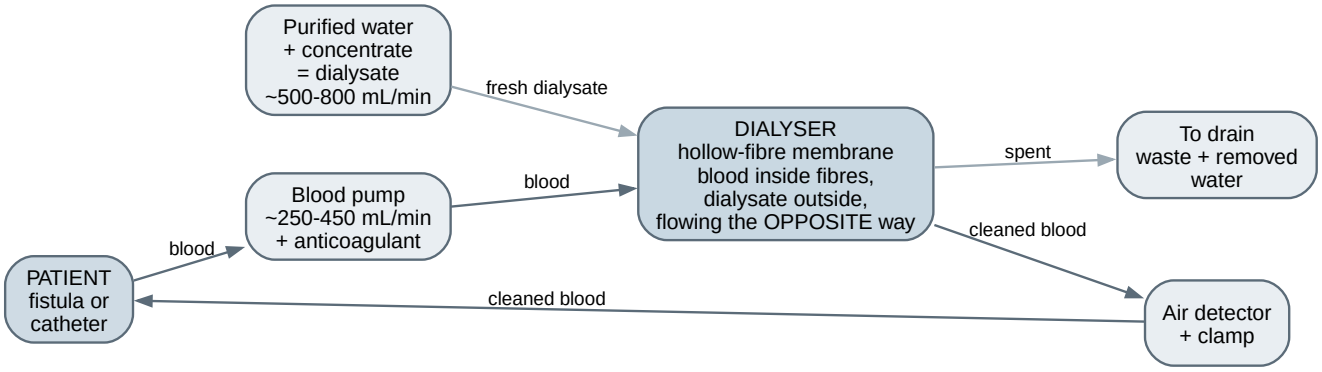
An **implantable cardioverter-defibrillator** (ICD) adds detection of dangerous fast rhythms and can respond with rapid pacing to interrupt them or with an internal shock of up to a few tens of joules. Shocks delivered to a conscious person are painful and frightening; inappropriate shocks — from misread signals or lead faults — are a recognised harm, and the detection algorithm is the heart of the device.

Implanted cardiac devices are now **networked**: they transmit to home monitors and on to the clinic. That brings real benefit and a new hazard class. In 2017 a firmware update was issued for roughly half a million pacemakers after a cybersecurity vulnerability was identified — the subject of Chapter 20. They also carry strong restrictions for MRI (Chapter 16).

17.6. Dialysis and extracorporeal circuits

17.6.1. Haemodialysis

When the kidneys fail (Chapter 10), waste products, potassium, acid, and water accumulate. **Haemodialysis** removes them by passing blood outside the body through a **dialyser** — a cartridge containing thousands of hollow fibres made of a semipermeable membrane, with a total surface area of around 1.5–2 m².



Three physical processes do the work:

- **Diffusion** — small solutes (urea, potassium) move across the membrane down their concentration gradient from blood into dialysate. Running dialysate **countercurrent** to the blood keeps a gradient along the whole length of the fibre, exactly as in a countercurrent heat exchanger.
- **Ultrafiltration** — a pressure difference across the membrane pushes water out of the blood, removing the fluid the patient cannot excrete. Too much, too fast, and blood pressure falls.
- **Convection** — dissolved substances carried along with that water (exploited deliberately in **haemodiafiltration**).

A typical outpatient schedule is about four hours, three times a week, through a surgically created **arteriovenous fistula** that makes a forearm vein large and fast-flowing enough to needle repeatedly. Each session uses well over a hundred litres of dialysate, so **water treatment** — reverse osmosis and continuous quality monitoring — is a major part of every dialysis unit's engineering. **Peritoneal dialysis** uses the lining of the abdomen as the membrane instead, with fluid exchanged through a catheter, often at home overnight by machine.

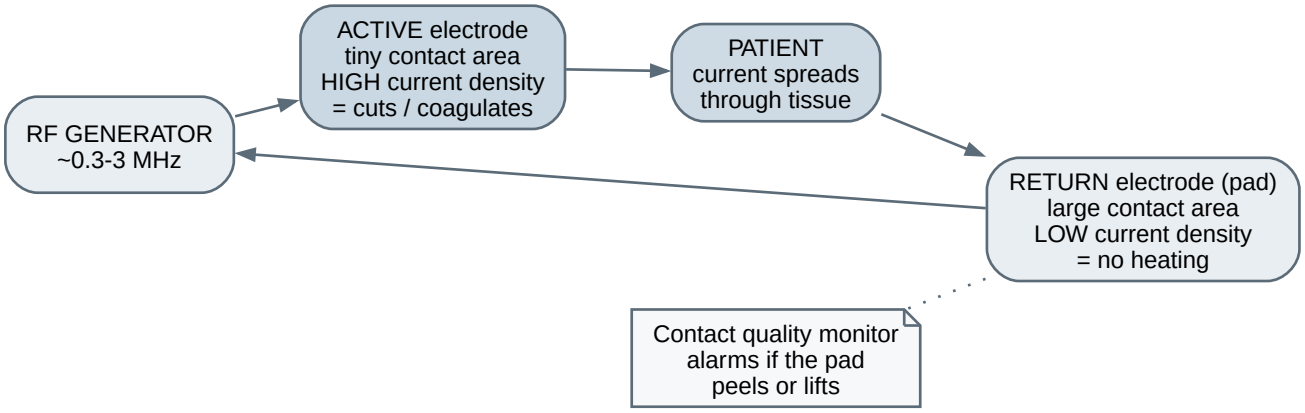
17.6.2. ECMO and cardiopulmonary bypass

The **heart–lung machine** used in cardiac surgery takes over both pumping and gas exchange: blood is drained from the body, pumped through an **oxygenator** — another hollow-fibre membrane, with gas on one side — and returned. **ECMO** (extracorporeal membrane oxygenation) applies the same idea for days or weeks in intensive care: **veno-venous** ECMO supports the lungs alone; **veno-arterial** ECMO supports heart and lungs. These circuits are run by **perfusionists** and specialist teams. Clotting in the circuit, bleeding from anticoagulation, air entry, and circuit disconnection are the defining hazards.

17.7. Energy in the operating theatre

17.7.1. Electrosurgery

Electrosurgery (often called **diathermy**) cuts and coagulates tissue with radiofrequency current, typically in the range of a few hundred kilohertz to a few megahertz. At those frequencies the current heats tissue without stimulating nerves and muscle — which would happen at mains frequency (Chapter 7).



In **monopolar** electrosurgery the same current that enters at a tiny active tip must leave through a large **return electrode** pad on the patient's skin. Heating depends on **current density** — current per unit area — so the tip burns and the pad, spreading the same current over a large area, does not. If the pad partly peels off, the current concentrates in what remains and the patient is burned **at the pad**, out of sight under the drapes. Modern generators monitor pad contact and stop if it deteriorates. Current can also find alternative paths — through ECG electrodes or metal contact points — which is why patient positioning in theatre includes keeping skin away from metal.

Bipolar instruments pass the current only between the two jaws of a forceps, so no pad is needed. **Vessel-sealing** devices combine pressure and controlled bipolar energy to fuse vessel walls. Ultrasonic instruments cut and coagulate by vibrating a blade tens of thousands of times per second.

All of these produce **surgical smoke**, which contains particles and chemicals and is extracted.

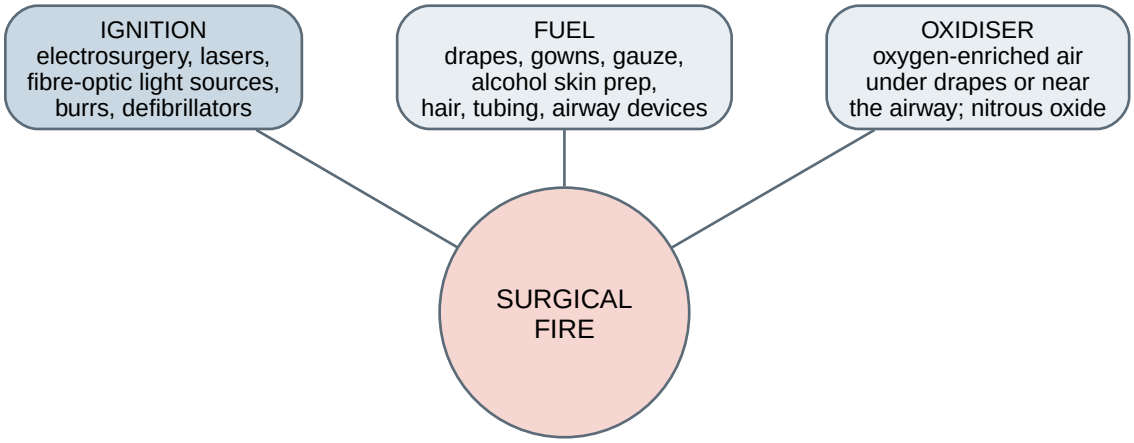
17.7.2. Lasers

Surgical lasers are chosen by how their wavelength is absorbed:

- **CO_2** (10.6 μm) — strongly absorbed by water; vaporises tissue at the surface; precise cutting.
- **Nd:YAG** (1064 nm) — penetrates deeper; coagulates.
- **Green (532 nm) lasers** — absorbed by haemoglobin; vessels, the retina, the prostate.
- **Excimer (193 nm)** — ultraviolet; removes corneal tissue a fraction of a micrometre at a time for refractive surgery (Chapter 8).

Everyone in the room wears eye protection **rated for that wavelength**; protection for one laser may be useless against another (Appendix D).

17.7.3. Surgical fire



Every theatre fire needs the three sides of the triangle at once, and each belongs to a different team member — the surgeon controls ignition, the nursing team much of the fuel, the anaesthetist the oxidiser. Fires are prevented by procedures that make all three visible: letting alcohol preparations dry, reducing oxygen concentration near the surgical site where possible, and keeping energy devices holstered when not in use. It is a pure systems problem, and a good model for any hazard that only exists when several people's separate decisions coincide.

17.7.4. Minimally invasive and robotic surgery

Laparoscopic (keyhole) surgery works through ports a few millimetres wide, with the abdomen inflated by carbon dioxide — typically at around 12–15 mmHg — to create working space, a camera for vision, and long instruments that pivot at the abdominal wall. The surgeon loses direct touch, depth perception (with a 2D camera), and natural hand–eye alignment; the training box in Chapter 16 exists to build exactly those skills.

Robotic surgery puts the instruments on robot arms controlled from a console. The robot does nothing on its own: it filters tremor, scales large hand movements to small instrument movements, adds wristed instruments that bend inside the body, and gives 3D vision. Haptic feedback — the sense of how hard you are pulling — has historically been limited, so surgeons learn to judge force visually.

Surgical instruments themselves are a vast and conservative field of precision mechanical design, almost all of it constrained by the need to be **reprocessed** (Chapter 13): every hinge, lumen, and crevice is a cleaning problem.

17.8. Radiotherapy

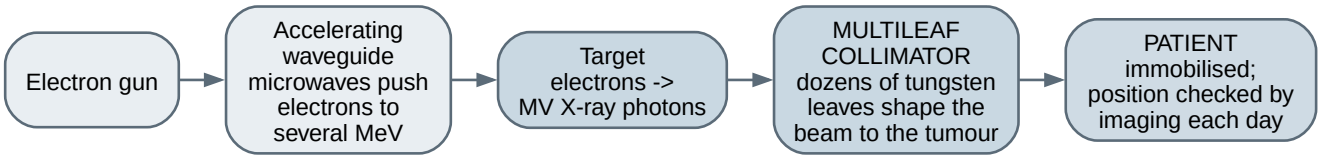
17.8.1. Killing cells with dose

Radiotherapy uses ionising radiation to damage the DNA of cancer cells so that they die when they try to divide. Dose is measured in **gray** (Gy): one joule of energy absorbed per kilogram of tissue.

The numbers are worth placing beside Chapter 16. A CT scan delivers an organ dose of the order of **tens of milligray**. A curative course of radiotherapy might deliver **60–70 Gy** to a tumour — thousands of times more — typically split into daily **fractions** of about 2 Gy over six to seven weeks, though shorter **hypofractionated** schedules with larger fractions are now common.

Fractionation exploits a biological difference: normal tissues repair sublethal damage between fractions better than many tumours do.

17.8.2. The linear accelerator



A **linear accelerator** (linac) accelerates electrons with microwaves along a waveguide and steers them into a metal target, producing a megavoltage X-ray beam — far more penetrating than diagnostic X-rays. The whole head rotates around the patient on a **gantry**. A **multileaf collimator** of motor-driven tungsten leaves shapes the beam to the outline of the tumour as seen from each angle.

Modern techniques shape the dose in three dimensions:

- **IMRT** and **VMAT** (intensity-modulated and volumetric-modulated arc therapy) move the leaves during treatment so that the dose wraps around the tumour and spares neighbouring organs.
- **Image guidance** (IGRT) uses an on-board cone-beam CT to check position before each fraction — a shift of a few millimetres can move dose from tumour to spinal cord.
- **Stereotactic** treatments deliver very high doses in a few fractions to small targets with sub-millimetre setup.
- **Proton therapy** exploits the **Bragg peak**: protons deposit most of their energy at a depth set by their energy and almost nothing beyond it, so there is little "exit dose".
- **Brachytherapy** places sealed radioactive sources inside or against the tumour; dose falls very steeply with distance from the source.

17.8.3. Immobilisation: the maker-shaped object in the room

Because millimetres matter, patients are held in reproducible positions for every fraction. For head and neck treatment, a sheet of **thermoplastic** mesh is softened in warm water and moulded over the face and shoulders to make a custom **mask** that clips to the couch. Vacuum bags, knee rests, breast boards, and arm supports do the same elsewhere.

These are exactly the objects a maker would recognise — and they sit **in the beam**. Every material in them attenuates and scatters the radiation and can change the skin dose, so radiotherapy accessories are specified, measured, and commissioned by **medical physicists** before use. A custom positioning aid — thermoformed, machined, cast, or printed — can be a genuine contribution, but only through that route. Masks are also a well-known source of claustrophobia and distress, which is a design problem as much as a clinical one.



Therac-25: the radiotherapy machine that taught software engineering.

Between 1985 and 1987, a computer-controlled radiotherapy machine delivered massive overdoses to six patients, several of whom died. Earlier models had

> independent **hardware** interlocks that physically prevented a dangerous beam configuration; the new model relied on software, which contained a timing fault (a **race condition**) triggered when an experienced operator edited the treatment settings quickly. The machine showed a cryptic error and allowed the treatment to be resumed. Reports of patients' pain were initially doubted [61].
>
> The lessons — independent safety mechanisms, never reusing software in a new context without re-verifying it, error messages that mean something, and taking incident reports seriously — are the foundation of Chapter 20 and Chapter 19.

Radiotherapy is delivered by **therapeutic radiographers** (radiation therapists), planned by **dosimetrists** and **medical physicists** with **clinical oncologists**, and every plan is checked independently before treatment.

17.9. Implants

An implant is a device intended to remain in the body. Everything Chapter 3 said about corrosion, leachables, and the finished device in its intended use applies — for decades, under millions of loading cycles.



Figure 31. Radiograph after hip replacement. The cemented metal stem sits inside the femur; the ball articulates with a polyethylene cup in the pelvis, outlined by thin wire markers because polyethylene does not show on X-ray. Image: Mikael Häggström, CC0. Credits: Appendix H

Implant	How it works	Characteristic failure modes
Hip and knee replacement	Metal femoral component (cobalt–chromium or titanium alloy), a ball of metal or ceramic, and a bearing surface of highly cross-linked **polyethylene**. Fixed with **acrylic bone cement** or pressed in with a porous or coated surface into which bone grows	**Wear particles** trigger inflammation that dissolves bone (**osteolysis**) and loosens the implant; infection; dislocation. **Stress shielding** — stiff metal carries load the bone used to, so the bone thins (Chapter 6' stiffness table). Some metal-on-metal designs released metal ions and were recalled
Coronary and vascular stents	Mesh tubes expanded inside a narrowed artery — balloon-expanded (cobalt–chromium) or **self-expanding** (superelastic nickel–titanium). **Drug-eluting** stents release a drug that reduces re-narrowing	Re-narrowing (**restenosis**); clot formation inside the stent; fracture under repeated bending
Heart valves	**Mechanical** (pyrolytic carbon leaflets) or **biological** (animal tissue on a frame). Now often delivered by catheter (**TAVI**) rather than open surgery	Mechanical: clotting — lifelong anticoagulation needed. Biological: gradual degeneration over years to a couple of decades
Neurostimulators	Implanted pulse generators and leads: spinal cord stimulation for pain, deep brain stimulation, vagus and sacral nerve stimulation (Chapter 7)	Lead migration or fracture; infection; battery depletion; MRI restrictions
Cochlear implants	Chapter 8, Chapter 7	Electrode position; device failure; MRI restrictions
Dental implants	Titanium screws into the jaw; bone bonds directly to the titanium surface (**osseointegration**)	Failure to integrate; infection around the implant

Implants are tracked: each carries a unique device identifier, many countries keep **implant registries**, and patients carry implant cards. Registries are how failing designs are spotted — statistically, years after approval, from thousands of patients — and they are a striking example of post-market surveillance (Appendix D) as a measurement system.

17.10. Drug delivery

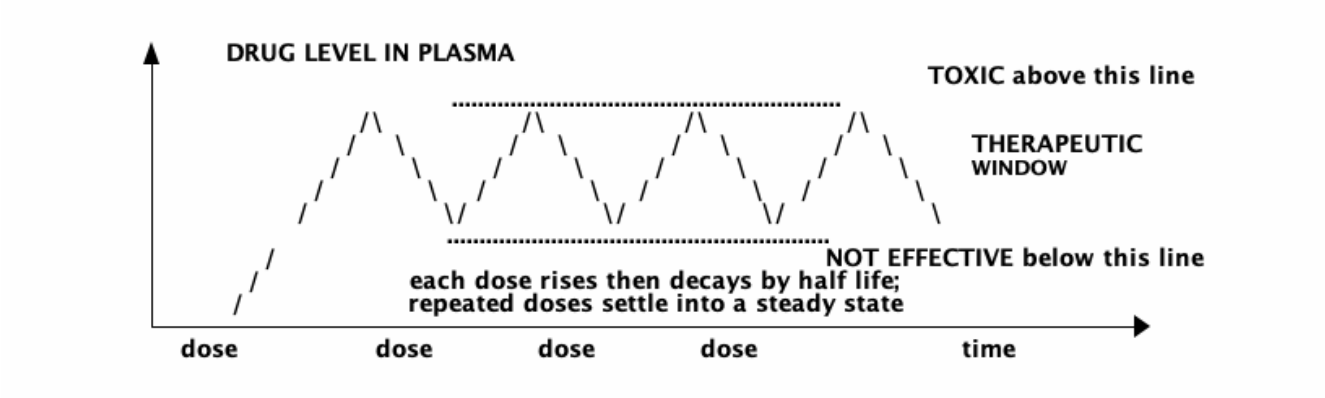
Most treatment is a drug, and many devices exist only to get a drug into the right place at the right rate. A small amount of pharmacology makes those devices make sense.

17.10.1. Pharmacokinetics in four words

Pharmacokinetics is what the body does to a drug: **absorption**, **distribution**, **metabolism**, and **excretion**.

- **Absorption** depends on the **route**. Intravenous drugs are 100% available immediately. Oral drugs must survive the stomach, cross the gut wall, and pass through the liver before reaching the circulation (**first-pass metabolism**), so only a fraction — the **bioavailability** — arrives.
- **Distribution** — into blood, tissues, fat — sets how much reaches the target.
- **Metabolism** (mostly the liver) and **excretion** (mostly the kidneys) remove it, which is why liver and kidney function change doses (Chapter 10), and why children and older adults are dosed differently (Chapter 12).

Many drugs are eliminated with **first-order kinetics**: a constant **fraction** per unit time, summarised by the **half-life**. After one half-life 50% remains; after four or five, only about 3–6%. On regular dosing, the level builds up to a **steady state** over about four to five half-lives.



The **therapeutic window** is the range between the concentration that works and the concentration that harms. For drugs with a **narrow** window — some anticoagulants, heart drugs, anti-epileptics, aminoglycoside antibiotics, insulin, opioids — small dosing or delivery errors cause harm, and these are exactly the drugs given by pump, with drug libraries, and measured by blood tests (Chapter 15).

17.10.2. Routes and the devices that serve them

Route	Device examples	Engineering and usability notes
Intravenous	Cannulas, infusion pumps, syringe drivers	Fastest and most controllable; most dangerous when wrong. Connector standards stop non-IV lines being connected to veins (Appendix D)
Subcutaneous / intramuscular	Syringes, pen injectors, **autoinjectors**, insulin pumps	Autoinjectors hide the needle and fire on pressure. Users have injected their own thumbs by holding the device the wrong way round, which drove redesign of grip, colour, and shape
Inhaled	Pressurised metered-dose inhalers, **spacers**, dry-powder inhalers, nebulisers	Technique dominates dose: studies commonly find a large share of users make errors

Route	Device examples	Engineering and usability notes
Transdermal	Patches	Only small, potent, fat-soluble drugs cross skin (Chapter 5). Heat increases absorption; some patches have metallic backings that heat in MRI (Chapter 16)
Oral and sublingual	Tablets, liquids, dose dispensers, pill organisers	Swallowing (Chapter 8); dexterity; memory; liquid measuring errors with household spoons
Intrathecal / epidural	Implanted pumps; epidural infusion sets	Tiny doses act directly on the spinal cord; wrong-route errors are catastrophic, hence dedicated connectors

The inhaler deserves a closer look because it is a pure usability problem. A pressurised metered-dose inhaler needs the user to press the canister **while** breathing in **slowly and deeply**, then hold their breath. Many people cannot coordinate that, so most of the dose lands on the back of the throat. A **spacer** — a chamber between inhaler and mouth — removes the timing requirement. A **dry-powder inhaler** needs the opposite technique: a **fast, forceful** breath to break up the powder. A person moved from one to the other without being retaught may use the right drug with the wrong breath. Every step is a row in Chapter 7' decomposition table.

17.11. Rehabilitation and assistive therapy devices

- **Prostheses** replace limbs. The **socket** — the interface between residual limb and device — determines comfort and use more than anything else, and it is Chapter 5' contact problem in its hardest form: a changing limb, sustained load, shear, sweat. Upper-limb prostheses are **body-powered** (cables and harness) or **myoelectric** (driven by surface EMG, Chapter 7); lower-limb knees range from mechanical to microprocessor-controlled. Prosthetists design and fit them.
- **Orthoses** support, align, or restrict joints — ankle-foot orthoses, wrist splints, spinal braces (Chapter 6). Orthotists design and fit them.
- **Wheelchairs and seating** are mobility devices and pressure-injury prevention devices at once (Chapter 5). Seating assessment is specialist work.
- **Negative pressure wound therapy** (Chapter 5), **TENS**, and **functional electrical stimulation** (Chapter 7) are therapeutic devices covered earlier.

This is the territory where maker volunteer networks are most visible — upper-limb assistive devices for children, adapted cutlery and tools, sewn and fabric supports, welded and machined one-off equipment, switch-adapted toys — and where Chapter 18's warning matters most: a device that bears load, restricts movement, or contacts compromised skin is prescribed clinical equipment, and good intentions do not substitute for the prosthetist's or therapist's assessment.

17.12. Engineering bridge: the numbers inside the machines

17.12.1. The lung as an RC circuit

During passive exhalation the lung behaves like a capacitor (compliance **C**) discharging through a resistor (airway resistance **R**). The **equation of motion** of the respiratory system, familiar from Chapter 9's test lung, is:

$$\mathbf{P_{airway} = V / C + R \times flow + PEEP}$$

and the time constant of emptying is $\tau = R \times C$.

Take a resistance of 10 cmH_2O per L/s and a compliance of 0.05 L per cmH_2O: $\tau = 10 \times 0.05 =$ **0.5 s**. Exhalation is about 95% complete after $3\tau = 1.5$ s.

Now raise the resistance to 30 (severe airway narrowing, as in an asthma attack): $\tau =$ **1.5 s**, and 95% emptying needs 4.5 s. If the ventilator allows only 2 s for exhalation, the lungs do not empty before the next breath arrives. Air is trapped, pressure builds breath by breath, and the heart's filling is squeezed (**auto-PEEP** or **breath stacking**). An electrical engineer would predict the problem from the time constant before a clinician saw it on the screen.

17.12.2. Defibrillator energy

A capacitor stores $E = \tfrac{1}{2} C V^2$. A 200 J shock from a 150 µF capacitor needs:

$$V = \sqrt{(2E / C)} = \sqrt{(400 / 0.000150)} \approx 1600 \text{ V}$$

Into a chest impedance of about 75 Ω, the initial current is in the region of 1600 / 75 ≈ **20 A**. The figures are illustrative — real waveforms are shaped and truncated — but the order of magnitude is the point: **a defibrillator capacitor holds a lethal charge for a long time after it is switched off**. Never open, salvage, or experiment with one.

17.12.3. Infusion arithmetic and the unit hazard

An invented example. A drug is prescribed at **0.1 microgram per kilogram per minute** for an 80 kg patient. The syringe contains **4 mg in 50 mL**.

1. Dose per minute: 0.1 × 80 = 8 µg/min.
2. Dose per hour: 8 × 60 = 480 µg/h.
3. Concentration: 4 mg / 50 mL = 4000 µg / 50 mL = 80 µg/mL.
4. Rate: 480 / 80 = **6 mL/h**.

Four steps, three unit conversions, and the drug is ordered in µg/kg/min while the pump runs in mL/h. Confuse mg with µg at step 3 and the rate is wrong by a factor of 1000; forget the "per minute" and it is wrong by 60; enter 60 instead of 6.0 and it is wrong by 10. Every one of these has happened. If you ever write software that does this, the units must be part of the values, and the conversions must be tested at their boundaries (Chapter 20, Appendix C).

17.12.4. Half-life arithmetic

Remaining amount after time **t**: **$A = A_0 \times 2^{-t / t½}$**.

A 400 mg dose of a drug with a 6-hour half-life leaves $400 \times 2^{-24/6} = 400 / 16 =$ **25 mg** after a day. Doubling the dose does not double the **duration** of effect — it adds just one half-life. That single observation explains why infusions, rather than larger single doses, are used for short-half-life drugs.

17.12.5. Dialysis clearance

A dialyser's **clearance** is the volume of blood completely cleared of a solute per minute:

$K = Q_b \times (C_{in} - C_{out}) / C_{in}$

With blood flow 300 mL/min, urea entering at 30 mmol/L and leaving at 9: $K = 300 \times 21 / 30 =$ **210 mL/min**. Over a 4-hour session $K \times t = 0.21 \times 240 \approx 50$ L. Divided by the volume the urea is distributed in (roughly total body water, about 42 L for a 70 kg adult) that gives a **Kt/V** of about 1.2 — the dimensionless "dose of dialysis" that units monitor. It is a pharmacokinetic calculation running in reverse.

17.12.6. Dose fall-off

From a point source, dose rate falls with the **inverse square** of distance: doubling the distance quarters the dose. Brachytherapy exploits this — very high dose at the tumour, low dose a few centimetres away — and radiation protection relies on it: **time, distance, shielding** (Chapter 16).

17.13. Devices, software, and training technology

This chapter, like Chapter 16, is itself a device survey. This section concentrates on where someone with engineering skills — electronics, software, machining, sheet metal, moulding and casting, textiles — can contribute without touching the therapy itself.

17.13.1. Clinical devices and the maker's entry points

Area	What surrounds the device	Maker-relevant work
Ventilation and oxygen	Circuits, masks, humidifiers, mounting, test lungs	**Test lungs and respiratory mechanics demonstrators** with adjustable R and C (Chapter 9): bellows or bags, machined orifices, springs, and a pressure-sensor readout. Mounting and cable management. Never anything in the gas path

Area	What surrounds the device	Maker-relevant work
Infusion	Poles, mounts, drug libraries, labels	Machined or sheet-metal pole mounts and tidy cable management; **syringe and line labelling** aids for teaching; drug-library analysis on anonymised logs with pharmacy
Defibrillation and resuscitation	AED cabinets, signage, trainers	**AED trainer** shells and electronics, **CPR feedback** devices (force or accelerometer sensing, firmware, display) for training — no energy storage, ever; weatherproof AED cabinets and signage (sheet metal, heaters, monitoring)
Dialysis	Chairs, tables, water monitoring	Patient comfort aids (fabric arm supports, warming covers); glassware and tubing teaching models of the dialyser and countercurrent exchange
Theatre	Positioning aids, instrument trays, holders	Instrument organisation and tray layout (Chapter 13) — stainless wire and sheet-metal trays are the established form; positioning-aid prototypes (foam, sewn or welded covers, machined frames) for review — load, pressure, and cleanability all apply
Radiotherapy	Immobilisation, positioning, bolus material	Custom positioning aids — thermoformed, machined, or vacuum-bag — and **bolus** (tissue-equivalent material placed on skin to shift dose), cast from silicone or wax or printed — a published research area, **only** with medical physics commissioning
Drug delivery	Inhalers, spacers, autoinjectors, pill organisers	**Trainer devices** and teaching aids; adaptive grips and openers (moulded, machined, or overmoulded), and electronic reminder or monitoring aids, developed with OTs and pharmacists
Rehabilitation	Prostheses, orthoses, seating	Cosmetic covers, accessory adaptors, fabric liners and bags, wheelchair accessories, and training aids alongside the prosthetist or seating team — **not** sockets or load-bearing parts

17.13.2. Software

Software	Does	Note
Ventilator and respiratory mechanics simulators	Simulate pressure, flow, and volume waveforms for set R, C, and modes	Several free and browser-based ones exist; writing your own from the equation of motion is an excellent exercise
Pharmacokinetic modelling — Open Systems Pharmacology (PK-Sim), mrgsolve, simple one-compartment models in Python	Predict drug concentration over time	One-compartment, first-order models are a few lines of code and illustrate the whole of this chapter's pharmacology
matRad	Open-source radiotherapy treatment planning for research and teaching	Lets you see optimisation, beam angles, and dose–volume histograms on public data
EGSnrc, TOPAS	Monte Carlo simulation of radiation transport	The engines behind dose calculation research
Automated insulin delivery software — community open-source projects; commercial and regulated algorithms	Adjust insulin from CGM data	The open-source community demonstrated the concept before regulated products caught up; one open project later received regulatory clearance. A remarkable — and safety-critical — example of patient-led engineering (Chapter 20)
Infusion pump drug libraries and DERS analytics	Configure limits; analyse overrides and alerts	Hospital pharmacy and clinical engineering own these; data analysis is a genuine contribution
Intensive care datasets — MIMIC (via PhysioNet)	De-identified ICU records: ventilator settings, infusions, vital signs	Credentialed access and training required; the realistic route to ICU data science (Chapter 20)
OpenSim [37]	Musculoskeletal simulation	Gait with prostheses and orthoses (Chapter 6)

17.13.3. Training and simulation

Training technology	Teaches	Form
Test lungs and ventilator simulators	Modes, alarms, troubleshooting, the RC time constant	Mechanical lungs with adjustable compliance and resistance. **Makeable as a teaching aid** from bellows, springs, machined restrictors, and a little electronics
High-fidelity patient simulators	Crisis management with real ventilators, pumps, and monitors	Manikins that breathe, have pulses, and respond to drugs via a programmed model (Chapter 21)
CPR manikins with feedback	Compression depth, rate, recoil, ventilation	Instrumented manikins. **Feedback devices are a good electronics-and-firmware project**: a force or displacement sensor, a microcontroller, and clear feedback
AED trainers	Opening, pad placement, following prompts	A dedicated trainer that **cannot** store or deliver energy, with scenario remote controls
Infusion pump training	Programming, drug libraries, troubleshooting	Real pumps in training mode with water; e-learning simulations
Dialysis and ECMO wet labs	Circuit priming, alarms, emergencies	Real circuits running water or blood substitute on a mock patient
Laparoscopic box and robotic simulators	Instrument and camera skills	Chapter 16; virtual-reality consoles for robotic surgery
Radiotherapy virtual environments	Linac setup, beam geometry, patient positioning	Full-scale virtual treatment rooms used in therapeutic radiographer training
Inhaler and autoinjector trainers	Technique, with no drug or needle	Placebo trainers, some with feedback on inhalation flow or hold time



Three good projects in this chapter.

- **A test lung with adjustable resistance and compliance** — a bag or bellows, interchangeable restrictors, a spring or elastic load — that lets a learner **see** the time constant and air trapping. It never connects to a person, and the honest limitations are easy to state.
- **A CPR or inhaler technique feedback device** on a trainer: a sensor, a microcontroller, and a clear display of rate, depth, or inhalation flow — the four-layer model of Chapter 9 applied to teaching.
- **A one-compartment pharmacokinetic simulator** that shows repeated dosing, steady state, and the therapeutic window, with units carried in the types.

17.14. Worked example: the request for an emergency ventilator

In a crisis, a local hospital asks the makerspace: "Could you build simple ventilators? Squeeze a bag with a motor." You have a capable workshop, a motor, a cam, and a stack of bag-valve-mask resuscitators. What do you say?

Start with what a ventilator has to do. Regulators issued minimum specifications for exactly this situation in 2020 [64]. Even a stripped-down version reads, in outline:

Requirement	Why it exists	Bag-squeezer status
Set and deliver tidal volume or pressure accurately	Too little suffocates; too much injures the lung	A cam squeezes a bag by a fixed stroke; delivered volume depends on bag wear, leaks, and the patient's compliance. **Not met**
Limit maximum pressure	Barotrauma; resistance can change suddenly	No pressure sensing at all. **Not met**
PEEP	Prevents alveolar collapse	Possible with a valve on the bag, but not set or monitored. **Partly**
Adjustable FiO_2	Oxygen requirement varies hugely between patients	Uncontrolled. **Not met**
Adjustable rate and inspiratory:expiratory ratio	Chapter 9; the time constant above	Motor speed only. **Partly**
Support patient-triggered breaths	Awake patients fight a machine that ignores them	**Not met**
Alarms: disconnection, high pressure, low volume, power and gas failure	Life-supporting device; staff cannot watch every breath	**None**
Single-fault safety and failure to a safe state	A motor stall or cam failure must not stop breathing silently	**Not met**
Run continuously for days to weeks; cleanable; electrically safe; documented	Real use; Chapter 13; IEC 60601-1	**Unknown**

The physics answer. A bag squeezer delivers **a volume of motion**, not a breath. Whether that becomes a safe breath depends on the patient's resistance and compliance, which change over hours. Without sensing and closed-loop control it cannot know what it is doing — and a person

connected to it cannot tell anyone.

The honest answer. "We can't responsibly build a ventilator in the time available. Here is what we can do instead": mounting hardware for the ventilators you have; **test lungs** for staff retraining; labelling and quick-reference cards; **PPE** items that sit within an established review process; and, if an engineering team with a regulatory route is building a real device, **test fixtures** and bench support for them.

The lesson. In 2020 many maker projects followed the first path; very few produced devices used on patients, and the most useful contributions were often the less visible ones. The skill this book is trying to teach is **recognising which request you have been given** — and saying so early, with reasons, and with an alternative.

17.15. Workshop exercise

1. **An RC lung on the bench.** Build the electrical analogue: a capacitor (compliance) discharging through a resistor (airway resistance), logged by a microcontroller. Measure τ and compare with R × C. Then halve the "exhalation time" with a timed switch and watch the voltage fail to return to baseline — breath stacking in electronics. Low voltage only; no mains, no large capacitors.
2. **A mechanical test lung.** With a bag or bellows, an elastic or spring load, and interchangeable restrictors, show qualitatively that a narrower restrictor slows emptying. Label it as a teaching aid with its limitations. It never connects to a person or a real ventilator without review.
3. **Infusion drills with units in the types.** Write a small program (Python with a units library is fine) that converts a dose in µg/kg/min, a weight, and a concentration into mL/h, and **refuses** ambiguous input. Write tests for the 10×, 60×, and 1000× errors described above.
4. **Half-lives on a spreadsheet.** Model repeated doses every 8 hours of a drug with a 6-hour half-life. Plot the concentration over three days and identify when steady state is reached. Then change the half-life to 24 hours and observe how long steady state takes.
5. **Inhaler task analysis.** Using a **placebo trainer** inhaler only, decompose correct technique for a metered-dose inhaler, with and without a spacer, into Chapter 7' perception, cognition, action, and feedback stages. Repeat for a dry-powder trainer. List every point where moving from one device to the other could produce an error.
6. **Fire triangle walk-through.** Write out a short invented scenario — facial surgery, oxygen by mask, alcohol skin preparation, electrosurgery — and mark each element as ignition, fuel, or oxidiser, and which team member controls it.
7. **Dialysis numbers.** Calculate clearance and Kt/V for blood flows of 250, 350, and 450 mL/min, assuming the same inlet concentration and an outlet concentration of 20%, 35%, and 50% of inlet respectively — faster blood spends less time in the dialyser, so a smaller fraction is removed. Use a 4-hour session and V = 42 L, as in the text. Comment on diminishing returns.



Do not build or salvage anything that stores high energy or delivers therapy. No defibrillator capacitors, no high-voltage supplies, no electrosurgical or laser sources, no oxygen equipment, no pumps connected to anything that could reach a person. Every exercise above is either a low-voltage analogue, a calculation, or a

passive teaching aid.

17.16. Questions

1. What makes a device "life-supporting", and name three engineering requirements that follow.
2. Why can starting positive-pressure ventilation lower blood pressure?
3. In volume-controlled ventilation, what happens if the patient's lungs suddenly stiffen? In pressure control?
4. What is PEEP, and what does it do?
5. Calculate the respiratory time constant for R = 20 $cmH_2O/(L/s)$ and C = 0.04 L/cmH_2O. How long does near-complete exhalation take, and what happens if less time is allowed?
6. Name four failure modes of infusion pumps or syringe drivers and the design response to each.
7. What can dose error reduction software catch, and what can it not?
8. Why do modern defibrillators measure transthoracic impedance?
9. Why is the AED a good example of designing a high-risk device for untrained users?
10. Explain how countercurrent flow improves dialysis efficiency.
11. Why does a partly detached electrosurgical return pad cause a burn?
12. Name the three elements of the surgical fire triangle and who in the team usually controls each.
13. Why is radiotherapy given in fractions, and how does its dose compare with a diagnostic CT?
14. Why must a radiotherapy immobilisation device or bolus be commissioned by medical physics before use?
15. A drug has a half-life of 8 hours. Roughly how long until steady state on regular dosing, and what fraction of a single dose remains after 24 hours?
16. What made the Therac-25 accidents possible, and what general lessons followed?

17.17. Answers and reasoning

1. One whose failure or misuse can kill within minutes — ventilators, critical infusions, defibrillators, pacemakers, dialysis and bypass circuits. Requirements: **single-fault safety** (safe or alarmed after any one failure); **reliable, meaningful alarms**; rigorous **usability engineering**, because the operator is part of the system. Also: continuous operation, cleanability, electrical safety, and documented verification.
2. Positive pressure in the chest compresses the great veins and reduces venous return to the heart, so cardiac output and blood pressure can fall — especially in someone already low on circulating volume.
3. **Volume control:** the machine still delivers the set volume, so airway pressure rises — risking lung injury and triggering a high-pressure alarm. **Pressure control:** the machine still delivers the set pressure, so the delivered volume falls — risking under-ventilation and triggering a low-volume alarm.
4. Positive end-expiratory pressure: pressure held in the airway at the end of exhalation by not

letting the lungs empty completely. It keeps small airways and alveoli open, improving oxygenation.

5. $\tau = 20 \times 0.04 = 0.8$ s; about $3\tau \approx 2.4$ s for ~95% emptying. If less time is allowed, gas is trapped, pressure builds with each breath (auto-PEEP), and venous return and blood pressure can fall.
6. Any four of: **free flow** — anti-free-flow clamps; **occlusion** — pressure sensing and alarm (slow at low rates; post-occlusion bolus); **air in line** — ultrasonic detection; **siphoning or height-change bolus** in syringe drivers — secure mounting, anti-siphon valves, keeping pump near patient height; **start-up delay** — priming and purging; **programming error** — drug libraries, hard and soft limits, barcode or electronic order transfer.
7. It catches doses outside the configured limits for that drug and care area — gross errors such as ten-fold mis-keying. It cannot catch a wrong drug selected, a wrong patient, an error within the allowed range, or programming outside the library.
8. The current that flows — and therefore the effect on the heart — depends on the chest's impedance. Measuring it lets the device adjust voltage or pulse duration to deliver an effective dose of current across a wide range of patients.
9. It is designed so that a person with no training can use it correctly under stress: it switches on when opened, gives voice and visual prompts, shows pad placement, decides itself whether a shock is needed, and refuses to shock non-shockable rhythms. Every stage of perception, cognition, action, and feedback is designed for, so safety is achieved by design rather than training.
10. Blood and dialysate flow in opposite directions, so blood meets progressively fresher dialysate as it is cleaned. A concentration gradient is maintained along the whole membrane, instead of the two streams equalising part-way along, as they would in parallel flow.
11. Heating depends on current density. The same current that cuts at the active tip must leave through the pad; if the pad's contact area shrinks, current density under the remaining contact rises until it heats the skin.
12. **Ignition** (electrosurgery, lasers, light sources) — usually the surgeon. **Fuel** (drapes, gauze, alcohol preparation) — largely the nursing and scrub team. **Oxidiser** (oxygen-enriched atmosphere, nitrous oxide) — the anaesthetist.
13. Normal tissues repair sublethal radiation damage between fractions better than many tumours, so splitting the dose spares normal tissue. A curative course is typically 60–70 Gy; a CT gives organ doses in the tens of milligray — thousands of times less.
14. It sits in the beam. Its material attenuates and scatters radiation, changing the dose to the tumour and the skin, and its geometry determines positional reproducibility. Those effects must be measured and accounted for in the plan by people qualified to do it.
15. About four to five half-lives — 32 to 40 hours. After 24 hours (three half-lives), 1/8 — 12.5% — of a single dose remains.
16. A software race condition, triggered by rapid editing, could configure a high-power beam without the beam-flattening hardware in place; earlier hardware interlocks had been removed and reliance put on software; the error message was cryptic and treatment could be resumed; and incident reports were initially disbelieved. Lessons: independent safety mechanisms; re-verify reused software in each new context; meaningful error messages; and treat incident reports as data.

17.18. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Test lung and breathing mechanics — the ventilator equation of motion, time constants and breath stacking.
- Half-lives, decay and steady state — repeated dosing to steady state, with invented drugs.
- Unit-safe clinical arithmetic — the infusion-rate error drill (10×, 60×, 1000×).
- Diffusion, dialysis and clearance — dialyser clearance and Kt/V.
- Capacitors — defibrillator energy, $E = \frac{1}{2}CV^2$.
- Ultrasound, Doppler and dose numbers — attenuation and inverse square.

17.19. Video companion and reading

- **Medication safety:** the Institute for Safe Medication Practices [63] publishes accessible analyses of real medication and pump errors. Read three; for each, find the step where a design change would have stopped it.
- **The classic software-safety case:** Leveson and Turner's investigation of the Therac-25 accidents [61]. Long, readable, and still the best single document on why safety-critical software is different.
- **Alarm safety:** the sentinel event alert on medical device alarm safety [62].
- **What a ventilator specification looks like:** the emergency ventilator specification published in 2020 [64]. Read it alongside the worked example above.
- **Resuscitation:** current resuscitation guidelines [65], for how defibrillation fits into the chain of survival. Learn CPR and AED use on a recognised course; it is the one piece of this chapter everyone should be able to do.
- **Patient-level explanations:** RadiologyInfo [43] covers radiotherapy and interventional procedures alongside imaging.
- **Core reading:** OpenStax [1] covers the physiology these devices replace — 19.2 Cardiac Muscle and Electrical Activity, 22.3 The Process of Breathing, and the kidney anatomy and urine formation sections of chapter 25 — rather than the devices themselves.

17.20. Exit criteria

- ▢ For ventilators, infusion pumps, defibrillators, dialysis, electrosurgery, and radiotherapy, state in two sentences each the physical principle, what the device is for, and its characteristic hazard.
- ▢ Explain the difference between volume and pressure control, and calculate a respiratory time constant and its consequence for exhalation time.
- ▢ Work an infusion rate calculation from a weight-based dose and show where each of the 10×, 60×, and 1000× errors would enter.

- ▢ Calculate defibrillator capacitor voltage from energy and capacitance, and state why you would never handle one.
- ▢ Model repeated dosing to steady state and identify the therapeutic window on your plot.
- ▢ Produce a requirements table for one life-supporting device and mark which requirements a naive maker design meets — and write the "here is what we can do instead" answer.
- ▢ Complete a recognised CPR and AED course. (This is the one practical skill in the chapter that applies to a person.)

17.21. What this chapter does not establish

- **You cannot set, operate, or troubleshoot any of these devices on a person.** Each is operated by trained, often registered professionals — intensivists, anaesthetists, nurses, respiratory therapists, perfusionists, renal teams, surgeons, therapeutic radiographers, medical physicists, prosthetists — under clinical governance.
- **You cannot calculate or check doses for real patients.** The pharmacology and infusion arithmetic here explain why errors happen; prescribing, dispensing, and administration belong to clinicians and pharmacists with their own checks.
- **You cannot build therapy devices or accessories in the therapy path.** Anything in the gas path, fluid path, energy path, beam, or body is Tier 3 (Appendix I). Test fixtures, trainers, and teaching aids are the legitimate route, and even those need review.
- **The modes, energies, and doses quoted are orientation figures**, chosen to make the physics concrete. Real settings are individual clinical decisions and change with evidence.
- **Knowing the incidents is not the same as preventing them.** The lessons of alarm fatigue, unit errors, and Therac-25 are applied through the processes of Chapter 19 and Chapter 20, with independent review — not by individual vigilance.

Chapter 18. Discovering a clinical need

Weeks 21–22 · approximately 8–12 hours · but begin outreach in week 1

18.1. Why this chapter exists

A request arrives: **"Could you make us a holder for these?"**

Or: **"Could you write us an app for this?" "Could you put a sensor on it so it beeps?" "Could you sew a cover for it?" "Could you weld up a frame like the old one?"** Requests to someone known for making things almost always arrive shaped like the requester's guess at what that person can make.

Each of those sentences contains a solution and conceals a need. Take the first. Perhaps the educator needs a part to stay visible. Perhaps they need a free hand. Perhaps things keep getting lost. Perhaps setup takes too long between sessions. Each of those implies a different design, and three of them are not a holder.

The skill this chapter teaches is to hear a proposed solution, set it gently aside, and find the problem underneath — without being tiresome about it, and without substituting your own guess for their expertise.

It is also where you learn that **finding a collaborator takes longer than learning the biology**, which is why it should have started in week 1.

18.2. Who is who

A clinical environment contains many professions and the distinctions matter. Addressing the wrong one wastes their time and signals that you have not done the basics.

Role	What they do	What they can tell you
Nurse	Delivers and coordinates most direct care; often the person who actually uses equipment all day	How the task really happens, where workflow breaks down, what gets improvised, what the night shift does differently
Doctor / physician	Diagnosis, treatment decisions, procedures	Clinical context, why a thing is done, what a result means, what is and is not negotiable
Occupational therapist (OT)	Enables participation in meaningful daily activities; adaptive equipment; function in context	**The key partner for most assistive technology.** Functional assessment, what the person actually wants to do, what has already been tried
Physiotherapist / physical therapist (PT)	Movement, strength, mobility, rehabilitation	Range of motion, strength, safe loading, movement quality, progression

Role	What they do	What they can tell you
Speech and language therapist	Communication and swallowing	Communication aids; feeding and swallowing safety
Prosthetist / orthotist	Designs and fits prostheses and orthoses	Interface design, casting and fitting, materials — the closest professional equivalent to a skilled maker in clinical practice
Simulation educator	Designs and runs training; defines learning objectives and assessment	**The key partner for training aids.** What learners must be able to do, what misconceptions arise, what the aid must show
Clinical engineer / biomedical engineering technologist	Maintains, tests, and manages medical equipment in the institution	Local technical standards, electrical safety, asset management, what breaks, who approves what
Infection prevention and control (IPC)	Policy and practice for preventing transmission	**Must review anything shared or reusable.** Cleaning process, material compatibility, acceptance into the environment
Pharmacist	Medicines	Anything touching medication handling, labelling, or delivery
Procurement / supply chain	Buys things	What already exists, what it costs, why the existing product was chosen, what the replacement cycle is
Information governance / data protection	Patient data	**Must approve anything involving real data, photography, or records**
Quality / regulatory affairs	Compliance, risk, device management	Whether and how a locally made item can be used at all
Research and development / innovation office	Brokers external collaboration	Often your actual route in
The intended user	The person who will use it — patient, resident, family carer, or staff member	**The only expert on their own experience, priorities, context, and what "better" means**



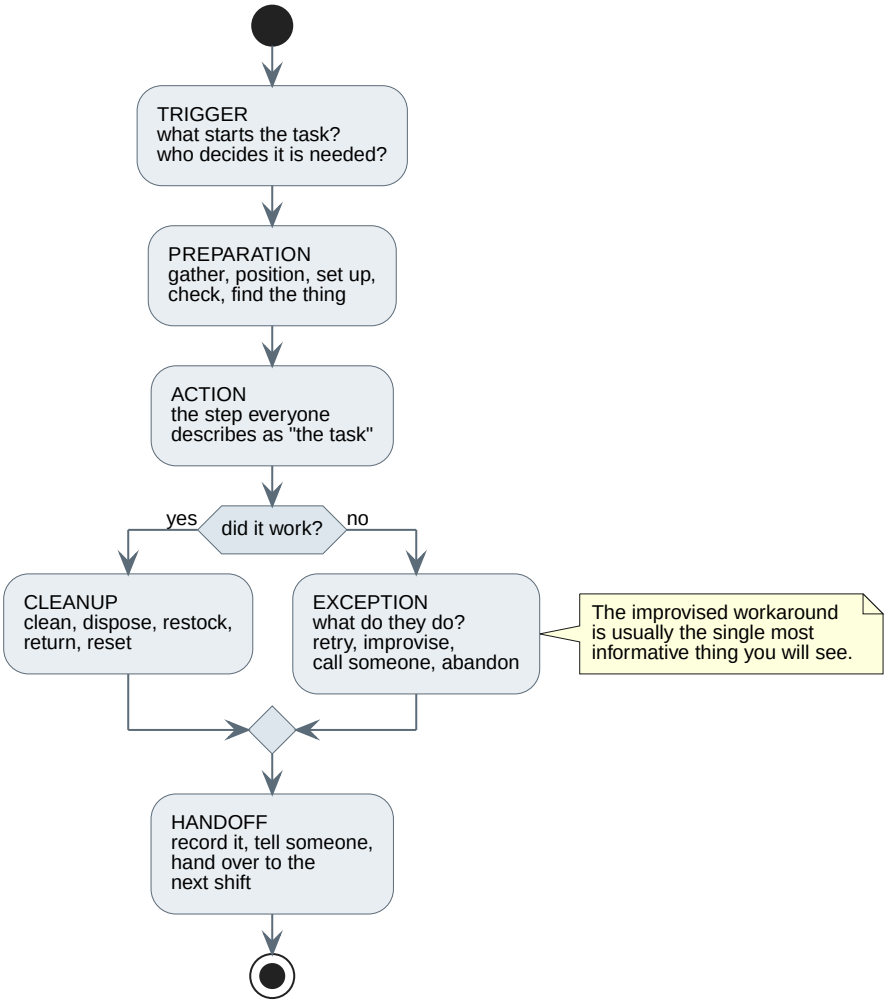
Two inputs, not one. The clinician brings professional assessment. The user brings lived experience, priorities, and context. They are not interchangeable and neither is sufficient alone.

A therapist can tell you a person has limited pinch strength. Only the person can tell you that the thing they actually care about is being able to fasten their own

coat before their grandchild arrives, and that a device requiring two hands to put on defeats the purpose entirely.

18.3. Start with the task as it is performed

18.3.1. Map the whole task, not the difficult moment



A newcomer watches the **action** step and designs for it. The difficulty is almost always somewhere else: in preparation (it takes four minutes to find), in cleanup (nobody reprocesses it so it is never available), in the exception path (it fails one time in five and the recovery is what costs the time), or in handoff.

For each stage ask: what happens, who does it, how long, how often, what can go wrong, what do they do when it does, and **what have they already improvised**?



Look for the tape. The improvised workaround — tape, an elastic band, a cardboard wedge, a bent paperclip, a second person holding something — is a solution someone already built under real constraints. It tells you the real requirement, the real environment, and the acceptable cost. It is more informative than any amount of asking.

18.3.2. Interviewing

The core move is to ask for **a specific recent instance** rather than a general opinion.

Instead of	Ask	Because
"Would this help?"	"Tell me about the last time this happened."	The first invites politeness; the second produces an account you can analyse
"Is the current one bad?"	"Walk me through how you did it this morning."	Concrete, recent, and recallable in detail
"Do you need a holder?" / "Would an app help?"	"What makes this part of the job difficult?"	Does not pre-commit to a solution
"How often does it happen?"	"When did it last happen? And before that?"	People estimate frequency badly and recall instances well
"Would you use it?"	"What do you do now instead?"	The current workaround is the real baseline your idea must beat
"What do you want?"	"What would a useful improvement let you do that you can't do now?"	Keeps the focus on outcome, not artefact

Then — and this is the step people skip — **summarise back and invite correction**. "So if I've understood: the problem is X, it happens about Y, you currently do Z, and what would help is being able to W. Have I got that wrong?" People correct a wrong summary far more readily than they volunteer a right one.



Observation and interviews happen through the host's authorised process. Always.

That means their arrangements for: access and supervision, consent, confidentiality, photography and recording, data handling and storage, and who may be present during care. These are not formalities. In many jurisdictions they are legal requirements, and breaching them can end a collaboration permanently and harm the person whose information was exposed.

Practical rules for your notes:

- Record **no patient identifiers**. Not names, not dates of birth, not record numbers, not incidental details that in combination identify someone.
- Photograph nothing without explicit permission through the host's process. Ask about **what is in the background** as well as the subject.
- If you are unsure whether something is permitted, it is not — until you have asked.
- For your first portfolio, **fictional scenarios and non-identifiable notes are usually entirely sufficient** to demonstrate your process.

18.4. Write a need without choosing the answer

A good need statement names **who**, **what difficulty**, and **what outcome** — and stops. It does not name a technology, a material, or a mechanism.

> **A way to] *to] *that [[overcomes what constraint]]**.

Compare:

	Statement	Problem
✗	"We need a motorised elbow model with an app."	Commits to actuation, electronics, and software before any of them is justified. Rules out every simpler answer
✗	"Learners don't understand the elbow."	Too broad to design against or to assess. What specifically, and how would you know?
✗	"Nurses need an app."	Names a technology and no problem at all
✔	"Beginning learners need a way to distinguish muscle force **generation** from tendon force **transmission** during an elbow movement demonstration."	Names the population, the specific difficulty, and the desired understanding. Many solutions remain open
✔	"A way for a person with limited pinch strength to open their own medication packaging independently, without requiring a second person."	Names who, outcome, and constraint. Says nothing about how

18.4.1. Then gather the constraints

The need says what. The constraints say what will and will not survive contact with reality, and they are **information to verify** rather than assumptions to bury in your CAD.

Constraint area	Questions
Environment	Where is it used? Lighting, noise, space, surfaces, water, interruptions, who else is present
Users	Who, how many, trained or not, turnover, gloved hands, one hand free or two
Frequency and duration	How often, for how long, how many cycles over its life
Setup and reset	How long is available between uses? Is there a gap at all?

Constraint area	Questions
Portability and storage	Carried between rooms? Stored where? Does it have to fit a specific space or trolley?
Power	Mains, battery, none? Who charges it? What happens when it fails?
Cleaning	Which process, by whom, how often (Chapter 13)
Maintenance and consumables	Who maintains it? What wears out? Where do replacements come from in two years?
Existing solutions	What is on the market? Why was it rejected — cost, availability, fit, policy? Have they tried it?
Cost and procurement	Who pays? Does it have to go through a purchasing process? What is the realistic ceiling?
Training	Who teaches people to use it? Are there instructions? Who maintains them?
Ownership	**Who owns and is responsible for this after you leave?** The most commonly skipped question, and the one that determines whether the project survives

18.5. Separate evidence from interpretation

This is a discipline, and it is visible in your notes or it does not exist.

Record type	Example	What to do next
Observation	"A learner pointed at the cord and called it a muscle."	Record the task and the explanation they had been given. Facts only
Interpretation	"The label may be unclear."	Generate at least one alternative explanation — perhaps the demonstration, or the sequence, or a prior lesson
Assumption	"A coloured sleeve will fix it."	Mark it as untested. Build a cheap comparison and check understanding
Requirement	"The aid must let a learner distinguish the force source from the transmission element."	Agree an observable task that would demonstrate it (Chapter 19)



Do not turn one observation into "all users need this."

Early observations are excellent at **generating** the next question and poor at **establishing** a general fact. One learner's confusion tells you a confusion is possible, is worth investigating, and has a plausible mechanism. It does not tell you how common it is, whether your explanation is the right one, or whether your fix works.

Write "one learner, on one occasion, said X" in your notes. Later, when you are

> tempted to write "learners find this confusing", the earlier note will keep you honest.

18.6. A sample conversation

Speaker	What they said
Educator	"We need a better elbow model."
Maker	"What's difficult with the one you use now?"
Educator	"Learners watch me pull the cable, and they think the cable is a muscle."
Maker	"So the main issue is understanding where the force comes from, and what the tendon's actually doing?"
Educator	"Yes. They can name the bones fine. That relationship is what's unclear."
Maker	"When did that last come up?"
Educator	"Last Tuesday. Two out of six described the cord as contracting when I asked them to explain it back."
Maker	"And what do you do when that happens?"
Educator	"I explain it again, verbally. It usually lands, but it costs me five minutes I don't have, and I'm not sure it sticks."

Look at what that exchange produced in six questions:

- The **opening request** ("a better model") was a solution with an unstated problem.
- The **actual difficulty** is a specific misconception about force generation versus transmission.
- A **baseline**: roughly a third of learners, in one observed session.
- The **current workaround**: a verbal re-explanation.
- Its **cost**: five minutes of scarce teaching time, with uncertain retention.
- A **constraint you did not ask for**: time is the scarce resource.

And crucially, it has **ruled out** a plausible but wrong direction. A more anatomically detailed shell — the obvious "better model" — would not address a misconception about **which element generates force**. Separating the force input, the transmission element, and the moving part visually might. That is a completely different design, and it is probably simpler and cheaper.

18.7. Engineering bridge: filtering needs

If you ever have more candidate needs than time — a good problem — you need a way to choose. The Stanford Biodesign process [14] formalises this as **need filtering**. A workable lightweight version:

Criterion	Ask
Is the need real?	Have at least two independent people described the same difficulty? Is there an improvised workaround?
Does it matter?	What is the consequence of not solving it — time, safety, dignity, cost, effectiveness? How often does it occur?
Is the outcome observable?	Could you tell whether it had been improved? If not, you cannot evaluate anything
Can you scope it?	Can a meaningful version be attempted in your available time, with your skills and your supervision?
Is it the right risk class?	Does it stay in territory appropriate for a supervised learning project (Chapter 5)? Walk away from anything invasive, life-supporting, or decision-making
Is there a sponsor?	Is there a named person who wants this, will review it, and will own it afterwards?
Does something already exist?	If a product exists and was rejected, understand **why** before rebuilding it. The reason is usually cost, availability, or policy — none of which you fix by making one yourself, by any method
Is there a route to use?	Could the result actually be used in that setting, or is there an institutional reason it never could?



The honest answer to the last two is often "no", and that is useful information, not failure.

A project you decline because an existing product solves it, or because there is no route to deployment, is a project you did not spend three months on. Recognising that early is a professional skill. Say so plainly to the collaborator — they will usually agree, and they will respect the reasoning.

18.7.1. Intellectual property, briefly

It will come up, so know the shape of it. If the work is done with or at an institution, **ask at the start** who owns what. Universities and hospitals typically have IP policies that apply to work done using their facilities or with their staff. Clarify before you build, not after it works.

For most volunteer maker projects the honest and simplest answer is an open licence and a clear statement that the host owns the deployment decision. Say it explicitly in the brief so nobody discovers a disagreement later.

18.8. Worked example: from request to brief

What arrived: "Can you make us a holder for these?"

What the conversation found:

- The items are set down on a trolley during a procedure and sometimes roll off.

- It happens maybe twice a week. Nobody is harmed; the item is re-fetched, which takes about ninety seconds and breaks concentration.
- Current workaround: a folded blue roll wedge, remade each time.
- There **is** a commercial tray, but it does not fit this trolley and the department cannot justify replacing the trolleys.
- Three different staff members use the trolley; each arranges it differently.
- It must be cleanable between uses with the department's existing wipes.
- Nobody has time to assemble anything between cases.

Need statement: "A way to keep these items stable and reachable on an existing trolley during a procedure, cleanable with the department's current process, requiring no setup time."

Notice what it does **not** say: nothing about a holder, nothing about a manufacturing method, nothing about geometry. "Stable and reachable" might be:

- a **change to how the trolley is laid out**, agreed among the three users — no making at all;
- a **non-slip silicone mat** — bought, or cast to fit;
- a **bent and welded stainless wire rack** or a **folded sheet-metal tray** clipped to the trolley rail — the established form for equipment that is wiped many times a day;
- a **machined or thermoformed insert** with pockets sized to the items;
- a **moulded or printed tray**, if the quantity and cleanability case supports it.

You now know enough to compare those honestly, and the comparison — which is mostly about cleanability, durability, and who replaces it — belongs to Chapter 14.

Requirements that fall out, each with how it would be checked:

#	Requirement	How it would be checked
1	Items remain in place when the trolley is moved at normal speed over a door threshold	Observed trial with the actual trolley, agreed in advance
2	Requires no assembly or setup between cases	Timed observation; educator or user confirms
3	Cleanable with the department's existing wipe process, all surfaces reachable	IPC review (Chapter 13)
4	Fits the existing trolley without modifying it	Physical fit check on the actual trolley, not a drawing
5	Any of the three users can position it correctly without instruction	Each tries it unprompted; record prompts given

Pending, and recorded as pending: IPC review; sponsor confirmation; who owns and replaces it after the project.

18.9. Workshop exercise

1. **Rewrite a solution as questions.** Take "build an app for nurses" and write three unanswered questions about the underlying need. Then do the same for a request you have actually received in any context.
2. **Write a need statement** for the elbow model using the template, without specifying any material or mechanism.
3. **Evidence separation.** From a real brief conversation or a fictional scenario, record one observation, one interpretation, and one assumption, each labelled as such. Then write the next step for each.
4. **Task map.** Map one task end to end — trigger, preparation, action, exception, cleanup, handoff — for something you can actually observe. It need not be clinical; mapping your own workshop's "set up the lathe for a batch of parts" or "solder and test a circuit board" works and teaches the same skill.
5. **Constraint sweep.** Take your own project and answer every row of the constraints table. Mark honestly which answers are **known**, which are **assumed**, and which are **unknown**.
6. **Get the brief corrected.** Write a one-page brief and ask the collaborator to correct it **before you build anything**. If you have no collaborator yet, record this as pending and use a clearly labelled fictional brief to practise. Do not invent agreement.
7. **Outreach, this week.** Write and send one email to a potential host. See Chapter 23 for what to say. This is not an exercise you can do at the end — the reply may take a month.

18.10. Questions

1. Rewrite "build an app for nurses" as three unanswered questions about a need.
2. Why is "Tell me about the last time this happened" better than "Would this help?"
3. What makes a need statement bad if it names a technology — a printed part, an app, a sensor?
4. Give one observation, one interpretation, and one assumption from the sample conversation, and say what distinguishes them.
5. Why is an improvised workaround so informative?
6. In the sample conversation, what constraint emerged that the maker never asked about?
7. A product already exists that solves the need, but the department rejected it. What do you do?
8. Why is "who owns this after you leave?" a design question rather than an admin question?
9. One learner was confused. What may you write in your notes, and what may you not?
10. You want to photograph the workspace to understand the layout. What must happen first?
11. Name the professional you would approach for: an assistive grip aid, a training model, a shared reusable object, anything involving real patient data.
12. Why should outreach start in week 1 rather than week 21?

18.11. Answers and reasoning

1. Examples: **which task** is difficult, and at which step? **Who** performs it, in what setting, how often? **What would a useful improvement let them do** that they cannot do now? You cannot choose a meaningful piece of software until the task is understood — and the answer may not be software.
2. Because the first asks for a recallable factual account you can analyse, while the second invites a polite prediction. People are poor at predicting their own future use of a hypothetical object and good at describing what actually happened on Tuesday.
3. Because it commits to a solution and a manufacturing process before either is justified, and in doing so silently rules out every simpler or better answer — including changing the workflow, buying something, using a different technology entirely, or doing nothing.
4. **Observation:** "two of six learners described the cord as contracting." **Interpretation:** "the force source and the transmission element are visually merged." **Assumption:** "separating them visually will fix it." The distinction is that the observation is what happened, the interpretation is one explanation among several, and the assumption is an untested prediction about a fix.
5. Because it is a working solution already built under the real constraints — real environment, real time budget, real materials to hand, real acceptable cost. It encodes the actual requirement more reliably than anyone's description of it.
6. **Time.** The educator volunteered that the workaround costs five minutes they do not have. That makes setup and reset time a first-class requirement, and it was never asked about — it emerged because the maker asked what they do now.
7. Find out **why** it was rejected before building anything. The usual reasons are cost, availability, fit, or policy — and making a copy, by any method, fixes none of those, while potentially introducing a worse version of something that already works. If the reason is fit, you may have a genuine project; if it is cost, you may be about to make an unsupported substitute for a supported product.
8. Because it determines maintenance, consumables, replacement, instructions, and whether the object is still in use in a year or in a cupboard. Those drive real design decisions: part count, repairability, replaceable wear items, documentation. A design with no owner should be designed differently — or not at all.
9. You may write: **"On [date], one learner in a group of six, during [task], described the cord as contracting."** You may not write "learners find this confusing" or "this is a common misconception". One instance establishes possibility, not prevalence.
10. Permission through the host's authorised process, covering consent, confidentiality, and data handling — and explicitly covering what is in the **background** of the photograph as well as the subject. If you are unsure whether it is permitted, it is not, until you have asked.
11. Assistive grip aid → **occupational therapist** (with the user, and a physiotherapist if movement or loading is involved). Training model → **simulation educator**. Shared reusable object → **infection prevention and control**. Real patient data → **information governance**, before anything else happens.
12. Because arranging a clinical partnership routinely takes longer than the study itself. Emails go unanswered, the right person is on leave, onboarding has a process, and the first conversation

often redirects you. Starting in week 1 means the partnership matures in parallel with your learning rather than after it.

18.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Project worksheets — the brief, evidence separation, task map and constraints sweep.

18.13. Video companion and reading

- **Core resource:** the Stanford *Student Guide to Biodesign* [14]. Work through the toolkits on choosing a need area, clinical observation, interviews, and need criteria. Its accompanying videos are in the video gallery [15]; registration or login may be required.
- **Alternative access:** Stanford's Textbook and Videos page [16] links public video resources and the companion book.
- **Optional book:** *Biodesign: The Process of Innovating Medical Technologies*, 2nd edition [17]. Borrow it if you want a deeper treatment of project structure. Use current official sources for anything regulatory — books date quickly there.
- **Collaborative practice:** Tikkun Olam Makers' product development approach [19], as a model of working with the person who holds the unmet need.

18.14. Exit criteria

- ☐ Map a current task end to end and distinguish facts, assumptions, and unanswered questions in writing.
- ☐ Write a need statement that names who, what difficulty, and what outcome, and does **not** presuppose a solution.
- ☐ Answer every row of the constraints table for your project, marking each as known, assumed, or unknown.
- ☐ Record at least five requirements and how each might be assessed.
- ☐ Name the specific professional roles who must review your project and why.
- ☐ Identify a clinical or educational sponsor, and the review route for the proposed work.
- ☐ Obtain the collaborator's confirmation that your brief reflects their need — or record it explicitly as **pending** and proceed with a clearly labelled fictional brief.
- ☐ Have sent at least one real outreach message.

18.15. What this chapter does not establish

- **You cannot observe clinical work on your own initiative.** Access runs through the host's process, every time, without exception.

- **You cannot handle patient data.** Not real data, not de-identified data, not photographs, without formal information governance approval.
- **One conversation is not requirements gathering.** It is a start. Needs get refined, contradicted, and occasionally abandoned as you learn more.
- **A confirmed brief is not permission to deploy.** It means you agree on the problem. Whether the result may be used in that setting is a separate decision by other people.
- **You are not the person who decides what matters.** Your judgement about which problem is worth solving is an input. The clinician's and the user's are the ones that count.

Chapter 19. Requirements, hazards, and useful tests

Weeks 23–24 · approximately 8–12 hours

19.1. Why this chapter exists

To turn **"it seems to work"** into an explicit question, a reproducible test, a recorded result including the failures, and an honest account of what remains unknown.

That transformation is most of what separates a professional engineering culture from an enthusiastic one. It is not bureaucracy. Every element of it exists because somebody was harmed by its absence: a requirement nobody wrote down and so nobody checked; a hazard that was obvious in hindsight; a test result carried forward to a version it did not apply to; a prototype that met every dimension on its drawing and was unusable.

19.2. Intended use, users, and environment

Before any requirement can be written, three things must be stated. They constrain everything else, and leaving them implicit is the most common cause of a design that is right for a situation nobody is in.

Statement	Must say
Intended use	What the thing is for, in one sentence, including what it is **not** for. "A tabletop teaching aid demonstrating elbow flexion and the role of a tendon. Not for attachment to a person, not for clinical use, not for load bearing."
Intended users	Who operates it. Their training, familiarity, frequency of use, and relevant capabilities. "Simulation educators and first-year learners, no prior exposure to the model."
Use environment	Where and under what conditions. Lighting, noise, interruptions, gloves, wet hands, available surface, available time, storage, who else is present.



Include the unglamorous conditions. Interruptions, gloved hands, poor lighting, a cluttered surface, someone else needing the trolley, a user who has done this nine times today and is not reading the label. Human factors guidance makes this central [20]: a device is used by real people in real conditions, not by an attentive expert in a quiet room.

And include **reasonably foreseeable misuse**. Not deliberate abuse — foreseeable misuse. If a part can be assembled backwards, someone will assemble it backwards. If a cord can be pulled hard, someone will pull it hard. The question is not whether that is their fault; it is whether your design tolerates it.

19.3. Turning adjectives into requirements

"Easy to use", "durable", "accurate", "robust", "intuitive" are conversation starters. None of them can be checked.

A usable requirement states **who**, does **what**, under **what conditions**, to **what criterion** — and the criterion is agreed **before** the test, with the person who cares about it.

Vague	Workable	How it would be checked
"Easy to reset"	"After one demonstration, the intended educator can restore the starting position without tools."	The actual educator does it, unprompted. Prompts given are recorded
"Durable"	"Completes 50 handling cycles under the agreed gentle manual operation without component separation or loss of function."	Counted cycles with inspection at defined intervals; failures recorded
"Clearly labelled"	"A learner who has not seen the model can state which element represents force generation and which represents transmission, without prompting."	Ask three unfamiliar learners; record responses and any prompts
"Portable"	"Fits within 400 × 300 × 150 mm and weighs under 2 kg, carried in one hand."	Measure
"Washable" (a sewn strap or sleeve)	"After 50 wash cycles at the laundry's stated programme and temperature, no seam has separated, the hook-and-loop fastening still holds the agreed test load, and the label is legible."	Wash in the actual process, not a domestic machine; inspect seams and pull-test the fastening at defined intervals
"Strong" (a welded or brazed frame)	"Each joint carries the agreed proof load for the agreed time without permanent deformation or visible cracking."	Proof-load test, then visual inspection — and dye-penetrant inspection if the consequences of a hidden crack justify it
"Safe to charge" (a battery-powered indicator)	"Uses a cell with built-in protection, charged only by the specified charger; the cell cannot be reached without a tool; the enclosure stays below an agreed temperature during a full charge."	Inspection, plus a monitored charge with a temperature reading at the warmest point
"Reliable" (firmware)	"After power loss at any point, the device restarts into the defined safe state and reports its firmware version."	Interrupt power at varied points, scripted, with each outcome logged
"Safe"	Not a requirement. Decompose it into specific hazards and controls (below)	—

Vague	Workable	How it would be checked
"Accurate"	Accurate to what, measured how, against what reference, to what tolerance?	Only answerable once those are named



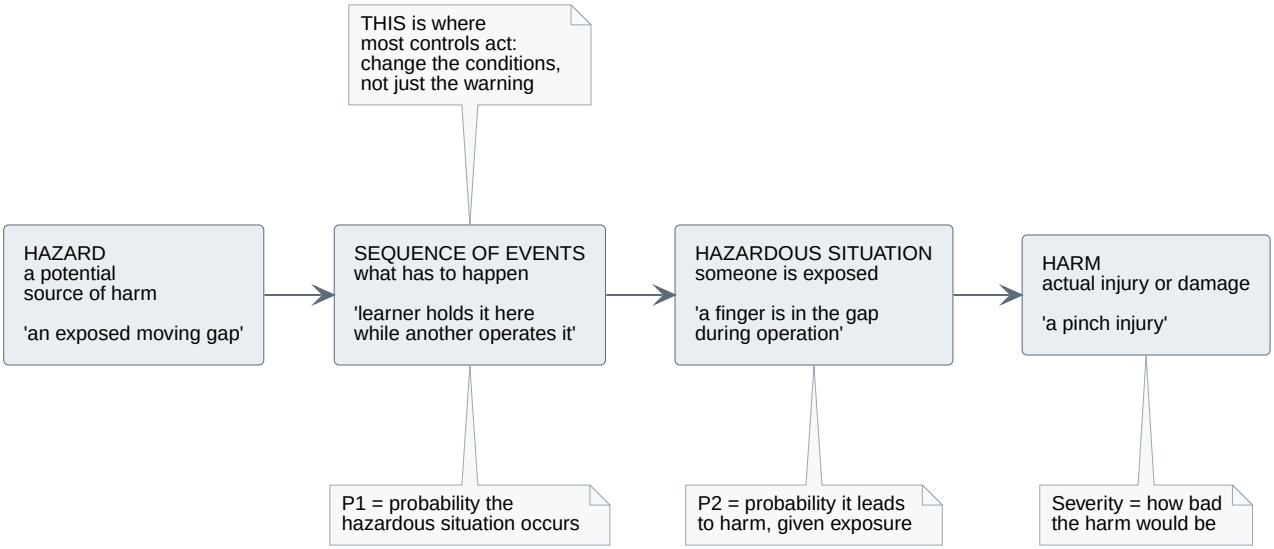
Do not select numbers because they are convenient to pass.

Choosing "10 cycles" because you are fairly sure it will manage 10 is not a requirement; it is a formality dressed as evidence. The number should come from the intended use — how many times will it actually be handled in a term? — and should be agreed with the person who will rely on it.

If no principled number is available, say so: **"50 cycles chosen as an exploratory check; not derived from expected service life; does not establish durability."** That sentence is more valuable than a fabricated justification.

19.4. Hazards, hazardous situations, and harm

The international standard for medical device risk management, ISO 14971, uses a chain that is worth adopting whatever you are building, because it directs attention at the place where you can actually intervene.



The chain matters because it shows there are **three** places to act. You can remove the hazard, prevent the sequence of events, or reduce the severity if harm occurs. A warning label acts only at the weakest point in the chain, which is why it is the last resort and not the first.

19.4.1. The risk control hierarchy

Priority	Control	Example on the elbow model
1. Inherently safe design	Eliminate the hazard by design	Change the geometry so no pinch gap exists at any point in the travel. Limit the travel mechanically so the strips cannot close

Priority	Control	Example on the elbow model
2. Protective measures	Guard, interlock, or limit	A guard over the pivot; a mechanical stop; a cord that detaches at a low force rather than transmitting a jerk
3. Information for safety	Labels, instructions, training	"Keep fingers clear of the pivot." Necessary, and **least effective**, because it depends on someone reading, remembering, and complying



The order is not a preference; it is an effectiveness ranking. A label is the control you add when you cannot eliminate or guard, not the one you add first because it is cheapest.

This is the same principle as Chapter 7' point about asymmetric connectors: **make the wrong thing impossible** rather than labelling it as wrong.

19.4.2. Every control creates new questions

This is the part people miss. Adding a guard is a design change, and design changes have consequences:

- Does the guard **hide the anatomy** the model exists to show? (For a teaching aid, that may defeat the purpose entirely.)
- Does it **obstruct the reset**, making a different requirement fail?
- Can it **loosen** during use and become a loose component — a new hazard?
- Does it create a **new crevice** that cannot be cleaned (Chapter 13)?
- Does it **add mass** at the end of a moving part?

So: implement the control, then **re-test everything the change could have affected**. Do not assume that adding a guard ends the discussion. It opens a new one.

19.4.3. Non-physical harms count

For a teaching aid, the most likely harm is not a pinch. It is **teaching something false**.

A cord labelled ambiguously produces learners who believe tendons contract. That belief then shapes how they understand injury, rehabilitation, and every subsequent explanation they hear. The harm is educational, delayed, and invisible at the bench — and it belongs in your hazard analysis.

The same logic applies to software: a misleading display, a value whose staleness is not shown, a label that implies a judgement you were not qualified to make (Chapter 9).

19.4.4. A lightweight hazard table

Hazard	Sequence of events	Hazardous situation	Possible harm	Control, and how its effectiveness is checked
Exposed moving gap at the pivot	Learner holds the model here while another operates it	Finger in the gap during operation	Pinch injury	Limit travel so the gap cannot close below 20 mm; verify by measuring the gap at every point in the travel, not just at rest
Cord under tension can release suddenly	Cord is pulled hard; attachment fails	Component springs back toward a face	Impact injury	Low-force attachment that detaches before storing significant energy; verify the release force by measurement
Card tears at the attachment	Repeated handling concentrates load locally	Model fails mid-demonstration	Loss of function; loose fragment	Reinforce and distribute the attachment; re-test the handling cycles after the change
Ambiguous cord labelling	Learner observes the demonstration and forms an explanation	Learner concludes the tendon generates force	**Educational harm: a durable misconception**	Redundant labelling (Chapter 7); verify by asking unfamiliar learners to explain it back, unprompted
Reassembled incorrectly after cleaning	Parts are separable and symmetrical	Model operates in an unintended configuration	Misleading demonstration; mechanical failure	Asymmetric geometry so wrong assembly is impossible; verify by asking someone to reassemble it from parts



On probability. Formal risk management quantifies severity and probability and combines them. For a learning project, **do not invent numbers you cannot support**. Rank hazards by judgement, state that the ranking is judgement, and focus your effort on control and verification rather than on a risk score that has the appearance of rigour and none of the substance.

> A made-up probability is worse than no probability, because it looks like evidence.

19.5. Designing a test around one question

A test answers **one question**. A test that answers five questions badly is worse than three tests that each answer one.

Element	Content
Version	Exactly which configuration was tested. See traceability below
Question	The single thing this test is intended to answer
Requirement and acceptance criterion	**Agreed before the test.** What counts as a pass
Setup	Equipment, conditions, who performed it, who observed
Procedure	Enough detail that someone else could repeat it
Result	What actually happened, including everything unexpected
Conclusion	What the result supports — stated narrowly
Limits	**What the result does not establish**
Action	What changes, and what must be re-tested as a result

19.5.1. Worked example, with an invented result

You set a prototype target of ten consecutive tool-free resets for a tabletop demonstration. On attempt seven, the cord slips out of its guide.

Honest record	Why the alternatives are wrong
"**Failed** at attempt 7 of 10. Cord slipped from the guide during reset. Guide retention needs revision."	—
☐ "Worked six times."	True and misleading. It states the successes and omits the criterion and the failure
☐ "Mostly successful."	Not a result. A criterion was defined; it was not met
☐ "Passed after fixing the guide."	Conflates two versions. The original version failed; a **different** version may pass, and must be tested as such
☐ "10/10 after retest."	Only acceptable if the retest actually happened on the new version and is recorded separately with its own entry

Then: revise the guide, and **consider what else the change affects**. Does the new guide alter the cord's line of action, and therefore the motion and the force (Chapter 6)? Does it create a crevice that complicates cleaning (Chapter 13)? Does it obscure the force path the model exists to teach (Chapter 21)? Each of those is a test that may need repeating.



Never carry a result forward to a version it was not obtained on.

This is the single most common documentation failure in amateur and professional work alike. The old result applied to the old object. When the object changes, ask what the change could have affected, and re-run those checks. A test record that cannot be tied to a specific version is not a test record.

19.6. Verification and validation

	Verification	Validation
Question	"Did we build the thing right?"	"Did we build the right thing?"
Formally	Confirms that an output meets its specified requirements	Confirms that the result meets user needs and its intended use
Example	The part measures 200 ± 1 mm as drawn	The educator can actually teach the intended concept with it
Evidence	Measurement, inspection, analysis, test against the spec	Use by the intended users, in a representative environment, doing the real task
Failure mode	Meets every dimension and is unusable	Users love it and it does not meet a requirement that matters

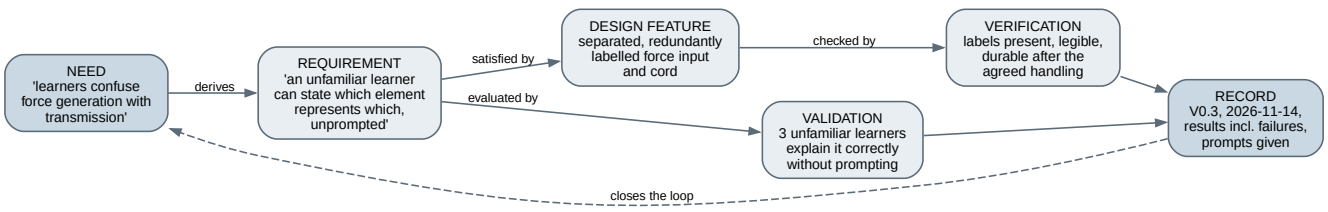


An informal tryout is not a validation study.

Handing the model to a colleague and watching them like it is useful — it generates hypotheses and finds obvious problems cheaply. It is not formal validation, which for a medical device involves representative users, representative use environments, representative tasks, defined success criteria, and an adequate number of participants.

Call the thing you did by its real name. "Informal tryout with one colleague, who suggested X" is honest and useful. "Validated with users" is neither.

19.7. Traceability: linking object to evidence



Every requirement should trace back to a need and forward to evidence. If a requirement traces to no need, ask why it exists. If it traces to no evidence, it is an intention.

19.7.1. Versioning a physical object

Software people version instinctively; physical makers often do not. The rule is the same: **a version is a configuration you can reproduce**.

Record	Detail
Version identifier	V0.1, V0.2… On the object itself, durably, surviving cleaning
Date	When this configuration came into existence
Design files and revision	CAD models, drawings, schematics and PCB layouts, sewing patterns, firmware and software source — each with a file name and commit or revision identifier
Materials	Type, grade, supplier, and batch, lot, or spool where it could matter; certificates where they exist
Process settings	For each process used, anything differing from a named default: CNC program and tooling; laser power and speed; print profile (nozzle, layer height, temperatures, infill, supports); casting mix ratio, degassing, cure time and temperature; stitch type, thread, and seam allowance; welding or brazing process and filler; the firmware build actually flashed
Assembly	Fasteners, adhesives, torque or tightness if it matters, solder and connectors, order of operations
Post-processing	Deburring, sanding, smoothing, washing, curing, passivation, annealing, painting, heat-sealing — with the agents and conditions used
What changed from the previous version, and **why**	The "why" is what stops you re-making an old mistake



Put the whole project in version control, not just the code. Git handles STEP and mesh files, drawings, schematics and PCB layouts, firmware, sewing patterns, CNC programs, images, and markdown perfectly well, and the commit history gives you the "what changed and why" for free. Tag the commit that corresponds to each physical version, and put the tag on the physical object's label.

A photograph is useful and **not sufficient**. You cannot reconstruct a cure schedule, a stitch type, a weld procedure, an infill percentage, or a firmware build from a photograph.

19.8. Engineering bridge: how the professional process is shaped

You are not implementing a quality management system for a teaching model. But knowing the shape of one tells you which of your habits will transfer and when you have reached the edge of a learning project.

Element	What it is
Design controls	A structured process — design inputs, outputs, review, verification, validation, transfer, and changes — with records. US 21 CFR 820.30; internationally ISO 13485
Design History File (DHF)	The accumulated record showing the design was developed according to that process. Your project portfolio is a DHF in miniature
Risk management file	ISO 14971: the hazard analysis, controls, verification of controls, and residual risk evaluation, maintained across the product's life
Human factors / usability engineering	IEC 62366-1 and FDA guidance [20]: use-related risk analysis, formative evaluation during design, summative evaluation of the final design
Software lifecycle	IEC 62304: software safety classification, development planning, architecture, unit implementation and verification, integration, release, and maintenance — including management of SOUP (software of unknown provenance, meaning your dependencies)
Biological evaluation	ISO 10993 family: evaluation of material–tissue interaction appropriate to the nature and duration of contact
Electrical safety	IEC 60601-1 for patient-connected electrical equipment (Chapter 7)
Post-market surveillance	Monitoring real-world performance, complaints, and adverse events after release. Development does not end at launch

See Appendix D for a fuller reference list and what each is for.



The useful insight is the shape, not the detail. Notice that every element pairs an **intention** with a **record**: decide, then write down the decision and the evidence. That is a habit you can adopt entirely today, at no cost, on a cardboard model — and it is the single thing that most distinguishes a portfolio that gets you a supervised project from one that does not.

19.8.1. Software-specific practices

If your project is software, some additional discipline that maps directly onto the above:

Practice	Why
Synthetic or fictional data only	Real data needs information governance approval (Chapter 18). Generate plausible, obviously fictional data
Visible simulation status	Unmissable and undismissable (Chapter 9)
Units in the type system	Not in a comment, not in a variable name. If your language supports a unit type or a newtype wrapper, use it. Unit confusion is a documented source of real clinical incidents
Explicit input bounds	Reject out-of-range input loudly rather than clamping silently

Practice	Why
Distinguish no-data from zero	Different states, different displays, never conflated
Error and recovery paths tested	Test the invalid input, the interrupted interaction, the lost connection, the backgrounded app. A clear error state teaches you more about robustness than another successful click
Deterministic, seeded behaviour	So a test result can be reproduced
Version visible in the UI	So a reported observation can be traced to a build
Dependency inventory	What libraries, what versions, what licences. This is the SOUP concept, and it matters outside regulation too

19.8.2. Hardware-, material-, and process-specific practices

Each way of making things has its own characteristic failures, and each deserves a line in the hazard table and a test. A short orientation; Chapter 14 covers the methods themselves.

If the project involves…	Check
Electronics	Battery power at low voltage rather than mains; protected cells and a specified charger; no exposed conductors; a defined safe state on fault or power loss; the firmware version readable on the device. Nothing connected to a person (Chapter 7)
Machined or fabricated metal	Burrs and sharp edges removed and inspected; welds and brazed joints proof-tested; the grade recorded, because corrosion and cleaning behaviour depend on it (Chapter 3)
Moulded or cast parts	Material lot and cure recorded; voids, flash, and parting lines inspected; shrinkage measured against the drawing; release-agent residue removed before use
Fabrics and sewn items	Seam strength and fastening retention after the real laundering process; edge fraying; colour fastness of anything coded by colour; what the fabric absorbs (Chapter 13)
Glass and transparent rigid parts	Edges finished; annealing for worked glass; what happens when it breaks, and whether the pieces are contained; stress cracking of acrylic and polycarbonate by cleaning agents
Printed parts	Layer orientation against the load; porosity and cleanability; the profile, material, and machine recorded (Chapter 13)

19.9. Worked example: three requirements with their limits

Requirement	Evidence to seek	What it establishes	What it does not establish
Cord pull demonstrates visible flexion in the intended setup	Operate it; review the labelled anatomy with the educator	That the intended motion occurs and the educator accepts the anatomical labelling	That all joint mechanics are represented; that learners interpret it correctly
The model can be reset without tools	The **intended educator** performs the reset, unprompted, and the prompts given are recorded	That this person could do it on this occasion	That all users can; that it remains true after wear; that it is quick enough
The version label remains readable after the agreed handling trial	Inspect and photograph the label after the trial	Legibility after that specific handling	Long-term service life; survival of a cleaning process not yet performed

Notice that the third column is always narrower than you want it to be, and the fourth is always longer. That is what honest evidence looks like.

19.10. Workshop exercise

1. **De-adjective.** Take three vague statements about your project — "strong", "clear", "easy" — and rewrite each as a question that could lead to a requirement. Then turn one into a full requirement with an acceptance criterion, and get it agreed before testing.
2. **Hazard table.** Complete the five-column hazard table for your model with at least five entries, including at least one **non-physical** harm. For each, name a control at the highest feasible level of the hierarchy and say how you would check the control works.
3. **Run a real test.** Pick one requirement. Write the full test record **before** testing — version, question, criterion, setup, procedure. Then run it. Then fill in the result, including anything unexpected, and state what it does not establish.
4. **Break something on purpose.** Change one thing — the guide, a label, a dimension — and identify every previous test that the change could invalidate. Re-run two of them.
5. **Traceability matrix.** Build a simple table: need → requirement → design feature → verification → validation → record. Find the gaps. Gaps are normal; unrecorded gaps are not.
6. **Reproduce from the record.** Give your version record to someone else — or to yourself in three weeks — and see whether the configuration can be rebuilt from it alone. Note every thing you had to remember rather than read.

19.11. Questions

1. Rewrite "the model must be strong" as a question that could lead to a requirement.
2. Describe a hazard, a sequence of events, a hazardous situation, and a harm for your project.
3. Why is a warning label the **last** risk control rather than the first?
4. Write the conclusion for the seven-of-ten reset example.
5. A model meets every dimension on its drawing but learners misunderstand it. What has been established and what has not?
6. Your test succeeds, but you guided the user through every step. What must the record say?
7. You change a cord guide after a successful test. Can you keep the old passing result?
8. You chose "10 cycles" because you were confident it would pass. What is wrong with that?
9. Give an example of a harm from a teaching model that is not a physical injury.
10. What is the difference between verification and validation, in one sentence each?
11. Why is a photograph insufficient as a version record?
12. You add a guard to eliminate a pinch hazard. Name three things you must now re-check.

19.12. Answers and reasoning

1. **"What loads and what repeated handling should it withstand, in which intended use, and how would we observe failure?"** The educator and maker must agree the relevant conditions before any meaningful test can be defined — strength is meaningless without a load case.
2. Example: **Hazard** — exposed moving gap at the pivot. **Sequence** — a learner holds the model by the pivot region while another person operates it. **Hazardous situation** — a finger is in the gap during operation. **Harm** — a pinch injury. A good answer also names a control and how its effectiveness would be checked.
3. Because it is the **least effective** control: it depends on someone noticing, reading, understanding, remembering, and complying, under real conditions including distraction and time pressure. Eliminating the hazard by design works regardless of attention.
4. "V0.1 **failed** the defined criterion: 7 of 10 consecutive tool-free resets completed. The cord slipped from its guide on attempt 7. Guide retention requires revision. Retested after revision as a new version. This does **not** establish service-life reliability."
5. **Established:** dimensional conformance — the object was built as drawn (verification). **Not established:** suitability for the learning objective (validation). Meeting a specification says nothing about whether the specification was right.
6. It must state exactly what prompts and assistance were given. The observation then supports **assisted** completion, not independent use — a different and much weaker claim. Recording prompts is the whole discipline here.
7. **No.** The old result applies to the old configuration. Assess what the change could affect — the cord’s line of action, the force required, the motion, cleanability, whether the force path is still visible — and re-run those checks on the new version as a new record.

8. The number was chosen to be passed rather than derived from the intended use, so passing it is a formality rather than evidence. The honest alternative is either to derive a number from expected use and agree it, or to state explicitly that the figure is an exploratory check not derived from service life.
9. A durable misconception — for example, learners concluding that tendons generate force. It is delayed, invisible at the bench, and propagates into everything they subsequently learn. It belongs in the hazard analysis.
10. **Verification:** did we build the thing according to its specified requirements? **Validation:** does the result meet the users' actual needs in its intended use?
11. Because it cannot record the process settings — a print profile, a cure schedule, a stitch type, a weld procedure, a CNC program, the firmware build — nor the material grade or batch, assembly torque, adhesive type, post-processing steps, or the design revision. A photograph shows what it looked like, not what it was.
12. Any three of: does the guard obscure the anatomy the model exists to show; does it obstruct the reset; can it loosen and become a loose component; does it create a crevice that cannot be cleaned; does it add mass or change the balance; does it change the motion or the required force; is the version label still visible. And re-run the affected tests.

19.13. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Project worksheets — hazard chains ranked by judgement, requirements, test and version records, and traceability.
- Batch, tolerance and shelf life — tolerances and capability.
- Statistics — analysing test records.

19.14. Video companion and reading

- **Core:** Stanford's prototyping toolkits on questions, planning, testing, and iteration [14], with the corresponding video material. Pick a **single** prototype question and apply the method to it.
- **Reference:** FDA *Applying Human Factors and Usability Engineering to Medical Devices* [20]. Concentrate on the sections about users, use environments, and tasks, and on the distinction between formative and summative evaluation.
- **Optional video:** FDA CDRH Learn [21], human factors section, for a more formal view of the professional process.
- **Context:** skim the scope sections of ISO 14971 and IEC 62366-1 if you can access them through a library (Appendix D). You are not implementing them; you are learning the shape of the process your work will eventually sit inside.



These resources describe a professional process. They are **not** a checklist that grants approval, and completing the exercises in this chapter does not make your prototype compliant with anything. What it does is make your evidence legible to

> people who work this way — which is what gets you taken seriously.

19.15. Exit criteria

- ▢ State intended use, intended users, and use environment for your project in writing, including what it is **not** for.
- ▢ Link each of three requirements to a recorded test and a recorded outcome, with the acceptance criterion agreed before the test.
- ▢ Produce a hazard table with at least five entries tracing hazard → sequence → hazardous situation → harm, including at least one non-physical harm.
- ▢ For each hazard, name a control at the highest feasible level of the hierarchy and state how the control's effectiveness is checked.
- ▢ Complete a versioned test record that includes a **failure** or an unresolved issue.
- ▢ Reproduce a tested configuration from your records alone.
- ▢ Explain why meeting a specification is not the same as meeting the user's need.
- ▢ Identify at least five plausible failure scenarios and propose a response to each.
- ▢ Explain who decides whether the project may advance to its next stage.

19.16. What this chapter does not establish

- **Your project is not compliant with any standard.** Learning the shape of ISO 14971 or IEC 62366 does not mean you have applied them. They require competence, resources, and organisational context you do not have as an individual learner.
- **A bench test is not clinical evidence.** It characterises an object under conditions you chose. Clinical suitability is established differently, by other people.
- **An informal tryout is not validation.** Say what you actually did.
- **You cannot decide whether the project proceeds.** That decision belongs to the sponsor and the institution. Your job is to give them an accurate picture, including the parts that do not flatter your work.
- **Risk ranking by judgement is judgement.** Do not dress it as quantification. A fabricated probability is worse than an acknowledged absence of one.

Chapter 20. Health data, medical software, and security

Week 25 · approximately 6–8 hours

20.1. Why this chapter exists

Software is one tool in your kit, alongside the workshop and the electronics bench — and in a hospital it is among the most consequential. Hospitals run on software — the electronic health record, the picture archive, the laboratory system, the pump's drug library, the monitor's alarm logic, the scheduling system that decides whether a scan happens this week — and much of it is worse than it should be. Someone who can build software **and** hardware, and who understands the clinical context, is rare and genuinely useful.

It also arrives uninvited. Put a microcontroller in a task trainer, a sensor on an assistive aid, or a Bluetooth link in a wearable, and you have written firmware, you may be handling data about a person, and you have a device with an attack surface. Everything in this chapter applies to the code inside the things you make as much as to applications on a screen.

But health software is not ordinary software with a medical theme. Four things change, and each has its own body of regulation, standards, and hard-won lessons:

1. **Software can be a medical device.** Whether it is depends not on the code but on what you **say it is for**. Once it is, a defined lifecycle applies (IEC 62304) on top of the risk and requirements process of Chapter 19.
2. **The data are special.** Health data are among the most sensitive personal data there are, and they are protected by law in every jurisdiction you are likely to work in. De-identification is much harder than deleting the name column.
3. **Systems must interoperate**, through standards with long histories and sharp edges — HL7, FHIR, DICOM, and the terminologies that give codes and units their meaning.
4. **Security failures harm patients directly.** Ransomware has closed emergency departments and halted pathology services. A vulnerable pump or monitor is a safety problem, not an IT problem.

The habits you already have — version control, code review, automated testing, threat modelling — transfer well. What you need is the map of where they must be formalised, what the regulation asks of them, and where your instincts from other domains are wrong.



This chapter is orientation, not legal or regulatory advice. Laws and guidance cited here are summarised as they stood in 2026, and several were changing. What applies to a real project is established with the host's information governance, clinical safety, and regulatory staff — never by reading a summary, including this one.

20.2. When is software a medical device?

20.2.1. Intended use decides

A medical device, in every major framework, is defined by its **intended purpose**: what the manufacturer claims it is for. The same code can be a device or not depending on the words on the website, the labelling, and the instructions.

Consider one algorithm that analyses heart rate from a wrist sensor:

- Marketed as **"track your workouts"** — a general wellness product. Generally not a medical device.
- Marketed as **"notifies you of an irregular rhythm suggestive of atrial fibrillation"** — a medical device. Smartwatch AF notification features went through regulatory clearance for exactly this reason.
- Used inside a hospital to **decide who gets anticoagulation** — a higher-risk device still.

The intended purpose is not only what you write; regulators also look at how the product is promoted, what it is reasonably used for, and the claims in its user interface. "For educational purposes only" in a footer does not neutralise a screen that says "Diagnosis: atrial fibrillation".

20.2.2. SaMD and SiMD

Two terms you will meet constantly, from the International Medical Device Regulators Forum (IMDRF):

- **Software as a Medical Device (SaMD)** — software that **is** the device, running on general-purpose hardware: an app that grades diabetic retinopathy from photographs, an algorithm that flags intracranial haemorrhage on CT, an insulin dose calculator.
- **Software in a Medical Device (SiMD)** — software that drives or is part of a hardware device: the firmware in an infusion pump, the reconstruction software of a CT scanner (Chapter 16), the control loop of a ventilator.

Both are regulated, along different paths. Software that is **neither** — hospital scheduling, billing, staff rostering, most of the EHR's administrative functions — is usually not a device, though it may still be safety-critical and is still covered by data protection law. In England, clinical risk management standards for health IT (DCB0129 for manufacturers, DCB0160 for deploying organisations) apply to health IT systems whether or not they are devices, and require a named **Clinical Safety Officer** — a useful reminder that "not a device" does not mean "not a safety question".

20.2.3. The IMDRF risk framework

The IMDRF SaMD framework [55] classifies software by two questions: **how serious is the healthcare situation**, and **how much does the clinician or patient rely on the software's output?**

State of the healthcare situation	Treat or diagnose	Drive clinical management	Inform clinical management
Critical	IV	III	II
Serious	III	II	I
Non-serious	II	I	I

Category IV is the highest risk. "Inform" means the output is one input among many to a clinician's decision; "drive" means it guides the next step; "treat or diagnose" means the output itself is the diagnosis or therapy decision. Regulators do not all use this table directly, but most risk-based software classification echoes its logic, and it is the fastest way to reason about where a project sits.

20.2.4. The shape of the main regimes

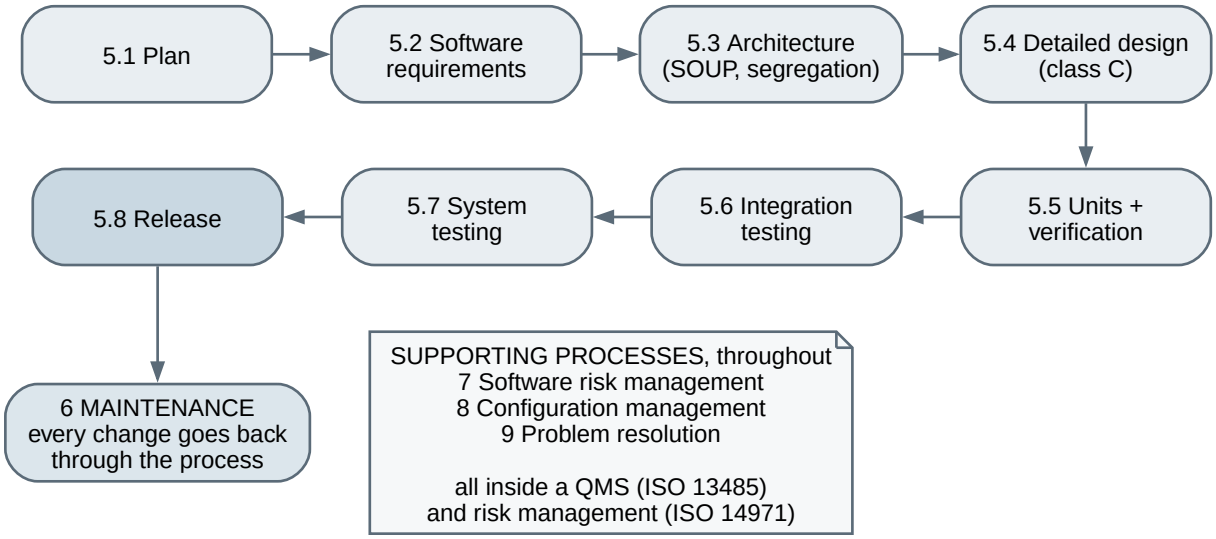
Jurisdiction	How software is caught, in outline
United States	The FDA regulates "device software functions". The 21st Century Cures Act (2016) excluded certain functions from the device definition — administrative support, general wellness, electronic records, and some **clinical decision support (CDS)** that lets a clinician independently review the basis of the recommendation [56]. FDA applies "enforcement discretion" to some low-risk functions. Premarket documentation scales with risk ("basic" or "enhanced" documentation, since 2023 guidance). From February 2026, the device quality system regulation (21 CFR 820) incorporates ISO 13485 by reference
European Union	The Medical Device Regulation (EU) 2017/745 (Appendix D). **Rule 11** classifies software that informs diagnostic or therapeutic decisions as at least Class IIa, rising to IIb or III where a wrong decision could cause serious deterioration or death; software monitoring vital physiological parameters where variation could cause immediate danger is IIb. In practice very little clinical software is Class I under the MDR. Guidance document MDCG 2019-11 [58] explains qualification and classification. The EU AI Act (Regulation 2024/1689) adds obligations for AI in devices that need notified-body assessment, phasing in to 2027
United Kingdom	The MHRA regulates software as a medical device under the UK Medical Devices Regulations, with a programme of reform for software and AI. NHS England adds DCB0129 and DCB0160 clinical safety standards and the Digital Technology Assessment Criteria (DTAC) for products used in the NHS



Clinical decision support is the grey zone a software project is most likely to wander into. A rule that says "consider sepsis screening" to a clinician who can see why it fired, and can easily ignore it, is treated very differently from one that says "start antibiotics" or one that patients see directly. The difference is not the code; it is who relies on the output, how much, and whether they can check its basis. Settle that with the host **before** you write the rule, because changing it later changes the regulatory question.

20.3. The software lifecycle: IEC 62304

Chapter 19 introduced risk management (ISO 14971) and usability engineering (IEC 62366-1). IEC 62304 is the third member of that family: the required **process** for developing and maintaining medical device software. It does not tell you how to write code. It tells you what you must plan, do, and record.



20.3.1. Safety classes

Each software system — and, with justification, each item within it — is assigned a **software safety class** according to the harm a failure could contribute to, **after** accounting for risk controls outside the software:

Class	If the software fails, the worst outcome is	Process burden
A	No injury or damage to health is possible	Planning, requirements, release, configuration management, problem resolution
B	Non-serious injury is possible	Adds architecture, unit verification, integration and system testing in defined depth
C	Death or serious injury is possible	Adds detailed design and the most rigorous verification of each unit

Two consequences follow, and both are architectural.

First, a **hardware** risk control outside the software — a mechanical pressure-relief valve, an independent hardware watchdog, a physical interlock — can lower the software's class, because the software can no longer cause the harm on its own. The Therac-25 story below is what happens when such a control is removed and the software is trusted to replace it.

Second, **segregation** lets you confine class C to a small, well-defined part of the system. Keep the dose calculation in a small, isolated, heavily verified component, and the user interface around it can be class B — provided the segregation is real and documented, not a folder name.

20.3.2. SOUP: your dependencies, formally

SOUP — software of unknown provenance — is any software item you did not develop under your 62304 process: the operating system, the database, the web framework, the npm package with eleven transitive dependencies. You must identify each one with its version, state the functional and performance requirements you rely on it for, specify the hardware and software it needs, and evaluate its **published anomalies** — known bugs — for effects on safety.

This is the regulatory form of something you already know: you are responsible for what you ship, including what you imported. It is also, now, a security requirement (below).

20.3.3. Configuration management and problem resolution

You must be able to say exactly what was in a release — every source file, configuration, SOUP version, and build tool — and reproduce it. Every reported problem must be recorded, investigated, assessed for safety relevance, resolved or justified, and the resolution verified. Your issue tracker and git history are most of this already, **if** you use them with that discipline: tagged releases, reproducible builds, linked issues, and no "quick fix" deployed outside the process.

20.3.4. Agile is allowed

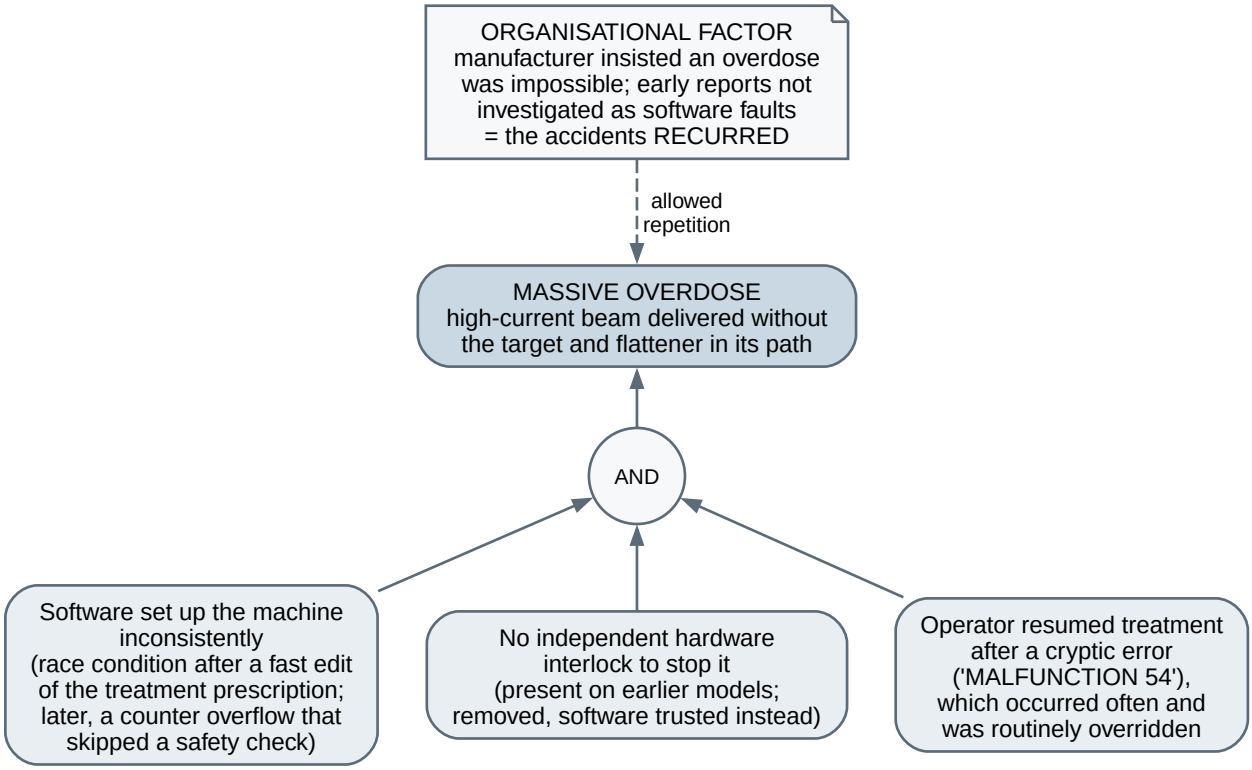
IEC 62304 specifies **activities and outputs**, not a waterfall. Iterative and agile development is acceptable and widely used; AAMI TIR45 is the standard guidance on how to map sprints, backlogs, and continuous integration onto the standard's requirements. The difference from ordinary agile is that "done" includes the records: requirement linked to test, risk control linked to verification, SOUP list updated, problem reports closed.

IEC 82304-1 sits above 62304 for **health software products** — standalone software that is a product in its own right — adding product-level requirements for validation, documentation, and accompanying information. **IEC 81001-5-1** adds the security lifecycle, covered later in this chapter.

20.4. What goes wrong: lessons written in harm

20.4.1. Therac-25

Between 1985 and 1987, the Therac-25 radiotherapy machine delivered massive radiation overdoses to at least six patients; several died. The investigation by Leveson and Turner [61] is required reading for anyone writing software that can hurt people, because almost every failure in it is a software-engineering failure rather than a coding slip.



The lessons, each of which has a modern equivalent:

- **Removing a hardware interlock and trusting software to replace it** moved the whole safety burden onto code that had not been designed or verified for it.
- **Concurrency bugs** depended on operator speed: they appeared only when an experienced operator edited the prescription quickly. Testing at normal pace did not reveal them.
- **Error messages were uninformative and frequent**, so operators learned to override them. An alarm that cries wolf trains its users to ignore it — the alarm fatigue of Chapter 9.
- **Reused code was assumed safe** because it had worked in an earlier machine — where the hardware interlocks had been masking its faults.
- **Overconfidence delayed the response.** Incident reports were met with assurances rather than investigation.

20.4.2. Infusion pumps and the arithmetic of a keypad

Infusion pumps deliver drugs, fluids, and nutrition at programmed rates (Chapter 17). Between 2005 and 2009 the FDA received tens of thousands of adverse event reports associated with them, and its 2010 improvement initiative identified **software defects** and **user-interface problems** as major contributors alongside hardware faults. Typical failure patterns:

- **Key bounce or a missed decimal point**: "1.0" entered as "10", a tenfold overdose.
- **Unit confusion**: mg versus micrograms, per hour versus per minute, per kilogram or not.
- **Silent clamping or silent acceptance** of an implausible value.
- **Mode confusion**: the user believes the pump is in one mode; it is in another.

The engineering response is **dose error reduction software**: a hospital-maintained **drug library**

with hard limits (cannot be overridden) and soft limits (can be, with a deliberate confirmation), per drug, per concentration, per care area. It works only if it is used, kept current, and not routinely bypassed — a sociotechnical control, not a purely technical one. Medication-safety practice also fixes notation: never a trailing zero ("5 mg", not "5.0 mg"), always a leading zero ("0.5 mg", not ".5 mg").

20.4.3. Units and silent conversions

Unit errors recur because they are invisible in code that passes numbers around. Glucose in mmol/L versus mg/dL differs by a factor of 18 (Appendix C); weight in pounds versus kilograms by 2.2; and outside medicine, NASA lost the Mars Climate Orbiter in 1999 to a pound-force/newton mismatch between two teams' software. The defence is structural: carry the unit with the value, convert explicitly at the boundary, and reject ambiguous input. The engineering bridge below shows how.

20.4.4. Alarms

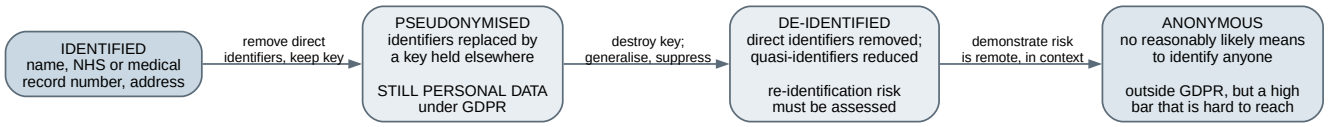
Studies of hospital monitoring have repeatedly found that the large majority of alarms — often quoted as 72–99% — are false or clinically irrelevant, and accreditation bodies have named alarm fatigue as a patient-safety priority. An alarm is a diagnostic test with a prevalence problem (Chapter 15): when the condition it detects is rare, even a good detector produces mostly false positives, and people respond accordingly. Alarm design is a safety function, not decoration: priority levels, escalation, delays and filters chosen with clinicians, and every alarm justified by an action it should prompt. IEC 60601-1-8 covers alarm systems in medical electrical equipment (Appendix D).

20.5. Health data

20.5.1. What counts

Health data are any data relating to a person's physical or mental health, including the provision of health care, that reveal information about their health status. That is wider than it looks. An appointment at a fertility clinic is health data. So is a dietary requirement on a hospital menu, a step count linked to a named person, a photograph of a wound, a voice recording that reveals a speech disorder, and a DICOM image header (Chapter 15). Genetic data reveal information about relatives who never consented to anything.

20.5.2. The identifiability spectrum



Three categories matter:

- **Direct identifiers** — name, address, record number, phone number, email, full-face photograph.
- **Quasi-identifiers** — attributes that are harmless alone but identifying in combination: date of birth, postcode, sex, admission date, rare diagnosis, occupation, ethnicity.

- **Sensitive attributes** — what you are trying to protect: the diagnosis, the test result.

Deleting direct identifiers is the easy part. In a widely cited study, Latanya Sweeney estimated that about 87% of the US population was likely to be uniquely identified by just three quasi-identifiers: five-digit ZIP code, sex, and full date of birth. A dataset with those three columns and a diagnosis is, for most people in it, an identified dataset.

Medical data add their own traps:

- **Free text** — clinical notes mention names, relatives, employers, and places.
- **Dates** — admission, procedure, and birth dates are strongly identifying; shifting them consistently per patient preserves intervals while hiding the calendar.
- **Images** — burned-in text on ultrasound and secondary-capture images; faces that can be reconstructed from head CT and MRI volumes (**defacing** tools exist for this); and the DICOM header, which may hold names, dates, institution, device serial numbers, and private vendor tags. The DICOM standard defines de-identification profiles (PS3.15 Annex E).
- **Rare conditions and small numbers** — "the only patient with condition X in this town".
- **Linkage** — a dataset safe on its own can identify people when joined with another public one, including one that does not exist yet.

20.5.3. The legal frameworks in outline

Framework	What a builder needs to know about its shape
GDPR (EU) and UK GDPR with the Data Protection Act 2018	Health, genetic, and biometric data are **special category data**: processing needs both an ordinary **lawful basis** (Article 6) and a specific **condition** (Article 9) — and consent is only one possible basis, often not the one used for direct care. **Pseudonymised data remain personal data.** High-risk processing requires a **Data Protection Impact Assessment** (DPIA) before it starts. **Data protection by design and by default** (Article 25) is a legal duty, not a slogan. Rights of access, rectification, and objection apply. The UK framework was amended in 2025; check current detail
Caldicott principles (UK health and care)	Eight principles for confidential information: justify the purpose; use it only when necessary; use the minimum necessary; access on a strict need-to-know basis; everyone understands their responsibilities; comply with the law; the duty to share can be as important as the duty to protect; inform people about how their data are used. Each organisation has a **Caldicott Guardian**

Framework	What a builder needs to know about its shape
HIPAA (US)	Applies to **covered entities** — health plans, clearinghouses, and providers conducting standard electronic transactions — and their **business associates**, who sign a business associate agreement. It protects **protected health information** (PHI). The **minimum necessary** standard limits uses and disclosures. De-identification is achieved either by **Safe Harbor** (removing 18 listed identifiers, with no actual knowledge that the remainder could identify anyone) or by **Expert Determination** [59]. A consumer health app not working for a covered entity may fall outside HIPAA entirely — and inside consumer protection and state privacy law instead

The 18 HIPAA Safe Harbor identifiers are worth reading once, because the list is a good checklist for any jurisdiction: names; geographic subdivisions smaller than a state (with a narrow exception for the first three ZIP digits of large areas); all elements of dates except the year, and ages over 89; telephone and fax numbers; email addresses; social security numbers; medical record numbers; health plan beneficiary numbers; account numbers; certificate and licence numbers; vehicle identifiers; device identifiers and serial numbers; URLs; IP addresses; biometric identifiers; full-face photographs; and any other unique identifying number, characteristic, or code.

20.5.4. Governance in practice

In a hospital, the people who decide what data you may see and how are the **information governance** (or privacy) team, often with the data protection officer, the Caldicott Guardian or privacy officer, and — for research — an ethics committee (Chapter 18). Expect questions about purpose, legal basis, minimisation, where data are stored and processed, who has access, retention, and how you will prove all of it. Expect them to take weeks.

That is why this book insists on **synthetic and public data** for learning. You can build, test, and demonstrate almost everything a health software project needs without a single real record — and arrive at governance with a working system, a data-flow diagram, and a clear, minimal request.



Never put real patient data into a tool the organisation has not approved — including a personal cloud drive, a chat assistant, a code-hosting service, a test fixture, a bug report, or a screenshot in a presentation. Data copied into an unapproved service is a breach the moment it leaves the approved boundary, whatever you intended, and it cannot be recalled.

20.6. Interoperability: making systems understand each other

A hospital may run hundreds of clinical systems. They exchange data through a small number of standards, and the meaning of the data depends as much on **codes** and **units** as on syntax.

Standard	What it is	What to know
HL7 version 2	Pipe-delimited text messages, each starting with an `MSH` header segment, sent over a simple TCP framing (MLLP)	Decades old and still carries most hospital traffic: admissions and transfers (ADT), orders (ORM/OML), results (ORU). Every site customises it; "HL7 v2 compliant" promises little
FHIR [60]	Resources (Patient, Observation, Encounter, MedicationRequest…) as JSON or XML over a REST API	The modern standard. R4 (2019) is the widely deployed release. National **profiles** (US Core, UK Core) constrain it for local use. **SMART on FHIR** adds OAuth 2.0-based app launch and authorisation, so a third-party app can run inside an EHR with the user's permissions. US regulation requires certified EHRs to offer standardised FHIR APIs
DICOM	File format **and** network protocol for medical images (Chapter 15)	Classic services (C-STORE, C-FIND, C-MOVE) between modalities and the PACS; **DICOMweb** (QIDO-RS, WADO-RS, STOW-RS) offers the same over HTTP
IHE profiles	Implementation guides combining standards for specific workflows	E.g. cross-enterprise document sharing (XDS), patient identity cross-referencing (PIX/PDQ). Tested at annual **Connectathons**

20.6.1. Codes and units carry the meaning

A number without a code and a unit is not clinical data. Four terminologies do most of the work:

- **SNOMED CT** — clinical findings, procedures, body structures; very large and hierarchical.
- **LOINC** — what was measured: "8867-4" is heart rate, "8480-6" systolic blood pressure, "2339-0" glucose (mass/volume) in blood.
- **ICD-10 and ICD-11** — diagnoses for statistics and billing; ICD-11 came into effect in 2022, with national adoption still in progress.
- **UCUM** — a formal, machine-readable grammar for units: `mm[Hg]`, `mmol/L`, `/min`, `kg`.

A heart-rate reading as a FHIR Observation, with fictional data:

```
{
  "resourceType": "Observation",
  "status": "final",
  "category": [{ "coding": [{
    "system": "http://terminology.hl7.org/CodeSystem/observation-category",
    "code": "vital-signs" }] }],
  "code": { "coding": [{
    "system": "http://loinc.org", "code": "8867-4", "display": "Heart rate" }] },
  "subject": { "reference": "Patient/fictional-0001" },
  "effectiveDateTime": "2026-10-03T09:15:00Z",
  "valueQuantity": {
    "value": 72, "unit": "beats/minute",
```

```
    "system": "http://unitsofmeasure.org", "code": "/min" }
}
```

Notice that the meaning is in the codes, not in the display text. "Heart rate" in the `display` field is for humans; the LOINC code is what a receiving system acts on. Notice too that `status`, the time, and the subject are part of the measurement. A value that is `preliminary`, from yesterday, or for a different patient is a different fact — the "stale value" problem of Chapter 9 in data form.

20.7. Security

20.7.1. Why hospitals are targets

Hospitals combine everything an attacker wants: valuable data, a critical service that must be restored quickly (and so may pay), large and flat networks, thousands of users under time pressure, and **long-lived devices** — scanners, monitors, analysers, pumps — running old operating systems that cannot easily be patched because the vendor's validation covers a specific configuration.

Some reference events:

- **WannaCry (May 2017)** spread across unpatched Windows systems worldwide. In England it disrupted a third of NHS hospital trusts; thousands of appointments and operations were cancelled and some ambulances diverted. No ransom was needed for harm: a worm reached everything that was unpatched and reachable.
- **Change Healthcare (February 2024)**, a US claims and payment clearinghouse, was hit by ransomware that disrupted pharmacy and payment processing nationally for weeks, in what became the largest health data breach reported in the US, affecting on the order of a hundred million or more people.
- **Synnovis (June 2024)**, a pathology provider to several London hospitals, was hit by ransomware that halted blood testing and transfusion services, forcing cancelled operations and emergency appeals for O-type blood donors.

The lesson is that **availability is a safety property**. A system that is confidential and intact but down is a system whose patients are not being treated.

20.7.2. What regulators now expect of device makers

- **United States.** Since March 2023, section 524B of the Food, Drug, and Cosmetic Act requires premarket submissions for **"cyber devices"** — devices with software that can connect to the internet — to include a plan to monitor and address vulnerabilities (including **coordinated vulnerability disclosure**), processes giving reasonable assurance of cybersecurity, and a **software bill of materials** (SBOM). FDA premarket cybersecurity guidance [57] describes threat modelling, security architecture views, testing, and labelling expected in submissions.
- **European Union.** The MDR's general safety and performance requirements include IT security for software and devices; MDCG 2019-16 is the guidance.
- **Standards. IEC 81001-5-1** defines a secure development lifecycle for health software, designed to sit alongside IEC 62304. The **MDS2** form (Manufacturer Disclosure Statement for Medical

Device Security) is how hospitals ask vendors standard security questions. ISO/IEC 29147 and 30111 cover vulnerability disclosure and handling.

20.7.3. Secure development, applied

Most of this is good practice you know, with three medical twists.

- **Threat model the clinical workflow, not only the code.** Who touches the system at 3 a.m. with gloves on? What happens if authentication takes 40 seconds in an emergency? Security controls that obstruct care get bypassed — shared logins, passwords on sticky notes, machines left unlocked — and the bypass becomes the real security posture.
- **Plan for patching from the start.** A device validated in a fixed configuration must still receive security updates. Design for updateability (signed updates, rollback, separation of security-relevant components) and know that regulators generally treat routine cybersecurity fixes as manageable without a new full submission — provided the manufacturer's process handles them.
- **Assume the network is hostile.** Hospital networks are large and segmented imperfectly. Authenticate and encrypt device communication; do not rely on being "inside".

The same applies to the embedded hardware you build. A microcontroller board with an open debug port, default credentials, an unauthenticated Bluetooth or Wi-Fi service, or a firmware update that anyone can push is a device with a published attack surface. Disable or protect debug interfaces in anything that leaves your bench, sign firmware updates, avoid hard-coded secrets, decide what the device should do when its link fails — the **safe state** — and record the firmware version on the object and in its output, so an observation can be traced to a build (Chapter 19).

20.8. AI and machine learning in medicine

By 2025 the FDA's public list of AI-enabled devices had passed a thousand entries, most of them in radiology. Building a model that performs well on a test set is now easy. Building one that **works in a different hospital next year** is not, for reasons that should be familiar from Chapter 15:

- **Dataset shift.** Patients, scanners, protocols, coding practice, and prevalence differ between sites and drift over time. A widely deployed proprietary sepsis prediction model, externally validated at one US health system in 2021, performed substantially worse than its developer's figures and missed most sepsis cases while alerting on many patients who did not have it.
- **Shortcut learning.** Models learn whatever predicts the label. A pneumonia model learned to recognise **which hospital** and **which X-ray machine** produced an image — portable bedside films come from sicker patients — rather than the lungs.
- **Subgroup performance.** Good average performance can hide poor performance in a subgroup: darker skin in dermatology datasets, women in cardiac datasets, children everywhere. The pulse oximeter's bias with skin pigmentation (Chapter 9) is the hardware precedent.
- **Label quality.** "Ground truth" is often one clinician's opinion, a billing code, or a report generated for a different purpose. A model cannot be more reliable than its labels.
- **Evaluation pitfalls.** Leakage (the same patient in training and test sets; future information in features), prevalence mismatch (a balanced test set inflating predictive values — Chapter 15's

PPV arithmetic), and accuracy reported without calibration or decision consequences.

Regulators distinguish **locked** algorithms, which do not change after release, from **adaptive** ones that learn in the field. A change to a locked model is a design change. The FDA's **predetermined change control plan** approach (final guidance December 2024) lets a manufacturer specify in advance which modifications it will make, how they will be validated, and how users will be told — so a defined class of retraining does not require a new submission each time.

20.9. Engineering bridge: safety-critical habits in code

20.9.1. Units in the type system

Chapter 19 said units belong in the type, not the comment. Here is the minimal version in Python; most languages can do better with newtypes or a units library.

```
from dataclasses import dataclass

MMOL_PER_MGDL = 1 / 18.0  # glucose only — molar mass ~180 g/mol

@dataclass(frozen=True)
class GlucoseMmolL:
    value: float
    def __post_init__(self):
        if not (0.5 <= self.value <= 50.0):           # plausibility, not a clinical
range
            raise ValueError(f"glucose {self.value} mmol/L outside plausible bounds")

def glucose_from_input(value: float, unit: str) -> GlucoseMmolL:
    if unit == "mmol/L":
        return GlucoseMmolL(value)
    if unit == "mg/dL":
        return GlucoseMmolL(value * MMOL_PER_MGDL)
    raise ValueError(f"unrecognised unit {unit!r}")   # never guess
```

Three habits are encoded here: the unit is part of the type, conversion happens once at the boundary, and implausible or ambiguous input is **rejected loudly**, not clamped or guessed. The bounds are a sanity check on data entry, chosen by the engineering team with clinical input — not a statement of what is normal.

20.9.2. Input validation that respects the user

Rejecting bad input is necessary but not sufficient. A pump that silently refuses "10" is as dangerous as one that accepts it, if the user does not notice. Good validation in a clinical context:

- states **what** is wrong and **what to do**, in the user's terms;
- distinguishes **impossible** (hard stop) from **unusual but possible** (deliberate confirmation, logged);

- never auto-corrects a clinical value;
- shows the value back with its unit before it is committed (**read-back**).

20.9.3. Audit trails

Clinical systems must record who did what, when, to which record — and the record must be trustworthy. A simple technique makes tampering detectable: chain each entry to the hash of the previous one.

```
import hashlib, json, time

def append_entry(log: list[dict], actor: str, action: str, record_id: str) -> dict:
    prev = log[-1]["hash"] if log else "0" * 64
    entry = {"t": time.time(), "actor": actor, "action": action,
             "record": record_id, "prev": prev}
    entry["hash"] = hashlib.sha256(
        json.dumps(entry, sort_keys=True).encode()).hexdigest()
    log.append(entry)
    return entry

def verify(log: list[dict]) -> bool:
    prev = "0" * 64
    for e in log:
        body = {k: v for k, v in e.items() if k != "hash"}
        if e["prev"] != prev or e["hash"] != hashlib.sha256(
                json.dumps(body, sort_keys=True).encode()).hexdigest():
            return False
        prev = e["hash"]
    return True
```

This detects alteration of the stored log; it does not prevent deletion of the whole log, or an attacker who can rewrite every entry. Production systems anchor the chain externally, write to append-only storage, and synchronise clocks. The point of the sketch is the property you want: **you can show that the record has not been changed**.

20.9.4. k-anonymity arithmetic

A dataset is **k-anonymous** if every combination of quasi-identifier values is shared by at least **k** records. Consider 1,000 fictional patients described by sex (2 values), year of birth (80 values), and a three-character postcode district (50 values). There are 2 × 80 × 50 = 8,000 possible combinations — eight times more cells than people. Most occupied cells will hold exactly one person: k = 1, and the dataset identifies nearly everyone in it to anyone who knows those three facts about them.

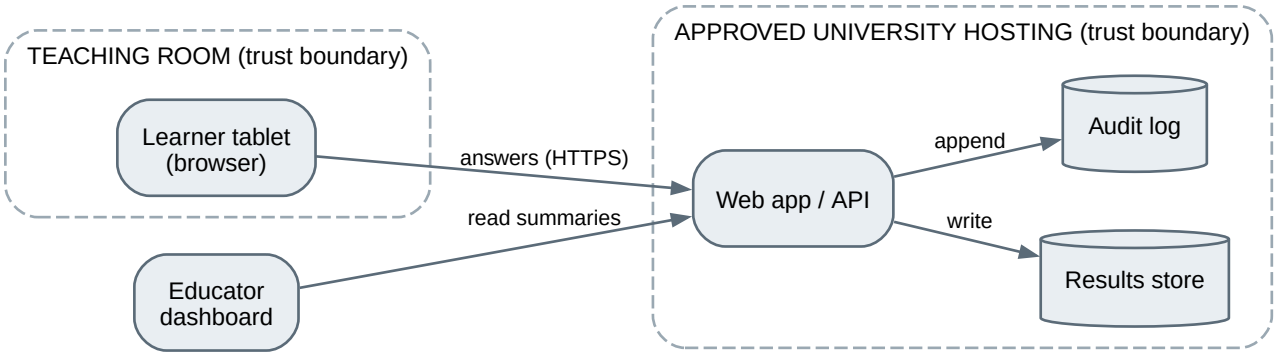
Generalise: 10-year age bands (8 values) and region instead of district (5 values) give 2 × 8 × 5 = 80 cells for 1,000 people — an average of 12.5 per cell, though some cells will still be small. Checking is a one-liner:

```
k = df.groupby(["sex", "age_band", "region"]).size().min()
```

k-anonymity is a **floor**, not a guarantee. If everyone in a cell has the same diagnosis, the diagnosis is disclosed anyway (the motivation for **l-diversity**); and linkage with outside data can still narrow things down. Formal de-identification is specialist work; the arithmetic is here so that you recognise when a dataset obviously fails.

20.9.5. Threat modelling with STRIDE

Draw the data flows, mark the trust boundaries, then ask six questions of each element:



Threat	Property violated	Example question for the diagram above
*S*poofing	Authentication	Can someone submit answers as another learner?
*T*ampering	Integrity	Can a result be changed after submission? Who could tell?
*R*epudiation	Non-repudiation	Can the educator show who changed a record, and when?
*I*nformation disclosure	Confidentiality	Can one learner see another's results? Is anything identifying in the URL or logs?
*D*enial of service	Availability	What happens to the session if the server is unreachable?
*E*levation of privilege	Authorisation	Can a learner reach the educator dashboard?

Each answer becomes either a control (with a test, Chapter 19) or an accepted, documented risk. The value is in being systematic; the threats you find are rarely exotic.

20.9.6. Fault trees for software

The Therac-25 figure above is a **fault tree**: a top undesired event decomposed through logic gates (AND, OR) into contributing causes — here, a single AND. The shape teaches the key point. When an AND gate joins a software fault to the **absence of an independent barrier**, adding a barrier — a hardware interlock, an independent check, a human read-back — turns a single point of failure

into a combination that must coincide. Software safety is designed at the architecture level, by asking what **else** must fail for the software's failure to matter.

20.10. Devices, software, and training technology

20.10.1. Clinical systems and software-intensive devices

System or device	Does	Engineering note
Electronic health record (EHR/EPR)	The longitudinal record, orders, results, documentation, prescribing	Mostly not a device as a whole, but contains device-like functions (dose calculators, CDS). Huge configuration surface; local build matters as much as vendor code
PACS, RIS, LIS	Images, radiology and laboratory workflow (Chapter 15)	HL7 v2 and DICOM integration; long data retention; large storage
Smart infusion pumps	Drug delivery with dose error reduction software	Drug library management, wireless library updates, and increasingly EHR–pump integration (auto-programming) to remove keypad entry
Patient monitors and central stations	Vital signs, alarms, networked surveillance	Alarm logic, latency, and the stale-value problem (Chapter 9)
Medical device integration middleware	Collects device data into the EHR	Time synchronisation, patient association ("is this monitor attached to this patient?"), and validation of automatically charted values
Clinical decision support	Alerts, risk scores, order sets, early warning systems	Alert fatigue; the CDS regulatory grey zone; local tuning
Insulin pumps, CGMs, automated insulin delivery	Closed-loop glucose control	Class C territory. Interoperable components from different makers must be verified together (Chapter 10)
AI-enabled imaging software	Triage, detection, measurement	Dataset shift, subgroup performance, post-market monitoring (Chapter 16)
Digital therapeutics and remote monitoring apps	Software-delivered treatment; home measurements sent to clinicians	Who responds to the data, and when, is part of the design
Telehealth platforms	Remote consultation	Security, identity verification, and accessibility (Chapter 7)

20.10.2. Tools you can use today

Tool	Does	Note
Synthea [39]	Generates synthetic patient records, exportable as FHIR	**Develop against this, never against real data**
HAPI FHIR	Open-source Java FHIR server and client library	Run a local FHIR server in a container in minutes
SMART Health IT sandbox; public FHIR test servers	Test SMART-on-FHIR app launch and API calls	Public test servers are public: put only synthetic data on them
Inferno	Open-source FHIR conformance test suite used for US certification	See what "conformant" actually tests
pydicom, dicom-anonymizer tools, defacing tools	Read, inspect, and de-identify DICOM (Chapter 15)	Inspect **every** header element, including private tags, before trusting a de-identifier
ARX, sdcMicro	Open-source statistical de-identification and risk measurement	k-anonymity, l-diversity, and re-identification risk metrics on tabular data
OWASP Threat Dragon, Microsoft Threat Modeling Tool	Data-flow diagrams and STRIDE threat enumeration	Free
Syft, CycloneDX and SPDX tools, OSV-Scanner, Dependency-Track	Generate SBOMs and match dependencies against known vulnerabilities	The SOUP list and the SBOM are the same discipline
PhysioNet, TCIA [36] [35]	Public physiological and imaging data	Some datasets, notably MIMIC, require credentialing, training, and a data use agreement — honour it exactly

20.10.3. Training and simulation

Training technology	Teaches	Form
EHR training environments	Clinicians and staff learn the live system safely	A copy of the production build with fictional patients. Keeping it in step with live configuration is a real engineering job
Educational EHRs	Documentation and record-keeping for students	Simplified EHRs populated with fictional cases

Training technology	Teaches	Form
FHIR and IHE Connectathons	Interoperability, by testing against other implementers	Regular events. **Possibly the fastest way to meet the health IT community**
Cybersecurity tabletop exercises	Incident response: "the lab system is encrypted; what now?"	Scenario-based, cross-disciplinary. Downtime procedures are as important as recovery
Medical device security labs and capture-the-flag events	Device vulnerability research, coordinated disclosure	E.g. the Biohacking Village at DEF CON. Always within disclosure norms and the law
Clinical safety officer training	Clinical risk management for health IT (DCB0129/0160 in England)	Short courses for clinicians and engineers
Simulated clinical scenarios with software	Using decision support and monitors under pressure	Combine with Chapter 21; build the mock monitor of Chapter 9



If your project involves software or connected electronics, start here.

Health systems need people who build software and devices and understand units, codes, governance, and alarms — and who will build with synthetic data, write the threat model before the code, and treat "not a device" as a question rather than an assumption. Run Synthea, stand up a HAPI FHIR server, write a SMART app that reads synthetic Observations and displays them with their units and timestamps, and threat-model it. Or build a battery-powered sensor on a teaching model that publishes its readings, with units and timestamps, to that same test server — and threat-model the radio link and the firmware update path too. Either is a portfolio piece a clinical informatics or clinical engineering team can evaluate in ten minutes.

20.11. Worked example: the companion app that grew

Your elbow teaching model (Chapter 22) is going well. The educator asks for a small web app: learners scan a QR code on the model, answer three questions about what the cord does, and the answers go into a results table so the educator can see which labels confuse people.

Version 1 — the brief as given. Learners enter a nickname and answer the questions. Results are stored with a timestamp. No real names, no health information, hosted on the university's approved platform.

- **Device?** No. It assesses learning about a teaching model; no clinical purpose.
- **Personal data?** Minimal. A nickname plus timestamps in a small class can be identifying, so treat it as personal data anyway: tell learners what is stored and why, keep it only as long as needed, and let them use an arbitrary nickname.

- **Security?** Ordinary: HTTPS, an educator login, learners cannot see each other's answers. The STRIDE table above was drawn for exactly this app.

Version 2 — "Could it also record whether the learner has joint hypermobility, so we can see whether it changes how they interpret the model?"

That is **health data** — special category under GDPR — about identifiable students, for a research-like purpose. Stop and change the plan:

- Is the data necessary for the stated purpose? (Minimisation.) Probably not for improving labels.
- If it is pursued, it is a **study**: it needs a lawful basis and an Article 9 condition, a DPIA, very likely ethics review, information governance approval, and a consent process in which students can genuinely decline without consequences from the person who grades them.
- You do not decide this. Write down the question, send it to the educator and the information governance team, and do not add the field.

Version 3 — "Patients in physiotherapy could use it at home: they'd measure their elbow angle with the phone camera and the app would tell them whether they're improving."

Now the intended use is **monitoring a patient's condition and informing their management**.

- **Device?** Very probably yes. Under the IMDRF framing it **informs** or **drives** management of a **non-serious to serious** situation — category I or II — and under EU MDR Rule 11 it would likely be at least Class IIa. A measurement claim ("your elbow bends 95°") brings accuracy requirements (Chapter 15): camera angle, lighting, skin tone, clothing, and validation against a goniometer measured by a therapist.
- **Software lifecycle?** IEC 62304 applies, with a safety class determined by what a wrong output could cause: a patient told they are improving when they are not, or pushing too hard after surgery. Class B is a plausible starting point — to be decided with the manufacturer's risk process, not by you.
- **Data?** Patient health data, at home, on personal phones. Privacy, security, and hosting questions multiply.
- **Who is the manufacturer?** Whoever places it on the market. A university, a hospital, or you — each with obligations.

The right response is not to refuse and not to build it quietly. It is to say: **"That is a different product, with a different regulatory status. I can build Version 1 now. Version 3 needs a sponsor, a regulatory decision, and a clinical lead. Here is what that would involve."** Then record the conversation (Chapter 19).

The general lesson: in software, **scope creep is regulatory creep**. Each innocent feature request can move the product across a line that changes everything about how it must be built. Keep the intended-use statement in the repository's README and re-read it whenever a feature request arrives.

20.12. Workshop exercise

1. **Classify five apps.** For each, write an intended-use sentence, place it in the IMDRF table, and say what you would need to ask to be sure: a step counter; a medication reminder; an app estimating heart rate from a fingertip on the camera, labelled "for fitness only"; a skin-lesion photo app that rates melanoma risk; a sepsis alert in an EHR.
2. **Build a synthetic FHIR pipeline.** Generate 100 patients with Synthea, load them into a local HAPI FHIR server, and write a script that retrieves all blood-pressure Observations (systolic and diastolic, by LOINC code) and prints them **with units and timestamps**. Then deliberately feed it one Observation with a missing unit and one with a different unit, and make the script reject or convert them explicitly.
3. **De-identify a DICOM file.** Download a public DICOM study (TCIA or the 3D Slicer samples). List every header element with pydicom, including private tags. Mark each as direct identifier, quasi-identifier, or safe. Apply a de-identification profile and check what remains — and whether any text is burned into the pixels.
4. **Measure k.** Using Synthea's patient table, compute k for (sex, birth date, postcode), then for (sex, birth year, region), then for (sex, ten-year band, region). Record how k changes and what each generalisation costs analytically.
5. **Threat-model the mock monitor.** Draw the data-flow diagram for the classroom mock monitor (Chapter 9), mark trust boundaries, and fill in a STRIDE table. Identify one threat that is really a **safety** hazard — for example, a tampered display shown as if real — and trace it into a hazard table entry (Chapter 19).
6. **Produce an SBOM.** Generate an SBOM for any project of yours with Syft or a CycloneDX tool, run a vulnerability scan against it, and write a short SOUP list for its three most important dependencies: version, what you rely on it for, and its known anomalies.

20.13. Questions

1. What determines whether a piece of software is a medical device?
2. Distinguish SaMD from SiMD, with an example of each.
3. Place a smartphone app that tells a patient with heart failure to "go to the emergency department now" on the IMDRF table, and justify the cell.
4. What are the IEC 62304 software safety classes, and how can a hardware design decision change the class of software?
5. What is SOUP, what must you record about it, and how does it relate to an SBOM?
6. Give three distinct software-engineering lessons from Therac-25.
7. Why is pseudonymised data still personal data, and what more is needed to reach anonymity?
8. Name four quasi-identifiers in a typical hospital dataset and explain why they matter.
9. Under GDPR, why is consent often not the lawful basis for processing health data for direct care, and what else is required for special category data?
10. What does a LOINC code add to a FHIR Observation that the display text does not?

11. Why is availability a patient-safety property? Give a real example.
12. What must a US premarket submission for a "cyber device" now include?
13. Describe two ways a medical AI model with excellent test-set performance can fail after deployment.
14. A dataset of 2,000 patients is described by sex, single year of age (0–99), and one of 40 districts. Without data, what can you say about k, and what would you change first?

20.14. Answers and reasoning

1. Its **intended purpose** as stated or implied by the manufacturer — labelling, claims, promotion, user interface — matched against the legal definition of a medical device. The code alone does not decide it; the same algorithm can be wellness or a device depending on what it is said to do.
2. **SaMD** is software that is itself the device, on general-purpose hardware (an app grading retinal photographs). **SiMD** is software that is part of a hardware device (infusion pump firmware, CT reconstruction software).
3. Heart failure decompensation is a **serious** to **critical** situation, and an instruction to attend the emergency department **drives** clinical management (arguably more). That places it at category II–III, and the patient-facing, action-directing nature pushes towards the higher end. A strong answer notes that the exact cell would be argued with regulators and that "drive" versus "inform" is the crux.
4. **A**: no injury possible; **B**: non-serious injury possible; **C**: death or serious injury possible — assessed after risk controls **external** to the software. An independent hardware control (pressure-relief valve, hardware interlock, watchdog) can mean a software failure can no longer cause the harm alone, lowering the class. Removing such a control does the opposite.
5. Software of unknown provenance: components not developed under your 62304 process. You record identity and version, the requirements you rely on it for, its hardware and software needs, and an evaluation of its published anomalies. An SBOM is the machine-readable inventory of the same components, now required for security monitoring; the SOUP list adds the safety reasoning.
6. Any three of: do not replace independent hardware interlocks with software without equivalent assurance; concurrency faults can depend on user speed and evade normal-pace testing; frequent cryptic error messages train users to override them; reused code is not proven safe by its history in a different system; incident reports must be investigated, not reassured away; a counter overflow can silently skip a safety check.
7. Because the key that re-links the pseudonyms to people still exists, and the remaining quasi-identifiers may allow re-identification by linkage. Anonymity requires that no one can reasonably identify anyone, by any means reasonably likely to be used, in context — in practice, destroying keys, generalising or suppressing quasi-identifiers, assessing linkage risk, and controlling the release environment.
8. Any four of date of birth, postcode, sex, admission or procedure dates, ethnicity, occupation, rare diagnosis, hospital or ward. Individually common, in combination they are frequently unique — about 87% of the US population by ZIP, sex, and full date of birth in Sweeney’s estimate — so they identify people when linked with other data.

9. Consent must be freely given and withdrawable, which fits poorly with care that must be recorded and shared, and where there is an imbalance of power; other bases such as a public task or the provision of health care are usually used. Special category data also require an Article 9 condition, and high-risk processing a DPIA. Confidentiality duties and national law add further requirements.
10. Machine-actionable meaning. The LOINC code identifies exactly what was measured (and, implicitly, how), so a receiving system can act on it; the display text is for humans and varies between systems. With a UCUM unit, the value becomes unambiguous.
11. Because a system that is down stops care: tests cannot be run, results cannot be seen, medicines cannot be checked. Examples: WannaCry (2017) cancelling NHS appointments and diverting ambulances; Synnovis (2024) halting London pathology and transfusion services.
12. A plan to monitor, identify, and address post-market vulnerabilities, including coordinated vulnerability disclosure; processes providing reasonable assurance that the device and related systems are cybersecure, including updates and patches; and a software bill of materials. FDA guidance also expects threat modelling, security architecture, and testing evidence.
13. Any two of: dataset shift (different population, scanner, protocol, or prevalence); shortcut learning of site or device artefacts; poor subgroup performance hidden by the average; unreliable labels; leakage that inflated test results; prevalence mismatch making predictive values much worse in practice.
14. There are 2 × 100 × 40 = 8,000 possible combinations for 2,000 people, so most occupied cells hold one person: k is almost certainly 1, and the dataset is effectively identified. Generalise age first (e.g. ten-year bands, 10 values → 800 cells) and districts into larger regions, then **compute** k, check the smallest cells, and consider suppression and l-diversity.

20.15. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Health data and medical software — software status, k-anonymity, Safe Harbor, FHIR, HL7 v2, audit logs and STRIDE.
- DICOM viewer and de-identification check — de-identification against DICOM Annex E.
- Unit-safe clinical arithmetic — units carried with every value.
- Test accuracy and risk — alarms as a test with a prevalence problem.
- Hashes — SHA-256.
- JSON — FHIR JSON.
- Data wrangler — grouping a CSV to check k.

20.16. Video companion and reading

- **Required:** Leveson and Turner, *An Investigation of the Therac-25 Accidents* [61]. Read the whole paper; it is the best case study in safety-critical software ever written.
- **Regulatory shape:** the IMDRF SaMD risk categorisation framework [55]; MDCG 2019-11 on

qualification and classification of software [58]; FDA guidance on clinical decision support software [56] and on cybersecurity [57]. Read for structure, not detail.

- **Data protection:** HHS guidance on HIPAA de-identification methods [59], which explains Safe Harbor and Expert Determination clearly; your national data protection regulator's guidance on health data and anonymisation.
- **Interoperability:** the FHIR specification [60] — start with the overview, the Observation resource, and the RESTful API page, then build something.
- **Tools:** Synthea [39], PhysioNet [36], TCIA [35].

20.17. Exit criteria

- ▢ Write an intended-use statement for a software idea and place it in the IMDRF table, naming the question that would decide the cell.
- ▢ Explain IEC 62304 safety classes, SOUP, and configuration management, and map them to practices you already use.
- ▢ Summarise three lessons from Therac-25 and say where each applies to your own work.
- ▢ Classify the fields of a sample dataset as direct identifiers, quasi-identifiers, or sensitive attributes, and compute k for a chosen generalisation.
- ▢ Inspect and de-identify a public DICOM file, documenting what remained.
- ▢ Build a small pipeline on synthetic FHIR data that handles codes and units explicitly and rejects ambiguous input.
- ▢ Produce a data-flow diagram, STRIDE table, and SBOM for one of your own projects.
- ▢ State in writing who, in a host organisation, decides data access, clinical safety, and regulatory status — and that it is not you.

20.18. What this chapter does not establish

- **You cannot determine regulatory status.** You can reason about where a product probably sits and which questions decide it. The determination belongs to the manufacturer's regulatory function and, ultimately, the regulator.
- **You are not compliant with IEC 62304, 81001-5-1, or any other standard.** Knowing their shape lets you work inside a process that applies them. Applying them requires a quality system, trained people, and records you do not have alone.
- **This is not legal advice.** Data protection and health privacy law differ by jurisdiction, change often, and depend on facts. Information governance staff and lawyers decide.
- **De-identification is specialist work.** The arithmetic here lets you recognise an obvious failure. Declaring a dataset anonymous requires formal risk assessment by people qualified to make it — and the declaration can be overturned by new linkable data.
- **Security testing of clinical systems requires authorisation.** Never probe, scan, or test a hospital system or a medical device you do not own without explicit written permission from its owner. Good intentions are not a defence, and a crashed device is a patient safety incident.

- **A model's metrics are not its clinical value.** Whether an algorithm helps patients is a clinical evaluation question, answered prospectively, in the setting of use.

Chapter 21. Making a model that teaches

Week 26 · approximately 4–6 hours

21.1. Why this chapter exists

A visually impressive model is useful only if it supports the educator's learning objective.

That sentence is easy to agree with and hard to act on, because the instinct of a skilled maker is to make the thing **better** — more detailed, more realistic, more finished. For a teaching aid, "better" frequently means **less**: fewer features, more visible mechanism, cruder appearance, clearer relationships.

This chapter is about designing the **learning**, and treating the object as the thing that serves it rather than the other way round.

21.2. Begin with what the learner must do

"Understand the elbow" cannot guide a build, cannot be assessed, and cannot fail. It is not an objective; it is a topic.

An objective names an **observable behaviour**, performed by a specified learner, under specified conditions, to a specified standard.

❌ Topic	✅ Objective
"Understand the elbow"	"After using the model, the learner will **identify** elbow flexion and **explain** that a tendon transmits force produced by muscle."
"Appreciate anatomy"	"The learner will **name** the humerus, radius, and ulna on the model and **state** which one rotates during pronation."
"Learn about joints"	"Given the model, the learner will **predict** the direction of movement before the cord is pulled, and **justify** the prediction by reference to the attachment point."



Use verbs you can observe. Identify, name, state, explain, describe, demonstrate, predict, compare, distinguish, justify, perform, select.

Avoid verbs you cannot observe. Understand, appreciate, know, be aware of, be familiar with, grasp. These describe an internal state you have no access to. If you cannot see it happen, you cannot tell whether your aid worked.

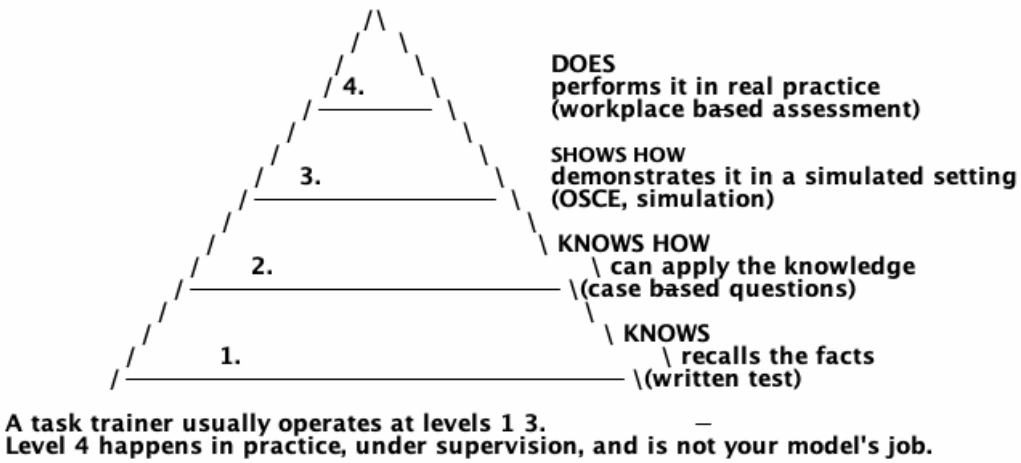
21.2.1. Bloom's taxonomy, briefly

A standard hierarchy of cognitive demand, useful because it stops you writing a recall-level objective and then testing at the analysis level, or vice versa.

Level	Learner can…	Elbow model example
Remember	Recall facts	Name the three bones at the elbow
Understand	Explain ideas	Explain what a tendon does
Apply	Use it in a new situation	Predict the motion when the attachment is moved
Analyse	Break it down, find relationships	Explain why a shorter moment arm needs more force
Evaluate	Judge against criteria	Identify what the model misrepresents and why it matters
Create	Produce something new	Design a modification that teaches pronation

Your assessment must sit at the same level as your objective. If the objective says **explain**, asking a yes/no recall question does not assess it.

21.2.2. Miller's pyramid: knowing is not doing



Knowing that a tendon transmits force is level 1–2. Demonstrating the relationship on a model is level 3 at best. Doing the right thing with a real patient is level 4, and no bench model establishes it. Be clear with yourself and the educator about which level your aid addresses.

21.3. Choose the realism that serves the objective

Fidelity is not one thing, and the forms of it are not interchangeable.

Dimension	What it means	When it matters
Visual / appearance	Looks like the real thing — colour, texture, proportion	When the objective is **recognition** — identifying a structure, spotting an abnormal appearance

Dimension	What it means	When it matters
Anatomical / spatial	Correct structures in correct relationships, correct scale	When the objective involves **locating** something, or understanding what lies next to what
Tactile	Feels right — compliance, texture, resistance, the "pop" of a layer	When the objective is a **procedural skill** guided by feel
Mechanical / functional	Behaves right — the motion, the force path, the cause and effect	When the objective is **how something works**. **This is the elbow model's dimension**
Environmental / contextual	The surroundings, the noise, the time pressure, the team	When the objective involves **performance under realistic conditions** — teamwork, crisis management
Emotional / psychological	The learner behaves as they would in the real situation	When the objective involves **decision-making under stress**



Choose your fidelity, and then deliberately sacrifice the rest.

For a force-path objective, skin-coloured paint adds nothing. A **transparent or open** construction, which looks less realistic, teaches the relationship far better because the mechanism is visible. Reduced visual fidelity is not a compromise here; it is the right design.

Conversely, a trainer for locating a structure by touch **needs** correct spatial relationships and plausible tactile behaviour. A generic hinge would be useless no matter how well made.

The question is never "how realistic should this be?" It is "**realistic in which dimension, and why that one?**"

21.3.1. Each kind of fidelity is made a different way

Once the dimension is chosen, the making method usually follows from it — and it is rarely one method. Most good trainers combine several.

Fidelity	Typically made by
Visual	Painting and moulage; pigmented cast silicone; printed or vacuum-formed surfaces; dyed and printed fabrics; photographs and displays
Anatomical / spatial	Structures cast or printed from imaging data (Chapter 15); machined or laser-cut location features; sewn compartments that hold parts in their correct relationship
Tactile	Layered cast silicone and gels of differing hardness; fabric and mesh layers for membranes and fascia; foams; embedded tubing for vessels — usually cast in moulds that were themselves machined or printed

Fidelity	Typically made by
Mechanical / functional	Turned pins and bushings; machined or laser-cut linkages; springs, cord, and cable; printed parts where geometry is complex; glass or acrylic tubing for fluid paths
Instrumented feedback	Sensors — force-sensing resistors, load cells, pressure transducers, angle sensors, inertial units — with a microcontroller and a display, light, or sound. The CPR manikin that reports compression depth is the familiar example. Firmware and data rules apply (Chapter 20)
Environmental	Props, garments, lighting, and sound; wearable trainers strapped to an actor (**hybrid simulation**), which are largely a sewing and fastening problem

Notice how often the printed or machined part is a **tool** — a mould, a jig, a fixture — rather than the thing the learner touches. Chapter 14 covers the methods and how to choose between them.

21.3.2. What our model needs, and does not

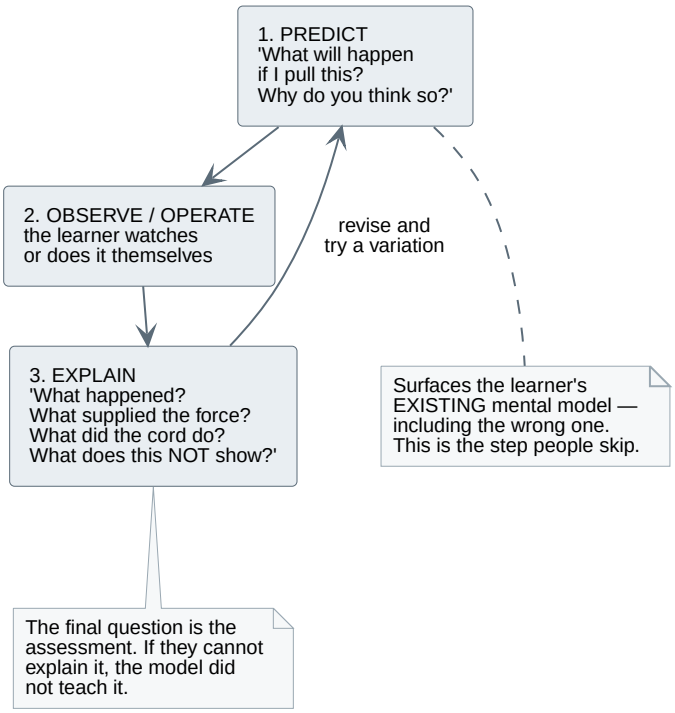
Feature	Needed?	Why
Visible pivot	**Yes**	The joint must be identifiable as the thing that moves
Obviously moving forearm segment	**Yes**	The motion is the point
Clearly labelled force input, distinct from the cord	**Yes**	This **is** the objective
Visible relationship between pull and movement	**Yes**	Cause and effect must be observable
Exact bone contours	No	Irrelevant to the force-path objective
Skin texture or colour	No	Actively unhelpful; obscures the mechanism
Motorisation	No	Hides the force source — it would **defeat** the objective
Electronic feedback (e.g. an angle readout)	Optional	Only if the objective needs a number; it must add feedback without supplying force or hiding the mechanism
Forearm rotation	**Deliberately omitted**	Out of scope — **and must be stated**
Anatomically correct tendon insertion	**Deliberately omitted**	The guide loop is a model convenience; label it as such

That table is also the beginning of your **limitations statement**. Deliberate omissions are not failures; undeclared omissions are.

21.4. Design the explanation as carefully as the object

The object does not teach. The **interaction** teaches, and you can design it.

21.4.1. Predict–observe–explain



The **predict** step is what makes this work. Asking for a prediction before the demonstration surfaces what the learner already believes. If they predict wrongly and then see the right thing happen, the mismatch is memorable in a way that watching a correct demonstration never is.

The **explain** step is the assessment. Note that it includes **"name something this model does not show"** — which assesses whether the learner has understood the model as a model.

21.4.2. A worked teaching script

Stage	Script
Before movement	"This is a simplified joint. Point to the part representing the upper arm, and the part representing the forearm. Now: predict what happens if I pull this cord. Why?"
During movement	"Watch where the cord pulls, relative to the pivot. What changes as it moves?"
After movement	"Which movement did you see? What supplied the force? What does the cord represent? Now name one thing about a real elbow that this model does **not** show."
If they say the tendon contracts	Do not correct and move on. Record it. Then change the labels, the demonstration, or the model, and try again with a different learner. Their error is **data about your design**.



Do not explain the answer while assessing whether the learner can produce it.

It is almost irresistible to help. You built the thing; you want it to work; a learner hesitating is uncomfortable. So you say "...and what's pulling it?" and they say "the muscle" and you record a success.

You have just recorded **your** explanation as evidence of **the object's** effectiveness.

If you prompt, **record the prompt**. "Completed with two prompts" is an honest and useful result. "Completed" is not, and the difference is exactly the one from Chapter 19: assisted performance is a weaker claim than independent performance, and the record must say which you observed.

21.5. Evaluate learning and maintenance separately

Keep two checklists. They answer different questions and a model can pass one and fail the other.

	Learning checklist	Object checklist
Asks	Did the learner achieve the objective?	Can this thing survive being used?
Items	Did they predict? Did their explanation match the objective? Did they identify a limitation? What misconception appeared? What prompts were given?	Setup time. Reset time and method. Damage after use. Missing or loose parts. Cleaning. Storage. Instructor visibility across a room. Who maintains it
Failure looks like	A robust, beautiful object that teaches the wrong thing	An excellent demonstration that takes six minutes to reset and so is never used



Reset time is the requirement educators care about most and mention least.

A teaching session has a fixed length and a queue of learners. An aid that takes ninety seconds to reset between attempts, in a class of twelve, consumes eighteen minutes. That aid will be used once and then left in the cupboard, regardless of how good it is.

Ask about it explicitly: "How long have you got between learners?" The answer is a first-class requirement (Chapter 19).

21.6. Misleading cues: the real hazard of a teaching model

From Chapter 19: the most likely harm from a teaching aid is not injury, it is **teaching something false**. A simplified model should not quietly teach an inaccurate relationship.

Simplification	What it might quietly teach	Mitigation
Single pivot at the elbow	That the elbow is one hinge joint with one degree of freedom	State the omission of pronation and supination explicitly, every time
Cord pulled by hand	That the tendon generates force	Label the force input and the transmission element separately and redundantly (Chapter 7)
Guide loop positioned for convenience	That this is where the biceps tendon actually runs	Label as "model convenience, not an anatomical route"
Manual reset	That extension is passive, or that the model returns itself	"Manual reset is an external action. The model does not represent the muscle activity involved in extension"
Rigid uniform strips	That bone is uniform and the limb is rigid	"Schematic, not to scale. Does not represent bone structure or soft tissue"
No cartilage, capsule, or ligaments	That a joint is just a pivot	List these among the stated omissions
Smooth, frictionless motion	That joints move freely without muscular control or stabilisation	Mention co-contraction and stabilisation verbally; it need not be modelled



An anatomically qualified educator should review the anatomy, the labels, and the limitations statement before the model is used for teaching.

You are not qualified to certify anatomical accuracy, and the misconceptions you are most likely to build in are the ones you share. This review is a hard requirement, not a courtesy — and if you have no educator yet, record it as **pending** and do not describe the model as anatomically reviewed.

21.7. Engineering bridge: materials and methods for simulation

If you move further into simulation, tissue simulants and trainer construction become a topic in their own right. A brief orientation:

Material	Used for	Notes
Platinum-cure silicone (various Shore A)	Soft tissue, skin, organ surfaces	Shore 00 for very soft tissue, Shore A 10–30 for firmer. Choose by the **behaviour** the objective needs, not by a number someone quoted
Ballistic / gelatin gels	Needle insertion, ultrasound phantoms	Cheap, easy to cast, degrade quickly; often single-session

Material	Used for	Notes
Hydrogels / agar	Ultrasound phantoms	Tunable acoustic properties; perishable
Foams (open and closed cell)	Bulk, padding, soft structure	Open-cell absorbs fluid and is usually unsuitable where cleaning matters (Chapter 13)
Textiles — stretch fabrics, meshes, felts, coated fabrics	Fascia and membrane layers, wearable trainers, garments, moulage bases, straps and covers	Sewing gives strength, flexibility, and easy replacement. Coated fabrics wipe clean; uncoated ones absorb fluid and must be laundered or replaced (Chapter 13)
Glass, acrylic, and polycarbonate tubing and sheet	Vessels and fluid paths, transparent chambers, phantom containers, viewing windows	Shows flow and level directly. Glass is chemically inert and scratch-resistant but breaks; acrylic is easier to work but crazes with some solvents and alcohol-based cleaners
Machined and laser-cut parts — aluminium, acetal, acrylic, plywood	Skeletal structure, joints, frames, bases, jigs	Accurate, durable, repeatable, and quick for flat or prismatic shapes. Acetal and brass make low-friction pivots; plywood must be sealed if it will be cleaned
Printed polymers, rigid and flexible	Complex skeletal shapes, housings, compliant features, and the **moulds** for casting silicone	Fast for complex one-offs. Layered and anisotropic: consider orientation against load (Chapter 6), and note that surface finish affects both feel and cleanability
Sensors and microcontrollers	Measuring what the learner does — force, depth, angle, pressure, timing — and reporting it	Battery-powered, low voltage, enclosed, and cleanable; record the firmware version. Feedback can change what is learned, so the educator agrees what is measured and shown



"Realistic feel" is an educator's judgement, not a durometer reading.

The correct process is: the educator tells you what the learner must be able to feel and discriminate; you make two or three candidates; they compare them; you record which they chose and why. Matching a published tissue modulus is not the objective and can actively mislead, since the relevant property for a procedural skill is usually the **contrast** between layers, not the absolute stiffness of any one.

21.8. Worked example: adapting the model to one objective

The educator's objective: "After using the model, a first-year learner will identify elbow flexion and explain that a tendon transmits force produced by muscle."

What that objective requires the model to make visible:

1. The **motion** — flexion, clearly and unambiguously.
2. The **force input** — identifiable as separate from the transmission.
3. The **transmission element** — identifiable as a transmitter, not a generator.
4. The **relationship** between pulling and moving.

What it permits the model to omit: everything else. Appearance, exact anatomy, pronation, extension mechanics, scale.

The learner task: "Before I pull the cord, predict which way the forearm part will move and why. Then watch. Then tell me what supplied the force, what the cord did, and one thing a real elbow has that this does not."

The observation checklist:

Observation	Y/N	Notes
Made a prediction before operation		What did they predict?
Correctly identified the movement as flexion		
Stated that an external force (hand/muscle) supplied the force		
Stated that the cord **transmitted** rather than generated		
Named at least one limitation of the model		Which one?
Prompts given		**Record each one verbatim**
Misconception observed		Record exactly what they said

The likely result, honestly recorded: the learner understands flexion but thinks the cord supplies its own energy. That is a **partial** result, not a failed project. The diagnosis is that the force source and the transmission element are still visually merged. The response is to separate them further — a visible handle on the pull end; a distinct "muscle" element where the force originates, such as a short sewn fabric sleeve or a cast silicone segment in a contrasting colour; a different colour, texture, and label for the cord; perhaps an engraved or painted arrow — and then check it again with a **different** learner, because the first one now knows the answer.

21.9. Workshop exercise

1. **Write one observable objective** for your model, using an observable verb, and state which level of Bloom's taxonomy it sits at.
2. **Write the matching assessment.** It must sit at the same cognitive level as the objective. Check this explicitly.
3. **Fidelity audit.** List every feature of your model. For each, state which fidelity dimension it serves and whether the objective requires it. Delete one feature that serves no dimension, and write down whether learning suffers.
4. **Method map.** For each feature you kept, name two different ways it could be made — for example a laser-cut and a machined forearm segment, or a sewn and a cast "muscle" — and say which better serves the fidelity dimension it exists for, and why.
5. **Write the teaching script** using predict–observe–explain, including the limitation question.
6. **Run it.** With an educator's permission and in an appropriate setting, have one willing learner attempt the task. Record their prediction, their explanation, every prompt you gave, and any confusion. Record no personal information you do not need.
7. **Build both checklists** — learning and object — and complete the object one yourself.
8. **Get the educator review.** Ask an educator to review the anatomy, the labels, and the limitations statement. Record their corrections. If you have no educator, mark this **pending** and do not claim the model is reviewed.
9. **Write the limitations statement.** Turn your omissions table into a paragraph that travels with the model.

21.10. Questions

1. Write a one-sentence observable objective for a model of your choosing.
2. Why is "appreciate anatomy" a poor objective?
3. Name the six fidelity dimensions and say which one the elbow model needs.
4. Would skin-coloured paint improve the elbow model? Why or why not?
5. Remove the visible cord attachment from the model. Does learning suffer? Why?
6. Write a question that checks understanding without giving away the answer.
7. A learner says "that was great." What have you assessed?
8. Why does the **predict** step matter more than it looks?
9. You help a learner through the explanation and they get it right. What may you record?
10. What is the most likely **harm** from a teaching model, and why is it easy to miss?
11. Why do educators care about reset time more than they say?
12. Who must review the anatomy, and what happens if that person is not available?

21.11. Answers and reasoning

1. A good answer specifies an observable behaviour, for example: "After using the model, the learner will identify elbow flexion and explain that a tendon transmits muscle-generated force." Weak answers use unobservable verbs.
2. Because "appreciate" describes an internal state you cannot observe, so you cannot tell whether it happened, cannot design an assessment for it, and cannot tell whether the aid worked. It also gives no guidance about what the model must show.
3. Visual/appearance, anatomical/spatial, tactile, mechanical/functional, environmental/contextual, emotional/psychological. The elbow model needs **mechanical/functional** fidelity — the force path and the cause-and-effect relationship.
4. **No, and it would probably make it worse.** It serves visual fidelity, which this objective does not need, while obscuring the mechanism, which this objective is entirely about. A transparent or open construction would be better despite looking less realistic.
5. **Yes, badly.** The visible attachment is what makes the relationship between the force input, the transmission element, and the motion observable. That relationship is the objective. Removing decorative texture costs nothing; removing the visible attachment removes the lesson.
6. "What provides the force in this demonstration, and what carries it to the moving part?" It checks the generation/transmission distinction without naming either answer. A poor version is "Does the cord transmit the force from the muscle?", which supplies the answer and invites agreement.
7. A reaction. Possibly engagement, possibly politeness. Not the intended knowledge or skill. Assessment requires a prediction, an explanation, or a demonstration aligned with the objective.
8. Because it surfaces the learner's **existing** mental model, including the wrong one, before the demonstration can overwrite it. A learner who predicts wrongly and then sees the correct behaviour experiences a memorable mismatch. A learner who simply watches a correct demonstration can agree with it without ever revising what they believed.
9. That the task was completed **with assistance**, and exactly what prompts were given. Recording it as a success attributes your explanation to the object's effectiveness. Assisted performance is a legitimate observation and a much weaker claim.
10. A **durable misconception** — teaching a relationship that does not exist. It is easy to miss because there is no visible failure at the bench, the harm is delayed, it propagates into everything the learner subsequently learns, and the model looks like it worked.
11. Because teaching time is fixed and the queue is not. An aid needing ninety seconds of reset in a class of twelve consumes eighteen minutes of a session, so it gets used once and then stays in the cupboard. It is a hard constraint that educators treat as background noise, so you have to ask.
12. An anatomically qualified educator must review the anatomy, labels, and limitations statement — because you are not qualified to certify anatomical accuracy and your most likely errors are the ones you cannot see. If no such person is available, record the criterion as **pending**, do not describe the model as reviewed, and do not use it for formal teaching.

21.12. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Project worksheets — learning objectives with observable verbs, and reset-time cost.

21.13. Video companion and reading

- **Core reading:** INACSL *Simfographics* [28], especially **Simulation Design** and **Outcomes and Objectives**. These are accessible one-page summaries of the healthcare simulation standards. The full standards page [29] gives deeper references; individual articles may have access restrictions.
- **Paired video exercise:** rewatch Crash Course #20 on joints [11] with a new question: **which single concept would my model teach, and which features must therefore be visible?** Use it as anatomy content to translate, not as a tutorial on simulation design.
- **For iteration:** revisit the Stanford prototype-testing material from Chapter 19 [14].



The teaching script and checklists in this chapter are original practice examples written for this book. They are not INACSL instruments and should not be presented as such. Any procedural trainer must follow the **host educator's** approved teaching method, not one you devised.

21.14. Exit criteria

- ☐ Write one observable learning objective and a matching assessment at the same cognitive level.
- ☐ Name the fidelity dimensions and justify, feature by feature, the realism your model includes and omits.
- ☐ Produce a written limitations statement listing every deliberate omission.
- ☐ Write and use a predict–observe–explain teaching script including a limitation question.
- ☐ Record learner responses **and any prompts given**, separately and verbatim.
- ☐ Maintain two distinct checklists: learning outcome and object maintenance.
- ☐ Demonstrate how the model is reset and maintained, with a measured reset time.
- ☐ Obtain educator feedback on anatomical accuracy and potential misconceptions — or record it explicitly as pending.
- ☐ Record at least one improvement prompted by a learner or an educator.

21.15. What this chapter does not establish

- **You cannot certify anatomical accuracy.** That review belongs to a qualified educator, every time.
- **You cannot design a procedural trainer unsupervised.** Anything teaching a clinical

procedure must follow the host educator's approved method, and the method comes first.

- **One learner is not evidence of effectiveness.** It is a useful observation that generates the next design question.
- **Learner satisfaction is not learning.** They are different measurements, and the literature on educational evaluation is clear that they correlate poorly.
- **A model that demonstrates is not a model that teaches.** Demonstration is the easy part. Whether the learner's explanation changed is the question, and it has to be asked.

Part III: Doing the work

Chapter 22. Building and reviewing your first project

Weeks 27–30 · approximately 12–18 hours

22.1. Why this chapter exists

Everything in the preceding nine chapters is now applied to one small object you can actually finish.

The point is not the object. The point is the **portfolio**: a documented piece of work whose reasoning is visible, whose evidence is honest, whose limitations are stated, and which a stranger could pick up and understand. That is what you take to a conversation about supervised work, and it is worth more than a stack of certificates.

A finished, humble, well-documented project beats an ambitious unfinished one every time.

22.2. Choose one route

Route	What it is	Key partner
Training model	An articulated anatomical model, or a fixture for an existing simulation component. The running elbow model is the default.	Simulation educator
Assistive technology prototype	A tabletop mock-up of an adaptation discussed with an occupational therapist and the intended user. **Any trial with the user follows the team's review process.**	Occupational therapist **and** the intended user
Software or instrumented training tool	An exercise interface using fictional data, or a battery-powered instrumented model — a sensor and microcontroller giving the learner feedback — with a clear learning objective and no connection to a person or to live clinical systems.	Educator, plus information governance if any data question arises

Whichever route you choose, the making methods are open: card, wood, metal, fabric, silicone, glass or acrylic, printed and machined parts, electronics, and code are all legitimate, and most good projects use several. The choice of method comes **after** the objective, and is recorded with its reason (Chapter 14).



If you have no partner yet, complete a clearly labelled **independent demonstration project**. Use the elbow model with a fictional brief. You can finish all the technical work and keep every sponsor-dependent criterion recorded as **pending**.

That is an honest and perfectly respectable portfolio. "I completed this

> independently; items 4, 7 and 9 await educator review" is a far better position than a project whose review was imagined.

22.3. Week 27: confirm the brief

22.3.1. The brief you will build against

> **Need.** Beginning learners need to distinguish a force **source** from a tendon's **transmission** role during a simplified elbow movement.
>
> **Scope.** A manually operated tabletop teaching aid. It demonstrates bending and returns to its starting position by manual repositioning. It does **not** attach to a person, support body weight, model a patient's range of motion, or reproduce every elbow structure.
>
> **Initial success criteria.** Movement is visible. Labels distinguish force generation from transmission. An unfamiliar learner can attempt the Chapter 21 questions. The model can be reset without tools.

These are learning-project criteria to discuss with an educator, not device specifications. Get them corrected before building (Chapter 18).

22.3.2. Confirm, in writing, before you cut anything

Item	Must be settled
The need	Stated without a solution in it, and confirmed by the sponsor
The scope	Including an explicit list of what it is **not**
Relevant anatomy	Which structures, at what level of simplification
Reviewers	Named people, by role: educator, and anyone else per Chapter 18
Requirements	Three to five, measurable, each with an acceptance criterion agreed **now**
Contact category	From Chapter 5. For this project: **no body contact**
Cleaning and storage	Who, how, where — agreed with whoever will host it
Ownership after handover	Named person or team
What "finished" means	The specific deliverables and the date

22.3.3. Compare two concepts, and record why

Do not design one thing. Design two, compare them explicitly, and record the reason for the choice. This takes an hour and is the difference between a decision and a default.

Criterion	Concept A: open card + cord	Concept B: enclosed fabricated mechanism (machined, moulded, or printed housing)	Comment
Makes the force path visible	Strong	Weak	**Decisive for this objective**
Build time	Hours	Days	—
Iteration cost	Trivial	High	Matters, because the first design will be wrong
Durability	Low	High	Not yet a requirement; no agreed service life
Cleanability	Poor (card)	Better (if designed for it)	Not yet required — tabletop only
Perceived "quality"	Low	High	**Irrelevant to the objective**; beware of being influenced by it
Risk of teaching the wrong thing	Lower — mechanism exposed	Higher — mechanism hidden	**Decisive**

Choice: Concept A, because the objective is the visibility of a force path, and the enclosed version hides exactly the relationship being taught. Durability and finish are not yet requirements and should not drive the decision.

Write that paragraph down. In three months you will not remember why, and neither will anyone reading your portfolio.

22.4. Weeks 28–29: build the first version cheaply



Begin with card, even though you have a workshop.

You have a well-equipped workshop — machine tools, a laser cutter, a printer, a sewing machine, an electronics bench — and the temptation is to use it. Do not, yet. The first question is whether the **relationship** can be understood, and that question is answered faster and more cheaply in card. A quick version lets you change the layout five times before you invest in a polished object that is polished in the wrong shape.

Moving to better fabrication is covered below, and it needs a **reason**.

22.4.1. Suggested materials

Dimensions are convenient classroom suggestions, **not anatomical measurements**.

- Two stiff card strips, approximately 200 × 35 mm.
- A paper fastener, or another lightweight freely rotating pivot.
- A short length of non-elastic cord.
- Tape, paper labels, a marker.
- A small smooth loop through which the cord can run — a bent paperclip or a small eyelet.

Use ordinary workshop precautions for cutting and hole-making.

22.4.2. Build sequence

1. **Make the two segments.** Mark a pivot hole about 20 mm from one end of each strip. Join them loosely enough to rotate freely. **Round or cover sharp corners and protrusions.**
2. **Orient the representation.** Support the upper-arm strip vertically above the pivot, with the other extending horizontally as the forearm. Label both as **schematic representations**, not anatomically accurate individual bones.
3. **Provide an attachment.** Attach one end of the cord to the forearm strip about 40 mm from the pivot. Reinforce the card locally — a washer, a folded tab, a strip of tape on both sides — because this is where it will tear.
4. **Guide the pull.** Fit the smooth guide loop to the upper-arm strip about 80 mm above the pivot. Pass the free cord end through it. Pulling gently draws the forearm attachment toward the guide and bends the model.
5. **Check the movement.** Operate it slowly through a modest visible arc, avoiding pinch points. Stop before the strips collide or the cord jams. **Do not add weights for the first demonstration.**
6. **Label the mechanism.** This is the most important step in the build:

Element	Label
Free pull end	"External force input representing muscle action"
The cord	"Tendon representation: transmits force, does not generate it"
The pivot	"Simplified joint — omits cartilage, capsule, ligaments, and forearm rotation"
The guide loop	"Model convenience, not an anatomical tendon route"
The whole object	"V0.1 — tabletop education prototype. Schematic, not to scale."

7. **Reset by hand.** State explicitly that manual reset is an **external** action. The model does not represent the muscle activity involved in extension.
8. **Version it.** Add "V0.1", the date, and "schematic, not to scale". Photograph it if useful and record the full build details per Chapter 19.



The guide loop and attachment locations create a convenient force path. They are **not** an anatomically accurate biceps insertion. That distinction belongs in the label and in your spoken explanation.

If the educator wants to teach **actual tendon attachments**, this design needs substantial revision and anatomical review. That is a different project, not an adjustment.

22.5. Troubleshooting: ask what changed

Observation	Possible explanation	Useful next check
The strip barely moves	Pivot friction, or the cord's line of action is nearly through the pivot	Move the strip by hand without the cord. Then examine the cord's line of action and its perpendicular distance from the pivot (Chapter 6)
The card tears at the attachment	Load concentrated at a point in a weak material	Inspect the attachment. Reinforce and distribute the connection over a larger area
The cord slips out of the guide	Guide geometry does not retain the cord through the full travel	Observe at which angle it escapes. Revise the guide. **Re-test from the beginning on the new version**
The learner calls the cord a muscle	The force source and the transmission element are visually merged	Separate their labels and appearance. Ask a **different** learner to explain it (Chapter 21)
The reset is confusing	The starting configuration is not evident from the object	Add a start-position drawing or a physical mark. Observe an **unprompted** reset
Motion looks right but the force "feels wrong"	Moment arm changes through the sweep	Expected behaviour — but say so in the explanation rather than letting it confuse
It works for you and not for the educator	You know things about it that the object does not communicate	This is the most valuable failure you will get. Find out what you knew



Change one important feature at a time where practical. Otherwise a successful revision leaves you unsure which change helped — and you will carry forward a change that did nothing while attributing the improvement to it.

Record **unexpected** effects too, including ones unrelated to the problem you were solving. Those are often the most informative observations you will make.

22.6. Run a short bench review

Inspect the pivot, edges, cord, and attachment. Perform the movement-and-reset sequence repeatedly at gentle manual force. Choose and record a modest trial count — say ten or fifty cycles — explicitly as an **exploratory check**. Stop if a part fails or creates an unsafe handling condition.

Write the result in this form:

> "V0.1 completed **[number]** of **[planned number]** cycles under **[stated conditions]**. We observed **[facts, including anything unexpected]**. This supports **[a narrow conclusion]**. This does **not** establish **[the things it does not]**. We still need **[the next evidence]**."



Never fill in a successful result before doing the work.

It sounds absurd written down. It happens constantly, usually not as deliberate falsification but as a template filled in optimistically "to be confirmed", and then never confirmed. Write the record **after** the test, or write it before with the result field conspicuously blank.

22.7. Review the learning claim

Two reviews, in this order:

1. Educator review of the anatomy. Ask an educator to review the anatomy, the labels, and the stated limitations. Record their corrections verbatim. If they say the labelling could still be misread, that is the finding — do not argue it away.

2. Learner trial. In an appropriate authorised setting, invite a willing learner to attempt the Chapter 21 task. Record:

- their **prediction** before the demonstration,
- their **explanation** afterward,
- **every prompt you gave**, verbatim,
- any **confusion or misconception**, in their words,
- whether they named a limitation of the model.

Record no personal information you do not need. A learner is "Learner 1", not a name.



A partial result is a result. A likely outcome is that the learner understands flexion and still thinks the cord supplies its own energy. That is not a failed project. It is a precise, actionable finding: the force source and the transmission element remain visually merged.

Revise the relevant representation, explain the change in your record, and check it again **with a different learner** — because the first one now knows the answer.

22.8. Decide when better fabrication is justified

Move beyond card when you have a **reason**:

- the card cannot survive the intended handling, demonstrated by a failure you recorded;
- a clearer or more precise geometry is needed, and you can say which;
- an agreed classroom requirement calls for a stable assembly;
- a cleaning requirement has emerged that card cannot meet (Chapter 13);
- the learner trial showed something the card version cannot make visible — for example, the force source and the cord still look like one thing.

"It would look more professional" is not a reason. A complicated part is not inherently a better teaching aid, whatever made it, and a polished object can teach worse than a crude one by hiding its mechanism (Chapter 21).

22.8.1. Choose a method for each element, not one for the whole

The elbow model is several elements with different jobs, and each has its own best method. Work through them and record the choice and the reason. Chapter 14 covers the methods in depth.

Element	Candidate methods	Choose it when	Watch for
Upper-arm and forearm segments	Laser-cut plywood or acrylic; machined aluminium bar; printed polymer; injection moulded (many copies only)	Laser-cut: flat parts, fast iteration, labels engraved in. Machined: durability and cleanability. Printed: a shaped segment the objective needs. Moulded: tens or hundreds of identical kits justify the tooling	Sharp acrylic edges and burrs on metal; unsealed plywood cannot be cleaned; printed layer lines and porosity; tooling cost and lead time for moulding
Pivot	Turned pin and bushing (brass, stainless steel, or acetal); shoulder bolt; printed pin	Low, repeatable friction and long life — almost always a turned or bought part	Wear and slop change the motion; a pinch gap at the joint (Chapter 19)
"Muscle" — the force source	Sewn fabric sleeve; cast silicone segment; a coloured handle	The learner trial showed the force source and the cord are being merged	It must remain visibly **separate** from the cord; fabric frays and absorbs; silicone needs a mould, which is itself a small project
"Tendon" — the transmission	Low-stretch braided cord; stainless cable with crimped ferrules	Cord for a light classroom model; cable where stretch or abrasion becomes a recorded failure	Sudden release under tension; crimp quality; a cut cable end is sharp

Element	Candidate methods	Choose it when	Watch for
Guide loop	Turned or bent metal eyelet; smooth printed or moulded loop	Low friction and no fraying point	A crevice that traps dirt; abrasion of the cord
Labels and cues	Laser-engraved or painted text and arrows; tactile features; contrasting materials	Always — redundant coding (Chapter 7)	Labels that wear off or fail cleaning; engraved lettering sized for the viewing distance (Chapter 8)
Optional electronic feedback	Angle sensor (potentiometer or Hall-effect) with an LED bar or small display, on battery power	Only if the objective needs a quantitative angle, agreed with the educator	It must not supply force or hide the mechanism; new failure modes, firmware, battery and cleaning questions (Chapter 19, Chapter 20)

A plausible V1.0, for example: laser-cut acrylic segments with engraved labels, a turned acetal pivot bush on a stainless pin, a sewn red fabric "muscle" sleeve clearly separate from a white braided cord, and no electronics. Every one of those choices should carry a sentence saying why — and one of them will turn out to be wrong, which is what V1.1 is for.



A material change is a new version, and it changes more than you expect.

New weight, new edges, new friction at the pivot, new fastening, new failure modes, new cleaning requirements, new cost to iterate, and possibly new visibility of the mechanism.

Treat it as V1.0, not as V0.4, and **repeat every check the change could have affected**. Do not carry forward the card version's results.

22.9. Week 30: review, revise, hand over

22.9.1. Assemble the handover

A future volunteer — or you in a year — should be able to pick this up and know what it is for and what has actually been established.

Item	Content
Need statement	Who, what difficulty, what outcome. No solution in it
Learning objective	Observable, with the matching assessment
Authorised scope	What it is for and explicitly what it is not for
Labelled photographs or drawings	With the anatomy or task annotated

Item	Content
Drawings and design files	Dimensioned drawings; CAD models; laser, CNC, and print files; sewing patterns; mould designs; schematics and PCB files; firmware and software source — all with revision identifiers
Bill of materials	Every part and material, with grade, supplier, and quantity; bought parts by part number
Materials and assembly notes	Enough to rebuild it, including the order of operations
Process notes	For each process: machine and settings, cure schedules, stitch and seam details, solder and connector choices, the firmware version flashed — anything non-default
Version history	What changed, when, and **why**
Requirements and acceptance criteria	As agreed, with dates
Test records	Including failures and unresolved items
Feedback log	Educator and learner responses, with prompts recorded separately
Limitations statement	The full list of deliberate omissions
Current status	The honest one-paragraph version, below
Next decision, and who makes it	Revise, stop, continue bench work, or seek specified review

Use photographs only where permitted, and include no identifiable patient information.

22.9.2. The status statement

> "**Tabletop demonstration prototype, V1.0.** It shows a selected bending motion and a simplified force-transmission relationship. It omits forearm rotation, exact attachment anatomy, active extension, and all living tissue behaviour. Bench and educational review results are recorded separately. Shared-classroom acceptance remains **pending** infection prevention review."

Update this statement only when the **evidence** changes — not when your confidence changes.

22.9.3. The five-minute explanation

You must be able to cover, in five minutes, to someone who has not followed the project:

1. The **need**, and who said so.
2. The **relevant anatomy**, correctly.
3. The **design decision** you made and the alternative you rejected, with the reason.
4. The **evidence** you have, including at least one failure.

5. The **limitations** — at least three.
6. The **next decision** and who makes it.

Then ask them: **"What's the biggest unresolved question here?"** Your ability to welcome that question, and to act on it rather than defend against it, is a substantial part of what readiness for collaborative work actually means.

22.10. Workshop exercise

1. **Week 27.** Write the one-page brief. Get it corrected by the sponsor, or mark it pending. Compare two concepts in a table and record the choice with its reason.
2. **Weeks 28–29.** Build V0.1 in card. Label every element per the table above. Troubleshoot. Run the bench review and write the result in the stated form, including the failure.
3. **Weeks 29–30.** Obtain educator review of the anatomy and labels. Run one learner trial. Record predictions, explanations, prompts, and confusions separately.
4. **Week 30.** Revise based on what you found. Decide — with a reason — whether better fabrication is justified. If yes, choose a method for each element with a recorded reason, build V1.0, and repeat the affected checks.
5. **Week 30.** Assemble the handover pack. Then hand it to someone who has not seen the project and ask them to tell you what the object is for and what has been established. Every gap they hit is a documentation defect.
6. **Week 30.** Deliver the five-minute explanation. Record the question you were asked and what you did about it.

22.11. Interactive tools

Browser tools on the companion website that go with this chapter. They are for learning: none of them is a medical device, and none is for assessing, diagnosing, treating or dosing anyone.

- Levers, moment arms and contact — the capstone geometry: moment arm and cord force across the sweep.
- Project worksheets — the project file, with the capstone as the worked example.
- Print tolerance — printed fits.

22.12. Exit criteria

- ▢ A sponsor confirms that the prototype addresses the agreed need — or this is explicitly recorded as **pending**.
- ▢ All agreed tests have recorded outcomes, **including failures and unresolved requirements**.
- ▢ The design can be reconstructed from your documentation alone, verified by someone else attempting it.
- ▢ You can give a five-minute explanation covering need, anatomy, design decisions, evidence, and limitations.

- ▢ You can state at least three limitations of the model unprompted.
- ▢ Educator and learner feedback are recorded, with prompts noted separately — or explicitly pending.
- ▢ The version history explains **why** each change was made, including why each element is made the way it is.
- ▢ The sponsor records the next step: revise, stop, continue bench work, or seek a specified further review.
- ▢ **No prototype is described as approved for patient use** unless the responsible organisation has actually authorised that use.

22.13. What this chapter does not establish

- **A completed capstone is not a qualification.** It is portfolio evidence of your process, which is exactly what it is meant to be.
- **A sponsor's confirmation that it meets the need is not approval to deploy.** Whether an object may be used in a setting is a separate institutional decision.
- **Your bench results characterise your object under your conditions.** Nothing more.
- **Finishing does not mean you may now work unsupervised.** It means you are ready for the conversation in Chapter 23.
- **The honest version is the valuable one.** A portfolio recording three failures, two pending reviews, and one unresolved requirement is more credible — and more useful to the person deciding whether to work with you — than one in which everything passed.

Chapter 23. Next steps, specialisms, and finding a place to contribute

23.1. Why this chapter exists

You have a portfolio. Now you need somewhere to use it, and a decision about how deep to go.

Both are easier than they look, and both are mostly a matter of asking specific questions of specific people rather than broadcasting general enthusiasm.

23.2. Finding a place to contribute

23.2.1. Where to look

Setting	What they do	Why they might want you
Hospital simulation centre	Runs clinical training with manikins, task trainers, and scenarios	Trainers wear out and need repair, modification, instrumentation, or replacement parts; scenarios need props, signals, and software; bespoke fixtures are constantly needed. **This is usually the easiest door.**
Occupational therapy / rehabilitation	Adaptive equipment, function in daily life	Custom adaptations are their daily work and commercial options frequently do not fit. Solutions span textiles, machined and moulded parts, electronics, and environmental controls
Clinical engineering / medical physics	Maintains and manages equipment	Test rigs and instrumentation, electronics repair and interfacing, machined jigs and fixtures, mounts, replacement parts for obsolete equipment, asset and maintenance software
Prosthetics and orthotics	Designs and fits devices to bodies	Already a making profession — thermoforming, lamination, machining, casting, textiles, scanning and additive manufacturing, sensors in limbs and braces
Innovation / R&D office	Brokers external collaboration	Their actual job is to connect people like you to people like them
Medical school / nursing school	Teaches anatomy and clinical skills	Anatomical models, teaching aids, demonstration rigs, interactive and instrumented teaching tools

Setting	What they do	Why they might want you
Research groups	Device development, biomechanics, imaging	Engineering capacity — mechanical, electronic, and software — is a perennial bottleneck
Hospice and community services	Care in home and community settings	Highly individual adaptations with little commercial provision

23.2.2. Organisations that exist for this

Organisation	What it is
Tikkun Olam Makers (TOM) [18]	An international community matching makers with people who have unmet needs ("need-knowers"), running makeathons and maintaining open designs. A well-structured route into collaborative assistive technology [19].
REMAP [31]	A UK charity whose volunteer engineers and makers build one-off custom equipment for disabled people, free of charge. Particularly well matched to practical fabrication skills, with an established referral and safety process.
UTMB MakerHealth [32]	An example of a hospital-embedded maker space. Access is described for staff and students; useful as a model of what such a facility looks like.
Local makerspaces and hackspaces	Often already have healthcare contacts, and sometimes run assistive technology groups.

Availability, eligibility, and local arrangements all need checking. Hospital maker facilities are frequently restricted to staff and students.

23.2.3. What to actually say

The failure mode is a general offer of enthusiasm, which is hard to act on and so gets filed. Be specific, be small, and be clear about supervision.

> **Subject: Volunteer engineer — small supervised project?**
>
> Dear [name],
>
> I'm an experienced technologist with software, electronics, and a full fabrication workshop (CAD, machining, metalwork, moulding and casting, glasswork, textiles, additive manufacturing). Over the past several months I've worked through a self-study foundation in anatomy, physiology, infection control, human factors, and device design process, and built a small documented teaching prototype — an elbow model demonstrating the distinction between muscle force generation and tendon force transmission. The write-up, including its test failures and its limitations, is attached.
>
> I'm looking for a **small supervised bench project** where those skills would be useful. I am

not looking to work with patients, and I understand that anything shared or reusable needs infection prevention review.

Three questions:

1. Do you accept external technical volunteers, and what onboarding would be required?
2. Is there a small, low-risk problem where an engineer with those skills would save you time?
3. If not you, is there someone else I should ask?

I'd be glad to come in for twenty minutes if that's easier than email.

Why that works:

- It says **what you can do**, concretely.
- It shows **evidence**, not claims.
- It **names its own limits** before being asked — which is the single most reassuring thing you can do.
- It asks a **small, specific, answerable** question.
- It offers an **exit**: "if not you, who?" — which converts a dead end into a referral.



Attach the portfolio that records failures. Counter-intuitively, a write-up that says "this failed at attempt seven and here is what I changed" is far more persuasive to a clinical professional than one where everything worked. It demonstrates that you report accurately, which is the quality they actually need from a collaborator and the one they cannot verify any other way.

23.2.4. Realistic expectations

Expect	Because
Slow replies	Clinical staff are busy and email is not their primary channel
Being redirected, repeatedly	You will rarely reach the right person first. Each redirect is progress
Onboarding requirements	Honorary contracts, occupational health, DBS or equivalent checks, mandatory training, induction. Weeks to months
A first project smaller than you hoped	Good. A finished small thing builds the trust that leads to a bigger one
Some doors simply closed	Policy, insurance, or capacity. Not personal
The question "who owns this?" early	It is the right question and they are right to ask it

23.3. Choosing a specialism after your first project

Depth comes after a project, not before. You now know which of these you actually enjoyed.

Direction	Deepen	Exit evidence
Assistive technology	Upper-limb anatomy and function, ergonomics and anthropometry, accessibility, therapist-led functional assessment, participatory design	Explain how **real user feedback** changed a design, and how suitability will be reviewed
Simulation and training	Relevant anatomy, educational design and assessment, material behaviour and tissue simulants, repeatability and maintenance	An educator-reviewed aid with an objective, an assessment, and a maintenance procedure
Medical software	Usability and accessibility, data governance, software quality processes (IEC 62304), interoperability (HL7, FHIR, DICOM), human factors in interface design	A documented training tool using fictional data, with tested error handling and a dependency inventory
Clinical engineering and instrumentation	Equipment management, electrical safety testing, test and calibration rigs, embedded systems, IEC 60601 basics, maintenance documentation	A documented test rig or fixture accepted by a clinical engineering team, with its calibration and limitations recorded
Design for manufacture and production	Process selection, tolerancing, injection moulding and machining for medical parts, textiles and soft goods, ISO 13485, process validation, supplier control	A design taken from one-off to a documented small batch, with process notes and inspection records
Anatomical models from imaging	Imaging modalities and orientation, DICOM handling, segmentation, scale and accuracy, anatomical verification with a specialist	An educational model whose source, processing, and **limitations** are documented and reviewed
Instrumentation and biosignals	Signal conditioning, amplifier design, noise and artefact, sampling and filtering, electrical safety (IEC 60601-1)	Analysis of **recorded or synthetic** data with a documented processing chain. **Never on a person**
Biomechanics	Statics and dynamics of the musculoskeletal system, inverse dynamics, gait analysis, materials and fatigue	A documented analysis with stated assumptions and a sensitivity check
Rehabilitation engineering	Seating and positioning, mobility, wheelchair and seating assessment, pressure management	Work alongside a clinical specialist; this field is heavily clinical



On imaging-derived models specifically. A reconstruction from a CT or MRI scan can look extraordinarily convincing and be wrong — through segmentation

threshold choices, partial volume effects, motion artefact, scaling errors, or simply a mis-set orientation.

Do not treat an attractive reconstruction as clinically accurate without validation by someone qualified. And imaging data is patient data: information governance applies from the first file.

23.4. Going further: formal study

If this turns into a career direction rather than a volunteering interest:

Route	Note
Biomedical engineering degree (BEng/MEng/MSc)	The standard route. An MSc conversion is realistic for someone with an existing technical degree. Core content: biomechanics, biomaterials, instrumentation, signals and systems, physiology, device design and regulation
Clinical engineering training schemes	Some health systems run structured training programmes for engineers working in healthcare — in the UK, for example, the NHS Scientist Training Programme
Regulatory affairs / quality qualifications	A distinct and employable specialism. Short courses and professional certifications exist and are often employer-funded
Human factors in healthcare	A growing field with its own societies, courses, and literature. Highly relevant to anyone who designs interfaces, controls, or physical products
Professional bodies	IPEM (UK), AAMI (US), IEEE EMBS (international), BMES (US). Membership gives access to literature, standards summaries, events, and people
Reading standards directly	ISO 14971, IEC 62366-1, IEC 62304, ISO 13485. Often available through a university or public library. Reading the **scope** and **terms and definitions** sections of each is a cheap and high-value afternoon



Have a concrete conversation about your portfolio rather than collecting unrelated certificates.

A finished, honestly documented small project plus a clear statement of what you do not yet know will open more doors than a list of completed MOOCs. Certificates demonstrate attendance. A portfolio demonstrates judgement, and judgement is the thing that is genuinely hard to acquire.

23.5. A compact resource shelf

Keep these, and resist the urge to build a library before you need one.

Resource	Use
OpenStax Anatomy and Physiology 2e [1]	Your main free reference for diagrams and deeper explanation. Keep it bookmarked
Crash Course Anatomy & Physiology [6]	Short video companion. Good for a first pass and for revisiting one topic
Michigan Anatomy Specialization [2]	Optional structure if you want longer lectures. Check current access and payment terms
Stanford Student Guide to Biodesign [14]	Keep it beside your project notes for need-finding and prototyping
Biodesign (book, 2nd ed.) [17]	Borrow when you want deeper development guidance. Use current official sources for regulatory detail
FDA and CDC pages cited here	Concept references. Your host determines what actually applies
Winter, Biomechanics and Motor Control of Human Movement [4]	If you go toward biomechanics
A specialist atlas	Only when a specific project justifies it. Not before

23.6. Exit criteria for this stage

- ▢ Identify **one target role**, **one potential host**, and **one specific knowledge gap** to address next.
- ▢ Send at least three specific outreach messages, with your portfolio attached.
- ▢ Have at least one concrete conversation about your portfolio with someone in the field.
- ▢ Be able to state, in one sentence each: what you can do unsupervised, what you can do supervised, and what you should not do at all.
- ▢ Choose a specialism direction, with a reason drawn from your capstone experience.

23.7. A closing note

The temptation after finishing a programme like this is to feel that you now know enough to be useful without supervision. You do not, and neither does anyone else at this point — including people with degrees, which is why clinical engineering, prosthetics, and device development all have supervised training structures.

What you **do** have is considerably more valuable than it sounds: the ability to recognise which questions are not yours to answer, and the habit of writing down what you actually established rather than what you hoped.

Those two things make you safe to work with. Safe to work with is how you get the second project.

Appendices

Appendix A: Study schedule and assessment method

A.1. The overall shape

Plan for approximately **130 to 190 hours** of core study and practical work over about **30 weeks at 4–6 hours per week**. The times include selected reading and the exercises, not completion of entire textbooks or courses.



Arranging a clinical partnership routinely takes longer than the study itself. That is why Chapter 18's outreach exercise appears in week 1 of the schedule below rather than week 21. Send the first email before you have finished Chapter 1.

A.2. Week by week

Weeks	Chapter	Focus	Evidence to keep
1–2	Chapter 1	Body language, levels of organisation, cells, homeostasis	Annotated body map; 20 vocabulary cards; terminology check with corrections
3	Chapter 2	Atoms and the periodic table, isotopes, gases and partial pressure, moles and concentrations, reactions, energy, kinetics, equilibrium, redox and electrochemistry, organic functional groups, chemical safety; the branches of chemistry and their instruments	Unit-checked concentration calculations; dilution-series calibration curve from a home-built or phone colorimeter; galvanic-cell and corrosion observations; plastic identification notes; SDS-based exposure assessment for five workshop substances; accelerated-ageing calculation
4–5	Chapter 3	Bonds, water, pH and buffers, macromolecules, enzymes, bioenergetics, polymer and surface chemistry	Bilayer model; pH measurements; physical molecule model with limitations statement; retained control sample

Weeks	Chapter	Focus	Evidence to keep
6	Chapter 4	Life and evolution; prokaryotic and eukaryotic cells; organelles; gene to protein and its regulation; cell cycle and stem cells; bacteria, viruses, immune cells; the branches of biology, their real-world uses, and the laboratory tools of each; biosafety	Labelled cell sketch from memory; diffusion, growth, PCR, and RCF calculations; DNA extraction and diffusion results; translated **INS** gene; cell count with threshold sensitivity; one real-world workflow mapped instrument by instrument
7	Chapter 5	Tissues, skin, contact mechanics, viscoelasticity	Layered skin model; pressure calculation with its limits; contact-design question list
8–9	Chapter 6	Bone, joints, movement, muscle, torque and moment arms	Articulated joint model; torque calculations; list of model omissions; functional brief
10	Chapter 7	Nervous system, sensation, proprioception, accessibility	Task decomposition; accessibility review with two documented changes; redundant coding
11	Chapter 8	Special senses; eye, ear, nose, throat, and mouth; swallowing; optics and acoustics	Light and sound path diagrams; airway–food crossing sketch; legibility calculation for a real label; sound-cue measurements at three distances; face pressure map
12	Chapter 9	Heart, circulation, breathing, gas exchange, vital signs	Annotated flow diagrams; MAP/CO calculations; mock monitor critique
13	Chapter 10	Kidneys and urinary tract, digestive system, endocrine glands; clearance and dialysis; hormone loops as control systems; units	Labelled nephron and GI tract drawings; clearance calculations including a unit error; bench dialysis results; toy glucose-loop plots labelled as such; unit-carrying glucose type with tests

Weeks	Chapter	Focus	Evidence to keep
14	Chapter 11	Cell injury, inflammation, infection and sepsis, shock, thrombosis, atherosclerosis, heart failure, airway disease, diabetes, cancer, autoimmunity, degeneration, genetics, risk arithmetic	Decompressed diagnosis with eight user implications; Virchow review of a fluid path; diabetic-foot diagram from memory; headline rewritten as ARR and NNT; worked pedigree
15	Chapter 12	Children, pregnancy, older adults, frailty and falls; cognition, cognitive load, dementia and delirium, mental health, neurodiversity, stress; ligature risk	BSA and airway calculations; fall-energy calculation with home hazard list; cognitive-load audit with redesign; tested easy-read rewrite; population review of one made object
16	Chapter 13	Infection control, cleaning/disinfection/sterilisation, materials across processes — machined, moulded, cast, printed, metal, glass, and textile	Contact map; three-material comparison with unknowns marked; traceability record
17	Chapter 14	Choosing the technology: alternatives to fabrication; software, electronics, additive, subtractive, moulding and casting, metalwork, glasswork, textiles; prototype to product; cost, tolerance, surface, battery, and capability arithmetic	Decision-diagram records for three needs; four-family method comparison; break-even plot; wipe-test results on five materials; C_{pk} for one dimension; a handover pack used by someone else
18	Chapter 15	Imaging modalities, the laboratory, measurement quality, diagnostic performance	Imaging-to-physical-model pipeline with every setting recorded; threshold sensitivity figures; ultrasound phantom; PPV calculation

Weeks	Chapter	Focus	Evidence to keep
19	Chapter 16	X-ray, CT, MRI, ultrasound, nuclear medicine and PET, endoscopy, physiological tests; radiation dose; MR safety; choosing a test	Two-sentence mechanism for each modality; HU measurements in a public CT; Doppler worked example; k-space exercise; MR-room materials review
20	Chapter 17	Ventilation, infusion, defibrillation and pacing, dialysis, theatre energy and fire, radiotherapy, implants, drug delivery, rehabilitation devices; alarm fatigue and incident lessons	Respiratory time-constant calculation and RC demonstration; unit-safe infusion calculator with tests; repeated-dosing plot; emergency-ventilator requirements table with the "instead" answer; inhaler task analysis; CPR/AED course certificate
21–22	Chapter 18	Clinical collaboration, need discovery, interviewing	Task map; need statement; interview notes; one-page brief; constraint sweep
23–24	Chapter 19	Requirements, hazards, risk control, testing, traceability	Hazard table; three requirements with acceptance criteria; versioned test record including a failure
25	Chapter 20	Software as a medical device, IEC 62304, classic failures, health data and de-identification, interoperability (FHIR, codes, units), cybersecurity, AI evaluation	IMDRF placement of five apps; synthetic FHIR pipeline with explicit units; de-identified DICOM with notes; k-anonymity calculation; STRIDE table and SBOM for one project
26	Chapter 21	Learning objectives, fidelity, teaching design, assessment	Observable objective and matching assessment; teaching script; two checklists; limitations statement
27–30	Chapter 22	Capstone: brief, build, test, review, handover	Full portfolio: brief, drawings, files, version history, test records, feedback log, status statement

Weeks	Chapter	Focus	Evidence to keep
after	Chapter 23	Outreach and specialism choice	Three sent outreach messages; one concrete conversation; chosen direction with a reason

Allow extra weeks wherever you need them. The schedule is a plan, not a deadline. The goal is retained understanding and a reviewable project, and both are harmed by rushing.

Two weeks each on the chapters with the most to absorb (Chapter 1, Chapter 3, Chapter 6) or the most that depends on other people (Chapter 18, Chapter 19), and four on the capstone. Chapter 2, Chapter 3, Chapter 4, Chapter 8, Chapter 10, Chapter 11, Chapter 12, Chapter 15, Chapter 16, Chapter 17, Chapter 20, and the device sections throughout are new material relative to the original programme — if you are short of time, Chapter 3's engineering bridge and Chapter 15's measurement-quality section are the parts that pay back fastest.

A.3. Assessment method

A.3.1. Per chapter

1. **Explain aloud without notes**, same day. Five minutes. To a person, a camera, or an empty room.
2. **Answer the chapter questions closed-book.** Aim for at least 8 in 10 correct.
3. **Correct every error in writing.** Not mentally — in writing, with the reasoning.
4. **Complete the practical exit criteria.** These are not optional and a quiz score does not substitute for them.
5. **Re-attempt the missed items one week later.**

A.3.2. Where another person is required

Several exit criteria require a clinician's or educator's review. Where that person is not yet available:

- Record the criterion as **pending**, with the date and what is needed.
- Continue the technical work.
- **Do not** mark it complete, and **do not** invent the feedback.

A portfolio with three clearly marked pending items is credible. One with imagined agreement is not, and the difference will be apparent to anyone who reads it carefully.

A.3.3. Final assessment

Appendix E contains a twelve-question final assessment with answers, and a readiness review to

discuss with a potential host.

A.4. A repeatable weekly routine

Time	Activity	Detail
~1 h	Watch and note	The chapter's companion video. Once for shape, once slower while drawing
~1–2 h	Read and draw	This book's chapter plus the named OpenStax sections. Redraw every diagram from memory at least once
~2 h	Build or apply	The workshop exercise. Hands on something physical or on code
~30 min	Explain aloud	Without notes. This is where you discover what you only recognise

Pause fast videos and replay difficult passages. Watching at a slower speed is a legitimate technique, not a weakness.

A.5. Spaced repetition

When	What	Form
Same day	Explain the chapter aloud	No notes, 5 minutes
+1 day	Ten-question check	Written, closed book
+1 week	Re-attempt only what you missed	Written
+1 month	Re-draw the chapter's main diagram	From memory, on paper
On project need	Re-read one section with a specific question	Targeted, not a re-read of the whole thing

A.6. Using one main textbook

OpenStax Anatomy and Physiology 2e [1] is the core free reference. Use its contents menu to reach the sections named in each chapter. Read chapter summaries first, then the relevant explanations and diagrams.



The chapter assignments in this book are a **selective study route**, not a recommendation to memorise the entire textbook. Reading all of OpenStax cover to cover would take far longer than this programme and would not make you a better collaborator. Depth where a project needs it; breadth only enough to know where to look.

Optional structured video course: the University of Michigan Anatomy Specialization [2]. Use it when you want longer explanations or external course structure. It should **replace** some video and reading time rather than becoming an additional obligation.

A.7. Cost and access

Most core resources are free to view at the time of writing.

Resource	Access
OpenStax	Free, openly licensed
Crash Course	Free; links lead to the creator's pages containing the videos
MedlinePlus	Free
CDC and FDA material	Free
Stanford Student Guide to Biodesign	Described as free; toolkit and video pages may require registration or login
Coursera (Michigan)	Access and certificates may involve payment
Optional books	Borrow them. You do not need to buy a library before starting
Standards (ISO, IEC)	Paid, but often accessible through a university or public library
3D Slicer, OpenSim, PyMOL, NVDA, Blender	Free and open source
TCIA, PhysioNet, Protein Data Bank	Free, open, properly cleared datasets



The software and datasets in Appendix I are all free and run on an ordinary laptop. They are the cheapest way to get genuine hands-on experience in imaging, biosignals, biomechanics, and molecular structure — with no patient contact, no materials cost, and no approvals needed.

Appendix B: Reusable worksheets

Copy these into your notebook, or keep them as files under version control. Complete a version for each project and **retain the earlier versions** — your reasoning over time is part of the evidence, and a superseded worksheet showing what you thought in week 3 is genuinely useful in week 14.

B.1. Learning record

Complete once per chapter.

Chapter and date	
Resources completed	
Three concepts I can explain without notes	
One diagram or object I made	
Questions I answered incorrectly	
Corrected explanation for each	
Exit criteria complete	
Exit criteria pending, and what is needed	
Feedback needed from a clinician or educator	
Date for one-week recall check	
Date for one-month diagram redraw	

B.2. One-page project brief

Complete before building anything. Get it corrected by the collaborator.

Need statement (who, what difficulty, what outcome — **no solution**)	
Who has the need	
Who will use the result	

The task, and how it is performed now	
Current workaround, and what it costs	
Desired improvement, stated as an outcome	
Relevant anatomy or clinical context	
Scope: what this is	
Scope: what this is explicitly NOT	
Contact category (Chapter 5) and the reviewer it implies	
Setting and environment	
Users, and their relevant characteristics	
Frequency, duration, and number of cycles	
Setup and reset time available	
Requirements (3–5), each with how it will be checked	
Acceptance criteria, agreed with whom, on what date	
Cleaning, storage, and maintenance questions (Chapter 13)	
Existing solutions, and why each was rejected	
Assumptions — explicitly labelled as assumptions	
Unresolved risks	
Clinical or educational sponsor	

Other reviewers required, by role	
Authorised scope and next review date	
Who owns and maintains this after I leave	

B.3. Evidence separation sheet

Use during and immediately after any observation or interview.

#	Observation (what was seen or said)	Interpretation (my explanation — and one alternative)	Assumption (untested prediction) and how I will check it
1			
2			
3			



Record **no patient identifiers**. "Learner 1", "Staff member A", "the educator". Nothing that identifies anyone individually, including in combination.

B.4. Task map

Stage	What happens	Who does it	How long / how often	What can go wrong, and what they do then
Trigger				
Preparation				
Action				
Exception path				
Cleanup				
Handoff				

B.5. Hazard analysis

#	Hazard	Sequence of events	Hazardous situation	Possible harm	Control (highest feasible level) and how its effectiveness is checked
1					
2					
3					
4					
5					

Include at least one **non-physical** harm — a misleading display, a misconception taught, a decision wrongly supported.

Control levels, in order: **1 inherently safe design → 2 protective measures → 3 information for safety**.

B.6. Requirement and acceptance criterion

One per requirement, agreed **before** any test.

Requirement ID	
Traces to which need	
Statement: who does what, under what conditions, to what criterion	
Acceptance criterion	
Agreed with whom, on what date	
How it will be checked (verification or validation?)	
What a pass would establish	
What a pass would NOT establish	

B.7. Test record

Prototype version and build date	
Requirement ID being tested	
The single question this test answers	
Acceptance criterion, agreed before the test, on what date	
Equipment, setup, and conditions	
Procedure (enough that someone else could repeat it)	
Who performed it; who observed	
Observed results, including failures and anything unexpected	
Prompts or assistance given — verbatim	
What the result supports (state narrowly)	
What the result does not establish	
Change made as a result	
Which other tests this change could invalidate	
Which of those were re-run, and their outcomes	
Reviewer and next action	

B.8. Version record

Version identifier (also on the object, durably)	
Date	

CAD / code revision or commit	
Materials, including batch or spool where relevant	
Machine, nozzle, profile, and any non-default settings	
Assembly details — fasteners, adhesives, order of operations	
Post-processing — sanding, washing, curing, coating, and agents used	
What changed from the previous version	
Why it changed	
Which tests from the previous version no longer apply	

B.9. Learning objective and assessment

For a teaching aid.

Learning objective (observable verb, specified learner, conditions, standard)	
Bloom level	
Matching assessment — at the same level	
Which fidelity dimension this aid serves	
Fidelity deliberately omitted, and why	
Teaching script: predict / observe / explain	
Observation checklist	

Reset procedure and **measured** reset time	
Maintenance and inspection points	
Educator who reviewed the anatomy, and date — or **pending**	
Misconceptions observed, in the learner’s own words	
Changes made in response	

B.10. Material comparison

Material	Intended setting	Fabrication method	Surface and geometry concerns	Cleaning questions, and who answers them	Evidence that is missing

Do not produce an overall ranking. Do not give any material a blanket "hospital safe" rating. The **missing evidence** column is the most important one.

B.11. Handover pack checklist

☐	Item	Note
☐	Need statement	No solution embedded
☐	Learning objective or functional outcome	Observable
☐	Authorised scope — what it is, and what it is not	Both halves
☐	Labelled photographs or drawings	Permitted images only; no identifiable information
☐	CAD / code files with revision identifiers	
☐	Materials list and assembly notes	

☐	**Item**	**Note**
☐	Build settings	
☐	Version history with reasons for each change	
☐	Requirements and acceptance criteria, with dates	
☐	Test records, including failures and unresolved items	
☐	Feedback log, with prompts recorded separately	
☐	Limitations statement	Every deliberate omission
☐	Current status statement	Honest; updated only when evidence changes
☐	Next decision and who makes it	
☐	Who owns and maintains it	

Appendix C: Physiological reference figures and units



Everything in this appendix is orientation, not a reference range.

These are typical adult textbook figures, collected so that you can check whether a number you have encountered is in the right decade. They vary with age, sex, build, activity, altitude, measurement method, and individual. Published "normal ranges" differ between sources and are usually constructed to contain roughly the central 95% of a reference population — which means around one in twenty healthy people fall outside them by construction.

You may not use these figures to assess anybody, to set a threshold, or to decide anything about a person. See Chapter 9.

C.1. Cardiovascular

Quantity	Typical resting adult	Note
Heart rate	~60–80 beats/min	Lower in trained individuals; rises with activity, anxiety, fever
Stroke volume	~70 mL	Volume ejected per beat
Cardiac output	~5 L/min	CO = HR × SV
End-diastolic volume	~120 mL	
End-systolic volume	~50 mL	
Ejection fraction	~55–70%	SV / EDV
Arterial pressure	~120/80 mmHg	Systolic / diastolic
Mean arterial pressure	~90 mmHg	≈ P_{dia} + □(P_{sys} − P_{dia})
Pulmonary artery pressure	~25/10 mmHg	About a fifth of systemic
Central venous pressure	~0–5 mmHg	
Capillary pressure	~35 mmHg arterial end → ~15 mmHg venous end	**Compare with external interface pressures**
Total blood volume	~5 L	~7% of body mass
Haematocrit	~40–50% (male), ~36–46% (female)	Proportion of blood volume that is red cells
Circulation time	~1 min	Whole volume circulates roughly once per minute at rest

C.2. Respiratory

Quantity	Typical adult	Note
Respiratory rate	~12–20 breaths/min	At rest
Tidal volume	~500 mL	One quiet breath
Anatomical dead space	~150 mL	Conducting zone; no gas exchange
Alveolar ventilation per breath	~350 mL	Tidal volume − dead space
Minute ventilation	~6 L/min	V_T × rate
Inspiratory reserve volume	~3000 mL	
Expiratory reserve volume	~1100 mL	
Residual volume	~1200 mL	Cannot be exhaled
Vital capacity	~4600 mL	Max exhalable after max inhalation
Total lung capacity	~5800 mL	VC + RV
Number of alveoli	~300–500 million	
Alveolar diameter	~0.2–0.3 mm	
Total alveolar surface area	~70 m²	Roughly a singles tennis court
Alveolar–capillary barrier	<1 µm	

Volume figures are typical young adult male values [1]; they vary substantially with sex, height, age, and fitness.

C.3. Gas partial pressures

Location	P_{O_2} (mmHg)	P_{CO_2} (mmHg)
Dry atmospheric air, sea level	~160	~0.3
Alveolar air	~104	~40
Arterial blood	~95–100	~40
Mixed venous blood	~40	~45

C.4. General chemistry

C.4.1. Constants

Quantity	Value	Note
Avogadro's number	6.022×10^{23} per mol	
Gas constant, R	8.314 J/(mol·K)	0.08206 L·atm/(mol·K)

Quantity	Value	Note
Faraday constant, F	96,485 C/mol	Charge of one mole of electrons
Molar volume of an ideal gas	~22.4 L (0 °C), ~24.0 L (20 °C), ~25.4 L (37 °C)	At 1 atm
Atmospheric pressure	101.3 kPa = 760 mmHg = 1.013 bar	
Water vapour pressure at 37 °C	~47 mmHg (6.3 kPa)	Saturated airway gas
Oxygen solubility in plasma	~0.003 mL per 100 mL per mmHg	Why haemoglobin is needed
Absolute zero	0 K = −273.15 °C	
pH electrode slope	~59 mV per pH unit at 25 °C	Nernst slope

C.4.2. Useful molar masses

Substance	g/mol
Water, H_2O	18.02
Sodium chloride, NaCl	58.44
Potassium chloride, KCl	74.55
Glucose, $C_6H_{12}O_6$	180.16
Urea, CH_4N_2O	60.06
Creatinine	113.12
Calcium (Ca)	40.08
Carbon dioxide, CO_2	44.01
Oxygen, O_2	32.00
Ethanol	46.07
Isopropanol	60.10

C.4.3. Concentration conversions and solutions

Relationship	Value
mg/dL → mmol/L	÷ (molar mass ÷ 10)
0.9% w/v NaCl	9 g/L ≈ 154 mmol/L; ~308 mOsm/L theoretical
5% w/v glucose (dextrose)	50 g/L ≈ 278 mmol/L
1 ppm in dilute water solution	≈ 1 mg/L
1,000 ppm available chlorine	≈ 1 g/L

Relationship	Value
Q_{10} = 2 accelerated ageing at 55 °C vs 25 °C	Ageing factor 8; one year ≈ 46 days

C.4.4. Standard electrode potentials (orientation)

Half-reaction	E° (V)
$Li^+ + e^- \rightarrow Li$	−3.04
$Zn^{2+} + 2e^- \rightarrow Zn$	−0.76
$Fe^{2+} + 2e^- \rightarrow Fe$	−0.44
$2H^+ + 2e^- \rightarrow H_2$	0.00 (reference)
$Cu^{2+} + 2e^- \rightarrow Cu$	+0.34
$AgCl + e^- \rightarrow Ag + Cl^-$	+0.22
$O_2 + 4H^+ + 4e^- \rightarrow 2H_2O$	+1.23

C.5. Chemistry and fluid

Quantity	Typical	Note
Arterial blood pH	~7.35–7.45	Tightly regulated; small deviations are physiologically large
Core body temperature	~36.5–37.5 °C	Varies by site and time of day; measurement site must be stated
Total body water	~50–60% of body mass	Intracellular ~2/3, extracellular ~1/3
Plasma osmolality	~275–295 mOsm/kg	
Blood glucose	~4–6 mmol/L fasting	≈ 70–110 mg/dL. **Unit confusion is a real hazard**

C.6. Scale reference

Structure	Typical scale	Maker reference
Protein molecule	~5 nm	
Cell membrane	~7–8 nm	
Red blood cell	~7–8 µm	~1/10 of a fine hair
Capillary lumen	~5–10 µm	Red cells pass nearly single-file
Skeletal muscle fibre	~10–100 µm	A human hair is ~50–100 µm; a printed layer ~100–300 µm

Structure	Typical scale	Maker reference
Epidermis	~0.05–0.1 mm	Thinner than printer paper; much thicker on palm and sole
Full skin (epidermis + dermis)	~1–4 mm by site	Thin plywood
Alveolus	~0.2–0.3 mm	
Nerve fascicle	~0.1–1 mm	
Small end mill, printer nozzle, fine laser kerf	~0.1–1 mm	**A 0.4 mm tool is larger than an alveolus**
Tympanic membrane thickness	~0.1 mm	About a sheet of printer paper
Central cornea	~0.5 mm	About 0.5 mm sheet aluminium

C.7. Cells and molecular biology

Quantity	Typical	Note
Human cells in an adult body	~30–40 trillion	Most by number are red blood cells
Bacteria in and on the body	Same order as human cells	Mostly in the large intestine
Recognised human cell types	Several hundred	Cell atlases resolve many more subtypes
Typical human cell diameter	~10–30 µm	
Red blood cell	~7–8 µm diameter, ~2 µm thick	
Bacterium (**E. coli**)	~1–2 µm long, ~0.5–1 µm wide	
Virus	~20–300 nm	
Ribosome	~20–30 nm	
Nucleus	~5–10 µm	
Human genome (haploid)	~3.1×10^9 base pairs	~775 MB at 2 bits per base
Protein-coding genes	~20,000	~1–2% of the genome is coding sequence
DNA per diploid cell, stretched out	~2 m	
Mitochondrial genome	~16.6 kb, 37 genes	Maternally inherited
E. coli genome	~4.6×10^6 base pairs	~4,300 genes
Replication error rate after repair	~10^{-9}–10^{-10} per base per division	
Difference between two people	~1 base in 1,000	

Quantity	Typical	Note
Genetic code	64 codons: 61 amino acid, 3 stop; 20 amino acids	
Cultured human cell cycle	~24 h	G1 ~11 h, S ~8 h, G2 ~4 h, M ~1 h (very approximate)
Hayflick limit	~40–60 divisions	
E. coli doubling, ideal conditions	~20 min	
M. tuberculosis doubling	~15–20 h	Why TB culture takes weeks
Distance from cells to nearest capillary	~100–200 µm	Diffusion limit
O_2 diffusion coefficient in water	~2×10^{-9} m²/s	$t \approx x^2 / 2D$
Light-microscope resolution	~200–250 nm	$d \approx \lambda / 2NA$
Culture incubator conditions	37 °C, ~5% CO_2, high humidity	
Storage temperatures	−20 °C, −80 °C, liquid nitrogen −196 °C	

C.7.1. Laboratory rules of thumb

- Exponential growth: $N = N_0 \times 2^{t/T}$. 30 doublings ≈ 10^9.
- PCR: copies ≈ $(1 + E)^n$ for efficiency E over n cycles. 30 cycles at 90%: ~2×10^8.
- Centrifuge: RCF ≈ $1.118 \times 10^{-5} \times r\,(cm) \times rpm^2$. Specify protocols in **g**, not rpm.
- Surface-to-volume of a sphere: 3 / r.

C.8. Biosignals

Signal	Amplitude	Bandwidth	Source
ECG	~0.5–4 mV	~0.05–150 Hz	Cardiac depolarisation
EMG (surface)	~0.05–5 mV	~10–500 Hz	Skeletal muscle fibres
EEG	~10–100 µV	~0.5–100 Hz	Summed cortical activity
EOG	~50–3500 µV	~0–30 Hz	Corneo-retinal potential
Nerve AP (surface)	~1–20 µV	~100–2000 Hz	Peripheral nerve conduction
PPG	Optical	~0.5–20 Hz	Pulsatile blood volume



Never connect anything you built to a person. Patient-connected equipment is governed by IEC 60601-1 (Appendix D). See Chapter 7.

C.9. Musculoskeletal

C.9.1. Approximate joint ranges

Joint and motion	Approximate adult range
Elbow flexion	~0–145°
Forearm pronation / supination	~80–90° each
Wrist flexion / extension	~70–80° / ~60–70°
Wrist radial / ulnar deviation	~20° / ~30°
Shoulder flexion / abduction	~0–180° each
Hip flexion	~0–120°
Knee flexion	~0–135°
Ankle dorsiflexion / plantarflexion	~20° / ~50°



Design to the **functional** range — what the person uses comfortably and repeatedly — not to the maximum, and have a therapist establish it for the individual. See Chapter 6.

C.9.2. Approximate material stiffness

Material	Elastic modulus	Note
Adipose tissue	~1–10 kPa	
Articular cartilage	~0.5–15 MPa	Strongly rate-dependent
Skin	~0.1–100 MPa	Depends entirely on strain level and method; non-linear
Silicone elastomer	~0.1–10 MPa	Grade-dependent
Trabecular bone	~0.01–3 GPa	Varies over orders of magnitude by site and density
Tendon (along fibres)	~1–2 GPa	Non-linear; compliant toe region at low strain
Polypropylene, moulded	~1.3–1.8 GPa	Living hinges; autoclavable grades exist
Nylon (PA6 / PA66)	~2–3 GPa	Absorbs water; properties shift with humidity
PLA, printed	~2–4 GPa	Anisotropic; weakest between layers
Acetal (POM), machined	~2.8–3.1 GPa	Dimensionally stable, low friction
PEEK	~3.6–4 GPa	Machined or moulded; tolerates repeated steam sterilisation
Cortical bone	~15–20 GPa	Anisotropic

Material	Elastic modulus	Note
Borosilicate glass	~63 GPa	Stiff and brittle; no yield before fracture
Aluminium 6061-T6	~69 GPa	
Ti-6Al-4V	~110 GPa	Closest common metal to bone stiffness
Stainless steel 316L	~190–200 GPa	~10× cortical bone

C.9.3. Two-point discrimination

Site	Approximate threshold
Fingertip	~2–4 mm
Palm	~10 mm
Forearm	~30–40 mm
Back	~40–50 mm or more

Tactile markers must be separated by considerably more than the threshold **at the site where they will be felt**. See Chapter 7.

C.10. Special senses

C.10.1. Eye

Quantity	Typical adult	Note
Axial length of eyeball	~24 mm	~1 mm extra ≈ 2.5–3 D of myopia
Central corneal thickness	~0.5 mm	
Total optical power	~60 D	Cornea ~43 D; lens ~20 D
Accommodation range	~10 D at 20 → ~1 D by 60	Presbyopia
Pupil diameter	~2–8 mm	
Intraocular pressure	~10–21 mmHg	
Rods / cones	~90–120 million / ~6 million	
Blind spot	~15° from fixation, toward the temple	
Dark adaptation	~20–30 min	Light adaptation takes seconds
Resting blink rate	~15–20 per minute	Falls sharply during visual concentration

Quantity	Typical adult	Note
Interpupillary distance	~54–74 mm	Adjustment range for eyewear and headsets
Standard (6/6) acuity	Letter subtends 5 arcmin	Detail of 1 arcmin

C.10.2. Legibility rule of thumb

Threshold letter height for standard acuity ≈ **viewing distance** × **0.00145**. Multiply by the acuity ratio for poorer vision (×2 for 6/12, ×3 for 6/18), then by a reserve of two to three for comfortable reading. See Chapter 8.

C.10.3. Ear and sound

Quantity	Typical	Note
Ear canal	~25 mm long, ~7 mm diameter	Quarter-wave resonance ~3 kHz
Tympanic membrane	~8–10 mm diameter, ~0.1 mm thick	
Middle-ear pressure gain	~22×, ~25–30 dB	Area ratio ~17 × lever ~1.3
Basilar membrane	~35 mm uncoiled	High frequencies at base, low at apex
Inner / outer hair cells	~3,500 / ~12,000	Do not regenerate
Audible range	~20 Hz–20 kHz	Upper limit falls with age; most sensitive ~2–5 kHz
Reference pressure, 0 dB SPL	20 µPa	
Conversation at 1 m	~60 dB(A)	
Occupational action level	85 dB(A), 8 h	Each +3 dB halves the allowable time
Threshold of pain	~120–140 dB	

Sound rules: two equal sources add 3 dB; level falls ~6 dB per doubling of distance in open space; +10 dB ≈ twice as loud.

C.10.4. Upper airway, voice, and mouth

Quantity	Typical adult	Note
Nasal share of airway resistance	~50%	Quiet nasal breathing
Peak inspiratory flow at rest	~25–30 L/min	Far above typical nasal cannula flows
Voice fundamental frequency	~85–155 Hz (men), ~165–255 Hz (women)	

Quantity	Typical adult	Note
Pharyngeal phase of swallow	< 1 s	Breathing stops during it
Oesophageal transit	~8–10 s	
Saliva production	~0.5–1.5 L/day	
Mouth opening (between incisors)	~40–55 mm	
Teeth	20 primary, 32 permanent	Enamel modulus ~80 GPa

C.11. Kidneys, gut, and hormones

Quantity	Typical adult	Note
Renal blood flow	~1–1.2 L/min	~20–25% of cardiac output
GFR	~120 mL/min (~180 L/day)	Normalised to 1.73 m² body surface area
Urine output	~1–2 L/day	Minimum to clear solutes ~0.4–0.5 L/day
Urine osmolality range	~50–1200 mOsm/kg	
Acid excreted by kidneys	~50–100 mmol/day	
Bladder capacity	~400–600 mL	First urge ~150–250 mL
Urethra length	~4 cm (female), ~20 cm (male)	
Stomach acid	pH ~1.5–3.5	
Gastric emptying	~2–4 h	
Small intestine length in life	~3–5 m	Longer post mortem
Large intestine length	~1.5 m	
Mouth-to-anus transit	~1–3 days	Highly variable
Absorptive surface of gut	~30 m²	Older texts quote 200–300 m²
Fluid entering gut daily	~9 L (2 L intake + ~7 L secretions)	All but ~0.1–0.2 L reabsorbed
Liver mass; blood supply	~1.5 kg; ~25% of cardiac output	~75% of it via the portal vein
Bile production	~0.5–1 L/day	
Fasting blood glucose	~4–6 mmol/L	Hypoglycaemia below ~4 mmol/L
Serum calcium (total)	~2.2–2.6 mmol/L	

Quantity	Typical adult	Note
Hormone half-lives	Adrenaline ~2–3 min; insulin ~4–6 min; cortisol ~1–1.5 h; T4 ~7 days	
Rapid-acting insulin (subcutaneous)	Onset ~10–20 min; peak ~1–2 h; duration ~3–5 h	
CGM lag behind blood glucose	~5–15 min	
Haemodialysis	Blood flow ~300–400 mL/min; ~4 h × 3 per week	

C.11.1. CKD stages by eGFR (mL/min/1.73 m²), after KDIGO

Stage	eGFR
G1	≥ 90
G2	60–89
G3a	45–59
G3b	30–44
G4	15–29
G5	< 15

C.12. Disease and risk arithmetic

Quantity	Typical	Note
Brain share of resting oxygen use	~20%	For ~2% of body mass; minutes of ischaemia tolerance
Heart muscle, irreversible ischaemic injury begins	~20–40 min	Complete occlusion; "time is muscle"
Neurons lost per minute, untreated large-vessel stroke	~1.9 million	Widely quoted estimate
Ischaemic share of strokes	~85%	Population-dependent
Type 2 share of diabetes	~90%	
HbA_{1c} averaging window	~2–3 months	Red cell lifespan ~120 days
Lifetime foot-ulcer risk with diabetes	~1 in 5 to 1 in 3	Published estimates vary
Doublings from one cell to ~1 cm tumour	~30	~10^9 cells; ~40 doublings to ~1 kg

Quantity	Typical	Note
Recessive inheritance, two carrier parents	1/4 affected, 1/2 carrier	Per child
Dominant inheritance, one affected parent	1/2	Per child
NNT	1 ÷ absolute risk reduction	State baseline risk alongside

C.13. Across the lifespan

C.13.1. Vital signs by age, orientation only

Age	Heart rate (beats/min)	Breathing rate (breaths/min)
Newborn	~110–160	~30–60
Infant (to 1 year)	~100–160	~30–40
Toddler (1–2 years)	~95–150	~25–35
Pre-school (2–5 years)	~95–140	~25–30
School age (5–12 years)	~80–120	~20–25
Adolescent and adult	~60–100	~12–20

Published paediatric ranges differ between sources; clinical tools state which reference they use. See Chapter 12.

C.13.2. Body size and composition

Quantity	Typical	Note
Newborn weight / length	~3.5 kg / ~50 cm	Weight triples in the first year
Head as fraction of body length	~1/4 at birth; ~1/8 in adults	
BSA, newborn / adult	~0.22 m^2 / ~1.8 m^2	~2.4× more surface per kg in the newborn
Body water, newborn / adult	~75–80% / ~50–60% of mass	
Blood volume, newborn	~80–90 mL/kg	~300 mL in a 3.5 kg baby
Working memory	~4 chunks	Cowan
People over 65 who fall each year	About a third	Rises with age

C.14. Making and manufacture

C.14.1. Process accuracy and surface, orientation only

Process	Typical accuracy	Typical Ra
FDM / FFF print	~±0.2–0.5 mm	~10–25 µm or worse
SLA / DLP print	~±0.05–0.15 mm	~1–5 µm
SLS / MJF print (nylon)	~±0.2–0.3 mm	~5–15 µm
Metal powder bed print, as built	~±0.05–0.2 mm	~5–15 µm
CNC milling / turning	~±0.02–0.1 mm	~0.8–3.2 µm
Polished or electropolished metal	—	~0.1–0.4 µm
Laser cutting (sheet)	~±0.1 mm; kerf ~0.1–0.3 mm	Edge-dependent
Waterjet	~±0.1–0.25 mm	Slightly rough, tapered edge
Injection moulding	~±0.05–0.2 mm, size-dependent	~0.05–0.4 µm with a polished tool

Hygienic-design guidance commonly targets around 0.8 µm Ra or smoother for product-contact surfaces. A bacterium is about 1 µm across. See Chapter 14.

C.14.2. Materials and methods notes

Quantity	Typical	Note
Injection-mould tooling	~£/$ thousands (aluminium) to hundreds of thousands (multi-cavity steel)	Lead time weeks
Moulded wall thickness	~1–3 mm, uniform	Draft ~0.5–2°
Borosilicate glass expansion	~3.3×10^{-6} /K	Soda-lime ~9×10^{-6} /K
Li-ion 18650 cell	~3.6–3.7 V nominal; ~9–13 Wh	~33–47 kJ per cell
Laundry thermal disinfection	65 °C for 10 min, or 71 °C for 3 min	Healthcare linen
Process capability target	$C_{pk} \geq 1.33$	1.0 means the process just fits
Tolerance stack, n equal parts ±t	Worst case ±n·t; statistical ±t·√n	

C.15. Diagnostic imaging

C.15.1. Effective dose, orientation only

Source	Typical effective dose
Dental bitewing	~0.005 mSv
DEXA	~0.001 mSv
Chest X-ray	~0.02–0.1 mSv
Mammogram	~0.4 mSv
CT head	~2 mSv
CT chest	~6 mSv
CT abdomen and pelvis	~8–10 mSv
PET-CT	~10–25 mSv
Natural background, per year	~2–3 mSv

C.15.2. Physical constants and rules

Quantity	Value	Note
Speed of sound assumed in soft tissue	1540 m/s	Depth = c × t / 2
Ultrasound attenuation	~0.5 dB/cm/MHz	Counted both ways
Reflection, soft tissue / air	> 99.9%	Why gel is used
Doppler shift	$\Delta f = 2 f_0 v \cos \theta / c$	
Hydrogen gyromagnetic ratio	42.58 MHz/T	64 MHz at 1.5 T; 128 MHz at 3 T
Clinical MRI field strengths	1.5 T, 3 T	Earth's field ~50 μT
Hounsfield scale	Water 0, air −1000	Fat ~−100; soft tissue +20–60; cortical bone +1000 or more
Tc-99m	140 keV gamma; half-life ~6 h	SPECT workhorse
F-18	Positron emitter; half-life ~110 min	FDG PET
Annihilation photons	2 × 511 keV, opposite directions	PET
Counting noise	SNR = √N	Halve noise → 4× counts

C.16. Treatment and life support



Orientation figures for understanding the devices in Chapter 17. They are not settings, doses, or targets for anybody.

Quantity	Typical figure	Note
Lung-protective tidal volume	~6 mL/kg predicted body weight	Severe lung injury; predicted, not actual, weight
Minimum PEEP in common use	~5 cmH_2O	
FiO_2 range	0.21–1.0	
Respiratory time constant	~0.5 s (healthy, ventilated)	$\tau = R \times C$; ~3τ to empty
High-flow nasal oxygen	up to ~60 L/min	Heated and humidified
Oxygen concentrator output	~90–95% O_2, a few L/min	
Oxygen cylinder pressure (full)	~140–230 bar	
Volumetric pump accuracy	~±5%	
External defibrillation energy	~120–200 J biphasic; up to 360 J monophasic	
Transthoracic impedance	~50–150 Ω	
Pacing pulse	a few volts, < 1 ms	
Haemodialysis blood flow	~250–450 mL/min	
Dialysate flow	~500–800 mL/min	
Dialyser membrane area	~1.5–2 m^2	
Typical haemodialysis schedule	~4 h, 3 times a week	
Kt/V (haemodialysis)	~1.2 or more	Dimensionless dose of dialysis
Electrosurgical frequency	~0.3–3 MHz	Above nerve-stimulating frequencies
Laparoscopic insufflation	~12–15 mmHg CO_2	
Curative radiotherapy course	~60–70 Gy in ~2 Gy fractions	Hypofractionated schedules common
CT organ dose, for comparison	tens of mGy	
Steady state on regular dosing	~4–5 half-lives	
Fraction remaining after n half-lives	2^{-n}	1 → 50%, 3 → 12.5%, 5 → ~3%

C.16.1. Rules

- Respiratory equation of motion: $P = V / C + R \times \text{flow} + PEEP$.
- Capacitor energy: $E = \frac{1}{2} C V^2$.
- Infusion rate (mL/h) = dose (µg/kg/min) × weight (kg) × 60 ÷ concentration (µg/mL).
- Dialyser clearance: $K = Q_b \times (C_{in} - C_{out}) / C_{in}$.

- Point-source dose rate falls as 1 / distance².

C.17. Unit conversions

From	To	Factor
mmHg	Pa	× 133.3
mmHg	kPa	× 0.1333
cmH_2O	Pa	× 98.1
cmH_2O	mmHg	× 0.735
kgf	N	× 9.81
psi	kPa	× 6.895
mmol/L glucose	mg/dL	× 18.0
mg/dL glucose	mmol/L	÷ 18.0
°C	°F	× 9/5, + 32
N/cm²	kPa	× 10
kPa	mmHg	× 7.50
Creatinine mg/dL	µmol/L	× 88.4
BUN mg/dL	urea mmol/L	× 0.357
HbA1c % (DCCT)	mmol/mol (IFCC)	(% − 2.15) × 10.929; 6.5% ≈ 48 mmol/mol

Useful anchors: **120 mmHg ≈ 16 kPa**. **20 kPa ≈ 150 mmHg**. **1 N/cm² = 10 kPa ≈ 75 mmHg**.



Carry units in the type, not in a comment. Unit mismatch — mmol/L versus mg/dL, kg versus lb, cm versus inches — is a documented cause of real clinical software incidents. If you write software handling physiological values, make the unit part of the value, reject ambiguous input rather than guessing, and display the unit beside every number, every time. See Chapter 19.

Appendix D: Standards and regulatory orientation



This appendix tells you what exists and roughly what each thing is for. It does not tell you what applies to you, and reading it does not make anything you build compliant with anything.

The requirements that actually apply depend on your jurisdiction, your host institution, the device classification, and the intended use. **Establish them with your host, early.** Most ISO and IEC standards are paid documents; many are accessible through a university or public library.

D.1. Why a maker should know the shape of this

Not to comply with it. To recognise it.

When a clinical engineer says "that would need a risk file", or an educator says "we'd have to run that past quality", you should know what they mean, why it is reasonable, and roughly how much work it implies. That recognition is the difference between a collaboration and a series of surprises.

It also tells you where the edges of a learning project are. Anything that triggers IEC 60601-1 or ISO 10993 in earnest has left the territory this book covers.

D.2. Core process standards

Reference	Title / subject	What it is for
ISO 13485	Medical devices — Quality management systems	The QMS framework for organisations that design and manufacture medical devices. Process, documentation, traceability, change control
ISO 14971	Application of risk management to medical devices	The risk process: identify hazards, estimate and evaluate risk, control it, evaluate residual risk, monitor in production and post-production. The hazard → hazardous situation → harm chain in Chapter 19 comes from here
IEC 62366-1	Application of usability engineering to medical devices	Use-related risk analysis, use specification, formative and summative evaluation. The human factors process
IEC 62304	Medical device software — software life cycle processes	Software safety classification (A/B/C), development planning, architecture, implementation and verification, release, maintenance, and management of SOUP (software of unknown provenance — your dependencies)

Reference	Title / subject	What it is for
IEC 82304-1	Health software — product safety requirements	Software that is a product in its own right, not embedded in a device
ISO 14155	Clinical investigation of medical devices for human subjects	Good clinical practice for device studies
ISO 20417	Information supplied by the manufacturer	Labelling and instructions for use
IEC 81001-5-1	Health software and health IT systems safety, effectiveness and security — security activities in the product life cycle	The secure development lifecycle for health software, designed to sit alongside IEC 62304: threat modelling, secure coding, security testing, vulnerability handling (Chapter 20)
AAMI TIR45	Guidance on the use of agile practices in the development of medical device software	How sprints, backlogs, and continuous integration map onto IEC 62304 activities and records
ISO/IEC 29147 and ISO/IEC 30111	Vulnerability disclosure; vulnerability handling processes	The basis of coordinated vulnerability disclosure programmes that regulators now expect
DCB0129 / DCB0160 (England)	Clinical risk management for the manufacture / deployment of health IT systems	Apply to health IT whether or not it is a medical device; require a named Clinical Safety Officer and a clinical safety case
MDS2 form	Manufacturer Disclosure Statement for Medical Device Security	Standard questionnaire hospitals use to compare device security features
GHTF/SG3/N99-10	Quality management systems — process validation guidance	The classic reference on IQ, OQ, and PQ and on when a process must be validated rather than its output inspected (Chapter 14)

D.3. Safety and biological evaluation

Reference	Subject	Note
IEC 60601-1	Medical electrical equipment — general requirements for basic safety and essential performance	Leakage currents, isolation, applied-part classification (B / BF / CF), mechanical and thermal safety. **This is why you do not connect self-built electronics to a person.** Many collateral and particular standards sit beneath it (60601-1-2 EMC, 60601-1-6 usability, 60601-1-8 alarms, and device-specific parts)

Reference	Subject	Note
ISO 10993 series	Biological evaluation of medical devices	Part 1 is the framework: evaluation is driven by the **nature and duration of contact**. Later parts cover cytotoxicity (-5), sensitisation and irritation (-10), systemic toxicity (-11), chemical characterisation (-18), and more. Note the framing: the **finished device** in its **intended use**, not a material in the abstract
ISO 17665	Sterilisation of health care products — moist heat	Development, validation, and routine control of steam sterilisation
ISO 11135 / ISO 11137	Sterilisation — ethylene oxide / radiation	The other common modalities
ISO 11607	Packaging for terminally sterilised medical devices	The package is part of the sterile barrier system
EN 556	Requirements for devices labelled "STERILE"	Where the 10^{-6} sterility assurance level comes from
AAMI ST79	Comprehensive guide to steam sterilization and sterility assurance in health care facilities	The practical reprocessing reference used in many facilities
IEC 60601-1-8	Alarm systems in medical electrical equipment	Alarm priorities, sound and light characteristics. Relevant background whenever a design uses sound as a cue (Chapter 8)
IEC 60601-2-33, -2-37, -2-44, -2 -54	Particular standards for MR equipment, ultrasound equipment, CT, and radiography and radioscopy	Where SAR and gradient limits (MRI), thermal and mechanical indices (ultrasound), and dose display requirements (CT, X-ray) are specified (Chapter 16)
ASTM F2503	Marking medical devices and other items for safety in the MR environment	The **MR Safe** / **MR Conditional** / **MR Unsafe** labels and their icons. Applies to anything, including teaching props, taken into an MRI magnet room (Chapter 16)
IEC 62471	Photobiological safety of lamps and lamp systems	Risk groups for UV, blue-light, and thermal hazards to the eye and skin from LEDs and lamps. Applies to the UV and bright-LED sources a maker builds with (Chapter 8, Chapter 13)
IEC 60825-1	Safety of laser products	Laser classes and exposure limits. Many low-cost laser modules are not what their label claims; treat any laser near eyes with suspicion
ISO 80369 series	Small-bore connectors for liquids and gases in healthcare	Connector geometries made deliberately incompatible between applications — enteral, respiratory, neuraxial, intravascular — so that a misconnection is physically impossible. Design-to-eliminate written into a standard (Chapter 9)

Reference	Subject	Note
EN 71 series / ASTM F963	Safety of toys — mechanical and physical properties, small parts, cords, entrapment	Not medical standards, but the reference for small parts, cord length, and finger entrapment in anything a child could handle (Chapter 12, Chapter 8)
Health Building Note 03-01 [51]	Design of adult acute mental-health units (UK)	Includes ligature-risk principles; an example of how a setting turns a clinical risk into a building and fittings requirement (Chapter 12)
ISO 80369-3	Small-bore connectors for enteral applications	The ENFit connector. Physically incompatible with intravascular Luer connectors (Chapter 10)
ISO 15197	In vitro diagnostic test systems — blood-glucose monitoring systems for self-testing	Accuracy requirements for home glucose meters (Chapter 10)
IEC 60601-2-16	Particular requirements for haemodialysis, haemodiafiltration, and haemofiltration equipment	Life-sustaining extracorporeal equipment
IEC 60601-2-24	Particular requirements for infusion pumps and controllers	Relevant to many insulin and enteral feed pumps
ISO 20696	Sterile urethral catheters for single use	Catheter performance and testing
US FDA iCGM, ACE pump, and iAGC classifications	Interoperable CGM, alternate controller enabled infusion pump, interoperable automated glycaemic controller	Regulatory pathways created so sensors, pumps, and algorithms from different makers can work together (Chapter 10)
ISO 80601-2-12	Particular requirements for critical care ventilators	Alarms, accuracy, and single-fault behaviour for life-supporting ventilators. Related parts cover home-care ventilators, sleep-apnoea devices, and oxygen concentrators (Chapter 17)
IEC 60601-2-4	Particular requirements for cardiac defibrillators	Energy accuracy, charge time, AED rhythm-recognition performance (Chapter 17)
IEC 60601-2-2	Particular requirements for high-frequency surgical equipment	Return electrode monitoring, output limits, and burn prevention (Chapter 17)
IEC 60601-2-1	Particular requirements for electron accelerators (radiotherapy linacs)	Dose monitoring and **independent** interlocks — the Therac-25 lesson in a standard (Chapter 17)
ISO 14708 series	Active implantable medical devices	Pacemakers, ICDs, neurostimulators, cochlear implants (Chapter 17)
IEC 60529	Degrees of protection provided by enclosures (IP code)	Ingress protection against solids and water; how "wipeable" and "splash-proof" are specified (Chapter 14)

Reference	Subject	Note
IEC 62133-2 / UN 38.3	Safety of portable sealed secondary lithium cells and batteries / transport testing of lithium batteries	The minimum to look for in any cell or pack you build into a device (Chapter 14)
IEC 60601-1-2	Medical electrical equipment — electromagnetic disturbances	EMC requirements; why hospital equipment must tolerate and not emit interference (Chapter 14)
ASTM A967 / ASTM A380	Chemical passivation treatments for stainless steel / cleaning, descaling, and passivation of stainless steel parts	Specify passivation after machining or fabrication (Chapter 14)

D.4. Regional regulation

Jurisdiction	Framework
EU	Medical Device Regulation (EU) 2017/745 (MDR) and In Vitro Diagnostic Regulation (EU) 2017/746 (IVDR). Risk-based classification (I, IIa, IIb, III), conformity assessment, notified bodies, CE marking, UDI, post-market surveillance and vigilance
UK	UK Medical Devices Regulations, administered by the MHRA, with UKCA marking and transitional arrangements. Check current status — this area has been in flux
US	FDA Center for Devices and Radiological Health. Classification I/II/III; most Class II devices via 510(k) premarket notification, Class III via PMA; Quality System Regulation 21 CFR 820 including design controls at 820.30
Elsewhere	Health Canada, TGA (Australia), PMDA (Japan), NMPA (China), and others, each with its own framework. The IMDRF works toward convergence

D.4.1. Concepts worth knowing by name

Term	Meaning
Intended use / intended purpose	What the manufacturer states the device is for. It drives classification, evaluation, and essentially everything else
Classification	Risk-based tiers determining the conformity assessment route. Higher class, more evidence and more oversight
Design controls	The structured design process with records: inputs, outputs, review, verification, validation, transfer, changes
Design History File (DHF)	The accumulated record demonstrating the design followed that process. Your project portfolio is a miniature one

Term	Meaning
Technical documentation / Technical File	The EU equivalent body of evidence
Essential principles / General Safety and Performance Requirements	The high-level safety and performance requirements a device must meet
Harmonised / recognised standards	Standards that, when followed, give presumption of conformity with specific requirements
UDI	Unique Device Identification — machine-readable identification supporting traceability and recall
Post-market surveillance / vigilance	Ongoing monitoring of real-world performance; reporting of incidents
Custom-made device	A specific legal category for a device made for an individual patient to a prescription. Relevant to one-off assistive devices — **and it has real legal requirements, not an exemption from them.** Do not assume it is a loophole
In-house / health institution exemption	Some jurisdictions allow institutions to manufacture devices for their own use under specified conditions. **Those conditions are substantial**, and the institution — not you — carries the obligations
SaMD / IMDRF risk categories	Software as a Medical Device, classified I–IV by situation seriousness × significance of the information (Chapter 20, [55])
EU MDR Rule 11	The MDR classification rule for software; guidance in MDCG 2019-11 [58]
Cyber device (US, section 524B)	Since March 2023, premarket submissions for internet-connectable devices must include vulnerability monitoring and disclosure plans, cybersecurity processes, and an SBOM [57]
Predetermined change control plan	Pre-authorised plan for specified future modifications of a device, notably AI-enabled software (FDA final guidance December 2024)
EU AI Act	Regulation (EU) 2024/1689. AI in medical devices that require notified-body assessment is high-risk; obligations phase in to 2027



The "custom-made" and "in-house manufacture" routes are the two that makers most often misread as permission.

They are not exemptions from safety obligations. They are alternative routes with their own documented requirements, and they are the **institution's** responsibility to discharge, not yours. If someone suggests a project proceed under either heading, that is a signal to involve the host's quality or regulatory function, not a signal that the paperwork has gone away.

D.5. Chemical safety

Reference	Subject
GHS [70]	UN system for classifying chemical hazards, labels, pictograms, and the sixteen-section Safety Data Sheet
CLP Regulation (EC) No 1272/2008; GB CLP	EU and UK implementation of GHS for classification, labelling, and packaging
COSHH Regulations 2002 (UK)	Control of Substances Hazardous to Health: assessment and control of exposure at work
OSHA Hazard Communication Standard, 29 CFR 1910.1200 (US)	Labels, SDSs, written programme, and training; aligned with GHS
ISO 8655	Piston-operated volumetric apparatus (pipettes): requirements and calibration
ASTM F1980	Accelerated ageing of sterile barrier systems and medical devices; the Q_{10} ageing-factor method (Chapter 2)

D.6. Laboratory biosafety and quality

Reference	Subject	Note
WHO Laboratory Biosafety Manual, 4th ed. [74]	Biosafety and biosecurity in laboratories	Risk-based approach to containment; the international reference (Chapter 4)
BMBL, 6th ed. [75]	Biosafety in Microbiological and Biomedical Laboratories (US)	Defines BSL-1 to BSL-4 practices and facilities (Chapter 4)
EN 12469 / NSF/ANSI 49	Performance criteria for microbiological safety cabinets (Europe / US)	How biological safety cabinets are tested and certified (Chapter 4)
ISO 15189	Medical laboratories — requirements for quality and competence	Accreditation standard for clinical laboratories; why diagnostic instruments must be validated (Chapter 4, Chapter 15)
ISO 10993-5	Biological evaluation of medical devices — tests for in vitro cytotoxicity	Cell culture testing of device materials; the most common biocompatibility test (Chapter 4, Chapter 3)

D.7. Health data protection

Reference	Subject
GDPR (EU) 2016/679; UK GDPR and Data Protection Act 2018	Personal data protection; health data are special category data. Article 9 conditions, DPIAs, data protection by design
Caldicott principles	Eight UK principles for confidential health and care information
HIPAA Privacy and Security Rules (US)	PHI protection for covered entities and business associates; de-identification by Safe Harbor or Expert Determination [59]
DICOM PS3.15 Annex E	Attribute confidentiality profiles for de-identifying DICOM objects
HL7 FHIR [60]	Modern health data exchange standard; resources over REST

D.8. Accessibility and human factors

Reference	Subject
WCAG 2.x / 3 [22]	Web Content Accessibility Guidelines. Perceivable, operable, understandable, robust
EN 301 549	European accessibility requirements for ICT products and services
ISO 9241 series	Ergonomics of human–system interaction
ISO 21542	Accessibility and usability of the built environment
AAMI HE75	Human factors engineering — design of medical devices
FDA human factors guidance [20]	Applying human factors and usability engineering to medical devices

D.9. Additive manufacturing

Reference	Subject
FDA, Technical Considerations for Additive Manufactured Medical Devices [26]	Device design, software workflow, material controls, post-processing, process validation, testing, and labelling considerations
FDA, 3D Printing of Medical Devices [25]	Overview of how additively manufactured devices are evaluated
ISO/ASTM 52900 series	Additive manufacturing terminology and process categories
ASTM F42 / ISO TC 261	The committees producing AM standards

D.10. Infection prevention

Reference	Subject
CDC, Guideline for Disinfection and Sterilization in Healthcare Facilities [24]	Definitions, Spaulding classification, process selection. Read the definitions; do not select a process from it
CDC Project Firstline / Inside Infection Control [23]	Accessible training material on infection control concepts
WHO hand hygiene guidance	The single most effective intervention
National and local decontamination guidance	**What actually applies to you.** Established with your host

D.11. How to read a standard without buying it

1. **Start with the scope and the terms-and-definitions sections.** They are often the most useful parts, and summaries of them are widely available.
2. **Look for free summaries** from regulators, notified bodies, and professional bodies. These are usually accurate about structure and intent.
3. **Use a library.** University and some public libraries hold BSI, ISO, or IEC subscriptions.
4. **Read the guidance, not the standard, first.** FDA guidance documents are free, readable, and explain the reasoning in a way standards do not.
5. **Ask your host.** If you are working with an institution, their quality or clinical engineering function already owns the relevant documents and knows which parts matter.



A realistic scope note. Nothing in this book brings a project within the scope of any of these standards, because nothing in this book involves a device placed on the market or used in patient care. The purpose of knowing them is to recognise the moment a project **would** — and to stop, and to ask.

Appendix E: Final assessment and readiness review

E.1. How to use this

Complete the assessment **without notes**, then compare your reasoning against the answers.

Each question is worth two points: **one for the correct conclusion, one for a sound explanation**. A suggested target is **20 out of 24**, followed by correcting every error in writing. Re-attempt any topic you missed a week later.



This is a self-study benchmark written for this book. It is not a recognised professional qualification and confers nothing. Its value is diagnostic: it tells you which chapters to revisit.

E.2. Written assessment

1. A right-arm model is rotated on the table. Do its proximal end and its anatomical right side change?
2. Explain the difference between a tendon, a ligament, and a muscle, using the elbow model as your example.
3. A broad pad has a lower calculated average pressure than a narrow pad under the same load. What does that establish, and what does it leave unanswered?
4. A 3 N load acts 0.10 m from a pivot. What ideal cable force balances it with a 0.03 m perpendicular moment arm? State your assumptions.
5. Why might a clearly labelled control still be inaccessible to someone?
6. Trace blood from the right ventricle to the aorta.
7. Explain why moving air is not the same as gas exchange.
8. A material supplier says a polymer is used in medical products. Is your moulded, machined, or printed object thereby established as suitable?
9. An educator asks for a motorised model. Give two questions to ask before designing the motor mechanism.
10. A prototype meets every dimension on its drawing, but learners give the wrong explanation. What has been demonstrated, and what has not?
11. Your test succeeds, but only after you guided the user through every step. What must the record say?
12. You change a cord guide after a successful test. Can you retain the old passing result for the new version?

E.3. Answers and reasoning

1. **No.** Anatomical references stay attached to the represented body. Proximal and distal refer to position along the limb relative to its attachment to the trunk, and anatomical left and right belong to the subject. The table's orientation is a separate coordinate system entirely, and conflating the two is the classic newcomer error. (Chapter 1)
2. **Muscle** generates contractile force. A **tendon** transmits that force from muscle to bone. A **ligament** typically connects bone to bone and contributes to joint stability and constraint. In the card model, an external hand supplies the force; the cord represents only the transmission element; nothing in the model represents a ligament at all — which is itself a stated limitation. (Chapter 5, Chapter 6)
3. It establishes a **lower calculated average** under the assumed load and contact area — one mechanical average, nothing more. It leaves unanswered: the actual pressure **distribution** and whether an edge or seam carries a disproportionate share; the **depth** at which peak stress occurs, which near a bony prominence is not at the surface; **duration**; the individual's tissue condition, perfusion, and sensation; **movement**, and therefore friction and shear; **microclimate**; and whether a competent clinician has reviewed it. (Chapter 5)
4. Load torque = 3 N × 0.10 m = **0.30 N·m**. Balancing cable force = 0.30 ÷ 0.03 = **10 N**. Assumptions: static equilibrium, forces exactly perpendicular, rigid massless beam, frictionless pivot, geometry fixed at this instant, no segment weight, no joint reaction force considered, single force path. (Chapter 6)
5. Because a label addresses only **perception of the label**, and the task has at least four independent demands. The person might not **see** it (acuity, contrast, lighting, angle); might not **understand** it (unfamiliar symbol, language, convention); might not **reach** it (range of motion, posture, obstruction); might not produce the required **force** or **precision**; might not **sustain** a required hold; or might not perceive the **feedback** confirming it worked. (Chapter 7)
6. Right ventricle → pulmonary valve → pulmonary arteries → lungs (gas exchange in the alveolar capillaries) → pulmonary veins → left atrium → mitral valve → left ventricle → aortic valve → aorta. (Chapter 9)
7. **Ventilation** moves air into and out of the lungs by changing thoracic volume and therefore pressure. **Gas exchange** is diffusion of O_2 and CO_2 across the alveolar–capillary barrier down partial-pressure gradients, and it additionally requires adequate **perfusion** matched to the ventilation. A model can inflate perfectly and demonstrate nothing about exchange. (Chapter 9)
8. **No.** That is a claim about a material as supplied, under stated conditions. Between the pellet and your part sit thermal history, additives and colourants, machine contamination history, layer porosity and anisotropy, surface roughness, support marks and residues, and post-processing chemicals. Evaluation is of a **finished device in its intended use**. Intended use, finished geometry, process, residues, performance, and relevant review all remain to be established. (Chapter 13)
9. Good questions address the need rather than the mechanism: **"What should the learner be able to do afterwards that they cannot now?"** and **"What is difficult about the current demonstration?"** A strong answer notes that these may reveal that motorisation would **defeat** the objective — if the point is to make the force source visible, hiding it in a motor is exactly wrong. (Chapter 18, Chapter 21)

10. **Demonstrated:** dimensional conformance — verification. The object was built as drawn. **Not demonstrated:** suitability for the learning objective — validation. Meeting a specification says nothing about whether the specification was correct. (Chapter 19)
11. It must describe **every prompt and all assistance given**, verbatim. The observation then supports **assisted** completion, which is a materially weaker claim than independent use. Recording it as a plain success attributes your explanation to the object's effectiveness. (Chapter 19, Chapter 21)
12. **No.** The result applies to the configuration it was obtained on. Assess what the change could affect — the cord's line of action and therefore the motion and required force, cleanability, whether the force path is still visible, whether the label is still readable — and repeat those checks on the new version as a new record. (Chapter 19)

E.4. Practical assessment

Show your project to someone who has **not** followed its development.

In five minutes, explain:

- the **need**, and who said so;
- the **relevant biology**, correctly;
- the **force path** or functional principle;
- the **learning task** or intended use;
- the **test evidence**, including at least one failure;
- the **limitations** — at least three.

Then ask them: **"What's the biggest unresolved question here?"**



Your ability to welcome and act on that question is part of the assessment.

Not your ability to answer it, and certainly not your ability to defend against it. A collaborator who hears a criticism, writes it down, and comes back having addressed it is the person a clinical team wants. A collaborator who explains why the criticism does not apply is not.

For a real host, ask an appropriate clinician or educator to assess anatomical accuracy and whether the proposed scope fits their setting. **Keep their judgement separate from your self-assessment** in the record — two columns, not one.

If you have not found a partner, your independent project remains useful portfolio evidence, with external review marked explicitly as **pending**.

E.5. Final readiness review

Discuss these with a potential host. A **yes** means **you can show evidence**, not that you have watched a video or read a chapter.

☐	**Statement**
☐	I can follow and clarify a basic anatomical explanation relevant to my project, and I ask when I do not understand a term.
☐	I can describe a location on the body unambiguously, from the subject's perspective.
☐	I can turn a user need into a written brief without assuming the solution.
☐	I can distinguish an observation from an interpretation from an assumption, in writing.
☐	I can identify when anatomy, infection control, usability, clinical engineering, information governance, or regulatory input is needed — and I ask before proceeding.
☐	I can write a requirement with an acceptance criterion agreed before the test.
☐	I can trace a hazard through to a possible harm and choose a control at the right level of the hierarchy.
☐	I can build a documented prototype and reproduce its configuration from my records.
☐	I report test failures honestly, including the ones that are my fault.
☐	I record prompts and assistance separately from independent performance.
☐	I can explain a model's limitations without being asked, and I do not overstate its accuracy.
☐	I understand that a material claim is not device evidence, and a bench result is not clinical evidence.
☐	I will not connect anything I built to a person.
☐	I can work within the host's privacy, consent, workshop, and review arrangements, and I ask rather than assume.
☐	I have a portfolio example I can show, and I can name the next supervision or training I need.
☐	I can say clearly what I can do unsupervised, what I can do supervised, and what I should not do at all.

E.6. Completion outcome

A foundation for contributing to **supervised** medical making, with a portfolio that makes both your capabilities and your remaining learning needs clear.

That second half is not a weakness in the portfolio. It is the most useful thing in it.

Appendix F: Working glossary

Use these as recall prompts, not as definitions to memorise. The drill: **explain the term, give an example, then say why it matters in your project.** If you cannot do the third part, you have not finished learning it.

F.1. Anatomy and orientation

Term	Meaning in this book
Anatomy	Body structures and their relationships
Physiology	How living structures function
Anatomical position	Standard reference pose: upright, facing forward, arms at sides, palms forward
Anterior / posterior	Toward the front / the back
Superior / inferior	Toward the head / the feet
Medial / lateral	Toward / away from the midline
Proximal / distal	Nearer / farther from a limb's attachment to the trunk
Superficial / deep	Nearer / farther from the body surface
Ipsilateral / contralateral	Same side / opposite side
Prone / supine	Lying face-down / face-up
Sagittal plane	Separates left from right
Frontal (coronal) plane	Separates anterior from posterior
Transverse (axial) plane	Separates superior from inferior

F.2. General chemistry

Term	Meaning in this book
Atomic number (Z)	Number of protons in the nucleus; defines the element
Mass number	Protons plus neutrons
Isotope	Atoms of one element with different numbers of neutrons; same chemistry, different mass, sometimes radioactive
Valence electrons	Outer-shell electrons that determine how an atom bonds
Ion / cation / anion	Charged atom or group / positive ion / negative ion
Electronegativity	How strongly an atom pulls shared electrons towards itself; differences predict bond polarity

Term	Meaning in this book
Periodic table: group / period	Column sharing outer-electron arrangement and chemistry / row adding an electron shell
Mole	6.022×10^{23} particles (Avogadro's number); the chemist's counting unit
Molar mass	Mass of one mole in grams; numerically equal to molecular mass in daltons
Molarity (mol/L, mmol/L)	Moles of solute per litre of solution
Osmolarity / osmolality	Dissolved particles per litre of solution / per kilogram of water; governs water movement across membranes
% w/v, % v/v	Grams per 100 mL / millilitres per 100 mL
ppm	Parts per million; in dilute water solutions ≈ mg/L
Dilution ($C_1V_1 = C_2V_2$)	Amount of solute is conserved when a solution is diluted
Serial dilution	Repeated fixed-ratio dilutions giving a geometric series of concentrations
Solubility / saturated / precipitate	How much dissolves / holding the maximum / solid coming out of solution
Stoichiometry	The mole arithmetic of a balanced chemical equation
Ideal gas law	PV = nRT; relates pressure, volume, amount, and temperature of a gas
Partial pressure (Dalton's law)	Each gas's share of the total pressure in a mixture
Henry's law	Dissolved gas is proportional to its partial pressure above the liquid
Vapour pressure	Pressure of a liquid's vapour; high in volatile, flammable solvents
Enthalpy change (ΔH)	Heat released (negative, exothermic) or absorbed (positive, endothermic) by a reaction
Activation energy	The energy barrier a reaction must overcome to proceed
Catalyst	Lowers the activation energy and speeds a reaction without being consumed or shifting equilibrium
Arrhenius equation / Q_{10}	Temperature dependence of reaction rate / factor by which rate rises per 10 °C (often ~2)
First-order kinetics	A constant fraction reacts per unit time: radioactive decay, many drug eliminations, microbial killing
D-value	Time at a given condition for a tenfold reduction in microbial population
Accelerated ageing	Elevated-temperature storage to simulate longer real-time ageing, within the limits of the material
Chemical equilibrium	Forward and reverse reactions at equal rates; composition steady
Le Chatelier's principle	A disturbed equilibrium shifts to partly counteract the disturbance
Oxidation / reduction	Loss / gain of electrons; always paired (redox)

Term	Meaning in this book
Oxidising agent	A substance that takes electrons — bleach, peroxide, peracetic acid
Anode / cathode	Electrode where oxidation / reduction occurs
Galvanic cell / electrolysis	Spontaneous redox producing current / current forcing a non-spontaneous reaction
Faraday's law	Mass deposited or dissolved is proportional to charge passed
Functional group	A small cluster of atoms that gives an organic molecule its characteristic chemistry
Enantiomers / chirality	Non-superimposable mirror-image molecules / the property of having them
GHS	Globally Harmonized System for classifying and labelling chemical hazards
Safety Data Sheet (SDS)	Standard sixteen-section hazard document supplied with a chemical
COSHH / HazCom	UK regulations and US OSHA standard for controlling and communicating chemical hazards at work
Sensitiser	A substance that primes the immune system so later small exposures cause reactions, often for life
Chromatography	Separation of mixtures between a stationary and a mobile phase: TLC, HPLC, GC
Mass spectrometry	Identification and quantification by mass-to-charge ratio; often coupled to chromatography
FTIR	Infrared spectroscopy identifying functional groups; quickly identifies unknown plastics and residues
LIMS / ELN	Laboratory information management system / electronic laboratory notebook

F.3. Cells, tissues, and regulation

Term	Meaning in this book
Cell	The smallest living unit
Tissue	Organised cells plus extracellular material, contributing to a function
Organ	Several tissues integrated into one function
Organ system	Organs cooperating
Diffusion	Net movement down a concentration gradient
Osmosis	Water movement across a selectively permeable membrane in response to effective solute differences
Tonicity	A solution's effect on cell volume: isotonic, hypotonic, hypertonic

Term	Meaning in this book
Active transport	Movement against a gradient, requiring cellular energy
ATP	The principal molecule through which cells transfer chemical energy
Homeostasis	Maintenance of internal conditions within workable ranges
Negative feedback	A response that opposes the change that triggered it. "Negative" means opposing, not harmful
Positive feedback	A response that amplifies the change; runs to a defined endpoint
Epithelial tissue	Covering and lining tissue; avascular; the body's primary barrier
Connective tissue	Supporting tissue characterised by abundant extracellular matrix — bone, cartilage, tendon, fat, blood
Collagen	The dominant tensile protein in tendon, ligament, skin, and bone's organic matrix
Elastin	Protein providing recoverable large strain
Extracellular matrix	The material secreted by cells that gives connective tissue its properties

F.4. Cell and molecular biology

Term	Meaning in this book
Cell theory	All living things are made of cells; the cell is the unit of life; cells arise only from existing cells
Natural selection	Heritable variants that survive and reproduce better become more common; the mechanism of evolution
Antibiotic resistance	Bacterial populations evolving to survive a drug; spread by mutation and by plasmid transfer
Prokaryote / eukaryote	Cell without a nucleus (bacteria, archaea) / cell with a nucleus and organelles (animals, plants, fungi, protists)
Organelle	Membrane-bound compartment with a specific job inside a eukaryotic cell
Nucleus / nucleolus	Compartment holding the chromosomes / region within it that assembles ribosomes
Ribosome	Molecular machine that translates mRNA into protein; 70S in bacteria, 80S in human cytoplasm
Endoplasmic reticulum (rough / smooth)	Membrane network making secreted and membrane proteins / lipids, drug detoxification
Golgi apparatus	Modifies, sorts, and packages proteins, including adding sugar chains (glycosylation)
Lysosome	Acidic compartment of digestive enzymes

Term	Meaning in this book
Mitochondrion	Organelle making most of the cell's ATP; has its own maternally inherited DNA
Cytoskeleton	Actin filaments, microtubules, and intermediate filaments: shape, transport, movement, mechanical sensing
Endocytosis / exocytosis	Taking in / releasing material by membrane vesicles; phagocytosis is endocytosis of particles
Genome	An organism's complete DNA; ~3.1 billion base pairs per human haploid set, ~20,000 protein-coding genes
Chromosome	One long DNA molecule with its packaging proteins; 46 in most human cells
Base pair	A–T or G–C pairing across the DNA double helix
Replication	Copying DNA before division; semi-conservative, proofread, ~10^{-9}–10^{-10} errors per base
Transcription	Copying a gene into RNA by RNA polymerase
Promoter / transcription factor	DNA region where transcription starts / protein that switches genes on or off
Splicing (exon / intron)	Removing non-coding introns from pre-mRNA and joining coding exons; alternative splicing makes protein variants
Translation	Ribosome building a protein from mRNA codons
Codon / genetic code	Three-base unit specifying an amino acid / the 64-codon table, redundant and nearly universal
tRNA	Adaptor RNA carrying the amino acid matching each codon
Epigenetics	Heritable patterns of gene use set by chemical marks on DNA and histones, not by sequence changes
Cell cycle (G1, S, G2, M, G0)	Growth, DNA synthesis, preparation, mitosis, and resting phases of a cell's life
Mitosis / meiosis	Division giving two identical cells / double division giving four gametes with half the chromosomes
Checkpoint (p53)	Control point stopping division if DNA is damaged; p53 is a key tumour suppressor
Telomere / Hayflick limit	Protective chromosome-end repeats that shorten with division / ~40–60 divisions before most human cells in culture stop
Senescence	Permanent arrest of cell division, with inflammatory signalling
Differentiation	A cell becoming a specialised type through stable changes in gene use
Stem cell (totipotent / pluripotent / multipotent)	Self-renewing cell able to form every type including placenta / every body type / a family of types
iPSC	Induced pluripotent stem cell, reprogrammed from an adult cell

Term	Meaning in this book
Gram-positive / Gram-negative	Bacteria with a thick cell wall that retains violet stain / thin wall plus outer membrane, stains pink
Plasmid	Small circular DNA in bacteria; carries and transfers genes, including resistance; a cloning vehicle
Bacterial growth curve	Lag, exponential, stationary, and death phases of a culture
Virus / envelope	Non-cellular particle of genome plus protein coat / outer lipid layer of some viruses, vulnerable to alcohol and detergent
Retrovirus / reverse transcriptase	RNA virus that writes its genome into host DNA / the enzyme that does it, also a lab tool
Innate / adaptive immunity	Rapid general response / slower, specific, remembered response by T and B cells
T cell / B cell / antibody	Lymphocytes that coordinate and kill / make antibodies / Y-shaped protein binding a specific target
Biosafety level (BSL-1 to 4)	Containment classification of laboratories by the risk of the organisms handled
Biological safety cabinet	HEPA-filtered enclosure protecting worker, sample, and room; not a fume hood
Relative centrifugal force (RCF)	Centrifugal acceleration in multiples of **g**; depends on radius and rpm²
Next-generation / long-read sequencing	Massively parallel short reads / single-molecule reads of tens of kilobases or more
CRISPR–Cas	RNA-guided system for cutting or editing DNA at a chosen sequence
Flow cytometry / FACS	Measuring cells single-file by laser scatter and fluorescence / sorting them
Organoid / organ-on-chip	Miniature self-organised 3D tissue / microfluidic device lined with living cells under flow

F.5. Skin and contact

Term	Meaning in this book
Epidermis	Outer, avascular layer of skin; the barrier
Dermis	Layer beneath; contains vessels, nerves, glands; mechanically dominant
Hypodermis / subcutaneous tissue	Adipose and loose connective tissue below the skin; thickness varies enormously
Stratum corneum	Outermost dead keratinised layer; the barrier proper
Pressure	Load normal to a surface; force divided by area gives the **average** only
Friction	Resistance to sliding at an interface; abrades the surface

Term	Meaning in this book
Shear	Load parallel to a surface, deforming tissue internally; distorts vessels
Microclimate	Local temperature and humidity at an interface
Viscoelastic	Mechanical response depends on rate and loading history
Creep	Continued deformation under constant load
Stress relaxation	Falling stress under constant deformation
Hysteresis	Energy lost per loading cycle
Anisotropic	Properties differ with direction
Maceration	Softening and weakening of skin from trapped moisture
Oedema	Excess interstitial fluid; swelling. A dimensional problem for anything worn

F.6. Musculoskeletal

Term	Meaning in this book
Tendon	Transmits muscle-generated force, usually to bone
Ligament	Usually connects bone to bone; contributes to joint stability
Articular cartilage	Smooth low-friction covering on bone ends within a synovial joint
Synovial joint	Freely movable joint with a capsule and fluid-filled cavity
Flexion / extension	Movements generally decreasing / increasing a joint angle
Abduction / adduction	Movement away from / toward the midline
Pronation / supination	Forearm rotation bringing the palm to face posteriorly / anteriorly
AROM / PROM	Active range of motion (own muscles) / passive range (moved externally)
Functional range	The range actually used comfortably and repeatedly for a real task
Moment arm	Perpendicular distance from a pivot to a force's line of action
Torque	The turning effect of a force; force × perpendicular moment arm
Isometric / concentric / eccentric	Contraction at constant length / while shortening / while lengthening
Motor unit	One motor neuron and all the muscle fibres it innervates
Agonist / antagonist	The muscle producing a movement / the one opposing and controlling it
Co-contraction	Agonist and antagonist acting together to stabilise a joint
Cortical / trabecular bone	Dense outer bone / open lattice inner bone
Stress shielding	Bone resorption caused by an implant carrying load the bone used to carry

Term	Meaning in this book
Anthropometry	Body dimension data. Percentiles are per-dimension; there is no "95th percentile person"

F.7. Nervous system and sensation

Term	Meaning in this book
CNS / PNS	Brain and spinal cord / the nerves connecting them to the body
Afferent / efferent	Carrying information toward / away from the CNS
Neuron	The signalling cell
Synapse	The junction at which one neuron signals the next cell
Neuromuscular junction	The synapse between a motor neuron and a skeletal muscle fibre
Action potential	The all-or-nothing electrical signal along an axon
Myelin	Insulating sheath enabling fast saltatory conduction
Proprioception	Sense of body position and movement
Nociception	Neural detection of potentially damaging stimuli — distinct from the experience of pain
Reflex	Rapid stereotyped involuntary response; can precede awareness
Two-point discrimination	Minimum separation at which two touches feel like two; varies hugely by site
Adaptation (sensory)	Reduced response to a sustained stimulus. Why you stop noticing a constant pressure
Redundant coding	Carrying the same information on more than one channel so no single impairment removes it
Perception / cognition / action / feedback	The four independent demands any interaction makes

F.8. Head, neck, and special senses

Term	Meaning in this book
Special senses	Vision, hearing, balance, smell, taste — dedicated organs in the head, served by cranial nerves
Cranial nerves	Twelve pairs of nerves leaving the brain directly; I–XII
Cornea	Clear front of the eye; avascular, densely innervated; ~2/3 of the eye’s focusing power
Lens / accommodation	The eye’s adjustable focusing element / its change of shape to focus near
Dioptre (D)	Optical power; reciprocal of focal length in metres

Term	Meaning in this book
Presbyopia	Age-related loss of accommodation; universal by the mid-forties to fifties
Retina / fovea	Light-sensitive layer at the back of the eye / its central pit of sharp, cone-rich vision
Rods / cones	Dim-light, colourless receptors / bright-light colour receptors (three types)
Intraocular pressure	Pressure of the eye's internal fluid; raised in many forms of glaucoma
Visual angle	Angle an object subtends at the eye; what actually determines legibility
Tympanic membrane	Eardrum; boundary between outer and middle ear
Ossicles	Malleus, incus, stapes; the middle ear's impedance-matching lever chain
Eustachian tube	Connects middle ear to nasopharynx; equalises pressure
Cochlea	Fluid-filled inner-ear spiral that sorts sound by frequency (tonotopy)
Hair cells	Mechanoreceptors of hearing and balance; do not regenerate in humans
Conductive / sensorineural hearing loss	Sound fails to reach the cochlea / cochlea or nerve damaged
Vestibular apparatus	Semicircular canals (rotation) plus otolith organs (linear acceleration, gravity)
Vestibulo-ocular reflex	Counter-rotation of the eyes that keeps vision stable as the head moves
dB SPL / dB(A) / dB HL	Sound level re 20 µPa / A-weighted toward human hearing / referenced to normal hearing per frequency
Turbinates (conchae)	Scroll-shaped projections in the nasal cavity that warm and humidify air
Paranasal sinuses	Air spaces in the skull draining into the nose
Pharynx	Muscular passage behind nose and mouth, shared by air and food
Larynx / glottis	Cartilaginous voice box guarding the airway / the gap between the vocal folds
Epiglottis	Flap that tilts over the airway entrance during swallowing
Dysphagia	Difficulty swallowing
Aspiration / silent aspiration	Material passing below the vocal folds / doing so without a cough
Tracheostomy / laryngectomy	Tube opening in the trachea / removal of the larynx, leaving a neck stoma and no upper-airway connection
IDDSI	International framework of eight food and drink texture levels with simple tests
Device-related pressure injury	Injury from a device's sustained load; common on nose, ears, and neck

F.9. Circulation and respiration

Term	Meaning in this book
Artery / vein	Carries blood away from / toward the heart. Defined by direction, not oxygen content
Capillary	Single-endothelial-layer vessel where exchange occurs
Systemic / pulmonary circuit	To the body / to the lungs
Cardiac output	Volume pumped per minute; CO = heart rate × stroke volume
Stroke volume	Volume ejected per beat
Systole / diastole	Contraction / relaxation phase of the cardiac cycle
Ventilation	Air movement into and out of the lungs
Gas exchange	Diffusion of O_2 and CO_2 across the alveolar–capillary barrier
Perfusion	Blood flow through tissue
V/Q matching	Ventilation and perfusion present in the same places
Dead space	Airway volume that moves air but performs no exchange (~150 mL anatomical)
Compliance	Change in volume per unit change in pressure
Surfactant	Substance reducing alveolar surface tension and stabilising small alveoli
Alveolus	The small air space where gas exchange occurs
SpO□	**Estimated** arterial haemoglobin oxygen saturation from a pulse oximeter
MAP	Mean arterial pressure; ≈ diastolic + □ pulse pressure

F.10. Kidneys, gut, and hormones

Term	Meaning in this book
Nephron	The kidney's working unit: a glomerular filter followed by a reabsorbing, secreting tubule; ~1 million per kidney
Glomerulus	Capillary knot where blood pressure filters plasma into the nephron
GFR / eGFR	Glomerular filtration rate (~120 mL/min) / its estimate from blood creatinine
Clearance	Volume of plasma cleared of a substance per minute; C = U × V / P
Creatinine	Muscle waste product used to estimate GFR; depends on muscle mass
ADH (vasopressin)	Pituitary hormone that makes the kidney retain water
Renin–angiotensin–aldosterone system	Hormone cascade that raises blood pressure and retains sodium
Erythropoietin	Kidney hormone that drives red cell production

Term	Meaning in this book
Detrusor	The bladder's muscular wall
Stress / urge / overflow / functional incontinence	Leakage from weak outlet / overactive bladder / incomplete emptying / barriers to reaching a toilet
Indwelling (Foley) catheter	Bladder catheter retained by a balloon, draining to a bag
CAUTI	Catheter-associated urinary tract infection
Peristalsis	Coordinated muscular waves that move contents along a tube
Chyme	Semi-liquid stomach contents passed to the small intestine
Villi / microvilli	Folds of the small-intestine lining that multiply absorptive surface
Portal vein	Vein carrying blood from the gut to the liver before the rest of the body
First-pass metabolism	Removal of an oral drug by gut and liver before it reaches the circulation
Bile	Liver secretion, stored in the gallbladder, that emulsifies fat
Microbiome	The microbial community of the gut
Enteral / parenteral nutrition	Feeding into the gut by tube / feeding intravenously
Gastrostomy (PEG)	Feeding tube through the abdominal wall into the stomach
Stoma	Bowel or urinary tract brought out through the abdominal wall: colostomy, ileostomy, urostomy
ISO 80369 / ENFit	Small-bore connector standards making different applications physically incompatible / the enteral connector
Hormone	Chemical messenger carried in blood to any cell with its receptor
Hypothalamic–pituitary axis	Three-level hormone cascade with negative feedback, e.g. TRH → TSH → T4
TSH	Pituitary hormone driving the thyroid; the usual first thyroid test
PTH	Parathyroid hormone; raises blood calcium
Insulin / glucagon	Pancreatic hormones that lower / raise blood glucose
CGM	Continuous glucose monitor; measures interstitial glucose, lagging blood by ~5–15 min
Automated insulin delivery (closed loop)	Pump dosing adjusted automatically from CGM readings
Haemodialysis / peritoneal dialysis	Blood cleaned across an artificial membrane / across the peritoneum

F.11. Disease mechanisms and genetics

Term	Meaning in this book
Aetiology / pathogenesis	The cause of a disease / the mechanism by which the cause produces damage
Sign / symptom	Found by an examiner or a measurement / reported by the person
Acute / chronic	Sudden and short / long-lasting. Not a statement of severity
Prognosis	Expected course and outcome
Incidence / prevalence	New cases over a period / existing cases at a point in time; prevalence ≈ incidence × duration
Comorbidity	Two or more conditions present together
Hypoxia / ischaemia	Too little oxygen / too little blood flow (oxygen, glucose, and waste clearance all lost)
Infarction	Tissue death from loss of blood supply
Necrosis / apoptosis	Uncontrolled cell death that provokes inflammation / programmed, orderly cell death that does not
Hypertrophy / atrophy	Cells growing larger / shrinking, as adaptations to load or disuse
Acute / chronic inflammation	Rapid vascular and neutrophil response / persistent response that tends to scar and destroy tissue
Sepsis	Life-threatening organ dysfunction from a dysregulated host response to infection
Shock	Circulatory failure: tissues not receiving enough oxygen. Hypovolaemic, cardiogenic, distributive, or obstructive
Thrombus / embolus	Clot formed inside a vessel / material carried in the blood until it lodges and blocks a vessel
Virchow's triad	Stasis, endothelial injury, hypercoagulability — the conditions that favour thrombosis
DVT / PE	Deep vein thrombosis / pulmonary embolism
Atherosclerosis	Fatty, inflammatory plaque in artery walls; narrows arteries and may rupture
Myocardial infarction	Heart attack: death of heart muscle from loss of blood supply
Ischaemic / haemorrhagic stroke	Brain injury from a blocked vessel / from a burst vessel
Heart failure	Heart cannot pump enough blood to meet needs at normal filling pressures
Asthma / COPD	Variable, reversible airway obstruction / persistent, largely irreversible obstruction
Type 1 / type 2 diabetes	Autoimmune loss of insulin production / insulin resistance with failing secretion
HbA_{1c}	Glycated haemoglobin; reflects average glucose over ~2–3 months

Term	Meaning in this book
Neuropathy	Nerve damage; in diabetes, loss of protective sensation in the feet
Benign / malignant	Local growth only / invasive growth able to spread
Metastasis	Spread of cancer to distant sites
TNM staging	Tumour extent, lymph node involvement, distant metastasis
Autoimmunity	Immune attack on the body's own tissues
Osteoarthritis / osteoporosis	Degeneration of a joint / loss of bone mass and strength
Mutation	A change in DNA sequence; inherited or acquired
Dominant / recessive / X-linked	One altered copy sufficient / two needed / gene on the X chromosome
Pharmacogenomics	Use of genetic variation to predict drug response
Relative / absolute risk reduction	Proportional change in risk / difference in risk in percentage points
Number needed to treat (NNT)	1 ÷ absolute risk reduction; people treated for one to benefit

F.12. Lifespan and the mind

Term	Meaning in this book
Neonate	A baby in the first four weeks of life
Preterm	Born before 37 weeks of pregnancy
Brown fat	Heat-producing fat tissue that keeps newborns warm; they cannot shiver effectively
Body surface area (BSA)	Estimated skin area; used to normalise some doses and physiological indices
Mosteller formula	BSA (m^2) = √(height cm × weight kg / 3600)
Kleiber's law	Metabolic rate scales roughly with mass to the power three-quarters
Centile (percentile) chart	Growth chart showing the distribution of a measurement by age; the trajectory matters more than one point
Assent	A child's own agreement, alongside a parent's or guardian's consent
Tenfold error	Dosing error from a misplaced decimal point or unit confusion; a recognised paediatric hazard
Physiological reserve	The margin between a system's everyday load and its maximum; falls with age
Sarcopenia	Age-related loss of muscle mass, strength, and power
Osteoporosis	Reduced bone mass and quality, making fractures from minor falls likely

Term	Meaning in this book
Frailty	Reduced reserve across several systems; vulnerability to disproportionate decline after a minor stress
Multimorbidity	Living with two or more long-term conditions
Polypharmacy	Taking many medicines at once, commonly defined as five or more
Attention	Selection of a small part of available information for processing
Vigilance decrement	Decline in detection performance during sustained monitoring
Working memory	Short-term mental scratchpad; about four meaningful chunks; fragile to interruption
Executive function	Planning, sequencing, switching, and inhibiting actions
Cognitive load	The demand a task places on attention and working memory
Hick's law	Decision time grows with the logarithm of the number of choices
Delirium	Acute, fluctuating disturbance of attention and awareness; a medical emergency
Dementia	Progressive decline in cognition affecting daily life; not normal ageing
Learning disability	Reduced ability to understand complex information and learn new skills, starting in childhood (UK usage; "intellectual disability" elsewhere)
Autism	Lifelong difference in communication, interaction, and sensory experience
Neurodiversity	Variation in how brains work, seen as part of human diversity
Easy read	Accessible information format: short sentences, common words, supporting pictures
Ligature point	A fixture something could be looped around that could bear a person's weight
Anti-ligature design	Fittings and furniture designed to remove ligature points, e.g. by collapsing under load

F.13. Infection control and materials

Term	Meaning in this book
Cleaning	Physical removal of soil and contaminants. Must come first
Disinfection	A process inactivating many disease-causing microorganisms. Not equivalent to sterilisation
Sterilisation	A validated process intended to eliminate all forms of microbial life, including spores
Spaulding classification	Critical / semi-critical / non-critical, determined by **use**, not material
Fomite	An inanimate object acting as a vehicle for indirect contact transmission

Term	Meaning in this book
Biofilm	Surface-attached microbial community protected by a self-produced matrix; resistant to removal
Spore	Dormant, highly resistant bacterial form
Chain of infection	Agent → reservoir → portal of exit → transmission → portal of entry → susceptible host
Biocompatibility	Appropriate material–tissue interaction for a given nature and duration of contact
Traceability	The ability to connect a result to the exact configuration that produced it

F.14. Making and manufacture

Term	Meaning in this book
Fixed / variable cost	One-off cost of tooling, programming, and setup / cost per unit made
Break-even quantity	Quantity at which two methods' total costs are equal: $(F_A - F_B) / (v_B - v_A)$
Tooling	Moulds, dies, formers, and fixtures that shape parts; a fixed cost repaid over quantity
Material extrusion (FDM/FFF)	Additive process depositing molten thermoplastic layer by layer
Vat photopolymerisation (SLA/DLP)	Additive process curing liquid resin with light
Powder bed fusion (SLS, MJF, LPBF)	Additive process fusing powder with a laser or heat; polymer or metal
CNC	Computer numerical control: machining by programmed toolpaths
CAM	Computer-aided manufacturing: generating toolpaths and machine programs from CAD
Injection moulding	Molten polymer forced into a closed tool; the dominant process for disposables
Draft	Slight taper on moulded faces so parts release from the tool
Overmoulding	Moulding one material over another, e.g. a soft grip on a rigid core
Thermoforming / vacuum forming	Heating a plastic sheet and drawing it over a former
Platinum- / tin-cure silicone	Addition-cure (low shrink, cure-inhibited by sulphur and some materials) / condensation-cure silicone rubbers
Shore hardness	Indentation hardness scale for elastomers (00, A, D scales)
Ra	Arithmetic mean surface roughness, in micrometres
Burr	Sharp raised edge left by machining or cutting

Term	Meaning in this book
304 / 316L stainless	General-purpose stainless / molybdenum-bearing, low-carbon grade with better chloride resistance
Anodising	Electrochemical oxide layer on aluminium; attacked by alkaline detergents
Borosilicate glass	Low-expansion, heat- and thermal-shock-resistant glass used in laboratories
Spacer fabric	Three-dimensional knitted textile that resists compression and lets air and moisture through
Welded seam	Seam joined by heat or ultrasound rather than stitching; can be fluid-proof
Thermal disinfection (laundry)	Washing at a specified time and temperature, e.g. 65 °C for 10 min or 71 °C for 3 min
IP rating	Ingress protection code (IEC 60529) for solids and water, e.g. IPX4, IP67
EMC	Electromagnetic compatibility: neither disturbing nor disturbed by other equipment
Thermal runaway	Self-sustaining overheating of a battery cell that can spread and ignite
Firmware	Software embedded in a device's microcontroller; subject to software lifecycle controls
DFMA	Design for manufacture and assembly: fewer, standard, foolproof, inspectable parts
Tolerance stack-up	Accumulation of part tolerances along a chain; worst case or root sum of squares
Process capability (C_p, C_{pk})	Ratio of tolerance to process spread; C_{pk} also accounts for centring
Process validation (IQ/OQ/PQ)	Documented evidence that a process consistently produces acceptable output: installation, operational, performance qualification
Certificate of conformity / mill certificate	Supplier's document stating a material batch's identity and properties
Change notification	Supplier commitment to warn before changing a material's formulation
Bill of materials (BOM)	Complete list of parts, materials, grades, and suppliers
STEP / DXF / Gerber	Neutral file formats for solid models / 2D profiles / circuit board manufacture

F.15. Design, evidence, and process

Term	Meaning in this book
Intended use	What the thing is for, including explicitly what it is not for
Use environment	The real conditions of use, including interruptions, gloves, and lighting
Requirement	A stated condition or capability the design must meet, checkable
Acceptance criterion	The pre-agreed basis for deciding whether a check passes
Verification	Checking conformance with specified requirements — "did we build it right?"
Validation	Evaluating suitability for intended needs and use — "did we build the right thing?"
Hazard	A potential source of harm
Hazardous situation	Circumstances in which someone is exposed to a hazard
Harm	Actual injury or damage, physical or otherwise
Risk control hierarchy	Inherently safe design → protective measures → information for safety
Reasonably foreseeable misuse	Use that is not intended but is predictable, and must be tolerated
Use error	An error arising from the interaction of a person with a design — a design property, not a character flaw
Formative / summative evaluation	Usability evaluation during design / of the final design
Design controls	The structured design process with records
Design History File	The accumulated record demonstrating the process was followed
SOUP	Software of unknown provenance — your dependencies
Version	A configuration you can reproduce
Need statement	Who, what difficulty, what outcome — with no solution embedded
Pending	A criterion that requires another person's review you have not yet obtained. An honest status

F.16. Education and simulation

Term	Meaning in this book
Learning objective	An observable outcome the learner should achieve, with a verb you can see
Fidelity	How a model represents relevant features of the real task or body. Multiple independent dimensions
Task trainer	A model for practising a specific skill or relationship
Predict–observe–explain	A teaching sequence that surfaces the learner's existing mental model first

Term	Meaning in this book
Prompt	Assistance given during assessment. Must be recorded separately
Bloom's taxonomy	Hierarchy of cognitive demand: remember, understand, apply, analyse, evaluate, create
Miller's pyramid	Knows → knows how → shows how → does
Misleading cue	A feature of a simplified model that quietly teaches something false. The main hazard of a teaching aid
Reset time	How long to restore the starting state between learners. The constraint educators mention least and care about most

F.17. Chemistry and molecular biology

Term	Meaning in this book
Covalent / ionic / hydrogen bond	Shared electrons / electrostatic attraction / attraction involving a polarised hydrogen
Hydrophobic effect	Water excluding nonpolar material to preserve its own hydrogen bonding. Not an attraction between the solutes
pH	Negative log of hydrogen ion activity. A log scale: one unit is a tenfold change
Buffer	A weak acid and its conjugate base together, resisting pH change
Macromolecule	Protein, polysaccharide, lipid, or nucleic acid
Denaturation	Loss of a protein's three-dimensional shape without breaking its sequence
Protein adsorption	The layer of protein that coats any material within seconds of contacting blood or tissue. The body responds to this layer, not to your material
Enzyme	A protein catalyst; specific, and sensitive to temperature and pH
Saturation kinetics	Reaction rate plateauing when all enzyme active sites are occupied. Sets a biosensor's upper linear range
Phospholipid bilayer	The self-assembling basis of every cell membrane
ATP	The molecule through which cells transfer chemical energy
Agonist / antagonist	A molecule that activates a receptor / one that blocks it
Pharmacokinetics / pharmacodynamics	What the body does to a drug / what the drug does to the body
Narrow therapeutic index	A small gap between an effective and a harmful dose. Why dose-handling software is safety-critical
Glass transition (T_g)	Temperature at which an amorphous polymer turns from glassy to rubbery

Term	Meaning in this book
Extractable	What can be pulled out of a material under exaggerated conditions — the worst case
Leachable	What actually migrates out in normal use — the real exposure
Hydrolytic degradation	Water cleaving a polymer backbone. How PLA and PLGA break down
Environmental stress cracking	Stress plus a chemical agent cracking a polymer that neither alone would
Passivation	A thin self-repairing oxide film giving stainless steel and titanium their corrosion resistance
Galvanic corrosion	Preferential corrosion of the less noble of two dissimilar metals in electrical contact in an electrolyte
Surface energy	Why clean metals bond well and polypropylene, silicone, and PTFE do not

F.18. Measurement, imaging, and diagnostics

Term	Meaning in this book
Attenuation	Reduction of X-ray intensity passing through tissue. What radiography measures
Acoustic impedance	Property whose mismatch at an interface produces an ultrasound echo
Partial volume effect	One voxel spanning two tissues reporting their average. Limits the smallest recoverable feature
Segmentation	Deciding which voxels belong to a structure. The decisive and most error-prone step from scan to model
DICOM	The medical imaging format and protocol. Carries patient identifiers, including in private tags and burned-in pixels
PACS	The hospital system that stores and distributes medical images
ALARA	As Low As Reasonably Achievable. The governing principle for ionising radiation dose
Accuracy	Closeness to the true value. Freedom from systematic error
Precision	Closeness of repeated measurements to one another. Freedom from random error
Bias	Difference between the average measurement and the true value. Precise-but-biased is the dangerous failure
Repeatability / reproducibility	Precision under identical conditions / across operators, instruments, days, and sites
Calibration	Establishing the relationship between instrument output and known standards

Term	Meaning in this book
Traceability (metrological)	An unbroken chain of calibrations linking a result to a reference standard
Uncertainty	The quantified doubt about a result. A result without it is incomplete
Interference	Response to something other than the intended analyte
Sensitivity (diagnostic)	Of those with the condition, the proportion the test detects
Specificity (diagnostic)	Of those without it, the proportion the test correctly clears
PPV / NPV	Probability the condition is present given a positive result / absent given a negative one. **Both depend on prevalence**
Prevalence	How common the condition is in the population tested. Why a 99% accurate test can still mostly produce false positives
ROC curve	Sensitivity against 1 − specificity across all thresholds. Describes the trade-off; does not choose it
Point-of-care testing	Testing near the patient. Trades analytical performance and traceability for speed and access
Immunoassay	Detection using antibody–antigen binding. ELISA in the lab, lateral flow on a strip
PCR	Exponential amplification of a target DNA sequence by thermal cycling
Beer–Lambert law	Relates absorbance to concentration. Behind both spectrophotometry and pulse oximetry

F.19. Diagnostic devices

Term	Meaning in this book
Ionising radiation	Photons energetic enough to break chemical bonds: X-rays and gamma rays
Effective dose (mSv)	Weighted whole-body estimate of radiation risk, for comparing procedures
kVp / mAs	X-ray tube voltage (photon energy, penetration) / tube current × time (photon number, noise, dose)
Photoelectric absorption	Photon fully absorbed; strongly dependent on atomic number. The source of X-ray contrast
Compton scatter	Photon deflected by an electron; degrades contrast
Contrast agent	Substance that changes image contrast: iodine and barium (X-ray, CT), gadolinium (MRI), microbubbles (ultrasound)
Fluoroscopy	Real-time continuous X-ray imaging
DEXA	Dual-energy X-ray absorptiometry; bone mineral density

Term	Meaning in this book
Filtered back-projection / iterative reconstruction	Methods of computing CT slices from projections
Hounsfield unit (HU)	Calibrated CT attenuation scale; water 0, air −1000
Window level / width	Display mapping of a band of HU to grey levels. Not the data
Larmor frequency	Precession frequency of nuclei in a magnetic field; 42.58 MHz/T for hydrogen
T1 / T2 relaxation	Recovery of magnetisation along the field / decay of the transverse signal; tissue-specific
TR / TE	MRI repetition time / echo time; set the image weighting
k-space	MRI raw data: the image's spatial-frequency spectrum
DWI	Diffusion-weighted MRI; shows acute stroke within minutes
SAR	Specific absorption rate (W/kg); RF energy deposited during MRI
MR Safe / Conditional / Unsafe	ASTM F2503 labels for items in the MRI environment
Quench	Sudden loss of superconductivity, venting the magnet's helium
Piezoelectric transducer	Converts voltage to pressure and back; the ultrasound probe
B-mode / M-mode	Ultrasound brightness image / one line plotted over time
Doppler shift	Frequency change from moving reflectors; gives blood velocity
Elastography	Measurement of tissue stiffness, e.g. by shear-wave speed
MI / TI	Ultrasound mechanical index (cavitation) / thermal index (heating)
Radiopharmaceutical	Biologically targeted molecule carrying a radioactive label
Half-life	Time for half of a radioactive sample to decay
SPECT	Tomographic gamma-camera imaging of single-photon emitters such as Tc-99m
PET	Imaging by coincidence detection of the paired photons from positron annihilation
FDG	Radioactive glucose analogue; the main PET tracer
SUV	Standardised uptake value: PET tracer concentration normalised to dose and body size
Coincidence detection	Two detectors firing within nanoseconds, defining a line of response
Endoscope	Flexible or rigid instrument carrying light, camera, and working channels into the body
Holter monitor	Ambulatory ECG over 24–48 hours

Term	Meaning in this book
Oscillometric blood pressure	Automated method estimating pressures from cuff oscillations; peak ≈ MAP
FEV_1 / FVC	Volume exhaled in one second / total forced exhalation; their ratio identifies obstruction
Polysomnography	Multichannel overnight sleep study
AHI	Apnoea–hypopnoea index: breathing events per hour of sleep
Biopsy	Removal of tissue for microscopic diagnosis
√N law	In counting measurements, SNR rises only as the square root of counts

F.20. Treatment and life-support devices

Term	Meaning in this book
Life-supporting device	One whose failure or misuse can kill within minutes; highest risk classes
Single-fault safety	Remaining safe, or failing safe with an alarm, after any one component failure
Positive-pressure ventilation	Pushing gas into the lungs, the reverse of natural breathing
Volume control / pressure control	Ventilator delivers a set volume (pressure varies) / a set pressure (volume varies)
PEEP	Positive end-expiratory pressure; keeps alveoli open at the end of exhalation
FiO_2	Fraction of inspired oxygen, 0.21 (air) to 1.0
Tidal volume	Volume of one breath; lung-protective ventilation uses small volumes scaled to predicted body weight
Auto-PEEP (breath stacking)	Gas trapped because exhalation time is short relative to the time constant
Time constant ($\tau = R \times C$)	Characteristic emptying time of the lungs; ~3τ for near-complete exhalation
CPAP / bilevel NIV	One constant airway pressure by mask / alternating higher and lower pressures assisting each breath
High-flow nasal oxygen	Heated, humidified gas at up to ~60 L/min through nasal prongs
Oxygen concentrator	Device using pressure-swing adsorption to produce ~90–95% oxygen from air
Volumetric pump / syringe driver	Infusion by squeezing an administration set / by driving a syringe plunger
Free flow	Uncontrolled gravity flow when a set is removed unclamped

Term	Meaning in this book
Occlusion alarm	Alarm on rising line pressure; slow at low rates; release can give a bolus
Siphoning	Uncontrolled syringe emptying under gravity when mounted high and unsecured
Dose error reduction software (drug library)	Pump software with soft and hard dose limits per drug and care area
Smart pump	Infusion pump with a drug library, often networked
Automated insulin delivery	Insulin pump adjusted by an algorithm from CGM readings; "closed loop"
Defibrillation	Shock that depolarises enough heart muscle to stop ventricular fibrillation
Biphasic waveform	Defibrillation current that reverses direction; effective at lower energy
Transthoracic impedance	Electrical resistance of the chest to a shock; typically tens to ~150 Ω
Cardioversion	Smaller shock synchronised to the R wave to reset an organised abnormal rhythm
AED	Automated external defibrillator, designed for untrained users
Pacemaker / pacing code	Implanted pulse generator that paces when needed / letter code (e.g. VVI, DDD) for chamber paced, sensed, response
ICD	Implantable cardioverter-defibrillator; detects and treats dangerous fast rhythms
Haemodialysis	Blood cleaned outside the body across a semipermeable membrane
Countercurrent flow	Blood and dialysate flowing in opposite directions to maintain a gradient
Ultrafiltration	Removal of water from blood by a pressure difference across a membrane
Clearance (K) / Kt/V	Blood volume cleared of a solute per minute / dimensionless dialysis dose
ECMO / cardiopulmonary bypass	External oxygenation and pumping for days to weeks / during heart surgery
Electrosurgery (diathermy)	Cutting and coagulating with radiofrequency current
Monopolar / bipolar	Current returns through a patient pad / passes only between forceps jaws
Return electrode	Large pad that spreads current to avoid heating; partial detachment causes burns
Surgical fire triangle	Ignition, fuel, oxidiser; each usually controlled by a different team member
Gray (Gy)	Absorbed dose: one joule per kilogram

Term	Meaning in this book
Fractionation	Splitting radiotherapy into many smaller doses to spare normal tissue
Linac / multileaf collimator	Radiotherapy linear accelerator / tungsten leaves that shape the beam
IMRT / VMAT / IGRT	Intensity-modulated, arc, and image-guided radiotherapy
Bragg peak	Depth at which protons deposit most of their energy, with little beyond
Brachytherapy	Radiotherapy from sealed sources placed in or against the tumour
Immobilisation mask	Moulded thermoplastic mask holding the head reproducibly for radiotherapy
Bolus (radiotherapy)	Tissue-equivalent material placed on skin to shift dose
Osteolysis / stress shielding	Bone loss from wear-particle inflammation / bone thinning when a stiff implant carries load
Drug-eluting stent	Stent releasing a drug that reduces re-narrowing
Pharmacokinetics (ADME)	Absorption, distribution, metabolism, excretion
Bioavailability / first-pass metabolism	Fraction of a dose reaching circulation / loss in gut and liver before it does
Therapeutic window	Concentration range between ineffective and toxic; narrow for high-risk drugs
Steady state	Stable average drug level on regular dosing, reached after ~4–5 half-lives
Spacer	Chamber between inhaler and mouth removing the need to coordinate actuation with breath
Autoinjector	Spring-driven injector with a hidden needle, for lay use

F.21. Devices and health informatics

Term	Meaning in this book
Task trainer	A model for practising one isolated skill
High-fidelity simulator	A full-body manikin with programmable physiology
Phantom	A tissue-mimicking object used to train or calibrate an imaging modality
Moulage	Simulated wounds and appearance for scenario realism
EHR / EPR	The electronic health record; the system everything else integrates with
HL7 v2 / FHIR	Health data messaging standards. FHIR is the modern REST/JSON one
SNOMED CT / LOINC / ICD	Terminologies for clinical findings, observations, and diagnoses
Early warning score	Aggregated vital signs triggering escalation. A prompt to a human, not a decision

Term	Meaning in this book
Alarm fatigue	Desensitisation caused by frequent or false alarms. A recognised hazard, not a user failing
AAC	Augmentative and alternative communication: speech output for people who cannot speak
Access method	How a person operates a device — switch, gaze, touch, voice. Usually the core design problem in assistive technology
Synthetic data	Realistic but entirely fictional records. What you develop against instead of real data

F.22. Medical software, health data, and security

Term	Meaning in this book
Intended purpose (software)	What the manufacturer says the software is for; decides whether it is a medical device
SaMD	Software as a Medical Device: software that is itself the device, on general-purpose hardware
SiMD	Software in a Medical Device: firmware or software that is part of a hardware device
IMDRF SaMD categories (I–IV)	Risk tiers from the seriousness of the healthcare situation × whether the output treats/diagnoses, drives, or informs management
Clinical decision support (CDS)	Software that gives clinicians patient-specific recommendations; may or may not be a device depending on its function and transparency
General wellness product	Low-risk product for healthy lifestyle without disease claims; generally not regulated as a device
MDR Rule 11	EU classification rule for software; most decision-informing software is Class IIa or higher
Clinical Safety Officer	Named clinician responsible for clinical risk management of health IT (DCB0129/DCB0160 in England)
Software safety class (A/B/C)	IEC 62304 class by worst harm a software failure could contribute to: none / non-serious / serious injury or death
Segregation (software)	Architectural separation allowing parts of a system to have different safety classes
Configuration management	Knowing and being able to reproduce exactly what was in each release
Problem resolution	Recording, investigating, and closing every reported software problem, with safety assessment
SBOM	Software bill of materials: machine-readable inventory of components and versions

Term	Meaning in this book
Dose error reduction software (DERS)	Infusion pump drug library with hard and soft limits per drug and care area
Hard limit / soft limit	A bound that cannot be overridden / one that can, with deliberate confirmation
Health data	Data revealing information about a person's physical or mental health or care
Special category data	GDPR class including health, genetic, and biometric data; needs an Article 9 condition
Lawful basis	One of the GDPR Article 6 grounds for processing personal data; consent is only one
DPIA	Data protection impact assessment, required before high-risk processing
Direct identifier / quasi-identifier	Identifies alone (name, record number) / identifies in combination (birth date, postcode, sex)
Pseudonymisation	Replacing identifiers with a key held separately; the data remain personal data
Anonymisation	Processing so that no one is reasonably likely to be identified; a high, context-dependent bar
Re-identification	Recovering identities from supposedly de-identified data, often by linkage
k-anonymity	Every combination of quasi-identifier values shared by at least k records
l-diversity	Each k-anonymous group also contains varied sensitive values
Caldicott principles / Guardian	UK principles for confidential health information / the senior person responsible for them
PHI	HIPAA protected health information
Covered entity / business associate	HIPAA-regulated providers, plans, clearinghouses / their contractors handling PHI
Safe Harbor / Expert Determination	The two HIPAA de-identification methods
Minimum necessary	Use or disclose only the information needed for the purpose
Information governance	The organisational function deciding lawful, safe use of information
SMART on FHIR	OAuth 2.0-based standard for launching apps inside EHRs with delegated access
DICOMweb	DICOM services over HTTP (QIDO-RS, WADO-RS, STOW-RS)
UCUM	Unified Code for Units of Measure; machine-readable unit expressions
IHE profile	Implementation guide combining standards for a specific clinical workflow

Term	Meaning in this book
Cyber device	US term for a device with software that can connect to the internet; triggers section 524B duties
Coordinated vulnerability disclosure	Agreed process for reporting and fixing security flaws before public release
MDS2	Manufacturer Disclosure Statement for Medical Device Security: standard vendor security questionnaire
STRIDE	Threat categories: spoofing, tampering, repudiation, information disclosure, denial of service, elevation of privilege
Dataset shift	Difference between training and deployment data that degrades model performance
Shortcut learning	A model exploiting spurious cues (site, scanner, markers) rather than the clinical signal
Locked / adaptive algorithm	A model fixed after release / one that changes with new data
Predetermined change control plan (PCCP)	Pre-specified, pre-validated plan for future modifications to an AI-enabled device

Appendix G: The resource shelf

G.1. Core, free, and sufficient

Resource	How to use it
OpenStax Anatomy and Physiology 2e [1]	Your main reference. Openly licensed, free, well-illustrated. Read the chapter summary first, then the named sections, then the diagrams. Do not read it cover to cover
Crash Course Anatomy & Physiology [6]	Short video companion. One pass for shape, one slower pass while drawing. Pause and replay; reduced playback speed is a legitimate technique
MedlinePlus, Understanding Medical Words [3]	Systematic practice with word roots. Complete it once in week 1
Stanford Student Guide to Biodesign [14]	Keep it beside your project notes from Chapter 18 onward. Toolkits on need finding, observation, interviews, need criteria, and prototyping
FDA, 3D Printing of Medical Devices [25]	Short and genuinely clarifying about how a regulator thinks about a printed object
FDA, Applying Human Factors and Usability Engineering to Medical Devices [20]	The users / tasks / environments framing, and the formative versus summative distinction
CDC, Inside Infection Control [23]	Accessible video explanation of infection control concepts
INACSL Simfographics [28]	One-page summaries of healthcare simulation standards. Simulation Design and Outcomes and Objectives are the two to read

G.2. Optional, when a project justifies it

Resource	When
University of Michigan Anatomy Specialization [2]	When you want external course structure and longer lectures. It should **replace** study time, not add to it. Check current access and payment terms
Biodesign: The Process of Innovating Medical Technologies, 2nd ed. [17]	When you want a deeper treatment of device development. **Borrow it.** Use current official sources for anything regulatory, which dates fast in books
Winter, Biomechanics and Motor Control of Human Movement [4]	If the statics in Chapter 6 interested you. The standard reference for segment parameters and inverse dynamics

Resource	When
A clinical anatomy atlas	Only when a specific project needs detail OpenStax does not carry. Not before
WCAG [22]	If your project has any interface at all, physical or digital

G.3. Where to find people

Organisation	Note
Tikkun Olam Makers [18]	International community matching makers with people holding unmet needs
REMAP [31]	UK charity; volunteer makers build custom equipment for disabled people
UTMB MakerHealth [32]	Example of a hospital-embedded maker space
Local hospital simulation centres	Usually the easiest first approach
Local makerspaces	Often already have healthcare contacts
IPEM, AAMI, IEEE EMBS, BMES	Professional bodies; literature, events, people

G.4. What not to buy yet

- Specialist anatomical atlases.
- Segmentation or CAD software licences for imaging work.
- Tissue simulant materials in quantity.
- Standards documents. Use a library first.
- Any certification course, until you have had a conversation with a potential host about which one would actually help.

Everything needed to complete this book is free to view or available from a library. The binding constraint on your progress will be **finding a collaborator**, not acquiring resources — which is why Chapter 18's outreach exercise belongs in week 1.

G.5. A note on currency

Resource links were checked through their public pages on **2 October 2026**. Access arrangements, course offerings, URLs, and institutional processes all change.

US agency references are included because they are freely readable and clearly written. **The host organisation identifies the requirements that actually apply.** Study times and exit criteria

throughout this book are recommendations for this learning plan, not requirements imposed by any of the linked organisations.

Appendix H: Figure credits and licences

Most figures in this book are **generated at build time** from PlantUML, Graphviz, and ditaa source embedded in the chapter files, and are covered by the same terms as the text.

The anatomical illustrations are **reused from openly licensed sources**, listed below. They are reproduced unmodified apart from downscaling for print, and, for `oa-pet.png`, taking the first frame of an animation.



Why these are reused rather than drawn or generated. Anatomical accuracy is not something this book is in a position to certify, and Chapter 21 argues at length that a representation which quietly teaches a false relationship is worse than no representation at all. That argument applies to pictures as much as to physical models.

These figures come from professionally drawn and reviewed open sources — principally OpenStax *Anatomy and Physiology* [1] — so the accuracy rests on their review process rather than on ours.

H.1. Reused figures

File	Original title	Licence	Creator	Ch.
`oa-planes.jpg`	Planes of Body	CC BY 4.0	Connexions	1
`oa-directions.jpg`	Directional Terms	CC BY 3.0	OpenStax. Anatomy & Physiology, Connexions Web	1
`oa-cavities.jpg`	Dorsal Ventral Body Cavities	CC BY 4.0	Connexions	1
`oa-chem-periodic.png`	Simple Periodic Table Chart-en (rendered from SVG to PNG)	Public domain	Offnfopt	2
`oa-cell-animal.png`	Animal cell structure en (SVG, rendered to PNG)	Public domain	LadyofHats (Mariana Ruiz)	4
`oa-cell-cycle.jpg`	0329 Cell Cycle	CC BY 4.0	OpenStax	4
`oa-tissues.jpg`	401 Types of Tissue	CC BY 3.0	OpenStax	5
`oa-skin.jpg`	501 Structure of the skin	CC BY 3.0	OpenStax	5
`oa-longbone.jpg`	603 Anatomy of a Long Bone	CC BY 4.0	OpenStax Anatomy and Physiology	6
`oa-armbones.png`	Human arm bones diagram	Public domain	LadyofHats Mariana Ruiz Villarreal	6
`oa-synovial.jpg`	907 Synovial Joints	CC BY 3.0	OpenStax	6
`oa-jointtypes.jpg`	909 Types of Synovial Joints	CC BY 3.0	OpenStax	6
`oa-movements.jpg`	Body Movements II	CC BY 4.0	Tonye Ogele CNX	6

File	Original title	Licence	Creator	Ch.
oa-musclefibre.jpg	1022 Muscle Fibers (small)	CC BY 4.0	OpenStax	6
oa-neuron.png	Blausen 0657 MultipolarNeuron	CC BY 3.0	BruceBlaus	7
oa-eye.jpg	1413 Structure of the Eye	CC BY 3.0	OpenStax College	8
oa-ear.jpg	1404 The Structures of the Ear	CC BY 4.0	OpenStax	8
oa-heartflow.jpg	2101 Blood Flow Through the Heart	CC BY 3.0	OpenStax	9
oa-respiratory.png	Respiratory system complete en	Public domain	LadyofHats, Jmarchn	9
oa-int-nephron.jpg	2618 Nephron Secretion Reabsorption	CC BY 3.0	OpenStax College	10
oa-int-digestive.jpg	2401 Components of the Digestive System	CC BY 3.0	OpenStax College	10
oa-int-endocrine.jpg	1801 The Endocrine System	CC BY 3.0	OpenStax College	10
oa-chestxray.jpg	Normal posteroanterior (PA) chest radiograph (X-ray)	CC0	Mikael Häggström	16
oa-ct-brain.jpg	Computed tomography of human brain - large	CC0	Department of Radiology, Uppsala University Hospital	16
oa-mri-weighting.jpg	T1t2PD	Public domain	KieranMaher	16
oa-doppler.jpg	CarotidDoppler1	CC BY 3.0	Etan J. Tal	16
oa-pet.png	PET-MIPS-anim (first frame, converted from GIF)	Public domain	Jens Maus	16
oa-ecg12.jpg	12 lead ECG of a 26 year old male	Public domain	MoodyGroove	16
oa-ther-aed.jpg	Modern AED Automated external defibrillator (converted from PNG to JPEG)	CC0	VoLid27	17
oa-ther-hip.jpg	Postoperative radiograph of hip prosthesis - anteroposterior view	CC0	Mikael Häggström	17

All originals are hosted on Wikimedia Commons; each is reachable at `https://commons.wikimedia.org/wiki/File:` followed by the original title above, with spaces replaced by underscores.

H.2. What the licences require

Licence	Obligation
CC BY 3.0 https://creativecommons.org/licenses/by/3.0	Credit the creator, name the licence, link to it, and say whether you changed the work. Commercial use and redistribution permitted. The table above discharges this.
CC BY 4.0 https://creativecommons.org/licenses/by/4.0	As above, with clearer wording and an explicit 30-day window to cure a breach.
Public domain	No obligation. Credited here as good practice.



If you modify any of these figures — recolour, crop beyond trimming, relabel, or combine them with other material — change the credit to read **"adapted from"**. That is the part people most often forget, and it is the part the licence actually requires.

Every figure here is CC BY or public domain, which keeps the licensing of the whole book uniform and simple. If you later add a **CC BY-SA** figure, note that ShareAlike imposes further conditions on adaptations. It is worth preserving the current position.

Appendix I: Device and software index, by maker accessibility

The chapters list devices, software, and training technology by **body system**. This appendix re-cuts the same material by the question you will actually ask: **what could I work on?**



Accessibility here means **technical and regulatory reachability for a supervised learner**. It is not permission. Every tier still requires a sponsor, a review route, and the host's agreement — and the boundary between tiers is easy to cross by accident. "Could you make it a bit stronger so they can lean on it?" moves a project up two tiers in one sentence.

I.1. Tier 1 — start here

No patient contact. No clinical claim. Real demand.

Project	Why it works	Chapter
Simulator spare parts and consumables	Manikin skins, airways, jaw hinges, replaceable inserts. Expensive, slow to source, constantly needed	Chapter 1, Chapter 9
Suture and skin pads	Consumed by the hundred in surgical skills teaching; layered silicone is squarely within a maker's competence	Chapter 5
Ultrasound and needle-guidance phantoms	Gelatin or agar with embedded targets; cheap, effective, perishable so continually re-needed	Chapter 15
UV hand hygiene training box	Light-tight enclosure plus UV-A source; makes an invisible problem visible	Chapter 13
Articulated anatomical and mechanism models	Your capstone's category. Joints, levers, force paths	Chapter 6, Chapter 21
Mechanical test lung	Settable compliance and resistance; teaches respiratory mechanics honestly with no electronics	Chapter 9
Molecular and anatomical teaching models	Protein Data Bank and imaging-derived geometry made physical — printed, laser-cut in stacked layers, milled, or cast from a mould. Educators want them	Chapter 3
Moulage and scenario props	Realism for team training. A craft skill with steady demand	Chapter 5

Project	Why it works	Chapter
Airway and surgical-airway trainer consumables	Replaceable necks, cricothyroid membranes, tongues, and teeth. Used up on every attempt	Chapter 8
Ear and eye examination trainers	Model ears with interchangeable eardrums; model eyes with interchangeable retinal images. Small, precise, and an honest optics project	Chapter 8
Laparoscopic box trainers	A box, a camera, ports, and tasks. Cheap, wanted, and well documented	Chapter 16, Chapter 21
Imaging physics demonstrations and software	CT reconstruction from simulated projections, k-space exercises, Doppler and pulse-echo demonstrators on objects — all on public or synthetic data	Chapter 16
Legibility and alarm-audibility checks	Visual-angle and decibel calculations on existing labels, screens, and cues — then a test in the real room. No materials at all	Chapter 8, Chapter 7
Bench fixtures, jigs, and storage	For clinical engineering, labs, and sterile services. Unglamorous and welcome	Chapter 13, Chapter 15
Accessibility auditing and remediation	No workshop needed, free tools, enormous unmet need	Chapter 7
Software on synthetic or public data	Synthea records, PhysioNet signals, TCIA images, open DICOM	Chapter 9, Chapter 15
Examination models with embedded findings	Breast, skin, and diabetic-foot models with lesions of graded size, depth, and stiffness. Silicone casting; the finding is specified with an educator	Chapter 11, Chapter 5
Inhaler technique trainers	Placebo devices and simple flow-feedback aids for teaching inhalation timing	Chapter 11
Easy-read and dementia-friendly information and signage	Plain-language rewrites, pictorial instructions, ward clocks and signs. Test with intended readers	Chapter 12

Project	Why it works	Chapter
Paediatric manikin replacement parts	Small skins, airways, and limbs are scarce and expensive	Chapter 12, Chapter 21
Dialysis and diffusion demonstrators	Dialysis tubing, salt solutions, conductivity meter; countercurrent versus co-current. Cheap, safe, and genuinely illuminating	Chapter 10
Catheterisation and stoma trainer consumables	Replaceable urethral inserts and silicone stoma models for skills teaching	Chapter 10, Chapter 21
Unit-safe data handling for glucose and kidney values	Software on synthetic or public data; units carried in the type	Chapter 10
Synthetic-data health software	Synthea records in a local FHIR server; SMART apps; unit- and code-aware pipelines. No governance needed, real skills, portfolio-ready	Chapter 20
Threat models, SBOMs, and security reviews of your own projects	Data-flow diagrams, STRIDE tables, dependency inventories. Free tools, transferable habits	Chapter 20
Test lungs and respiratory mechanics demonstrators	Adjustable resistance and compliance; shows time constants and air trapping. Never in a gas path to a person	Chapter 17, Chapter 9
CPR, AED, and inhaler technique trainers and feedback devices	Training-only hardware with no energy storage, drug, or needle; feedback on rate, depth, or inhalation flow	Chapter 17, Chapter 21
Pharmacokinetic and infusion-calculation teaching software	One-compartment models and unit-safe calculators on invented data, clearly labelled as teaching tools	Chapter 17

Project	Why it works	Chapter
Batches of teaching aids made by the right method	Machined, cast, thermoformed, sewn, or printed — chosen against quantity and cleaning, with jigs, inspection, and a handover pack	Chapter 14, Chapter 21
Bench test fixtures, jigs, and signal simulators	Metalwork, machining, and electronics for clinical engineering and laboratories; no patient contact	Chapter 14
Soft goods for equipment, not people	Bags, covers, and carriers for pumps, monitors, and kits, in wipeable laminated fabrics	Chapter 14
Teaching instruments for chemistry	LED colorimeters, simple spectrophotometers, pH and conductivity logging rigs, electrochemistry demonstrators: electronics, optics, machined housings, software. Teaching only	Chapter 2
Workshop chemical-safety improvements	Fume extraction, solvent storage, labelled dispensing, written exposure assessments — your own workshop first	Chapter 2
Open-source teaching microscopes and imaging accessories	Optics, machined or printed flexure stages, illumination, cameras, and software; wanted for teaching and outreach	Chapter 4
Biological image analysis on public data	Counting, segmenting, and tracking cells in Fiji, CellProfiler, or Python; no lab or samples needed	Chapter 4
Teaching-lab instruments and accessories	Blue-light gel imagers, low-voltage education kits, labware racks and holders, temperature loggers — for BSL-1 teaching labs	Chapter 4, Chapter 14

I.2. Tier 2 — with a named clinical partner

Still no patient contact, but the output informs clinical work or carries an accuracy claim someone must verify.

Project	What it requires	Chapter
Anatomical models from patient imaging	Information governance approval, verified anonymisation, anatomical verification, documented pipeline and limitations	Chapter 15
Procedural task trainers	Educator-defined objective and approved teaching method; tactile fidelity agreed by judgement, not durometer	Chapter 21
Imaging phantoms for quality control or teaching on clinical scanners	CT, MRI, and nuclear medicine phantoms used on real equipment need the medical physicist's specification and approval	Chapter 16
Interface pressure and shear demonstration rigs	Teaching only. Must not be presented as a measurement of safety	Chapter 5
Clinical training software	Fictional data, unmissable simulation labelling, educator-defined behaviour	Chapter 9
Adapted equipment mock-ups	Occupational therapist **and** the intended user. Any trial follows the team's review process	Chapter 6, Chapter 7
Feeding and drinking aids	Adapted cutlery, plate guards, cut-away cups. Speech and language therapist or OT input, because swallowing safety and diet texture are clinical decisions	Chapter 8
Switch and alternative access hardware	Highly individual; OT-led. Electrically isolated, battery powered, no patient connection	Chapter 7
Workflow and logistics software	Real operational value, comparatively low clinical risk — but real data means real governance	Chapter 1
Daily-living aids for people with chronic conditions	e.g. foot-inspection aids for diabetes. Requirements derived from disease mechanisms, verified with the clinician and the users; no wound contact and no clinical claim	Chapter 11

Project	What it requires	Chapter
Objects for children's wards or older-adult groups	Play specialist or therapist input; small-parts, cord, entrapment, and cleanability review	Chapter 12
Catheter drainage-bag holders and continence aids	Infection prevention review; nursing input on emptying and measurement; OT and user for self-catheterisation aids	Chapter 10
Educational and workflow apps handling learner or staff data	Approved hosting, information governance review, privacy notice, retention plan; health data about learners turns it into a study	Chapter 20
Integration and interface work on test systems	Vendor test environments, synthetic data, and the host's integration team; never live interfaces without change control	Chapter 20
Radiotherapy positioning aids and bolus	Medical physics specification, measurement, and commissioning; they sit in the beam	Chapter 17
Prosthetic and orthotic accessories and cosmetic covers	The prosthetist or orthotist and the user; nothing load-bearing	Chapter 17, Chapter 6
Drug-library and alarm data analysis	Pharmacy, clinical engineering, and information governance; anonymised logs	Chapter 17, Chapter 20
Battery-powered accessibility electronics	Switch-adapted toys and controls, timers, prompts; OT and user involvement; protected cells, no patient electrical connection, cleaning review	Chapter 14, Chapter 7
Textile positioning and carrying aids	Therapist input on pressure and fit; laundering and flammability checked against the actual process	Chapter 14, Chapter 5
Modifications to existing equipment	Clinical engineering approval; the manufacturer's position on the modification	Chapter 14
Laboratory fixtures, racks, and automation aids	The lab's agreement on chemical compatibility, cleaning, and — in regulated labs — validation	Chapter 2, Chapter 15

Project	What it requires	Chapter
Disinfectant preparation and dispensing aids	Infection prevention ownership, COSHH assessment, material compatibility, and measured verification of the concentration achieved	Chapter 2, Chapter 13
Research-lab instrumentation and automation	A host laboratory, its biosafety rules, and its scientists' specification; open liquid handlers, environmental control, custom fixtures, microfluidic moulds	Chapter 4
Teaching bioreactors and instrumented culture systems	Supervised laboratory setting; sensors, control, and data logging	Chapter 4

I.3. Tier 3 — professional territory

Do not approach these as learning projects. They are listed so you recognise them.

Category	Why
Anything patient-connected and electrical	IEC 60601-1: leakage current, isolation, applied parts. The failure mode is electrocution (Chapter 7)
Life-supporting or life-sustaining devices	Ventilators, defibrillators, infusion pumps, ECMO, oxygen delivery, pacing
Anything entering the body	Implants, catheters, needles, airway devices, surgical instruments. Sterilisation, biocompatibility, and classification all engage
Anything in or on the eye, ear canal, nose, mouth, or airway	Contact lenses, earmoulds and hearing devices, nasal interfaces, mouthpieces, tracheostomy equipment, feeding devices. Thin, sensitive, easily injured tissue, and the airway–food crossing (Chapter 8)
Anything for the MRI magnet room, or that handles radioactive material	MR safety determination and radiation protection are made by named specialists under departmental procedures (Chapter 16)
Light sources and lasers aimed where eyes may be	Photobiological and laser safety standards apply; retinal injury is painless (Chapter 8, Appendix D)
Anything that bears load or restricts movement	Orthoses, prostheses, mobility aids, seating, hoists. Prescribed clinical devices

Category	Why
Wound-contact and compromised-skin products	Dressings, NPWT, skin substitutes (Chapter 5)
Diagnostic devices and software that output a clinical conclusion	In vitro diagnostics and clinical decision software are regulated as devices
Medication-handling software	Dose calculation, prescribing, administration. Narrow therapeutic indices (Chapter 3)
Anything using real patient data without approval	Not a technical question. Information governance precedes access (Chapter 18)
Sterile-field items	Materials, sterilisation validation, packaging, and regulatory status all change
Disease-management devices that dose, detect, or decide	Insulin pumps and closed-loop systems, anticoagulation meters, automated sepsis or arrhythmia alerts, offloading footwear. Control of a failed physiological loop, or a clinical claim (Chapter 11, Chapter 20)
Anything for a mental-health ward	Ligature risk assessment by the ward team before introduction; cords and straps are ligatures (Chapter 12)
Paediatric dose calculation and neonatal life support	Tenfold-error risk; incubators and warmers are life-supporting (Chapter 12)
Insulin delivery, dosing software, and closed-loop algorithms	A wrong dose can kill within hours; regulated as medical devices, with cybersecurity obligations (Chapter 10, Chapter 20)
Dialysis, feeding-line, and catheter components	Life-sustaining or body-entering; ISO 80369 connector rules for anything in a feeding line (Chapter 10)
Software with a clinical intended purpose	Anything that diagnoses, monitors, predicts, or recommends treatment for a patient is probably a medical device: IEC 62304, a QMS, and a regulatory route apply (Chapter 20)
Security testing of clinical systems or devices	Only with explicit written authorisation from the owner. A crashed device is a patient safety incident
Ventilators, infusion pumps, defibrillators, dialysis and bypass circuits	Life-supporting: single-fault safety, alarms, and particular IEC/ISO standards. Includes emergency "simple ventilator" designs (Chapter 17)

Category	Why
Energy-delivering surgical and therapy equipment	Electrosurgery, lasers, radiotherapy; anything in the beam or energy path (Chapter 17)
Products placed on the market	A quality management system, validated processes, and a regulatory route; a manufacturer's job, not a maker's (Chapter 14)
Analysers or sensors reporting patient chemistry	In vitro diagnostic devices: validated methods, calibration traceability, quality control, and regulation (Chapter 2, Chapter 15)
Anything that cultures organisms, handles clinical samples, or reports patient results	Biosafety containment, laboratory accreditation (ISO 15189), and in vitro diagnostic regulation apply; centrifuge rotors and electrophoresis power supplies carry serious energy hazards (Chapter 4)

I.4. Open tools and data worth knowing

Everything here is free, runnable on your own machine, and touches no patient.

Tool or dataset	What it is	Use it for
3D Slicer [34]	Medical image computing	DICOM → segmentation → STL
ITK-SNAP, InVesalius	Segmentation	Approachable alternatives to Slicer
pydicom, SimpleITK, VTK, ITK	Python and C++ libraries	Building an imaging pipeline in code
Orthanc, dcm4che, OHIF Viewer	DICOM server, toolkit, web viewer	A complete local imaging stack
The Cancer Imaging Archive [35]	Large anonymised imaging datasets	Practice data that is properly cleared
PhysioNet [36]	Published physiological signal recordings	ECG, PPG, EEG, respiratory signal processing
NeuroKit2, MNE-Python, EEGLAB	Biosignal analysis libraries	Signal processing without hardware
OpenSim [37]	Musculoskeletal modelling	Turning Chapter 6' statics into real analysis
OpenPose, MediaPipe	Markerless pose estimation	Movement analysis from ordinary video

Tool or dataset	What it is	Use it for
Protein Data Bank [33]	Open structural biology database	Molecular models for printing
PyMOL, ChimeraX, Mol, Avogadro, RDKit*	Molecular visualisation and cheminformatics	Chemistry models and analysis
Synthea [39]	Synthetic patient record generator	Realistic fictional data for development
HAPI FHIR, OpenMRS, openEHR	Open health record toolkits and platforms	Learning clinical data standards safely
NVDA	Open-source screen reader	Accessibility testing (Chapter 7)
axe, WAVE, Lighthouse	Accessibility audit tools	Finding interface barriers
Z-Anatomy, BodyParts3D	Open 3D anatomy models	Teaching models and reference
Blender, FreeCAD, OpenSCAD	Open 3D modelling and CAD	Everything downstream



If you want one afternoon that moves you furthest: install 3D Slicer, load a public dataset from TCIA, segment a bone, export a surface model, and make it physical — print it, mill it, or use it to make a mould and cast it — then write down every threshold and measure one dimension against the source images.

That single exercise touches imaging physics, segmentation, mesh processing, fabrication, **and** the documentation discipline the whole book argues for. It produces a portfolio item. And it uses nobody's data but the one person who consented to it being public.

Bibliography

Core texts and courses

- [1] OpenStax. *Anatomy and Physiology 2e*. Rice University, openly licensed. https://openstax.org/books/anatomy-and-physiology-2e/pages/1-introduction
- [2] University of Michigan. *Anatomy Specialization*. Coursera. https://www.coursera.org/specializations/anatomy
- [3] MedlinePlus. *Understanding Medical Words*. US National Library of Medicine. https://medlineplus.gov/medwords/
- [4] Winter, D. A. *Biomechanics and Motor Control of Human Movement*, 4th ed. Wiley.
- [5] Published anthropometric data sets — for example national anthropometric surveys and the standard ergonomics data compilations. Consult the edition and population relevant to your users.

Video companions

- [6] *Crash Course Anatomy & Physiology* — series index. https://thecrashcourse.com/topic/anatomy/
- [7] *Introduction to Anatomy & Physiology* — Crash Course A&P #1. https://thecrashcourse.com/courses/introduction-to-anatomy-physiology-crash-course-ap-1/
- [8] *Tissues, Part 1* — Crash Course A&P #2. https://thecrashcourse.com/courses/tissues-part-1-crash-course-ap-2/
- [9] *The Integumentary System, Part 1: Skin Deep* — Crash Course A&P #6. https://thecrashcourse.com/courses/the-integumentary-system-part-1-skin-deep-crash-course-ap-6/
- [10] *The Nervous System, Part 1* — Crash Course A&P #8. https://thecrashcourse.com/courses/the-nervous-system-part-1-crash-course-ap-8/
- [11] *Joints* — Crash Course A&P #20. https://thecrashcourse.com/courses/joints-crash-course-ap-20/
- [12] *The Heart, Part 1: Under Pressure* — Crash Course A&P #25. https://thecrashcourse.com/courses/the-heart-part-1-under-pressure-crash-course-ap-25/
- [13] *The Respiratory System, Part 1* — Crash Course A&P #31. https://thecrashcourse.com/courses/respiratory-system-part-1-crash-course-ap-31/

Design process and biodesign

- [14] Stanford Byers Center for Biodesign. *Student Guide to Biodesign*. https://biodesignguide.stanford.edu/
- [15] Stanford Biodesign. *Video Gallery*. https://biodesignguide.stanford.edu/video-gallery/
- [16] Stanford Biodesign. *Textbook and Videos*. https://biodesign.stanford.edu/resources/learning/

textbook-videos.html

- [17] Yock, P. G., Zenios, S., Makower, J., et al. *Biodesign: The Process of Innovating Medical Technologies*, 2nd ed. Cambridge University Press.
- [18] Tikkun Olam Makers. https://www.tikkunolammakers.org/
- [19] Tikkun Olam Makers. *Product Development*. https://www.tikkunolammakers.org/product-development

Human factors, usability, and accessibility

- [20] US Food and Drug Administration. *Applying Human Factors and Usability Engineering to Medical Devices* — guidance for industry and FDA staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/applying-human-factors-and-usability-engineering-medical-devices
- [21] US Food and Drug Administration. *CDRH Learn*. https://www.fda.gov/Training/CDRHLearn/
- [22] World Wide Web Consortium. *Web Content Accessibility Guidelines*. https://www.w3.org/WAI/standards-guidelines/wcag/

Infection control and materials

- [23] US Centers for Disease Control and Prevention. *Project Firstline / Inside Infection Control*. https://cdc.gov/project-firstline/hcp/training/covid-19/index.html
- [24] US Centers for Disease Control and Prevention. *Guideline for Disinfection and Sterilization in Healthcare Facilities*. https://www.cdc.gov/infection-control/hcp/disinfection-sterilization/
- [25] US Food and Drug Administration. *3D Printing of Medical Devices*. https://www.fda.gov/medical-devices/3d-printing-medical-devices/process-3d-printing-medical-devices
- [26] US Food and Drug Administration. *Technical Considerations for Additive Manufactured Medical Devices*. https://www.fda.gov/media/97633/download

Measurement and devices

- [27] US Food and Drug Administration. *Pulse Oximeter Accuracy and Limitations* — safety communication and related guidance. https://www.fda.gov/medical-devices/safety-communications/pulse-oximeter-accuracy-and-limitations-fda-safety-communication

Simulation and education

- [28] International Nursing Association for Clinical Simulation and Learning. *Simfographics*. https://www.inacsl.org/simfographics
- [29] INACSL. *Healthcare Simulation Standards of Best Practice*. https://inacsl.memberclicks.net/healthcare-simulation-standards

Tissue and contact

- [30] European Pressure Ulcer Advisory Panel, National Pressure Injury Advisory Panel, and Pan Pacific Pressure Injury Alliance. *Prevention and Treatment of Pressure Ulcers/Injuries: Clinical Practice Guideline*. Cited to show the shape and ownership of the professional literature, not as a source of design values.

Routes to contribute

- [31] REMAP. *Engineer / Maker volunteering*. https://remap.org.uk/volunteering-roles/engineer-maker/
- [32] University of Texas Medical Branch. *MakerHealth*. https://www.utmb.edu/maker/home

Open tools, data, and reference platforms

- [33] RCSB Protein Data Bank. Open structural biology archive. https://www.rcsb.org/ — see also the *Molecule of the Month* series.
- [34] 3D Slicer. Free, open-source medical image computing platform. https://www.slicer.org/
- [35] The Cancer Imaging Archive (TCIA). Openly available, de-identified medical imaging datasets. https://www.cancerimagingarchive.net/
- [36] PhysioNet. Repository of physiological signal recordings and related software. https://physionet.org/
- [37] OpenSim. Open-source musculoskeletal modelling and simulation. https://opensim.stanford.edu/
- [38] Radiopaedia. Open-edit radiology reference and case library. https://radiopaedia.org/
- [39] Synthea. Synthetic patient population and health-record generator. https://synthetichealth.github.io/synthea/

Special senses, head and neck

- [40] World Health Organization. *World report on hearing*. 2021. https://www.who.int/publications/i/item/9789240020481
- [41] World Health Organization. *World report on vision*. 2019. https://www.who.int/publications/i/item/9789241516570
- [42] International Dysphagia Diet Standardisation Initiative. *The IDDSI Framework*. https://iddsi.org/

Diagnostic imaging and devices

- [43] Radiological Society of North America and American College of Radiology. *RadiologyInfo.org* — patient information on imaging tests and procedures. https://www.radiologyinfo.org/
- [44] Elster, A. D. *Questions and Answers in MRI*. https://mriquestions.com/

- [45] *Image Wisely* — radiation safety in adult medical imaging. https://www.imagewisely.org/
- [46] Bushberg, J. T., Seibert, J. A., Leidholdt, E. M., and Boone, J. M. *The Essential Physics of Medical Imaging*, 4th ed. Wolters Kluwer.

Lifespan and the mind

- [47] World Health Organization. *WHO Child Growth Standards*. https://www.who.int/tools/child-growth-standards
- [48] World Health Organization. *Ageing and health* (fact sheet). https://www.who.int/news-room/fact-sheets/detail/ageing-and-health
- [49] Fried, L. P., Tangen, C. M., Walston, J., et al. Frailty in older adults: evidence for a phenotype. *Journals of Gerontology Series A* 56(3), M146–M156, 2001.
- [50] Cowan, N. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 24(1), 87–114, 2001.
- [51] NHS England. *Health Building Note 03-01: Adult acute mental health units*. Department of Health / NHS England.

Kidneys, gut, and hormones

- [52] Kidney Disease: Improving Global Outcomes (KDIGO). *Clinical Practice Guideline for the Evaluation and Management of Chronic Kidney Disease*. https://kdigo.org/guidelines/ckd-evaluation-and-management/
- [53] OpenAPS. *Open Artificial Pancreas System* — reference design and documentation. https://openaps.org/
- [54] Global Enteral Device Supplier Association (GEDSA). *Stay Connected* — ENFit and ISO 80369 information. https://www.stayconnected.org/

Medical software, health data, and security

- [55] International Medical Device Regulators Forum. *Software as a Medical Device: Possible Framework for Risk Categorization and Corresponding Considerations* (IMDRF/SaMD WG/N12FINAL). 2014. https://www.imdrf.org/
- [56] US Food and Drug Administration. *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff*. 2022. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software
- [57] US Food and Drug Administration. *Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions*. Final guidance 2023, since updated. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions
- [58] Medical Device Coordination Group. *MDCG 2019-11: Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745 — MDR and Regulation (EU) 2017/746 — IVDR*. European Commission. https://health.ec.europa.eu/medical-devices-sector/new-regulations/guidance-mdcg-endorsed-documents-and-other-guidance_en

- [59] US Department of Health and Human Services. *Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the HIPAA Privacy Rule*. 2012. https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification/
- [60] HL7 International. *FHIR — Fast Healthcare Interoperability Resources*. https://hl7.org/fhir/

Treatment and life-support devices

- [61] Leveson, N. G., and Turner, C. S. "An Investigation of the Therac-25 Accidents." *IEEE Computer* 26(7), 1993, 18–41.
- [62] The Joint Commission. *Sentinel Event Alert 50: Medical device alarm safety in hospitals*. 2013. https://www.jointcommission.org/
- [63] Institute for Safe Medication Practices. Medication safety alerts and resources. https://www.ismp.org/
- [64] Medicines and Healthcare products Regulatory Agency (UK). *Specification for ventilators to be used in UK hospitals during the coronavirus (COVID-19) outbreak* (Rapidly Manufactured Ventilator System). 2020. https://www.gov.uk/
- [65] European Resuscitation Council. *European Resuscitation Council Guidelines 2021*. https://www.erc.edu/

General chemistry

- [66] OpenStax. *Chemistry 2e*. Rice University, openly licensed. https://openstax.org/details/books/chemistry-2e
- [67] *Crash Course Chemistry* — series index. https://thecrashcourse.com/topic/chemistry/
- [68] National Center for Biotechnology Information. *PubChem*. https://pubchem.ncbi.nlm.nih.gov/
- [69] National Institute of Standards and Technology. *NIST Chemistry WebBook*. https://webbook.nist.gov/chemistry/
- [70] United Nations Economic Commission for Europe. *Globally Harmonized System of Classification and Labelling of Chemicals (GHS)*. https://unece.org/about-ghs

Cell and molecular biology

- [71] OpenStax. *Biology 2e*. Rice University, openly licensed. https://openstax.org/details/books/biology-2e
- [72] Alberts, B., Heald, R., Johnson, A., Morgan, D., Raff, M., Roberts, K., and Walter, P. *Molecular Biology of the Cell*, 7th ed. W. W. Norton, 2022.
- [73] Milo, R., and Phillips, R. *Cell Biology by the Numbers*. Garland Science, 2015. Free online edition: http://book.bionumbers.org/
- [74] World Health Organization. *Laboratory Biosafety Manual*, 4th ed. 2020. https://www.who.int/publications/i/item/9789240011311
- [75] US Centers for Disease Control and Prevention and National Institutes of Health. *Biosafety*

in Microbiological and Biomedical Laboratories (BMBL), 6th ed. 2020. https://www.cdc.gov/labs/bmbl/

- [76] National Center for Biotechnology Information. Databases and tools, including Gene, Nucleotide, and BLAST. https://www.ncbi.nlm.nih.gov/

Note on citation

Links were checked through their public pages on 2 October 2026. Standards referenced in Appendix D (ISO, IEC, AAMI, ASTM, EN) are paid documents and are cited by number rather than by link; consult a library or your host institution for access.