Clinical Data Science in R

Methods in Clinical AI and Translational Trauma Research

R. Heller

Table of contents

Preface **1**
Who this book is for . . . 1
How the book is organised . . . 1
Data status . . . 2
Reading paths in the course . . . 3
Companion books and courses . . . 4
Programming . . . 5
History of this book . . . 5
Citation and licence . . . 5

I. Part I · The question before the model **7**

1. Clinical decisions and the translational chain **9**
1.1. Case and decision . . . 9
1.2. Core reading . . . 10
1.2.1. An output is worth the action it changes . . . 10
1.2.2. The decision map . . . 11
1.2.3. From rat to patient: the translational chain . . . 13
1.2.4. Two teaching datasets that must stay apart . . . 14
1.3. Lab . . . 14
1.3.1. Step 1: R as a calculator . . . 15
1.3.2. Step 2: a table of Ms K.'s pathway . . . 15
1.3.3. Step 3: the first graphic . . . 17
1.3.4. Step 4: the decision map as data . . . 19
1.3.5. Step 5: two datasets, two units . . . 20

1.4. Counter-check 21
1.5. Consequences for patients and operations 22
1.6. Look up and transfer 23
1.7. Further reading 24

2. Describe, predict, explain **27**
2.1. Case and decision 27
2.2. Core reading 28
2.2.1. Three questions about the same variables 28
2.2.2. The target quantity has to be written down 30
2.2.3. A small causal diagram 30
2.2.4. Why the sign flips 32
2.2.5. What would answer the head of department 32
2.3. Lab 33
2.3.1. Step 1: the graph as an R object 33
2.3.2. Step 2: data that follow the graph 36
2.3.3. Step 3: three analyses, three answers 37
2.3.4. Step 4: explaining the sign 38
2.3.5. Step 5: is it chance? 39
2.3.6. Step 6: the prognostic question uses the same data differently 40
2.4. Counter-check 41
2.5. Consequences for patients and operations 42
2.6. Look up and transfer 43
2.7. Further reading 44

3. How data arise: missingness and denominators **45**
3.1. Case and decision 45
3.2. Core reading 46
3.2.1. Every row has a history 46
3.2.2. Different reasons, different consequences 48
3.2.3. Many rows, few units 49
3.2.4. Rare combinations identify people 49

3.3. Lab 50
3.3.1. Step 1: the patient table and the planned-visit table 50
3.3.2. Step 2: the data dictionary 52
3.3.3. Step 3: counts and denominators 54
3.3.4. Step 4: patients, measurements and linked records, counted twice 55
3.3.5. Step 5: first trajectory figure, individual courses . . 57
3.3.6. Step 6: second trajectory figure, means with their denominators 58
3.3.7. Step 7: saving the tables for the next chapter 61
3.3.8. Step 8: pseudoreplication in the HSI rat study . . . 61
3.3.9. Step 9: rare combinations in a registry extract . . . 63
3.4. Counter-check 65
3.5. Consequences for patients and operations 66
3.6. Look up and transfer 68
3.7. Further reading 69

4. Reconstructing a research project in R 71
4.1. Case and decision 71
4.2. Core reading 72
4.2.1. Reproducible is a modest, checkable promise 72
4.2.2. A project that runs with one command 73
4.2.3. What a colleague needs to rerun it 74
4.2.4. Versions of code, data and software 74
4.2.5. Every exported file carries its passport 75
4.3. Lab 76
4.3.1. Step 1: R objects in five minutes 76
4.3.2. Step 2: reading a file with a relative path 78
4.3.3. Step 3: a base graphic 78
4.3.4. Step 4: long and wide 81
4.3.5. Step 5: start instructions and data schema 81
4.3.6. Step 6: the scripts 82
4.3.7. Step 7: the one command, run inside this chapter . . 86
4.3.8. Step 8: software version and report 89

4.3.9. Step 9: version control with Git 91
4.3.10. Step 10: the artefact register 91
4.4. Counter-check 92
4.5. Consequences for patients and operations 93
4.6. Look up and transfer 94
4.7. Further reading 95

II. Part II · How models learn and how evidence arises 97

5. From baseline to learning model 99
5.1. Case and decision 99
5.2. Core reading 100
5.3. Lab 105
5.3.1. Step 1 · Generate the synthetic cohort 105
5.3.2. Step 2 · Split once, by patient, and assign the inner folds 107
5.3.3. Step 3 · Learn the preprocessing from development data only 108
5.3.4. Step 4 · Null model, points rule and logistic regression 109
5.3.5. Step 5 · Elastic net, tuned in the inner folds 109
5.3.6. Step 6 · Gradient boosting, tuned in the inner folds 110
5.3.7. Step 7 · Compare all candidates on the inner folds . 111
5.3.8. Step 8 · Final fits and one look at the test set . . . 112
5.3.9. Step 9 · Export the teaching file 115
5.4. Counter-check 116
5.5. Consequences for patients and operations 119
5.6. Look up and transfer 120
5.7. Further reading 121

6. Accuracy, calibration and clinical utility 123
6.1. Case and decision 123
6.2. Core reading 124

6.3. Lab . . . 128
6.3.1. Step 1 · Read the teaching file and confirm the test patients . . . 128
6.3.2. Step 2 · Refit the prespecified logistic regression . . . 129
6.3.3. Step 3 · The 2 × 2 table at a threshold of 20% . . . 129
6.3.4. Step 4 · AUC with an uncertainty interval . . . 130
6.3.5. Step 5 · Calibration: plot, intercept and slope . . . 131
6.3.6. Step 6 · From one threshold to many . . . 132
6.3.7. Step 7 · Net benefit and the decision curve . . . 133
6.3.8. Step 8 · Two groups, with denominators . . . 136
6.4. Counter-check . . . 136
6.5. Consequences for patients and operations . . . 138
6.6. Look up and transfer . . . 139
6.7. Further reading . . . 141

7. Reading an AI study critically 143
7.1. Case and decision . . . 143
7.2. Core reading . . . 144
7.3. Lab . . . 150
7.3.1. Part A · Reconstruct the 2 × 2 table from the published abstract . . . 150
7.3.2. Part B · A journal-club card for a fictitious abstract 152
7.4. Counter-check . . . 155
7.5. Consequences for patients and operations . . . 156
7.6. Look up and transfer . . . 158
7.7. Further reading . . . 159

8. Language models in clinic and science 161
8.1. Case and decision . . . 161
8.2. Core reading . . . 162
8.3. Lab . . . 167
8.3.1. Step 1 · Store the two answers with their provenance 167
8.3.2. Step 2 · A bigram model from twelve sentences . . . 168
8.3.3. Step 3 · Temperature and generation . . . 171

8.3.4. Step 4 · The evidence card: existence and support, judged separately . . . 173
8.3.5. Step 5 · Check the patient letter for readability . . 174
8.3.6. Step 6 · Document origin, model version and use . 176
8.4. Counter-check . . . 178
8.5. Consequences for patients and operations . . . 179
8.6. Look up and transfer . . . 180
8.7. Further reading . . . 182

III. Part III · Diagnostics and measurement in four case studies 183

9. A fracture alert in the emergency department 185
9.1. Case and decision . . . 185
9.2. Core reading . . . 186
9.3. Lab . . . 191
9.4. Counter-check . . . 196
9.5. Consequences for patients and operations . . . 198
9.6. Look up and transfer . . . 200
9.7. Further reading . . . 201

10. Hyperspectral imaging 203
10.1. Case and decision . . . 203
10.2. Core reading . . . 204
10.3. Lab . . . 210
10.4. Counter-check . . . 217
10.5. Consequences for patients and operations . . . 219
10.6. Look up and transfer . . . 220
10.7. Further reading . . . 221

11. Measuring gait from video 223
11.1. Case and decision . . . 223
11.2. Core reading . . . 224

11.3. Lab . . . 231
11.4. Counter-check . . . 237
11.5. Consequences for patients and operations . . . 238
11.6. Look up and transfer . . . 240
11.7. Further reading . . . 241

12. Spatial biology and image analysis 243
12.1. Case and decision . . . 243
12.2. Core reading . . . 244
12.3. Lab . . . 253
12.4. Counter-check . . . 262
12.5. Consequences for patients and operations . . . 263
12.6. Look up and transfer . . . 264
12.7. Further reading . . . 265

IV. Part IV · Complex models for trauma research 267

13. Longitudinal biomarkers 269
13.1. Case and decision . . . 269
13.2. Core reading . . . 270
13.3. Lab . . . 274
13.3.1. Step 1: write down the simulation assumptions . . . 274
13.3.2. Step 2: document the denominators at person level . 276
13.3.3. Step 3: three estimates of the interaction . . . 277
13.3.4. Step 4: sensitivity analysis for late and missing visits 280
13.4. Counter-check . . . 284
13.5. Consequences for patients and operations . . . 286
13.6. Look up and transfer . . . 287
13.7. Further reading . . . 288

14. Cluster, gradient or artefact 291
14.1. Case and decision . . . 291
14.2. Core reading . . . 292

14.3. Lab 295
14.3.1. Step 1: simulate a gradient without groups 296
14.3.2. Step 2: three scalings and a PCA 297
14.3.3. Step 3: k-means with two to five clusters 297
14.3.4. Step 4: a Gaussian mixture model 298
14.3.5. Step 5: bootstrap stability 299
14.3.6. Step 6: look for a trough instead of a cluster count . 300
14.3.7. Step 7: the continuous baseline 301
14.3.8. Step 8: collect the variants and check the target values 301
14.3.9. Step 9: an external test plan 302
14.4. Counter-check 304
14.5. Consequences for patients and operations 307
14.6. Look up and transfer 308
14.7. Further reading 309

15. Mechanistic models and pharmacometrics 311
15.1. Case and decision 312
15.2. Core reading 312
15.3. Lab 317
15.3.1. Step 1: parameters with units 317
15.3.2. Step 2: the numerical solver against the analytical solution 317
15.3.3. Step 3: structural identifiability at a single time . . 319
15.3.4. Step 4: expected precision of each design 319
15.3.5. Step 5: the model-free baseline 321
15.3.6. Step 6: interindividual variability and the two-stage shortcut 322
15.4. Counter-check 325
15.5. Consequences for patients and operations 326
15.6. Look up and transfer 327
15.7. Further reading 328

16. Bayes, small samples and the 3Rs 329
16.1. Case and decision 329

16.2. Core reading 330
16.3. Lab 334
16.3.1. Step 1: the data with their origin 334
16.3.2. Step 2: the frequentist baseline 334
16.3.3. Step 3: two priors and their posteriors 335
16.3.4. Step 4: event rate is not effect 336
16.3.5. Step 5: design of the next experiment 336
16.3.6. Step 6: planning and reporting are different tasks . . 338
16.3.7. Step 7: decision exercise – a letter to the editor . . . 339
16.4. Counter-check 340
16.5. Consequences for patients and operations 341
16.6. Look up and transfer 342
16.7. Further reading 344

V. Part V · Care systems and translation into practice 345

17. Forecasting patient flow 347
17.1. Case and decision 347
17.2. Core reading 348
17.3. Lab 352
17.3.1. The synthetic arrival series 353
17.3.2. Rolling-origin backtest 354
17.3.3. Evaluation per method, week and shift 356
17.3.4. A shift map 359
17.3.5. Weighted error cost 361
17.4. Counter-check 362
17.5. Consequences for patients and operations 363
17.6. Look up and transfer 364
17.7. Further reading 365

18. Mass casualty incidents, operating capacity and simulation 367
18.1. Case and decision 367
18.2. Core reading 368

18.3. Lab . . . 372
18.3.1. Step 1: draw and store the event lists . . . 373
18.3.2. Step 2: the model . . . 374
18.3.3. Step 3: verification against a hand calculation . . . 375
18.3.4. Step 4: run the scenarios and check resource release 376
18.3.5. Step 5: variation between runs, paired and unpaired 378
18.4. Counter-check . . . 380
18.5. Consequences for patients and operations . . . 381
18.6. Look up and transfer . . . 382
18.7. Further reading . . . 383

19. From one centre or species to the next 385
19.1. Case and decision . . . 385
19.2. Core reading . . . 386
19.3. Lab . . . 390
19.3.1. Raw values and the recall model . . . 391
19.3.2. Image-wise split: recognition . . . 392
19.3.3. Animal-wise split: new animals . . . 393
19.3.4. A split plan with zero ID overlap . . . 394
19.3.5. Self-contained variant: a nearest-neighbour model . 395
19.3.6. Variation between independent animals . . . 399
19.3.7. Deepening: a missing modality as a stress test . . . 401
19.4. Counter-check . . . 403
19.5. Consequences for patients and operations . . . 403
19.6. Look up and transfer . . . 404
19.7. Further reading . . . 406

20. Pilot operation, impact and monitoring 409
20.1. Case and decision . . . 409
20.2. Core reading . . . 410
20.3. Lab . . . 413
20.3.1. Step 1: simulate the pilot . . . 414
20.3.2. Step 2: a weekly control chart per shift type . . . 415
20.3.3. Step 3: the per-shift stop rule . . . 416

20.3.4. Step 4: performance with honest denominators . . . 417
20.3.5. Step 5: small aggregates and disclosure 418
20.3.6. The monitoring sheet 420
20.3.7. Decision exercise at the shift handover 425
20.4. Counter-check 426
20.5. Consequences for patients and operations 427
20.6. Look up and transfer 428
20.7. Further reading 430

VI. Part VI · Ethics, law and the protection of data and systems 431

21. Ethics at the bedside and in study design 433
21.1. Case and decision 433
21.2. Core reading 434
21.3. Decision exercise 438
21.4. Counter-check 442
21.5. Consequences for patients and operations 443
21.6. Look up and transfer 445
21.7. Further reading 447

22. Regulating research software and clinical AI 449
22.1. Case and decision 449
22.2. Core reading 450
22.3. Decision exercise 456
22.4. Counter-check 459
22.5. Consequences for patients and operations 459
22.6. Look up and transfer 462
22.7. Further reading 463

23. Data protection along the entire data flow 465
23.1. Case and decision 465
23.2. Core reading 466

23.3. Decision exercise 471
23.4. Counter-check 475
23.5. Consequences for patients and operations 476
23.6. Look up and transfer 478
23.7. Further reading 479

24. Information security and recovery **481**
24.1. Case and decision 481
24.2. Core reading 482
24.3. Lab 488
24.4. Counter-check 495
24.5. Consequences for patients and operations 496
24.6. Look up and transfer 498
24.7. Further reading 499

VII. Part VII · Special topics and project workshop **501**

25. Security risks of learning and generative systems **503**
25.1. Case and decision 503
25.2. Core reading 504
25.2.1. Instructions and data share one channel 504
25.2.2. Trust boundaries 505
25.2.3. Three levels, kept apart 505
25.2.4. A defensive tabletop exercise 509
25.3. Lab 510
25.4. Counter-check 516
25.5. Consequences for patients and operations 519
25.6. Look up and transfer 520
25.7. Further reading 522

26. Secondary use and responsible open science **523**
26.1. Case and decision 523

26.2. Core reading 524
26.2.1. Secondary use is a change of purpose 524
26.2.2. Openness in tiers 525
26.2.3. Anonymity depends on the recipient 527
26.2.4. Reproducibility without exporting raw data 528
26.3. Lab 528
26.4. Counter-check 535
26.5. Consequences for patients and operations 536
26.6. Look up and transfer 539
26.7. Further reading 541

27. Model passport, study report and decision 543
27.1. Case and decision 543
27.2. Core reading 544
27.2.1. One target statement, one passport 544
27.2.2. Three rules for filling in 545
27.2.3. A separate decision sheet 545
27.2.4. Choosing the reporting path 546
27.2.5. The counter-review 548
27.3. Lab 549
27.3.1. The filled passport 552
27.3.2. The decision sheet 556
27.3.3. When an alert goes wrong: who acts next 558
27.3.4. A method-specific short report 559
27.3.5. Counter-review 560
27.3.6. The blank passport and decision sheet 561
27.4. Counter-check 563
27.5. Consequences for patients and operations 565
27.6. Look up and transfer 566
27.7. Further reading 568

28. Case atlas, glossary and transfer 569
28.1. Case and decision 569

28.2. Core reading 570
28.2.1. Start from the decision, not from the data 570
28.2.2. The route by risk 571
28.2.3. Two reading routes as examples 572
28.2.4. Combining routes: the coordinator's three requests . 573
28.2.5. Working with the atlas 573
28.3. Lab 574
28.4. Counter-check 577
28.5. Consequences for patients and operations 578
28.6. Look up and transfer 578
28.6.1. Case atlas 579
28.6.2. Formative model answers 583
28.6.3. Navigator by goal and by data type 586
28.6.4. Glossary (English–German) 588
28.6.5. Register of common mistakes 593
28.6.6. The next experiment on one page 596
28.7. Further reading 597

References **599**

Appendices **625**

A. Edition notes, software and imprint **625**
A.1. Edition 625
A.2. Software 626
A.3. Legal status 627
A.4. Use of AI tools 627
A.5. Imprint 627

B. Chapter register **629**

C. Checklist for contributions and revisions **653**

Preface

R. Heller · Date: 2026-11-08 · Last changed: 2026-11-09

A clinical team receives a fracture alert. A research group measures hundreds of spectra from a handful of animals. An emergency department plans its next shift. This book follows the decision that each of these situations is meant to support. Readers should find three things in every chapter: a calculation they can retrace, a way to try to refute their own claim, and a reasoned answer to the question of who may do what with the result.

Who this book is for

The book is written for medical students, doctoral researchers, clinical research teams and people who review models. Newcomers start from the clinical question and learn the R they need along the way (Chapter 4). Part I assumes basic medical knowledge but no programming experience. Readers who want to check a specific result can enter by data type, method or failure pattern; the navigator in Chapter 28 leads from the starting question to the right chapter.

How the book is organised

The 28 chapters are grouped into seven parts. The chapter number doubles as a fixed identifier used by the course and by teaching teams: chapter 9

is B09. Every chapter follows the same pattern:

1. **Case and decision.** One person or role, one question and an action with consequences.
2. **Core reading.** A labelled figure or table, the key assumption and a question for discussion. The core reading is clearly marked, runs to about 800–1,200 words (8–12 minutes) and is the required reading in the course.
3. **Lab or decision exercise.** In computational chapters, the complete route in R from inputs to result; in legal and communication chapters, reading of primary sources, a role assignment and a reasoned decision.
4. **Counter-check.** A parameter, an assumption or the target population is changed. Predict the direction first, then compare with the result.
5. **Consequences for patients and operations.** Ethics, regulation, data protection and security, ordered by their relevance to the case.
6. **Look up and transfer.** A reviewer question, a self-test with explanation, annotated primary sources, the link to the course and further reading in this book and its companion volumes.

Data status

All computational examples use synthetic data generated in the chapter with a fixed seed, or explicitly identified published figures. Every number in the text belongs to one of four categories:

Label	Meaning
synthetic	generated in the book; describes no real population
published	taken from the cited primary source

Label	Meaning
assumption	value set for teaching purposes that would need local verification in practice
open	a question the chapter raises but does not answer

No number in this book describes a commercial product, a real hospital or unpublished project results.

Reading paths in the course

The course Clinical Data Science in R uses selected chapters of this book as required reading. Readers working through the book on their own do not need the course.

Course stage	Required reading per workshop	Supplementary
I · Entry	W1: Chapter 1, Chapter 3 · W2: Chapter 7, Chapter 8	Chapter 23, Chapter 25
II · Application	W3: Chapter 6, Chapter 9 · W4: Chapter 10, Chapter 19	Chapter 11, Chapter 12
III · Advanced	W5: Chapter 2, Chapter 13 · W6: Chapter 17 and Chapter 15 or Chapter 18	Chapter 14, Chapter 16, Chapter 27

Course stage	Required reading per workshop	Supplementary
IV · Translation	W7: Chapter 22, Chapter 23 · W8: Chapter 27, Chapter 28	Chapter 20, Chapter 21, Chapter 24, Chapter 25

For a quick look-up, the first question decides the route. “How reliable is the alert for this patient?” leads from Chapter 9 via Chapter 6 to Chapter 21. “May our test set contain images from the same animal?” leads from Chapter 10 via Chapter 19 to Chapter 27. “The support log contains case text” leads through Chapter 23 to Chapter 25 and on to the restart decision in Chapter 20.

Companion books and courses

General statistics is repeated here only as far as a case requires. For the methods themselves, the chapters link to the relevant sections of:

- Methods in R – a method reference for researchers and reviewers, with diagnostics, red flags and reporting templates (including TRIPOD-AI and ARRIVE);
- Scientometrics in R – case studies, labs and reference material on research literature, reproducible pipelines and responsible metrics;
- Strategy in R – simulation with explicit assumptions and decisions under uncertainty;
- CTTIR Courses – a four-part biostatistics curriculum in R whose progression the course for this book follows;
- CTTIR Tutorials – stand-alone tutorials on foundations, regression, time series, Bayesian statistics, machine learning and clinical biostatistics.

Each chapter ends with a “Further reading” section that points to the matching sections.

Programming

The complete teaching path is written in R (R Core Team 2026). Base R is sufficient for the elementary examples; mixed models, differential equations and process simulation use well-established packages such as **lme4**, **deSolve** and **simmer**. Python or Julia appear only where a data pipeline, research weights, solvers or interfaces offer a concrete advantage; a box then describes the original implementation and what must be identical for a fair comparison. Software and legal status of this edition are recorded in Appendix A.

History of this book

Work on the cases, worked examples and chapter structure began in February 2025 and was extended and revised regularly until July 2026. This edition reflects the legal, literature and software status of November 2026.

Citation and licence

> Heller, R. (2026). *Clinical Data Science in R: Methods in Clinical AI and Translational Trauma Research.* https://cttir.github.io/clinical-data-science-in-r/

Text and figures are licensed under CC BY 4.0, code under the MIT licence. Please report corrections through the repository issues.

This book does not replace individual medical, legal or statistical advice.

Part I.

Part I · The question before the model

1. Clinical decisions and the translational chain

R. Heller · Date: 2026-07-25 · Last changed: 2026-11-09

At a glance

Guiding question: Who will do something differently because of the output?
For: medical students, clinical researchers and doctoral candidates who are new to data analysis · **Prerequisites:** basic clinical knowledge; no programming experience
In the course: W1 · From clinical case to question · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

1.1. Case and decision

Ms K., 84 years old, falls in her kitchen early in the morning. She is a fictional patient. In the emergency department a hip radiograph is taken 30 minutes after arrival. Four minutes later, image-analysis software marks a possible femoral neck fracture. The radiologist reports 40 minutes after that, and the orthopaedic team plans surgery for the next morning. After the operation a nurse screens her for delirium, a physiotherapist gets her walking on the second postoperative day, and on day 5 a case manager decides whether to apply for geriatric rehabilitation. Every step produces

information. This chapter asks a narrower question: which of these outputs changes what a named person does, and what happens when the output is wrong or missing?

A second, unrelated case runs alongside. A research group measures tissue oxygenation with hyperspectral imaging (HSI) in six rats after a standardised soft-tissue injury. They want to know whether the signal justifies a study in patients.

1.2. Core reading

Core reading · required for W1 · about 10 minutes · ends before 'Lab'

1.2.1. An output is worth the action it changes

A test result, a risk score or a coloured box on an image is not useful in itself. It becomes useful when a person sees it in time, understands it and does something different from what they would otherwise have done. Fryback and Thornbury described this as a hierarchy for diagnostic imaging (Fryback and Thornbury 1991). Level 1 is technical image quality. Level 2 is diagnostic accuracy, such as sensitivity and specificity. Level 3 asks whether the information changes the clinician's diagnostic thinking. Level 4 asks whether it changes the management plan, level 5 whether patients are better off, and level 6 whether the benefit justifies the costs to society. Each level is necessary for the next, but it does not guarantee it.

For Ms K. this hierarchy is concrete. Suppose the fracture flag is highly accurate (level 2). If the radiologist would have seen the same fracture 40 minutes later and the surgeon would have planned the same operation, nothing changes at level 4. The same flag can be valuable in a different setting: at night, when a junior doctor reads the image alone and might

send a patient home with an undisplaced fracture. Accuracy belongs to the software. Decision value belongs to the software *in a particular pathway*.

1.2.2. The decision map

The simplest tool for this chapter is a decision map (Table 1.1). Each row is one output in Ms K.'s pathway. The columns force four answers: who acts, which action changes, in which population the output is meant to be used, and what happens when the output is wrong or unavailable. The last column, the *fallback*, is often forgotten. It describes how care continues if the software is switched off, gives no answer or gives a wrong one.

Table 1.1.: Decision map for the fictional patient Ms K. (n = 1, *synthetic* case) and for the preclinical HSI study (n = 6 animals, *synthetic*). Each row names one output, the person who acts on it and the route back to usual care.

Output	Who acts differently	Changed action	Population	Harm if wrong	Fallback
Fracture flag on the hip radiograph	Emergency physician, radiologist	Second look; CT if the radiograph is equivocal	Adults with a fall and hip pain in the emergency department	Missed fracture: delayed surgery. False alarm: extra CT, time	Usual radiologist report

Output	Who acts differently	Changed action	Population	Harm if wrong	Fallback
Delirium screening score after surgery	Ward nurse	Prevention bundle (orientation, sleep, glasses, hearing aids)	Older inpatients after hip surgery	Missed delirium; time spent on low-risk patients	Daily clinical assessment
Predicted need for inpatient rehabilitation on day 5	Case manager	Early application for geriatric rehabilitation	Inpatients after hip fracture surgery	Wrong placement; delayed discharge	Team discussion on the ward round
Gait speed in rehabilitation (m/s)	Physiotherapist	Adjust therapy; discuss discharge home	Patients in geriatric rehabilitation	Discharge too early or too late	Therapist's clinical judgement
HSI oxygenation index in the rat model	Research team	Plan the next experiment, or stop	Rats in one injury model, not patients	Wasted animals and funding; premature human study	No clinical use; standard assessment continues

Three entries need particular care. The **population** is where the output was developed and checked, not where it might conceivably run. A fracture

flag trained on emergency radiographs of adults says little about children or about follow-up images after surgery. The **action** must be specific. "Better decisions" is not an action; "request CT before discharge" is. The **fallback** must be a route that already works without the tool, so that a failure of the software does not become a failure of care.

The map rests on one assumption that is easy to overlook: **the person in the 'who acts' column sees the output before the decision is taken, and has the time, the competence and the authority to act differently.** If any of these conditions fails, accuracy is irrelevant. The Lab and the Counter-check below test the timing part of this assumption with a few lines of R.

1.2.3. From rat to patient: the translational chain

The HSI study looks similar but has a different addressee. Nobody treats a patient on the basis of six rats. The research team decides whether the signal deserves the next experiment. The output is a reason to formulate a hypothesis, not a reason to change care.

Between the two lies a chain of evidence. A preclinical signal must first be shown to be robust in animals: random allocation, blinded assessment and complete reporting, for which the ARRIVE guidelines give a checklist (Percie du Sert et al. 2020). Next comes feasibility in patients, then diagnostic accuracy in the target population, which STARD-AI helps to report (Sounderajah et al. 2025). Early clinical evaluation of a decision-support tool, with real users in a real workflow, is the subject of DECIDE-AI (Vasey et al. 2022). Only comparative studies can show that patients are better off. Each link answers a different question, and none can be skipped by a larger sample at an earlier link.

The chain often breaks. Van der Worp and colleagues cite a review of animal studies from leading high-impact journals: about one-third of the

studies translated to the level of human randomised trials, and about one-tenth of the interventions were eventually approved for use in patients (*published*) (Worp et al. 2010). The same authors list typical reasons: weak internal validity (no randomisation, no blinding), limited external validity (young, otherwise healthy animals versus older patients with several diseases) and selective publication of positive results. Ms K. is 84 and takes several medicines. The rats are young, healthy and injured in a standardised way. That gap is not a flaw of the rat study. It is the reason why its result is a hypothesis.

1.2.4. Two teaching datasets that must stay apart

This book uses both cases repeatedly, so their data are kept strictly separate. The **fracture cohort** describes 12 fictional patients (identifiers `P01` to `P12`; Ms K. is `P01`). Its unit of analysis is the patient. The **HSI rat study** describes six fictional animals (`A1` to `A6`). It has thousands of rows, one per image pixel, but only six independent units. Mixing the two, or counting pixels as if they were animals, would produce numbers that look precise and mean nothing. Chapter 3 returns to this.

Question for the workshop: Choose one row of Table 1.1. At what time, and on which screen or sheet of paper, must the output appear so that the named person can still act on it? What would make the row worthless?

1.3. Lab

This Lab is the first contact with R. Every line is explained, and the expected output is printed below each command. You can follow the Lab on paper without running anything. If you want to run it, install R and RStudio first; CTTIR Course 1, Week 1 Session 2 shows the set-up step by step.

1.3.1. Step 1: R as a calculator

Text after `#` is a comment; R ignores it. The arrow `<-` stores a value under a name.

```
xray_min   <- 30        # radiograph taken, minutes after arrival
flag_min   <- 34        # fracture flag available
report_min <- 74        # radiologist's report available
report_min - flag_min   # how many minutes earlier the flag arrives
```

```
#> [1] 40
```

The line starting with `#>` is R's answer: the flag is available 40 minutes before the report. The `[1]` only means "first element of the result".

1.3.2. Step 2: a table of Ms K.'s pathway

A table in R is called a *data frame*. Each column is a variable, each row an observation. Here one row is one step in Ms K.'s pathway.

```
pathway <- data.frame(
  pid  = "P01",                                    # pseudonymous patient identifier
  step = c("Radiograph and fracture flag", "Surgery", "Delirium screening",
           "First walk, gait test", "Rehabilitation decision", "Transfer to rehabilitation",
           "Gait test in rehabilitation", "Follow-up visit"),
  day  = c(0, 1, 2, 3, 5, 8, 16, 92),              # days since arrival (day 0 = arrival)
  role = c("Emergency department", "Surgery", "Ward nursing", "Physiotherapy",
           "Case management", "Case management", "Physiotherapy", "Outpatient clinic")
)
pathway
```

```
#>   pid                          step day                 role
#> 1 P01 Radiograph and fracture flag   0 Emergency department
#> 2 P01                      Surgery   1              Surgery
#> 3 P01          Delirium screening    2         Ward nursing
#> 4 P01       First walk, gait test    3        Physiotherapy
#> 5 P01     Rehabilitation decision    5      Case management
#> 6 P01  Transfer to rehabilitation    8      Case management
#> 7 P01 Gait test in rehabilitation   16        Physiotherapy
#> 8 P01             Follow-up visit   92    Outpatient clinic
```

How to read this output:

- The numbers 1 to 8 on the left are row numbers. R adds them; they are not data.
- `pid` repeats the same value in every row. It is the *identifier* (ID) that links these rows to Ms K. in other tables. It is a pseudonym, not her name. In Chapter 3 the same ID appears in the gait-speed table.
- `day` has a unit: days since arrival in the emergency department. This starting point is called *time zero*. A number without its unit and time zero cannot be checked by anyone else. Chapter 3 counts the same gait tests from the day of surgery (day 1 here), so the first test appears there as day 2, not day 3.
- `role` names who acts. It is the column that the decision map is about.

Before we draw anything, R checks two facts about the table. `stopifnot()` stops with an error if a statement is false, so silence means "all true".

```
stopifnot(nrow(pathway) == 8,                # eight steps
          all(diff(pathway$day) >= 0))       # days never go backwards
```

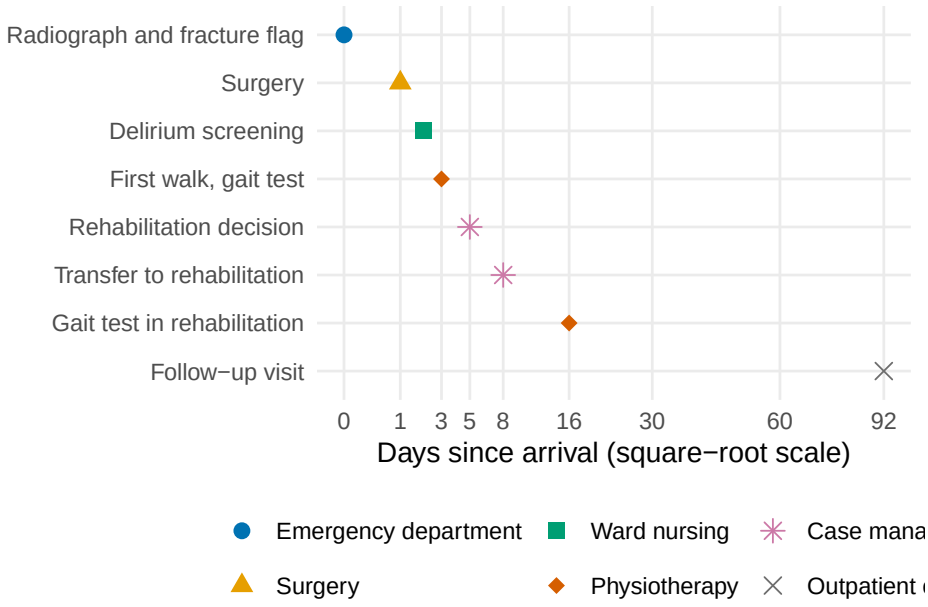
1.3.3. Step 3: the first graphic

The package **ggplot2** builds a graphic from three parts: the data, the *mapping* of columns to positions or colours (`aes()`), and the kind of mark (`geom_point()` for points). `library()` loads the package once per session.

```
library(ggplot2)

pathway$step <- factor(pathway$step, levels = rev(pathway$step))   # keep the order top to b
pathway$role <- factor(pathway$role, levels = unique(pathway$role))

ggplot(pathway, aes(x = day, y = step, colour = role, shape = role)) +
  geom_point(size = 3) +
  scale_x_continuous(transform = "sqrt", breaks = c(0, 1, 3, 5, 8, 16, 30, 60, 92)) +
  scale_colour_manual(values = unname(cdsr_colours[c("blue", "orange", "green", "red", "purp
  scale_shape_manual(values = c(16, 17, 15, 18, 8, 4)) +
  labs(x = "Days since arrival (square-root scale)", y = NULL, colour = NULL, shape = NULL,
       caption = "Synthetic case, n = 1 patient.") +
  theme_cdsr() +
  guides(colour = guide_legend(nrow = 2))
```



Synthetic case, n = 1 patient.

Figure 1.1.: Pathway of the fictional patient Ms K. (`P01`) from the emergency department to the follow-up visit. Horizontal axis: days since arrival (day 0) on a square-root scale; one point per step, coloured by the role that acts. n = 1 patient, 8 steps; *synthetic* case.

What you should see

Eight points, one per row of the table, with the radiograph at the top left (day 0) and the follow-up visit at the bottom right (day 92). The first few days take up most of the width because the axis is drawn on a square-root scale; equal distances on the axis are therefore *not* equal numbers of days. Two pairs of points share a colour and shape:

case management (days 5 and 8) and physiotherapy (days 3 and 16). If your figure shows fewer points, a row was lost; if the steps appear in alphabetical order, the `factor()` line was skipped.

1.3.4. Step 4: the decision map as data

The map in Table 1.1 can be stored as a data frame too. This makes a simple quality check possible: no cell may be empty, because an empty 'fallback' is exactly the gap that causes harm when software fails.

```
decision_map <- data.frame(
  output     = c("Fracture flag", "Delirium score", "Rehabilitation prediction", "Gait speed
  who_acts   = c("Emergency physician, radiologist", "Ward nurse", "Case manager", "Physioth
                 "Research team"),
  action     = c("Second look; CT if equivocal", "Prevention bundle", "Early application",
                 "Adjust therapy", "Plan or stop next experiment"),
  population = c("Adults, fall, hip pain, ED", "Older inpatients after hip surgery",
                 "Inpatients after hip surgery", "Patients in geriatric rehabilitation", "Ra
  fallback   = c("Usual report", "Daily assessment", "Ward-round discussion", "Clinical judg
                 "No clinical use")
)
stopifnot(all(nzchar(as.matrix(decision_map))))   # nzchar = "not an empty string"
table(clinical = decision_map$population != "Rats, one injury model")
```

```
#> clinical
#> FALSE  TRUE
#>     1     4
```

Four outputs concern patients and one concerns animals. The `table()` output counts them.

1.3.5. Step 5: two datasets, two units

The last step builds only the *design* of the two teaching datasets, without measured values, to show the difference between rows and independent units. `expand.grid()` creates every combination of its arguments.

```
cohort_design <- expand.grid(window = c("V1", "V2", "V3", "V4", "V5"),    # fi
                             pid = sprintf("P%02d", 1:12))                 # l
hsi_design <- expand.grid(pixel = 1:200, image = 1:2, limb = c("control", "in
                          animal = sprintf("A%d", 1:6))                    # A

units <- data.frame(
  dataset           = c("Fracture cohort (clinical)", "HSI rat study (preclin
  rows              = c(nrow(cohort_design), nrow(hsi_design)),
  independent_units = c(length(unique(cohort_design$pid)), length(unique(hsi_
  unit              = c("patient", "animal")
)
units
```

```
#>                       dataset rows independent_units    unit
#> 1  Fracture cohort (clinical)   60                12 patient
#> 2 HSI rat study (preclinical) 4800                 6  animal
```

```
stopifnot(units$rows == c(60, 4800), units$independent_units == c(12, 6),
          length(intersect(cohort_design$pid, hsi_design$animal)) == 0)    # n
```

The cohort has 60 rows for 12 patients; the rat study has 4,800 rows for six animals. The last check confirms that no identifier occurs in both datasets, so they cannot be merged by accident.

1.4. Counter-check

The key assumption of the decision map was timing. Suppose the hospital decides to run the fracture software as a nightly batch to save computing costs, so that the flag for Ms K. appears 18 hours after arrival instead of after 34 minutes. **Before you run the code, predict:** how many of the four emergency-department decisions below can the flag still influence?

```
ed_decisions <- data.frame(
  decision = c("Radiologist's report", "CT if radiograph equivocal", "Plan surgery", "Anaest
  hour     = c(74, 80, 120, 150) / 60                # hours since arrival
)
flag_hour_now   <- 34 / 60                           # flag after 34 minutes
flag_hour_batch <- 18                                # nightly batch: 18 hours

ed_decisions$flag_in_time_now   <- flag_hour_now   < ed_decisions$hour
ed_decisions$flag_in_time_batch <- flag_hour_batch < ed_decisions$hour
ed_decisions
```

```
#>                     decision     hour flag_in_time_now flag_in_time_batch
#> 1       Radiologist's report 1.233333             TRUE              FALSE
#> 2 CT if radiograph equivocal 1.333333             TRUE              FALSE
#> 3               Plan surgery 2.000000             TRUE              FALSE
#> 4     Anaesthesia assessment 2.500000             TRUE              FALSE
```

```
stopifnot(sum(ed_decisions$flag_in_time_now) == 4, sum(ed_decisions$flag_in_time_batch) == 0
```

With immediate processing the flag precedes all four decisions; as a nightly batch it precedes none. Its sensitivity and specificity are unchanged, yet its decision value for Ms K. falls to zero. The same logic applies to every row of Table 1.1: a delirium score that appears after the ward round, or a rehabilitation prediction delivered after the application deadline, cannot change the action it was built for.

1.5. Consequences for patients and operations

Patients. A missed fracture means pain, a second visit and delayed surgery; a false alarm means an extra CT, more radiation and more time in the emergency department. Both errors fall on the patient, but they are not equal, and the decision map makes the trade-off visible before any accuracy figure is discussed. Chapter 9 turns these errors into counts per shift, and Chapter 21 asks who carries them.

Operations. Every output needs a named owner and a fallback that works without it. If the fracture software fails at 3 a.m., the radiologist's report remains the basis of care. A tool whose failure goes unnoticed is more dangerous than no tool, because staff may rely on a silence that only means "not processed". Monitoring and stopping rules are the subject of Chapter 20.

Regulation. Whether fracture-flagging software is a medical device depends on its intended purpose, as set by the manufacturer, under the EU Medical Device Regulation (MDR, Regulation (EU) 2017/745) (European Parliament and Council of the European Union 2017a). Annex VIII, Rule 11 of this binding text classifies software that provides information used for diagnostic or therapeutic decisions. In which class a given product falls, and which obligations of the EU Artificial Intelligence Act (Regulation (EU) 2024/1689, amended by Regulation (EU) 2026/1744) apply from which date (European Parliament and Council of the European Union 2024b, 2026), is worked through in Chapter 22 (status: November 2026). A research prototype that is quietly used for patient care changes its regulatory position; that is a question for the institution, not for the analyst alone (*open*, to be clarified locally with the regulatory affairs or medical device officer).

Animal research. The rat study is subject to Directive 2010/63/EU on the protection of animals used for scientific purposes (European Parliament and Council of the European Union 2010), published on 20 October 2010 and in force since 9 November 2010. Its Article 4 obliges Member

States to ensure replacement, reduction and refinement (the 3Rs) wherever possible; the national implementing provisions apply from 1 January 2013. Which national authority approves a given project is a local question (*open*, to be clarified with the institution's animal welfare body). ARRIVE 2.0 governs how the study is reported (Percie du Sert et al. 2020), and the NC3Rs explain the 3Rs for planning (NC3Rs 2026). A within-animal control limb, as in our design, is one way to reduce animal numbers, but it does not replace a justified sample size (Chapter 16).

Data protection. Not directly touched here, because all data are *synthetic*. It becomes relevant as soon as real radiographs or registry data leave the treating hospital (Chapter 23).

1.6. Look up and transfer

- **Reviewer question:** For the output you evaluate, which named role changes which action, at what time after time zero, in which population — and what is the fallback when the output is missing or wrong?
- **Self-test:** A model predicts on day 7 after hip surgery whether a patient will need inpatient rehabilitation. In the hospital, rehabilitation applications must be submitted by day 5. The model's discrimination is good. What is its decision value in this pathway, and what would you change first: the model, its timing or its addressee?

  Answer

  As deployed, the decision value is zero: the output arrives after the decision it is meant to inform. Improving the model's accuracy does not help. The first change is the timing (predict on day 3 or 4 with the information then available, which may lower accuracy) or the addressee (use the prediction for a

later decision, such as the timing of discharge from rehabilitation). Fryback and Thornbury's levels 3 and 4 are not reached, whatever level 2 shows.

- **Annotated sources:**
  - Fryback and Thornbury (1991) — the hierarchy from technical quality to societal value; the basis for separating accuracy from decision value in this chapter.
  - Worp et al. (2010) — reasons why animal findings fail to translate, with the published estimate quoted in the Core reading.
  - Percie du Sert et al. (2020) — ARRIVE 2.0, the reporting checklist for the rat study.
  - Sounderajah et al. (2025) — STARD-AI, reporting of diagnostic accuracy studies of AI tools (level 2).
  - Vasey et al. (2022) — DECIDE-AI, reporting of early clinical evaluation in real workflows (levels 3 and 4).
  - European Parliament and Council of the European Union (2010) — the binding EU basis for the 3Rs (Article 4).

1.7. Further reading

- Chapter 2 — the same pathway seen as three different questions: describe, predict, explain.
- Chapter 9 — what the fracture flag means in counts of alarms per shift at different prevalences.
- Chapter 19 — when a result from one centre or one species may be transferred to the next.
- Methods in R: The research workflow — 1. Question and 9. Knowledge / decisions — the two ends of the chain in this chapter.
- CTTIR Course 1, Week 1 Session 1: Scientific process and research workflow — a lab on turning a question into a plan.

- CTTIR Course 3, Week 1 Session 4: Bench and translational design — design of the preclinical link of the chain.
- Methods in R: ARRIVE — animal research — a reporting template for the rat study.
- CTTIR Course 1, Week 1 Session 5: ggplot2 grammar — the next step after the first graphic.

2. Describe, predict, explain

R. Heller · Date: 2025-03-09 · Last changed: 2026-11-09

At a glance

Guiding question: Which target quantity answers the actual question?
For: clinicians and researchers who plan an analysis of routine or registry data · **Prerequisites:** Chapter 1 (decision map, first steps in R)
In the course: W5 · Causality, trajectories and phenotypes (required core reading); the three questions are first drafted in W1 · From clinical case to question · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 11 minutes

2.1. Case and decision

A geriatric trauma team reviews its results after hip fracture surgery. The in-house standard aims to operate within 24 hours of admission (an *assumption* of this case, not a guideline statement). A colleague presents an analysis of the intensive care database. Among older patients with hip fracture who needed intensive care, those whose surgery was delayed beyond 24 hours were *less* often unable to walk independently at day 90. Someone suggests that the push for early surgery could be relaxed for complex patients. The head of department has to decide whether to change the standard. Before looking at any model, the team needs to know which

question the analysis answered, and whether it is the question the decision requires.

2.2. Core reading

Core reading · required for W5 · about 11 minutes · ends before 'Lab'

2.2.1. Three questions about the same variables

The same variables can serve three different purposes. Hernán and colleagues call them description, prediction and counterfactual prediction, the last being the task of causal inference (Hernán et al. 2019). Shmueli showed that a model built to explain is judged by different criteria from a model built to predict, even if the software call looks the same (Shmueli 2010). Table 2.1 applies this to the case.

Table 2.1.: Three questions about the same variables in the hip fracture case. The data and the software can be identical; time zero, target quantity and the role of each variable differ.

	Describe	Predict	Explain (causal)
Example	What proportion of patients aged 70 or over with hip fracture could not walk independently at day 90, by timing of surgery?	For a patient like Ms K., what is the probability at admission of not walking independently at day 90?	How would the proportion not walking at day 90 change if all patients were operated within 24 hours rather than later?
Time zero	Admission	Moment of prediction (here: admission)	Moment of the decision on timing (admission)
Target quantity	Proportions with their denominators	An individual risk; judged by accuracy and calibration	A risk difference between two hypothetical strategies
'Surgery delayed'	Grouping variable	Predictor, if known at the moment of prediction	The intervention
'ICU admission'	Defines a subgroup and changes the population	Not yet known at admission; a valid predictor only for a later time zero	A consequence of the intervention; must not be selected or adjusted for

A descriptive answer is correct if the counts are right and the denominators are named. A prognostic answer is useful if its predictions are accurate in the population where they will be used. A causal answer requires something more: an assumption about how the variables influence each other.

2.2.2. The target quantity has to be written down

For the causal question, the target quantity, the *estimand*, has to be specified before the analysis. The ICH E9(R1) addendum lists its attributes: the treatment conditions to be compared, the population, the variable measured in each patient, the handling of *intercurrent events* (events after the start of treatment, such as death or transfer, that affect the measurement or its meaning) and the population-level summary (International Council for Harmonisation 2019). For the case: older patients admitted with hip fracture; surgery within 24 hours versus later; not walking independently at day 90; death before day 90 counted as not walking (a composite strategy, one of several possible choices); the difference in proportions. Hernán and Robins recommend describing the randomised trial one would ideally run, the *target trial*, and then aligning eligibility, assignment and time zero in the observational data with it (Hernán and Robins 2016). In such a trial, nobody would be enrolled only after reaching the intensive care unit.

2.2.3. A small causal diagram

Figure 2.1 draws the assumptions as a directed acyclic graph (DAG). Each arrow means "may directly influence"; a missing arrow is a strong claim that there is no direct influence.

Three roles matter. **Frailty** is a *confounder*: it makes delay more likely (frail patients often need medical optimisation first) and independently

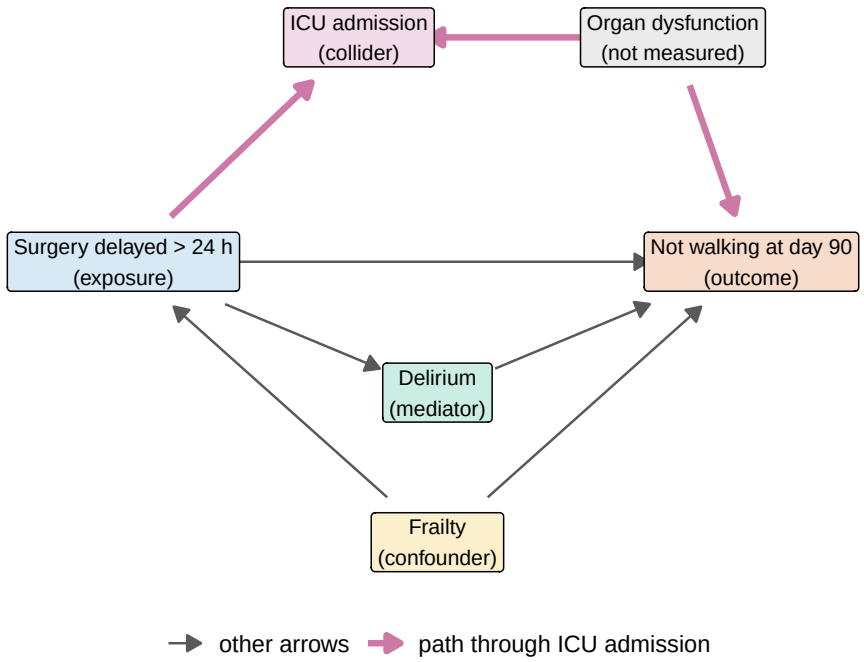


Figure 2.1.: Causal assumptions for the hip fracture case (directed acyclic graph). Each arrow is an *assumption* of this teaching example, not a finding. The role of each variable for the effect of delayed surgery on walking is given in brackets. Thick purple arrows: the path through ICU admission that opens when only ICU patients are analysed. Organ dysfunction is not recorded in the database. No sample size: the graph describes assumed structure, not data.

worsens walking. Comparing delayed with early surgery without accounting for frailty mixes the two. **Delirium** is a *mediator*: it lies on a path from delay to walking. For the total effect of delay, one does not adjust for it, because that would remove part of the effect we want to estimate. **ICU admission** is a *collider* on the path delay → ICU ← organ dysfunction → walking: two arrows collide in it. That path is closed as long as we leave ICU admission alone. Selecting only ICU patients, or adjusting for ICU admission, opens it (Hernán et al. 2004). The R package **dagitty** reads such a graph and lists the variables that are sufficient for adjustment (Textor et al. 2016); here the answer is frailty alone.

2.2.4. Why the sign flips

In the Lab we simulate 4,000 patients whose data follow exactly this graph. The true effect of delay is built in: 7.5 percentage points more patients not walking at day 90 (*assumption*). In the whole cohort, after adjusting for frailty, the estimate is +7.2 points. Among the 767 ICU patients alone it is −9.9 points: delay seems protective (both *synthetic*).

The explanation lies in why patients reach the ICU. Patients operated early usually got there because of organ dysfunction; in the simulation this applies to 81% of them. Patients with delayed surgery often got there because of complications of the waiting time, and only 40% had organ dysfunction. Organ dysfunction strongly impairs recovery. Within the ICU, the early-surgery group is therefore sicker, and delay looks favourable. The same mechanism has been described for studies that sampled only hospitalised or tested patients during the COVID-19 pandemic (Griffith et al. 2020).

2.2.5. What would answer the head of department

The ICU analysis is not useless; it correctly describes patients who reached intensive care. It cannot answer the question on the table, because the

decision about timing is taken at admission and applies to all admitted patients. An analysis that could inform the standard starts at admission (time zero), includes every older patient admitted with hip fracture, compares surgery within 24 hours with later surgery, adjusts for the confounders named in the graph (here frailty), treats death before day 90 according to the estimand, and does not select or adjust on anything that happens after the timing decision, including ICU admission and delirium. If organ dysfunction at admission were recorded, it would be measured before the decision and could be used as well; recorded after the decision, it would be a consequence. The whole difference lies in the order of events, which is why time zero comes first.

The DAG is an assumption, not a result. The data alone cannot show whether delay causes ICU admission; clinical knowledge has to justify that arrow, and with it the decision not to analyse ICU patients separately.

Question for the workshop: Which of the three questions does the head of department need answered? Write down its time zero, its population and the one variable you would refuse to select on — and say why.

2.3. Lab

2.3.1. Step 1: the graph as an R object

The DAG in Figure 2.1 was defined with the following code (the drawing part is ordinary **ggplot2**).

```
library(ggplot2)
library(dagitty)

dag_b02 <- dagitty('dag {
  delay    [exposure, pos="0,0"]
```

```
  walk     [outcome,  pos="4,0"]
  frailty  [pos="2,1.5"]
  delirium [pos="2,0.7"]
  icu      [pos="1.5,-1.2"]
  organ    [latent, pos="3.5,-1.2"]
  frailty -> delay
  frailty -> walk
  delay -> walk
  delay -> delirium -> walk
  delay -> icu
  organ -> icu
  organ -> walk
}')

# node positions (dagitty counts y downwards, so the sign is flipped for plot
pos   <- coordinates(dag_b02)
nodes <- data.frame(name = names(pos$x), x = pos$x, y = -pos$y)
labels <- c(delay = "Surgery delayed > 24 h\n(exposure)", walk = "Not walking
            frailty = "Frailty\n(confounder)", delirium = "Delirium\n(mediato
            icu = "ICU admission\n(collider)", organ = "Organ dysfunction\n(m
nodes$label <- labels[nodes$name]

# arrows run from box edge to box edge; box half-width grows with the longest
longest  <- sapply(strsplit(nodes$label, "\n"), function(l) max(nchar(l)))
nodes$hw <- 0.026 * longest + 0.06          # half-width in plot units
hh <- 0.21                                  # half-height in plot units
e <- as.data.frame(edges(dag_b02))[, c("v", "w")]
from <- nodes[match(e$v, nodes$name), ]
to   <- nodes[match(e$w, nodes$name), ]
dx <- to$x - from$x; dy <- to$y - from$y; len <- sqrt(dx^2 + dy^2)
ux <- dx / len; uy <- dy / len
gap_from <- pmin(from$hw / abs(ux), hh / abs(uy)) + 0.05
gap_to   <- pmin(to$hw / abs(ux), hh / abs(uy)) + 0.05
```

```
arrows_df <- data.frame(x = from$x + gap_from * ux, y = from$y + gap_from * uy,
                        xend = to$x - gap_to * ux, yend = to$y - gap_to * uy,
                        path = ifelse(e$v == "organ" | e$w == "icu",
                                      "path through ICU admission", "other arrows"))

ggplot() +
  geom_segment(data = arrows_df,
               aes(x = x, y = y, xend = xend, yend = yend, colour = path, linewidth = path),
               arrow = arrow(length = unit(2.4, "mm"), type = "closed")) +
  geom_label(data = nodes, aes(x = x, y = y, label = label), size = 3.2, linewidth = 0.3,
             fill = c(delay = "#D6E9F5", walk = "#F8DCCB", frailty = "#FBEFCC", delirium = "
                      icu = "#F2DCE9", organ = "grey92")[nodes$name]) +
  scale_colour_manual(values = c("other arrows" = "grey35",
                                 "path through ICU admission" = cdsr_colours[["purple"]]), n
  scale_linewidth_manual(values = c("other arrows" = 0.5, "path through ICU admission" = 1.3
  coord_cartesian(xlim = c(-0.7, 4.7), ylim = c(-1.7, 1.5)) +
  theme_void(base_size = 12) + theme(legend.position = "bottom")
```

dagitty can now answer questions about the graph. `adjustmentSets()` lists the minimal sets of variables that block every non-causal path between exposure and outcome; `isAdjustmentSet()` tests a proposed set.

```
adjustmentSets(dag_b02, effect = "total")
```

```
#> { frailty }
```

```
checks <- c(
  frailty              = isAdjustmentSet(dag_b02, "frailty"),
  frailty_and_icu      = isAdjustmentSet(dag_b02, c("frailty", "icu")),
  frailty_and_delirium = isAdjustmentSet(dag_b02, c("frailty", "delirium"))
)
checks
```

```
#>              frailty      frailty_and_icu frailty_and_delirium
#>                 TRUE                FALSE                FALSE
```

```
stopifnot(identical(unname(checks), c(TRUE, FALSE, FALSE)))
```

Adding ICU admission breaks the adjustment because it opens the collider path. Adding delirium breaks it for the *total* effect because it blocks part of the effect itself. `paths()` shows which paths are open after selecting on ICU admission:

```
paths(dag_b02, Z = c("frailty", "icu"))
```

```
#> $paths
#> [1] "delay -> delirium -> walk"     "delay -> icu <- organ -> walk"
#> [3] "delay -> walk"                 "delay <- frailty -> walk"
#>
#> $open
#> [1]  TRUE  TRUE  TRUE FALSE
```

The path `delay -> icu <- organ -> walk` is listed as open (`TRUE`). The back-door path through frailty is closed.

2.3.2. Step 2: data that follow the graph

The function below simulates patients exactly as the graph says. Every number in it is a teaching *assumption*. Each line draws one variable with `rbinom()`, which returns 1 (yes) with the stated probability and 0 (no) otherwise. The seed inside the function makes the result identical on every computer.

```r
simulate_cohort <- function(n = 4000, delay_to_icu = 0.25, seed = 2402) {
  set.seed(seed)
  frailty  <- rbinom(n, 1, 0.40)                                    # 40% frail
  organ    <- rbinom(n, 1, 0.20)                                    # 20% acute organ dysfunc
  delay    <- rbinom(n, 1, 0.25 + 0.20 * frailty)                   # frailty -> delay
  delirium <- rbinom(n, 1, 0.15 + 0.15 * delay)                     # delay -> delirium
  icu      <- rbinom(n, 1, 0.03 + delay_to_icu * delay + 0.45 * organ)   # delay, organ -> I
  not_walk <- rbinom(n, 1, 0.10 + 0.06 * delay + 0.10 * delirium + 0.15 * frailty + 0.35 * o
  data.frame(frailty, organ, delay, delirium, icu, not_walk)
}
cohort <- simulate_cohort()
true_effect <- 0.06 + 0.10 * 0.15          # direct part + part through delirium
c(patients = nrow(cohort), icu = sum(cohort$icu), true_effect = true_effect)
```

```
#>    patients         icu true_effect
#>    4000.000     767.000       0.075
```

The true total effect is 0.06 (direct) plus 0.10×0.15 (the extra delirium caused by delay, times the harm of delirium) $= 0.075$, i.e. 7.5 percentage points.

2.3.3. Step 3: three analyses, three answers

A linear model with a 0/1 outcome, `lm(not_walk ~ delay + frailty)`, returns the risk difference for delay directly as its coefficient. This is enough for a simulation in which all effects add up; it is not a general recommendation for binary outcomes (Methods in R: Multiple Linear Regression explains the model).

```
risk_diff <- function(formula, data) coef(lm(formula, data = data))[["delay"]
estimates <- c(
  describe_all          = risk_diff(not_walk ~ delay, cohort),
  adjust_frailty        = risk_diff(not_walk ~ delay + frailty, cohort),
  adjust_frailty_delir  = risk_diff(not_walk ~ delay + frailty + delirium, co
  icu_only              = risk_diff(not_walk ~ delay + frailty, cohort[cohort
)
round(100 * estimates, 1)   # percentage points
```

```
#>         describe_all       adjust_frailty adjust_frailty_delir
#>                 10.4                  7.2                  5.8
#>             icu_only
#>                 -9.9
```

```
stopifnot(round(100 * estimates[["adjust_frailty"]], 1) == 7.2,
          round(100 * estimates[["icu_only"]], 1) == -9.9,
          sum(cohort$icu) == 767)
```

The crude difference (+10.4 points) is a correct *description* of this cohort but overstates the effect, because frail patients are delayed more often. Adjusting for frailty gives +7.2 points, close to the built-in 7.5. Adding delirium gives +5.8 points: roughly the direct part only. Restricting to ICU patients reverses the sign.

2.3.4. Step 4: explaining the sign

```
organ_share <- rbind(
  all_patients = tapply(cohort$organ, cohort$delay, mean),
  icu_only     = tapply(cohort$organ[cohort$icu == 1], cohort$delay[cohort$ic
)
```

```
colnames(organ_share) <- c("early surgery", "delayed surgery")
round(organ_share, 2)
```

```
#>              early surgery delayed surgery
#> all_patients          0.20             0.2
#> icu_only              0.81             0.4
```

```
stopifnot(round(organ_share["icu_only", ], 2) == c(0.81, 0.40))
```

In the whole cohort, organ dysfunction is equally common with and without delay (about 20%), as the graph demands. Among ICU patients, it is twice as common after early surgery. Because organ dysfunction adds 35 percentage points to the risk of not walking, the early-surgery ICU patients do worse, and delay appears protective. The negative sign is not a property of delay. It is a property of the selection.

2.3.5. Step 5: is it chance?

One simulated cohort could be unlucky. Repeating the simulation with 100 different seeds shows how stable the sign is.

```
icu_estimates <- sapply(1:100, function(s) {
  d <- simulate_cohort(seed = s)
  risk_diff(not_walk ~ delay + frailty, d[d$icu == 1, ])
})
share_negative <- mean(icu_estimates < 0)
round(c(median = median(icu_estimates), share_negative = share_negative), 3)
```

```
#>         median share_negative
#>         -0.068          0.960
```

```
stopifnot(share_negative >= 0.90)
```

In 96 of 100 repetitions the ICU-only estimate is negative, although delay harms every simulated patient group. Selection, not chance, drives the result.

2.3.6. Step 6: the prognostic question uses the same data differently

For prediction, the question is not whether a coefficient is causal but whether predicted risks are accurate for new patients. A logistic model with the two variables that the database records at admission, frailty and timing of surgery, gives a risk for a patient like Ms K.: frail and operated early.

```
fit_prog <- glm(not_walk ~ frailty + delay, family = binomial, data = cohort)
ms_k <- data.frame(frailty = 1, delay = 0)
risk_ms_k <- predict(fit_prog, newdata = ms_k, type = "response")
round(risk_ms_k, 3)
```

```
#>     1
#> 0.329
```

```
built_in <- 0.10 + 0.15 + 0.10 * 0.15 + 0.35 * 0.20   # baseline + frailty +
stopifnot(abs(risk_ms_k - built_in) < 0.03)
```

The predicted risk of about 32.9% is close to the average risk built into the simulation for such patients (33.5%). The model does not know about organ dysfunction, yet it can still be a reasonable *average* prediction. Whether it is good enough for a decision is a question of calibration and clinical utility (Chapter 6), not of causal structure.

> **Common mistake**
>
> Reporting every coefficient of one adjusted model as if each were a causal effect (sometimes called the "Table 2 fallacy"). In Step 3, the model `not_walk ~ delay + frailty` is built for the effect of delay. Its frailty coefficient does not estimate the effect of frailty, because part of that effect runs through delay. Each causal question needs its own adjustment set (Methods in R: DAGs).

2.4. Counter-check

Change one assumption: delay no longer causes ICU admission (`delay_to_icu = 0`), for example because the ward now has a dedicated orthogeriatric team that manages complications without intensive care. ICU admission then depends only on organ dysfunction. **Predict first:** what happens to the ICU-only estimate — does it stay negative, move towards +7.5, or overshoot?

```
cohort_no_arrow <- simulate_cohort(delay_to_icu = 0)
counter <- c(
  icu_only_with_arrow    = estimates[["icu_only"]],
  icu_only_without_arrow = risk_diff(not_walk ~ delay + frailty, cohort_no_arrow[cohort_no_a
)
round(100 * counter, 1)
```

```
#>    icu_only_with_arrow icu_only_without_arrow
#>                   -9.9                    6.2
```

```
stopifnot(counter[["icu_only_without_arrow"]] > 0)
```

Without the arrow from delay to ICU, ICU admission is no longer a collider between delay and organ dysfunction, and the ICU-only estimate is positive again (6.2 points, with more random noise because only about 760 patients remain). The data could not have told us which of the two worlds we are in; only the clinical argument about the arrow can.

2.5. Consequences for patients and operations

Patients. Read causally, the ICU analysis would suggest that delaying surgery is harmless or even beneficial for complex patients. In the simulated world the opposite is true, and frail patients would bear the harm of a relaxed standard. An analysis restricted by its data source (here: the ICU database) has selected its population before any model was fitted. That restriction must be reported as prominently as the result.

Operations. A department that changes a standard on the basis of routine data should ask for three things before the meeting: the written estimand, the DAG with a justification for each arrow, and the inclusion criteria of the data source. Prediction models may use post-admission information if their time zero is later; such a model should state its time zero in its name and documentation.

Ethics and regulation. Presenting a prognostic association as a treatment recommendation is an error with ethical weight, because it shifts risk to patients without evidence. If a prognostic model were used to steer individual treatment decisions, its intended purpose could make it a medical device under the MDR (European Parliament and Council of the European Union 2017a); this is examined in Chapter 22. Data protection is touched only indirectly; secondary use of routine data for such analyses is the subject of Chapter 23 and Chapter 26.

2.6. Look up and transfer

- **Reviewer question:** Which variables were conditioned on by the design or the data source (inclusion criteria, registry, database) rather than in the model — and could any of them be a consequence of the exposure?

- **Self-test:** A geriatric rehabilitation clinic analyses its own patients after hip fracture. Patients with dementia regain walking faster than patients without dementia. The clinic admits patients only if they had good mobility before the fracture *or* good cognitive function. Why might dementia look favourable, and what is the collider?

  Answer

  Admission to the clinic is a collider: both pre-fracture mobility and cognitive function lead to admission. Among admitted patients, those with dementia must have had very good pre-fracture mobility, otherwise they would not have been admitted. Good pre-fracture mobility predicts faster recovery, so dementia appears favourable. The sign is produced by the admission rule, not by dementia. A population-based cohort starting at the fracture (time zero) avoids this selection.

- **Annotated sources:**

  - Hernán et al. (2019) — the three tasks description, prediction and counterfactual prediction.
  - Shmueli (2010) — why explanatory and predictive modelling are judged differently.
  - International Council for Harmonisation (2019) — the attributes of an estimand and the strategies for intercurrent events such as death.

  - Hernán and Robins (2016) — the target trial as a template for time zero and eligibility.
  - Hernán et al. (2004) — selection bias as conditioning on a common effect; the structure behind the ICU example.
  - Textor et al. (2016) — the R package used to read the DAG and find adjustment sets.
  - Griffith et al. (2020) — a *published* example of collider bias from samples of hospitalised or tested patients.

2.7. Further reading

- Chapter 3 — how data arise, and why missing values can follow the same selection logic.
- Chapter 5 and Chapter 6 — building and evaluating the prognostic model of Step 6.
- Chapter 13 — repeated measurements, where time zero and informative drop-out matter again.
- CTTIR Course 3, Week 3 Session 3: DAGs with dagitty and ggdag — a full lab on drawing and testing DAGs.
- CTTIR Course 3, Week 3 Session 1: Time-varying covariates and immortal-time bias — what goes wrong when time zero is misplaced.
- CTTIR Course 2, Week 4 Session 5: Explanation vs prediction; reporting — the distinction of this chapter in a regression lab.
- Methods in R: Explanation vs prediction (Shmueli) — a one-page summary.
- Methods in R: Propensity scores / IPTW — how adjustment is done once the adjustment set is known.

3. How data arise: missingness and denominators

R. Heller · Date: 2025-04-21 · Last changed: 2026-11-09

At a glance

Guiding question: Which patient, animal or sample is missing from the table?
For: students and clinical researchers who collect data in routine care or experiments, or receive them as an export · **Prerequisites:** Chapter 1 (the two teaching datasets and the first R steps)
In the course: W1 · From clinical case to question (required core reading) and W5 · Causality, trajectories and phenotypes · **Data:** synthetic (generated in the chapter; download b03-visits.csv, b03-patients.csv, b03-dictionary.csv) · **Core reading time:** about 11 minutes

3.1. Case and decision

The lead physiotherapist of a fictional orthopaedic department prepares a slide for the hospital board. It shows that the mean gait speed of patients after hip fracture surgery rose from 0.33 m/s in the first days to 0.75 m/s three months later. The board is to decide whether the new rehabilitation pathway works and whether to fund two additional therapist posts. The

slide is based on an export from the measuring devices. Before it is presented, the department's data manager asks a simple question: who is in each of these means, and who is not? Of the 12 patients in the cohort, 7 contribute to the three-month value. The others died, were transferred or missed the visit. The decision is how to present the rise: as an effect of rehabilitation, or as a description of the patients who could still be measured.

3.2. Core reading

Core reading · required for W1 · about 11 minutes · ends before 'Lab'

3.2.1. Every row has a history

Clinical data are a by-product of care. A gait-speed value exists because a test was scheduled, the patient could stand and walk four metres, a therapist had time, the device worked and the export ran. Every missing value also has a reason, and most of these reasons are clinical events: pain, delirium, a transfer, a discharge home, death. An analysis that sees only the exported values sees the end of this process, not the process itself.

The first tool of this chapter is therefore a **planned-visit table**: one row for every planned visit of every patient, whether or not a measurement was taken, with a *status* that says what happened. For the 12 patients of the fracture cohort (Ms K. from Chapter 1 is patient `P01`) and five planned visit windows, this table has 60 rows. The device export has only 42. Table 3.1 summarises the 60 rows.

Table 3.1.: Status of all 60 planned visits in the fracture cohort (n = 12 patients, *synthetic*). Only the first row contributes to the mean in the last row. Gait speed: 4-metre walk at usual pace, metres per second.

Visit window (days since surgery)	V1 (1–3)	V2 (4–9)	V3 (10–20)	V4 (21–35)	V5 (80–100)
Measured	9	11	8	7	7
Unable to walk 4 m	3	1	0	1	0
Device fault	0	0	1	0	0
Visit missed	0	0	0	0	1
At home, no visit planned	0	0	2	2	0
Transferred	0	0	0	1	2
Died	0	0	1	1	2
Total patients	**12**	**12**	**12**	**12**	**12**
Mean gait speed of those measured (m/s)	0.33	0.40	0.45	0.56	0.75

3.2.2. Different reasons, different consequences

The reasons in Table 3.1 are not interchangeable. Rubin's classification describes how a missing value relates to the data (Rubin 1976):

- **Missing completely at random (MCAR).** The sensor mat failed on day 14 for patient `P12`. The fault has nothing to do with her walking. Such gaps cost precision but do not distort the mean.
- **Missing at random (MAR).** Missingness depends only on information that is already recorded. Patient `P03` walked slowly at day 14 (0.38 m/s, recorded) and was transferred to a nursing home; patients `P04` and `P09` walked fast in the first week (recorded) and went home without inpatient rehabilitation. Because the reason is in the data, methods that use the observed history can, in principle, account for it.
- **Missing not at random (MNAR).** "Unable to walk 4 m" means the test was not done *because* the patient was too slow or unsafe. The missing value depends on itself. No analysis of the observed values alone can repair this without an extra assumption about the values nobody measured.

Two further reasons are not "missing data" in the usual sense. **Death** ends the existence of a gait speed; there is no value to be missing. The question must say what it means: the gait speed of survivors at day 90, or a combined outcome such as "alive and walking at least 0.6 m/s". ICH E9(R1) calls death an *intercurrent event* that must be addressed in the definition of the question itself (International Council for Harmonisation 2019), and Kurland and colleagues show how different analyses answer different questions about data truncated by death (Kurland et al. 2009). **Structural gaps** arise from the pathway: a patient who went home has no visit in the rehabilitation clinic because none was planned.

The consequence for the board slide is uncomfortable. In windows V3 and V4 the two *fastest* patients of the first week are missing (they went home), together with some of the slowest. In V5 mainly the *slowest* are

missing (they died or were transferred). The rise from 0.56 to 0.75 m/s between V4 and V5 is partly a change in who is measured. **Every mean in a longitudinal table is a mean of those measured at that time; without the planned-visit table nobody can tell whether a change over time is a change in patients or a change in who is counted.** The Lab draws this as a figure with the denominator written above every point.

3.2.3. Many rows, few units

The second trap is the opposite of a missing row: too many rows. In the HSI rat study of Chapter 1, each of six animals contributes two limbs, two images per limb and 200 pixels per image: 4,800 rows. Treating the pixels as independent observations is *pseudoreplication*, defined by Hurlbert as testing a treatment effect with replicates that are not statistically independent (Hurlbert 1984). In a survey of one issue of a neuroscience journal, Lazic found pseudoreplication in 12% of papers and suspected it in a further 36% (*published*) (Lazic 2010). In the Lab, the pixel analysis of our synthetic rats gives an interval of −5.1 to −4.1 percentage points for the difference in oxygenation between injured and control limb. Analysed per animal, the interval is −11.3 to +2.1: the data are compatible with no difference at all. The same applies to the fracture cohort: 42 gait-speed values come from 11 patients, not from 42. ARRIVE 2.0 therefore asks every animal study to name its experimental unit (Percie du Sert et al. 2020).

3.2.4. Rare combinations identify people

The third question concerns a teaching dataset leaving the hospital. Individual attributes such as age, sex, injury, accident mechanism, admission date or postal district look harmless. Their combination is not. In a synthetic registry extract of 403 patients built in the Lab, every record is

unique on these six attributes. One record describes a 34-year-old woman admitted on 14 August after a paragliding accident with a scapulothoracic dissociation, a rare injury. The injury alone is unique; a short local news report would be enough to name her. Rocher and colleagues estimated that 99.98% of Americans would be correctly re-identified in any dataset using 15 demographic attributes (*published*) (Rocher et al. 2019). Sweeney's k-anonymity formalises the defence: each released record should be indistinguishable from at least $k - 1$ others on the attributes an outsider could know (Sweeney 2002). For teaching, the remedy is mostly subtraction: keep only the columns the exercise needs, coarsen the rest, and leave out what remains unique.

Question for the workshop: Rewrite the board slide. Under each of the five points, write the denominator and the reason why patients are missing. Which point would you remove or label differently, and why?

3.3. Lab

3.3.1. Step 1: the patient table and the planned-visit table

The data come in two tables, as they would in a real project. `patients` has one row per patient; `visits` has one row per planned visit. The column `pid` links them. `rbind()` stacks the rows of several small tables; `NA` is R's code for "not available".

```
patients <- data.frame(
  pid        = sprintf("P%02d", 1:12),
  age        = c(84, 79, 91, 76, 93, 82, 80, 88, 74, 86, 90, 78),       # yea
  sex        = c("f", "f", "f", "m", "f", "f", "m", "f", "f", "m", "f", "f")
  fracture   = c("neck", "pertroch", "pertroch", "neck", "pertroch", "neck",
                 "neck", "pertroch", "neck", "pertroch", "neck", "neck"),
  end_day    = c(91, 88, 20, 90, 11, 90, 93, 33, 90, 90, 52, 89),       # day
  end_reason = c("follow-up complete", "follow-up complete", "transferred to
```

```
                "follow-up complete", "died", "day-90 visit missed", "follow-up complete",
                "transferred to acute hospital", "follow-up complete", "follow-up complete"
                "died", "follow-up complete")
)

windows <- c("V1", "V2", "V3", "V4", "V5")
visit_rows <- function(pid, day, gait_ms, status) {          # five planned visits of one pa
  data.frame(pid = pid, window = windows, day = day, gait_ms = gait_ms, status = status)
}
m <- "measured"
visits <- rbind(
  visit_rows("P01", c(2, 6, 15, 27, 91), c(.31, .42, .55, .63, .78), rep(m, 5)),
  visit_rows("P02", c(1, 5, 13, 29, 88), c(.28, .35, .47, .58, .70), rep(m, 5)),
  visit_rows("P03", c(2, 6, 14, NA, NA), c(NA, .31, .38, NA, NA),
             c("unable", m, m, "transferred", "transferred")),
  visit_rows("P04", c(3, 8, NA, NA, 90), c(.40, .52, NA, NA, .85),
             c(m, m, "at home", "at home", m)),
  visit_rows("P05", c(1, 4, NA, NA, NA), c(NA, NA, NA, NA, NA),
             c("unable", "unable", "died", "died", "died")),
  visit_rows("P06", c(2, 7, 16, 30, NA), c(.33, .41, .50, .57, NA),
             c(m, m, m, m, "visit missed")),
  visit_rows("P07", c(2, 6, 12, 26, 93), c(.36, .47, .58, .66, .81), rep(m, 5)),
  visit_rows("P08", c(3, 8, 17, 31, NA), c(NA, .30, .29, NA, NA),
             c("unable", m, m, "unable", "transferred")),
  visit_rows("P09", c(2, 5, NA, NA, 90), c(.38, .49, NA, NA, .88),
             c(m, m, "at home", "at home", m)),
  visit_rows("P10", c(2, 9, 18, 28, 90), c(.27, .33, .41, .46, .52), rep(m, 5)),
  visit_rows("P11", c(1, 6, 15, 24, NA), c(.30, .37, .45, .44, NA), c(m, m, m, m, "died")),
  visit_rows("P12", c(3, 7, 14, 28, 89), c(.35, .44, NA, .61, .74),
             c(m, m, "device fault", m, m))
)
visits$method <- ifelse(visits$window %in% c("V3", "V4"), "sensor mat", "stopwatch")
visits$method[is.na(visits$day)] <- NA     # no visit took place, so no method
```

```
head(visits, 7)
```

```
#>   pid window day gait_ms   status     method
#> 1 P01     V1   2    0.31 measured  stopwatch
#> 2 P01     V2   6    0.42 measured  stopwatch
#> 3 P01     V3  15    0.55 measured sensor mat
#> 4 P01     V4  27    0.63 measured sensor mat
#> 5 P01     V5  91    0.78 measured  stopwatch
#> 6 P02     V1   1    0.28 measured  stopwatch
#> 7 P02     V2   5    0.35 measured  stopwatch
```

The first seven rows show patient `P01` completely and the start of `P02`. Each row answers four questions: who (`pid`), which planned visit (`window`), when it actually happened (`day`, days since surgery) and what came out (`gait_ms` and `status`). This layout, one row per observation and one column per variable, is called *long* or *tidy*; Chapter 4 explains why it is the right format for analysis.

3.3.2. Step 2: the data dictionary

A table without a description is a puzzle for the next person. A data dictionary names, for each column, its meaning, type, unit, time reference and the plausible reasons for a missing value (Broman and Woo 2018).

```
dictionary <- data.frame(
  table    = c(rep("patients", 6), rep("visits", 6)),
  variable = c("pid", "age", "sex", "fracture", "end_day", "end_reason",
               "pid", "window", "day", "gait_ms", "status", "method"),
  meaning  = c("Pseudonymous patient identifier (P01-P12); key", "Age at adm
               "Sex as recorded (f, m)", "Fracture type: femoral neck (neck)
               "Last day of follow-up", "Why follow-up ended",
```

```
                "Link to the patients table", "Planned window: V1 1-3, V2 4-9, V3 10-20, V4 2
                "Actual day of the visit", "Gait speed, 4-metre walk at usual pace",
                "measured, unable, device fault, visit missed, at home, transferred, died",
                "stopwatch (ward, outpatient) or sensor mat (rehabilitation)"),
  type     = c("text", "number", "text", "text", "number", "text",
               "text", "text", "number", "number", "text", "text"),
  unit     = c("-", "years", "-", "-", "days since surgery", "-",
               "-", "days since surgery", "days since surgery", "m/s", "-", "-"),
  missing_when = c("never", "never", "never", "never", "never", "never",
                   "never", "never", "no visit took place", "status is not 'measured'",
                   "never", "no visit took place")
)
knitr::kable(dictionary)
```

Table 3.2.: Data dictionary of the fracture cohort (*synthetic*). Time zero for all day counts is the day of surgery (day 0).

table	variable	meaning	type	unit	missing_when
patients	pid	Pseudonymous patient identifier (P01-P12); key	text	-	never
patients	age	Age at admission	number	years	never
patients	sex	Sex as recorded (f, m)	text	-	never
patients	fracture	Fracture type: femoral neck (neck) or pertrochanteric (pertroch)	text	-	never
patients	end_day	Last day of follow-up	number	days since surgery	never
patients	end_reason	Why follow-up ended	text	-	never
visits	pid	Link to the patients table	text	-	never

Table 3.2.: Data dictionary of the fracture cohort (*synthetic*). Time zero for all day counts is the day of surgery (day 0).

table	variable	meaning	type	unit	missing_when
visits	window	Planned window: V1 1-3, V2 4-9, V3 10-20, V4 21-35, V5 80-100	text	days since surgery	never
visits	day	Actual day of the visit	number	days since surgery	no visit took place
visits	gait_ms	Gait speed, 4-metre walk at usual pace	number	m/s	status is not 'measured'
visits	status	measured, unable, device fault, visit missed, at home, transferred, died	text	-	never
visits	method	stopwatch (ward, outpatient) or sensor mat (rehabilitation)	text	-	no visit took place

Two entries deserve a comment. The `method` changes from stopwatch to sensor mat between V2 and V3, because the rehabilitation clinic uses a different device. A jump in the curve at that point may be a measurement artefact; how large the systematic difference between the two methods is remains *open* in this cohort (Chapter 11 treats method comparison). And `status` is never missing: it is the column that makes every gap explainable.

3.3.3. Step 3: counts and denominators

`table()` counts combinations; `tapply()` applies a function, here the mean, within groups. `na.rm = TRUE` tells `mean()` to ignore missing values, which is exactly the step that hides the denominator.

```
status_levels <- c("measured", "unable", "device fault", "visit missed", "at home", "transfe
visits$status <- factor(visits$status, levels = status_levels)
status_table <- table(visits$status, visits$window)
status_table
```

```
#>
#>                V1 V2 V3 V4 V5
#>   measured      9 11  8  7  7
#>   unable        3  1  0  1  0
#>   device fault  0  0  1  0  0
#>   visit missed  0  0  0  0  1
#>   at home       0  0  2  2  0
#>   transferred   0  0  0  1  2
#>   died          0  0  1  1  2
```

```
mean_by_window <- tapply(visits$gait_ms, visits$window, mean, na.rm = TRUE)
round(mean_by_window, 2)
```

```
#>   V1   V2   V3   V4   V5
#> 0.33 0.40 0.45 0.56 0.75
```

```
stopifnot(
  all(colSums(status_table) == 12),                                    # every window holds a
  status_table["measured", ] == c(9, 11, 8, 7, 7),                     # as in the Core readi
  round(mean_by_window, 2) == c(0.33, 0.40, 0.45, 0.56, 0.75)
)
```

3.3.4. Step 4: patients, measurements and linked records, counted twice

Counting looks trivial, and that is why errors survive. Below, each number is computed on two independent paths: path A from the visit table, path

B from other objects or with other functions. The numbers are accepted only if both paths agree.

```
# path A: from the visit table
a <- c(patients      = length(unique(visits$pid)),
       measurements  = sum(visits$status == "measured"),
       with_value    = length(unique(visits$pid[!is.na(visits$gait_ms)])),
       linked        = nrow(merge(visits[!is.na(visits$gait_ms), ], patients
# path B: from the patient table, per-patient sums and a matching test
per_patient <- tapply(!is.na(visits$gait_ms), visits$pid, sum)
b <- c(patients      = nrow(patients),
       measurements  = sum(per_patient),
       with_value    = sum(per_patient > 0),
       linked        = sum(visits$pid[!is.na(visits$gait_ms)] %in% patients$p
rbind(path_A = a, path_B = b)
```

```
#>        patients measurements with_value linked
#> path_A       12           42         11     42
#> path_B       12           42         11     42
```

```
stopifnot(identical(a, b), b[["patients"]] == 12, b[["measurements"]] == 42,
          b[["with_value"]] == 11, b[["linked"]] == 42)
names(per_patient)[per_patient == 0]     # the patient who never appears in t
```

```
#> [1] "P05"
```

Twelve patients, 42 measurements, all 42 linked to a patient, but only 11 patients with any value. Patient `P05` was unable to walk at both early visits and died on day 11. She exists in the patient table and in the planned-visit table, but not in any device export: precisely the person the board slide does not show.

3.3.5. Step 5: first trajectory figure, individual courses

A *spaghetti plot* draws one line per patient. Colour, line type and point shape show how follow-up ended; crosses at the bottom mark visits where the test was attempted but not possible.

```
library(ggplot2)
course_of <- c("follow-up complete" = "observed to day 90", "day-90 visit missed" = "day-90
               "transferred to nursing home" = "transferred", "transferred to acute hospital
               "died" = "died")
measured <- merge(visits[visits$status == "measured", ], patients[, c("pid", "end_reason")],
measured$course <- factor(course_of[measured$end_reason],
                          levels = c("observed to day 90", "day-90 visit missed", "transferr
unable <- visits[visits$status == "unable", ]

ggplot(measured, aes(x = day, y = gait_ms, group = pid, colour = course, linetype = course,
  geom_line() +
  geom_point(size = 2) +
  geom_point(data = unable, aes(x = day, y = 0.05), inherit.aes = FALSE,
             shape = 4, size = 2.5, colour = cdsr_colours[["grey"]]) +
  scale_x_continuous(transform = "sqrt", breaks = c(1, 3, 7, 14, 21, 28, 35, 60, 90)) +
  scale_y_continuous(limits = c(0, 1)) +
  scale_colour_manual(values = unname(cdsr_colours[c("blue", "grey", "orange", "red")])) +
  scale_linetype_manual(values = cdsr_linetypes) +
  scale_shape_manual(values = cdsr_shapes) +
  labs(x = "Days since surgery (square-root scale)", y = "Gait speed (m/s)",
       colour = NULL, linetype = NULL, shape = NULL,
       caption = "Synthetic data. n = 12 patients; 42 measurements in 11 patients; 5 tests r
  theme_cdsr()
```

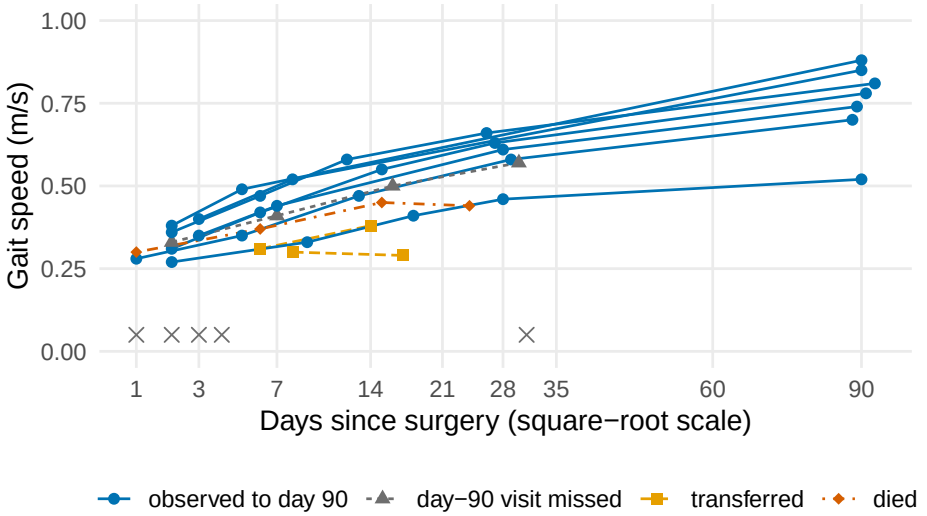


Figure 3.1.: Gait speed (m/s) after hip fracture surgery, one line per patient, by how follow-up ended. Crosses at 0.05 m/s: test attempted but patient unable to walk 4 m. Horizontal axis on a square-root scale. n = 12 patients; 42 measurements from 11 patients; *synthetic*.

3.3.6. Step 6: second trajectory figure, means with their denominators

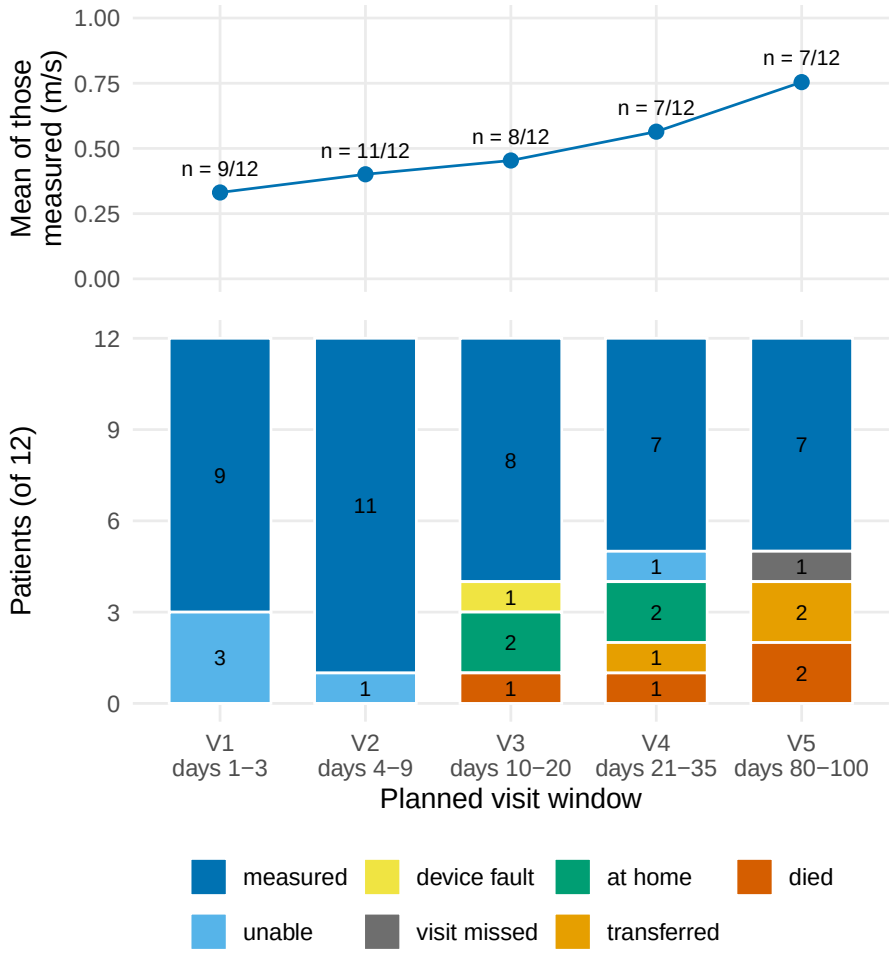
The second figure shows what the board slide showed, the mean per window, but writes the denominator above every point and shows, underneath, what happened to all 12 patients in that window. **ggplot2** draws the two panels; **patchwork** stacks them.

```
library(patchwork)
summary_win <- data.frame(window = names(mean_by_window),
                          mean_ms = as.vector(mean_by_window),
                          n = as.vector(status_table["measured", ]))
window_labels <- c(V1 = "V1\ndays 1-3", V2 = "V2\ndays 4-9", V3 = "V3\ndays 10-20",
                   V4 = "V4\ndays 21-35", V5 = "V5\ndays 80-100")

top <- ggplot(summary_win, aes(x = window, y = mean_ms, group = 1)) +
  geom_line(colour = cdsr_colours[["blue"]]) +
  geom_point(colour = cdsr_colours[["blue"]], size = 2.5) +
  geom_text(aes(label = paste0("n = ", n, "/12")), vjust = -1.1, size = 3.3) +
  scale_y_continuous(limits = c(0, 1)) +
  scale_x_discrete(labels = NULL) +
  labs(x = NULL, y = "Mean of those\nmeasured (m/s)") +
  theme_cdsr()

counts <- as.data.frame(status_table, responseName = "patients")
names(counts)[1:2] <- c("status", "window")
bottom <- ggplot(counts, aes(x = window, y = patients, fill = status)) +
  geom_col(width = 0.7, colour = "white") +
  geom_text(aes(label = ifelse(patients > 0, patients, "")),
            position = position_stack(vjust = 0.5), size = 3.2) +
  scale_fill_manual(values = unname(cdsr_colours[c("blue", "skyblue", "yellow", "grey", "gre
  scale_x_discrete(labels = window_labels) +
  scale_y_continuous(breaks = c(0, 3, 6, 9, 12)) +
  labs(x = "Planned visit window", y = "Patients (of 12)", fill = NULL,
       caption = "Synthetic data. Each bar contains all 12 patients; only the 'measured' seg
  theme_cdsr() +
  guides(fill = guide_legend(nrow = 2))

top / bottom + plot_layout(heights = c(1, 1.4))
```



Synthetic data. Each bar contains all 12 patients; only the 'measured' segment enters tl

Figure 3.2.: Top: mean gait speed (m/s) of the patients measured in each planned visit window, with the observed denominator. Bottom: status of all 12 patients in each window (number of patients written in each segment). n = 12 patients, 60 planned visits; *synthetic*.

The upper panel alone is the board slide. With the lower panel, the reader sees that the five means describe five different groups of patients.

3.3.7. Step 7: saving the tables for the next chapter

The three tables are written as plain-text CSV files into the book's `data` folder, using a path *relative* to this chapter (`..` means "one folder up"). Chapter 4 rebuilds them from raw exports and checks that it gets exactly these files.

```
visits_out <- visits
visits_out$status <- as.character(visits_out$status)    # store plain text, not R factor cod
write.csv(patients,   "../data/b03-patients.csv",   row.names = FALSE)
write.csv(visits_out, "../data/b03-visits.csv",     row.names = FALSE)
write.csv(dictionary, "../data/b03-dictionary.csv", row.names = FALSE)
stopifnot(nrow(read.csv("../data/b03-visits.csv")) == 60)
```

3.3.8. Step 8: pseudoreplication in the HSI rat study

The rat data are generated by a function with its own seed. Oxygen saturation (StO_2) is given in percent; each animal has its own baseline and its own injury effect, and each image adds a small shift.

```
make_hsi <- function(seed = 7) {
  set.seed(seed)
  animal_diff <- rnorm(6, mean = -3, sd = 5)       # injured minus control, percentage point
  animal_base <- rnorm(6, mean = 70, sd = 4)       # StO2 of the control limb, percent (assu
  grid <- expand.grid(pixel = 1:200, image = 1:2, limb = c("control", "injured"),
                      animal = sprintf("A%d", 1:6), stringsAsFactors = FALSE)
  i <- match(grid$animal, sprintf("A%d", 1:6))
  image_shift <- rnorm(24, 0, 1.5)[as.integer(factor(paste(grid$animal, grid$limb, grid$imag
  mu <- animal_base[i] + (grid$limb == "injured") * animal_diff[i] + image_shift
```

```
  grid$sto2 <- round(rnorm(nrow(grid), mu, 6), 1)  # pixel noise
  grid[, c("animal", "limb", "image", "pixel", "sto2")]
}
hsi <- make_hsi()
c(rows = nrow(hsi), animals = length(unique(hsi$animal)))
```

```
#>    rows animals
#>    4800       6
```

```
# wrong unit: every pixel treated as an independent observation
pixel_ci <- t.test(hsi$sto2[hsi$limb == "injured"], hsi$sto2[hsi$limb == "co
# right unit: one difference per animal (injured minus control), six values
animal_means <- tapply(hsi$sto2, list(hsi$animal, hsi$limb), mean)
animal_diff  <- animal_means[, "injured"] - animal_means[, "control"]
animal_ci <- t.test(animal_diff)$conf.int
round(rbind(pixels_as_units = pixel_ci, animals_as_units = animal_ci), 1)
```

```
#>                   [,1] [,2]
#> pixels_as_units   -5.1 -4.1
#> animals_as_units -11.3  2.1
```

```
stopifnot(pixel_ci[2] < 0, animal_ci[1] < 0, animal_ci[2] > 0,
          round(c(pixel_ci, animal_ci), 1) == c(-5.1, -4.1, -11.3, 2.1))
```

Both intervals are centred near the same difference (about -4.6 percentage points), but the pixel interval is about 13 times narrower and excludes zero, while the animal interval includes it. The pixel analysis answers a question nobody asked: how precisely the 4,800 pixel values of *these six rats* are known. The research question concerns rats in general, and its sample size is six. Mixed models that keep the pixel level are introduced in Chapter 13; Chapter 10 and Chapter 19 build on this with spectra and splits by animal.

3.3.9. Step 9: rare combinations in a registry extract

The last step builds a *synthetic* registry extract of 400 common fractures and adds three rare, severe injuries. The function `k_size()` returns, for each record, how many records share its values on the chosen columns: the k of k-anonymity. A record with $k = 1$ is unique.

```
make_registry <- function(seed = 3107, n = 400) {
  set.seed(seed)
  injury <- sample(c("distal radius fracture", "proximal femur fracture", "ankle fracture",
                     "proximal humerus fracture", "vertebral fracture"), n, replace = TRUE,
                   prob = c(0.30, 0.25, 0.20, 0.15, 0.10))
  age_mean <- c("distal radius fracture" = 62, "proximal femur fracture" = 83, "ankle fractu
                "proximal humerus fracture" = 72, "vertebral fracture" = 68)[injury]
  age <- pmin(99, pmax(18, round(rnorm(n, age_mean, 12))))
  sex <- ifelse(runif(n) < ifelse(injury == "ankle fracture", 0.5, 0.7), "f", "m")
  mechanism <- ifelse(age >= 70, "fall from standing height",
                      sample(c("fall from standing height", "road traffic", "sport", "fall f
                             n, replace = TRUE, prob = c(0.40, 0.25, 0.25, 0.10)))
  admission <- as.Date("2025-01-01") + sample(0:364, n, replace = TRUE)
  district <- sample(c("890", "891", "892", "893", "894", "895"), n, replace = TRUE)
  common <- data.frame(rid = sprintf("R%03d", 1:n), age, sex, injury, mechanism, admission,
  rare <- data.frame(rid = c("R401", "R402", "R403"), age = c(34, 41, 57), sex = c("f", "m",
                     injury = c("scapulothoracic dissociation", "open pelvic fracture", "ope
                     mechanism = c("paragliding", "road traffic", "fall from height"),
                     admission = as.Date(c("2025-08-14", "2025-03-02", "2025-11-20")),
                     district = c("893", "890", "895"))
  rbind(common, rare)
}
k_size <- function(data, columns) {
  key <- do.call(paste, c(data[columns], sep = "|"))
  as.vector(table(key)[key])
}
registry <- make_registry()
```

```
all_six <- c("age", "sex", "injury", "mechanism", "admission", "district")
c(records = nrow(registry), unique_on_all_six = sum(k_size(registry, all_six)
```

```
#>           records unique_on_all_six
#>               403               403
```

Every one of the 403 records is unique on the six attributes. The table below shows how many records share their values with record `R401`, the paragliding accident, when only some attributes are released.

```
registry$admission_month <- format(registry$admission, "%m")
registry$age_band  <- cut(registry$age, c(17, 49, 69, 84, 120), labels = c("
registry$injury_group <- ifelse(registry$injury %in% c("scapulothoracic disso
                                "other severe injury", registry$injury)
released <- list("injury"                           = "injury",
                 "sex + mechanism"                  = c("sex", "mechanism"),
                 "age + sex"                        = c("age", "sex"),
                 "sex + district + admission month" = c("sex", "district", "a
                 "age band + sex + injury group"    = c("age_band", "sex", "i
k_r401 <- sapply(released, function(cols) k_size(registry, cols)[registry$ric
k_r401
```

```
#>                           injury                   sex + mechanism
#>                                1                                 1
#>                       age + sex sex + district + admission month
#>                                2                                 4
#>    age band + sex + injury group
#>                                1
```

```
stopifnot(k_r401 == c(1, 1, 2, 4, 1))
```

Even after grouping the rare injuries and banding the age, `R401` remains unique, because she is the only woman under 50 with a severe injury other than a common fracture. The teaching purpose decides what can go. Suppose the exercise is to compare length of stay between injury groups in younger and older patients. Then sex, mechanism, date and district are not needed at all (data minimisation), age needs only two bands, and records that remain rare are left out of the teaching file.

```
registry$age_2 <- ifelse(registry$age < 70, "18-69", "70+")
k_teaching <- k_size(registry, c("age_2", "injury_group"))
teaching <- registry[k_teaching >= 5, c("rid", "age_2", "injury_group")]
c(records_kept = nrow(teaching), records_left_out = sum(k_teaching < 5),
  smallest_k_after = min(k_size(teaching, c("age_2", "injury_group"))))
```

```
#>     records_kept records_left_out smallest_k_after
#>              400                3                5
```

```
stopifnot(sum(k_teaching < 5) == 3, min(k_size(teaching, c("age_2", "injury_group"))) >= 5,
          !"R401" %in% teaching$rid)
```

The reduced teaching file keeps 400 of 403 records on two columns, and every record shares its values with at least four others. The three severe rare injuries are left out; the file says so in its documentation, because their absence is itself a missing-data mechanism for anyone studying severe trauma with it. Whether such a reduced extract would be anonymous in the legal sense is a separate question (see below).

3.4. Counter-check

Because the cohort is synthetic, we know the gait speeds that could not be measured: the values the simulation assumed for the five “unable” visits.

Real data never offer this. **Predict first:** if these patients had been measured, would the means in V1, V2 and V4 rise or fall, and in which window would the change be largest?

```
hidden <- data.frame(pid    = c("P03", "P05", "P05", "P08", "P08"),
                     window = c("V1", "V1", "V2", "V1", "V4"),
                     gait_ms_hidden = c(0.15, 0.12, 0.16, 0.17, 0.21))   # kn
completed <- merge(visits, hidden, by = c("pid", "window"), all.x = TRUE)
completed$gait_ms_all <- ifelse(is.na(completed$gait_ms), completed$gait_ms_h
mean_if_tested <- tapply(completed$gait_ms_all, completed$window, mean, na.rm
round(rbind(observed = mean_by_window, if_unable_were_tested = mean_if_tested
```

```
#>                           V1    V2    V3    V4    V5
#> observed                0.331 0.401 0.454 0.564 0.754
#> if_unable_were_tested 0.285 0.381 0.454 0.520 0.754
```

```
stopifnot(all(mean_if_tested[c("V1", "V2", "V4")] < mean_by_window[c("V1", "V
          which.max(mean_by_window - mean_if_tested) == 1)
```

All three means fall, most in V1 (by 0.046 m/s), where three of twelve patients could not be tested. The rise across the whole follow-up becomes steeper, not flatter, because the early values were too optimistic. Nothing comparable is possible for V5: the two patients who died have no gait speed at day 90, so the question itself must be redefined (survivors only, or a combined outcome), not the data completed.

3.5. Consequences for patients and operations

Patients. Means of "those measured" systematically leave out the frailest: those unable to walk, transferred or dead. A board slide that hides them can lead to funding decisions that look successful because the patients

who fared worst have vanished from the graph. Showing denominators is a matter of fairness to those patients, not only of statistics.

Operations. The planned-visit table needs a status column that clinical staff fill in at the time, with a short list of allowed values. Adding the reasons months later from memory turns MNAR into guesswork. A change of measuring device must be recorded in the data, with its date. Reports should show per-window denominators by default.

Data protection. Gait speed is data concerning health, a special category under Article 9(1) of the General Data Protection Regulation (GDPR, Regulation (EU) 2016/679) (European Parliament and Council of the European Union 2016), published on 4 May 2016, in force since 24 May 2016 and applicable since 25 May 2018. Replacing names by identifiers such as `P01` is *pseudonymisation* in the sense of Article 4(5); Recital 26 states that pseudonymised data which could be attributed to a person by additional information remain information on an identifiable person, while anonymous information falls outside the Regulation. Article 5(1)(c) requires data to be limited to what is necessary (data minimisation), and Article 89(1) requires safeguards for research, including that research purposes which can be fulfilled without identifying data subjects shall be fulfilled in that manner (status: November 2026). Whether a reduced extract like the one in Step 9 is anonymous depends on all means reasonably likely to be used for identification, including local news reports and other registries. That assessment is *open* for any real dataset and belongs to the data protection officer of the institution, not to the analyst (Chapter 23, Chapter 26).

Animal research. Reporting 4,800 pixels as the sample size would overstate the evidence from six animals and could justify further experiments on a false basis. ARRIVE 2.0 asks for the experimental unit to be stated for each experiment (item 1b); see also Chapter 16 for small samples and the 3Rs.

3.6. Look up and transfer

- **Reviewer question:** For every mean or proportion reported over time: what is the denominator at that time, who is missing, and for which documented reason?
- **Self-test:** On an intensive care unit, C-reactive protein (CRP) is measured daily in unstable patients and twice a week in stable ones. A study reports the mean of all CRP values per patient-week and finds that "CRP is higher in the first week". What is the problem with the denominator, and how would you analyse it instead?

  Answer

  The number of measurements depends on how ill a patient is: unstable patients contribute more values, so a mean over all measurements is weighted towards the sickest patients and towards the days on which they were sickest. The visit process is informative. Summarise per patient first (for example one mean or the last value per patient and week), report how many patients contribute to each week, and treat patients who died or were discharged as an explicit status, not as missing values.

- **Annotated sources:**
  - Rubin (1976) — the original definitions behind MCAR, MAR and MNAR.
  - International Council for Harmonisation (2019) — death and transfer as intercurrent events that belong in the question.
  - Kurland et al. (2009) — which analysis answers which question when follow-up is truncated by death.
  - Hurlbert (1984) — the definition of pseudoreplication.
  - Lazic (2010) — how often pseudoreplication occurs in laboratory research (*published* figures).

– Sweeney (2002) — k-anonymity as a criterion for released records.
– Rocher et al. (2019) — how few attributes suffice to re-identify people (*published* estimate).
– Broman and Woo (2018) — practical rules for tables and data dictionaries.
– European Parliament and Council of the European Union (2016) — GDPR Articles 4(5), 5(1)(c), 9(1) and 89(1), and Recital 26, as cited above.

3.7. Further reading

- Chapter 4 — the same tables rebuilt as a reproducible R project from raw exports.
- Chapter 13 — longitudinal models for repeated measurements with informative drop-out.
- Chapter 23 — data protection along the whole data flow, including pseudonymisation and the key file.
- CTTIR Course 1, Week 1 Session 3: Data types, tidy data, accuracy and precision — the long format in practice.
- CTTIR Course 1, Week 1 Session 4: Import, joins, and missingness with dplyr — linking tables and checking for lost rows.
- CTTIR Course 3, Week 2 Session 1: MCAR, MAR, MNAR — the three mechanisms simulated in detail.
- Methods in R: Missing Data in RCTs and Multiple Imputation — what to do once the mechanism is understood.
- Methods in R: Repeated-Measures Designs — designs in which the unit is measured several times.

4. Reconstructing a research project in R

R. Heller · Date: 2026-06-28 · Last changed: 2026-11-09

At a glance

Guiding question: Can a colleague regenerate the number in the figure?
For: anyone who hands over, reviews or reuses an analysis; readers taking their first steps in R · **Prerequisites:** Chapter 3 (the fracture cohort tables); the Lab of Chapter 1
In the course: W1 · From clinical case to question (a script that rebuilds the table in a fresh session), W6 · Pharmacometrics, forecasting and simulation (another group checks one number) and W8 · Colloquium and decision; recommended reading · **Data:** synthetic (built in the chapter; compared with b03-visits.csv) · **Core reading time:** about 10 minutes

4.1. Case and decision

Six months after the board meeting in Chapter 3, a fictional rehabilitation physician is asked to include one number in a grant report: mean gait speed 0.56 m/s in weeks three to five after surgery, n = 7 of 12 patients. She asks the physiotherapy team to regenerate it. They find a spreadsheet called `final_v3_corrected.xlsx`, a figure pasted into a slide and

no script. A colleague remembers that the sensor-mat values were “divided by 100 somewhere”, that one patient identifier was fixed by hand and that a duplicate row was deleted. Nobody can say in which order, or whether anything else happened. The physician has to decide whether the number may go into the report. The department decides on a rule for the future: a number leaves the team only if one command can regenerate it from the raw exports.

4.2. Core reading

Core reading · recommended for W1, W6 and W8 · about 10 minutes · ends before ‘Lab’

4.2.1. Reproducible is a modest, checkable promise

A result is *reproducible* when the same data and the same code produce the same number. This is less than *replication*, which asks whether new data lead to the same conclusion. Peng described reproducibility as the minimum standard for judging a computational claim when full replication is not possible (Peng 2011). It is modest because it says nothing about whether the analysis is right. It is valuable because anyone can check it, and because every error that is found can be traced to a line of code.

The spreadsheet in the case fails this standard not because spreadsheets are bad, but because the steps that changed the data were performed by hand and left no trace. Such *hidden interactive steps* are common: editing cells, filtering in a viewer and copying the visible rows, retyping a number from one window into another. Some are not even chosen by the user. Ziemann and colleagues found that about one-fifth of genomics papers with supplementary Excel gene lists contained gene names that the

software had silently converted into dates or numbers (*published*) (Ziemann et al. 2016).

Sandve and colleagues condensed the remedy into ten rules (Sandve et al. 2013). Five of them carry this chapter: keep track of how every result was produced; avoid manual data manipulation; archive the exact versions of the software used; put all custom scripts under version control; and, for analyses with randomness, record the seed.

4.2.2. A project that runs with one command

Wilson and colleagues describe "good enough" practices for researchers without formal training in programming (Wilson et al. 2017). Applied to the fracture cohort, they lead to a small folder structure:

```
hip-gait-project/
   README.md               start instructions, data schema, target value
   ARTEFACTS.csv           provenance, rights and data status of every exported file
   run.R                   the one command:  Rscript run.R
   R/                      code, numbered in the order it runs
      00_visits.R          planned-visit table of chapter B03 (synthetic)
      01_simulate_raw.R    raw exports (in a real project: delivered, read-only)
      02_clean.R           documented cleaning instead of manual edits
      03_check.R           counts on two independent paths
      04_figure.R          figure and target value
   data/raw/               written once, never edited
   data/derived/           always regenerated, never edited by hand
   output/                 figure, session information
   report/report.qmd       Quarto report: text and numbers from the same code
```

Three conventions do most of the work. First, *raw data are read-only*; every change happens in code. Second, every path is *relative* to the project folder, such as `data/raw/sensormat.csv`, never `C:/Users/anna/Desktop/...`,

so the folder runs on any computer. Third, the project states a **target value**, a number that must come out, and stops with an error if it does not. Here it is the number from the case: in visit window V4, 7 of 12 patients measured, mean gait speed 0.56 m/s.

4.2.3. What a colleague needs to rerun it

A colleague who receives the folder should not need to ask anything. The README answers four questions in this order: how to start (one command), what the input files contain (the *data schema*: one line per file with its unit of observation, its key and the unit of every measured column), which number must come out (the target value), and under which terms the files may be used. The data schema is where the case went wrong: the stopwatch export gives seconds for four metres, the sensor mat gives centimetres per second, and only the conversion written in code makes both comparable. The second safeguard is the count on two paths, as in Chapter 3. Patients, measurements and linked records are computed once from the visit table and once from the patient table with different functions; any lost or duplicated row shows up as a disagreement. Both checks are cheap, and both fail loudly instead of producing a plausible figure.

4.2.4. Versions of code, data and software

A project changes over time. **Git**, a version control system, records every change to text files with author, date and a short message, so that the state of the code on the day a figure was made can be restored (Bryan 2018). The R function `sessionInfo()` records the versions of R and of every loaded package; it belongs in every report. **Quarto** turns a text file with embedded R code into a report in which every number in the text is computed, not typed (Quarto, n.d.). Text and figure can then no longer drift apart.

4.2.5. Every exported file carries its passport

Files travel: into slides, e-mails, other projects. Once they leave the project folder, they lose their context. Each exported artefact therefore needs three pieces of information that travel with it: its *provenance* (which code made it from which input), its *rights* (who may use it under which licence) and its *data status* (synthetic, published, real and pseudonymised, aggregate). Table 4.1 shows them for the project in this chapter.

Table 4.1.: Provenance, rights and data status of every exported artefact of the teaching project (n = 12 *synthetic* patients). In a project with real data, the rights column would name the data use agreement and ethics approval, and the raw files would not be released.

Artefact	Produced by	Provenance	Rights	Data status
`data/raw/*.csv` (4 files)	`R/01_simulate_raw.R`	Generated from the planned-visit table of Chapter 3	CC BY 4.0	*synthetic*
`data/derived/visits.csv`	`R/02_clean.R`	From the four raw files; cleaning log: 1 duplicate removed, 1 identifier corrected	CC BY 4.0	*synthetic*
`output/figure_1.png`	`R/04_figure.R`	From `data/derived/visits.csv`	CC BY 4.0	*synthetic*, aggregate

Artefact	Produced by	Provenance	Rights	Data status
`output/sessionInfo.txt`	`run.R`	Software versions of the run	CC BY 4.0	technical metadata
`report/report.html`	Quarto, `report/report.qmd`	Text, figure and numbers from the files above	Text CC BY 4.0; code MIT	*synthetic*, aggregate
`R/*.R`, `run.R`	Authors	—	MIT	code, no data

A number counts as reproduced only if code alone regenerates it from the raw exports in a fresh R session and it matches the stated target value. A number that is retyped or copied counts as unreproduced, however plausible it looks.

Question for the workshop: Give your script and raw files to another group. They start a fresh R session, run one command and report the target value. Which file, package or step did they need that was not in your project?

4.3. Lab

4.3.1. Step 1: R objects in five minutes

R works with *objects* that have a name. The most important are *vectors* (a sequence of values of one type) and *data frames* (several vectors of equal length side by side). Ms K.'s five gait-speed values from Chapter 3 serve as example.

```
gait <- c(0.31, 0.42, 0.55, 0.63, 0.78)    # numeric vector, metres per second
day  <- c(2, 6, 15, 27, 91)                 # days since surgery
length(gait)                                # how many values
```

```
#> [1] 5
```

```
mean(gait)                                  # one number from five
```

```
#> [1] 0.538
```

```
gait[day > 10]                              # values after day 10: logical selection with [
```

```
#> [1] 0.55 0.63 0.78
```

```
class(gait)                                 # the type of the object
```

```
#> [1] "numeric"
```

```
with_gap <- c(0.31, NA, 0.55)               # NA = not available
mean(with_gap)                              # NA is contagious ...
```

```
#> [1] NA
```

```
mean(with_gap, na.rm = TRUE)                # ... unless it is removed on purpose
```

```
#> [1] 0.43
```

The last two lines are the whole of Chapter 3 in miniature: R refuses to average over a missing value until you tell it to, and from then on the denominator is smaller without anyone seeing it.

4.3.2. Step 2: reading a file with a relative path

The planned-visit table was saved by Chapter 3 in the book's `data` folder. This chapter lives in the `chapters` folder, so the relative path goes one level up (`..`) and then into `data`.

```
path_b03 <- file.path("..", "data", "b03-visits.csv")
stopifnot(file.exists(path_b03))
visits_b03 <- read.csv(path_b03, stringsAsFactors = FALSE)
str(visits_b03)
```

```
#> 'data.frame':    60 obs. of  6 variables:
#>  $ pid    : chr  "P01" "P01" "P01" "P01" ...
#>  $ window : chr  "V1" "V2" "V3" "V4" ...
#>  $ day    : int  2 6 15 27 91 1 5 13 29 88 ...
#>  $ gait_ms: num  0.31 0.42 0.55 0.63 0.78 0.28 0.35 0.47 0.58 0.7 ...
#>  $ status : chr  "measured" "measured" "measured" "measured" ...
#>  $ method : chr  "stopwatch" "stopwatch" "sensor mat" "sensor mat" ...
```

`str()` shows the *structure*: 60 rows (observations), 6 columns (variables), the type of each column (`chr` for text, `int` or `num` for numbers) and the first values. Missing values appear as `NA`. If R reports "cannot open file", the working directory is not the folder you think; `getwd()` shows where R is looking.

4.3.3. Step 3: a base graphic

Base R draws graphics without additional packages. `cdsr_par()` sets the book's margins and fonts; `plot()` draws points and lines; `text()` writes labels at given coordinates. The function below is the figure of the project, and it is used again in Step 7.

```
summarise_windows <- function(visits) {
  data.frame(window  = sort(unique(visits$window)),
             n       = as.vector(tapply(visits$status == "measured", visits$window, sum)),
             of      = as.vector(table(visits$window)),
             mean_ms = as.vector(tapply(visits$gait_ms, visits$window, mean, na.rm = TRUE)))
}
draw_windows <- function(s) {
  plot(seq_len(nrow(s)), s$mean_ms, type = "b", pch = 16, ylim = c(0, 1), xlim = c(0.5, nrow
       xlab = "Planned visit window", ylab = "Mean gait speed (m/s)")
  axis(1, at = seq_len(nrow(s)), labels = s$window)
  text(seq_len(nrow(s)), s$mean_ms, paste0("n = ", s$n, "/", s$of), pos = 3, cex = 0.85)
}
```

```
cdsr_par()
draw_windows(summarise_windows(visits_b03))
```

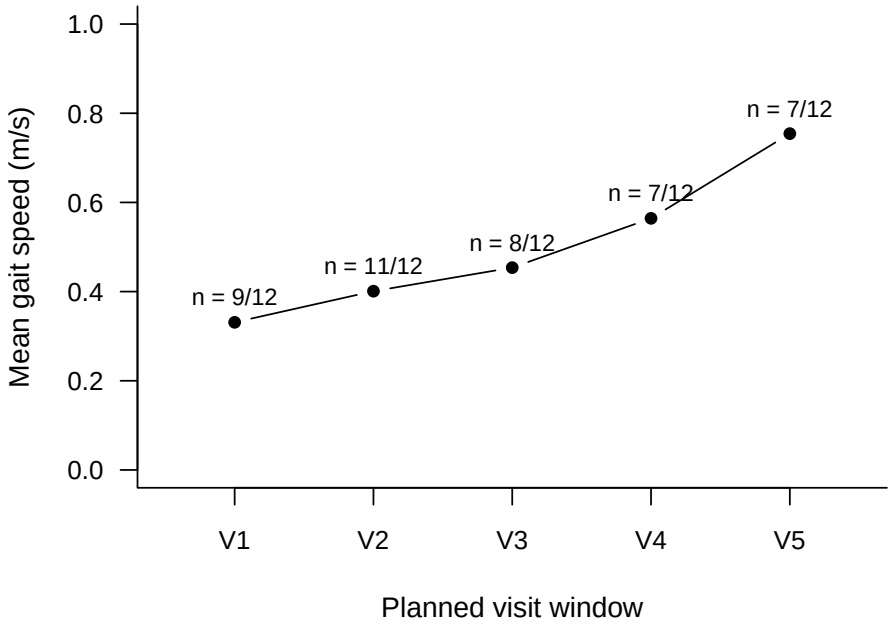


Figure 4.1.: Mean gait speed (m/s) of the patients measured in each planned visit window (V1 days 1–3 to V5 days 80–100 after surgery), with the observed denominator above each point. Base R graphic. n = 12 patients; *synthetic*.

What you should see

Five points joined by lines, rising from left to right; the vertical axis runs from 0 to 1 m/s. The labels above the points read `n = 9/12`, `n = 11/12`, `n = 8/12`, `n = 7/12` and `n = 7/12`. The fourth point, V4, lies just above 0.55 m/s: this is the target value of the project.

4.3.4. Step 4: long and wide

The table has one row per planned visit: the *long* format. Many spreadsheets use the *wide* format instead, with one row per patient and one column per visit. Base R's `reshape()` converts between the two.

```
wide <- reshape(visits_b03[, c("pid", "window", "gait_ms")],
                idvar = "pid", timevar = "window", direction = "wide")
head(wide, 4)
```

```
#>    pid gait_ms.V1 gait_ms.V2 gait_ms.V3 gait_ms.V4 gait_ms.V5
#> 1  P01       0.31       0.42       0.55       0.63       0.78
#> 6  P02       0.28       0.35       0.47       0.58       0.70
#> 11 P03         NA       0.31       0.38         NA         NA
#> 16 P04       0.40       0.52         NA         NA       0.85
```

```
long_again <- reshape(wide, direction = "long")
stopifnot(dim(wide) == c(12, 6), nrow(long_again) == 60,
          sum(!is.na(wide[, -1])) == 42)      # still 42 measured values
```

The wide table looks tidier on screen but has lost two things: the actual day of each visit and the reason for each missing value. In the wide table, `NA` for patient `P04` in V3 (at home) and `NA` for `P05` in V3 (died) look identical. Wickham's definition of *tidy data* — each variable a column, each observation a row, each type of observational unit a table — is the reason why the project keeps the long format and a separate patient table (Wickham 2014).

4.3.5. Step 5: start instructions and data schema

The README is the first file a colleague opens. It needs only four parts.

```
# hip-gait-project (synthetic teaching data)

## Start
1. Install R (version used: see output/session-info.txt) and Quarto.
2. Open a terminal in this folder and run:   Rscript run.R
3. Expected: "Done. Target value V4: 0.56 m/s (n = 7 of 12)". Any other resul

## Data schema (data/raw/, read-only)
patients.csv   one row per patient     pid, age (years), sex, fracture, end_
visit_log.csv  one row per planned visit  pid, window (V1-V5), day (days sin
stopwatch.csv  one row per ward/outpatient test  pid, day, time_4m_s (seconds
sensormat.csv  one row per rehabilitation test   pid, day, speed_cm_s (centi
Key: pid (patients) = pid (all other files); pid + day identifies a test.

## Target value
Visit window V4: 7 of 12 patients measured, mean gait speed 0.564 m/s (rounde

## Rights and data status
Synthetic data, CC BY 4.0. Code under MIT licence. See ARTEFACTS.csv.
```

The schema states the unit of every measurement column, and the units differ: seconds for 4 metres on the ward, centimetres per second on the sensor mat. The former manual division by 100 lived exactly here.

4.3.6. Step 6: the scripts

Each script defines one or two *functions*: named recipes that take an input and return an output. The first one is the planned-visit table of Chapter 3, copied unchanged; Step 7 checks that it really is unchanged.

```
# R/00_visits.R -- planned-visit table of chapter B03 (synthetic); one row =
build_visits <- function() {
```

```
  windows <- c("V1", "V2", "V3", "V4", "V5")
  visit_rows <- function(pid, day, gait_ms, status) {
    data.frame(pid = pid, window = windows, day = day, gait_ms = gait_ms, status = status)
  }
  m <- "measured"
  visits <- rbind(
    visit_rows("P01", c(2, 6, 15, 27, 91), c(.31, .42, .55, .63, .78), rep(m, 5)),
    visit_rows("P02", c(1, 5, 13, 29, 88), c(.28, .35, .47, .58, .70), rep(m, 5)),
    visit_rows("P03", c(2, 6, 14, NA, NA), c(NA, .31, .38, NA, NA),
               c("unable", m, m, "transferred", "transferred")),
    visit_rows("P04", c(3, 8, NA, NA, 90), c(.40, .52, NA, NA, .85),
               c(m, m, "at home", "at home", m)),
    visit_rows("P05", c(1, 4, NA, NA, NA), c(NA, NA, NA, NA, NA),
               c("unable", "unable", "died", "died", "died")),
    visit_rows("P06", c(2, 7, 16, 30, NA), c(.33, .41, .50, .57, NA),
               c(m, m, m, m, "visit missed")),
    visit_rows("P07", c(2, 6, 12, 26, 93), c(.36, .47, .58, .66, .81), rep(m, 5)),
    visit_rows("P08", c(3, 8, 17, 31, NA), c(NA, .30, .29, NA, NA),
               c("unable", m, m, "unable", "transferred")),
    visit_rows("P09", c(2, 5, NA, NA, 90), c(.38, .49, NA, NA, .88),
               c(m, m, "at home", "at home", m)),
    visit_rows("P10", c(2, 9, 18, 28, 90), c(.27, .33, .41, .46, .52), rep(m, 5)),
    visit_rows("P11", c(1, 6, 15, 24, NA), c(.30, .37, .45, .44, NA), c(m, m, m, m, "died")),
    visit_rows("P12", c(3, 7, 14, 28, 89), c(.35, .44, NA, .61, .74),
               c(m, m, "device fault", m, m))
  )
  visits$method <- ifelse(visits$window %in% c("V3", "V4"), "sensor mat", "stopwatch")
  visits$method[is.na(visits$day)] <- NA
  visits
}
```

The second script turns this table into the files the devices and the clinic would deliver: a patient list, a visit log, a stopwatch export in seconds and

a sensor-mat export in centimetres per second, including one duplicated row and one mistyped identifier, as in the case.

```
# R/01_simulate_raw.R -- raw exports (synthetic). In a real project: delivere
simulate_raw <- function() {
  patients <- data.frame(
    pid = sprintf("P%02d", 1:12),
    age = c(84, 79, 91, 76, 93, 82, 80, 88, 74, 86, 90, 78),
    sex = c("f", "f", "f", "m", "f", "f", "m", "f", "f", "m", "f", "f"),
    fracture = c("neck", "pertroch", "pertroch", "neck", "pertroch", "neck",
                 "neck", "pertroch", "neck", "pertroch", "neck", "neck"),
    end_day = c(91, 88, 20, 90, 11, 90, 93, 33, 90, 90, 52, 89),
    end_reason = c("follow-up complete", "follow-up complete", "transferred t
                   "follow-up complete", "died", "day-90 visit missed", "foll
                   "transferred to acute hospital", "follow-up complete", "fo
                   "died", "follow-up complete"))
  visits <- build_visits()
  m <- visits[visits$status == "measured", ]
  ward <- m[m$method == "stopwatch", ]
  stopwatch <- data.frame(pid = ward$pid, day = ward$day,
                          time_4m_s = round(4 / ward$gait_ms, 2))        # se
  rehab <- m[m$method == "sensor mat", ]
  sensormat <- data.frame(pid = rehab$pid, day = rehab$day,
                          speed_cm_s = round(100 * rehab$gait_ms))       # ce
  sensormat <- rbind(sensormat, sensormat[sensormat$pid == "P02" & sensormat$
  sensormat$pid[sensormat$pid == "P07" & sensormat$day == 12] <- "P7"
  list(patients = patients, visit_log = visits[, c("pid", "window", "day", "s
       stopwatch = stopwatch, sensormat = sensormat)
}
```

The third script replaces the hidden spreadsheet step. Every change is written down, and the function stops if the result is implausible or if a value cannot be linked to exactly one planned visit.

```
# R/02_clean.R -- documented cleaning instead of manual edits
id_corrections <- c(P7 = "P07")    # known typing error at the sensor mat, confirmed against

clean_visits <- function(raw) {
  sw <- raw$stopwatch
  sw$gait_ms <- round(4 / sw$time_4m_s, 2)           # 4 metres / seconds = metres per seco
  sm <- raw$sensormat
  n_dup <- sum(duplicated(sm))
  sm <- sm[!duplicated(sm), ]                         # remove exact duplicate exports
  fix <- sm$pid %in% names(id_corrections)
  sm$pid[fix] <- id_corrections[sm$pid[fix]]          # correct known identifier errors
  sm$gait_ms <- sm$speed_cm_s / 100                   # centimetres per second -> metres pe
  values <- rbind(sw[c("pid", "day", "gait_ms")], sm[c("pid", "day", "gait_ms")])
  v <- raw$visit_log
  measured <- v$status == "measured"
  key_visit <- paste(v$pid, v$day)
  key_value <- paste(values$pid, values$day)
  stopifnot(!anyDuplicated(key_value),                # one value per patient and day
            all(key_value %in% key_visit[measured]),  # every value belongs to a measured v
            all(key_visit[measured] %in% key_value))  # every measured visit has a value
  v$gait_ms <- ifelse(measured, values$gait_ms[match(key_visit, key_value)], NA)
  stopifnot(all(v$gait_ms[measured] > 0.05 & v$gait_ms[measured] < 2.5))   # plausible range
  v$method <- ifelse(v$window %in% c("V3", "V4"), "sensor mat", "stopwatch")
  v$method[is.na(v$day)] <- NA
  v <- v[, c("pid", "window", "day", "gait_ms", "status", "method")]
  attr(v, "cleaning_log") <- c(duplicates_removed = n_dup, ids_corrected = sum(fix))
  v
}
```

The fourth script counts patients, measurements and linked records on two independent paths and stops if they disagree. The fifth script is the pair of figure functions from Step 3.

```
# R/03_check.R -- counts on two independent paths
check_counts <- function(visits, patients) {
  a <- c(patients     = length(unique(visits$pid)),                    # pat
         measurements = sum(visits$status == "measured"),
         linked       = sum(visits$pid[!is.na(visits$gait_ms)] %in% patients$
  per_patient <- tapply(!is.na(visits$gait_ms), visits$pid, sum)       # patl
  b <- c(patients     = nrow(patients),
         measurements = sum(per_patient),
         linked       = nrow(merge(visits[!is.na(visits$gait_ms), ], patients
  stopifnot(identical(a, b), b[["patients"]] == 12, b[["measurements"]] == 42
  b
}
```

4.3.7. Step 7: the one command, run inside this chapter

`run.R` calls the scripts in order and writes the files. Because a book chapter must not write project folders, the same sequence runs here in memory; the listing shows the file version.

```
# run.R -- one command, run from the project folder:  Rscript run.R
for (f in sort(list.files("R", pattern = "[.]R$", full.names = TRUE))) source
dir.create(file.path("data", "raw"), recursive = TRUE, showWarnings = FALSE)
dir.create(file.path("data", "derived"), showWarnings = FALSE)
dir.create("output", showWarnings = FALSE)

raw <- simulate_raw()
for (nm in names(raw)) write.csv(raw[[nm]], file.path("data", "raw", paste0(n
raw_in <- lapply(setNames(names(raw), names(raw)), function(nm)
  read.csv(file.path("data", "raw", paste0(nm, ".csv")), stringsAsFactors = F

visits <- clean_visits(raw_in)
write.csv(visits, file.path("data", "derived", "visits.csv"), row.names = FA
```

```
counts <- check_counts(visits, raw_in$patients)
s <- summarise_windows(visits)
stopifnot(s$n[s$window == "V4"] == 7, abs(s$mean_ms[s$window == "V4"] - 0.564) < 0.001)   #

png(file.path("output", "figure-windows.png"), width = 1400, height = 900, res = 200)
draw_windows(s)
invisible(dev.off())
writeLines(capture.output(sessionInfo()), file.path("output", "session-info.txt"))
stopifnot(system2("quarto", c("render", file.path("report", "report.qmd"))) == 0)
cat("Done. Target value V4:", round(s$mean_ms[s$window == "V4"], 2), "m/s (n = 7 of 12)\n")
```

In memory, the CSV round trip is simulated with a *text connection*, so that the cleaning step sees exactly what it would read from disk.

```
as_csv_text <- function(df) capture.output(write.csv(df, row.names = FALSE))
from_csv_text <- function(lines) read.csv(text = lines, stringsAsFactors = FALSE)

raw <- simulate_raw()
raw_in <- lapply(raw, function(df) from_csv_text(as_csv_text(df)))
c(device_rows = nrow(raw_in$stopwatch) + nrow(raw_in$sensormat),
  unknown_ids = length(setdiff(c(raw_in$stopwatch$pid, raw_in$sensormat$pid), raw_in$patient
```

```
#> device_rows unknown_ids
#>          43           1
```

```
setdiff(raw_in$sensormat$pid, raw_in$patients$pid)    # which identifier cannot be linked
```

```
#> [1] "P7"
```

```
visits_new <- clean_visits(raw_in)
attr(visits_new, "cleaning_log")
```

```
#> duplicates_removed      ids_corrected
#>                  1                  1
```

```
counts <- check_counts(visits_new, raw_in$patients)
counts
```

```
#>     patients measurements       linked
#>           12           42           42
```

```
s <- summarise_windows(visits_new)
v4 <- s[s$window == "V4", ]
stopifnot(v4$n == 7, v4$of == 12, abs(v4$mean_ms - 0.564) < 0.001)    # targe
```

The raw exports contain 43 device rows, one of them with an identifier (`P7`) that does not exist in the patient list. After cleaning, the log reports one duplicate removed and one identifier corrected; the counts agree on both paths (12 patients, 42 measurements, 42 linked records), and the target value holds: in V4, 7 of 12 patients, mean 0.56 m/s.

The last check compares the result with the file that Chapter 3 wrote, cell by cell and as a cryptographic checksum (SHA-256, computed with the **digest** package). If a single character differed, the checksums would differ.

```
library(digest)
new_csv <- as_csv_text(visits_new)
b03_csv <- readLines(path_b03)
checksums <- c(rebuilt = digest(paste(new_csv, collapse = "\n"), algo = "sha2
               b03_file = digest(paste(b03_csv, collapse = "\n"), algo = "sha
substr(checksums, 1, 16)                                          # first 16 of 64 h
```

```
#>            rebuilt           b03_file
#> "1b6e35d066ddaa31" "1b6e35d066ddaa31"
```

```
stopifnot(identical(from_csv_text(new_csv), visits_b03),   # same table, cell by cell
          checksums[["rebuilt"]] == checksums[["b03_file"]])
```

The table rebuilt from raw exports is identical to the one in Chapter 3. The number in the grant report can now be regenerated by anyone with the project folder.

4.3.8. Step 8: software version and report

`sessionInfo()` lists the R version, the operating system and the attached packages. Its output belongs at the end of every report.

```
print(sessionInfo(), locale = FALSE)
```

```
#> R version 4.6.1 (2026-06-24)
#> Platform: aarch64-apple-darwin27.0.0
#> Running under: macOS Golden Gate 27.0.1
#>
#> Matrix products: default
#> BLAS:   /opt/homebrew/Cellar/openblas/0.3.34/lib/libopenblasp-r0.3.34.dylib
#> LAPACK: /opt/homebrew/Cellar/r/4.6.1_1/lib/R/lib/libRlapack.dylib;  LAPACK version 3.12.1
#>
#> attached base packages:
#> [1] stats     graphics  grDevices utils     datasets  methods   base
#>
#> other attached packages:
#> [1] digest_0.6.39
#>
#> loaded via a namespace (and not attached):
```

```
#>  [1] compiler_4.6.1  fastmap_1.2.0   cli_3.6.6       tools_4.6.1
#>  [5] htmltools_0.5.9 otel_0.2.0      yaml_2.3.12     rmarkdown_2.32
#>  [9] knitr_1.52      jsonlite_2.0.0  xfun_0.62       rlang_1.3.0
#> [13] evaluate_1.0.5
```

The Quarto report computes every number it mentions. The `r` inside back-ticks in the text below is evaluated when the report is rendered.

````
---
title: "Gait speed after hip fracture surgery (synthetic teaching data)"
format: html
---

```{r}
#| include: false
source("../R/04_figure.R")
visits <- read.csv("../data/derived/visits.csv", stringsAsFactors = FALSE)
s <- summarise_windows(visits)
v4 <- s[s$window == "V4", ]
```

In visit window V4 (days 21-35), `r v4$n` of `r v4$of` patients were measured
Their mean gait speed was `r sprintf("%.2f", v4$mean_ms)` m/s (@fig-windows)

```{r}
#| label: fig-windows
#| fig-cap: "Mean gait speed per visit window; n = patients measured of 12 (s
#| fig-alt: "Line chart of mean gait speed in five visit windows with the nu
draw_windows(s)
```

```{r}
sessionInfo()
```
````

4.3.9. Step 9: version control with Git

Git is run in a terminal, not in R. The commands below put the project under version control; the file `.gitignore` lists what Git should never record. Derived files are excluded because `run.R` regenerates them. In a project with real patient data, the raw data folder belongs in `.gitignore` too: a file that was committed once remains in the history even after it is deleted.

```
git init                                        # start recording in this folder
git add README.md run.R R/ report/report.qmd .gitignore
git commit -m "Rebuild gait figure from raw exports; target value V4 = 0.56 m/s"
git log --oneline                               # list of recorded versions
```

```
# .gitignore
data/derived/
output/
report/report.html
report/report_files/
```

4.3.10. Step 10: the artefact register

The passport of Table 4.1 can itself be produced by code, with a checksum for each data file so that a later reader can verify that the file in hand is the file described.

```
sha <- function(lines) substr(digest(paste(lines, collapse = "\n"), algo = "sha256", seriali
artefacts <- data.frame(
  artefact    = c(paste0("data/raw/", names(raw_in), ".csv"), "data/derived/visits.csv"),
  produced_by = c(rep("R/01_simulate_raw.R", 4), "R/02_clean.R"),
  rights      = "CC BY 4.0",
  data_status = "synthetic",
```

```
  rows        = c(sapply(raw_in, nrow), nrow(visits_new)),
  sha256_start = c(sapply(raw, function(df) sha(as_csv_text(df))), sha(new_cs
)
artefacts
```

```
#>                          artefact         produced_by    rights data_statu
#> patients    data/raw/patients.csv R/01_simulate_raw.R CC BY 4.0   synthet
#> visit_log  data/raw/visit_log.csv R/01_simulate_raw.R CC BY 4.0   synthet
#> stopwatch  data/raw/stopwatch.csv R/01_simulate_raw.R CC BY 4.0   synthet
#> sensormat  data/raw/sensormat.csv R/01_simulate_raw.R CC BY 4.0   synthet
#>           data/derived/visits.csv        R/02_clean.R CC BY 4.0   synthet
#>           rows sha256_start
#> patients    12 1b4cbc97ca03
#> visit_log   60 ba8abf3cbfe3
#> stopwatch   27 1a5c3dbdeda2
#> sensormat   16 42ee1164c4e8
#>             60 1b6e35d066dd
```

```
stopifnot(artefacts$rows == c(12, 60, 27, 16, 60))
```

The register shows 12 patients, 60 planned visits, 27 stopwatch rows and 16 sensor-mat rows (15 tests plus the duplicate). Saved as `ARTEFACTS.csv` next to the README, it answers the three questions of the passport for every data file.

4.4. Counter-check

A firmware update of the sensor mat changes its export from centimetres per second to metres per second, without notice. The project is run again unchanged. **Predict first:** does the run finish with a wrong target value, finish with the right one, or stop — and at which line?

```
raw_updated <- raw_in
raw_updated$sensormat$speed_cm_s <- raw_updated$sensormat$speed_cm_s / 100   # now really m/
outcome <- tryCatch({
  clean_visits(raw_updated)
  "run finished"
}, error = function(e) conditionMessage(e))
outcome
```

```
#> [1] "all(v$gait_ms[measured] > 0.05 & v$gait_ms[measured] < 2.5) is not TRUE"
```

```
stopifnot(outcome != "run finished")
```

The run stops inside `clean_visits()` at the plausibility check: after the now superfluous division by 100, rehabilitation values of about 0.005 m/s fall outside the range of 0.05 to 2.5 m/s. Without that line, the target value check would have caught the error one step later, because the V4 mean would have collapsed. In the spreadsheet version of the project, the same firmware update would have produced a plausible-looking figure with nonsense in two of its five points. A check that can fail is what turns a script into evidence.

4.5. Consequences for patients and operations

Patients. An unreproducible number in a grant report or a publication can steer resources and clinical practice. Patients are not harmed by a missing script directly, but by a decision based on a number nobody can check. The target value and the two-path count make a silent error visible before it travels.

Operations. A department can adopt three simple rules: raw exports are stored read-only; every reported number comes from a script with a

target value; and every exported file has a passport like Table 4.1. Device changes such as the firmware update need a change log that the analysis team can see.

Data protection and information security. Version control, cloud storage and code-sharing platforms are copies. For real patient data, the rule is: code and documentation may be shared under the conditions agreed, data are not committed. A raw file that once entered a Git history remains there until the history is rewritten, which is a security incident to be handled, not a clean-up task (Chapter 24). Pseudonymised identifiers such as `P01` remain personal data under the GDPR (European Parliament and Council of the European Union 2016); the key that links them to names is kept apart from the project (Chapter 23). Which platform an institution permits for which data class is a local decision (*open*, to be clarified with the data protection officer and IT security).

Rights. In this book, text and figures are licensed under CC BY 4.0 and code under the MIT licence. In a research project the rights column must name the actual terms: data use agreements, ethics approval, and third-party material such as device manuals or published figures.

4.6. Look up and transfer

- **Reviewer question:** Can the number in the main figure be regenerated from the raw data with the published code and a single command, and is there a stated target value that the code checks?
- **Self-test:** A colleague sends you `analysis.R`. It starts with `setwd("C:/Users/maria/Documents/thesis")`, loads `data_final_clean.csv` and draws the figure. There is no other file. List three things you need before you can say whether the figure is reproducible.

Answer

(1) The raw data and the code that produced `data_final_clean.csv` from them; "final clean" suggests hidden steps.
(2) A project folder with relative paths instead of `setwd()` to a personal directory, so that it runs on another computer. (3) A target value (the number the figure should show) and the software versions (`sessionInfo()`); with random elements also the seed. Useful in addition: a data dictionary with units and the Git history.

- **Annotated sources:**
  - Peng (2011) — reproducibility as the minimum standard for computational claims.
  - Sandve et al. (2013) — ten rules, five of which structure the project in this chapter.
  - Wilson et al. (2017) — practical project organisation for researchers without programming training.
  - Ziemann et al. (2016) — a *published* example of silent spreadsheet conversions in the literature.
  - Wickham (2014) — the definition of tidy (long) data used for the visit table.
  - Bryan (2018) — Git for statistical work, written for beginners.
  - Quarto (n.d.) and R Core Team (2026) — the publishing system and the R software used.

4.7. Further reading

- Chapter 3 — where the tables and their denominators come from.

- Chapter 24 — information security, backups and recovery from a verified copy.
- Chapter 27 — the model passport, which builds on the artefact passport of this chapter.
- CTTIR Course 1, Week 1 Session 2: R, RStudio, Quarto, and renv toolchain — installing and using the tools step by step.
- Methods in R: Common errors — Git and Quarto — the error messages beginners meet first.
- Methods in R: Writing a report — Reproducibility — what a reproducible report contains.
- Scientometrics in R: Reproducible Pipelines — the next step: pipelines that rerun only what changed.
- Strategy in R: R Refresher — Data structures — vectors, lists and data frames in more depth.

Part II.

Part II · How models learn and how evidence arises

5. From baseline to learning model

R. Heller · Date: 2026-02-08 · Last changed: 2026-11-14

At a glance

Guiding question: Which additional effort delivers an improvement that matters for a decision?
For: clinicians, doctoral researchers and reviewers who build or appraise a first prediction model · **Prerequisites:** Chapter 2 (prediction versus explanation), Chapter 3 (missing values and denominators), Chapter 4 (a reproducible R project)
In the course: W1 · From clinical case to question and W3 · Diagnostics, prediction and evidence · **Data:** synthetic (generated in this chapter; teaching file `data/b05-delirium-cohort.csv`, reused in Chapter 6) · **Core reading time:** about 10 minutes

5.1. Case and decision

The orthogeriatric team of a fictitious trauma department can offer proactive geriatric co-management to about 30 of every 100 patients aged 65 or older who are admitted with a hip fracture. The aim is to prevent postoperative delirium, an acute confusional state that prolongs recovery and is often missed. Today the ward uses a three-item points rule. A colleague proposes a gradient-boosting model that draws on 19 routine variables from the hospital information system. The head of the team has to decide which approach goes into a six-month pilot. The question is not

which method is more modern. It is whether the extra effort — a data feed, tuning, maintenance and explanation to staff — finds more patients who will develop delirium within the same capacity. If the team chooses wrongly, either geriatric time is spent on low-risk patients or preventable delirium goes unnoticed.

5.2. Core reading

Core reading · recommended for W1 and W3 · about 10 minutes · ends before 'Lab'

A prediction model is a rule that turns information available at one point in time into an estimated risk for a later event. The point in time matters. Here it is the first morning after admission. Everything the model uses must be known by then, and the outcome — delirium within five days after surgery — must lie in the future. With this framing fixed, the methods form a ladder. Each rung adds capacity to learn from data and adds effort.

The ladder of candidates. The *null model* predicts the same risk for everyone, namely the proportion of patients with delirium in the development data. It cannot rank anybody, but it sets the floor: a model that does not beat it has learned nothing useful. A *points rule* adds one point per risk factor, in our case age 85 or older, pre-existing cognitive impairment and an anaesthetic risk class (ASA) of 3 or higher. Clinicians can apply it at the bedside. *Logistic regression* estimates a weight for each of a small set of prespecified predictors and returns a probability (details in Methods in R: Logistic Regression). *Regularised regression* such as the elastic net offers the same model many more candidate predictors but shrinks their weights towards zero and can drop some entirely (Zou and Hastie 2005); the amount of shrinkage is a tuning parameter (Methods in R: Elastic

Net). *Tree ensembles and boosting* build many small decision trees in sequence and can capture thresholds, curved relationships and interactions without being told about them (Chen and Guestrin 2016). They have several tuning parameters. The top rung, *representation learning* with neural networks, learns its own features from raw signals such as images or spectra. It plays no part here because our data are 19 numbers per patient; Chapter 9 and Chapter 10 return to it.

Moving up the ladder does not guarantee better predictions. A systematic review of 71 studies found no average performance benefit of machine learning over logistic regression among comparisons at low risk of bias; an advantage for machine learning appeared only among comparisons at high risk of bias (Christodoulou et al. 2019). The useful question is therefore not "which method wins?" but "which improvement survives a fair comparison and changes what the team does?"

Three places, three jobs. A fair comparison needs to keep three jobs apart (Table 5.1). *Preprocessing* — here the replacement of missing laboratory values by a median — is learned from development data only and then applied unchanged to anybody else. *Tuning* — the choice of shrinkage for the elastic net and of the number of boosting rounds — uses an inner cross-validation within the development data: the 800 development patients are divided into five folds, each fold is predicted once by a model trained on the other four, and the setting with the best average performance is kept (Methods in R: Cross-Validation). The *independent test* uses 400 patients who were set aside before anything else happened and who are predicted once, by every candidate, at the very end.

Table 5.1.: Where each step of the lab takes place. Counts are *synthetic*.

Step	Data used	What happens	Do test patients take part?
Split	all 1,200 patients	random allocation of patient IDs: 800 development, 400 test	allocation only
Preprocessing	800 development patients	medians for missing haemoglobin and sodium	no
Tuning	800 development patients, 5 inner folds	shrinkage and mixing (elastic net), number of rounds (boosting)	no
Choice of family	inner cross-validation results	prespecified: regression unless a flexible model is better by 0.02 in AUC	no
Final fit	800 development patients	each candidate refitted once with its chosen settings	no
Independent test	400 test patients	one prediction per candidate, same patients for all	yes, once

The validation in this chapter is therefore exactly this: **a single random split into development and test data (internal hold-out validation), with an inner five-fold cross-validation used only for tuning and for the choice of model family.** It is not nested cross-validation, not temporal validation and not external validation. A nested cross-validation would repeat the tuning inside each of several outer folds and so estimate the performance of the whole procedure without a separate test set (Varma and Simon 2006) (Methods in R: Nested Cross-Validation). Temporal and external validation ask a different question — does the model hold in a later period or another hospital? — and are the subject of Chapter 19.

Two observations follow from Figure 5.1, which is computed in the lab. First, all models beat the null model by a wide margin, and the regression models rank patients somewhat better than the points rule. Second, the differences between regression, elastic net and boosting are small compared with the width of the intervals. With 400 test patients and about 100 events, an AUC difference of 0.02 cannot be distinguished from chance. The decision-relevant reading is the label: within the capacity of 120 patients, the candidates find between 59 and 66 of 103 future delirium cases. A handful of additional cases per 400 patients may matter — but only if the gain is real, holds in the next year's patients, and is worth a data feed of 19 variables.

The most important assumption of this comparison is that the 400 test patients come from the same hospital, period and data-generating process as the development patients. Under this assumption a random hold-out set estimates how the model performs for similar patients; it says nothing about another hospital, a new protocol or next winter's case mix.

A last point concerns honesty with the test set. If the team had compared ten models on the 400 test patients and reported the best, the winning number would be optimistic, because the test set would have been used

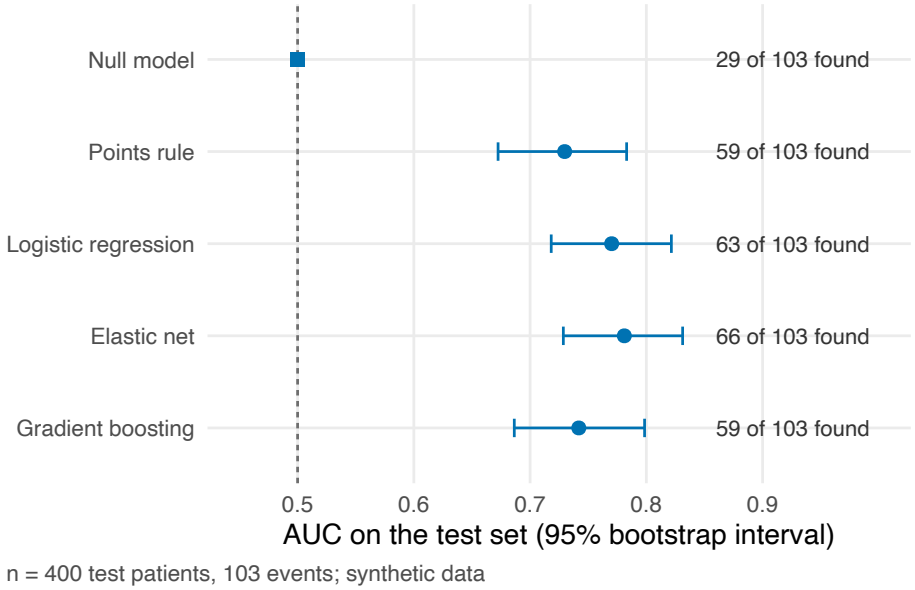


Figure 5.1.: Discrimination of five candidates on the same 400 synthetic test patients (103 with delirium). Points: area under the ROC curve (AUC); bars: 95% percentile bootstrap intervals (2,000 resamples of patients). Labels: delirium cases among the 120 patients ranked highest (30% capacity). Unit: patient; data status: synthetic.

for selection. In this chapter the choice of model family is fixed in advance and made on the inner cross-validation; the test set only confirms it.

Question for the workshop: In the test set the elastic net finds a few more cases than the rule and the regression, while boosting finds none more than the rule. Which additional evidence would you need before you accept the cost of a 19-variable data feed — and which of these questions could the hold-out test answer at all?

5.3. Lab

The lab runs in a fresh R session in well under a minute. It generates the cohort, splits it once, learns the preprocessing, tunes two models in the inner folds, compares all five candidates on the inner folds and finally predicts the test patients once.

5.3.1. Step 1 · Generate the synthetic cohort

Every value below is *synthetic*. The coefficients in the data-generating model are teaching *assumptions*, chosen to be clinically plausible in direction; they are not estimates from any study. Because the data are simulated we also know each patient's true risk, which no real data set contains. We keep it for the counter-check in Chapter 6.

```
set.seed(2605)                                          # fixed start value for the whole la
n <- 1200
id <- sprintf("P%04d", seq_len(n))                      # constant patient IDs, reused in BO

# Eight clinically motivated predictors
age <- pmin(101, 65 + round(rgamma(n, shape = 3.5, scale = 5)))           # years
cognitive_impairment <- rbinom(n, 1, plogis(-3.0 + 0.09 * (age - 65)))   # pre-existing, 0/1
asa <- as.integer(cut(rnorm(n, 0.03 * (age - 80), 1),                    # ASA class 1-4
```

```
                      c(-Inf, -1.7, 0, 1.25, Inf)))
haemoglobin <- round(rnorm(n, 12.0 - 0.02 * (age - 80), 1.6), 1)          # g
sodium <- round(rnorm(n, 138.5, 4))                                       # m
n_medications <- rpois(n, 3 + 0.08 * (age - 65))                          # m
hours_to_surgery <- round(rgamma(n, shape = 2, scale = 14))               # h
sensory_impairment <- rbinom(n, 1, 0.3)                                   # w

# Eleven further routine variables; only CRP and albumin carry a weak signal
potassium <- round(rnorm(n, 4.2, 0.5), 1)        # mmol/L
creatinine <- round(rlnorm(n, log(1.0), 0.3), 2) # mg/dL
crp <- round(rlnorm(n, log(15), 1), 0)           # mg/L
leukocytes <- round(rlnorm(n, log(9), 0.3), 1)   # 10^9/L
glucose <- round(rlnorm(n, log(120), 0.25))      # mg/dL
albumin <- round(rnorm(n, 36, 5))                # g/L
bmi <- round(rnorm(n, 24, 4), 1)                 # kg/m^2
heart_rate <- round(rnorm(n, 85, 14))            # beats/min
systolic_bp <- round(rnorm(n, 140, 22))          # mmHg
temperature <- round(rnorm(n, 36.9, 0.4), 1)     # degrees C
spo2 <- pmin(100, round(rnorm(n, 95, 2.5)))      # %

# True linear predictor (assumption): includes a U-shaped sodium effect that
lp <- -1.6 + 0.065 * (age - 80) + 1.1 * cognitive_impairment + 0.35 * (asa -
  0.18 * (haemoglobin - 12) + 0.012 * ((sodium - 139)^2 - 16) + 0.09 * (n_med
  0.012 * (hours_to_surgery - 28) + 0.35 * sensory_impairment + 0.004 * (crp
true_risk <- plogis(lp)
delirium <- rbinom(n, 1, true_risk)              # postoperative delirium wit

# Missing laboratory values (assumption: missing completely at random, see BO
haemoglobin[sample(n, 70)] <- NA
sodium[sample(n, 45)] <- NA

cohort <- data.frame(id, age, cognitive_impairment, asa, haemoglobin, sodium
                     hours_to_surgery, sensory_impairment, potassium, creatin
```

```
                      glucose, albumin, bmi, heart_rate, systolic_bp, temperature, spo2,
                      true_risk = round(true_risk, 4), delirium)
stopifnot(nrow(cohort) == 1200, !anyDuplicated(cohort$id), sum(cohort$delirium) == 320)
```

The cohort has 1,200 patients, 320 of whom develop delirium (26.7%).

5.3.2. Step 2 · Split once, by patient, and assign the inner folds

The split is made on patient IDs before any model sees the data. Each patient has one admission here; with repeated admissions the split would have to be made on the person, not the admission (Chapter 19).

```
test_ids <- sort(sample(cohort$id, 400))                          # independent test patients
cohort$split <- ifelse(cohort$id %in% test_ids, "test", "development")
dev_ids <- cohort$id[cohort$split == "development"]
inner_fold <- setNames(sample(rep(1:5, length.out = length(dev_ids))), dev_ids)
cohort$fold <- ifelse(cohort$split == "development", inner_fold[cohort$id], NA)  # 1-5; NA f

dev <- cohort[cohort$split == "development", ]
test <- cohort[cohort$split == "test", ]
stopifnot(length(intersect(dev$id, test$id)) == 0,                # no patient in both sets
          nrow(dev) == 800, nrow(test) == 400,
          sum(dev$delirium) == 217, sum(test$delirium) == 103,
          all(table(dev$fold) == 160))
table(set = cohort$split, delirium = cohort$delirium)
```

```
#>              delirium
#> set             0   1
#>   development 583 217
#>   test        297 103
```

5.3.3. Step 3 · Learn the preprocessing from development data only

```
core <- c("age", "cognitive_impairment", "asa", "haemoglobin", "sodium", "n_m
          "hours_to_surgery", "sensory_impairment")
extra <- c("potassium", "creatinine", "crp", "leukocytes", "glucose", "albumi
           "heart_rate", "systolic_bp", "temperature", "spo2")
all_pred <- c(core, extra)

medians <- sapply(dev[all_pred], median, na.rm = TRUE)          # learned on
impute <- function(d, m) {                                      # applies st
  for (v in names(m)) d[[v]][is.na(d[[v]])] <- m[[v]]
  d
}
dev_i <- impute(dev, medians)
test_i <- impute(test, medians)
stopifnot(!anyNA(dev_i[all_pred]), !anyNA(test_i[all_pred]))
medians[c("haemoglobin", "sodium")]
```

```
#> haemoglobin      sodium
#>        11.9       139.0
```

One simplification is deliberate and must be stated: the medians come from all 800 development patients, so inside the inner cross-validation the held-out fold has contributed a little to its own imputation. For a median this effect is negligible; for steps that learn more from the data, such as selecting predictors, it is not (see the box "Common mistake" below). A pipeline that refits every preprocessing step within each fold removes the issue.

5.3.4. Step 4 · Null model, points rule and logistic regression

```
auc <- function(p, y) {                          # rank formula (Mann-Whitney), equals the ROC
  r <- rank(p); n1 <- as.double(sum(y)); n0 <- length(y) - n1
  (sum(r[y == 1]) - n1 * (n1 + 1) / 2) / (n1 * n0)
}
brier <- function(p, y) mean((p - y)^2)          # mean squared error of the predicted risk

rule_points <- function(d) (d$age >= 85) + d$cognitive_impairment + (d$asa >= 3)   # 0-3 poi
fit_rule <- function(d) tapply(d$delirium, factor(rule_points(d), levels = 0:3), mean)
predict_rule <- function(tab, d) as.numeric(tab[as.character(rule_points(d))])

f_lr <- reformulate(core, response = "delirium")              # eight prespecified predict
fit_rule(dev_i)                                               # observed risk per point (d
```

```
#>         0         1         2         3
#> 0.1050420 0.1993671 0.3987342 0.7500000
```

5.3.5. Step 5 · Elastic net, tuned in the inner folds

`cv.glmnet()` (Friedman et al. 2010) runs the inner cross-validation for a sequence of shrinkage values . We pass our own fold vector so that every candidate sees the same folds, and repeat it for five values of the mixing parameter (0 = ridge, 1 = lasso). The combination with the lowest cross-validated deviance is kept.

```
library(glmnet)
x_dev <- as.matrix(dev_i[all_pred]); x_test <- as.matrix(test_i[all_pred])
alphas <- c(0, 0.25, 0.5, 0.75, 1)
cv_fits <- lapply(alphas, function(a)
  cv.glmnet(x_dev, dev_i$delirium, family = "binomial", alpha = a,
```

```
            foldid = dev_i$fold, keep = TRUE))                    # keep = ou
best <- which.min(sapply(cv_fits, function(cv) min(cv$cvm)))
en_cv <- cv_fits[[best]]
en_alpha <- alphas[best]; en_lambda <- en_cv$lambda.min
en_coef <- coef(en_cv, s = "lambda.min")
en_nonzero <- sum(en_coef[-1, 1] != 0)                          # predictors
c(alpha = en_alpha, lambda = round(en_lambda, 4), predictors_kept = en_nonzer
```

```
#>           alpha          lambda predictors_kept
#>          0.2500          0.0244         15.0000
```

5.3.6. Step 6 · Gradient boosting, tuned in the inner folds

Depth, learning rate and minimum leaf size are fixed in advance; only the number of boosting rounds is tuned, with early stopping on the same five folds. With one thread and no subsampling the result is reproducible.

```
library(xgboost)
xgb_params <- list(objective = "binary:logistic", max_depth = 2, eta = 0.05,
                   min_child_weight = 5, eval_metric = "logloss", nthread = 1
fold_index <- lapply(1:5, function(f) which(dev_i$fold == f))
xgb_cv <- xgb.cv(params = xgb_params, data = xgb.DMatrix(x_dev, label = dev_i
                 nrounds = 1000, folds = fold_index, early_stopping_rounds =
xgb_rounds <- xgb_cv$early_stop$best_iteration
xgb_rounds
```

```
#> [1] 110
```

5.3.7. Step 7 · Compare all candidates on the inner folds

All five candidates are now evaluated on the same inner folds of the development data. AUC and Brier score are computed per fold and averaged, so that the null model receives its correct AUC of 0.5. The test patients are still untouched.

```
oof <- data.frame(id = dev_i$id, fold = dev_i$fold, y = dev_i$delirium,
                  null = NA, rule = NA, lr = NA, en = NA, xgb = NA)
for (f in 1:5) {
  tr <- dev_i$fold != f; ho <- dev_i$fold == f                  # training folds and held-ou
  oof$null[ho] <- mean(dev_i$delirium[tr])
  oof$rule[ho] <- predict_rule(fit_rule(dev_i[tr, ]), dev_i[ho, ])
  oof$lr[ho] <- predict(glm(f_lr, binomial, dev_i[tr, ]), dev_i[ho, ], type = "response")
  booster <- xgb.train(params = xgb_params, nrounds = xgb_rounds, verbose = 0,
                       data = xgb.DMatrix(x_dev[tr, ], label = dev_i$delirium[tr]))
  oof$xgb[ho] <- predict(booster, xgb.DMatrix(x_dev[ho, ]))
}
oof$en <- plogis(en_cv$fit.preval[, en_cv$index["min", 1]])     # out-of-fold, at the chosen

models <- c(null = "Null model", rule = "Points rule", lr = "Logistic regression",
            en = "Elastic net", xgb = "Gradient boosting")
fold_mean <- function(metric, m) mean(sapply(1:5, function(f) {
  s <- oof$fold == f; metric(oof[[m]][s], oof$y[s]) }))
inner <- data.frame(model = models,
                    auc = sapply(names(models), function(m) fold_mean(auc, m)),
                    brier = sapply(names(models), function(m) fold_mean(brier, m)))
inner$auc[1] <- 0.5                                             # constant prediction: exact
stopifnot(isTRUE(all.equal(fold_mean(auc, "null"), 0.5)))

# Prespecified choice of family: regression unless a flexible model is better by >= 0.02 AUC
gain_flex <- max(inner["en", "auc"], inner["xgb", "auc"]) - inner["lr", "auc"]
chosen <- if (gain_flex >= 0.02) "flexible model" else "logistic regression"
knitr::kable(inner, digits = 3, row.names = FALSE,
```

```
            col.names = c("Candidate", "AUC (inner CV)", "Brier score (inner
            caption = "Inner five-fold cross-validation on 800 synthetic dev
```

Table 5.2.: Inner five-fold cross-validation on 800 synthetic development patients (217 with delirium); metrics averaged over folds.

Candidate	AUC (inner CV)	Brier score (inner CV)
Null model	0.500	0.198
Points rule	0.746	0.160
Logistic regression	0.775	0.150
Elastic net	0.773	0.152
Gradient boosting	0.757	0.157

On the development data the best flexible model differs from logistic regression by -0.002 in AUC, so the prespecified criterion selects **logistic regression**. The cross-validated values of the two tuned models are slightly optimistic, because the same folds served for tuning and for this comparison; nested cross-validation would remove that optimism.

5.3.8. Step 8 · Final fits and one look at the test set

```
lr_fit <- glm(f_lr, binomial, dev_i)
booster <- xgb.train(params = xgb_params, nrounds = xgb_rounds, verbose = 0,
                     data = xgb.DMatrix(x_dev, label = dev_i$delirium))
pred <- data.frame(id = test_i$id, y = test_i$delirium,
                   null = mean(dev_i$delirium),
                   rule = predict_rule(fit_rule(dev_i), test_i),
                   lr = predict(lr_fit, test_i, type = "response"),
                   en = as.numeric(predict(en_cv, newx = x_test, s = "lambda
                   xgb = predict(booster, xgb.DMatrix(x_test)))
```

```
stopifnot(identical(pred$id, test$id))                          # same test patients for eve

# Paired bootstrap: resample test patients, recompute every AUC on the same resample
set.seed(5051)
boot <- replicate(2000, {
  i <- sample(nrow(pred), replace = TRUE)
  sapply(names(models)[-1], function(m) auc(pred[[m]][i], pred$y[i]))
})

# Decision-oriented reading: cases among the 120 highest-ranked patients (30% capacity).
# Ties (rule: four risk levels; null model: one) are broken in a fixed random order.
set.seed(5052)
tie_breaker <- runif(nrow(pred))
found_in_top <- function(p, y, k = 120) sum(y[order(-p, tie_breaker)[seq_len(k)]])

res <- data.frame(model = models,
                  auc = sapply(names(models), function(m) auc(pred[[m]], pred$y)),
                  lower = c(NA, apply(boot, 1, quantile, 0.025)),
                  upper = c(NA, apply(boot, 1, quantile, 0.975)),
                  brier = sapply(names(models), function(m) brier(pred[[m]], pred$y)),
                  found = sapply(names(models), function(m) found_in_top(pred[[m]], pred$y)),
                  predictors = c(0, 3, length(core), en_nonzero, length(all_pred)))
res$auc[1] <- 0.5
diff_xgb_lr <- quantile(boot["xgb", ] - boot["lr", ], c(0.025, 0.975))
diff_en_lr <- quantile(boot["en", ] - boot["lr", ], c(0.025, 0.975))
stopifnot(abs(res["lr", "auc"] - 0.770) < 0.002, abs(res["rule", "auc"] - 0.730) < 0.002,
          abs(res["en", "auc"] - 0.781) < 0.003, abs(res["xgb", "auc"] - 0.742) < 0.01,
          res["lr", "found"] == 63, sum(pred$y) == 103)
knitr::kable(res, digits = 3, row.names = FALSE,
             col.names = c("Candidate", "AUC", "2.5%", "97.5%", "Brier score",
                           "Cases in top 120", "Predictors used"),
             caption = "Single evaluation on 400 synthetic test patients (103 with delirium)
```

Table 5.3.: Single evaluation on 400 synthetic test patients (103 with delirium). Interval: paired percentile bootstrap, 2,000 resamples of patients.

Candidate	AUC	2.5%	97.5%	Brier score	Cases in top 120	Predictors used
Null model	0.500	NA	NA	0.191	29	0
Points rule	0.730	0.672	0.783	0.161	59	3
Logistic regression	0.770	0.718	0.822	0.157	63	8
Elastic net	0.781	0.729	0.831	0.152	66	15
Gradient boosting	0.742	0.686	0.798	0.163	59	19

The following code draws Figure 5.1 in the core reading from the objects `res` and `pred`.

```
library(ggplot2)
plot_df <- res
plot_df$model <- factor(plot_df$model, levels = rev(models))
plot_df$label <- paste0(plot_df$found, " of ", sum(pred$y), " found")
ggplot(plot_df, aes(x = auc, y = model)) +
  geom_vline(xintercept = 0.5, linetype = cdsr_linetypes[2], colour = cdsr_co
  geom_errorbar(aes(xmin = lower, xmax = upper), width = 0.2, orientation = "
                colour = cdsr_colours[["blue"]], na.rm = TRUE) +
  geom_point(aes(shape = model == "Null model"), size = 2.6, colour = cdsr_co
  geom_text(aes(x = 0.86, label = label), hjust = 0, size = 3.3, colour = "#3
  scale_shape_manual(values = c(`FALSE` = cdsr_shapes[1], `TRUE` = cdsr_shape
  scale_x_continuous(limits = c(0.45, 1.0), breaks = seq(0.5, 0.9, 0.1)) +
  labs(x = "AUC on the test set (95% bootstrap interval)", y = NULL,
       caption = "n = 400 test patients, 103 events; synthetic data") +
  theme_cdsr()
```

The test results agree with the inner cross-validation in two respects: the regression ranks better than the rule, and no flexible model is clearly better than the regression. Boosting tends to rank slightly worse: the paired bootstrap interval for its AUC difference to the regression runs from -0.056 to 0.000. The elastic net tends to rank slightly better, with an interval from -0.001 to 0.024. Both intervals reach zero, and in cases found within the capacity the three models differ by at most 7 patients. Brier scores — the mean squared distance between predicted risk and outcome — tell the same story.

> ⚠ Common mistake
>
> **Learning from the test patients before the test.** Typical forms are: choosing the ten predictors most strongly associated with delirium on the full data set and then cross-validating; standardising or imputing with statistics of all patients; trying several models on the test set and reporting the best. Each step lets information about the held-out outcomes leak into the model, and the reported performance becomes optimistic. PROBAST+AI therefore asks whether data leakage was avoided when data were split and whether all model development steps were repeated in each resampling round (Moons et al. 2025, model evaluation, signalling questions 4.5 and 4.6).

5.3.9. Step 9 · Export the teaching file

The cohort with its split and fold columns is written to `data/b05-delirium-cohort.csv`. Chapter 6 reads exactly this file, so both chapters use identical patients and identical test IDs.

```
write.csv(cohort, "../data/b05-delirium-cohort.csv", row.names = FALSE)
check <- read.csv("../data/b05-delirium-cohort.csv")
stopifnot(nrow(check) == 1200, sum(check$split == "test") == 400,
          identical(sort(check$id[check$split == "test"]), test_ids))
```

Other implementation: Python/scikit-learn

Original implementation. scikit-learn (Pedregosa et al. 2011), stable release 1.9.1 (documentation checked November 2026): `Pipeline`, `PredefinedSplit`, `GroupKFold`, `LogisticRegression`.
Concrete advantage. A `Pipeline` of imputer, scaler and model, tuned with `GridSearchCV`, refits the preprocessing inside every inner fold and so removes the simplification of Step 3. `GroupKFold` keeps all admissions of a patient in one fold.
What must be identical. The input file `data/b05-delirium-cohort.csv`, its split column and its inner folds (imported with `PredefinedSplit(test_fold = fold - 1)`, not regenerated), the unit (patient), the metrics and an agreed tolerance, for example ±0.005 in AUC. According to the scikit-learn user guide the penalty is divided by $n \cdot C$, so glmnet's corresponds to $1/(n \cdot C)$ and to `l1_ratio` (Friedman et al. 2010); glmnet standardises internally, scikit-learn only with a scaler in the pipeline. Elastic-net penalties need `solver = "saga"`, the `penalty` argument is deprecated since version 1.8, and `LogisticRegression` is penalised by default (`C = 1.0`), so the unpenalised `glm()` of Step 4 corresponds to `C = np.inf`.
Status: conceptual comparison, not executed. No statement about speed is made.

5.4. Counter-check

Change one thing: the size of the development data. Keep the 400 test patients as they are, but develop the regression, the elastic net (with the mixing parameter chosen above) and the boosting model on a random subset of only 200 of the 800 development patients. Because a single subset is a single draw, repeat this 25 times.

Before you run the code, write down your prediction: Which candidate loses most discrimination? And in which direction does the calibration slope of each model move? The calibration slope is the coefficient obtained when the outcome is regressed on the logit of the predicted risk. A slope below 1 means that predictions are too extreme (high risks too high, low risks too low), which is typical of overfitting; a slope above 1 means that they are too cautious. Chapter 6 develops calibration in detail.

```
cal_slope <- function(p, y) unname(coef(glm(y ~ qlogis(p), family = binomial))[2])
one_small_run <- function() {
  small <- dev_i[dev_i$id %in% sample(dev_i$id, 200), ]        # 200 patients, inner folds k
  x_small <- as.matrix(small[all_pred])
  p_lr <- predict(glm(f_lr, binomial, small), test_i, type = "response")
  en_s <- cv.glmnet(x_small, small$delirium, family = "binomial", alpha = en_alpha, foldid =
  p_en <- as.numeric(predict(en_s, newx = x_test, s = "lambda.min", type = "response"))
  xgb_s <- xgb.cv(params = xgb_params, data = xgb.DMatrix(x_small, label = small$delirium),
                  nrounds = 1000, folds = lapply(1:5, function(f) which(small$fold == f)),
                  early_stopping_rounds = 30, verbose = 0)
  p_xgb <- predict(xgb.train(params = xgb_params, verbose = 0, nrounds = xgb_s$early_stop$be
                             data = xgb.DMatrix(x_small, label = small$delirium)), xgb.DMatr
  c(events = sum(small$delirium),
    auc_lr = auc(p_lr, pred$y), auc_en = auc(p_en, pred$y), auc_xgb = auc(p_xgb, pred$y),
    slope_lr = cal_slope(p_lr, pred$y), slope_en = cal_slope(p_en, pred$y),
    slope_xgb = cal_slope(p_xgb, pred$y))
}
set.seed(5053)
runs <- t(replicate(25, one_small_run()))                      # 25 subsets x 7 results
q <- function(v) quantile(v, c(0.1, 0.5, 0.9))
counter <- data.frame(
  model = c("Logistic regression", "Elastic net", "Gradient boosting"),
  auc_800 = c(res["lr", "auc"], res["en", "auc"], res["xgb", "auc"]),
  auc_200 = sapply(c("auc_lr", "auc_en", "auc_xgb"), function(v) q(runs[, v])[2]),
  slope_800 = c(cal_slope(pred$lr, pred$y), cal_slope(pred$en, pred$y), cal_slope(pred$xgb,
  slope_200 = sapply(c("slope_lr", "slope_en", "slope_xgb"), function(v) q(runs[, v])[2]),
```

```
  slope_range = sapply(c("slope_lr", "slope_en", "slope_xgb"), function(v)
    paste(sprintf("%.2f", q(runs[, v])[c(1, 3)]), collapse = " to ")))
stopifnot(q(runs[, "slope_lr"])[3] < 1,                          # regression
          q(runs[, "auc_xgb"])[2] < q(runs[, "auc_lr"])[2])
knitr::kable(counter, digits = 2, row.names = FALSE,
             col.names = c("Candidate", "AUC, 800 dev.", "AUC, 200 dev. (medi
                           "Slope, 800 dev.", "Slope, 200 dev. (median)", "Sl
             caption = "Counter-check: development on 25 random subsets of 20
```

Table 5.4.: Counter-check: development on 25 random subsets of 200 instead of 800 synthetic patients; evaluation always on the same 400 test patients.

Candidate	AUC, 800 dev.	AUC, 200 dev. (median)	Slope, 800 dev.	Slope, 200 dev. (median)	Slope, 10th to 90th percentile
Logistic regression	0.77	0.76	0.87	0.73	0.57 to 0.91
Elastic net	0.78	0.76	1.07	1.13	0.85 to 1.55
Gradient boosting	0.74	0.72	0.90	0.97	0.72 to 1.45

With 200 development patients the subsets contain a median of 54 events for eight prespecified predictors. Discrimination drops only modestly for the regression models and somewhat more for boosting, which has more to learn from fewer cases. The clearer change is in calibration. The unpenalised regression becomes systematically overconfident: its median slope is 0.73 and stays below 1 in at least 90% of the runs. The tuned models scatter in both directions, because the amount of shrinkage or the number of boosting rounds is itself estimated from few events — sometimes too

much, sometimes too little. This is why the sample size of a prediction model is planned to limit overfitting and not only to estimate the AUC precisely (Riley et al. 2020). An AUC alone would have hidden the problem, which is the starting point of Chapter 6.

5.5. Consequences for patients and operations

Patients. The model decides who is offered prevention, so its errors fall on people: a patient ranked below the capacity line loses a chance of prevention. The points rule allocates care partly by age (85 or older), which needs a clinical justification and a check of how it treats younger patients with other risk factors. Whether errors fall unevenly on groups is a question for the test set, reported with denominators (Chapter 6, Chapter 21).

Operations. The rule needs three items the admitting clinician already knows; the regression eight values, two of them laboratory results with missing values; the flexible models a reliable feed of 19 variables, a documented imputation, a versioned model object and a plan for monitoring and updating (Chapter 20, Chapter 27). A gain of a few cases per 400 patients has to be weighed against these costs and against a feed that fails silently.

Regulation. As soon as the output is used to allocate care to individual patients, the intended purpose of the software has to be examined. If it qualifies as a medical device, Rule 11 in Annex VIII of the Medical Device Regulation classifies software that provides information used to take decisions with diagnostic or therapeutic purposes as class IIa or higher (European Parliament and Council of the European Union 2017a); the guidance of the Medical Device Coordination Group (Medical Device Coordination Group 2025) — guidance, not binding law — describes how qualification is assessed. Whether an in-house pilot needs this assessment is a question for the local regulatory function (*open*; Chapter 22).

Data protection and security. Developing the model from routine data is a secondary use of health data and needs its own legal basis and documentation (Chapter 23, Chapter 26). The stored model and its preprocessing statistics are artefacts whose integrity has to be protected, because a changed median or coefficient changes every prediction (Chapter 24).

5.6. Look up and transfer

Reviewer question: "Which data informed the preprocessing, which the tuning, which the choice of model family and which the reported performance — and did any test patient contribute to the first three?" TRIPOD+AI asks authors to describe how the data were partitioned and to specify all model-building steps, including hyperparameter tuning and the method of internal validation (Collins et al. 2024, items 12a and 12c).

Self-test: A research group reports a random forest for delirium with an AUC of 0.86 from ten-fold cross-validation on all 600 patients of their hospital. Before the cross-validation they selected the ten predictors most strongly associated with delirium in the full data set. Which validation type is this, and what does the AUC of 0.86 estimate?

Answer

It is internal validation by cross-validation, but with predictor selection outside the resampling loop. The selection has already seen the outcomes of every patient who is later held out, so the AUC of 0.86 is an optimistic estimate of the performance of the whole procedure. A correct estimate would repeat the selection inside each training fold (nested resampling) or confirm the final model on patients who played no part in development. Even then the result would describe new patients from the same hospital and period, not another hospital.

Annotated sources:

- Zou and Hastie (2005) introduces the elastic net, the penalty that mixes ridge and lasso and is tuned in Step 5.
- Friedman et al. (2010) describes the coordinate-descent algorithm and parameterisation of glmnet used in the lab.
- Chen and Guestrin (2016) describes the gradient-boosting system behind the flexible candidate.
- Christodoulou et al. (2019) is the systematic review showing no average benefit of machine learning over logistic regression in comparisons at low risk of bias.
- Varma and Simon (2006) shows that cross-validated error of a tuned classifier is biased and that nested cross-validation reduces the bias.
- Riley et al. (2020) supports the counter-check: sample size for prediction models must limit overfitting, not only ensure precise performance estimates.
- Moons et al. (2025) and Collins et al. (2024) provide the appraisal and reporting items cited for leakage, tuning and internal validation.
- Pedregosa et al. (2011) is the reference for the scikit-learn implementation in the comparison box.
- The Medical Device Regulation (European Parliament and Council of the European Union 2017a) and MDCG 2019-11 Rev. 1 (Medical Device Coordination Group 2025) are the regulatory starting points for software that allocates care.

5.7. Further reading

- Chapter 6 — the same test patients, now with calibration, thresholds and consequences for two groups.
- Chapter 19 — what changes when the test patients come from another centre or a later period.
- Methods in R: Nested Cross-Validation — how to estimate the performance of a tuned procedure without a separate test set.

- Methods in R: Gradient Boosting — the tuning parameters fixed in Step 6 and their effect.
- CTTIR Courses, Course 4 · Cross-validation, nested CV, bootstrap .632+ — a full lab on resampling designs.
- CTTIR Courses, Course 4 · Regularisation: ridge, lasso, elastic net — practising the shrinkage path behind Step 5.
- CTTIR Courses, Course 2 · Logistic regression (binomial GLM) — the regression baseline in more depth.
- CTTIR Tutorials: XGBoost — a stand-alone tutorial on the boosting library used here.

6. Accuracy, calibration and clinical utility

R. Heller · Date: 2025-09-07 · Last changed: 2026-11-14

At a glance

Guiding question: When does a good ranking lead to the wrong action?
For: clinicians, nurses, doctoral researchers and reviewers who have to turn a predicted risk into a rule for action · **Prerequisites:** Chapter 5 (the delirium cohort and its validation design); the 2×2 table is introduced here and applied again in Chapter 9
In the course: W3 · Diagnostics, prediction and evidence and W8 · Colloquium and decision · **Data:** synthetic (the 400 test patients of Chapter 5, read from `data/b05-delirium-cohort.csv`) · **Core reading time:** about 11 minutes

6.1. Case and decision

The orthogeriatric team from Chapter 5 has chosen the prespecified logistic regression for its pilot. Now the standard operating procedure needs one number: from which predicted risk of postoperative delirium is a patient offered proactive geriatric co-management? The ward manager remembers a slide that said "AUC 0.77" and asks whether that settles it. The orthogeriatrician worries about something else: the oldest patients seem

to be flagged almost always, while a 74-year-old with several risk factors was not flagged last month and developed delirium. The team has to fix a threshold, say how many alarms per week it will produce, and explain to the nursing staff which patients will be missed. A threshold that is too low exhausts geriatric capacity; one that is too high leaves preventable delirium unaddressed — and the two errors do not fall on the same patients.

6.2. Core reading

Core reading · required for W3 · about 11 minutes · ends before 'Lab'

A prediction model produces two kinds of information. It *ranks* patients from lower to higher risk, and it attaches a *number* to each patient, the predicted probability. Different measures judge different parts of this output, and a model can be good at one and poor at the other.

Discrimination: does the ranking work? The area under the ROC curve (AUC, also called the c-statistic) is the probability that a randomly chosen patient who develops delirium receives a higher predicted risk than a randomly chosen patient who does not. For our model on the 400 test patients it is 0.77. The AUC uses only the order of the predictions; multiplying every risk by two would not change it (Methods in R: ROC and AUC for Logistic Models). It therefore cannot tell whether a predicted risk of 20% means that about one in five such patients develops delirium, and it does not involve any threshold.

Accuracy at a threshold: what happens to people? Once the team decides to act on predicted risks of 20% or more, every test patient falls into one cell of a 2×2 table (Table 6.1). The table is the most direct translation of a model into consequences, and it is clearer in absolute numbers than in percentages (Hoffrage and Gigerenzer 1998). Of 1,000 patients like these, about 258 will develop delirium. At the 20% threshold the model flags

about 456. Of the flagged patients, about 188 will develop delirium and about 268 will not. About 70 patients who will develop delirium are not flagged. Sensitivity is 72.8%, the positive predictive value 41.2%.

Table 6.1.: The model of Chapter 5 at a threshold of 20% on the same 400 test patients. Unit: patient; counts are *synthetic*.

At a threshold of 20%	Delirium develops	No delirium	Total
Flagged (risk 20%)	75	107	182
Not flagged	28	190	218
Total	103	297	400

Calibration: can the number be taken at face value? A model is calibrated when, among patients with a predicted risk of about 20%, about 20% develop the outcome. Figure 6.1 groups the test patients into ten groups of 40 by predicted risk and compares the mean prediction with the observed proportion. Points near the diagonal mean that the numbers can be read as risks. Two summaries go with the plot. The *calibration intercept* (here -0.01) shows whether risks are too high or too low on average; zero is ideal. The *calibration slope* (here 0.87) shows whether they are too extreme or too timid; one is ideal (Van Calster et al. 2019). Our model is close to the ideal on average, with a slope slightly below 1 — a mild sign of overfitting that the counter-check in Chapter 5 made more visible.

The threshold is a value judgement. Choosing 20% means that the team accepts flagging up to four patients without delirium (80 : 20) for each patient with delirium whom it reaches. This is not a statistical property; it expresses how the team weighs a missed case against an unnecessary consultation. Decision curve analysis makes this explicit by computing the *net benefit* of the model across thresholds (Vickers and Elkin 2006), and it belongs in every evaluation of a model meant for clinical decisions (Steyerberg et al. 2010). The lab computes it.

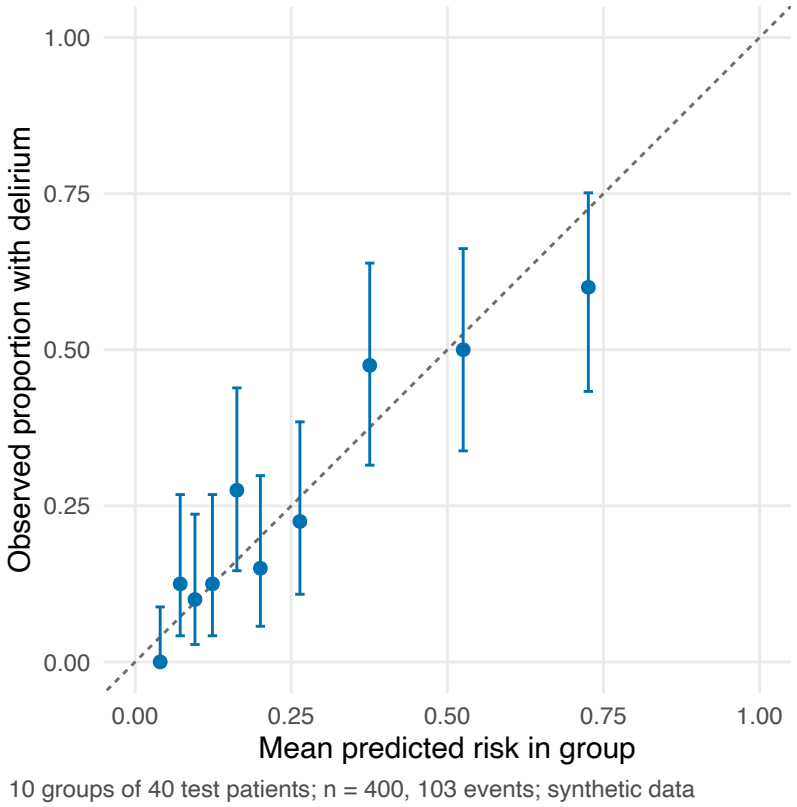


Figure 6.1.: Calibration of the logistic regression on 400 synthetic test patients (103 with delirium). Points: observed proportion with delirium in ten groups of 40 patients ordered by predicted risk; vertical bars: exact binomial 95% intervals; dashed line: perfect calibration. Unit: patient; data status: synthetic.

Same model, same threshold, two groups. Table 6.2 splits the test patients by age. The model is calibrated in both groups: the mean predicted risk is close to the observed proportion in each. Yet the consequences differ sharply. Among patients aged 80 or older, 74% are flagged and 11 of 78 delirium cases are missed. Among patients aged 65–79, only 22 of 183 are flagged and 17 of 25 delirium cases are missed. The reason is not a faulty model but the lower base rate in the younger group, which places most of its future cases below the threshold.

Table 6.2.: Consequences of the 20% threshold in two age groups of the 400 synthetic test patients. Unit: patient.

Age group	Patients	With delirium	Flagged	Cases found	Cases missed	Flagged without delirium	Mean predicted	Observed
65–79 years	183	25	22	8	17	14	11.9%	13.7%
80 years or older	217	78	160	67	11	93	37.6%	35.9%

Four measures, four questions. It helps to keep the measures apart by the question each answers. The AUC answers "does the model put patients in the right order?" Calibration answers "can I read the number as a risk?" The 2×2 table at a threshold answers "how many patients are flagged, found and missed, and how much work follows?" Net benefit answers "is acting on the model better than treating everybody or nobody, given how we weigh the two errors?" A report that gives only the first answer leaves the clinical team to guess the other three. Overall measures such as

the Brier score, the mean squared difference between predicted risk and outcome, combine ranking and calibration in one number and are useful for comparing models, but they do not replace the others (Steyerberg et al. 2010) (Methods in R: Diagnostic Test Accuracy).

The most important assumption of this chapter is that the predicted risks are calibrated in the population and setting where the threshold is applied. A threshold on the probability scale is only meaningful if the probability is right. If the outcome becomes more common, or is recorded more thoroughly, the ranking may stay as good as before while every number is too low — and the threshold quietly stops meaning what the team decided. The counter-check shows this.

Question for the workshop: The 20% threshold flags 46% of patients, more than the 30% capacity of the team. Would you raise the threshold, accept a lower sensitivity in the younger group, or argue for more capacity — and how would you explain your choice to a 74-year-old patient who is not flagged?

6.3. Lab

6.3.1. Step 1 · Read the teaching file and confirm the test patients

The lab reads the cohort written by Chapter 5. A checksum of the sorted test IDs proves that the test patients are the same as there.

```
cohort <- read.csv("../data/b05-delirium-cohort.csv")
dev <- cohort[cohort$split == "development", ]
test <- cohort[cohort$split == "test", ]
stopifnot(nrow(cohort) == 1200, nrow(test) == 400, sum(test$delirium) == 103
          length(intersect(dev$id, test$id)) == 0,
          digest::digest(sort(test$id)) == "9b43cfb98549b2a4c17e4feaca07de46
```

6.3.2. Step 2 · Refit the prespecified logistic regression

Preprocessing and model are exactly those of Chapter 5: medians from the development data, eight predictors, linear terms.

```
core <- c("age", "cognitive_impairment", "asa", "haemoglobin", "sodium", "n_medications",
          "hours_to_surgery", "sensory_impairment")
medians <- sapply(dev[core], median, na.rm = TRUE)          # learned on development data or
impute <- function(d) { for (v in core) d[[v]][is.na(d[[v]])] <- medians[[v]]; d }
dev <- impute(dev); test <- impute(test)
fit <- glm(reformulate(core, "delirium"), family = binomial, data = dev)
p <- predict(fit, test, type = "response")                  # predicted risk, test patients
y <- test$delirium                                          # observed outcome, 0/1
auc_test <- local({ r <- rank(p); n1 <- as.double(sum(y)); n0 <- length(y) - n1
  (sum(r[y == 1]) - n1 * (n1 + 1) / 2) / (n1 * n0) })
stopifnot(abs(auc_test - 0.770) < 0.001)                    # identical to B05
per1000 <- function(x) round(1000 * x / length(y))          # natural frequency per 1,000 pa
```

6.3.3. Step 3 · The 2 × 2 table at a threshold of 20%

```
threshold <- 0.20
flag <- p >= threshold
tab <- c(tp = sum(flag & y == 1), fp = sum(flag & y == 0),
         fn = sum(!flag & y == 1), tn = sum(!flag & y == 0))
sens <- tab["tp"] / sum(y)                    # flagged among patients with delirium
spec <- tab["tn"] / sum(1 - y)                # not flagged among patients without deliri
ppv <- tab["tp"] / sum(flag)                  # delirium among flagged patients
npv <- tab["tn"] / sum(!flag)                 # no delirium among non-flagged patients
stopifnot(tab["tp"] == 75, tab["fp"] == 107, tab["fn"] == 28, tab["tn"] == 190)
exact_ci <- function(x, n) binom.test(x, n)$conf.int   # Clopper-Pearson interval
round(rbind(sensitivity = c(sens, exact_ci(tab["tp"], sum(y))),
```

```
            specificity = c(spec, exact_ci(tab[["tn"]], sum(1 - y))),
            ppv = c(ppv, exact_ci(tab[["tp"]], sum(flag))),
            npv = c(npv, exact_ci(tab[["tn"]], sum(!flag)))), 3)
```

```
#>                [,1]  [,2]  [,3]
#> sensitivity 0.728 0.632 0.811
#> specificity 0.640 0.582 0.694
#> ppv         0.412 0.340 0.487
#> npv         0.872 0.820 0.913
```

The intervals are wide because the denominators are small: 103 patients with delirium for the sensitivity, 182 flagged patients for the positive predictive value. Every percentage in a report should come with its denominator (CTTIR Courses, Course 1 · Diagnostic testing: Se, Sp, PPV, NPV, LR).

6.3.4. Step 4 · AUC with an uncertainty interval

Two common intervals are compared: the analytical DeLong interval (DeLong et al. 1988) as implemented in **pROC** (Robin et al. 2011), and a percentile bootstrap over patients.

```
library(pROC)
roc_obj <- roc(y, p, levels = c(0, 1), direction = "<", quiet = TRUE)
ci_delong <- as.numeric(ci.auc(roc_obj, method = "delong"))[c(1, 3)]
set.seed(6062)
boot_auc <- replicate(2000, {
  i <- sample(length(y), replace = TRUE)
  as.numeric(auc(roc(y[i], p[i], levels = c(0, 1), direction = "<", quiet = T
})
ci_boot <- quantile(boot_auc, c(0.025, 0.975))
```

```
stopifnot(abs(as.numeric(auc(roc_obj)) - auc_test) < 1e-12)
round(rbind(DeLong = ci_delong, bootstrap = ci_boot), 3)
```

```
#>             2.5% 97.5%
#> DeLong    0.718 0.822
#> bootstrap 0.718 0.820
```

Both intervals run from about 0.72 to 0.82. With 103 events, a second model would need a clearly larger AUC before any difference could be claimed — one reason why Chapter 5 did not crown a winner.

6.3.5. Step 5 · Calibration: plot, intercept and slope

```
lp <- qlogis(p)                                              # linear predictor (log-odds)
cal_int <- unname(coef(glm(y ~ offset(lp), family = binomial))[1])   # calibration-in-the-la
cal_slope <- unname(coef(glm(y ~ lp, family = binomial))[2])         # calibration slope
o_e <- sum(y) / sum(p)                                       # observed / expected events
bins <- cut(rank(p, ties.method = "first"), 10, labels = FALSE)      # ten groups of 40
cal <- do.call(rbind, lapply(split(seq_along(p), bins), function(i) {
  ci <- binom.test(sum(y[i]), length(i))$conf.int
  data.frame(n = length(i), predicted = mean(p[i]), observed = mean(y[i]),
             lower = ci[1], upper = ci[2])
}))
stopifnot(all(cal$n == 40), abs(cal_int) < 0.05, cal_slope > 0.8, cal_slope < 0.95)
round(c(intercept = cal_int, slope = cal_slope, observed_expected = o_e), 3)
```

```
#>         intercept             slope observed_expected
#>            -0.007             0.865             0.996
```

The following code draws Figure 6.1 in the core reading.

```
library(ggplot2)
ggplot(cal, aes(x = predicted, y = observed)) +
  geom_abline(intercept = 0, slope = 1, linetype = cdsr_linetypes[2], colour
  geom_errorbar(aes(ymin = lower, ymax = upper), width = 0.015, colour = cdsr
  geom_point(size = 2.4, shape = cdsr_shapes[1], colour = cdsr_colours[["blue
  coord_equal(xlim = c(0, 1), ylim = c(0, 1)) +
  labs(x = "Mean predicted risk in group", y = "Observed proportion with deli
       caption = "10 groups of 40 test patients; n = 400, 103 events; synthet
  theme_cdsr()
```

A smooth calibration curve is an alternative to grouped points but needs larger samples to be stable (Van Calster et al. 2019); with 400 patients the grouped version with intervals is the more honest display (Methods in R: Calibration Plots).

6.3.6. Step 6 · From one threshold to many

```
at_threshold <- function(t) {
  f <- p >= t
  c(threshold = t, flagged = sum(f), flagged_pct = mean(f), tp = sum(f & y ==
    fp = sum(f & y == 0), fn = sum(!f & y == 1), sensitivity = sum(f & y == 1
    ppv = sum(f & y == 1) / sum(f), alarms_per_case = sum(f) / sum(f & y == 1
}
thr <- as.data.frame(t(sapply(c(0.10, 0.15, 0.20, 0.30, 0.40), at_threshold))
stopifnot(thr$flagged[thr$threshold == 0.30] <= 0.32 * length(y))   # 30% thr
knitr::kable(thr, digits = c(2, 0, 2, 0, 0, 0, 2, 2, 1), row.names = FALSE,
             col.names = c("Threshold", "Flagged", "Share flagged", "Cases fo
                           "Cases missed", "Sensitivity", "PPV", "Alarms per
             caption = "Consequences of five thresholds for the same model on
```

Table 6.3.: Consequences of five thresholds for the same model on 400 synthetic test patients (103 with delirium).

Threshold	Flagged	Share flagged	Cases found	Flagged, no delirium	Cases missed	Sensitivity	PPV	Alarms per case found
0.10	292	0.73	96	196	7	0.93	0.33	3.0
0.15	235	0.59	88	147	15	0.85	0.37	2.7
0.20	182	0.46	75	107	28	0.73	0.41	2.4
0.30	126	0.32	65	61	38	0.63	0.52	1.9
0.40	92	0.23	50	42	53	0.49	0.54	1.8

The last column is the number of flagged patients per delirium case reached, the reciprocal of the positive predictive value. Of the five thresholds, 30% is the lowest that keeps the share of flagged patients near the team's capacity of about 30%; it does so at the price of 38 missed cases instead of 28.

6.3.7. Step 7 · Net benefit and the decision curve

Net benefit counts correctly flagged patients and subtracts flagged patients without delirium, weighted by the odds of the threshold: $\mathrm{NB} = \mathrm{TP}/\mathrm{n} - \mathrm{FP}/\mathrm{n} \times \mathrm{t}/(1 - \mathrm{t})$ (Vickers and Elkin 2006). It is compared with flagging everybody and flagging nobody (net benefit 0).

```
ts <- seq(0.05, 0.50, by = 0.01)
nb_model <- sapply(ts, function(t) { f <- p >= t
  sum(f & y == 1) / length(y) - sum(f & y == 0) / length(y) * t / (1 - t) })
prev <- mean(y)
nb_all <- prev - (1 - prev) * ts / (1 - ts)
dca <- rbind(data.frame(threshold = ts, net_benefit = nb_model, strategy = "Model"),
             data.frame(threshold = ts, net_benefit = nb_all, strategy = "Flag everybody"),
```

```
            data.frame(threshold = ts, net_benefit = 0, strategy = "Flag nob
nb20 <- nb_model[which.min(abs(ts - 0.20))]
stopifnot(nb20 > nb_all[which.min(abs(ts - 0.20))], nb20 > 0)
```

```
library(ggplot2)
ggplot(dca, aes(x = threshold, y = net_benefit, colour = strategy, linetype =
  geom_line(linewidth = 0.9) +
  geom_vline(xintercept = 0.20, colour = cdsr_colours[["grey"]], linetype = d
  scale_colour_manual(values = c(Model = cdsr_colours[["blue"]], `Flag everyb
                                 `Flag nobody` = cdsr_colours[["grey"]])) +
  scale_linetype_manual(values = c(Model = cdsr_linetypes[1], `Flag everybody
                                   `Flag nobody` = cdsr_linetypes[3])) +
  coord_cartesian(ylim = c(-0.05, 0.27)) +
  labs(x = "Threshold probability", y = "Net benefit", colour = NULL, linetyp
       caption = "Vertical line: threshold of 20%. n = 400 test patients; syn
  theme_cdsr()
```

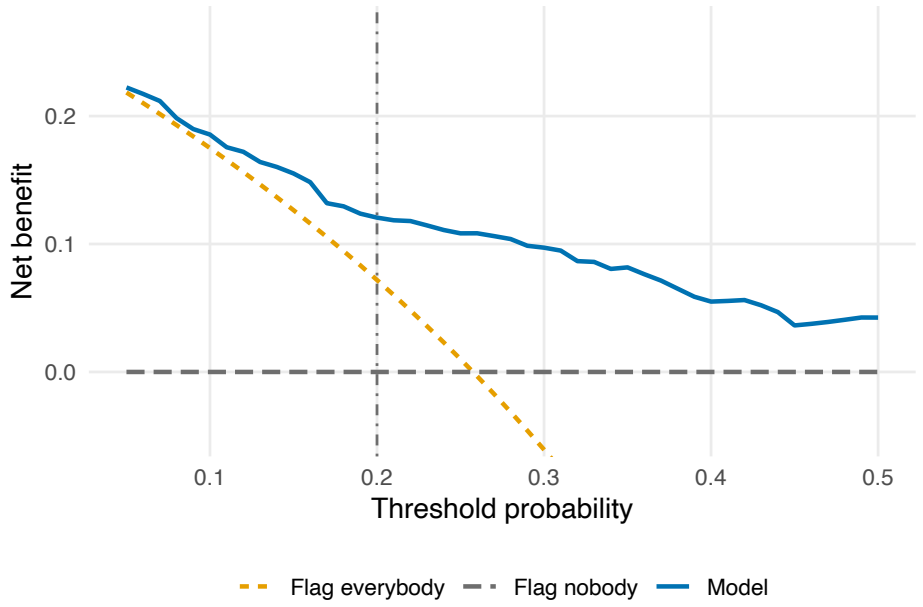


Figure 6.2.: Decision curve for the logistic regression on 400 synthetic test patients (103 with delirium): net benefit of acting on the model compared with flagging everybody or nobody, for thresholds from 5% to 50%. Unit: patient; data status: synthetic.

At 20% the model's net benefit is 0.121: the equivalent of 121 correctly flagged patients per 1,000 with no unnecessary consultations at all. Decision curves and related measures are practised in CTTIR Courses, Course 2 · Decision curves, NRI, IDI.

6.3.8. Step 8 · Two groups, with denominators

```
age_group <- ifelse(test$age >= 80, "80+", "65-79")
grp <- t(sapply(c("65-79", "80+"), function(g) {
  i <- age_group == g
  c(n = sum(i), cases = sum(y[i]), alarms = sum(flag[i]), tp = sum(flag[i] &
    fn = sum(!flag[i] & y[i] == 1), fp = sum(flag[i] & y[i] == 0),
    mean_pred = mean(p[i]), observed = mean(y[i]))
}))
sens_ci <- t(sapply(rownames(grp), function(g) exact_ci(grp[g, "tp"], grp[g,
stopifnot(sum(grp[, "n"]) == 400, sum(grp[, "tp"]) == tab[["tp"]], sum(grp[,
          grp["65-79", "fn"] == 17, grp["80+", "fn"] == 11)
round(cbind(sensitivity = grp[, "tp"] / grp[, "cases"], lower = sens_ci[, 1]
```

```
#>       sensitivity lower upper
#> 65-79        0.32  0.15  0.54
#> 80+          0.86  0.76  0.93
```

In the younger group the sensitivity is 32% with an interval from 15% to 54%, based on only 25 cases. The difference to the older group is large, but the younger group's estimate is imprecise. TRIPOD+AI asks for performance estimates with confidence intervals, including for key subgroups (Collins et al. 2024, item 23a).

6.4. Counter-check

Change one assumption: how delirium is recorded. Suppose that the ward introduces a structured delirium screen twice a day instead of relying on notes. More mild episodes are recorded. We represent this as an *assumption*: the log-odds of recorded delirium rise by 0.8 for every patient,

while the patients and their predictors stay the same. Because the data are synthetic, each patient's true risk is known and the new outcomes can be simulated.

Before you run the code, predict: Does the AUC fall, stay or rise? Does the calibration intercept move, and in which direction? How many patients whose (new) risk is at least 20% will the unchanged 20% threshold miss?

```
set.seed(6061)
risk_b <- plogis(qlogis(test$true_risk) + 0.8)              # true risk under the new record
y_b <- rbinom(length(risk_b), 1, risk_b)                    # recorded delirium, setting B
auc_b <- as.numeric(auc(roc(y_b, p, levels = c(0, 1), direction = "<", quiet = TRUE)))
cal_int_b <- unname(coef(glm(y_b ~ offset(lp), family = binomial))[1])
missed_high_a <- sum(test$true_risk >= threshold & !flag)   # true risk >= 20%, not flagged,
missed_high_b <- sum(risk_b >= threshold & !flag)           # same, setting B
sens_b <- sum(flag & y_b == 1) / sum(y_b)

# Intercept update: estimate the shift on 200 patients, check it on the other 200
half <- seq_along(p) <= 200
shift <- unname(coef(glm(y_b[half] ~ offset(lp[half]), family = binomial))[1])
p_upd <- plogis(lp + shift)
cal_int_upd <- unname(coef(glm(y_b[!half] ~ offset(qlogis(p_upd[!half])), family = binomial)
stopifnot(abs(auc_b - auc_test) < 0.05, cal_int_b > 0.5, missed_high_b > 3 * missed_high_a,
          abs(cal_int_upd) < abs(cal_int_b))
knitr::kable(data.frame(
  quantity = c("Recorded delirium", "AUC", "Calibration intercept", "Sensitivity at 20%",
               "Patients with true risk  20% not flagged"),
  setting_a = c(paste0(sum(y), "/400"), sprintf("%.2f", auc_test), sprintf("%.2f", cal_int),
                cdsr_percent(sens), paste0(missed_high_a, " of ", sum(test$true_risk >= thre
  setting_b = c(paste0(sum(y_b), "/400"), sprintf("%.2f", auc_b), sprintf("%.2f", cal_int_b)
                cdsr_percent(sens_b), paste0(missed_high_b, " of ", sum(risk_b >= threshold)
  col.names = c("Quantity", "Setting A (as developed)", "Setting B (structured screen)"),
  caption = "Counter-check: same model, same 400 synthetic patients, outcome recorded more t
```

Table 6.4.: Counter-check: same model, same 400 synthetic patients, outcome recorded more thoroughly (assumption: +0.8 on the log-odds scale).

Quantity	Setting A (as developed)	Setting B (structured screen)
Recorded delirium	103/400	154/400
AUC	0.77	0.79
Calibration intercept	-0.01	0.77
Sensitivity at 20%	72.8%	70.8%
Patients with true risk 20% not flagged	31 of 194	128 of 310

The ranking survives: the AUC is 0.79 instead of 0.77. The numbers do not. The calibration intercept of 0.77 means that the model systematically underestimates risk in setting B, and 128 patients whose risk now exceeds the 20% that the team agreed on are not flagged — compared with 31 in setting A. A team that monitored only the AUC would notice nothing. Updating the intercept on 200 patients from the new setting and checking it on the other 200 reduces the calibration intercept to -0.19. TRIPOD+AI asks authors to describe such model updating and to report the updated model (Collins et al. 2024, items 12f and 24). A different hospital with a higher baseline risk would produce the same pattern; Chapter 19 and Chapter 20 take this further.

6.5. Consequences for patients and operations

Patients and fairness. The age-group table is the core ethical finding of this chapter. With a single threshold, patients aged 65–79 who will develop delirium are missed far more often than older patients, although the model is calibrated in both groups. Lowering the threshold for the

younger group would find more of their cases at the cost of more consultations without delirium; it is a value decision that should be made openly, with the affected clinicians and, where possible, patient representatives, and not hidden in code (Chapter 21; Strategy in R: Fairness in Machine Learning). A patient who is not flagged must still receive the usual clinical attention; the threshold allocates an additional offer, it does not withdraw care.

Operations. If about eight eligible patients are admitted per week (*assumption*), the 20% threshold produces about 3.6 alarms per week and the 30% threshold about 2.5. The procedure should state the threshold, the expected alarms per week, the expected missed cases per 100 patients and the person who reviews these numbers each month. Because calibration can drift unnoticed, monitoring must include observed versus predicted events, not only alarms and AUC (Chapter 20).

Regulation and documentation. The threshold is part of how the software is used. If a manufacturer's documentation specifies performance at a particular threshold, a local change of threshold changes the conditions under which that claim was made; who may change it, and how this is documented, is a question for the intended purpose and quality management (Chapter 22, *open* for the local setting).

Data protection. The analyses here use aggregated test results. Subgroup tables with very small cells can make individuals identifiable when published; the rules in Chapter 26 apply before such tables leave the team.

6.6. Look up and transfer

Reviewer question: "Do the authors report calibration (plot, intercept and slope) in the population where the model is meant to be used, and do they translate their chosen threshold into absolute numbers of flagged,

found and missed patients — with denominators and intervals for key subgroups?" TRIPOD+AI asks for the measures and plots used to evaluate discrimination, calibration and clinical utility, and for subgroup estimates with confidence intervals (Collins et al. 2024, items 12e and 23a).

Self-test: A colleague reports that a delirium model "performs well" in a rehabilitation clinic: AUC 0.80. In this clinic delirium occurs in 8% of patients, but the mean predicted risk is 24%. The clinic plans to use a 20% threshold. What do you expect for the number of alarms and the positive predictive value, and what would you ask for before the threshold is adopted?

Answer

The model overestimates risk by a factor of about three on average (calibration-in-the-large is poor), so a 20% threshold on these inflated numbers will flag many patients whose true risk is far below 20%. Expect a high share of flagged patients and a low positive predictive value — at an 8% base rate, most alarms will concern patients without delirium. The AUC of 0.80 does not contradict this, because it ignores the absolute level of the predictions. Before adoption, ask for a calibration plot with intercept and slope in this clinic, a recalibration (at least of the intercept) on local data, and a table of flagged, found and missed patients per 100 at the proposed threshold.

Annotated sources:

- Van Calster et al. (2019) explains why poor calibration can make an algorithm misleading for decisions and how to assess calibration at validation.
- Steyerberg et al. (2010) sets out the framework of overall performance, discrimination, calibration and decision-analytic measures used in this chapter.
- Vickers and Elkin (2006) introduces decision curve analysis and the link between threshold probability and the weighing of false-positive

and false-negative decisions.
- DeLong et al. (1988) and Robin et al. (2011) are the method and software for the analytical AUC interval in Step 4.
- Hoffrage and Gigerenzer (1998) shows that natural frequencies improve diagnostic reasoning compared with probabilities, supporting the per-1,000 translation.
- Collins et al. (2024) provides the reporting items for performance measures, subgroup estimates and model updating.

6.7. Further reading

- Chapter 5 — how the model, the split and the test patients used here were built.
- Chapter 9 — the same 2 × 2 logic for a diagnostic alert at two different prevalences.
- Chapter 21 — who bears the consequences of missed cases and false alarms, and how to discuss them with patients.
- Methods in R: Week 3, Session 5 — Calibration, discrimination, ROC/AUC, Brier score — a complete lab on the measures of Steps 3 to 5.
- CTTIR Courses, Course 2 · Calibration, discrimination, ROC/AUC, Brier score — the same topic as a course session with exercises.
- CTTIR Courses, Course 4 · Time-dependent Brier, IPA, external validation — performance measures for external validation.
- CTTIR Tutorials: Platt Scaling — recalibration methods beyond the intercept update of the counter-check.
- Methods in R: Cutpoint Selection — why data-driven "optimal" cutoffs need the same caution as data-driven model choice.

7. Reading an AI study critically

R. Heller · Date: 2026-07-26 · Last changed: 2026-11-09

At a glance

Guiding question: Which findings support measurement quality, and which support better care?
For: clinicians preparing a journal club, members of committees that decide on clinical software, doctoral researchers and peer reviewers · **Prerequisites:** Chapter 5 (validation designs) and Chapter 6 (2 × 2 table, calibration, thresholds)
In the course: W2 · Language models and evidence and W3 · Diagnostics, prediction and evidence · **Data:** *published* figures from one external validation study and a *synthetic* (fictitious) abstract; no patient data · **Core reading time:** about 11 minutes

7.1. Case and decision

The committee for clinical decision support of a fictitious hospital meets in two weeks. On the agenda is a request to switch on a sepsis warning that comes with the hospital's electronic health record and, according to the supplier, is used in many hospitals. A resident who runs the internal medicine journal club is asked to prepare a one-page brief. The committee wants to know two things: how well does the score recognise sepsis, and does its use lead to better care for patients? The resident finds one

independent external validation of a widely implemented proprietary sepsis model, a set of reporting guidelines with unfamiliar acronyms and a supplier brochure. The recommendation — switch on, pilot under observation, or decline — will change the work of every ward. A brief that confuses accuracy with benefit could commit the hospital to thousands of alerts without any evidence that patients gain.

7.2. Core reading

Core reading · required for W2 · about 11 minutes · ends before 'Lab'

A ladder of evidence. Studies of medical tests and models answer questions at different levels. Fryback and Thornbury described a hierarchy for diagnostic imaging that still serves well: technical quality, diagnostic accuracy, effect on the clinician's diagnostic thinking, effect on management, effect on patient outcomes, and societal costs and benefits. Evidence at a lower level is necessary but not sufficient for a claim at a higher level (Fryback and Thornbury 1991). An AI system climbs the same ladder. A *reader study* or *diagnostic accuracy study* compares the system's output with a reference standard. An *external validation* applies a prediction model to patients from another place or time and measures discrimination and calibration there (Chapter 6). An *early clinical evaluation* places the system into a real workflow on a small scale and asks how clinicians use it, when they override it and what goes wrong. Only a *comparative study*, ideally randomised, can show whether patients are better off with the system than without it.

The higher rungs are sparsely populated. A systematic review of deep learning studies in medical imaging that compared algorithms with clinicians found only 10 records of randomised trials, two of them published; of 81 non-randomised studies only nine were prospective, 58 were at high

risk of bias, and 61 stated in their abstract that the algorithm performed at least as well as clinicians (Nagendran et al. 2020). Claims in abstracts can therefore run ahead of the design.

Four tools, four tasks. Reporting guidelines list what a paper must report so that a reader can judge it; an appraisal tool guides the judgement itself. Table 7.1 assigns each of the four AI-specific tools to one concrete check. The item numbers refer to the published checklists.

Table 7.1.: Four tools for four questions about clinical AI. For randomised trials of AI interventions, CONSORT-AI adds 14 items to CONSORT 2010 (Liu et al. 2020).

Study type and question	Tool	One concrete check	What the study cannot show
Diagnostic accuracy: does the output agree with the reference standard?	STARD-AI, reporting (Sounderajah et al. 2025)	How were the data partitioned into training, validation and test sets (item 15b), and does the test set represent the target condition and population (items 25, 28)?	Whether clinicians decide differently or patients do better

Study type and question	Tool	One concrete check	What the study cannot show
Prediction model development or evaluation: how well does the model discriminate and calibrate?	TRIPOD+AI, reporting (Collins et al. 2024)	Are all measures for discrimination, calibration and clinical utility specified (item 12e) and reported with confidence intervals, including key subgroups (item 23a)?	Whether acting on the prediction changes outcomes
Same studies, appraised: can the result be trusted and does it apply here?	PROBAST+AI, quality, risk of bias and applicability (Moons et al. 2025)	Model evaluation, domain 4: was data leakage avoided (4.5), were all development steps repeated in resampling (4.6), were calibration, discrimination and net benefit evaluated (4.7)?	Anything beyond the population and setting of the evaluation

Study type and question	Tool	One concrete check	What the study cannot show
Early clinical evaluation in live use: how is the system used, and is it safe?	DECIDE-AI, reporting (Vasey et al. 2022)	Is user agreement reported, with instances of and reasons for variation from the recommendation (item 12), and are errors and risks to patient safety listed (item 13)?	Comparative effectiveness

A missing item is not a verdict. It is a question to the authors, the supplier or the local team. If a paper on a prediction model reports an AUC but no calibration, the right response is "How do predicted and observed risks compare in your test set?", not silent acceptance or silent rejection. Such questions can go to the authors through the journal, to the supplier as part of procurement, or to the local team as the first item of a plan for a silent evaluation in the hospital's own patients. The answer may change the verdict in either direction.

Five questions for the journal club. Whatever the study type, five questions organise the reading, and the lab turns them into a card. *Recruitment:* who was included, from where and when, and who was left out? A test set collected from patients with clear-cut disease and clear-cut health tends to overstate accuracy for the mixed patients of everyday care. *Reference:* how was the target condition or outcome established, by whom, and without knowledge of the model's output? Discharge codes, for example, capture only episodes that were documented and coded. *Analysis:* were tuning, model choice and testing kept apart (Chapter 5), and are

discrimination, calibration and a threshold reported with denominators (Chapter 6)? *Generalisation:* does the evaluation population resemble the patients on whom the model would be used here — case mix, measurement devices, time, care pathway? *Decision:* does the conclusion stay at the level the design reached, or does an accuracy study recommend implementation "to improve outcomes"? Each question can be answered from the paper, from a supplement, or not at all; the last answer is itself a finding.

One statement, checked against its source. The committee brief will contain numbers, and each number needs a source that can be checked. Here is one statement, traced to its origin.

Table 7.2.: A checked statement: claim, original, locator, unit and verdict. Numbers are *published* (Wong et al. 2021).

Element	Content
Statement in the brief	In an independent external validation, the sepsis model discriminated poorly (AUC 0.63) and missed two thirds of sepsis cases while alerting in almost one in five hospitalisations.
Original study	Wong et al. (2021), retrospective cohort at one US academic health system, December 2018 to October 2019

Element	Content
Locator	Abstract, “Results”: AUC 0.63 (95% CI 0.62–0.64) at hospitalisation level; the model “did not identify 1709 patients with sepsis (67%) despite generating alerts for an ESM score of 6 or higher for 6971 of all 38 455 hospitalized patients (18%)”
Unit and denominator	38,455 hospitalisations of 27,697 patients; 2,552 with sepsis (7%)
Data status	*published*; abstract read in PubMed, DOI checked against Crossref; the article carries a correction (2 August 2021) of the number needed to evaluate in Table 2 and the Results, so only the corrected version may be quoted for that number
Verdict	Supported for measurement quality at this centre. Not evidence about patient outcomes, nor about other hospitals or later versions of the model

The lab reconstructs the 2×2 table behind this statement: a sensitivity of about 33%, a positive predictive value of about 12%, and roughly 8 hospitalisations with an alert for each hospitalisation with sepsis that the alert caught. The study also compared the score with contemporary care: it identified 183 of the 2,552 patients with sepsis (7%) who did not receive timely antibiotics. The authors read this as low sensitivity compared with contemporary clinical practice: the score seldom flagged sepsis that clini-

cians had not already treated in time. That is a question one level above accuracy on the Fryback–Thornbury ladder — whether the information changes what clinicians do.

The most important assumption when reading any AI study is that a study supports only the claim its design can answer: accuracy in a test population is not evidence of better care, and better care in one hospital is not evidence for another. The supplier's statement that the model is "used in many hospitals" is evidence of dissemination, not of benefit.

Question for the workshop: The committee asks for a single sentence. Write one sentence about the sepsis warning that is fully supported by Table 7.2, and one sentence that the brochure might contain but that the study cannot support. Which study type would be needed for the second sentence?

7.3. Lab

7.3.1. Part A · Reconstruct the 2 × 2 table from the published abstract

All inputs are *published* counts from the abstract of Wong et al. (2021). Combining them into one table rests on an *assumption* that must be stated: the counts refer to the same unit (hospitalisation), the same alert threshold (score 6 or higher) and the same time window. The abstract places them in one sentence but does not say this explicitly; the full text (Table 2, corrected version) is where a reader would confirm it.

```
n_hosp <- 38455          # hospitalisations analysed (published)
n_sepsis <- 2552         # hospitalisations with sepsis (published)
missed <- 1709           # sepsis not identified by the alert (published, 67%
alerted <- 6971          # hospitalisations with an alert, score >= 6 (publis
```

```
tp <- n_sepsis - missed                       # sepsis with alert
fp <- alerted - tp                            # alert without sepsis
fn <- missed                                  # sepsis without alert
tn <- n_hosp - n_sepsis - fp                  # neither
sens <- tp / n_sepsis
spec <- tn / (n_hosp - n_sepsis)
ppv <- tp / alerted
npv <- tn / (tn + fn)
alerts_per_case <- alerted / tp               # alerts per sepsis case caught

# Consistency with the percentages printed in the abstract
stopifnot(tp + fp + fn + tn == n_hosp,
          round(100 * missed / n_sepsis) == 67, round(100 * alerted / n_hosp) == 18,
          round(100 * n_sepsis / n_hosp) == 7, tp == 843)
round(c(sensitivity = sens, specificity = spec, ppv = ppv, npv = npv,
        alerts_per_case = alerts_per_case), 3)
```

```
#>     sensitivity     specificity             ppv             npv alerts_per_case
#>           0.330           0.829           0.121           0.946           8.269
```

```
two_by_two <- data.frame(
  row = c("Alert (score  6)", "No alert", "Total"),
  sepsis = c(tp, fn, n_sepsis),
  no_sepsis = c(fp, tn, n_hosp - n_sepsis),
  total = c(alerted, fn + tn, n_hosp))
knitr::kable(two_by_two, col.names = c("", "Sepsis", "No sepsis", "Total"),
             format.args = list(big.mark = ","), row.names = FALSE)
```

Table 7.3.: Reconstructed 2 × 2 table for the sepsis alert (score 6) at the hospitalisation level, from the published abstract of Wong et al. (2021). Unit: hospitalisation; n = 38,455 from 27,697 patients; counts derived from published figures under the stated assumption.

	Sepsis	No sepsis	Total
Alert (score 6)	843	6,128	6,971
No alert	1,709	29,775	31,484
Total	2,552	35,903	38,455

Read the result as a committee member would. Of 1,000 hospitalisations, about 181 receive an alert; about 22 of these involve sepsis and about 159 do not; about 44 hospitalisations with sepsis receive no alert. The unit matters: 38,455 hospitalisations come from 27,697 patients, so some patients contribute more than one hospitalisation, and the hospitalisations are not fully independent observations (Chapter 19).

7.3.2. Part B · A journal-club card for a fictitious abstract

The second exercise uses a *synthetic* abstract written for this book. It describes no real study, product or hospital.

Fictitious abstract. *Machine learning predicts postoperative delirium after hip fracture.* We developed a gradient-boosting model using 46 routinely collected variables from 2,140 admissions of patients aged 65 or older at two hospitals (2019–2023). Delirium was identified from discharge diagnosis codes. Admissions were randomly split 80 : 20 into training and test sets; hyperparameters were optimised on the test set. The model achieved an AUC of 0.89 (95% CI 0.86–0.92) and outperformed a clinical points rule (AUC 0.71). Feature

> importance identified postoperative opioid dose, age and dementia as the strongest predictors. The model reliably identifies patients at risk and should be implemented to reduce delirium.

The card has five rows — recruitment, reference (here: outcome definition), analysis, generalisation and decision. Each row names the tool item it draws on and turns a missing or doubtful item into a question.

```
card <- data.frame(
  domain = c("Recruitment", "Reference / outcome", "Analysis", "Analysis", "Generalisation",
  abstract_says = c("2,140 admissions, two hospitals, 2019-2023",
                    "delirium from discharge diagnosis codes",
                    "random 80:20 split of admissions",
                    "hyperparameters optimised on the test set; AUC only",
                    "test set from the same two hospitals and period",
                    "'should be implemented to reduce delirium'"),
  item = c("PROBAST+AI 1.3 (representative inclusions)",
           "PROBAST+AI 3.1 (outcome defined and assessed appropriately)",
           "PROBAST+AI 4.5 (data leakage avoided); 2.4 (predictors available at prediction t
           "TRIPOD+AI 12c, 12e; PROBAST+AI 4.7 (calibration, net benefit)",
           "STARD-AI 28 / TRIPOD+AI 23a (target population, subgroups)",
           "DECIDE-AI; CONSORT-AI (impact needs live use and comparison)"),
  rating = c("unclear", "doubtful", "doubtful", "not reported", "not reported", "not support
card$question <- ifelse(card$rating == "reported", "",
  c("How many eligible admissions were excluded, and why?",
    "How sensitive are discharge codes for delirium compared with a structured bedside asses
    "Were repeat admissions of the same patient kept on one side of the split? Was the posto
    "What was the performance on data not used for tuning, and how well calibrated are the p
    "How does the model perform in another hospital, a later period and in key subgroups?",
    "Which study showed that acting on the model reduces delirium?"))
n_questions <- sum(card$question != "")
stopifnot(n_questions == nrow(card))        # every row of this abstract triggers a question
knitr::kable(card, col.names = c("Domain", "Abstract says", "Tool item", "Rating", "Question
```

```
                row.names = FALSE)
```

Table 7.4.: Journal-club card for the fictitious delirium abstract. Every row rated other than ‘reported’ produces a question to the authors.

Domain	Abstract says	Tool item	Rating	Question to the authors
Recruitment	2,140 admissions, two hospitals, 2019-2023	PROBAST+AI 1.3 (representative inclusions)	unclear	How many eligible admissions were excluded, and why?
Reference / outcome	delirium from discharge diagnosis codes	PROBAST+AI 3.1 (outcome defined and assessed appropriately)	doubtful	How sensitive are discharge codes for delirium compared with a structured bedside assessment?
Analysis	random 80:20 split of admissions	PROBAST+AI 4.5 (data leakage avoided); 2.4 (predictors available at prediction time)	doubtful	Were repeat admissions of the same patient kept on one side of the split? Was the postoperative opioid dose known at the time of prediction?
Analysis	hyperparameters optimised on the test set; AUC only	TRIPOD+AI 12c, 12e; PROBAST+AI 4.7 (calibration, net benefit)	not reported	What was the performance on data not used for tuning, and how well calibrated are the predicted risks?

Table 7.4.: Journal-club card for the fictitious delirium abstract. Every row rated other than ‘reported’ produces a question to the authors.

Domain	Abstract says	Tool item	Rating	Question to the authors
Generalisation	data from the same two hospitals and period	STARD-AI 28 / TRIPOD+AI 23a (target population, subgroups)	not reported	How does the model perform in another hospital, a later period and in key subgroups?
Decision	‘should be implemented to reduce delirium’	DECIDE-AI; CONSORT-AI (impact needs live use and comparison)	not supported	Which study showed that acting on the model reduces delirium?

Two findings stand out. Tuning on the test set removes the independence of the test, so the AUC of 0.89 is an optimistic estimate (Chapter 5). And postoperative opioid dose cannot be known on the first morning after admission: if the model is meant to trigger prevention before surgery, one of its strongest predictors is not available at the time of use (PROBAST+AI, signalling question 2.4). The concluding sentence makes an impact claim from an accuracy study.

7.4. Counter-check

Change the target population. Suppose the committee considers the alert only for an intensive care unit, where sepsis is far more common — say 20% of stays instead of 7% (*assumption*).

Before you compute, predict: With the published sensitivity and specificity, does the positive predictive value rise or fall? And what happens to the number of alerts per sepsis case caught?

```
ppv_at <- function(prev, se = sens, sp = spec) se * prev / (se * prev + (1 -
prevalences <- c(published = n_sepsis / n_hosp, icu_assumed = 0.20)
cc <- data.frame(setting = names(prevalences), prevalence = prevalences,
                 ppv = ppv_at(prevalences), alerts_per_case = 1 / ppv_at(prev
stopifnot(abs(cc$ppv[1] - ppv) < 1e-9, cc$ppv[2] > cc$ppv[1])
knitr::kable(cc, digits = 3, row.names = FALSE,
             col.names = c("Setting", "Prevalence", "PPV", "Alerts per case c
             caption = "Counter-check: published sensitivity and specificity
```

Table 7.5.: Counter-check: published sensitivity and specificity applied to an assumed higher prevalence.

Setting	Prevalence	PPV	Alerts per case caught
published	0.066	0.121	8.269
icu_assumed	0.200	0.326	3.067

The positive predictive value rises to about 33% and the alerts per case caught fall to about 3.1. The calculation, however, assumes that sensitivity and specificity carry over unchanged to intensive care. That is rarely true: patients there are sicker, measured more often and treated earlier, so the score's behaviour can change in either direction. The counter-check therefore yields a hypothesis for a local evaluation, not a result (Methods in R: Predictive Values and Prevalence).

7.5. Consequences for patients and operations

Patients. Two harms sit on either side of a sepsis alert. A missed case may delay treatment, although clinicians often recognise sepsis without

the alert, as the comparison with timely antibiotics in Wong et al. (2021) shows. A false alert can lead to unnecessary tests, antibiotics and interruptions, and many false alerts teach staff to ignore all of them. Which harm dominates depends on how the alert is embedded in the workflow — a question for an early clinical evaluation (DECIDE-AI), not for the accuracy study.

Operations. A brief for the committee should give the expected number of alerts per ward and day, the person who responds to an alert, and the rule for stopping the pilot (Chapter 20). Local data from a silent, non-alerting phase can show discrimination and calibration in the hospital's own population before any alert reaches a clinician.

Regulation. Software intended to support the diagnosis of sepsis in individual patients is typically assessed as medical device software under the Medical Device Regulation (European Parliament and Council of the European Union 2017a); the qualification guidance MDCG 2019-11 Rev. 1 is orientation, not binding law (Medical Device Coordination Group 2025). Under Article 6(1) of the AI Act (European Parliament and Council of the European Union 2024b), an AI system is high-risk if it is a safety component of a product, or itself a product, covered by the legislation listed in Annex I — which includes the Medical Device Regulation — *and* that product has to undergo a third-party conformity assessment. After the amendment by Regulation (EU) 2026/1744 (European Parliament and Council of the European Union 2026), which entered into force on 27 July 2026, Chapter III, Sections 1 to 3 apply to such systems from 2 August 2028 (Article 113, as amended; read on EUR-Lex in the English text of the amendment; status: November 2026). Whether a particular warning function is a medical device, and who is its manufacturer, depends on its intended purpose and is a question for the hospital's regulatory function (*open*; Chapter 22).

Data protection. A local validation uses patient data for a purpose other than treatment. Its legal basis, minimisation and access rules must be settled before data are extracted (Chapter 23, Chapter 26).

7.6. Look up and transfer

Reviewer question: "For each claim in your abstract, which result in which table supports it, at which level of evidence — accuracy, effect on decisions, or effect on patient outcomes — and in which population?"

Self-test: A paper reports that after an AI fracture-detection tool was introduced in an emergency department, the rate of missed fractures in the following six months was lower than in the six months before. The authors conclude that the tool improves patient outcomes. Which level of evidence does the study reach, and what is the main alternative explanation?

Answer

It is an uncontrolled before–after comparison of a process measure (missed fractures), so it sits between accuracy and effect on management, and it does not measure patient outcomes such as complications or function. The main alternative explanations are secular trends and co-interventions (new staff, training, awareness after the introduction), and changes in how missed fractures were detected or recorded. A concurrent comparison — ideally a cluster-randomised or stepped-wedge design reported according to CONSORT-AI — and a patient-relevant outcome would be needed. An early clinical evaluation reported according to DECIDE-AI would describe how clinicians used the tool.

Annotated sources:

- Fryback and Thornbury (1991) defines the hierarchy of diagnostic efficacy that structures the core reading.
- Nagendran et al. (2020) documents how few prospective and randomised studies exist for deep learning in imaging and how often abstracts claim parity with clinicians.

- Sounderajah et al. (2025) (STARD-AI; an author correction of July 2026 concerns only the consortium list), Collins et al. (2024) (TRIPOD+AI), Moons et al. (2025) (PROBAST+AI) and Vasey et al. (2022) (DECIDE-AI) are the four tools assigned to concrete checks in Table 7.1.
- Liu et al. (2020) is the reporting extension for randomised trials of AI interventions, the top rung of the ladder.
- Wong et al. (2021) is the external validation from which the checked statement and the reconstructed 2×2 table come.
- The AI Act (European Parliament and Council of the European Union 2024b) and its 2026 amendment (European Parliament and Council of the European Union 2026) define the high-risk timetable cited in the consequences; the Medical Device Regulation (European Parliament and Council of the European Union 2017a) and MDCG 2019-11 (Medical Device Coordination Group 2025) the device question.

7.7. Further reading

- Chapter 6 — the measures (discrimination, calibration, net benefit) that the checks in Table 7.1 ask for.
- Chapter 20 — how a silent phase, a pilot and monitoring turn an accuracy result into evidence about use.
- Chapter 27 — the model passport that collects these answers for one's own model.
- Methods in R: TRIPOD-AI — prediction models — a reporting template to compare a paper against.
- Methods in R: CONSORT — randomised trials — the core items that CONSORT-AI extends.
- CTTIR Courses, Course 4 · TRIPOD-AI, fairness auditing, reproducibility at scale — a lab that applies TRIPOD-AI to a model.

- CTTIR Tutorials: Diagnostic Test Accuracy — the accuracy measures behind Part A.
- Scientometrics in R: Retraction Analysis — checking corrections and retractions before a paper is cited.

8. Language models in clinic and science

R. Heller · Date: 2026-08-23 · Last changed: 2026-11-09

At a glance

Guiding question: Which statement can be checked after generation?
For: clinicians, researchers, teachers and students who use, review or publish text produced with language models · **Prerequisites:** Chapter 7 (tracing a statement to its source); no programming experience is needed for the core reading
In the course: W2 · Language models and evidence · **Data:** synthetic (two stored answers and a twelve-sentence corpus written for this book; teaching file `data/b08-stored-answers.csv`); *published* findings only where cited · **Core reading time:** about 11 minutes

8.1. Case and decision

The daughter of a 78-year-old woman writes through the patient portal of a fictitious trauma department. Her mother has a hip fracture and will be operated on tomorrow. “Can anything be done so that she does not get confused after the operation? Should she take melatonin?” The resident on duty uses the hospital’s writing assistant, which stores two answers to the same question. Both sound competent, both cite literature, and they

disagree. The resident has twenty minutes before the evening round. Every sentence that goes into the reply will be read as the hospital's advice and may shape what the family does tonight. The decision is which generated sentences, if any, may be used — and the resident is entitled to use none of them.

8.2. Core reading

Core reading · required for W2 · about 11 minutes · ends before 'Lab'

How a language model writes. A language model splits text into *tokens* — words or pieces of words — and learns, from very large collections of text, how likely each token is to follow the tokens before it. To write, it repeatedly picks a next token from this probability distribution and appends it. Modern models compute these probabilities with the transformer architecture (Vaswani et al. 2017), which lets every token draw on a long stretch of preceding text (Methods in R: Transformers). The principle can be seen in a model that is tiny by comparison: a *bigram model*, which predicts the next word only from the current word, using counts from a corpus (Jurafsky and Martin 2026, chap. 3). The lab builds one from twelve synthetic sentences about delirium. One of them states that melatonin does *not* prevent delirium after hip surgery; others say that various measures help prevent delirium after hip surgery, and that melatonin helps patients sleep. Following the most frequent continuations, the model writes: "melatonin helps prevent delirium after hip surgery". Every word pair in this sentence occurs in the corpus; the sentence as a whole does not, and it reverses the only statement the corpus makes on the question.

This is the central point of the chapter. The model is trained to produce *probable* continuations, not *true* ones. Plausibility and truth often coincide, which is why language models are useful, but nothing in the mechanism

checks the second. Fluent, confident and wrong text is a known property of such systems, discussed as "hallucination" (Ji et al. 2023).

Temperature. Before sampling, the probabilities can be reshaped. Dividing the model's internal scores by a *temperature* below 1 sharpens the distribution towards the most probable tokens; a temperature above 1 flattens it. In general, emphasising the most probable tokens gives text that people rate as more accurate and coherent but also more repetitive, while more weight on less probable tokens gives more diverse but less factual text (Jurafsky and Martin 2026, sec. 7.6.3); the decoding strategy alone can change the quality of generated text markedly (Holtzman et al. 2019). But probability is not truth. Figure 8.1 shows what happens in our small model, where the most frequent continuation is the false one: a low temperature makes the error more likely, not less.

Two different failures with sources. Generated text often carries references, and they fail in two distinct ways. A reference can be *fabricated*: the authors, title or journal do not exist. In one study of 636 references produced by two versions of a widely used chatbot for short literature reviews, 55% of the references from the older version and 18% of those from the newer version were fabricated, and many of the real ones contained substantive errors (Walters and Wilder 2023). A reference can also be *real but not supporting*: the paper exists, is on the right topic, and says something different from the sentence it is attached to. The second failure is harder to catch, because a quick check of title and DOI succeeds. Systems that first *retrieve* documents and then generate an answer from them (Lewis et al. 2020) reduce the first failure but not the second; a retrieved paper can still be summarised wrongly.

The two checks must therefore be kept apart. *Existence* is bibliographic: does the DOI resolve, do authors, title, journal and year match? *Support* is substantive: does the cited passage say what the sentence claims, for the population and outcome in question? Table 8.1 applies both checks to the five claims in the two stored answers.

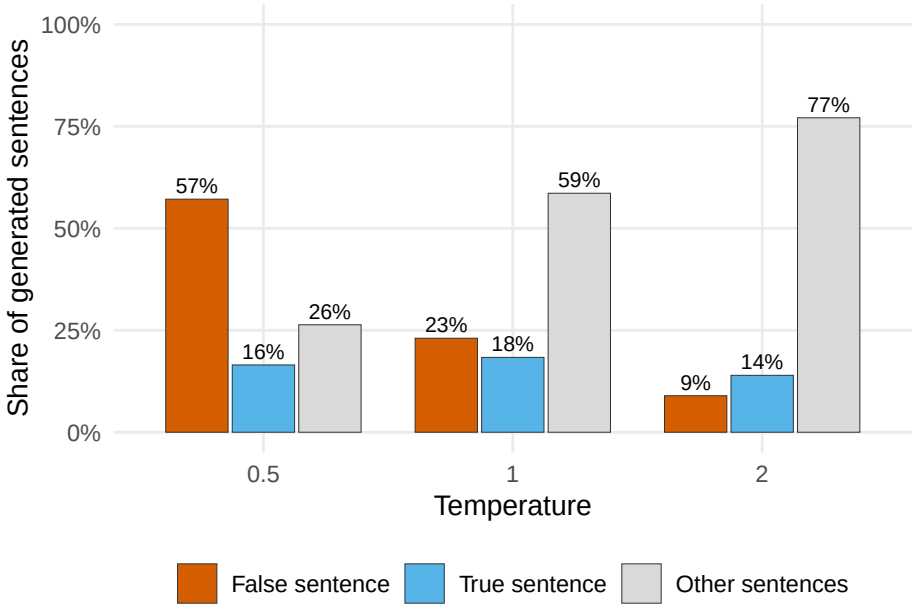


Figure 8.1.: Sentences generated by the bigram model from the start word "melatonin" at three temperatures (2,000 generated sentences per temperature). Bars: share that reproduces the true corpus sentence, share that produces the false sentence "melatonin helps prevent delirium after hip surgery", and all other sentences. Unit: generated sentence; data status: synthetic.

Table 8.1.: Evidence card for the two stored answers: existence and support are judged separately; the action follows from both. Data status: the answers are synthetic; the findings named under 'support' are published.

Claim (answer)	Source given	Source exists?	Source supports claim?	Verdict	Action in the reply
A1 Melatonin reliably prevents postoperative delirium	de Jonghe et al. 2014	yes	no: no reduction in incidence (29.6% vs 25.5%)	real source, contradicts claim	reject
A2 Glasses, hearing aids and family visits help orientation	none	not applicable	not assessable without a source	unsourced	adapt only after finding a source
B1 Measures for orientation, sleep, vision, hearing, fluids and mobility reduce delirium in older hospital patients	Inouye et al. 1999	yes	partly: 852 general-medicine patients aged 70+, not surgical	supported with caveat	adapt: name the population
B2 Daily geriatrician visits reduced delirium from 50% to 32% in hip fracture patients	Marcantonio et al. 2001	yes	yes: 20/62 vs 32/64 patients, one centre, 126 patients	supported	adopt in plain language

Table 8.1.: Evidence card for the two stored answers: existence and support are judged separately; the action follows from both. Data status: the answers are synthetic; the findings named under 'support' are published.

Claim (answer)	Source given	Source exists?	Source supports claim?	Verdict	Action in the reply
B3 Lavender aromatherapy halves the risk of delirium after hip surgery	Hartmann and Weiss 2021	no	not assessable	fabricated reference	reject

The melatonin claim in answer A is the instructive case. The cited trial exists and is exactly on topic: older patients with hip fracture, melatonin or placebo, delirium as primary outcome. Its result was that melatonin did not reduce the incidence of delirium — 29.6% with melatonin versus 25.5% with placebo (de Jonghe et al. 2014). The generated sentence turns this result into its opposite. Nothing in this misuse is the trial authors' doing; the error lies entirely in the generated sentence. The lavender claim in answer B cites a paper and a journal that could not be found in Crossref or the NLM Catalog, and its DOI does not resolve; it was invented for this book to show what a fabricated reference looks like. The two remaining claims of answer B are supported, one of them with a caveat about the population studied.

The most important assumption for using generated text is that a generated sentence is a hypothesis about the literature, not a report of it; it becomes usable only after its source has been found and read. In clinical communication, a further condition applies: the person who signs the letter is responsible for every sentence in it, whatever produced the first draft. Benchmarks of clinical knowledge show that models can answer many examination-style questions while human

evaluation still finds important gaps (Singhal et al. 2023); such results describe average behaviour, not the reliability of the sentence in front of you.

Question for the workshop: For each of the five claims in Table 8.1, decide whether you would adopt, adapt or reject it in the reply to the daughter. Is there a claim you would reject even though its source is real and supportive — and is rejecting the whole generated text a defensible choice?

8.3. Lab

8.3.1. Step 1 · Store the two answers with their provenance

The two answers below are *synthetic*. They were written for this book to reproduce typical patterns of generated text; they are not output of any language model, and the "settings" describe a fictitious assistant. Each answer is stored with a SHA-256 checksum so that later readers can confirm that the text they assess is the text that was assessed.

```
library(digest)
question <- paste("My mother is 78 and will have surgery for a hip fracture tomorrow.",
                  "Can anything be done so that she does not get confused after the operatio
                  "Should she take melatonin?")
answer_a <- paste(
  "Confusion after surgery, called delirium, is common in older adults with a hip fracture."
  "The most effective step is melatonin: a randomised trial in hip fracture patients showed
  "melatonin reliably prevents postoperative delirium (de Jonghe et al., CMAJ 2014).",
  "Your mother should take melatonin every evening, starting the evening before surgery.",
  "Bringing her glasses and hearing aids and having family visit can also help her stay orie
answer_b <- paste(
  "Confusion after an operation, called delirium, affects many older patients.",
  "Measures that address orientation, sleep, vision, hearing, fluids and mobility have been
```

```
  "to reduce delirium in older hospital patients (Inouye et al., NEJM 1999).",
  "In a randomised trial in hip fracture patients, daily visits by a geriatri
  "from 50% to 32% (Marcantonio et al., J Am Geriatr Soc 2001).",
  "A 2021 study also found that aromatherapy with lavender oil halves the ris
  "surgery (Hartmann and Weiss, Journal of Orthogeriatric Nursing 2021; doi:1
  "Please ask the ward team which of these measures they offer.")

stored <- data.frame(
  answer = c("A", "B"),
  setting_described = c("fictitious assistant, temperature 0.2", "fictitious
  question = question,
  text = c(answer_a, answer_b),
  data_status = "synthetic; written for this book; not produced by a language
stored$sha256 <- vapply(stored$text, digest, character(1), algo = "sha256", s
write.csv(stored, "../data/b08-stored-answers.csv", row.names = FALSE)
stopifnot(nrow(read.csv("../data/b08-stored-answers.csv")) == 2,
          all(nchar(stored$sha256) == 64))
substr(stored$sha256, 1, 16)        # first 16 characters of each checksum
```

```
#> [1] "79592eb3bc08717b" "06eb23553cd7c8ca"
```

8.3.2. Step 2 · A bigram model from twelve sentences

```
corpus <- c(
  "delirium after hip surgery is common in older patients",
  "geriatric visits help prevent delirium after hip surgery",
  "early mobilisation helps prevent delirium after hip surgery",
  "regular sleep helps prevent delirium in older patients",
  "melatonin does not prevent delirium after hip surgery",
  "melatonin helps some patients sleep",
  "melatonin helps older patients sleep",
```

```
  "glasses and hearing aids help patients stay oriented",
  "family visits help patients stay oriented",
  "delirium is often missed in older patients",
  "pain control helps prevent delirium after hip surgery",
  "older patients often sleep poorly in hospital")

# Tokens are words here; <s> and </s> mark the start and end of a sentence
bigram_table <- function(sentences) {
  pairs <- do.call(rbind, lapply(strsplit(sentences, " "), function(w) {
    w <- c("<s>", w, "</s>")
    data.frame(from = head(w, -1), to = tail(w, -1))
  }))
  table(pairs$from, pairs$to)                      # counts: rows = current word, columns =
}
counts <- bigram_table(corpus)

# Next-word probabilities, optionally reshaped by a temperature T: p^(1/T), renormalised.
# This equals dividing the log-probabilities ("logits") by T before the softmax.
next_word <- function(word, temperature = 1, cnt = counts) {
  k <- cnt[word, ]; k <- k[k > 0]
  p <- k / sum(k)
  q <- p^(1 / temperature)
  q / sum(q)
}
greedy <- function(start, cnt = counts, max_len = 15) {   # always take the most probable ne
  out <- start; w <- start
  repeat {
    p <- next_word(w, cnt = cnt); w <- names(p)[which.max(p)]
    if (w == "</s>" || length(out) >= max_len) break
    out <- c(out, w)
  }
  paste(out, collapse = " ")
}
```

```
greedy_sentence <- greedy("melatonin")
false_sentence <- "melatonin helps prevent delirium after hip surgery"
true_sentence <- "melatonin does not prevent delirium after hip surgery"
stopifnot(greedy_sentence == false_sentence, !(false_sentence %in% corpus), t
round(next_word("melatonin"), 2); round(next_word("helps"), 2)
```

```
#>  does helps
#>  0.33  0.67
```

```
#>   older prevent    some
#>     0.2     0.6     0.2
```

The model's probability for the false sentence (start word given) is the product of its word-to-word probabilities:

```
sentence_prob <- function(s, cnt = counts) {
  w <- c(strsplit(s, " ")[[1]], "</s>")
  prod(mapply(function(a, b) next_word(a, cnt = cnt)[b], head(w, -1), tail(w
}
probs <- c(false = sentence_prob(false_sentence), true = sentence_prob(true_s
stopifnot(probs[["false"]] > probs[["true"]])
round(probs, 3)
```

```
#> false  true
#> 0.229 0.190
```

The false sentence is more probable than the true one, because "melatonin helps" and "helps prevent" are frequent pairs in the corpus. Nothing in the counts encodes which statement about melatonin and delirium is correct.

8.3.3. Step 3 · Temperature and generation

```
generate <- function(start, temperature, cnt = counts, max_len = 15) {
  out <- start; w <- start
  repeat {
    p <- next_word(w, temperature, cnt)
    w <- sample(names(p), 1, prob = p)
    if (w == "</s>" || length(out) >= max_len) break
    out <- c(out, w)
  }
  paste(out, collapse = " ")
}
set.seed(8081)
temperatures <- c(0.5, 1, 2)
runs <- lapply(temperatures, function(t) replicate(2000, generate("melatonin", t)))
temp_summary <- data.frame(
  temperature = temperatures,
  false_share = sapply(runs, function(s) mean(s == false_sentence)),
  true_share = sapply(runs, function(s) mean(s == true_sentence)),
  distinct = sapply(runs, function(s) length(unique(s))))
temp_summary$other_share <- 1 - temp_summary$false_share - temp_summary$true_share
stopifnot(temp_summary$false_share[1] > 0.5, temp_summary$false_share[3] < temp_summary$fals
          temp_summary$distinct[3] > temp_summary$distinct[1])
knitr::kable(temp_summary, digits = 3, row.names = FALSE,
             col.names = c("Temperature", "False sentence", "True sentence", "Distinct sente
             caption = "2,000 sentences generated from "melatonin" at each temperature (synt
```

Table 8.2.: 2,000 sentences generated from “melatonin” at each temperature (synthetic).

Temperature	False sentence	True sentence	Distinct sentences	Other
0.5	0.572	0.165	94	0.263
1.0	0.230	0.184	200	0.586
2.0	0.090	0.140	275	0.771

The following code draws Figure 8.1 in the core reading.

```
library(ggplot2)
plot_df <- data.frame(
  temperature = factor(rep(temp_summary$temperature, 3)),
  outcome = factor(rep(c("False sentence", "True sentence", "Other sentences"
                   levels = c("False sentence", "True sentence", "Other sente
  share = c(temp_summary$false_share, temp_summary$true_share, temp_summary$o
ggplot(plot_df, aes(x = temperature, y = share, fill = outcome)) +
  geom_col(position = position_dodge(width = 0.8), width = 0.75, colour = "#3
  geom_text(aes(label = cdsr_percent(share, 0)), position = position_dodge(wi
            vjust = -0.4, size = 3.2) +
  scale_fill_manual(values = c(`False sentence` = cdsr_colours[["red"]],
                               `True sentence` = cdsr_colours[["skyblue"]],
                               `Other sentences` = "#D9D9D9")) +
  scale_y_continuous(labels = function(x) cdsr_percent(x, 0), limits = c(0, 1
  labs(x = "Temperature", y = "Share of generated sentences", fill = NULL,
       caption = "2,000 generated sentences per temperature; start word 'mela
  theme_cdsr()
```

8.3.4. Step 4 · The evidence card: existence and support, judged separately

Each claim receives two separate judgements. The verdict follows from their combination, and the action follows from the verdict. "Reject" is a legitimate action — for a single sentence or for the whole answer.

```
card <- data.frame(
  claim = c("A1 Melatonin reliably prevents postoperative delirium",
            "A2 Glasses, hearing aids and family visits help orientation",
            "B1 Measures for orientation, sleep, vision, hearing, fluids and mobility reduce
            "B2 Daily geriatrician visits reduced delirium from 50% to 32% in hip fracture p
            "B3 Lavender aromatherapy halves the risk of delirium after hip surgery"),
  source = c("de Jonghe et al. 2014", "none", "Inouye et al. 1999", "Marcantonio et al. 2001
             "Hartmann and Weiss 2021"),
  exists = c("yes", "not applicable", "yes", "yes", "no"),
  supports = c("no: no reduction in incidence (29.6% vs 25.5%)",
               "not assessable without a source",
               "partly: 852 general-medicine patients aged 70+, not surgical",
               "yes: 20/62 vs 32/64 patients, one centre, 126 patients",
               "not assessable"))
verdict <- function(exists, supports) {
  ifelse(exists == "no", "fabricated reference",
  ifelse(exists == "not applicable", "unsourced",
  ifelse(startsWith(supports, "no"), "real source, contradicts claim",
  ifelse(startsWith(supports, "partly"), "supported with caveat", "supported"))))
}
card$verdict <- verdict(card$exists, card$supports)
action_for <- c("real source, contradicts claim" = "reject",
                "fabricated reference" = "reject",
                "unsourced" = "adapt only after finding a source",
                "supported with caveat" = "adapt: name the population",
                "supported" = "adopt in plain language")
card$action <- unname(action_for[card$verdict])          # the action follows from the verdi
```

```
stopifnot(card$verdict[1] == "real source, contradicts claim",
          card$verdict[5] == "fabricated reference",
          sum(card$action == "reject") == 2)
```

The following code prints Table 8.1 in the core reading.

```
knitr::kable(card[, c("claim", "source", "exists", "supports", "verdict", "ac
             col.names = c("Claim (answer)", "Source given", "Source exists?"
                           "Verdict", "Action in the reply"), row.names = FAL
```

The two judgements in the card rest on checks that every reader can repeat. Existence: the DOIs of the three real studies were resolved through Crossref and their abstracts read in PubMed (de Jonghe et al. 2014; Inouye et al. 1999; Marcantonio et al. 2001); the DOI given for the lavender study returns "not found" at doi.org and Crossref, and neither the paper nor a journal of that name was found in Crossref or the NLM Catalog (checked 8 November 2026). Support: the numbers in the card come from the abstracts' results sections. Tools for the existence check at scale — DOI lookups, retraction databases — are covered in Scientometrics in R: Retraction Analysis; the support check cannot be automated away, because it requires reading.

8.3.5. Step 5 · Check the patient letter for readability

The resident's first draft is accurate but written for colleagues. The revision keeps the supported content, removes the rejected claims, offers a choice and ends with a question that checks understanding (teach-back, Chapter 21). Both letters are *synthetic*; their clinical content would still have to be confirmed against the local protocol (*open*).

```
draft <- paste(
  "Dear [patient's name], postoperative delirium is an acute neuropsychiatric syndrome chara
  "fluctuating disturbances of attention and cognition, with a considerable incidence follow
  "surgery in geriatric populations. Pharmacological prophylaxis with melatonin has not demo
  "in a randomised controlled trial in this population. Multicomponent non-pharmacological i
  "addressing orientation, sensory impairment, sleep, hydration and mobilisation are associa
  "reduced incidence in hospitalised older patients, and proactive geriatric consultation re
  "incidence in a randomised trial in patients with hip fracture.")
revised <- paste(
  "Dear [patient's name], after a hip operation some older people become confused for a few
  "Doctors call this delirium. It usually gets better, but it can slow down recovery.",
  "In a study of almost four hundred older people with a broken hip, melatonin did not preve
  "Simple things can help. Please bring your glasses and hearing aids. Ask your family to vi
  "Our team can also ask a doctor for older people to see you every day. You can say yes or
  "Could you tell us in your own words what you will bring and whom you will call if you fee

count_syllables <- function(word) {                 # approximation: groups of vowels, silen
  w <- tolower(gsub("[^A-Za-z]", "", word))
  if (nchar(w) == 0) return(0)
  n <- lengths(regmatches(w, gregexpr("[aeiouy]+", w)))
  if (grepl("[^l]e$", w) && n > 1) n <- n - 1
  max(1, n)
}
readability <- function(text) {
  sentences <- unlist(strsplit(text, "(?<=[.?!])\\s+", perl = TRUE))
  words <- unlist(strsplit(gsub("[^A-Za-z' -]", " ", text), "\\s+")); words <- words[words !
  syl <- vapply(words, count_syllables, numeric(1))
  c(sentences = length(sentences), words = length(words),
    words_per_sentence = length(words) / length(sentences),
    syllables_per_word = mean(syl),
    long_words = mean(syl >= 3),
    flesch = 206.835 - 1.015 * length(words) / length(sentences) - 84.6 * mean(syl))  # Fles
}
```

```
jargon <- c("postoperative", "neuropsychiatric", "incidence", "pharmacologica
            "efficacy", "multicomponent", "geriatric", "hospitalised", "mobil
jargon_hits <- function(text) sum(sapply(jargon, function(j) grepl(j, tolowe
letters_tab <- rbind(draft = c(readability(draft), jargon_terms = jargon_hits
                     revised = c(readability(revised), jargon_terms = jargon_
stopifnot(letters_tab["revised", "flesch"] > letters_tab["draft", "flesch"] +
          letters_tab["revised", "jargon_terms"] < letters_tab["draft", "jarg
round(letters_tab, 2)
```

```
#>         sentences words words_per_sentence syllables_per_word long_words
#> draft           3    82              27.33                2.5       0.49
#> revised        10   108              10.80                1.4       0.08
#>         jargon_terms
#> draft             10
#> revised            0
```

The Flesch Reading Ease score rises from -32 to 78 (higher means easier; the formula combines sentence length and word length) (Flesch 1948); scores below zero occur for text with very long sentences and long words. The syllable count is an approximation, and no formula shows whether a reader has understood the letter. That is what the final question in the revised letter is for: the reply of the patient or her daughter, in their own words, is the real test.

8.3.6. Step 6 · Document origin, model version and use

The reply is clinical documentation; the generated drafts are part of its origin. A short record, stored with the case, makes the process reviewable. For research, the reporting guideline TRIPOD-LLM covers studies that evaluate or use large language models (Gallifant et al. 2025), and medical journals following the ICMJE recommendations expect authors to disclose at submission whether and how AI-assisted technologies were used; the

humans remain responsible for the content (International Committee of Medical Journal Editors 2026).

```
library(jsonlite)
record <- list(
  task = "draft reply to a patient-portal question on delirium prevention",
  tool = "fictitious hospital writing assistant (teaching case; in real use: product, model
  deployment = "in real use: local, institutional or external service; whether personal data
  date_of_generation = "not applicable (synthetic answers written for this book)",
  input = question,
  personal_data_in_input = "none (age and diagnosis only, no identifiers)",
  settings = stored$setting_described,
  stored_output_sha256 = stored$sha256,
  sentences_rejected = card$claim[card$action == "reject"],
  sentences_adapted = card$claim[startsWith(card$action, "adapt")],
  sources_checked = c("10.1503/cmaj.140495", "10.1056/NEJM199903043400901",
                      "10.1046/j.1532-5415.2001.49108.x", "10.5555/jon.2021.0147 (not found)
  checked_by = "resident on duty; countersigned by consultant",
  disclosure = "noted in the case record; the letter is signed by the responsible physician"
stopifnot(length(record$sentences_rejected) == 2)
cat(toJSON(record, pretty = TRUE, auto_unbox = TRUE))
```

```
#> {
#>   "task": "draft reply to a patient-portal question on delirium prevention",
#>   "tool": "fictitious hospital writing assistant (teaching case; in real use: product, mo
#>   "deployment": "in real use: local, institutional or external service; whether personal
#>   "date_of_generation": "not applicable (synthetic answers written for this book)",
#>   "input": "My mother is 78 and will have surgery for a hip fracture tomorrow. Can anythi
#>   "personal_data_in_input": "none (age and diagnosis only, no identifiers)",
#>   "settings": ["fictitious assistant, temperature 0.2", "fictitious assistant, temperatur
#>   "stored_output_sha256": ["79592eb3bc08717bbcb61a29408416d98b76bcb0a1478766096b62be1b249
#>   "sentences_rejected": ["A1 Melatonin reliably prevents postoperative delirium", "B3 Law
#>   "sentences_adapted": ["A2 Glasses, hearing aids and family visits help orientation", "B
```

```
#>   "sources_checked": ["10.1503/cmaj.140495", "10.1056/NEJM199903043400901"
#>   "checked_by": "resident on duty; countersigned by consultant",
#>   "disclosure": "noted in the case record; the letter is signed by the res
#> }
```

8.4. Counter-check

Change the evidence the model can draw on. Retrieval-augmented systems add documents to what the model sees before it writes (Lewis et al. 2020). Our bigram analogue: add two sentences that restate the trial result to the corpus — "melatonin does not prevent delirium in older patients" and "melatonin does not reduce delirium after hip surgery".

Before you run the code, predict: What will the greedy sentence be? At temperature 1, will the false sentence disappear?

```
corpus_retrieved <- c(corpus, "melatonin does not prevent delirium in older p
                      "melatonin does not reduce delirium after hip surgery")
counts_retrieved <- bigram_table(corpus_retrieved)
greedy_retrieved <- greedy("melatonin", cnt = counts_retrieved)
set.seed(8082)
gen_retrieved <- replicate(2000, generate("melatonin", 1, cnt = counts_retrie
share_false_retrieved <- mean(gen_retrieved == false_sentence)
stopifnot(greedy_retrieved == true_sentence, share_false_retrieved > 0,
          share_false_retrieved < temp_summary$false_share[2])
c(greedy = greedy_retrieved, false_share_t1 = round(share_false_retrieved, 3)
```

```
#>                                                  greedy
#> "melatonin does not prevent delirium after hip surgery"
#>                                           false_share_t1
#>                                                   "0.13"
```

With the added evidence, the most probable sentence becomes the true one. Yet at temperature 1 the false sentence is still generated in 13% of runs, compared with 23% before. Retrieval shifts the probabilities; it does not remove the recombination that produced the error. The check of every sentence against its source therefore remains necessary even with a retrieval system — a point that holds, with far more complex models, in practice.

8.5. Consequences for patients and operations

Patients. The harm of answer A would have been concrete: a family relying on melatonin, possibly neglecting measures with better support. The revised letter also changes the relationship: it offers a choice and asks for feedback instead of instructing. Patients are entitled to know when text addressed to them was drafted with a language model and who checked it; how this is communicated is part of the hospital's policy (Chapter 21). The WHO guidance on the ethics and governance of large multi-modal models in health is a starting point for these questions (World Health Organization 2024).

Data protection. The question as entered contained age and diagnosis but no name. Entering identifiable patient data into an external service is a processing of health data — a special category under Article 9 of the General Data Protection Regulation (GDPR) (European Parliament and Council of the European Union 2016) (English text on EUR-Lex). It needs a legal basis, a processor agreement under Article 28 where a provider processes data on the hospital's behalf, and an assessment under Chapter V if data leave the European Economic Area. Which services are approved is a local decision (*open*; Chapter 23).

Security. Text that a model reads can contain instructions. A retrieved document or an attached file can steer the output — indirect prompt

injection — and an assistant with access to records can leak them. These risks and their controls are treated in Chapter 25.

Regulation. Under Article 50 of the AI Act (European Parliament and Council of the European Union 2024b), applicable from 2 August 2026, providers must ensure that people are informed when they interact directly with an AI system, unless this is obvious, and must mark generated text and other synthetic content in a machine-readable way; for systems already on the market before that date, the marking obligation applies from 2 December 2026 (Article 111(4), inserted by Regulation (EU) 2026/1744 (European Parliament and Council of the European Union 2026)). Deployers who publish generated text to inform the public on matters of public interest must disclose this unless the text has undergone human review and someone holds editorial responsibility. Article 4, as replaced by the 2026 amendment, requires providers and deployers to take measures to support the AI literacy of their staff. (Read on EUR-Lex; status: November 2026.) Whether an assistant that answers patient-specific medical questions is itself a medical device depends on its intended purpose (*open*; Chapter 22).

8.6. Look up and transfer

Reviewer question: "For every reference in the manuscript, has someone confirmed both that it exists and that the cited passage supports the sentence it is attached to — and do the authors disclose which AI-assisted tools were used, for what, and with which version?"

Self-test: A generated discussion paragraph for a manuscript reads: "Deep learning has been shown to outperform clinicians in image-based diagnosis [Nagendran et al., BMJ 2020]." The DOI resolves and the paper is a systematic review of studies comparing deep learning with clinicians. Would you keep the sentence?

Answer

No. The source exists, so the existence check passes, but it does not support the sentence. The review examined the design, reporting standards and claims of such studies; it found few prospective studies and randomised trials, a high risk of bias in most non-randomised studies, and abstracts that often claimed performance at least comparable to clinicians (Nagendran et al. 2020). It is a critique of the very claim the sentence makes. A defensible sentence would report what the review found — and the decision to drop the generated sentence altogether is equally defensible.

Annotated sources:

- Jurafsky and Martin (2026) explains n-gram language models (chapter 3) and temperature sampling (section 7.6.3), the two mechanisms of the lab.
- Vaswani et al. (2017) introduces the transformer architecture on which current language models are built.
- Holtzman et al. (2019) shows how decoding choices change the quality and diversity of generated text.
- Ji et al. (2023) surveys hallucination, the production of fluent but unsupported content.
- Walters and Wilder (2023) quantifies fabricated and erroneous references produced by a chatbot.
- Lewis et al. (2020) describes retrieval-augmented generation, the idea behind the counter-check.
- Singhal et al. (2023) evaluates clinical knowledge of large language models and finds gaps on human evaluation.
- de Jonghe et al. (2014), Inouye et al. (1999) and Marcantonio et al. (2001) are the three real studies checked in the evidence card; the first is the real source that does not support the claim attached to it.
- Flesch (1948) defines the Reading Ease formula used for the patient

letter.
- Gallifant et al. (2025) and International Committee of Medical Journal Editors (2026) set the reporting and disclosure expectations for AI use.
- The AI Act (European Parliament and Council of the European Union 2024b), its 2026 amendment (European Parliament and Council of the European Union 2026) and the GDPR (European Parliament and Council of the European Union 2016) are the binding texts behind the consequences; the WHO guidance (World Health Organization 2024) is orientation.

8.7. Further reading

- Chapter 7 — how to trace a statement from a brief or brochure back to the original study.
- Chapter 25 — prompt injection, data leakage and other security risks of generative systems.
- Chapter 21 — teach-back and shared decisions when explaining uncertainty to patients.
- Scientometrics in R: Text Mining Bibliographic Corpora — tokens, n-grams and word counts on real literature corpora.
- Scientometrics in R: Embeddings for Scientometrics — the vector representations that retrieval systems use to find documents.
- CTTIR Tutorials: Transformers: Overview — the architecture behind current language models.
- Strategy in R: LLM Agents and Strategic Reasoning — what changes when language models act rather than only write.
- Methods in R: Use of LLM tools — an example of a disclosure statement that names tools, tasks and the checks applied.

Part III.

Part III · Diagnostics and measurement in four case studies

9. A fracture alert in the emergency department

R. Heller · Date: 2026-07-26 · Last changed: 2026-11-09

At a glance

Guiding question: What does an alert mean when the prevalence changes?
For: clinicians, medical students and reviewers who have to act on a positive image-analysis alert; the core reading needs no programming. · **Prerequisites:** the fourfold table from Chapter 6; the idea of a target population from Chapter 1.
In the course: W3 · Diagnostics, prediction and external evidence · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

9.1. Case and decision

It is the middle of a night shift in a fictitious emergency department. A junior doctor has just looked at the wrist radiograph of a 67-year-old patient who fell on an outstretched hand. She sees no fracture. The image-analysis tool that the department is piloting has placed a box over the distal radius with the label “possible fracture – second look recommended”. The information sheet states a sensitivity and specificity of 90% each; in this chapter both values are *assumptions*. The doctor now has to choose:

call the on-call senior to review the image, splint the wrist and book a follow-up radiograph, or discharge the patient with advice. Before she decides, she asks the question this chapter answers: how often is a box like this one right? The answer depends less on the tool than on who comes through the door.

9.2. Core reading

Core reading · required for W3 · about 10 minutes · ends before 'Lab'

Two settings, one tool. We compare two hypothetical care settings. In setting A, an emergency department radiographs every patient with a limb injury; 5% of the patients examined actually have a fracture. In setting B, patients are referred from a trauma clinic because a clinician already suspects a fracture; here 20% have one. Each setting examines 1,000 patients. For this calculation we give the tool the same sensitivity (the share of fractures that it flags) and the same specificity (the share of patients without fracture that it leaves unflagged) in both settings: 90% each. This is of the same order as pooled estimates from a meta-analysis of AI fracture detection, in which about half of the included studies had a high risk of bias [Kuo et al. (2022); *published*]. **The key assumption is that sensitivity and specificity do not change between the two settings.** In practice this assumption often fails, because the mix of injuries, patients and imaging differs between settings. A meta-epidemiological study of 23 diagnostic meta-analyses found that sensitivity or specificity often varied with prevalence; specificity tended to be lower where prevalence was higher [Leeflang et al. (2013); *published*]. Any real deployment would therefore have to measure both values locally.

Absolute numbers first. Before any percentage, count people. In setting A, 50 of the 1,000 patients have a fracture. The tool flags 45 of them

and misses 5. Of the 950 patients without a fracture, the tool flags 95 and leaves 855 unflagged. In setting B, 200 patients have a fracture: 180 are flagged and 20 are missed. Of the 800 without a fracture, 80 are flagged and 720 are not. Counts of this kind, called natural frequencies, are easier to reason with than conditional probabilities (Hoffrage et al. 2000). Table 9.1 shows them as two fourfold tables and Figure 9.1 shows every patient as one symbol.

Table 9.1.: Fourfold tables for 1,000 patients examined in each setting. Sensitivity and specificity of 90% are assumptions; all counts are synthetic. Independent unit: patient (n = 1,000 per setting).

Setting	Tool output	Fracture	No fracture	Total	Share with fracture
A · 5% prevalence	alert	45	95	140	45/140 = 32.1%
A · 5% prevalence	no alert	5	855	860	5/860 = 0.6%
B · 20% prevalence	alert	180	80	260	180/260 = 69.2%
B · 20% prevalence	no alert	20	720	740	20/740 = 2.7%

Then the denominator. The positive predictive value (PPV) answers the junior doctor's question: of all patients with an alert, how many have a fracture? Its denominator is every patient with an alert, not every patient with a fracture. In setting A, 140 patients receive an alert and 45 of them have a fracture: the PPV is 45/140 = 32.1%. In setting B the same tool gives 180/260 = 69.2%. Nothing about the software has changed; only the patients have. The negative predictive value (NPV) moves in the other direction, from 99.4% in A to 97.3% in B. A missing alert is therefore not an all-clear: in B, 20 of 1,000 patients have a fracture that the tool did not

flag. The general relationship between prevalence and predictive values is worked through in Methods in R: Predictive Values and Prevalence and in the CTTIR lab Diagnostic testing: Se, Sp, PPV, NPV, LR (Altman and Bland 1994).

```
library(ggplot2)
# One row per patient; categories in the order alert/fracture, alert/no fract
categories <- c("Alert, fracture (true positive)", "Alert, no fracture (false
                "No alert, fracture (false negative)", "No alert, no fracture
icons <- do.call(rbind, lapply(1:2, function(i) {
  counts <- c(tp[i], fp[i], fn[i], tn[i])
  data.frame(
    setting = sprintf("%s · prevalence %s · PPV %d/%d = %s", names(prev)[i],
                      tp[i], alerts[i], cdsr_percent(ppv[i])),
    category = factor(rep(categories, times = counts), levels = categories),
    position = 0:(n_examined - 1)
  )
}))
icons$column <- icons$position %% 50
icons$row <- icons$position %/% 50
ggplot(icons, aes(column, -row, colour = category, shape = category, size = c
  geom_point() +
  facet_wrap(~ setting, ncol = 1) +
  scale_colour_manual(values = unname(c(cdsr_colours["blue"], cdsr_colours["r
                                        cdsr_colours["orange"], "grey70"))) +
  scale_shape_manual(values = c(16, 15, 17, 1)) +
  scale_size_manual(values = c(1.9, 1.9, 1.9, 0.9)) +
  guides(colour = guide_legend(ncol = 2), shape = guide_legend(ncol = 2), siz
  coord_equal() +
  labs(x = NULL, y = NULL, colour = NULL, shape = NULL, size = NULL,
       caption = "Synthetic data · n = 1,000 patients per setting · sensitivi
  theme_cdsr() +
  theme(axis.text = element_blank(), panel.grid = element_blank())
```

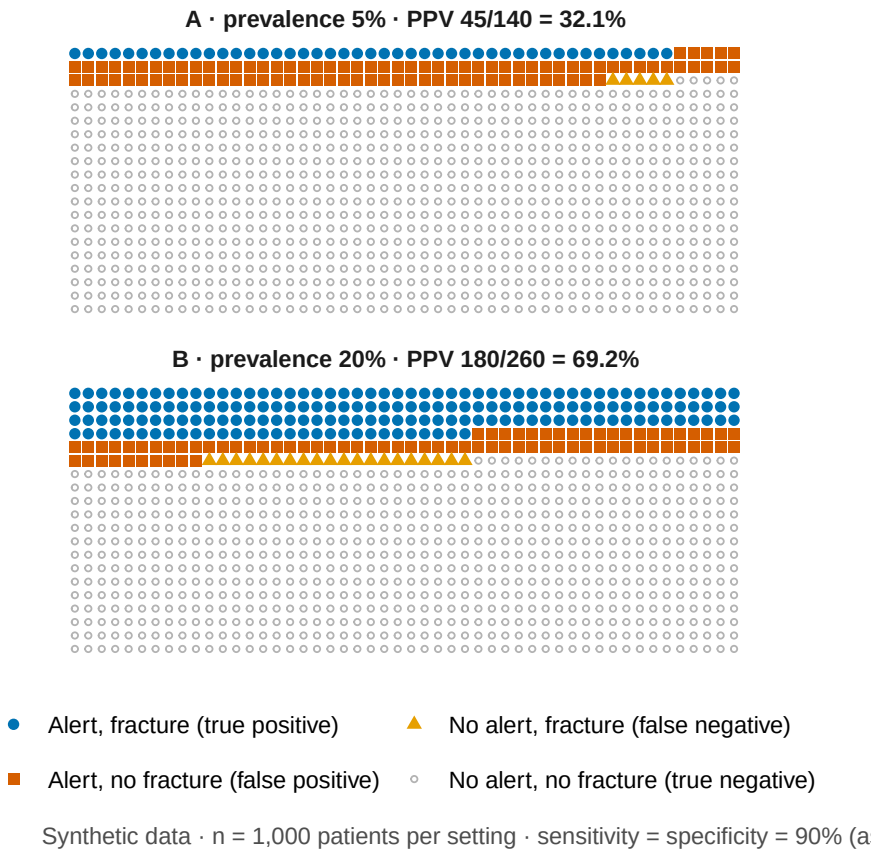


Figure 9.1.: Every symbol is one of 1,000 patients examined per setting, ordered by tool output and reference result. The alerts (first two blocks) form the denominator of the positive predictive value. Synthetic counts; sensitivity and specificity of 90% are assumptions. Independent unit: patient (n = 1,000 per setting).

Three research questions, not one. The fourfold table invites a single

verdict: "the tool works" or "it does not". In fact three separate questions are at stake, and each needs its own design.

1. *Model performance.* How well does the tool separate patients with and without fracture, compared with a defined reference standard, in the population where it will be used? This is a diagnostic accuracy study, reported according to STARD-AI (Sounderajah et al. 2025). It yields sensitivity, specificity and predictive values with confidence intervals.
2. *Patient benefit.* Do patients whose images are assessed with the tool have fewer missed fractures with consequences, without more harm from extra imaging, immobilisation or waiting? Accuracy cannot answer this. It needs a comparison group and outcomes that matter to patients, for example in a cluster-randomised or stepped-wedge study.
3. *Organisational burden.* How many additional second looks does the tool create per shift, and who absorbs them? If 1,000 examinations fall into 90 shifts (30 days with three shifts each; an *assumption*), setting A produces about 1.6 alerts per shift, 1.1 of them false. Setting B produces 2.9 alerts per shift, of which 0.9 are false. Early clinical evaluations of this kind are reported along the lines of DECIDE-AI (Vasey et al. 2022).

Reader study and local pilot answer different questions. In a reader study, several clinicians read the same set of images with and without the tool. Such a study can show whether the tool changes readers' accuracy under controlled conditions. Its case mix is usually chosen by design: in one multi-reader study of fracture detection on radiographs, the fracture prevalence was fixed at 50% [Guermazi et al. (2022); *published*]. A PPV from such a set says little about setting A. A local pilot, for example in silent mode, records how often the tool alerts in the department's own case mix, what the local PPV is and how long second looks take. Neither design on its own shows that patients are better off.

A lower PPV alone is no reason to reject a product. Setting A's

PPV of 32.1% sounds poor. Whether it is acceptable depends on what a false alert costs (a few minutes of a senior clinician's time, perhaps an extra view) and on what a true alert gains. The gain is not 45 detected fractures: it is the fractures among them that the clinicians would otherwise have missed. Weighing these consequences is a decision analysis, not a property of the PPV (Vickers and Elkin 2006); Chapter 6 develops it with thresholds and net benefit.

Saying it at handover. A spoken handover in setting A could be: *"The tool has flagged the distal radius. In a patient group like ours, only about one in three of these flags turns out to be a real fracture, so I would like a second pair of eyes on the image before we decide."* In real operation, the "one in three" has to come from local data, not from this exercise. The patient may then ask: *"If the computer is wrong two times out of three, why listen to it at all, and would no box have meant that I could go home?"* A good answer explains that the box asks for a second look but does not decide, and that a missing box lowers the probability of a fracture without ruling it out. The clinical findings still count.

Question for the workshop: Your team works in setting A. Which single number would you put on the whiteboard for the night shift – the PPV, the false alerts per shift or the missed fractures per month – and who should be told the other two?

9.3. Lab

The lab recomputes every number from the core reading in base R, from the inputs to the predictive values. In W3, one team works out setting A and another setting B by hand before anyone runs the code. Start a fresh R session in the book folder. The following single line already reproduces both PPVs and is a useful first check that R is running:

```
# Start command: PPV = true positives / all alerts, for 5% and 20% prevalence
round(100 * 0.9 * c(0.05, 0.20) / (0.9 * c(0.05, 0.20) + 0.1 * c(0.95, 0.80)))
```

```
#> [1] 32.1 69.2
```

Step 1 – inputs. Each row describes one setting. Sensitivity and specificity are inputs, not results; in this chapter they are assumptions.

```
cohorts <- data.frame(
  setting = c("A", "B"),
  description = c("unselected emergency department", "referred with clinical
  examined = c(1000L, 1000L),      # patients examined per setting
  prevalence = c(0.05, 0.20),      # share with fracture according to the refe
  sensitivity = c(0.90, 0.90),     # assumption: identical in both settings
  specificity = c(0.90, 0.90)      # assumption: identical in both settings
)
cohorts
```

```
#>   setting                      description examined prevalence sensitivity
#> 1       A  unselected emergency department     1000       0.05         0.9
#> 2       B referred with clinical suspicion     1000       0.20         0.9
#>   specificity
#> 1         0.9
#> 2         0.9
```

Step 2 – absolute frequencies. People are counted in whole numbers, so every count is rounded explicitly.

```
cohorts$fracture <- round(cohorts$examined * cohorts$prevalence)
cohorts$no_fracture <- cohorts$examined - cohorts$fracture
cohorts$true_pos <- round(cohorts$sensitivity * cohorts$fracture)        # al
```

```
cohorts$false_neg <- cohorts$fracture - cohorts$true_pos                # fracture without
cohorts$false_pos <- round((1 - cohorts$specificity) * cohorts$no_fracture) # alert without
cohorts$true_neg <- cohorts$no_fracture - cohorts$false_pos             # neither
cohorts[, c("setting", "fracture", "true_pos", "false_neg", "no_fracture", "false_pos", "tru
```

```
#>   setting fracture true_pos false_neg no_fracture false_pos true_neg
#> 1       A       50       45         5         950        95       855
#> 2       B      200      180        20         800        80       720
```

Step 3 – denominators and predictive values. The alerts form the denominator of the PPV; the patients without an alert form the denominator of the NPV.

```
cohorts$alerts <- cohorts$true_pos + cohorts$false_pos
cohorts$no_alerts <- cohorts$false_neg + cohorts$true_neg
cohorts$ppv <- cohorts$true_pos / cohorts$alerts
cohorts$npv <- cohorts$true_neg / cohorts$no_alerts
data.frame(setting = cohorts$setting,
           ppv = paste0(cohorts$true_pos, "/", cohorts$alerts, " = ", cdsr_percent(cohorts$p
           npv = paste0(cohorts$true_neg, "/", cohorts$no_alerts, " = ", cdsr_percent(cohort
```

```
#>   setting             ppv             npv
#> 1       A  45/140 = 32.1% 855/860 = 99.4%
#> 2       B 180/260 = 69.2% 720/740 = 97.3%
```

```
# Target values from the core reading
stopifnot(
  cohorts$true_pos == c(45, 180),
  cohorts$false_pos == c(95, 80),
  cohorts$alerts == c(140, 260),
  isTRUE(all.equal(cohorts$ppv, c(45 / 140, 180 / 260))),
```

```
  cdsr_percent(cohorts$ppv) == c("32.1%", "69.2%"),
  # the lab and the hidden calculation behind the core reading agree
  cohorts$true_pos == tp, cohorts$false_pos == fp, cohorts$alerts == alerts
)
```

> **Common mistake**
>
> "Sensitivity 90%" is often read as "90% of alerts are right". Sensitivity is conditioned on the fracture (45 of 50 fractures are flagged); the PPV is conditioned on the alert (45 of 140 alerts are fractures). Swapping the two conditions turns 32.1% into 90% for the same 1,000 patients.

Step 4 – the workload per shift. The number of shifts is an *assumption* and should be replaced by the local rota. The calculation separates the alerts that need a second look from the false alerts among them.

```
shifts_per_1000 <- 90  # assumption: 30 days x 3 shifts for 1,000 examinatio
cohorts$alerts_per_shift <- cohorts$alerts / shifts_per_1000
cohorts$false_alerts_per_shift <- cohorts$false_pos / shifts_per_1000
cohorts$missed_per_1000 <- cohorts$false_neg
cohorts[, c("setting", "alerts_per_shift", "false_alerts_per_shift", "missed_
```

```
#>   setting alerts_per_shift false_alerts_per_shift missed_per_1000
#> 1       A         1.555556              1.0555556              5
#> 2       B         2.888889              0.8888889             20
```

```
stopifnot(round(cohorts$alerts_per_shift, 2) == c(1.56, 2.89))
```

Setting B produces more alerts per shift but fewer false ones: 80 against 95 per 1,000 examinations. Workload and PPV are different quantities, and a team planning staffing needs both.

Step 5 – raw values for download. The fourfold counts are stored as a small table so that W3 groups can start from the file instead of the code.

```
raw <- cohorts[, c("setting", "description", "examined", "prevalence", "sensitivity", "speci
                   "true_pos", "false_pos", "false_neg", "true_neg")]
write.csv(raw, "../data/b09-fracture-alert-cohorts.csv", row.names = FALSE)
# Re-reading the file must give the same PPVs
check <- read.csv("../data/b09-fracture-alert-cohorts.csv")
stopifnot(isTRUE(all.equal(check$true_pos / (check$true_pos + check$false_pos), cohorts$ppv)
```

The file `data/b09-fracture-alert-cohorts.csv` contains only synthetic counts and the two assumptions.

> **Other implementation: Python with PyTorch**
>
> A fracture tool of this kind is usually a convolutional network trained on radiographs. The PyTorch tutorial on transfer learning shows the standard route: start from a ResNet-18 with ImageNet weights (`weights='IMAGENET1K_V1'`), resize to 256 pixels, crop to 224, normalise with the ImageNet channel means and standard deviations and retrain the final layer (Paszke et al. 2019). The concrete advantage is access to pretrained weights and to research code that already works on images; R would only consume the network's output. For a fair comparison with the counts in this chapter the following must be identical and written down: the image set and the reference standard, the image orientation (training with random horizontal flips mixes left and right), the conversion of greyscale radiographs to three channels, the weights file and software version (torch 2.15.0 and torchvision 0.29.2 were current in November 2026), and the threshold that turns a probability into an alert. Logit, probability and binary alert should be compared separately. Patients often have two or three views; the counting unit for the fourfold table (patient or image)

must be the same on both sides. Status: conceptual comparison, not run.

9.4. Counter-check

Change one assumption and predict the direction before you compute. Suppose the department installs new radiography equipment and the tool's specificity in setting A drops from 90% to 80%, while sensitivity stays at 90%. Does the PPV halve, fall a little or stay the same? And what happens to the false alerts per shift?

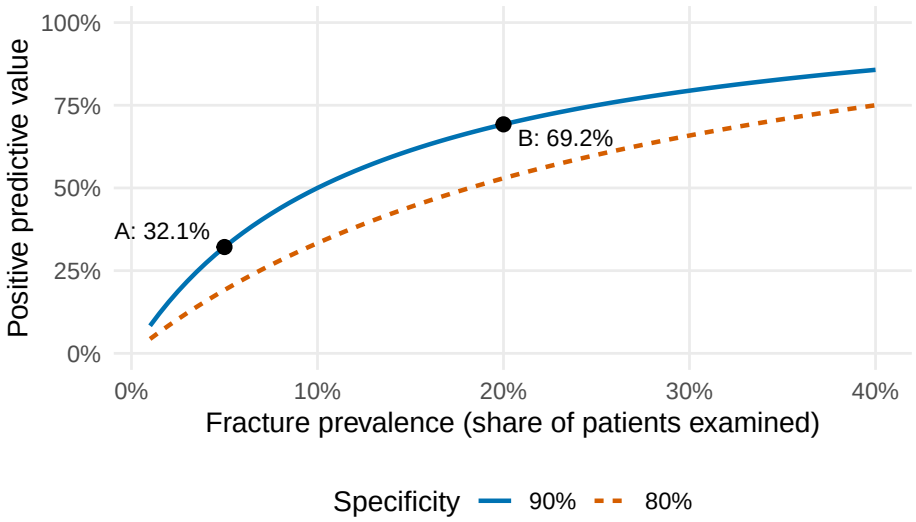
```
# Function: predictive values and counts for any combination of inputs
fourfold <- function(prevalence, sensitivity, specificity, examined = 1000) {
  fracture <- round(examined * prevalence)
  true_pos <- round(sensitivity * fracture)
  false_pos <- round((1 - specificity) * (examined - fracture))
  data.frame(prevalence, sensitivity, specificity, true_pos, false_pos,
             alerts = true_pos + false_pos, ppv = true_pos / (true_pos + fals
}
changed <- rbind(fourfold(0.05, 0.90, 0.90), fourfold(0.05, 0.90, 0.80))
changed$false_alerts_per_shift <- changed$false_pos / shifts_per_1000
changed
```

```
#>   prevalence sensitivity specificity true_pos false_pos alerts       ppv
#> 1       0.05         0.9         0.9       45        95    140 0.3214286
#> 2       0.05         0.9         0.8       45       190    235 0.1914894
#>   false_alerts_per_shift
#> 1               1.055556
#> 2               2.111111
```

```
stopifnot(changed$false_pos == c(95, 190), round(changed$ppv[2], 3) == 0.191)
```

The false alerts double, from 95 to 190 per 1,000 examinations, while the PPV falls from 32.1% to 19.1%. At low prevalence, specificity drives the workload. Figure 9.2 shows the same reasoning over a range of prevalences.

```
grid <- expand.grid(prevalence = seq(0.01, 0.40, by = 0.005), specificity = c(0.90, 0.80))
grid$ppv <- 0.90 * grid$prevalence / (0.90 * grid$prevalence + (1 - grid$specificity) * (1 -
grid$specificity <- factor(cdsr_percent(grid$specificity, 0), levels = c("90%", "80%"))
marks <- data.frame(prevalence = prev, ppv = ppv, label = paste0(names(prev), ": ", cdsr_per
                    hjust = c(1.15, -0.15), vjust = c(-0.5, 1.3))  # label A above left, B b
ggplot(grid, aes(prevalence, ppv, linetype = specificity, colour = specificity)) +
  geom_line(linewidth = 0.9) +
  geom_point(data = marks, aes(prevalence, ppv), inherit.aes = FALSE, size = 2.5) +
  geom_text(data = marks, aes(prevalence, ppv, label = label, hjust = hjust, vjust = vjust),
            inherit.aes = FALSE, size = 3.5) +
  scale_x_continuous(labels = function(x) cdsr_percent(x, 0)) +
  scale_y_continuous(labels = function(x) cdsr_percent(x, 0), limits = c(0, 1)) +
  scale_colour_manual(values = unname(cdsr_colours[c("blue", "red")])) +
  scale_linetype_manual(values = cdsr_linetypes[1:2]) +
  labs(x = "Fracture prevalence (share of patients examined)", y = "Positive predictive valu
       colour = "Specificity", linetype = "Specificity",
       caption = "Synthetic calculation · sensitivity 90% (assumption) · denominator: patien
  theme_cdsr()
```



Synthetic calculation · sensitivity 90% (assumption) · denominator: patients with an alert

Figure 9.2.: Positive predictive value as a function of fracture prevalence for sensitivity 90% and specificity 90% or 80% (assumptions). The two settings of the core reading are marked. Synthetic calculation; denominator: patients with an alert.

A second counter-check changes the target population instead of the tool. If the tool were used for all radiographs in a children's emergency department, which numbers in the table would you have to measure again before reusing any of them? (All four cells: prevalence, injury pattern, image appearance and the reference standard can all differ.)

9.5. Consequences for patients and operations

Ethics. False alerts are not free. Patients may wait longer, receive an extra view or a CT scan with additional radiation, or leave with a splint

they did not need. Clinicians who see many uninformative alerts start to ignore them: in one primary care study, acceptance of reminders fell with each additional reminder per encounter and with the share of repeated reminders [Ancker et al. (2017); *published*]. The opposite risk is automation bias: a missing box may be read as an all-clear. Who carries which share of these harms, and how uncertainty is explained to a patient, is the subject of Chapter 21.

Regulation. The intended purpose decides the regulatory question, not the calculation. Software "intended to provide information which is used to take decisions with diagnosis or therapeutic purposes" is classified at least as class IIa under Rule 11 of Annex VIII of the Medical Device Regulation (EU) 2017/745 [European Parliament and Council of the European Union (2017a); binding text]. An AI system that is itself a medical device and has to undergo third-party conformity assessment is a high-risk AI system under Article 6(1) of the AI Act, Regulation (EU) 2024/1689 [European Parliament and Council of the European Union (2024b); binding text]. Since the amending Regulation (EU) 2026/1744 (published 24 July 2026, in force since 27 July 2026), the obligations of Chapter III, Sections 1 to 3, apply to such systems from 2 August 2028 [European Parliament and Council of the European Union (2026); binding text]. Whether the fictitious tool in this chapter is such a device, and who would be manufacturer and who deployer, is examined in Chapter 22 by reading its intended purpose. Status: November 2026.

Data protection. The local pilot needs images, tool outputs and reference diagnoses. Whether an external service provider needs the original images for this evaluation, or whether derived and minimised data suffice, is decided in Chapter 23.

Information security and operations. If the tool fails or an unapproved model version is installed, reporting must continue by the regular route without delay. Chapter 24 designs a failure test and that fallback; Chapter 20 uses alerts per shift as a monitoring signal.

The case returns four times. In Chapter 21 a patient, a nurse and a researcher judge the same alert and the distribution of false alerts. In Chapter 22 teams read the fictitious intended purpose. In Chapter 23 they decide whether a service provider needs the original images. In Chapter 24 they design a failure test and a safe route back to regular reporting.

9.6. Look up and transfer

- **Reviewer question:** The manuscript reports a PPV of 70%. In which population was it measured, what was the prevalence, and was that prevalence fixed by the study design?
- **Self-test:** A smaller emergency department wants to use the tool only at night. It performs 400 limb radiographs per month on 30 night shifts. Assume a fracture prevalence of 10%, a sensitivity of 90% and, because of older equipment, a specificity of 80%. How many alerts and how many false alerts arise per night, what is the PPV, and should the department switch the tool off because the PPV is “only one third”?

Answer

Of 400 patients, 40 have a fracture: 36 are flagged and 4 are missed. Of the 360 without a fracture, 72 are flagged and 288 are not. There are 108 alerts, so the PPV is 36/108 = 33.3% and the NPV 288/292 = 98.6%. Per night shift this means 3.6 alerts, 2.4 of them false. The PPV alone does not justify switching the tool off. The department needs to know how long a second look takes at night, how many of the 36 flagged fractures the night team would otherwise have missed, and whether the specificity of 80% reflects a local equipment problem that could be fixed. With `fourfold(0.10, 0.90, 0.80, examined = 400)` the lab function reproduces these numbers.

- **Annotated sources:**
  - Leeflang et al. (2013) supports the warning that sensitivity and specificity need not be transferable between settings with different prevalence.
  - Hoffrage et al. (2000) is the reason why the chapter counts people before it computes percentages.
  - Altman and Bland (1994) is the classic short explanation of predictive values and their dependence on prevalence.
  - Guermazi et al. (2022) is the published example of a reader study with prevalence fixed by design.
  - Vickers and Elkin (2006) introduces net benefit, the decision-analytic view behind "a lower PPV is not a verdict".
  - Ancker et al. (2017) provides published evidence that more and repeated alerts are accepted less often.
  - Sounderajah et al. (2025) and Vasey et al. (2022) are the reporting guidelines for accuracy studies and early clinical evaluation.
  - European Parliament and Council of the European Union (2017a), European Parliament and Council of the European Union (2024b) and European Parliament and Council of the European Union (2026) are the binding legal texts behind the regulatory paragraph.

9.7. Further reading

- Chapter 6 – thresholds, calibration and net benefit for the same fourfold logic.
- Chapter 7 – how to read a diagnostic AI study with STARD-AI before trusting its PPV.
- Chapter 20 – alerts per shift as a monitoring signal during a pilot.
- Chapter 21 – the same alert from the perspectives of patient, nurse and researcher.

- Methods in R: Diagnostic Test Accuracy – sensitivity, specificity and their confidence intervals.
- Methods in R: Power for Diagnostic Accuracy – how many patients a local accuracy study would need.
- CTTIR Courses: Diagnostic testing: Se, Sp, PPV, NPV, LR – a lab with likelihood ratios as the next step.
- CTTIR Tutorials: Diagnostic Test Accuracy – a compact tutorial for revision.

10. Hyperspectral imaging

R. Heller · Date: 2026-03-08 · Last changed: 2026-11-09

At a glance

Guiding question: What separates a biological signal from illumination and device characteristics?
For: doctoral researchers and preclinical teams who analyse spectral images, and reviewers of such studies. · **Prerequisites:** analysis units and denominators from Chapter 3; baseline thinking from Chapter 5.
In the course: W4 · What imaging, gait and tissue models measure · **Data:** synthetic (generated in the chapter; fictitious animals) · **Core reading time:** about 10 minutes

10.1. Case and decision

A doctoral researcher in a fictitious preclinical group analyses hyperspectral images of skin flaps in eight rats. The question is whether a fictitious substance improves the oxygenation of the flap tip on day 3. The four control animals were imaged in the first week, the four treated animals in the second. Between the two weeks the camera lamp was replaced. The camera records a new white reference at every start-up, but the analysis script still loads the file from the first day. The oxygenation index of the treated animals is clearly higher, and a t-test over all pixels returns a p-value far below 0.001. The supervisor wants to submit an abstract this

week. The researcher has to decide whether to report a treatment effect, or first to check the calibration and count how many independent units the study really has. The first error in this chapter is not statistical. It is the light.

10.2. Core reading

Core reading · required for W4 · about 10 minutes · ends before ‘Lab’

What the camera records. A hyperspectral camera stores a cube of numbers: two spatial axes and one spectral axis. Every pixel holds a spectrum, here 101 values from 500 to 1,000 nm in steps of 5 nm. The raw value in each band is a count from the sensor (digital numbers, DN). It is the product of three things: how much light the lamp emits at that wavelength, how much of it the tissue reflects, and how sensitive the sensor is at that wavelength. Only the middle factor is biology (Lu and Fei 2014).

Calibration turns counts into reflectance. Two reference images separate the tissue from the device. A dark reference, taken with the shutter closed, records the sensor’s offset. A white reference, taken of a standard tile that reflects almost all light (here 99%), records lamp and sensor together. The reflectance of a pixel is then $R(\lambda) = 0.99 \cdot (\text{raw} - \text{dark})/(\text{white} - \text{dark})$, a dimensionless share between 0 and 1. **The key assumption is that the white reference was recorded under the same light as the tissue image.** If the lamp changes in between, the division no longer cancels it.

From spectrum to index. Deoxygenated haemoglobin absorbs red light around 660 nm much more strongly than oxygenated haemoglobin, whereas around 800 nm the two absorb similarly. The ratio $R(660)/R(800)$ therefore rises with oxygenation. Real perfusion parameters are more

elaborate and are derived from validated absorption models (Holmer et al. 2018); the index in this chapter is a didactic simplification built on synthetic absorption profiles. A ratio has one useful property: it cancels anything that scales the whole spectrum, such as the distance between camera and tissue. It does not cancel anything that changes the *shape* of the spectrum.

Two kinds of reference. The word "reference" means two different things in this chapter. The white and dark references are *technical* references: they turn the number on the screen into a reflectance, and a grey tile of known reflectance can check them at every session. A *biological* reference answers another question: does the index track tissue oxygenation at all? That would need an independent measurement in the same tissue, for example an invasive oxygen probe or a histological marker of ischaemia, read without knowledge of the index. The synthetic study has no such reference; its absence is recorded as *open*. Without it, a higher index means "the tissue reflects relatively more red light", not "the tissue is better oxygenated". The physical units change along the chain: wavelength in nanometres, raw signal in counts, reflectance as a share between 0 and 1, and the index as a dimensionless ratio. A difference of 0.05 index units has no clinical meaning until a biological reference has given it one.

What went wrong. The replacement lamp is brighter and emits relatively more red light. Brightness alone would not matter for a ratio. The change in colour does: with the old white reference, the reflectance at 660 nm is inflated relative to 800 nm by a factor of 1.117, and so is every index measured in week 2 (*synthetic*). Because all treated animals were imaged in week 2, lamp and treatment are completely confounded. No statistical model can separate them; only a calibration with the matching white reference can. A gain of about 12% is roughly the difference between the poorly perfused tip and the well perfused base of a flap in the synthetic data, so the artefact is large enough to look like biology.

The measurement chain. Figure 10.1 lays out each step from lamp

to decision. It marks where the references enter, where the physical unit changes and where a learning model could replace the simple index. Drawing the chain before analysing data is the quickest way to find the step at which a device property can masquerade as a treatment effect.

```
library(ggplot2)
# Generic drawing of a measurement chain: one box per step, a note per step,
draw_chain <- function(steps, note_title = "Unit or reference") {
  steps$y <- rev(seq_len(nrow(steps)))
  types <- c("measurement", "reference", "model", "decision")
  ggplot(steps) +
    geom_segment(data = steps[-nrow(steps), ], aes(x = 1, xend = 1, y = y - 0
                 arrow = grid::arrow(length = grid::unit(0.15, "cm")), colour
    geom_tile(aes(x = 1, y = y, fill = type, linetype = type), width = 1.7, h
              colour = "grey20", linewidth = 0.5) +
    geom_text(aes(x = 1, y = y, label = label), size = 3.4) +
    geom_text(aes(x = 2.0, y = y, label = note), hjust = 0, size = 3.1, colou
    annotate("text", x = 2.0, y = nrow(steps) + 0.55, label = note_title, hju
             size = 3.2) +
    scale_fill_manual(values = c(measurement = "white", reference = "grey85",
                                 decision = "grey70"), breaks = types, name =
    scale_linetype_manual(values = c(measurement = "solid", reference = "soli
                                     decision = "solid"), breaks = types, nam
    scale_x_continuous(limits = c(0.1, 4.6)) +
    scale_y_continuous(limits = c(0.5, nrow(steps) + 0.8)) +
    theme_void(base_size = 11) +
    theme(legend.position = "bottom", plot.caption = element_text(hjust = 0,
}
chain_b10 <- data.frame(
  label = c("Lamp and tissue", "Camera: raw counts", "Dark and white referenc
            "Region mean (acquisition)", "Index or learning model", "Animal m
  note = c("wavelength 500-1,000 nm; lamp spectrum varies", "counts (DN) per
           "white tile 99% reflectance; same lamp required", "share 0-1, dime
           "48 region means (8 animals x 2 regions x 3)", "index R(660)/R(800
```

```
           "independent unit: animal, n = 8", "4 vs 4 animals; check order of imaging"),
  type = c("measurement", "measurement", "reference", "measurement", "measurement", "model",
           "decision")
)
draw_chain(chain_b10) +
  labs(caption = "Synthetic case · fictitious animals · independent unit: animal (n = 8)")
```

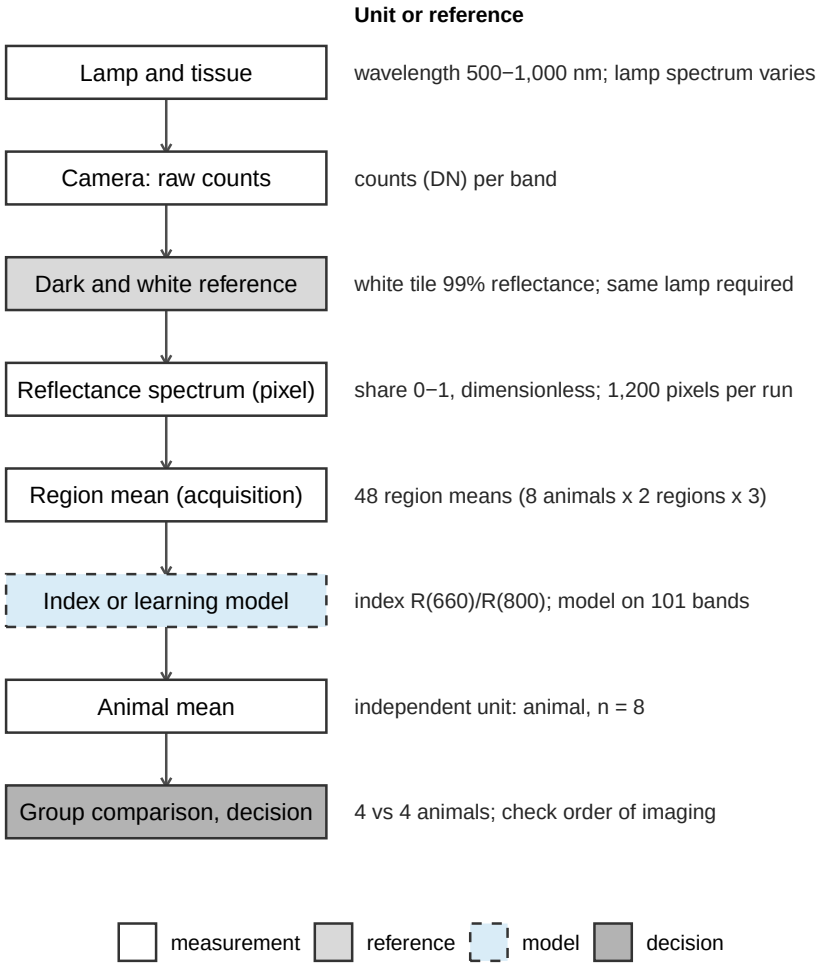


Figure 10.1.: Measurement chain for the hyperspectral case. Each box is one step, the right-hand column gives the unit or the reference at that step. The dashed box marks where a learning model can replace the index. Independent unit for the treatment question: animal (n = 8, four per group). Synthetic case.

Count the units before testing. The study produces many numbers, but few independent ones (Table 10.1). Pixels within one region share the same animal, the same lamp and the same acquisition; they are pseudoreplicates, not replicates (Lazic 2010). The experimental unit is the animal, as the ARRIVE guidelines ask authors to state explicitly (Percie du Sert et al. 2020).

Table 10.1.: Analysis levels in the hyperspectral case. Synthetic design; independent unit for the treatment question: animal (n = 8).

Level	Count in the case	What varies at this level	Unit for the treatment test?
Pixel	1,200 (8 × 2 × 3 × 25)	noise, small-scale tissue texture	no
Region in one acquisition	48	positioning, brightness	no
Animal	8 (4 per group)	biology, treatment	**yes**
Imaging week	2	lamp, calibration file	confounded with group

Where the learning model sits. A classifier trained on full spectra could replace the index in step six of the chain, for instance to tell well and poorly perfused tissue apart. It sits *after* calibration and inherits every calibration error. It therefore has to be compared with the simple index on new animals, not on new pixels of known animals, and it has to pass the same illumination check.

Question for the workshop: Which single additional image, taken at every imaging session, would have revealed the lamp problem before any animal was analysed, and what value would you expect it to show?

10.3. Lab

The lab simulates the whole chain: synthetic tissue spectra, lamp, sensor, dark and white references, calibration, the index, a grouped test and a comparison of the index with a learning model on new animals. All absorption profiles are synthetic and qualitative; they are not tabulated extinction coefficients.

Step 1 – tissue, lamps and sensor.

```
wl <- seq(500, 1000, by = 5)                        # wavelengths in nm (101 ba
bump <- function(centre, width) exp(-0.5 * ((wl - centre) / width)^2)
# Synthetic absorption profiles: deoxygenated absorbs strongly around 660 nm
abs_deoxy <- 1.00 * bump(555, 22) + 0.32 * bump(650, 60) + 0.10 * bump(760, 1
abs_oxy   <- 0.85 * bump(542, 11) + 0.90 * bump(577, 10) + 0.05 * bump(900, 9
# Tissue reflectance (share 0-1) for oxygen saturation 'sat', blood volume fa
tissue_reflectance <- function(sat, blood, scatter = 1) {
  background <- (0.55 + 0.00025 * (wl - 500)) * scatter
  background * exp(-blood * (sat * abs_oxy + (1 - sat) * abs_deoxy))
}
lamp_A <- 0.45 + 0.55 * (wl - 500) / 500                  # original lamp, ri
lamp_B <- 1.2 * lamp_A * (1 + 0.25 * (800 - wl) / 300)    # replacement: 20%
sensor <- 0.95 - 0.45 * ((wl - 720) / 300)^2              # relative sensor s
full_scale <- 20000                                       # counts for a peri
dark_level <- 300                                         # sensor offset in
# Raw counts of one spectrum: signal plus offset plus photon noise
raw_counts <- function(reflectance, lamp) {
  signal <- full_scale * lamp * sensor * reflectance
  signal + dark_level + rnorm(length(signal), 0, sqrt(signal + dark_level))
}
# Index of one spectrum (vector) or of every pixel (matrix with one column pe
oxygen_index <- function(R) {
  if (is.matrix(R)) R[wl == 660, ] / R[wl == 800, ] else R[wl == 660] / R[wl
}
```

```
round(c(poorly_oxygenated = oxygen_index(tissue_reflectance(0.45, 1.5)),
        well_oxygenated = oxygen_index(tissue_reflectance(0.70, 1.5))), 3)
```

```
#> poorly_oxygenated   well_oxygenated
#>             0.750             0.847
```

Step 2 – dark and white references and a check. The white reference is the mean of ten frames of the 99% tile under each lamp. A grey tile with a flat reflectance of 50% serves as a check: after calibration its spectrum must be flat at 0.5 and its index must be 1.00 within a tolerance of ±0.02, set before the study (*assumption*).

```
set.seed(1010)
lamps <- list(week1 = lamp_A, week2 = lamp_B)
white_ref <- lapply(lamps, function(l) rowMeans(sapply(1:10, function(i) raw_counts(rep(0.99
calibrate <- function(raw, white) 0.99 * (raw - dark_level) / (white - dark_level)
# Grey check tile: mean of 100 pixels, calibrated with the matching and with the stale white
grey_tile <- function(week, white) {
  raw <- sapply(1:100, function(i) raw_counts(rep(0.50, length(wl)), lamps[[week]]))
  rowMeans(calibrate(raw, white))
}
check <- data.frame(
  week = c("week 1", "week 2", "week 2"),
  white_reference = c("matching", "matching", "stale (week 1)"),
  mean_reflectance = c(mean(grey_tile(1, white_ref$week1)), mean(grey_tile(2, white_ref$week
                       mean(grey_tile(2, white_ref$week1))),
  grey_index = c(oxygen_index(grey_tile(1, white_ref$week1)), oxygen_index(grey_tile(2, whit
                 oxygen_index(grey_tile(2, white_ref$week1)))
)
check$within_tolerance <- abs(check$grey_index - 1) <= 0.02
check
```

```
#>     week white_reference mean_reflectance grey_index within_tolerance
#> 1 week 1        matching        0.5000383  0.9945404             TRUE
#> 2 week 2        matching        0.5000635  0.9912745             TRUE
#> 3 week 2  stale (week 1)        0.6251338  1.1137797            FALSE
```

```
stopifnot(check$within_tolerance == c(TRUE, TRUE, FALSE),
          abs(check$mean_reflectance[1:2] - 0.5) < 0.01)
```

The stale white reference fails the grey check in week 2. The check needs no animal and no statistics.

Step 3 – eight fictitious animals. In the synthetic truth the substance has no effect: both groups draw their oxygen saturation from the same distribution. Each animal has two regions (flap base and flap tip), three acquisitions per region and 25 pixels per region of interest.

```
set.seed(2018)
animals <- data.frame(animal = sprintf("R%02d", 1:8), group = rep(c("control"
                      week = rep(1:2, each = 4))
animals$blood <- rnorm(8, 1.5, 0.15)                        # blood volume fac
animals$scatter <- rnorm(8, 1, 0.05)                        # scattering facto
animals$sat_base <- rnorm(8, 0.70, 0.03)                    # no treatment ef
animals$sat_tip <- animals$sat_base - rnorm(8, 0.25, 0.03) # the tip is less
pixel_rows <- list()
spectrum_rows <- list()
for (a in 1:8) for (region in c("base", "tip")) for (acq in 1:3) {
  sat <- if (region == "base") animals$sat_base[a] else animals$sat_tip[a]
  sat_acq <- sat + rnorm(1, 0, 0.01)                        # small change bet
  brightness <- rnorm(1, 1, 0.05)                           # distance and ang
  raw <- sapply(1:25, function(p) raw_counts(
    brightness * tissue_reflectance(sat_acq + rnorm(1, 0, 0.02), animals$bloo
    lamps[[animals$week[a]]]))
  for (cal in c("stale", "matching")) {
```

```
    white <- if (cal == "matching") white_ref[[animals$week[a]]] else white_ref$week1
    R <- calibrate(raw, white)                                  # 101 bands x 25 pixels
    info <- data.frame(animal = animals$animal[a], group = animals$group[a], week = animals$
                       region = region, acquisition = acq, calibration = cal)
    pixel_rows[[length(pixel_rows) + 1]] <- cbind(info, pixel = 1:25, index = oxygen_index(R
    spectrum_rows[[length(spectrum_rows) + 1]] <- cbind(info, t(rowMeans(R)))
  }
}
pixels <- do.call(rbind, pixel_rows)
spectra <- do.call(rbind, spectrum_rows)
names(spectra)[-(1:6)] <- paste0("nm", wl)
spectra$index <- spectra$nm660 / spectra$nm800
c(pixels_per_calibration = nrow(pixels) / 2, region_means = nrow(spectra) / 2, animals = nro
```

```
#> pixels_per_calibration          region_means               animals
#>                   1200                    48                     8
```

Step 4 – the grouped test. The treatment question concerns the flap tip. We test it twice, once with every pixel as if it were independent and once with one mean per animal, and we do this for both calibrations.

```
tip <- pixels[pixels$region == "tip", ]
grouped_tests <- do.call(rbind, lapply(c("stale", "matching"), function(cal) {
  d <- tip[tip$calibration == cal, ]
  per_animal <- aggregate(index ~ animal + group, data = d, FUN = mean)   # one value per an
  data.frame(calibration = cal,
             n_pixels = nrow(d), p_pixels = t.test(index ~ group, data = d)$p.value,
             n_animals = nrow(per_animal), p_animals = t.test(index ~ group, data = per_anim
             mean_control = mean(per_animal$index[per_animal$group == "control"]),
             mean_treatment = mean(per_animal$index[per_animal$group == "treatment"]))
}))
grouped_tests
```

```
#>   calibration n_pixels      p_pixels n_animals   p_animals mean_control
#> 1       stale      600 5.340962e-173         8 0.001330005    0.7544171
#> 2    matching      600  3.513837e-12         8 0.352695787    0.7544171
#>   mean_treatment
#> 1      0.8313038
#> 2      0.7416595
```

```
stopifnot(
  grouped_tests$n_animals == 8,
  grouped_tests$p_pixels < 0.001,                                  # pixel tests
  grouped_tests$p_animals[grouped_tests$calibration == "stale"] < 0.05,
  grouped_tests$p_animals[grouped_tests$calibration == "matching"] > 0.05
)
```

Two lessons sit in this table. With the stale white reference, the "effect" survives even the correct animal-level test ($p = 0.0013$ with 4 vs 4 animals): the right denominator does not protect against confounding. With the matching reference the mean index of the treated animals is 0.013 index units lower than that of the controls, and the animal-level test gives $p = 0.35$. The pixel-level test still reports $p = 3.5e\text{-}12$ for this chance difference, because it treats 600 tip pixels from eight animals as 600 independent observations. Figure 10.2 shows both levels.

```
acq_tip <- spectra[spectra$region == "tip", ]
animal_tip <- aggregate(index ~ animal + group + calibration, data = acq_tip
facet_names <- c(stale = "Stale white reference (week 1)", matching = "Matchi
acq_tip$calibration <- factor(facet_names[acq_tip$calibration], levels = face
animal_tip$calibration <- factor(facet_names[animal_tip$calibration], levels
ggplot(acq_tip, aes(group, index, colour = group, shape = group)) +
  geom_point(position = position_jitter(width = 0.12, height = 0, seed = 1),
  geom_point(data = animal_tip, size = 3.6, position = position_jitter(width
  facet_wrap(~ calibration) +
  scale_colour_manual(values = unname(cdsr_colours[c("blue", "orange")])) +
```

```
  scale_shape_manual(values = cdsr_shapes[1:2]) +
  labs(x = NULL, y = "Oxygenation index R(660)/R(800) (dimensionless)", colour = NULL, shape
       caption = "Synthetic data · flap tip · n = 8 animals (4 per group), 3 acquisitions ea
  theme_cdsr()
```

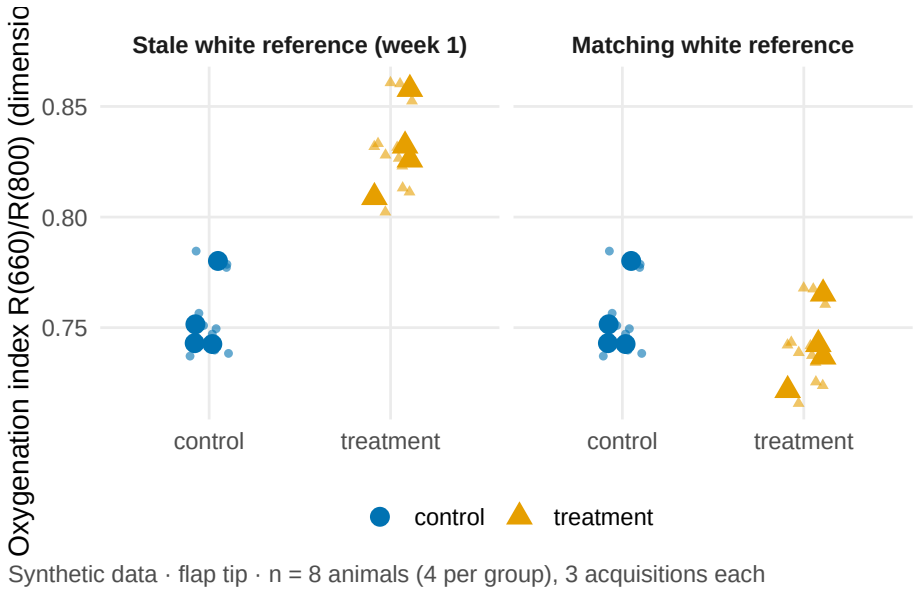


Figure 10.2.: Oxygenation index R(660)/R(800) of the flap tip, dimensionless. Small symbols: region means per acquisition (three per animal); large symbols: animal means. Left: week-2 images calibrated with the stale week-1 white reference; right: with the matching reference. Synthetic data; independent unit: animal (n = 8, four per group).

Step 5 – index against a learning model on new animals. Before a learning model earns a place in the chain, it has to beat the simple index on animals it has never seen. We develop both methods on the four

week-1 animals and test them on the four week-2 animals. The model is a k-nearest-neighbour classifier ($k = 3$) on the first three principal components of the full region-mean spectra; the principal components are computed from the training animals only. The task is to tell flap base from flap tip.

```
library(class)
bands <- paste0("nm", wl)
train <- spectra[spectra$week == 1 & spectra$calibration == "matching", ]
# Index baseline: threshold halfway between the mean index of base and tip in
threshold <- mean(tapply(train$index, train$region, mean))
# Learning model: principal components from training spectra only, then k-nea
pca <- prcomp(train[, bands], center = TRUE)
evaluate <- function(cal) {
  test <- spectra[spectra$week == 2 & spectra$calibration == cal, ]
  pred_index <- ifelse(test$index > threshold, "base", "tip")
  scores_test <- predict(pca, test[, bands])[, 1:3]
  pred_model <- as.character(knn(pca$x[, 1:3], scores_test, cl = train$regio
  data.frame(calibration = cal, test_animals = length(unique(test$animal)), t
             accuracy_index = mean(pred_index == test$region), accuracy_model
}
comparison <- rbind(evaluate("matching"), evaluate("stale"))
stopifnot(length(intersect(unique(train$animal), unique(spectra$animal[spectr
comparison
```

```
#>   calibration test_animals test_spectra accuracy_index accuracy_model
#> 1    matching            4           24            1.0      0.9166667
#> 2       stale            4           24            0.5      0.5000000
```

```
stopifnot(comparison$accuracy_index[1] >= comparison$accuracy_model[1],   # t
          comparison$accuracy_index[2] == 0.5, comparison$accuracy_model[2] =
```

With the matching calibration, the index classifies 100% of the 24 test spectra correctly and the model 92%. The more flexible model brings no

gain here, so the index remains the reference method in the model passport (Chapter 27). With the stale calibration both drop to 50%: every flap tip now looks like a flap base. A learning model does not repair an illumination error; it learns on top of it. The independent test denominator is four animals, not 24 spectra.

i Other implementation: Python with scikit-image and PyTorch

Hyperspectral pipelines are often written in Python. scikit-image (version 0.26.0 in November 2026) provides tested image operators such as filtering, registration and region measurements; most of them accept a `channel_axis` argument, and for a cube stored as rows × columns × bands this must be set to the spectral axis, otherwise the filter mixes neighbouring wavelengths (Walt et al. 2014). PyTorch connects to pretrained networks (Paszke et al. 2019). Networks with ImageNet weights, as in the transfer learning tutorial, expect three colour channels normalised with ImageNet means; a 101-band cube needs a new first layer, whose weights then cannot be transferred. Pretrained hyperspectral weights from research code are only usable if wavelength grid, calibration (reflectance, not raw counts), spatial resolution and band order match. For a fair comparison with this lab, both implementations must receive the identical calibrated reflectance table, the same animal-wise split (week-1 animals for development, week-2 animals for testing), the same index definition and the same accuracy computed per spectrum and per animal. Status: conceptual comparison, not run.

10.4. Counter-check

Predict first, then compute. Suppose the replacement lamp had been only brighter, with exactly the same spectral shape as the original. Would the stale white reference still inflate the index? And would the grey check tile

still fail?

```
set.seed(3010)
lamp_C <- 1.2 * lamp_A                                  # brighter only, same
raw_tip <- raw_counts(tissue_reflectance(0.45, 1.5), lamp_C)
index_stale_brightness <- oxygen_index(calibrate(raw_tip, white_ref$week1))
grey_C <- rowMeans(calibrate(sapply(1:100, function(i) raw_counts(rep(0.5, le
c(true_index = oxygen_index(tissue_reflectance(0.45, 1.5)),
  stale_reference_brighter_lamp = index_stale_brightness,
  grey_index = oxygen_index(grey_C), grey_mean_reflectance = mean(grey_C))
```

```
#>                    true_index stale_reference_brighter_lamp
#>                     0.7499384                     0.7407447
#>                    grey_index         grey_mean_reflectance
#>                     1.0003271                     0.6001668
```

```
stopifnot(abs(index_stale_brightness / oxygen_index(tissue_reflectance(0.45,
          abs(oxygen_index(grey_C) - 1) <= 0.02, mean(grey_C) > 0.55)
```

The ratio index survives a pure change in brightness: it stays within a few per cent of the true value, and the grey tile's index stays within tolerance. The grey tile's mean reflectance, however, rises to about 0.60 instead of 0.50, so the check still flags a problem with absolute reflectance. A change in the lamp's colour affects the index; a change in brightness affects absolute values. The grey tile should be read on both quantities.

> **Common mistake**
>
> "A ratio index needs no calibration." A ratio cancels everything that scales all wavelengths alike, such as distance, exposure time or a brighter lamp of the same colour. It does not cancel anything that changes the spectral shape: a new lamp type, an ageing bulb, daylight from a window or a different filter. A white reference recorded under

> the actual light remains necessary.

A design counter-check needs no code: if animals of both groups had been imaged in alternating order in both weeks, as ARRIVE item 4b recommends for the order of measurements (Percie du Sert et al. 2020), the lamp change would have shifted both groups alike and could not have produced a group difference.

10.5. Consequences for patients and operations

Wrong tissue interpretation. The same artefact in a clinical setting would make poorly oxygenated tissue look better oxygenated than it is. A surgeon judging a flap by such a value might delay a revision. The measurement assumption that protects against this is checkable: *a grey reference tile imaged at every session has an index of 1.00 ± 0.02 and a mean reflectance of 0.50 ± 0.01.* The failure mode is equally concrete: a lamp change without a new white reference shifts every index of that session by the same factor.

Animals and the 3Rs. A false positive treatment effect in an animal study has consequences beyond the abstract. It invites follow-up experiments, and therefore more animals, to pursue an effect that was never there. Reduction in the sense of the 3Rs starts with analyses that use the animal as the unit and designs that avoid confounding (NC3Rs 2026). Reporting the experimental unit and the order of measurements is part of ARRIVE (Percie du Sert et al. 2020).

Regulation. In this preclinical study the camera is a research instrument. If the same index were offered to clinicians to assess flap perfusion, its intended purpose would decide whether the software is a medical device and which class applies; Chapter 22 works through such an intended purpose.

Data protection and security. Animal data are not personal data. In a human study, wound and skin images can show tattoos, faces or rare findings and become identifiable; Chapter 23 covers this. The white reference and calibration files are part of the measurement. Storing them with a version and a checksum, so that the analysis can prove which file it used, is an integrity control of the kind discussed in Chapter 24.

10.6. Look up and transfer

- **Reviewer question:** The manuscript reports $p < 0.001$ for a difference in an oxygenation index between two groups of animals. What was the experimental unit, were both groups imaged in the same sessions, and which reference target shows that the calibration was stable across sessions?
- **Self-test:** A follow-up study images six animals per group, with four acquisitions per animal and 50 pixels per region of interest. A colleague proposes to divide every spectrum by its own mean, "to remove the lamp". How many independent units does the group comparison have, does the proposal remove a change in lamp colour, and which design change prevents the confounding in this chapter?

Answer

The group comparison has 12 independent units, six animals per group; the 2,400 pixels per region and the 48 acquisitions are repeated measurements within animals. Dividing by the mean removes a change in brightness but not a change in spectral shape, so a lamp with a different colour still distorts the ratio $R(660)/R(800)$; only a white reference recorded under the same lamp corrects it. Imaging animals of both groups in alternating or randomised order within every session prevents the lamp from coinciding with the group, and a grey reference tile at every session detects a calibration problem

before any animal is analysed.

- **Annotated sources:**
  - Lu and Fei (2014) introduces the hypercube, calibration and medical applications of hyperspectral imaging.
  - Holmer et al. (2018) describes how real perfusion parameters are computed from spectra and why the chapter's index is only a simplification.
  - Lazic (2010) explains pseudoreplication with examples and is the basis for testing per animal.
  - Percie du Sert et al. (2020) (items 1b and 4b) asks for the experimental unit and the order of measurements to be reported.
  - NC3Rs (2026) frames the consequence of a false positive result for further animal use.
  - Walt et al. (2014) and Paszke et al. (2019) document the Python implementations named in the box.

10.7. Further reading

- Chapter 19 – what happens when images of known animals end up in the test set.
- Chapter 16 – how much uncertainty remains with four animals per group.
- Chapter 27 – the model passport, where the index belongs as the reference method.
- Methods in R: Principal Component Analysis – the method behind the features of the learning model.
- CTTIR Courses: PCA, factor analysis, CCA, LDA – a lab on dimension reduction for many correlated variables such as spectral bands.
- Methods in R: Repeated-Measures Designs – how repeated acquisitions per animal enter the design.

- CTTIR Courses: Data types, tidy data, accuracy and precision – the difference between a precise and an accurate measurement, which the lamp error illustrates.

11. Measuring gait from video

R. Heller · Date: 2025-05-18 · Last changed: 2026-11-09

At a glance

Guiding question: Are two measurement methods similar enough for this rehabilitation decision?
For: physiotherapists, rehabilitation and trauma clinicians, and researchers who compare a new measurement method with an established one. · **Prerequisites:** repeated measurements and the unit of analysis from Chapter 3; the measurement chain from Chapter 10. **In the course:** W4 · What imaging, gait and tissue models measure (gait station) and W5 · Causal questions, trajectories and phenotypes · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

11.1. Case and decision

A physiotherapist in a fictitious rehabilitation clinic follows patients after surgery for a tibial plateau fracture. One criterion for moving a patient to the next stage of the gait programme is the peak knee flexion during the swing phase. Until now it has been measured in the gait laboratory with a marker-based motion capture system, which is booked out for weeks. A smartphone video analysed by a pose-estimation model would take five minutes on the ward. In the first trial recordings the video values are

systematically lower for some patients. Looking at the setup, the physiotherapist notices that the phone stood at an angle to the walkway because the corridor is narrow. She has to decide whether the video method may replace the laboratory measurement for this decision, and under which conditions. The first error in this chapter is the perspective.

11.2. Core reading

> **Core reading** · recommended reading for W4 (gait station) · about 10 minutes · ends before ‘Lab’

From pixels to degrees. A pose-estimation model is a learning model that finds anatomical landmarks in each video frame, for example hip, knee and ankle, and returns their pixel coordinates with a confidence score (Mathis et al. 2018). Everything after that is geometry. The knee angle at one moment is the angle between the thigh (knee to hip) and the shank (knee to ankle). Knee flexion is 180° minus that angle: 0° for a straight leg. Over a gait cycle the flexion rises to a peak in the swing phase; that peak, in degrees, is the quantity the clinic wants. Figure 11.1 shows the whole chain and where the learning model sits.

```
library(ggplot2)
# Generic drawing of a measurement chain: one box per step, a note per step,
draw_chain <- function(steps, note_title = "Unit or reference") {
  steps$y <- rev(seq_len(nrow(steps)))
  types <- c("measurement", "reference", "model", "decision")
  ggplot(steps) +
    geom_segment(data = steps[-nrow(steps), ], aes(x = 1, xend = 1, y = y - 0
                 arrow = grid::arrow(length = grid::unit(0.15, "cm")), colour
    geom_tile(aes(x = 1, y = y, fill = type, linetype = type), width = 1.7, h
              colour = "grey20", linewidth = 0.5) +
    geom_text(aes(x = 1, y = y, label = label), size = 3.4) +
```

```
    geom_text(aes(x = 2.0, y = y, label = note), hjust = 0, size = 3.1, colour = "grey15") +
    annotate("text", x = 2.0, y = nrow(steps) + 0.55, label = note_title, hjust = 0, fontfac
             size = 3.2) +
    scale_fill_manual(values = c(measurement = "white", reference = "grey85", model = "#D9EC
                                 decision = "grey70"), breaks = types, name = NULL) +
    scale_linetype_manual(values = c(measurement = "solid", reference = "solid", model = "da
                                     decision = "solid"), breaks = types, name = NULL) +
    scale_x_continuous(limits = c(0.1, 4.6)) +
    scale_y_continuous(limits = c(0.5, nrow(steps) + 0.8)) +
    theme_void(base_size = 11) +
    theme(legend.position = "bottom", plot.caption = element_text(hjust = 0, colour = "#4D4D
}
chain_b11 <- data.frame(
  label = c("Patient walking", "Camera: position, frame rate", "Video frames", "Pose estimat
            "Knee angle from 3 points", "Peak flexion per step", "Mean per patient", "Refere
  note = c("knee flexion in the sagittal plane", "perpendicular to walkway? height? distance
           "pixels; face and body visible", "hip, knee, ankle in pixels + confidence",
           "degrees; 0 = straight leg", "degrees; 10 steps per patient", "independent unit:
           "marker-based motion capture; tolerance +/-5 degrees"),
  type = c("measurement", "measurement", "measurement", "model", "measurement", "measurement
           "decision")
)
draw_chain(chain_b11) +
  labs(caption = "Synthetic case · independent unit: patient (n = 24)")
```

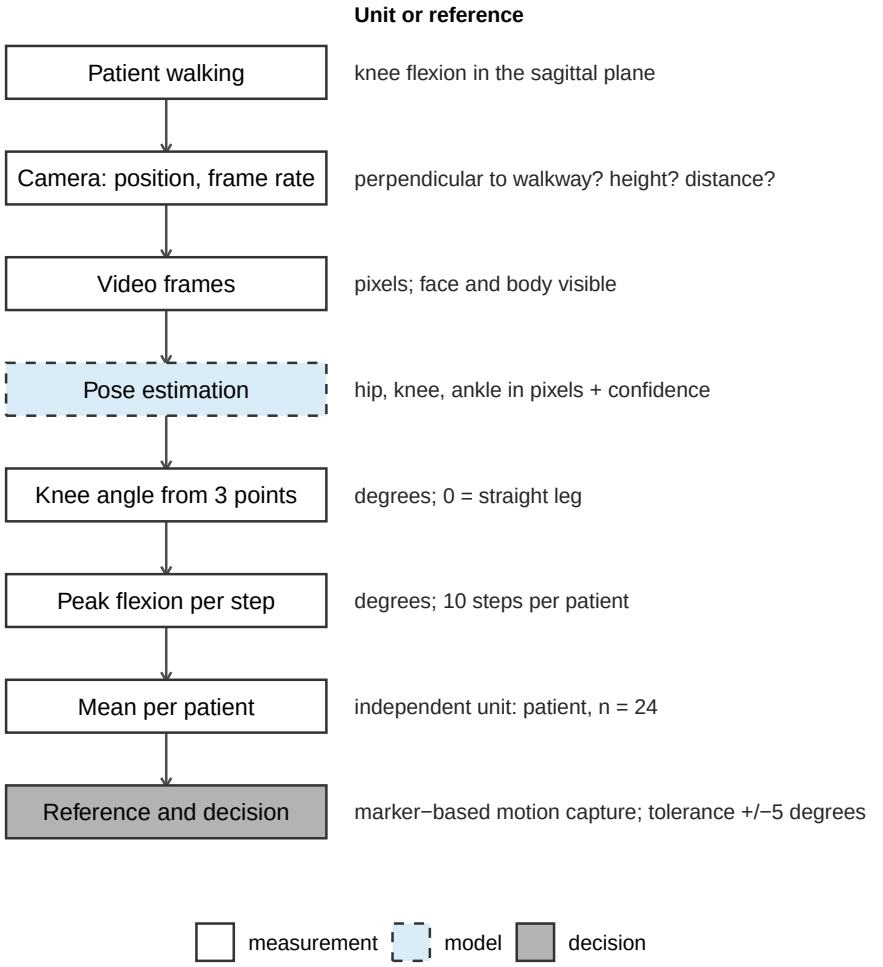


Figure 11.1.: Measurement chain for knee flexion from video. The dashed box marks the learning model (pose estimation); the dark grey box marks the comparison with the reference method and the decision. Unit: degrees from step 5 onwards. Independent unit for the agreement question: patient (n = 24 in the lab, ten steps each). Synthetic case.

```
# Geometry used in the core reading (the lab explains the same functions step by step)
knee_flexion <- function(hip, knee, ankle) {
  hip <- matrix(hip, ncol = 2); knee <- matrix(knee, ncol = 2); ankle <- matrix(ankle, ncol
  v1 <- hip - knee; v2 <- ankle - knee
  cos_angle <- rowSums(v1 * v2) / sqrt(rowSums(v1^2) * rowSums(v2^2))
  180 - acos(pmin(pmax(cos_angle, -1), 1)) * 180 / pi
}
landmarks <- function(hip_flex, knee_flex, thigh = 0.43, shank = 0.42) {
  r <- pi / 180
  knee <- cbind(thigh * sin(hip_flex * r), -thigh * cos(hip_flex * r))
  ankle <- knee + cbind(shank * sin((hip_flex - knee_flex) * r), -shank * cos((hip_flex - kn
  list(hip = cbind(0 * hip_flex, 0 * hip_flex), knee = knee, ankle = ankle)
}
project <- function(p, yaw) cbind(p[, 1] * cos(yaw * pi / 180), p[, 2])
seen <- function(knee_flex, yaw, hip_flex = 20) {
  L <- landmarks(hip_flex, knee_flex)
  knee_flexion(project(L$hip, yaw), project(L$knee, yaw), project(L$ankle, yaw))
}
proj_30_60 <- seen(60, 30)
```

The perspective error. A single camera sees a flat projection of a three-dimensional movement. If it looks at the walkway at a right angle, the sagittal plane is projected without distortion. If the phone is turned by a yaw angle, all forward-backward distances in the image shrink by the cosine of that angle while vertical distances stay. Thigh and shank then look more vertical than they are, and the knee looks straighter. Figure 11.2 computes this for a simplified setup without lens distortion: at a yaw of 30°, a true flexion of 60° appears as about 53.5°. The error is systematic, not random. Repeating the recording does not average it away, and it is larger in patients who bend the knee more.

```
grid_yaw <- expand.grid(yaw = 0:45, true_flexion = c(40, 60))
grid_yaw$error <- mapply(function(f, y) seen(f, y) - f, grid_yaw$true_flexion
grid_yaw$true_flexion <- factor(paste0(grid_yaw$true_flexion, "°"))
ggplot(grid_yaw, aes(yaw, error, linetype = true_flexion, colour = true_flexi
  geom_hline(yintercept = c(-5, 0), colour = "grey50", linetype = c("dotted"
  geom_line(linewidth = 0.9) +
  annotate("text", x = 1, y = -5.6, label = "tolerance −5°", hjust = 0, size
  scale_colour_manual(values = unname(cdsr_colours[c("blue", "red")])) +
  scale_linetype_manual(values = cdsr_linetypes[1:2]) +
  labs(x = "Camera yaw relative to the perpendicular (degrees)", y = "Video r
       colour = "True knee flexion", linetype = "True knee flexion",
       caption = "Synthetic geometry · thigh 0.43 m, shank 0.42 m · orthograp
  theme_cdsr()
```

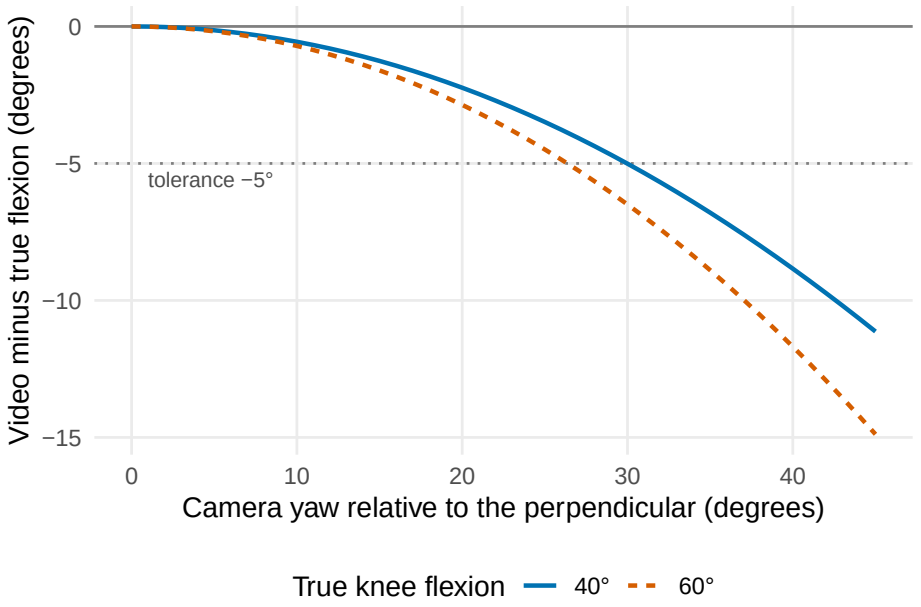


Figure 11.2.: Systematic error of video-based knee flexion when the camera is turned away from the perpendicular, for true flexion of 40° and 60° (hip flexion 20°; simplified orthographic projection). Unit: degrees. Synthetic geometry, no patients involved.

What the pose model knows. The pose-estimation model was trained on labelled images, most of them of people in everyday scenes. It has not necessarily seen crutches, knee braces, compression stockings or a far leg hidden behind the near one (occlusion). Each of these can shift the knee landmark by a centimetre or two, and with a thigh of about 43 cm that already moves the angle by a few degrees. A low confidence score is a warning, not a correction. The model is the only learned step in the chain; everything after it is geometry that can be checked by hand, as the lab does with two angles worked out on paper.

Reference and unit. The comparison method is the laboratory's marker-based motion capture. It is itself a measurement with its own error, not the truth, so an agreement study compares two methods and assumes neither to be perfect. The physical unit is the degree throughout. The independent unit is the patient: 24 patients with ten steps each in the lab.

Agreement, not correlation. To decide whether the video can replace the laboratory, the clinic records the same steps with both methods and looks at the differences. The Bland–Altman approach plots the difference of the two methods against their mean, estimates the average difference (bias) and the limits of agreement, within which about 95% of differences are expected (Bland and Altman 1986). A high correlation would say nothing useful: two methods can correlate perfectly while one reads 10° lower. Details are in Methods in R: Bland-Altman Limits of Agreement.

Set the tolerance before looking. Whether the limits of agreement are good enough is a clinical judgement, not a statistical one, and it has to be made before the data are seen. In the fictitious protocol the progression criteria are set in steps of 10°, so the team decides that a patient's value may differ from the laboratory value by at most half a step. **The key assumption is a tolerance of ±5° for the value on which the decision is based** (*assumption*, to be agreed locally). Published comparisons of two-dimensional video with motion capture report mean absolute errors of several degrees for sagittal knee angles, so the question is realistic [Stenum et al. (2021); *published*].

Which value is the decision based on? Each patient contributes ten steps. Steps of the same patient are not independent: they share the person, the camera setup and the clothing. The independent unit is the patient. Two analyses are legitimate, and they answer different questions. If the decision uses a single step, the limits of agreement must describe single steps; with repeated steps per patient they are estimated from the within-patient and between-patient variance components (Bland and Altman 2007). If the decision uses the mean of ten steps, the limits

must describe ten-step means, one per patient. Pooling all 240 steps as if they came from 240 patients is not legitimate: it makes the uncertainty of the bias look smaller than it is. Which level applies, and how many steps a recording must contain, is written into the protocol before the comparison starts.

Question for the workshop: Your clinic uses the mean of ten steps. Which of the two sets of limits of agreement belongs in the clinical protocol, and what would have to change in the protocol if the ward only recorded three steps?

11.3. Lab

The lab simulates 24 fictitious patients with ten steps each, measured simultaneously by a marker-based reference and by a video method whose landmarks carry noise and a small systematic offset. It then checks agreement at both decision levels.

Step 1 – the angle from three points. The function accepts single points or one row per step and is checked against two angles that can be worked out by hand.

```
# Knee flexion in degrees from 2D hip, knee and ankle coordinates (x forward, z up); 0 = str
knee_flexion <- function(hip, knee, ankle) {
  hip <- matrix(hip, ncol = 2); knee <- matrix(knee, ncol = 2); ankle <- matrix(ankle, ncol
  thigh_vec <- hip - knee                                      # from knee to hip
  shank_vec <- ankle - knee                                    # from knee to ankle
  cos_angle <- rowSums(thigh_vec * shank_vec) / sqrt(rowSums(thigh_vec^2) * rowSums(shank_ve
  180 - acos(pmin(pmax(cos_angle, -1), 1)) * 180 / pi       # flexion = 180 degrees minus th
}
stopifnot(
  abs(knee_flexion(c(0, 0), c(0, -1), c(0, -2)) - 0) < 1e-9,    # straight leg: 0 degrees
  abs(knee_flexion(c(0, 0), c(0, -1), c(-1, -1)) - 90) < 1e-9   # shank horizontal: 90 degre
```

```
)
knee_flexion(c(0, 0), c(0.15, -0.40), c(-0.20, -0.65))          # an arbitrar
```

```
#> [1] 75.01837
```

Step 2 – synthetic patients and steps. Landmarks are generated in the sagittal plane, projected for the camera's yaw, offset and blurred like a pose-estimation output, and then turned back into angles.

```
set.seed(1111)
n_patients <- 24
n_steps <- 10
patients <- data.frame(
  patient = sprintf("P%02d", 1:n_patients),
  mean_flexion = rnorm(n_patients, 55, 7),     # true peak swing flexion per p
  hip_flexion = rnorm(n_patients, 20, 3),      # hip flexion at that moment (c
  yaw = runif(n_patients, -5, 5),              # camera almost perpendicular
  knee_offset = rnorm(n_patients, 0.004, 0.006), # knee landmark placed sligh
  size = rnorm(n_patients, 1, 0.05)            # body size factor for segment
)
# Landmarks in metres for given hip and knee flexion; projection for a camera
landmarks <- function(hip_flex, knee_flex, thigh = 0.43, shank = 0.42) {
  r <- pi / 180
  knee <- cbind(thigh * sin(hip_flex * r), -thigh * cos(hip_flex * r))
  ankle <- knee + cbind(shank * sin((hip_flex - knee_flex) * r), -shank * cos
  list(hip = cbind(0 * hip_flex, 0 * hip_flex), knee = knee, ankle = ankle)
}
project <- function(p, yaw) cbind(p[, 1] * cos(yaw * pi / 180), p[, 2])
simulate_steps <- function(patients, yaw) {
  rows <- lapply(seq_len(nrow(patients)), function(i) {
    true <- patients$mean_flexion[i] + rnorm(n_steps, 0, 2.5)        # step-t
    hip <- patients$hip_flexion[i] + rnorm(n_steps, 0, 2)
```

```
    L <- landmarks(hip, true, 0.43 * patients$size[i], 0.42 * patients$size[i])
    L$knee[, 1] <- L$knee[, 1] + patients$knee_offset[i]              # systematic landmark
    blur <- function() matrix(rnorm(2 * n_steps, 0, 0.012), ncol = 2)  # landmark noise, met
    video <- knee_flexion(project(L$hip, yaw[i]) + blur(), project(L$knee, yaw[i]) + blur(),
                          project(L$ankle, yaw[i]) + blur())
    data.frame(patient = patients$patient[i], step = 1:n_steps, true = true,
               reference = true + rnorm(n_steps, 0, 1),               # marker-based referen
               video = video)
  })
  steps <- do.call(rbind, rows)
  steps$difference <- steps$video - steps$reference
  steps
}
steps <- simulate_steps(patients, patients$yaw)
gait_file <- steps[, c("patient", "step", "reference", "video")]
gait_file[, c("reference", "video")] <- round(gait_file[, c("reference", "video")], 2)    # d
write.csv(gait_file, "../data/b11-gait-steps.csv", row.names = FALSE)
c(patients = length(unique(steps$patient)), steps = nrow(steps))
```

```
#> patients    steps
#>       24      240
```

The file `data/b11-gait-steps.csv` holds the synthetic reference and video values for the gait station in W4; the true values stay in the simulation.

Step 3 – agreement at both decision levels. The single-step limits use the variance components of Bland and Altman's method for repeated measurements in which the true value varies (Bland and Altman 2007). The ten-step limits use one mean difference per patient.

```
tolerance <- 5                                          # degrees, set before the
# (a) Single steps: one-way ANOVA of the differences by patient
fit <- anova(lm(difference ~ patient, data = steps))
ms_within <- fit["Residuals", "Mean Sq"]
ms_between <- fit["patient", "Mean Sq"]
var_between <- max(0, (ms_between - ms_within) / n_steps)
bias <- mean(steps$difference)
sd_single <- sqrt(var_between + ms_within)
# (b) Ten-step means: one value per patient
per_patient <- aggregate(cbind(reference, video, difference) ~ patient, data
sd_mean10 <- sd(per_patient$difference)
# Confidence interval of the bias: patients are the independent unit (n = 24)
ci_patients <- bias + c(-1, 1) * qt(0.975, n_patients - 1) * sd_mean10 / sqrt
ci_steps_wrong <- bias + c(-1, 1) * qt(0.975, nrow(steps) - 1) * sd(steps$di
agreement <- data.frame(
  decision_level = c("single step", "mean of 10 steps"),
  bias = bias, lower_loa = bias - 1.96 * c(sd_single, sd_mean10), upper_loa =
)
agreement$within_tolerance <- agreement$lower_loa >= -tolerance & agreement$u
agreement
```

```
#>     decision_level     bias lower_loa upper_loa within_tolerance
#> 1      single step 1.290632 -6.030212  8.611475            FALSE
#> 2 mean of 10 steps 1.290632 -1.395141  3.976405             TRUE
```

```
rbind(ci_bias_patients = ci_patients, ci_bias_steps_wrong = ci_steps_wrong)
```

```
#>                           [,1]     [,2]
#> ci_bias_patients    0.7120080 1.869255
#> ci_bias_steps_wrong 0.8160208 1.765243
```

```
stopifnot(agreement$within_tolerance == c(FALSE, TRUE),
          diff(ci_patients) > diff(ci_steps_wrong))      # pooling steps understates the unc
```

For single steps the limits of agreement run from -6° to 8.6° and exceed the tolerance. For the ten-step mean they run from -1.4° to 4° and lie within it. The video method, recorded perpendicular to the walkway, is acceptable for a decision based on ten steps and not for one based on a single step. The confidence interval for the bias based on 24 patients (0.7° to 1.9°) is wider than the one obtained by pretending that 240 steps were 240 patients (0.8° to 1.8°).

> **Common mistake**
>
> Averaging the ten steps per patient first and then quoting the resulting limits of agreement as if they applied to single measurements. Limits of agreement belong to the quantity on which they were computed: limits for ten-step means are narrower than limits for single steps and must not be used for a protocol that relies on one step.

```
steps$mean_methods <- (steps$video + steps$reference) / 2
per_patient$mean_methods <- (per_patient$video + per_patient$reference) / 2
lines_ba <- data.frame(y = c(bias, agreement$lower_loa[2], agreement$upper_loa[2], agreement
                             agreement$upper_loa[1]),
                       type = c("bias", "limits, 10-step mean", "limits, 10-step mean", "lim
                                "limits, single step"))
ggplot() +
  annotate("rect", xmin = -Inf, xmax = Inf, ymin = -tolerance, ymax = tolerance, fill = "gre
  geom_point(data = steps, aes(mean_methods, difference), colour = "grey55", size = 0.9) +
  geom_point(data = per_patient, aes(mean_methods, difference), colour = cdsr_colours["blue"
             shape = cdsr_shapes[1]) +
  geom_hline(data = lines_ba, aes(yintercept = y, linetype = type)) +
  scale_linetype_manual(values = c("bias" = "solid", "limits, 10-step mean" = "dashed",
                                   "limits, single step" = "dotted"), name = NULL) +
```

```
  labs(x = "Mean of video and reference (degrees)", y = "Video minus referenc
       caption = "Synthetic data · n = 24 patients, 10 steps each · tolerance
  theme_cdsr()
```

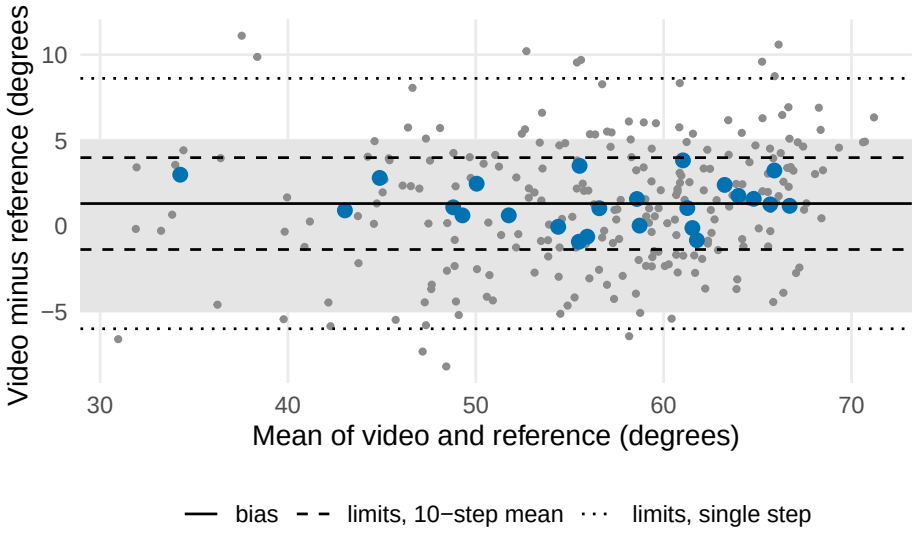


Figure 11.3.: Bland–Altman plot for peak knee flexion, video minus marker-based reference, in degrees. Large symbols: ten-step means per patient (n = 24 independent patients); small grey dots: single steps (240). Solid line: bias; dashed lines: limits of agreement for ten-step means; dotted lines: limits for single steps; shaded band: pre-specified tolerance of ±5° (assumption). Synthetic data.

Other implementation: Python with PyTorch

Pose estimation itself is rarely written in R. The PyTorch library torchvision offers a Keypoint R-CNN model with weights trained on the COCO person keypoints (`KeypointRCNN_ResNet50_FPN_Weights.COCO_V1`), which returns 17 keypoints per person, including both hips, knees and ankles, as `[x, y, visibility]` in pixel coordinates; inputs are expected as tensors scaled to the range 0–1 (Paszke et al. 2019). For landmarks that COCO does not define, such as markers on a brace, DeepLabCut trains a network on a few hundred hand-labelled frames (Mathis et al. 2018). The concrete advantage is a tested detector with known weights instead of hand-placed points. To compare it fairly with this lab, the following must be identical and recorded: the video frames (frame rate, resolution and the phone's rotation metadata), the mapping from keypoint index to hip, knee and ankle of the correct side, the handling of low-confidence keypoints, the definition of a step and its peak, the averaging over ten steps, the metric (difference to the reference in degrees) and the tolerance of ±5°. Weights and software version (torchvision 0.29.2 or DeepLabCut 3.0.2 in November 2026) belong in the model passport. Status: conceptual comparison, not run.

11.4. Counter-check

Now reproduce the error from the case. For the first eight patients the phone stands at 30° to the walkway; the remaining sixteen are recorded as before. Predict first: does the bias move up or down, and do the ten-step limits of agreement stay within ±5°?

```
set.seed(2111)
yaw_corridor <- patients$yaw
yaw_corridor[1:8] <- 30                                  # narrow corridor: camera
steps_corridor <- simulate_steps(patients, yaw_corridor)
pp_corridor <- aggregate(difference ~ patient, data = steps_corridor, FUN = m
pp_corridor$setup <- ifelse(pp_corridor$patient %in% patients$patient[1:8], "
bias_by_setup <- tapply(pp_corridor$difference, pp_corridor$setup, mean)
loa_corridor <- mean(pp_corridor$difference) + c(-1.96, 1.96) * sd(pp_corrido
round(c(bias_by_setup, lower_loa = loa_corridor[1], upper_loa = loa_corridor
```

```
#>     perpendicular turned 30 degrees         lower_loa         upper_loa
#>                 1.1              -4.9              -7.2               5.5
```

```
stopifnot(bias_by_setup["turned 30 degrees"] < bias_by_setup["perpendicular"]
          loa_corridor[1] < -tolerance)
```

The bias in the turned setup is -4.9° against 1.1° with a perpendicular camera, and the lower limit of agreement of the ten-step means falls to -7.2°, outside the tolerance. The perspective error is systematic, so averaging steps does not remove it. The plausible failure mode is therefore concrete: *a phone placed at an angle underestimates knee flexion, most in patients with the largest flexion, and may hold patients back in the programme.* The checkable measurement assumption is equally concrete: *the camera axis is perpendicular to the walkway within ±10°, verified with a floor marking and a photo of the setup for every recording.*

11.5. Consequences for patients and operations

An identifiable movement pattern. A gait video shows face, body and clothing, and it is health data. Removing the face is not enough: gait patterns are so individual that learning models can attribute them to single

persons from joint angles and forces alone [Horst et al. (2019); *published*]. The derived keypoint trajectories may therefore still identify a patient, especially combined with ward and date. The General Data Protection Regulation (EU) 2016/679 (GDPR) defines biometric data as personal data resulting from specific technical processing of physical, physiological or behavioural characteristics "which allow or confirm the unique identification" of a person (Article 4(14)); Article 9(1) lists biometric data "for the purpose of uniquely identifying a natural person" and data concerning health among the special categories [European Parliament and Council of the European Union (2016); binding text]. Whether the keypoint data in this case fall under the biometric category is a question of interpretation for the data protection officer; that in this clinical context they are data concerning health (Article 4(15)) is hardly in doubt. The decision the team has to take is practical: does the clinic need the raw video for its purpose, for example to audit implausible angles, and if so, for how long? Chapter 23 follows exactly this data flow from video to angle and deletion point.

Regulation. If the clinic used the app's output to decide a patient's progression, the intended purpose stated by the manufacturer would determine whether the software is a medical device and of which class (Rule 11 of Annex VIII of the MDR (European Parliament and Council of the European Union 2017a)). An app marketed for "fitness tracking" and used for a clinical decision raises its own question. Chapter 22 works through such intended purposes.

Ethics and operations. A systematic underestimation does not hit everyone alike: it is largest in patients with the largest flexion and in wards with the narrowest corridors. The patient sees only "not yet". A short note in the protocol on camera setup, and a re-measurement in the laboratory when the video value lies within 5° of a decision threshold, keep the faster method without shifting its errors onto patients.

11.6. Look up and transfer

- **Reviewer question:** The study reports limits of agreement of ±4° from 600 steps of 30 participants. Were the limits computed for single steps or for participant means, how were repeated steps handled, and was the tolerance defined before the analysis?
- **Self-test:** The ward wants to record only three steps per patient. Using the variance components from the lab, estimate the limits of agreement for a three-step mean. Is the video method still acceptable at a tolerance of ±5°?

Answer

For the mean of m steps the variance of a patient's mean difference is the between-patient component plus the within-patient component divided by m. With the lab's values (between-patient variance 0.54, within-patient variance 13.41, both in squared degrees) a three-step mean has a standard deviation of 2.24° and limits of agreement from -3.1° to 5.7°. In R: `bias + c(-1.96, 1.96) * sqrt(var_between + ms_within / 3)`. The upper limit exceeds +5°, so at this tolerance the video method is not acceptable for a three-step protocol, although it is for ten steps. A protocol based on three steps needs its own agreement statement and must not reuse the ten-step limits. The estimate also assumes that three steps on the ward vary like ten steps in the laboratory, which itself is an assumption to check.

- **Annotated sources:**
  - Bland and Altman (1986) introduced limits of agreement and explains why correlation is the wrong tool.
  - Bland and Altman (2007) provides the variance-components method for several observations per person used in the lab.
  - Stenum et al. (2021) compares two-dimensional video pose estimation with motion capture on a stride-by-stride basis and

shows realistic error sizes for sagittal knee angles.
  - Mathis et al. (2018) describes markerless pose estimation of user-defined landmarks with a learning model.
  - Horst et al. (2019) supports the statement that individual gait patterns are distinctive enough to identify people.
  - European Parliament and Council of the European Union (2016) is the binding text for the definitions of biometric and health data.

11.7. Further reading

- Chapter 23 – the data flow from gait video to angle, service provider and deletion point.
- Chapter 13 – repeated measurements over time and the person as the unit of analysis.
- Chapter 10 – the measurement chain of a second imaging method, where the first error is the light.
- Methods in R: Bland-Altman Plots – how to draw and read the difference plot.
- CTTIR Courses: Kappa, ICC, Bland–Altman – a lab that sets agreement next to reliability.
- CTTIR Tutorials: Bland-Altman Limits of Agreement – a short tutorial for revision.
- Methods in R: Power for Bland-Altman Studies – how many patients an agreement study needs.

12. Spatial biology and image analysis

R. Heller · Date: 2026-01-11 · Last changed: 2026-11-09

At a glance

Guiding question: How does a segmentation decision change the biological conclusion?
For: researchers in pathology and translational immunology, doctoral researchers who analyse multiplex tissue images, and reviewers of spatial analyses. · **Prerequisites:** the unit of analysis from Chapter 3; the measurement chain and pseudoreplication from Chapter 10.
In the course: W4 · What imaging, gait and tissue models measure (tissue station) and W5 · Causal questions, trajectories and phenotypes · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

12.1. Case and decision

A research pathologist in a fictitious translational group analyses multiplex immunofluorescence images of soft tissue taken from the fracture site during surgery. Twelve patients are included: six whose fractures healed on schedule and six with delayed healing. The hypothesis for a grant application is that T cells and macrophages are more often in direct contact

in delayed healing. The image-analysis pipeline finds nuclei with a pretrained segmentation model and then grows every nucleus outwards by a fixed distance to approximate the cell boundary. The lab's default is 4 µm; a colleague's older pipeline used 1 µm. With 4 µm, the share of T cells touching a macrophage is clearly higher in delayed healing, and a test over all T cells gives $p < 0.01$. With 1 µm, the difference almost disappears. The pathologist has to decide which result goes into the application, and what its sample size is. The first error in this chapter is the segment boundary.

12.2. Core reading

> **Core reading** · recommended reading for W4 (tissue station) · about 10 minutes · ends before 'Lab'

A chain of decisions. A spatial tissue analysis is a long measurement chain (Figure 12.1). Tissue is cut into sections a few micrometres thick and stained with several antibodies at once, for example a nuclear stain, CD3 for T cells and CD68 for macrophages. A scanner records each channel as an image whose pixel size is given in micrometres. A learning model, typically a pretrained network, then segments nuclei (Stringer et al. 2021; Greenwald et al. 2022). Because membranes are harder to see than nuclei, many pipelines approximate each cell by growing its nucleus outwards by a fixed distance. Marker intensities are averaged within this boundary, a threshold assigns a phenotype, and every cell becomes a point with a type. Only then does the neighbourhood analysis begin (Schapiro et al. 2017). Every step contains a parameter, and the boundary expansion is one of the least visible.

```
library(ggplot2)
# Generic drawing of a measurement chain: one box per step, a note per step,
draw_chain <- function(steps, note_title = "Unit or reference") {
```

```
  steps$y <- rev(seq_len(nrow(steps)))
  types <- c("measurement", "reference", "model", "decision")
  ggplot(steps) +
    geom_segment(data = steps[-nrow(steps), ], aes(x = 1, xend = 1, y = y - 0.24, yend = y -
                 arrow = grid::arrow(length = grid::unit(0.15, "cm")), colour = "grey30") +
    geom_tile(aes(x = 1, y = y, fill = type, linetype = type), width = 1.7, height = 0.5,
              colour = "grey20", linewidth = 0.5) +
    geom_text(aes(x = 1, y = y, label = label), size = 3.4) +
    geom_text(aes(x = 2.0, y = y, label = note), hjust = 0, size = 3.1, colour = "grey15") +
    annotate("text", x = 2.0, y = nrow(steps) + 0.55, label = note_title, hjust = 0, fontfac
             size = 3.2) +
    scale_fill_manual(values = c(measurement = "white", reference = "grey85", model = "#D9EC
                                 decision = "grey70"), breaks = types, name = NULL) +
    scale_linetype_manual(values = c(measurement = "solid", reference = "solid", model = "da
                                     decision = "solid"), breaks = types, name = NULL) +
    scale_x_continuous(limits = c(0.1, 4.6)) +
    scale_y_continuous(limits = c(0.5, nrow(steps) + 0.8)) +
    theme_void(base_size = 11) +
    theme(legend.position = "bottom", plot.caption = element_text(hjust = 0, colour = "#4D4D
}
chain_b12 <- data.frame(
  label = c("Tissue from fracture site", "Section and staining", "Scanner images", "Nucleus
            "Cell boundary (expansion)", "Phenotype by threshold", "Contact measure", "Summa
            "Group comparison, decision"),
  note = c("12 patients, 3 fields each", "section thickness in µm; staining run = batch",
           "pixel size in µm per pixel; one channel per marker", "reference: manual annotati
           "expansion e in µm (1 to 5 in the lab)", "reference: single-stain controls",
           "share of T cells in contact (%); permutation reference", "independent unit: pati
           "6 vs 6 patients"),
  type = c("measurement", "measurement", "measurement", "model", "measurement", "reference",
           "measurement", "decision")
)
draw_chain(chain_b12) +
```

```
  labs(caption = "Synthetic case · independent unit: patient (n = 12)")
```

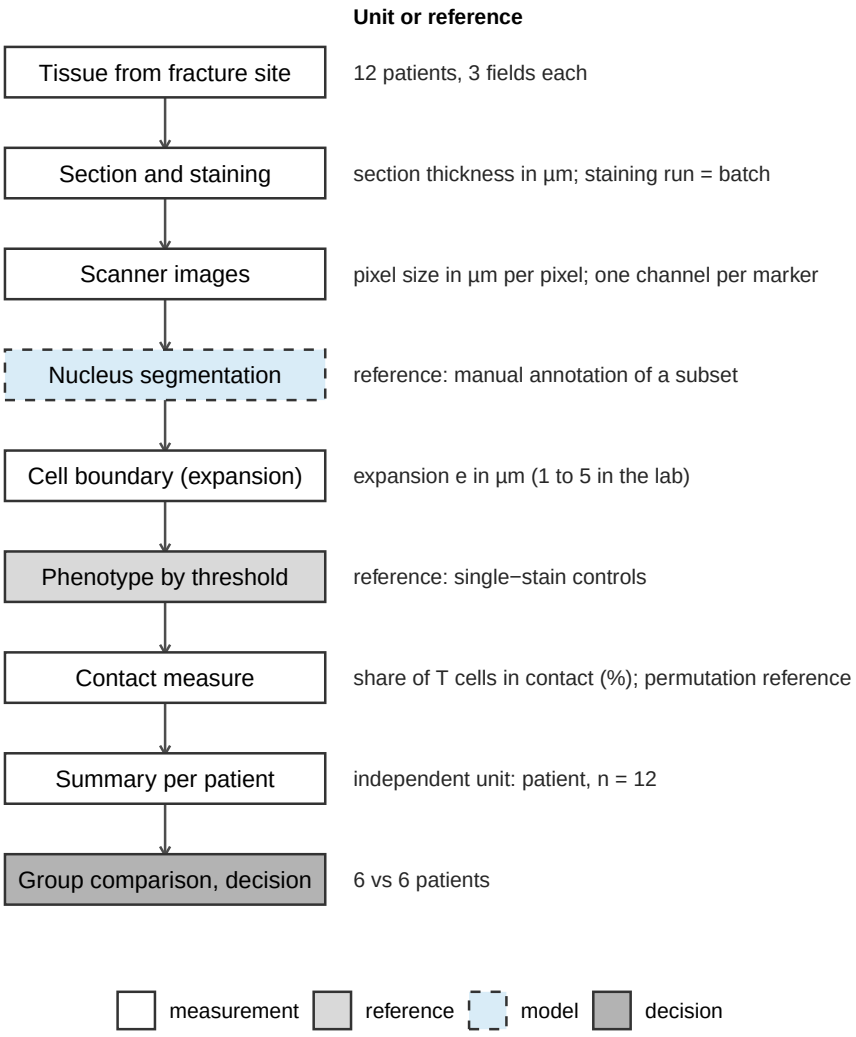


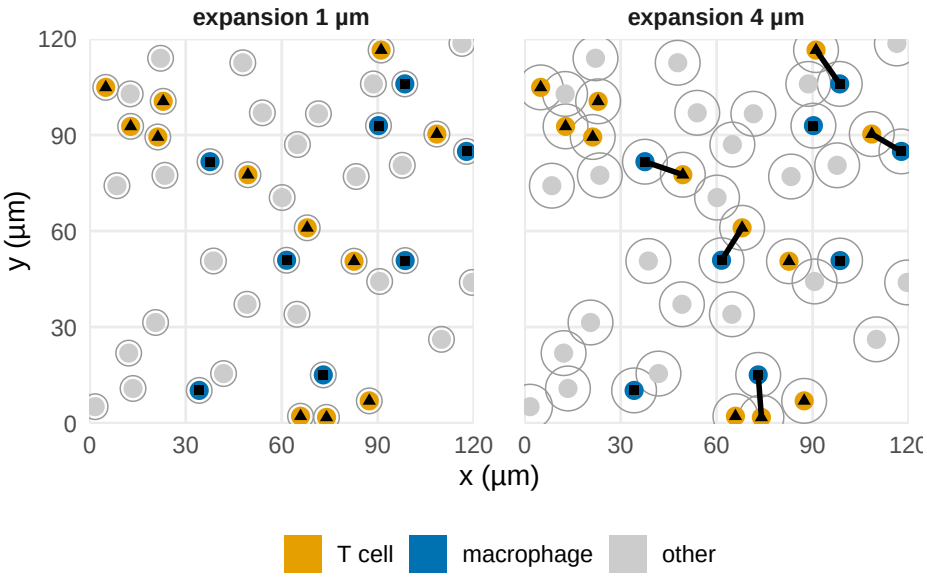
Figure 12.1.: Measurement chain for a cell-contact analysis in multiplex tissue images. The dashed box marks the learning model (nucleus segmentation); grey boxes mark the reference and the decision. Units change from micrometres to cell types to shares. Independent unit for the group question: patient (n = 12, six per group). Synthetic case.

Why the boundary matters. Two cells count as being in contact when their segmented boundaries touch. With nuclei of radius 3 µm and an expansion of e µm on each side, that happens when the centres are at most $2(3+e)$ µm apart: 8 µm for $e = 1$, 14 µm for $e = 4$. The area within reach of a cell grows with the square of that distance, so the expansion has a large effect on the number of contacts (Figure 12.2). It has a larger effect where cells are packed more densely. In the synthetic data, tissue from delayed healing contains more cells per square millimetre. A wider boundary therefore produces more chance contacts there, even though the T cells and macrophages are placed without any attraction to each other. **The key assumption is that a contact between segmented boundaries reflects a contact between real cells, to the same degree in both groups.** Denser tissue, thicker sections or a different staining run can each break this assumption for one group only.

```
# Choose the 120 µm x 120 µm detail of patient T07, field 1, that contains th
field_t07 <- cells[cells$patient == "T07" & cells$field == 1, ]
count_links <- function(w, e) {
  t_idx <- which(w$type == "T cell"); m_idx <- which(w$type == "macrophage")
  if (length(t_idx) == 0 || length(m_idx) == 0) return(0)
  g <- expand.grid(t = t_idx, m = m_idx)
  sum(sqrt((w$x[g$t] - w$x[g$m])^2 + (w$y[g$t] - w$y[g$m])^2) <= 2 * (nucleus
}
origins <- expand.grid(x0 = seq(0, field_size - 120, by = 20), y0 = seq(0, fi
origins$links <- apply(origins, 1, function(o) count_links(
  field_t07[field_t07$x >= o[1] & field_t07$x < o[1] + 120 & field_t07$y >= o
best <- origins[which.max(origins$links), ]
window <- field_t07[field_t07$x >= best$x0 & field_t07$x < best$x0 + 120 &
                    field_t07$y >= best$y0 & field_t07$y < best$y0 + 120, ]
window$x <- window$x - best$x0   # local coordinates of the detail
window$y <- window$y - best$y0
circle <- function(x, y, r, id) {
  a <- seq(0, 2 * pi, length.out = 41)
  data.frame(id = id, x = x + r * cos(a), y = y + r * sin(a))
```

```
}
panels <- lapply(c(1, 4), function(e) {
  label <- paste0("expansion ", e, " µm")
  outlines <- do.call(rbind, lapply(seq_len(nrow(window)), function(k)
    circle(window$x[k], window$y[k], nucleus_radius + e, k)))
  nuclei <- do.call(rbind, lapply(seq_len(nrow(window)), function(k)
    cbind(circle(window$x[k], window$y[k], nucleus_radius, k), type = window$type[k])))
  t_idx <- which(window$type == "T cell"); m_idx <- which(window$type == "macrophage")
  pairs <- expand.grid(t = t_idx, m = m_idx)
  pairs <- pairs[sqrt((window$x[pairs$t] - window$x[pairs$m])^2 + (window$y[pairs$t] - windo
                  2 * (nucleus_radius + e), ]
  links <- data.frame(x = window$x[pairs$t], y = window$y[pairs$t], xend = window$x[pairs$m]
                      yend = window$y[pairs$m])
  list(outlines = cbind(outlines, panel = label), nuclei = cbind(nuclei, panel = label),
       links = if (nrow(links) > 0) cbind(links, panel = label) else NULL,
       centres = cbind(window, panel = label))
})
pick <- function(part) do.call(rbind, lapply(panels, `[[`, part))
type_levels <- c("T cell", "macrophage", "other")
centres <- pick("centres"); centres$type <- factor(centres$type, levels = type_levels)
nuclei <- pick("nuclei"); nuclei$type <- factor(nuclei$type, levels = type_levels)
ggplot() +
  geom_polygon(data = pick("outlines"), aes(x, y, group = id), fill = NA, colour = "grey60",
  geom_polygon(data = nuclei, aes(x, y, group = id, fill = type), colour = NA) +
  geom_segment(data = pick("links"), aes(x, y, xend = xend, yend = yend), linewidth = 1.1) +
  geom_point(data = centres[centres$type != "other", ], aes(x, y, shape = type), size = 1.6,
  facet_wrap(~ panel) +
  scale_fill_manual(values = c("T cell" = unname(cdsr_colours["orange"]),
                               macrophage = unname(cdsr_colours["blue"]), other = "grey80"),
  scale_shape_manual(values = c("T cell" = 17, macrophage = 15), name = NULL) +
  guides(shape = "none") +                        # triangles = T cells, squares = macrophage
  coord_equal(xlim = c(0, 120), ylim = c(0, 120), expand = FALSE) +
  labs(x = "x (µm)", y = "y (µm)",
```

```
      caption = "Synthetic tissue · patient T07, field 1 · triangles: T cell
  theme_cdsr() +
  theme(panel.spacing = grid::unit(1.5, "lines"))
```



Synthetic tissue · patient T07, field 1 · triangles: T cells, squares: macrophages

Figure 12.2.: A 120 µm × 120 µm detail of one synthetic field of view from a patient with delayed healing, chosen as the detail with the most contacts at 4 µm. Filled discs: nuclei (radius 3 µm; triangles mark T cells, squares macrophages); grey outlines: cell boundaries after expansion by 1 µm (left) or 4 µm (right); thick lines: T cell–macrophage pairs counted as in contact. Unit: µm. Synthetic data; the independent unit for the study is the patient (n = 12).

Units and the learning step. The units change several times along the chain. Section thickness and distances are micrometres; positions in

the image are pixels and become micrometres only through the scanner's pixel size; cell density is given in cells per square millimetre; the contact measure is a share of T cells. Confusing pixels with micrometres is a silent error, because every number still looks plausible. The learning model sits at the segmentation step. It was trained on other images, often of other tissues and stains. In dense infiltrates, touching nuclei can be merged into one cell or split into two, and that changes cell counts and contacts precisely in the group with the higher density. The segmentation model therefore needs the same scrutiny as the expansion parameter that follows it: its name, version and weights belong in the methods, and its errors should be checked in both groups separately.

What counts as a reference. Each step needs its own check. Segmentation is checked against a pathologist's manual annotation of a subset of fields, made without knowing the patient's group. Phenotype thresholds are checked with single-stain controls. For the contact measure itself, the natural reference is what density alone would produce: if the cell types are shuffled at random among the same positions, how often would a T cell touch a macrophage? Comparing the observed share with this permutation expectation separates a specific neighbourhood preference from simple crowding (Schapiro et al. 2017).

Two pipelines, two conclusions. The colleague's 1 µm and the lab's 4 µm are both defensible defaults. Neither is right in itself; which one approximates real cell boundaries is an empirical question for the manual reference above. What is not defensible is choosing between them after seeing which gives the larger group difference. The expansion belongs in the analysis plan, fixed after a blinded check against manual boundaries, and the results for a range of expansions belong in the supplement as a sensitivity analysis.

Count patients, not cells. The study images tens of thousands of cells, but the independent unit is the patient. Cells from the same patient share genetics, treatment, surgery and staining; they are pseudoreplicates (Zimmerman et al. 2021). Tests that treat every cell as independent

produce false discoveries even in the absence of a biological difference (Squair et al. 2021). Two further pitfalls sit at the edges of the chain. Cells near the border of a field of view have neighbours outside the image that are not counted, which lowers their apparent number of contacts; spatial statistics handle this with edge corrections (Baddeley and Turner 2005). And if the two groups were stained in different runs, a batch effect in staining intensity can shift phenotypes in one group only (Leek et al. 2010), the same structure as the lamp in Chapter 10.

Question for the workshop: Which number would you report as the sample size for the contact comparison, and which single additional analysis would show whether more contacts in delayed healing mean more than more cells?

12.3. Lab

The lab generates twelve fictitious patients with three fields of view each, varies the boundary expansion, recomputes the contact share, compares it with a permutation reference and condenses everything per patient. In the synthetic truth there is no attraction between T cells and macrophages; only the cell density differs between the groups.

Step 1 – the synthetic tissue. The cells were already placed in a hidden chunk of the core reading, because the figure there needs them. Its code is repeated here without being run again: nucleus centres at least 7 µm apart (a simple hard-core process), 15% T cells, 15% macrophages and 70% other cells, assigned at random, and a higher cell density in delayed healing.

```
# Synthetic tissue: positions and cell types (the lab shows and checks this code step by ste
field_size <- 500        # field of view, µm per side
nucleus_radius <- 3      # µm
min_gap <- 7             # minimum distance between nucleus centres, µm
```

```
place_cells <- function(n, size = field_size, min_dist = min_gap) {
  x <- numeric(0); y <- numeric(0); tries <- 0
  while (length(x) < n && tries < 50 * n) {
    tries <- tries + 1
    cx <- runif(1, 0, size); cy <- runif(1, 0, size)
    if (length(x) == 0 || min((x - cx)^2 + (y - cy)^2) >= min_dist^2) { x <-
  }
  data.frame(x = x, y = y)
}
patients <- data.frame(patient = sprintf("T%02d", 1:12), group = rep(c("uneve
patients$density <- ifelse(patients$group == "delayed", rnorm(12, 3000, 300)
cell_list <- list()
for (i in 1:12) for (f in 1:3) {
  n <- rpois(1, patients$density[i] * (field_size / 1000)^2)
  p <- place_cells(n)
  p$type <- sample(c("T cell", "macrophage", "other"), nrow(p), replace = TRU
  p$patient <- patients$patient[i]; p$group <- patients$group[i]; p$field <-
  cell_list[[length(cell_list) + 1]] <- p
}
cells <- do.call(rbind, cell_list)
```

Then we check the design.

```
cells_per_field <- aggregate(x ~ patient + group + field, data = cells, FUN =
names(cells_per_field)[4] <- "cells"
density_per_patient <- aggregate(cells ~ patient + group, data = cells_per_fi
density_per_patient$cells_per_mm2 <- density_per_patient$cells / (field_size
aggregate(cells_per_mm2 ~ group, data = density_per_patient, FUN = mean)
```

```
#>        group cells_per_mm2
#> 1    delayed      2915.333
#> 2 uneventful      2227.111
```

```
min_distance <- tapply(seq_len(nrow(cells)), paste(cells$patient, cells$field), function(k)
  min(dist(cells[k, c("x", "y")])))
c(patients = length(unique(cells$patient)), fields = nrow(cells_per_field), cells = nrow(cel
  t_cells = sum(cells$type == "T cell"))
```

```
#> patients   fields    cells  t_cells
#>       12       36    23141     3478
```

```
stopifnot(length(unique(cells$patient)) == 12, nrow(cells_per_field) == 36, min(min_distance
```

Step 2 – contacts for a given boundary. For each field we list all pairs of cells whose centres are close enough to touch at the largest expansion. The contact share for any expansion is then a filter on this list.

```
expansions <- 1:5                                          # boundary expansion in µm
max_reach <- 2 * (nucleus_radius + max(expansions))        # centre distance for contact at 5
close_pairs <- function(d) {
  D <- as.matrix(dist(d[, c("x", "y")]))
  idx <- which(D <= max_reach & upper.tri(D), arr.ind = TRUE)
  data.frame(i = idx[, 1], j = idx[, 2], distance = D[idx])
}
# Share of T cells with at least one macrophage contact; 'edge' drops T cells whose reach cr
contact_count <- function(d, pairs, e, types = d$type, edge = FALSE) {
  reach <- 2 * (nucleus_radius + e)
  p <- pairs[pairs$distance <= reach, ]
  t_cells <- which(types == "T cell")
  if (edge) t_cells <- t_cells[d$x[t_cells] > reach & d$x[t_cells] < field_size - reach &
                               d$y[t_cells] > reach & d$y[t_cells] < field_size - reach]
  touching <- c(p$i[types[p$j] == "macrophage"], p$j[types[p$i] == "macrophage"])
  c(t_cells = length(t_cells), in_contact = sum(t_cells %in% touching))
}
```

Step 3 – shift the boundary and build the permutation reference. For every field and expansion we record the observed contacts, the contacts after excluding T cells near the border, and the share expected if the cell types were shuffled at random among the same positions (100 permutations per field).

```
set.seed(2212)
field_keys <- unique(cells[, c("patient", "group", "field")])
per_field <- do.call(rbind, lapply(seq_len(nrow(field_keys)), function(k) {
  d <- cells[cells$patient == field_keys$patient[k] & cells$field == field_ke
  pairs <- close_pairs(d)
  do.call(rbind, lapply(expansions, function(e) {
    observed <- contact_count(d, pairs, e)
    inner <- contact_count(d, pairs, e, edge = TRUE)
    shuffled <- replicate(100, { s <- contact_count(d, pairs, e, types = samp
    data.frame(field_keys[k, ], expansion = e, t_cells = observed[1], in_cont
               t_cells_inner = inner[1], in_contact_inner = inner[2], expecte
  }))
}))
# Condense per patient: pooled counts over the three fields
per_patient <- aggregate(cbind(t_cells, in_contact, t_cells_inner, in_contact
                         data = per_field, FUN = sum)
per_patient <- merge(per_patient, aggregate(expected_share ~ patient + expans
per_patient$share <- per_patient$in_contact / per_patient$t_cells
per_patient$share_inner <- per_patient$in_contact_inner / per_patient$t_cells
per_patient$ratio <- per_patient$share / per_patient$expected_share
# Contacts can only increase with a wider boundary, for every patient
monotone <- tapply(seq_len(nrow(per_patient)), per_patient$patient, function
  all(diff(per_patient$share[k][order(per_patient$expansion[k])]) >= 0))
stopifnot(all(monotone))
round(aggregate(cbind(share, expected_share, ratio) ~ group + expansion, data
```

```
#>    expansion share expected_share ratio
```

```
#> 1           1 0.022          0.022 1.023
#> 2           1 0.017          0.018 0.988
#> 3           2 0.066          0.072 0.918
#> 4           2 0.056          0.057 1.001
#> 5           3 0.117          0.127 0.923
#> 6           3 0.097          0.099 0.979
#> 7           4 0.173          0.186 0.927
#> 8           4 0.139          0.150 0.928
#> 9           5 0.234          0.249 0.933
#> 10          5 0.188          0.203 0.924
```

Step 4 – the same question at two levels. For the colleague's 1 µm and the default 4 µm we test the group difference once over all T cells, as if they were independent, and once with one value per patient.

```
compare_levels <- function(e) {
  d <- per_patient[per_patient$expansion == e, ]
  pooled <- aggregate(cbind(in_contact, t_cells) ~ group, data = d, FUN = sum)
  data.frame(expansion_um = e,
             share_delayed = mean(d$share[d$group == "delayed"]),
             share_uneventful = mean(d$share[d$group == "uneventful"]),
             n_t_cells = sum(d$t_cells),
             p_cells = prop.test(pooled$in_contact, pooled$t_cells)$p.value,   # every T cel
             n_patients = nrow(d),
             p_patients = t.test(share ~ group, data = d)$p.value,            # one value pe
             p_ratio = t.test(ratio ~ group, data = d)$p.value)                # density-adj
}
levels_test <- rbind(compare_levels(1), compare_levels(4))
levels_test
```

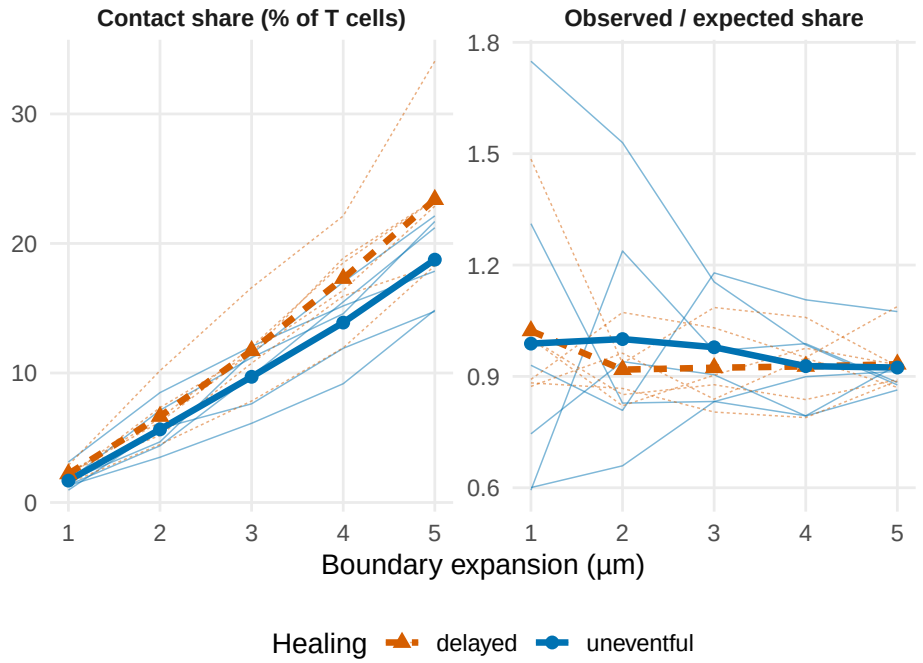
```
#>   expansion_um share_delayed share_uneventful n_t_cells     p_cells n_patients
#> 1            1    0.02210588       0.01688545      3478 0.273426883         12
#> 2            4    0.17294060       0.13889900      3478 0.003194273         12
```

```
#>   p_patients   p_ratio
#> 1 0.18559854 0.8741119
#> 2 0.09185391 0.9906893
```

```
gap <- levels_test$share_delayed - levels_test$share_uneventful
stopifnot(gap[2] > gap[1], levels_test$n_t_cells[2] > 3000,
          levels_test$p_cells[2] < 0.01, levels_test$p_patients[2] > 0.05, le
```

At 1 µm the groups differ by 0.5 percentage points; at 4 µm by 3.4 points. Only the boundary changed. The test over 3,478 T cells reports p = 0.0032 at 4 µm; the test over 12 patients reports p = 0.092. The density-adjusted ratio, observed divided by expected share, is close to 1 in both groups (p = 0.99): the extra contacts in delayed healing are what crowding alone produces. Figure 12.3 shows the whole range of expansions.

```
long <- rbind(
  data.frame(per_patient[, c("patient", "group", "expansion")], value = 100 *
             panel = "Contact share (% of T cells)"),
  data.frame(per_patient[, c("patient", "group", "expansion")], value = per_p
             panel = "Observed / expected share")
)
means <- aggregate(value ~ group + expansion + panel, data = long, FUN = mean
ggplot(long, aes(expansion, value, colour = group, linetype = group)) +
  geom_line(aes(group = patient), linewidth = 0.3, alpha = 0.5) +
  geom_line(data = means, linewidth = 1.3) +
  geom_point(data = means, aes(shape = group), size = 2.4) +
  facet_wrap(~ panel, scales = "free_y") +
  scale_colour_manual(values = unname(cdsr_colours[c("red", "blue")])) +
  scale_linetype_manual(values = cdsr_linetypes[c(2, 1)]) +
  scale_shape_manual(values = cdsr_shapes[c(2, 1)]) +
  labs(x = "Boundary expansion (µm)", y = NULL, colour = "Healing", linetype
       caption = "Synthetic data · n = 12 patients (6 per group) · 3 fields o
  theme_cdsr()
```



Synthetic data · n = 12 patients (6 per group) · 3 fields of 0.25 mm² each

Figure 12.3.: Left: share of T cells in contact with at least one macrophage, per patient (thin lines) and group mean (thick lines), against the boundary expansion in µm. Right: the same share divided by the share expected from random placement of cell types (permutation reference); 1 means no preference beyond density. Synthetic data; independent unit: patient (n = 12, six per group); 3 fields of 500 µm × 500 µm per patient.

> **Common mistake**
>
> Reporting "n = 3,478 T cells" as the sample size. Cells, fields and pixels are measurements within patients. The sample size for a comparison between groups of patients is the number of patients; the number of cells per patient describes the precision of each patient's value and belongs next to it.

Step 5 – what condensing per patient loses. Condensing to one share per patient is what makes the test honest, but it discards information that belongs in the report.

```
at4 <- per_field[per_field$expansion == 4, ]
at4$share <- at4$in_contact / at4$t_cells
loss <- aggregate(cbind(t_cells, share) ~ patient + group, data = at4,
                  FUN = function(v) c(min = min(v), max = max(v)))
loss <- data.frame(patient = loss$patient, group = loss$group,
                   t_cells_total = per_patient$t_cells[per_patient$expansion
                                     per_patient$patient[per_patient$expansi
                   field_share_min = round(loss$share[, "min"], 3), field_sha
head(loss[order(-loss$t_cells_total), ], 4)
```

```
#>   patient   group t_cells_total field_share_min field_share_max
#> 4     T10 delayed           452           0.178           0.255
#> 5     T11 delayed           355           0.164           0.204
#> 2     T08 delayed           354           0.137           0.200
#> 6     T12 delayed           308           0.168           0.196
```

```
range(loss$t_cells_total)
```

```
#> [1] 210 452
```

Three things disappear in the patient summary. First, precision: the shares rest on between 210 and 452 T cells per patient, so they are not equally reliable. Second, heterogeneity: the three fields of one patient can differ by several percentage points, and a single field with a dense aggregate can dominate. Third, the spatial arrangement itself: where in the tissue the contacts occur, for example near vessels or at the fracture gap, is no longer visible. Reporting the number of T cells and fields per patient, and the range across fields, keeps the loss transparent. Models that keep the cells but treat the patient as the unit, such as generalised linear mixed models, are the next step (Methods in R: glmer: Generalised Mixed Models).

i Other implementation: Python with scikit-image and Cellpose

In practice the segmentation runs in Python. scikit-image (0.26.0 in November 2026) provides `skimage.segmentation.expand_labels(label_image, distance)`, which "expands labels in label image by distance pixels without overlapping", and `skimage.measure.regionprops_table` for centroids and mean marker intensities (Walt et al. 2014). The distance is given in pixels: at 0.5 µm per pixel, the lab's 4 µm become 8 pixels, and a pipeline that forgets the conversion silently doubles or halves the boundary. Nuclei are often found with Cellpose (4.2.1 in November 2026), a PyTorch-based generalist model (Stringer et al. 2021; Paszke et al. 2019); its current model uses the first three image channels and was trained on objects of about 7.5 to 120 pixels in diameter, so channel order and pixel size must be documented. The concrete advantage is a tested segmentation of real images, which this lab replaces by synthetic points. For a fair comparison the following must be identical: pixel size, the label image, the expansion in µm, the contact definition (labels that touch after expansion), the field borders and edge rule, and the per-patient aggregation. Because `expand_labels` stops where two growing cells

meet, a third cell between two others can block a contact that the distance rule in this lab would count; the shares should therefore agree closely but not exactly, and the tolerance (for example one percentage point per patient) has to be set beforehand. Status: conceptual comparison, not run.

12.4. Counter-check

Predict first: if T cells whose reach crosses the border of the field are excluded, does the contact share rise or fall, and does the conclusion of the group comparison change?

```
edge_check <- aggregate(cbind(share, share_inner) ~ group, data = per_patient
                        FUN = mean)
edge_check$t_cells_kept <- aggregate(t_cells_inner ~ group, data = per_patien
                                     FUN = sum)$t_cells_inner
edge_check
```

```
#>        group     share share_inner t_cells_kept
#> 1    delayed 0.1729406   0.1802516         1767
#> 2 uneventful 0.1388990   0.1395331         1336
```

```
d4 <- per_patient[per_patient$expansion == 4, ]
p_inner <- t.test(share_inner ~ group, data = d4)$p.value
stopifnot(all(edge_check$share_inner > edge_check$share), p_inner > 0.05)
```

The border correction raises the share slightly in both groups, because cells at the border lose neighbours that lie outside the image. With 500 µm fields and a reach of 14 µm the effect is small and does not change the patient-level result ($p = 0.069$). It grows with smaller fields and larger

reach. A second counter-check is a thought experiment that the lab has already answered: if both groups had the same cell density, the gap in contact share would shrink towards zero at every expansion, which is what the observed/expected ratio in Figure 12.3 shows.

The checkable measurement assumption of this case is: *the boundary expansion, converted to pixels with the scanner's pixel size, reproduces a pathologist's manual cell boundaries in a blinded subset of fields of both groups equally well.* The plausible failure mode is: *a default expansion chosen for sparse tissue creates chance contacts in dense tissue and turns a difference in cellularity into an apparent difference in cell interaction.*

12.5. Consequences for patients and operations

A misleading sample size. The most direct harm in this case is a number. "$p < 0.01$ in more than 3,000 T cells" reads like strong evidence and can secure funding, steer a follow-up study and recruit further patients for tissue donation, all for a difference that crowding explains. The honest statement is "12 patients, six per group, no difference in density-adjusted contacts". Patients who donate tissue do so on the assumption that their samples are analysed in a way that can answer the question.

Wrong tissue interpretation. If a contact score were later used to classify patients, for example as at risk of delayed healing, the segmentation parameter would decide which patients receive closer follow-up or additional procedures. Software that analyses images of tissue specimens to provide information on a pathological state may fall under the In Vitro Diagnostic Medical Devices Regulation (EU) 2017/746, whose definition covers software intended to be used in vitro for the examination of specimens derived from the human body [European Parliament and Council of the European Union (2017b); binding text]. Whether a given tool is such a device depends on its intended purpose; the guidance MDCG 2019-

11 Rev. 1 helps with the qualification but is not binding (Medical Device Coordination Group 2025). Chapter 22 works through an IVD vignette.

Data protection. Whole-slide images linked with clinical data are health data. Rare combinations of diagnosis, age and date of surgery can identify a patient even without a name; Chapter 23 and Chapter 26 discuss minimisation and controlled access.

Operations and integrity. The expansion distance, the segmentation model and its version, the pixel size and the phenotype thresholds are part of the result. They belong in the model passport (Chapter 27) and should be stored with a checksum, so that a later re-analysis can show that it used the same parameters (Chapter 24).

12.6. Look up and transfer

- **Reviewer question:** The manuscript reports a significant increase in T cell–macrophage contacts “in 4,812 cells”. How many patients contributed, how was the cell boundary defined and in which unit, and was the contact frequency compared with what cell density alone would produce?
- **Self-test:** A second study uses the same pipeline but images only one field of 250 µm × 250 µm per patient, and the boundary expansion is set to 5 µm. Without computing: will the border effect be larger or smaller than in the lab, will the gap between dense and sparse tissue be larger or smaller, and what is the independent sample size if 20 patients contribute 9,000 T cells?

Answer

The border effect will be larger. A smaller field has more cells near its border relative to its area, and the reach grows to $2(3 + 5) = 16$ µm, so a larger share of T cells has neighbours outside the image.

The gap between dense and sparse tissue will also be larger, because a wider boundary produces more chance contacts where cells are packed more tightly. The independent sample size is 20 patients; the 9,000 T cells are repeated measurements within them. A density-adjusted measure, such as the observed share divided by a permutation expectation, and an edge correction belong in the analysis plan before the data are seen.

- **Annotated sources:**
  - Schapiro et al. (2017) introduced neighbourhood analysis of multiplexed tissue images against a permutation reference, the idea behind the observed/expected ratio.
  - Stringer et al. (2021) and Greenwald et al. (2022) describe pretrained segmentation models, the learning step of the chain.
  - Zimmerman et al. (2021) shows that cells from one individual are pseudoreplicates and that ignoring this inflates type I error.
  - Squair et al. (2021) demonstrates false discoveries when variation between biological replicates is ignored.
  - Baddeley and Turner (2005) documents the R tools for spatial point patterns, including edge corrections.
  - Leek et al. (2010) explains how batch effects can masquerade as biology.
  - European Parliament and Council of the European Union (2017b) and Medical Device Coordination Group (2025) are the binding text and the non-binding guidance for the regulatory paragraph.

12.7. Further reading

- Chapter 14 – when a pattern in high-dimensional data is a cluster, a gradient or an artefact.

- Chapter 26 – controlled access and minimisation for image data linked to patients.
- Chapter 22 – the IVD vignette and how an intended purpose changes the regulatory question.
- Methods in R: Spatial Transcriptomics – spatially resolved data beyond imaging.
- Methods in R: Batch Correction with ComBat – what batch correction can and cannot repair.
- CTTIR Courses: GLMMs and GEE – a lab on models that keep cells but use the patient as the unit.
- CTTIR Tutorials: Spatial Transcriptomics – a compact tutorial for revision.

Part IV.

Part IV · Complex models for trauma research

13. Longitudinal biomarkers

R. Heller · Date: 2025-03-30 · Last changed: 2026-11-09

At a glance

Guiding question: Whose course does the estimated slope describe?
For: clinical researchers, doctoral students and reviewers who work with repeated laboratory values · **Prerequisites:** estimands (Chapter 2), missing values and denominators (Chapter 3), simple linear regression
In the course: W5 · Causality, trajectories and phenotypes · **Data:** synthetic (generated in the chapter; download as CSV) · **Core reading time:** about 10 minutes

13.1. Case and decision

A doctoral researcher in a fictitious spinal cord injury centre has serum selenium values from 36 people after traumatic spinal cord injury: one sample at admission and up to five more over the following eight weeks. Before the next research meeting, the supervising consultant asks a single question. Does selenium recover more slowly after a complete injury than after an incomplete one? If it does, the team will draft a protocol for a supplementation trial restricted to people with complete injuries. Published observational work on trace elements after spinal cord injury motivates the question (Seelig et al. 2020); every value in this chapter is *synthetic*.

Not everybody stayed for all visits. Some people were transferred, others missed late appointments. Whether the protocol is justified therefore depends on a question that sounds technical but is clinical: whose course does the estimated slope describe?

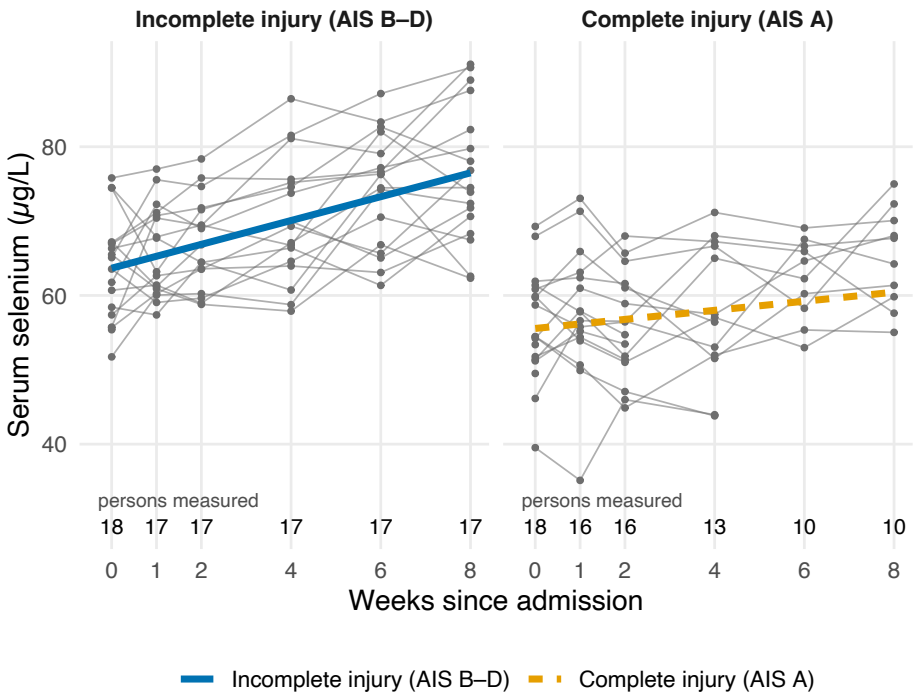
13.2. Core reading

Core reading · required for W5 · about 10 minutes · ends before ‘Lab’

One table, three levels. A longitudinal data set has three levels that are easily confused: the person, the visit and the measurement. Here, 36 people (18 per injury group) were due to be sampled six times, at admission and at weeks 1, 2, 4, 6 and 8. That gives 216 planned values; 186 were measured. The independent unit is the person. Any statement about “the course after a complete injury” is a statement about people, so its uncertainty must be driven by the number of people, not by the number of tubes in the freezer.

The quantity of interest. The consultant’s question concerns a time–group interaction: the difference between the mean change per week after a complete injury (ASIA Impairment Scale grade A) and after an incomplete injury (grades B–D). Its unit is µg/L per week. A value of -1 means that selenium rises by 1 µg/L per week less after a complete injury. The difference at admission is a separate quantity, the group main effect, and answers a different question.

What the figure shows. Figure 13.1 draws one thin line per person (a spaghetti plot) and the fitted mean course of each group. In the complete-injury group, the denominator falls from 18 at admission to 10 at week 8; in the incomplete group it falls to 17. People with low values left more often. The people still visible at week 8 are therefore not a random subset of those admitted.



Synthetic data; group means from a linear mixed model with random intercept and slope.

Figure 13.1.: Serum selenium (µg/L, *synthetic*) over eight weeks after traumatic spinal cord injury. Thin grey lines: one person each; thick line: fitted mean course of the group from the mixed model. The numbers along the bottom are the persons measured at each visit. n = 36 persons (18 per group), 186 of 216 planned measurements.

Three answers to one question. Table 13.1 compares three ways of estimating the interaction.

- *Pooled regression.* One straight line per group through all measured values, as if every measurement came from a different person. It ignores that values from the same person are correlated, so its interval is too narrow for a statement about people, and it lets persons with many visits dominate. Here, the late visits come mostly from people with a better course, and the estimated difference shrinks towards zero.
- *Slope per person (baseline).* For each person with at least two values, fit a straight line and keep its slope; then compare the mean slope between groups. This summary-measure approach (Matthews et al. 1990) is the biometrically sensible baseline: transparent, person-based and easy to check by hand. It weights a person with two values like a person with six, and it silently drops people with a single value.
- *Linear mixed model.* A model with a random intercept and a random slope per person (Laird and Ware 1982), fitted with **lme4** (Bates et al. 2015). Each person gets their own line, but lines from people with few values are pulled towards the group mean. All 36 persons and all measured values enter.

Table 13.1.: Difference in the weekly change of serum selenium, complete minus incomplete injury (µg/L per week, *synthetic*), with 95% confidence intervals. The simulation truth is known only because the data were generated in the chapter.

Approach	Units used	Difference	Lower 95%	Upper 95%
Pooled regression (ignores persons)	186 measurements	-0.54	-1.33	0.25
Slope per person (baseline)	33 persons	-1.13	-1.92	-0.34

Table 13.1.: Difference in the weekly change of serum selenium, complete minus incomplete injury (µg/L per week, *synthetic*), with 95% confidence intervals. The simulation truth is known only because the data were generated in the chapter.

Approach	Units used	Difference	Lower 95%	Upper 95%
Mixed model, random intercept and slope	36 persons, 186 measurements	-0.99	-1.67	-0.31
Simulation truth	36 persons, 216 planned values	-1.00		

Because the data are synthetic, the true difference (−1.0 µg/L per week) is known. In this cohort, the pooled regression underestimates the difference, the per-person slopes overestimate it and the mixed model comes closest. One cohort proves nothing about a method; the counter-check repeats the experiment a hundred times (Section 13.4).

Whose course? The three numbers describe different people. The pooled slope is the course of the *measurements*, and late measurements come from those who stayed. The mean of per-person slopes describes everybody with at least two values, each counted once, however noisy their line. The fixed slope of the mixed model describes a typical admitted person – under one condition.

The most important assumption. The mixed model's slope describes everyone who was admitted only if the chance of leaving the study depends on what had already been measured, not on the unmeasured current value or on the person's own unobserved rate of recovery. This condition is called missing at random (MAR) (Rubin 1976). In the simulation, transfers depend on the last measured selenium value, a stand-in for the observable clinical condition, so MAR holds by construction. In a real cohort, it is an assumption that no analysis of the observed data can prove. When leaving depends on what was not

observed (missing not at random, MNAR), every estimate in Table 13.1 can be biased, and the honest response is a sensitivity analysis (Diggle and Kenward 1994; Little et al. 2012). The mechanisms are explained in more depth in the Methods in R lab on MCAR, MAR and MNAR.

What to say in the meeting. A defensible sentence would be: "In 36 people, selenium rose by about 1.0 µg/L per week less after a complete injury (mixed model, 95% CI 0.3 to 1.7). This assumes that leaving the study depended on observed values. If the people who left had values about 18 µg/L higher than their own earlier course predicts, the difference would disappear." The last sentence comes from the sensitivity analysis in the lab. It turns a hidden assumption into a number that the clinical team can judge.

Question for the workshop. Which observable reason for leaving the study – transfer to intensive care, death, discharge home, a missed outpatient appointment – would make the MAR assumption more plausible, and which less? Sort the four reasons before you open the lab.

13.3. Lab

The lab runs in a fresh R session in about ten seconds. It needs **lme4** and **ggplot2**.

13.3.1. Step 1: write down the simulation assumptions

Every number in the simulation is an *assumption* that is set here and nowhere else. The dropout rule makes the chance of leaving before the next visit depend on the last measured value (MAR). A second rule, used only in the counter-check, makes it depend on the unmeasured current level and on the person's own recovery rate (MNAR).

```r
set.seed(1313)                                          # set again so this chunk always gives th
n_per_group <- 18                                       # persons per injury group
weeks <- c(0, 1, 2, 4, 6, 8)                            # planned visits, weeks since admission
true_start <- c(incomplete = 62, complete = 58)         # mean selenium at admission (µg/L)
true_slope <- c(incomplete = 1.6, complete = 0.6)       # mean change (µg/L per week)
sd_start <- 7                                           # between-person SD at admission (µg/L)
sd_slope <- 0.6                                         # between-person SD of the weekly change
sd_noise <- 4                                           # measurement and day-to-day noise (µg/L)

simulate_cohort <- function(mechanism = c("mar", "mnar", "none")) {
  mechanism <- match.arg(mechanism)
  ids <- sprintf("P%02d", seq_len(2 * n_per_group))
  group <- rep(names(true_start), each = n_per_group)
  dev_start <- rnorm(length(ids), 0, sd_start)          # each person's deviation at admission
  dev_slope <- rnorm(length(ids), 0, sd_slope)          # each person's deviation of the weekly c
  d <- data.frame(id = rep(ids, each = length(weeks)),
                  group = rep(group, each = length(weeks)),
                  week = rep(weeks, times = length(ids)))
  p <- match(d$id, ids)
  d$true_level <- true_start[d$group] + dev_start[p] + (true_slope[d$group] + dev_slope[p])
  d$selenium <- d$true_level + rnorm(nrow(d), 0, sd_noise)
  d$observed <- TRUE
  for (k in seq_along(ids)) {                           # monotone dropout: once gone, always gon
    r <- which(d$id == ids[k])
    for (j in 2:length(r)) {
      linear_predictor <- switch(mechanism,
        none = -Inf,
        mar  = -2.8 - 0.2 * (d$selenium[r[j - 1]] - 60),                # last measured value
        mnar = -2.8 - 0.2 * (d$true_level[r[j]] - 60) - 1.5 * dev_slope[k] / sd_slope)  # un
      if (runif(1) < plogis(linear_predictor)) {
        d$observed[r[j:length(r)]] <- FALSE
        break
      }
```

```
    }
  }
  d$group <- factor(d$group, levels = c("incomplete", "complete"))
  d
}

cohort_full <- simulate_cohort("mar")              # all planned visits, with
cohort <- cohort_full[cohort_full$observed, c("id", "group", "week", "seleniu
n_persons <- length(unique(cohort$id))
n_week8 <- table(cohort$group[cohort$week == 8])   # persons still measured a
stopifnot(nrow(cohort_full) == 216, n_persons == 36,
          all(table(cohort$id[cohort$week == 0]) == 1))  # everybody has an a
```

13.3.2. Step 2: document the denominators at person level

Before any model, count people, not rows. Table 13.2 gives the denominator for every visit and the number of people who contribute at least two values (needed for a per-person slope).

```
per_visit <- as.data.frame.matrix(table(cohort$group, cohort$week))
n_meas_person <- table(cohort$id)
group_of_id <- tapply(as.character(cohort$group), cohort$id, unique)
per_visit$`persons with  2 values` <- as.vector(tapply(n_meas_person >= 2, g
per_visit$`all measurements` <- as.vector(table(cohort$group)[rownames(per_vi
names(per_visit)[1:6] <- paste("week", weeks)
knitr::kable(per_visit)
```

Table 13.2.: Persons with a measurement at each visit, by injury group (*synthetic*). The last two columns count persons with at least two measurements and all measurements of the group.

	week 0	week 1	week 2	week 4	week 6	week 8	persons with 2 values	all measurements
incomplete	18	17	17	17	17	17	17	103
complete	18	16	16	13	10	10	16	83

```
stopifnot(per_visit$`week 0` == c(18, 18),                # nobody is missing at admission
          sum(per_visit$`all measurements`) == nrow(cohort))
```

The data set is saved for the workshop; it contains only the synthetic values that this chapter generates.

```
write.csv(cohort, "../data/b13-selenium-synthetic.csv", row.names = FALSE)
```

13.3.3. Step 3: three estimates of the interaction

```
library(lme4)
ctrl <- lmerControl(optimizer = "bobyqa")
# (1) Pooled regression: every measurement treated as if it were a separate person (a common
fit_pooled <- lm(selenium ~ week * group, data = cohort)

# (2) Baseline: one least-squares slope per person with at least two values
per_person <- do.call(rbind, lapply(split(cohort, cohort$id), function(d) {
  if (nrow(d) < 2) return(NULL)
  data.frame(id = d$id[1], group = d$group[1], n_values = nrow(d),
             slope = unname(coef(lm(selenium ~ week, data = d))["week"]))
```

```
}))
welch <- t.test(slope ~ group, data = per_person)   # Welch test; estimates a

# (3) Linear mixed model: random intercept and random slope for every person
fit_lmm <- lmer(selenium ~ week * group + (week | id), data = cohort, REML =
summary(fit_lmm)$coefficients
```

```
#>                      Estimate Std. Error   t value
#> (Intercept)        63.6672802  1.5549898 40.943856
#> week                1.6025084  0.2190990  7.314082
#> groupcomplete      -8.1093780  2.2107273 -3.668195
#> week:groupcomplete -0.9926302  0.3458073 -2.870472
```

```
VarCorr(fit_lmm)
```

```
#>  Groups   Name        Std.Dev. Corr
#>  id       (Intercept) 6.00271
#>           week        0.68483  -0.083
#>  Residual             4.07563
```

The coefficient `week` is the mean weekly change after an incomplete injury (µg/L per week); `week:groupcomplete` is the interaction, the amount by which the weekly change differs after a complete injury.

```
est_pooled <- unname(coef(fit_pooled)["week:groupcomplete"])
ci_pooled <- unname(confint(fit_pooled)["week:groupcomplete", ])
est_person <- unname(welch$estimate[2] - welch$estimate[1])      # complete r
ci_person <- -rev(unname(welch$conf.int))
est_lmm <- unname(fixef(fit_lmm)["week:groupcomplete"])
ci_lmm <- unname(confint(fit_lmm, parm = "week:groupcomplete", method = "Walc
fmt <- function(x) cdsr_number(x, 2)
```

```
estimates_table <- data.frame(
  Approach = c("Pooled regression (ignores persons)", "Slope per person (baseline)",
               "Mixed model, random intercept and slope", "Simulation truth"),
  `Units used` = c(paste(nrow(cohort), "measurements"), paste(nrow(per_person), "persons"),
                   paste(n_persons, "persons,", nrow(cohort), "measurements"), "36 persons,
  Difference = c(fmt(est_pooled), fmt(est_person), fmt(est_lmm), fmt(-1)),
  `Lower 95%` = c(fmt(ci_pooled[1]), fmt(ci_person[1]), fmt(ci_lmm[1]), ""),
  `Upper 95%` = c(fmt(ci_pooled[2]), fmt(ci_person[2]), fmt(ci_lmm[2]), ""),
  check.names = FALSE)
estimates_table
```

```
#>                                 Approach                   Units used
#> 1     Pooled regression (ignores persons)               186 measurements
#> 2             Slope per person (baseline)                     33 persons
#> 3 Mixed model, random intercept and slope   36 persons, 186 measurements
#> 4                         Simulation truth 36 persons, 216 planned values
#>   Difference Lower 95% Upper 95%
#> 1      -0.54     -1.33      0.25
#> 2      -1.13     -1.92     -0.34
#> 3      -0.99     -1.67     -0.31
#> 4      -1.00
```

The target values below are fixed for this seed. They guard against silent changes in the data or the packages.

```
stopifnot(
  nrow(cohort) == 186, nrow(per_person) == 33,           # measured values and persons with
  !isSingular(fit_lmm),                                  # both random-effect variances esti
  abs(est_lmm - (-1)) < 0.1,                             # mixed model close to the truth in
  est_pooled > est_lmm, est_person < est_lmm,            # pooled shrinks, per-person slopes
  ci_lmm[1] < -1, ci_lmm[2] > -1                         # interval covers the simulation tr
)
```

> **Common mistake**
>
> Fitting `lm(selenium ~ week * group)` to all rows treats 186 tubes as 186 people. The standard error then reflects the number of measurements, not the number of independent persons, and persons with many visits dominate the slope. The same pseudoreplication appears with spectra per animal (Chapter 10) and steps per walk (Chapter 11).

13.3.4. Step 4: sensitivity analysis for late and missing visits

Two analyses test how much the answer depends on the visits that were not made.

(a) Early slope only. Refit the mixed model with weeks 0–4, where most people are still present. If the interaction changes markedly, the conclusion rests on the late visits of a selected group.

(b) Delta adjustment and tipping point. Fill every missing planned visit with the mixed model's prediction for that person (the MAR value), shift the filled values by µg/L, and refit. A positive means that people who left were doing better than their own earlier course predicts (for example, discharged home after a good recovery); a negative means they were doing worse. The tipping point is the at which the group difference vanishes. Filling in a single value per visit understates the uncertainty, so only point estimates are reported here; for intervals, use multiple imputation (see the CTTIR lab on multiple imputation with mice).

```
fit_early <- lmer(selenium ~ week * group + (week | id), data = subset(cohort
est_early <- unname(fixef(fit_early)["week:groupcomplete"])

missing_visits <- cohort_full[!cohort_full$observed, c("id", "group", "week"
mar_value <- predict(fit_lmm, newdata = missing_visits)   # person-specific p
```

```
refit_with_delta <- function(delta) {
  filled <- rbind(cohort, data.frame(missing_visits, selenium = mar_value + delta))
  unname(fixef(lmer(selenium ~ week * group + (week | id), data = filled, control = ctrl))["
}
delta_grid <- seq(-8, 24, by = 4)
delta_est <- vapply(delta_grid, refit_with_delta, numeric(1))
tipping_delta <- uniroot(refit_with_delta, interval = c(0, 40))$root
missing_by_group <- table(missing_visits$group)
data.frame(delta = delta_grid, interaction = round(delta_est, 2))
```

```
#>   delta interaction
#> 1    -8       -1.43
#> 2    -4       -1.21
#> 3     0       -0.99
#> 4     4       -0.77
#> 5     8       -0.55
#> 6    12       -0.33
#> 7    16       -0.11
#> 8    20        0.11
#> 9    24        0.33
```

```
c(early_only = round(est_early, 2), tipping_point_delta = round(tipping_delta, 1))
```

```
#>          early_only tipping_point_delta
#>               -1.45               18.10
```

```
stopifnot(nrow(missing_visits) == 216 - nrow(cohort),
          tipping_delta > 0,                                    # the difference vanishe
          missing_by_group[["complete"]] > missing_by_group[["incomplete"]])
```

```
ggplot(data.frame(delta = delta_grid, interaction = delta_est), aes(delta, in
  geom_hline(yintercept = 0, linetype = cdsr_linetypes[[2]], colour = cdsr_co
  geom_hline(yintercept = est_lmm, linetype = cdsr_linetypes[[4]], colour = c
  geom_vline(xintercept = tipping_delta, colour = cdsr_colours[["red"]]) +
  geom_line(colour = cdsr_colours[["blue"]], linewidth = 1) +
  geom_point(colour = cdsr_colours[["blue"]], shape = cdsr_shapes[[1]], size
  annotate("text", x = tipping_delta, y = min(delta_est), hjust = 1.05, size
           label = paste0("tipping point:   = ", cdsr_number(tipping_delta, 1
  labs(x = "Assumed shift   of missing values (µg/L)", y = "Interaction (µg/L
       caption = "Single conditional-mean imputation; point estimates only.")
  theme_cdsr()
```

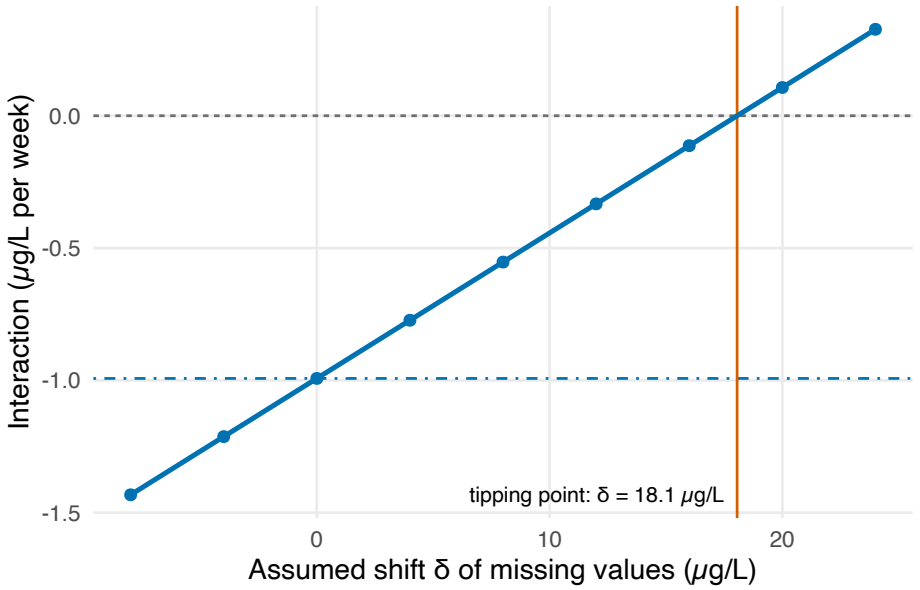


Figure 13.2.: Sensitivity of the interaction (complete minus incomplete, µg/L per week, *synthetic*) to the assumed shift of the missing values against the MAR prediction. Dashed line: no difference; dash-dotted line: estimate from the observed data. n = 36 persons, 30 missing planned visits filled.

With weeks 0–4 only, the interaction is -1.45 µg/L per week (all visits: -0.99). The tipping point lies at = 18.1 µg/L (Figure 13.2), about 4.5 times the assumed measurement noise. Most filled visits belong to the complete-injury group (25 of 30), so optimistic assumptions about the people who left move this group's slope upwards and the difference towards zero. Whether a shift of this size is plausible is a clinical judgement, not a statistical one. It should be made before the data are seen, by people who know why patients leave this kind of unit.

If the result stays unclear. When the tipping point lies within the range that clinicians find plausible, the honest summary is that the data do not settle the question. Report the MAR estimate, the tipping point and the reasons for leaving; do not search for the analysis that gives the desired answer. The next study can record reasons for missed visits, add a home or outpatient sample at week 8 and fix the sensitivity range in the analysis plan.

13.4. Counter-check

Change the dropout mechanism and keep everything else. **Predict first:** (1) If nobody drops out, will the three approaches agree? (2) Under MAR dropout, which approach is unbiased on average? (3) If leaving depends on the unmeasured current level and on the person's own recovery rate (MNAR), in which direction will the mixed-model estimate move? Then run one hundred cohorts per mechanism. The simulation follows the logic of Monte Carlo experiments with explicit assumptions (see Strategy in R: law of large numbers and Monte Carlo).

```
set.seed(1300)
n_rep <- 100
n_singular <- 0                                    # singular fits are counte
estimate_all <- function(obs) {
  per <- do.call(rbind, lapply(split(obs, obs$id), function(d)
    if (nrow(d) >= 2) data.frame(group = d$group[1], slope = coef(lm(seleniu
  m <- suppressMessages(lmer(selenium ~ week * group + (week | id), data = ob
  if (isSingular(m)) n_singular <<- n_singular + 1
  c(pooled = coef(lm(selenium ~ week * group, data = obs))[["week:groupcomple
    per_person = diff(tapply(per$slope, per$group, mean))[[1]],
    mixed = fixef(m)[["week:groupcomplete"]])
}
replicates <- lapply(c(none = "none", mar = "mar", mnar = "mnar"), function(m
```

```
  t(replicate(n_rep, { d <- simulate_cohort(mech); estimate_all(d[d$observed, ]) })))
summary_rep <- do.call(rbind, lapply(names(replicates), function(mech) {
  r <- replicates[[mech]]
  data.frame(mechanism = mech, approach = colnames(r), mean = colMeans(r),
             mc_se = apply(r, 2, sd) / sqrt(n_rep))
}))
rownames(summary_rep) <- NULL
summary_rep$mean <- round(summary_rep$mean, 2); summary_rep$mc_se <- round(summary_rep$mc_se
summary_rep
```

```
#>   mechanism   approach  mean mc_se
#> 1      none     pooled -0.95 0.029
#> 2      none per_person -0.95 0.029
#> 3      none      mixed -0.95 0.029
#> 4       mar     pooled -0.81 0.044
#> 5       mar per_person -1.34 0.072
#> 6       mar      mixed -1.02 0.036
#> 7      mnar     pooled -0.83 0.035
#> 8      mnar per_person -0.93 0.071
#> 9      mnar      mixed -0.89 0.032
```

```
m_of <- function(mech, appr) summary_rep$mean[summary_rep$mechanism == mech & summary_rep$ap
stopifnot(
  max(abs(replicates$none[, "pooled"] - replicates$none[, "mixed"])) < 1e-6,   # complete da
  max(abs(replicates$none[, "per_person"] - replicates$none[, "mixed"])) < 1e-6,
  abs(m_of("mar", "mixed") - (-1)) < 0.1,          # MAR: mixed model unbiased within simula
  m_of("mar", "pooled") > -0.9,                     # MAR: pooled estimate attenuated
  m_of("mnar", "mixed") > m_of("mar", "mixed")      # MNAR: mixed model moves towards zero
)
```

The results answer the three predictions. With complete data, all three approaches give identical estimates in every cohort: in a balanced design, the

mixed model's fixed slope equals the pooled and the per-person slope. Under MAR dropout, the mixed model is unbiased on average (-1.02 against -1.00), the pooled regression is pulled towards zero (-0.81) and the per-person slopes are pushed away from it (-1.34), because people who leave after a low value contribute short, steep lines. Under MNAR, the mixed model also drifts towards zero (-0.89): it can only use what was observed. In 11 of 300 fits, the random-slope variance was estimated at the boundary; these fits are counted and kept. The lesson for the consultant's question is that agreement between baseline and mixed model is reassuring only when few visits are missing. Disagreement is a signal to examine the reasons for leaving.

13.5. Consequences for patients and operations

Patients and ethics. The people who left were mostly from the complete-injury group, and they had low values. They are exactly the people a supplementation trial would target, and the estimate is least certain for them. A protocol should not rest on a slope that describes the people who managed to come back. Each extra blood draw is a burden for people with a new spinal cord injury; the next study should place samples where care already happens (ward rounds, rehabilitation visits) and record every reason for a missed visit. A group-level slope is a research finding, not an individual prognosis, and should not be communicated to patients as one (Chapter 21).

Data protection. A small cohort with a rare injury level, age, sex and exact visit dates can identify a person even without a name. The CSV file of this chapter is synthetic. With real data, share weeks since admission rather than dates, pseudonymised IDs and only the variables needed for the analysis; the questions of access and re-use are treated in Chapter 23 and Chapter 26.

Regulation and security. A slope used as a research end point does not make the analysis a medical device. The question would arise if software computed individual recovery forecasts from biomarker courses and clinicians used them to choose a treatment; Chapter 22 sets out how intended purpose drives that classification. Information security is indirectly relevant: the analysis data set, the code and the package versions must be versioned so that the numbers in the meeting can be reproduced (Chapter 4).

13.6. Look up and transfer

Reviewer question. How many persons, not measurements, contribute to the late part of each group's course, and how far does the interaction move under a plausible MNAR scenario?

Self-test. A second synthetic cohort studies zinc after spinal cord injury. Here, people who recover well are discharged home after week 4 and their week-6 and week-8 values are missing; people who do poorly stay on the ward. In which direction is the pooled regression biased for the group with most discharges, and what would a sensitivity analysis need to assume?

Answer

The late measurements now come from people with a poorer course, so the pooled slope of the group with most discharges is too flat. The mixed model under MAR corrects this only as far as the early values already reveal who is recovering well. If discharge depends on improvement that was not yet measured, the data are MNAR, and the mixed model is also too flat for that group. A sensitivity analysis would shift the missing late values upwards (positive) and report the at which the conclusion changes. Recording the reason "discharged home" makes the mechanism visible.

Annotated sources.

- Matthews et al. (1990) – the case for summary measures such as the slope per person; the source of the baseline.
- Laird and Ware (1982) – the random-effects model for longitudinal data that underlies the mixed model.
- Bates et al. (2015) – the implementation used in the lab, including the model syntax and estimation method.
- Rubin (1976) – the formal definition of missing at random and when the missing-data mechanism can be ignored.
- Diggle and Kenward (1994) – why informative dropout biases longitudinal analyses and how to model it.
- Little et al. (2012) – recommendations for preventing missing data and for sensitivity analyses in clinical studies.
- Seelig et al. (2020) – a published observational study of selenium and copper after spinal cord injury; it motivates the question only, no values are used.

13.7. Further reading

- Chapter 3 – how missing values and denominators arise in clinical data, before any model.
- Chapter 16 – how much a small sample can say, and how to state the remaining uncertainty.
- Methods in R: Random-Slope Models – model formulation, diagnostics and warning signs.
- Methods in R: Sensitivity Analyses in Clinical Trials – the wider family of sensitivity analyses for missing outcomes.
- CTTIR Course 2, Week 2, Session 3 – RCBD and repeated measures → mixed models – a step-by-step lab from repeated measures to mixed models.

- CTTIR Course 3, Week 2, Session 1 – MCAR, MAR, MNAR and Session 3 – Linear mixed models with lme4 – missing-data mechanisms and their consequences for longitudinal models.
- CTTIR Tutorials: Mixed-Effects Models – Introduction – a short refresher on random effects and shrinkage.

14. Cluster, gradient or artefact

R. Heller · Date: 2026-08-30 · Last changed: 2026-11-09

At a glance

Guiding question: When does a figure become a testable subtype hypothesis?
For: translational researchers, doctoral students and reviewers who meet cluster maps of patients · **Prerequisites:** denominators and units (Chapter 3), the person as independent unit (Chapter 13), basic idea of distances between observations
In the course: W5 · Causality, trajectories and phenotypes · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

14.1. Case and decision

A fictitious trauma research group has six routine values from the first day after major trauma in 300 people: C-reactive protein, interleukin-6, lactate, leukocytes, platelets and body temperature. A colleague shows a two-dimensional map in which the patients form three islands. The proposal is to call them “hyperinflammatory”, “intermediate” and “quiet” subtypes, to name them in the manuscript and to stratify a planned trial by subtype. The group leader has to decide whether the word “subtype” goes into the manuscript and the trial protocol. The data in this chapter are *synthetic* and were generated, on purpose, from a single continuous severity

gradient without any groups. The task is to find out which analysis would have revealed that – and what evidence would be needed before a subtype label may change what happens to a patient.

14.2. Core reading

Core reading · recommended for W5 · about 10 minutes · ends before ‘Lab’

What the table contains. One row per person, six columns with six units: mg/L, pg/mL, mmol/L, 10^9/L twice and °C. The independent unit is the person; n = 300. Any clustering method needs a distance between two people, and that distance depends on how the six columns are put on a common scale. Without scaling, platelets (standard deviation 59×10^9/L) and interleukin-6 (61 pg/mL) dominate, while lactate (0.8 mmol/L) and temperature (0.5 °C) hardly count. After standardisation to z-scores, each marker counts equally. After a log transformation of the skewed markers, the distances change again. Scaling is a modelling decision with clinical content, not a technical preliminary. In the lab, k-means on the original units agrees with the z-score solution only with an adjusted Rand index of 0.37 (1 means identical partitions, 0 chance agreement): changing the units reassigns a large share of the patients.

A gradient cut into slices. A principal component analysis (PCA) of the z-scores shows one dominant axis: the first component carries 62.5% of the variance and correlates at 0.93 with the hidden severity used to generate the data. k-means clustering with three clusters (Hartigan and Wong 1979) does what it is built to do: it minimises the sum of squared distances between each person and the nearest of three centres. On a gradient, it returns three slices (Figure 14.1). The left panel looks like three groups because the colours say so. The right panel shows the decisive

detail: the cluster boundaries do not sit in empty regions. They cut through the densest part of the distribution.

Methods that agree for the wrong reason. Three further analyses seem to confirm the islands. A Gaussian mixture model (**mclust**, Scrucca et al. (2016)) selects 6 components by the Bayesian information criterion (BIC) on z-scores and 3 on log-transformed markers. BIC counts how many bell-shaped components are needed to describe the shape of the cloud; a skewed or curved cloud needs several, even without groups. Bootstrap stability looks reassuring as well: when the 300 persons are resampled and re-clustered a hundred times, the three clusters are recovered with mean Jaccard similarities between 0.77 and 0.91 (Hennig 2007). By the usual guidance for this measure, values above 0.75 count as stable (Hennig 2026). Stability, however, shows that the same cut is made each time, not that there is a seam to cut along. A knife set at the same place gives reproducible slices of a loaf.

The question behind the question. Every clustering algorithm returns as many groups as it is asked for. Asking "how many clusters are there?" therefore presupposes that there are clusters. The first question is a different one: is there any discontinuity in the population, or do the patients vary along one or a few continuous axes? Only if the answer is the former does the choice of k become meaningful.

What would distinguish groups from a gradient. A real group structure leaves a trough in the density between groups. A simple diagnostic compares the density of the first component at the boundary between two k-means clusters with the density at the two cluster centres. Here, the ratio is 1.18: the boundary lies where people are *more* frequent than at the centres. The counter-check (Section 14.4) builds a data set with a real gap and shows how this ratio falls.

UMAP islands. Two-dimensional maps such as UMAP (McInnes et al. 2018) are built to keep each person's nearest neighbours close. The documentation of the reference implementation states that UMAP, like t-SNE,

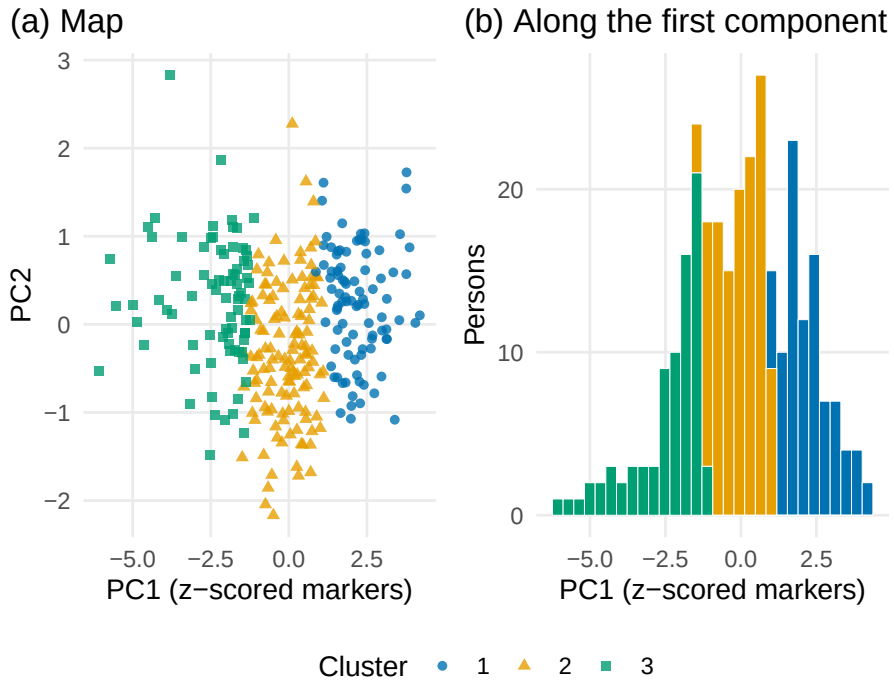


Figure 14.1.: k-means with k = 3 on z-scored markers (*synthetic*, n = 300 persons). (a) First two principal components, colour and shape by cluster. (b) Distribution of the first principal component, stacked by cluster: the clusters are adjacent slices without gaps between them.

does not completely preserve density and can create false tears in clusters, yielding a finer clustering than is present in the data (McInnes 2026). Island sizes, gaps and distances between islands are not measurements; Chari and Pachter (2023) show how strongly such maps can distort the structure of single-cell data. A map is a reason to look, not evidence of subtypes.

The baseline. The biometrically sensible baseline for a gradient is a continuous score, here the first principal component. In the lab, a model for a *synthetic* day-7 organ dysfunction outcome fits better with the continuous score (AIC 326.0) than with the three cluster labels (AIC 337.1). Cutting the gradient into slices throws information away. A label also hides similarity: two patients who lie just either side of a boundary receive different names, while two patients at opposite ends of the same slice share one.

The most important assumption. A cluster label claims a discontinuity in the population. It is supported only if a gap survives changes of scaling, algorithm and sample, reappears in an independent cohort and the label adds information beyond a continuous score. Without this, "three subtypes" describes the analysis, not the patients.

Question for the workshop. Which observation would convince you that the three islands are subtypes rather than slices: the same islands after log transformation, a trough in the density, reproduction in another centre or a different treatment effect in a randomised trial? Rank the four and justify the order.

14.3. Lab

The lab needs **mclust**, **cluster**, **ggplot2** and **patchwork** and runs in a few seconds.

14.3.1. Step 1: simulate a gradient without groups

All values are *synthetic assumptions*. A single latent severity score drives all six markers; three of them are right-skewed, as many inflammatory markers are. The day-7 outcome depends on the same score, continuously.

```
set.seed(1414)                                       # same data each time this
n <- 300
make_markers <- function(severity) {                 # one latent score drives a
  m <- length(severity)
  data.frame(
    crp_mg_l          = pmax(1, 60 + 35 * severity + rnorm(m, 0, 20)),
    il6_pg_ml         = exp(3.6 + 0.8 * severity + rnorm(m, 0, 0.45)),
    lactate_mmol_l    = exp(log(1.8) + 0.30 * severity + rnorm(m, 0, 0.25)),
    leukocytes_10e9_l = pmax(1, 11 + 2.5 * severity + rnorm(m, 0, 2.5)),
    platelets_10e9_l  = pmax(20, 220 - 35 * severity + rnorm(m, 0, 45)),
    temperature_c     = 37.4 + 0.4 * severity + rnorm(m, 0, 0.35))
}
severity <- rnorm(n)                                 # continuous, unimodal: no
markers <- make_markers(severity)
organ_dysfunction <- rbinom(n, 1, plogis(-1.2 + 1.3 * severity))   # synthet
round(sapply(markers, function(x) c(median = median(x), sd = sd(x))), 2)
```

```
#>        crp_mg_l il6_pg_ml lactate_mmol_l leukocytes_10e9_l platelets_10e9_
#> median    62.95     38.73           1.81             10.94           213.7
#> sd        39.66     61.02           0.84              3.56            59.0
#>        temperature_c
#> median         37.46
#> sd              0.53
```

14.3.2. Step 2: three scalings and a PCA

```
log_vars <- c("crp_mg_l", "il6_pg_ml", "lactate_mmol_l")
x_raw <- as.matrix(markers)                       # original units
x_z <- scale(markers)                             # z-scores: every marker counts equally
markers_log <- markers
markers_log[log_vars] <- lapply(markers_log[log_vars], log)
x_logz <- scale(markers_log)                      # log for skewed markers, then z-scores
pca <- prcomp(x_z)
pc1_share <- summary(pca)$importance["Proportion of Variance", "PC1"]
cor_pc1 <- cor(pca$x[, 1], severity)
round(summary(pca)$importance[, 1:3], 3)
```

```
#>                           PC1   PC2   PC3
#> Standard deviation     1.936 0.773 0.702
#> Proportion of Variance 0.625 0.100 0.082
#> Cumulative Proportion  0.625 0.724 0.806
```

14.3.3. Step 3: k-means with two to five clusters

The mean silhouette width compares, for every person, the distance to the own cluster with the distance to the nearest other cluster (Rousseeuw 1987); values near 1 indicate well-separated clusters.

```
library(cluster)
library(mclust)
set.seed(14141)
d_z <- dist(x_z)
km <- lapply(2:5, function(k) kmeans(x_z, centers = k, nstart = 25))
silhouette_mean <- sapply(km, function(m) mean(silhouette(m$cluster, d_z)[, "sil_width"]))
names(silhouette_mean) <- paste0("k=", 2:5)
```

```
cluster_ref <- km[[2]]$cluster                    # reference solution: z-sc
cluster_raw <- kmeans(x_raw, centers = 3, nstart = 25)$cluster
cluster_log <- kmeans(x_logz, centers = 3, nstart = 25)$cluster
ari_raw <- adjustedRandIndex(cluster_ref, cluster_raw)   # 1 = identical part
ari_log <- adjustedRandIndex(cluster_ref, cluster_log)
round(c(silhouette_mean, ARI_raw_vs_z = ari_raw, ARI_log_vs_z = ari_log), 3)
```

```
#>          k=2          k=3          k=4          k=5 ARI_raw_vs_z ARI_log_
#>        0.335        0.247        0.233        0.192        0.368        0
```

The adjusted Rand index (ARI) compares two partitions of the same people and is 0 for chance agreement (Hubert and Arabie 1985). Changing only the scaling reassigns a large share of the patients.

14.3.4. Step 4: a Gaussian mixture model

```
mc_z <- Mclust(x_z, G = 1:6, verbose = FALSE)
mc_log <- Mclust(x_logz, G = 1:6, verbose = FALSE)
c(z_scores = mc_z$G, log_z = mc_log$G)
```

```
#> z_scores    log_z
#>        6        3
```

```
c(model_z = mc_z$modelName, model_log = mc_log$modelName)
```

```
#>   model_z model_log
#>     "VVI"     "EEE"
```

14.3.5. Step 5: bootstrap stability

For each resample, the persons are drawn with replacement, k-means is rerun, and every original cluster is compared with its best-matching new cluster on the persons present in both (Jaccard similarity: shared persons divided by persons in either). This follows the logic of Hennig's cluster-wise stability assessment; the same procedure is available as `fpc::clusterboot()` (Hennig 2007, 2026).

```
bootstrap_stability <- function(x, k, n_boot = 100) {
  ref <- kmeans(x, centers = k, nstart = 25)$cluster
  jac <- matrix(NA_real_, n_boot, k)
  ari <- numeric(n_boot)
  for (b in seq_len(n_boot)) {
    idx <- sample(nrow(x), replace = TRUE)
    new <- kmeans(x[idx, ], centers = k, nstart = 10)$cluster
    present <- unique(idx)
    new_present <- new[match(present, idx)]       # new cluster of each person in the resamp
    for (j in seq_len(k)) {
      a <- present[ref[present] == j]
      jac[b, j] <- max(sapply(seq_len(k), function(m) {
        bset <- present[new_present == m]
        length(intersect(a, bset)) / length(union(a, bset))
      }))
    }
    ari[b] <- adjustedRandIndex(ref[present], new_present)
  }
  list(jaccard = colMeans(jac), ari_median = median(ari))
}
set.seed(14142)
stab <- lapply(2:4, function(k) bootstrap_stability(x_z, k))
jac <- list(k2 = stab[[1]]$jaccard, k3 = stab[[2]]$jaccard, k4 = stab[[3]]$jaccard)
lapply(jac, round, 2)
```

```
#> $k2
#> [1] 0.95 0.94
#>
#> $k3
#> [1] 0.82 0.91 0.77
#>
#> $k4
#> [1] 0.89 0.96 0.84 0.72
```

14.3.6. Step 6: look for a trough instead of a cluster count

```
valley_ratio_of <- function(x) {
  pc <- prcomp(x)$x[, 1]
  cl <- kmeans(x, centers = 2, nstart = 25)$cluster
  centres <- tapply(pc, cl, mean)
  dens <- density(pc, bw = "SJ")                  # kernel density of the fi
  f <- approxfun(dens$x, dens$y)
  f(mean(centres)) / mean(f(centres))             # density at the boundary /
}
set.seed(14143)
valley_ratio <- valley_ratio_of(x_z)
round(valley_ratio, 2)
```

```
#> [1] 1.18
```

A ratio well below 1 points to a trough between two groups; a ratio near or above 1 means that the boundary runs through a dense region. The ratio is a didactic diagnostic, not a formal test, and it looks only along the first component.

14.3.7. Step 7: the continuous baseline

```
fit_cluster <- glm(organ_dysfunction ~ factor(cluster_ref), family = binomial)
fit_score <- glm(organ_dysfunction ~ pca$x[, 1], family = binomial)
aic_cluster <- AIC(fit_cluster)
aic_pc1 <- AIC(fit_score)
c(AIC_three_clusters = round(aic_cluster, 1), AIC_continuous_score = round(aic_pc1, 1))
```

```
#>   AIC_three_clusters AIC_continuous_score
#>                337.1                326.0
```

14.3.8. Step 8: collect the variants and check the target values

```
variants <- data.frame(
  Variant = c("z-scores, k-means, k = 2", "z-scores, k-means, k = 3 (reference)", "z-scores,
              "original units, k-means, k = 3", "log + z-scores, k-means, k = 3",
              "z-scores, Gaussian mixture (BIC)", "log + z-scores, Gaussian mixture (BIC)"),
  Silhouette = c(cdsr_number(silhouette_mean[1:3], 2), "", "", "", ""),
  Jaccard = c(cdsr_number(sapply(jac, min), 2), "", "", "", ""),
  ARI = c("", "1.00", "", cdsr_number(ari_raw, 2), cdsr_number(ari_log, 2), "", ""),
  Components = c("2", "3", "4", "3", "3", as.character(mc_z$G), as.character(mc_log$G)))
knitr::kable(variants, align = "lrrrr")
```

Table 14.1.: Results of the clustering variants on the same 300 synthetic persons. Silhouette: mean silhouette width; Jaccard: smallest mean bootstrap Jaccard similarity of the clusters (100 resamples); ARI: agreement with the reference solution (z-scores, k-means, k = 3).

Variant	Silhouette	Jaccard	ARI	Components
z-scores, k-means, k = 2	0.34	0.94		2
z-scores, k-means, k = 3 (reference)	0.25	0.77	1.00	3
z-scores, k-means, k = 4	0.23	0.72		4
original units, k-means, k = 3			0.37	3
log + z-scores, k-means, k = 3			0.46	3
z-scores, Gaussian mixture (BIC)				6
log + z-scores, Gaussian mixture (BIC)				3

```
stopifnot(
  pc1_share > 0.5, abs(cor_pc1) > 0.9,            # one dominant axis that tr
  mc_z$G > 1, mc_log$G > 1,                       # BIC prefers several compo
  min(jac$k3) > 0.75,                             # "stable" by the usual gui
  valley_ratio > 0.8,                             # ... but no trough at the
  ari_raw < 0.6,                                  # scaling changes the parti
  aic_pc1 < aic_cluster                           # the continuous score is t
)
```

14.3.9. Step 9: an external test plan

A subtype hypothesis becomes testable when the analysis is frozen before new data are seen. Table 14.2 lists a minimal plan; each line names what would count against the hypothesis.

Table 14.2.: Minimal external test plan for a subtype hypothesis.

Element	Fixed in advance	Counts against the subtype hypothesis
Pipeline	markers, units, transformation, scaling, algorithm, k, seed, centres	any change after seeing the new cohort
Cohort	another centre or period, same assays and units, consecutive patients	selected or differently measured cohort
Reproduction	assign new patients to the frozen centres; recluster independently; compare (ARI, prediction strength (Tibshirani and Walther 2005))	low agreement between frozen and new partition
Discontinuity	trough between clusters along the main axis	boundary in a dense region (ratio near or above 1)
Added value	clusters versus continuous score for a pre-specified outcome	no gain over the continuous score
Consequence	differential treatment effect in randomised data	effect identical across clusters

The last line is what the trial stratification would need. A published example of the kind of evidence it requires is the latent class analysis of two randomised trials in acute respiratory distress syndrome, in which two

subphenotypes were found in both cohorts and the effect of the ventilation strategy differed by subphenotype in a secondary analysis (Calfee et al. 2014). Even there, the authors framed the result as useful for selecting patients for future trials, not as a basis for treatment choice.

> **Other implementation: Python/umap-learn**
>
> The reference implementation of UMAP is the Python package umap-learn (version 0.5.12, April 2026; documentation) (McInnes et al. 2018). Its concrete advantage is direct access to the original algorithm and its parameters (`n_neighbors`, `min_dist`, `metric`), which matters when a published map is to be reproduced. For a fair comparison with an R workflow, the input matrix (same persons, same transformation and scaling), the parameters, the random seed and the number of dimensions must be identical; even then, different random initialisations can move islands. Status: conceptual comparison, not run. Any cluster seen on such a map must pass the same checks as above on the original, unprojected data.

If the result stays unclear. "No evidence for discrete subtypes; the markers vary along one continuous axis" is a finding worth reporting. It comes with the continuous score, its distribution and the outcome model, which others can test. When some diagnostics support groups and others do not, report all of them and keep the language at "candidate clusters" until the external plan has been run.

14.4. Counter-check

Change the data, not the method: generate the same six markers from a severity score with a real gap – two groups of 150 people centred at -1.2 and $+1.2$ with a within-group spread of 0.45 (*assumption*). **Predict first:** which of the diagnostics will change markedly – the BIC component

count, the bootstrap Jaccard similarity, the silhouette width or the trough ratio?

```
set.seed(1415)
severity_gap <- c(rnorm(n / 2, -1.2, 0.45), rnorm(n / 2, 1.2, 0.45))
x_gap <- scale(make_markers(severity_gap))
km_gap <- kmeans(x_gap, centers = 2, nstart = 25)
km_grad2 <- km[[1]]
compare <- data.frame(
  Diagnostic = c("mean silhouette, k = 2", "smallest bootstrap Jaccard, k = 2", "trough rati
                 "Gaussian mixture components (BIC)"),
  Gradient = c(silhouette_mean[["k=2"]], min(jac$k2), valley_ratio, mc_z$G),
  Gap = c(mean(silhouette(km_gap$cluster, dist(x_gap))[, "sil_width"]),
          min(bootstrap_stability(x_gap, 2)$jaccard),
          valley_ratio_of(x_gap),
          Mclust(x_gap, G = 1:6, verbose = FALSE)$G))
compare$Gradient <- round(compare$Gradient, 2); compare$Gap <- round(compare$Gap, 2)
compare
```

```
#>                          Diagnostic Gradient  Gap
#> 1            mean silhouette, k = 2     0.34 0.49
#> 2 smallest bootstrap Jaccard, k = 2     0.94 0.99
#> 3                       trough ratio     1.18 0.35
#> 4 Gaussian mixture components (BIC)     6.00 4.00
```

```
stopifnot(compare$Gap[3] < 0.5, compare$Gradient[3] > 0.8,      # the trough separates the
          compare$Gap[1] > compare$Gradient[1],                   # silhouette rises, but m
          abs(compare$Gap[2] - compare$Gradient[2]) < 0.1)        # stability hardly differ
```

```
boundary_of <- function(x, cl) mean(tapply(prcomp(x)$x[, 1], cl, mean))
hist_df <- rbind(data.frame(scenario = "Gradient (no groups)", pc1 = prcomp(x_z)$x[, 1]),
                 data.frame(scenario = "Two groups with a gap", pc1 = prcomp(x_gap)$x[, 1]))
```

```
bounds <- data.frame(scenario = c("Gradient (no groups)", "Two groups with a
                     b = c(boundary_of(x_z, km_grad2$cluster), boundary_of(x_
ggplot(hist_df, aes(pc1)) +
  geom_histogram(bins = 30, fill = cdsr_colours[["skyblue"]], colour = "white
  geom_vline(data = bounds, aes(xintercept = b), colour = cdsr_colours[["red"
  facet_wrap(~ scenario) +
  labs(x = "PC1 (z-scored markers)", y = "Persons") +
  theme_cdsr()
```

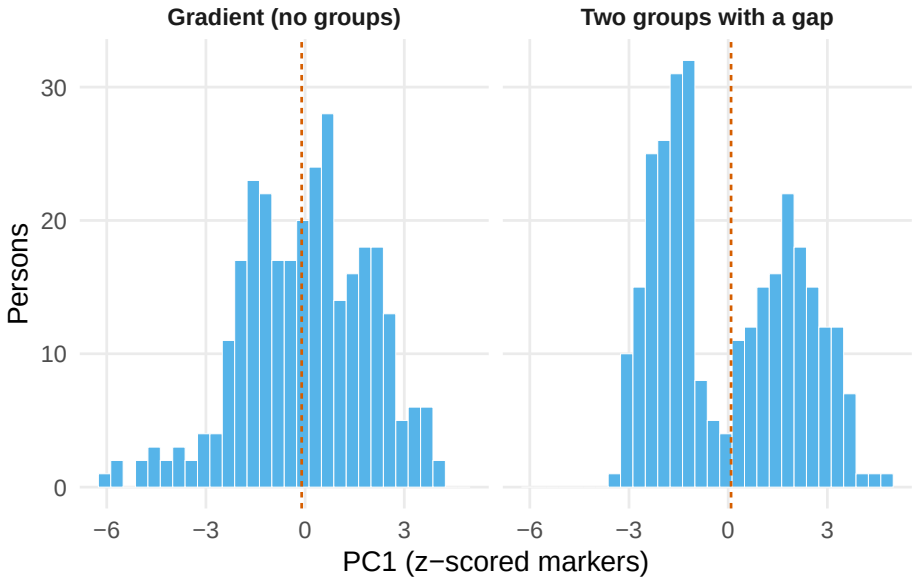


Figure 14.2.: First principal component of z-scored markers for a continuous gradient and for two groups with a real gap (*synthetic*, n = 300 persons each). Vertical line: boundary between the two k-means clusters.

The trough ratio falls from 1.18 to 0.35, and the silhouette width rises from

0.34 to 0.49. Bootstrap stability hardly changes (0.94 versus 0.99): it was already high for the gradient. The mixture model selects 4 components for two groups, which again shows that the BIC count describes the shape of the density, not the number of patient groups. Only the diagnostics that look for empty space between groups separate the two situations.

14.5. Consequences for patients and operations

Patients and ethics. A subtype label travels. It enters records, referral letters and conversations, and it can change expectations about recovery. A person labelled "hyperinflammatory" because their values lie just across an arbitrary boundary may receive a different treatment from a neighbour with almost identical values. A subtype name therefore needs a justified consequence: a treatment, monitoring or prognosis decision that differs by subtype and has been shown to do so, as in the external plan above. Without it, the continuous score is more honest and loses less information. If the markers differ systematically by sex, age or laboratory method, clusters can become proxies for these characteristics; check the distribution of such variables across clusters before naming them (Chapter 21).

Regulation. Exploratory clustering in a research data set is not a medical device. If software assigned patients to subtypes in order to select a treatment, its intended purpose would have to be assessed under the medical device rules (Chapter 22).

Data protection and security. A cluster label derived from personal data is itself personal data. Small clusters with unusual marker combinations can make individuals recognisable in published figures or supplementary tables (Chapter 23). Security is only indirectly relevant: the frozen pipeline of the external plan, including centres and seed, must be stored and versioned so that the test cannot be adjusted after the fact.

14.6. Look up and transfer

Reviewer question. Did the authors show a gap between clusters that survives a change of scaling and reappears in an independent cohort, and does the cluster label predict the outcome better than a continuous score?

Self-test. A manuscript reports four "immune subtypes" after polytrauma, found with k-means on z-scored markers. The authors cite mean bootstrap Jaccard similarities above 0.85 and a BIC-selected mixture model with four components as validation. What is missing?

Answer

Both arguments are compatible with a continuous gradient: k-means cuts a gradient reproducibly, and BIC adds components to describe a skewed or curved density. Missing are evidence of a discontinuity (a trough or clear separation that survives other scalings and algorithms), reproduction in an independent cohort with a frozen pipeline, and a comparison with a continuous score. A treatment consequence would additionally need a differential effect in randomised data.

Annotated sources.

- Hennig (2007) – cluster-wise bootstrap stability with the Jaccard similarity; the method behind Step 5.
- Hennig (2026) – documentation of `clusterboot()`, including the usual interpretation thresholds for the mean Jaccard similarity.
- Scrucca et al. (2016) – the Gaussian mixture models and BIC-based model choice used in Step 4.
- Hartigan and Wong (1979) – the k-means algorithm that R's `kmeans()` uses by default.
- Rousseeuw (1987) and Hubert and Arabie (1985) – silhouette width and adjusted Rand index.

- Tibshirani and Walther (2005) – prediction strength as a check of whether clusters reproduce in new data.
- McInnes et al. (2018), McInnes (2026) and Chari and Pachter (2023) – how UMAP works, the warning about false tears, and how far two-dimensional embeddings can distort structure.
- Calfee et al. (2014) – a published subphenotype analysis that tested reproduction in two cohorts and a differential treatment effect.

14.7. Further reading

- Chapter 12 – phenotypes from segmented tissue images, where cluster boundaries also depend on upstream choices.
- Chapter 19 – how to plan the external cohort that the test plan needs.
- Methods in R: Gaussian Mixture Models and k-Means Clustering – assumptions, diagnostics and warning signs of both methods.
- Methods in R: UMAP and t-SNE – Overview – what two-dimensional embeddings preserve and what they do not.
- Methods in R: Centring and Scaling Predictors – why the scale of a variable changes its weight.
- CTTIR Course 4, Week 1, Session 4 – Clustering: k-means, hierarchical, model-based and Session 5 – UMAP and t-SNE – hands-on labs with the same methods.
- CTTIR Tutorials: Cluster Validation – Silhouette – interpreting silhouette widths.

15. Mechanistic models and pharmacometrics

R. Heller · Date: 2026-04-12 · Last changed: 2026-11-09

At a glance

Guiding question: Which sampling times make it possible to separate clearance from volume?
For: clinical pharmacology trainees, doctoral students planning sampling schedules, reviewers of pharmacokinetic analyses · **Prerequisites:** exponential functions and logarithms, the idea of a likelihood (Chapter 16 uses it as well), repeated measurements per person (Chapter 13)
In the course: W6 · Pharmacometrics, forecasting and simulation · **Data:** synthetic, virtual parameters only (generated in the chapter) · **Core reading time:** about 11 minutes

⚠ Virtual parameters only

Every dose, volume and clearance in this chapter is a *virtual* model input for a fictitious compound. Nothing here is a dosing recommendation, and the numbers do not describe any real drug or patient group.

15.1. Case and decision

A doctoral researcher in clinical pharmacology is writing the sampling section of a protocol for a fictitious intravenous compound in injured adults. The ethics committee has asked the team to keep blood sampling to a minimum, and the trial unit can process two samples per person. The researcher must choose the two times. A senior colleague suggests taking both samples late, at 6 and 8 hours, "because by then the drug is evenly distributed". Another proposes one sample shortly after the injection and one late. The protocol will fix the times for every participant, and they cannot be changed after the first person is enrolled. Which schedule allows the clearance and the volume of distribution to be estimated separately – and how precisely? The chapter answers with a one-compartment model and *virtual* parameters only.

15.2. Core reading

Core reading · required for W6 (alternative to Chapter 18) · about 11 minutes · ends before 'Lab'

From mass balance to a concentration curve. A one-compartment model treats the body as one well-mixed volume. Let A(t) be the amount of compound in the body (mg) and C(t) the plasma concentration (mg/L). The volume of distribution V (L) links the two: C = A / V. It is a proportionality constant, not an anatomical space. The clearance CL (L/h) is the volume of plasma cleared of the compound per hour, so the elimination rate is CL × C, in L/h × mg/L = mg/h. The mass balance says that the amount in the body falls exactly by what is eliminated:

$$\frac{dA}{dt} = -CL \cdot C = -\frac{CL}{V}\,A, \qquad A(0) = \text{dose}.$$

The units check: (L/h) / L × mg = mg/h. After a single intravenous bolus, the solution is $C(t) = (\text{dose} / V) \cdot \exp(-(CL / V) \cdot t)$. The ratio $k = CL / V$ (1/h) is the elimination rate constant, the half-life is $\ln 2 \cdot V / CL$, and the area under the curve from zero to infinity is dose / CL (mg · h/L). With the virtual values dose = 100 mg, V = 20 L and CL = 5 L/h, the initial concentration is 5 mg/L, the half-life 2.77 h and the AUC 20 mg · h/L.

Two parameters, one sample. On a logarithmic scale the curve is a straight line: $\ln C(t) = \ln(\text{dose} / V) - k \cdot t$. The intercept carries the volume, the slope carries the ratio CL / V. One sample, or two samples at the same time, gives one equation for two unknowns. Panel (a) of Figure 15.1 shows two virtual parameter sets that produce exactly the same concentration at 4 h: V = 20 L with CL = 5 L/h, and V = 26 L with CL = 4.79 L/h. A design that measures only at 4 h cannot tell them apart. This is structural non-identifiability (Bellman and Åström 1970): no amount of measurement precision helps.

Where a sample is informative. Panel (b) shows the normalised sensitivities: how much ln C changes when ln V or ln CL changes. For the volume, the sensitivity is $-1 + k \cdot t$; it is zero at $t = V / CL = 4$ h. A sample taken then is blind to V, given CL. For the clearance, the sensitivity is $-k \cdot t$; early samples hardly inform it. The two parameters are separated best when the samples sit at times where the two curves differ in shape, early and late.

Two samples, four designs. With a proportional measurement error of 10% (*assumption*), the expected standard errors follow from these sensitivities (the Fisher information, linearised at the virtual values). Table 15.1 compares four designs. The early-and-late design A estimates both parameters to within about 7–11%. Two late samples (design B) leave the volume uncertain by about 50%, because the intercept must be extrapolated back over six hours. Design B is structurally identifiable, yet practically the data hardly constrain V – the distinction between structural and practical identifiability (Raue et al. 2009). Two early samples (design C) estimate

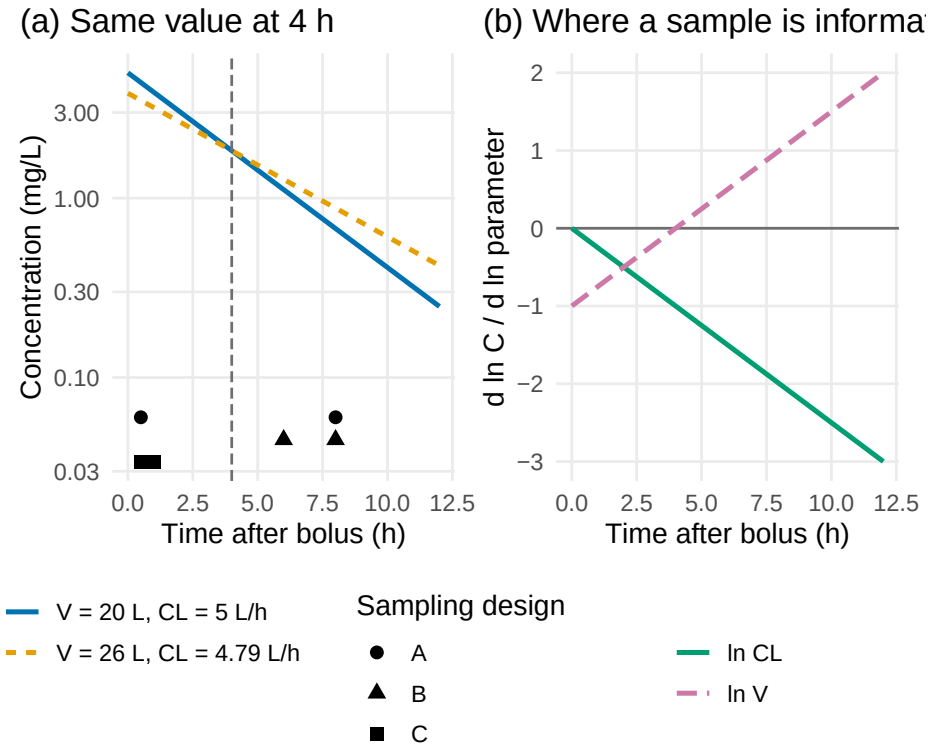


Figure 15.1.: Identifiability in a one-compartment model with *virtual* parameters (dose 100 mg intravenous bolus). (a) Concentration (mg/L, log scale) for two parameter sets that coincide at 4 h; symbols mark the sampling times of designs A–C. (b) Normalised sensitivity of ln C to ln V and ln CL; at t = V/CL = 4 h, a sample carries no information about V. One virtual person, no measurement error.

the volume reasonably but the clearance hardly at all. Two samples at the same time (design D) cannot separate the parameters at all. The simulated spread over 2,000 virtual repetitions agrees with the linearised values. Intuitively, a straight line through two points that are close together is poorly anchored: a 10% error at either point tilts the line, and the tilt is magnified when the line is extended back to time zero to read off the volume, or forward to infinity to obtain the area under the curve.

Table 15.1.: Approximate relative standard errors (%) of V and CL for four two-sample designs, *virtual* parameters (V = 20 L, CL = 5 L/h), proportional error 10% (*assumption*). Expected: from the linearised Fisher information; simulated: standard deviation of the estimates over 2,000 virtual repetitions of one person.

Design	Sampling times (h)	V expected	V simulated	CL expected	CL simulated
A: early and late	0.5 and 8	11	11	7	7
B: both late	6 and 8	50	50	22	21
C: both early	0.5 and 1	22	23	92	85
D: same time twice	4 and 4	not identifiable	not identifiable	not identifiable	not identifiable

The most important assumption. These standard errors hold only if the structural model is right and the parameters are near the values used to compute them. The best design depends on the very parameters it is meant to estimate; the counter-check shows what happens when the clearance is lower than expected. Model-based design is therefore local: it needs prior knowledge and should be checked across a plausible parameter range.

Checking the arithmetic before the biology. Most mechanistic models have no closed-form solution and are solved numerically. Before any parameter is estimated, the lab therefore checks that the numerical solver reproduces the analytical curve of this simple model within a tolerance fixed in advance. A closed mass balance – the amount in the body plus the amount eliminated always equals the dose – is necessary but does not prove that the solution is accurate.

A baseline without a model. With many samples per person, non-compartmental analysis (NCA) estimates the clearance without assuming a compartment model: the AUC is computed from the measured concentrations with the trapezoidal rule and extrapolated to infinity, and CL = dose / AUC. The lab uses it as the baseline for a virtual person with nine samples. With two samples per person, NCA is not available, and the model carries the inference.

Beyond one person. Real participants differ: the clearance and the volume vary between people (interindividual variability), and every measurement carries assay and handling error (residual variability). Population pharmacokinetic models estimate both at once and let different people be sampled at different times (Mould and Upton 2012, 2013). Automated searches over model structures, for example with genetic algorithms (Sale and Sherer 2015), are a further step. Neither repairs a design in which the data cannot separate the parameters.

Question for the workshop. The colleague's argument for late samples was that "the drug is evenly distributed by then". In a one-compartment model this is true from the start. Which feature of a real compound would make early samples misleading, and what would that change in the design?

15.3. Lab

The lab needs **deSolve**, **ggplot2** and **patchwork** and runs in about two seconds.

15.3.1. Step 1: parameters with units

```
dose_mg <- 100           # virtual intravenous bolus (mg) - model input, not a dosing recomm
V_l <- 20                # virtual volume of distribution (L)
CL_l_h <- 5              # virtual clearance (L/h)
k_per_h <- CL_l_h / V_l  # elimination rate constant (1/h)
sigma_log <- 0.1         # proportional measurement error, SD on the log scale (assumption)
conc_analytic <- function(t, dose = dose_mg, V = V_l, CL = CL_l_h) dose / V * exp(-CL / V *
c(C0_mg_l = dose_mg / V_l, half_life_h = log(2) * V_l / CL_l_h, AUC_mg_h_l = dose_mg / CL_l_
```

```
#>     C0_mg_l half_life_h  AUC_mg_h_l
#>    5.000000    2.772589   20.000000
```

15.3.2. Step 2: the numerical solver against the analytical solution

The model is written as a mass balance with two states: the amount in the body (A) and the cumulative amount eliminated (E). Before running the solver, the tolerance is fixed: the largest relative deviation from the analytical solution at the output times must not exceed 10^{-4}. This is a thousand times smaller than the assumed 10% measurement error and below the precision of three significant digits used for reporting, so solver error cannot change any conclusion.

```
library(deSolve)
pk_rhs <- function(t, state, parms) {
  with(as.list(c(state, parms)), {
    conc <- A / V                                  # mg/L
    list(c(A = -CL * conc,                         # mg/h leaving the body
           E = CL * conc))                         # mg/h arriving in the eli
  })
}
times <- seq(0, 24, by = 0.25)                     # output grid (h)
state0 <- c(A = dose_mg, E = 0)
parms <- c(CL = CL_l_h, V = V_l)
solve_with <- function(method, ...) ode(y = state0, times = times, func = pk_
solutions <- list(
  `lsoda, rtol = atol = 1e-8` = solve_with("lsoda", rtol = 1e-8, atol = 1e-8)
  `lsoda, defaults (1e-6)` = solve_with("lsoda"),
  `rk4, fixed step 0.25 h` = solve_with("rk4"),
  `Euler, fixed step 0.25 h` = solve_with("euler"))
tolerance <- 1e-4
reference <- conc_analytic(times)
solver_check <- data.frame(
  solver = names(solutions),
  max_relative_error = sapply(solutions, function(s) max(abs(s[, "A"] / V_l -
  max_mass_balance_error_mg = sapply(solutions, function(s) max(abs(s[, "A"]
solver_check$passes <- solver_check$max_relative_error <= tolerance
rownames(solver_check) <- NULL
solver_check
```

```
#>                     solver max_relative_error max_mass_balance_error_mg p
#> 1 lsoda, rtol = atol = 1e-8       4.464623e-10              8.526513e-14
#> 2    lsoda, defaults (1e-6)       2.921983e-08              7.105427e-14
#> 3    rk4, fixed step 0.25 h       8.037602e-07              4.263256e-14
#> 4  Euler, fixed step 0.25 h       1.777395e-01              2.842171e-14
```

```
stopifnot(solver_check$passes[1:3],                          # adaptive and rk4 solutions with
          !solver_check$passes[4],                           # the check detects the inadequat
          all(solver_check$max_mass_balance_error_mg < 1e-8))  # mass is conserved by every
```

All four methods conserve mass, yet the Euler solution deviates by up to 17.8% from the analytical curve. A closed mass balance is necessary, not sufficient; only the comparison with an independent reference reveals the error. The adaptive solver **lsoda** (Soetaert et al. 2010; Petzold 1983) passes with both tolerance settings.

15.3.3. Step 3: structural identifiability at a single time

```
V_alt <- 26                                                  # a second virtual volume (L)
CL_alt <- uniroot(function(cl) conc_analytic(4, V = V_alt, CL = cl) - conc_analytic(4),
                  interval = c(0.1, 20))$root                # clearance that gives the same C at 4
c(CL_alt_l_h = CL_alt, C4_set1 = conc_analytic(4), C4_set2 = conc_analytic(4, V = V_alt, CL
  C05_set1 = conc_analytic(0.5), C05_set2 = conc_analytic(0.5, V = V_alt, CL = CL_alt))
```

```
#> CL_alt_l_h    C4_set1    C4_set2   C05_set1   C05_set2
#>   4.794632   1.839397   1.839397   4.412485   3.507380
```

```
stopifnot(abs(conc_analytic(4) - conc_analytic(4, V = V_alt, CL = CL_alt)) < 1e-8,
          abs(conc_analytic(0.5) / conc_analytic(0.5, V = V_alt, CL = CL_alt) - 1) > 0.1)  #
```

15.3.4. Step 4: expected precision of each design

For a log-normal error with standard deviation , the covariance of the estimates of (ln V, ln CL) is approximately ² (S S) ¹, where each row of S holds the sensitivities d ln C / d ln V = −1 + k · t and d ln C / d ln CL

$= -k \cdot t$ at one sampling time. The diagonal gives the squared standard errors on the log scale, which are approximately relative standard errors. A design whose matrix S has rank below 2 is not identifiable.

```
sens_log <- function(t, V = V_l, CL = CL_l_h) cbind(lnV = -1 + CL / V * t, ln
expected_se <- function(t, V = V_l, CL = CL_l_h, sigma = sigma_log) {
  S <- sens_log(t, V, CL)
  if (qr(S)$rank < 2) return(c(lnV = Inf, lnCL = Inf))  # structurally non-id
  sqrt(diag(sigma^2 * solve(crossprod(S))))
}
designs <- data.frame(label = c("A", "B", "C", "D"), t1 = c(0.5, 6, 0.5, 4),
                      description = c("early and late", "both late", "both ea
# Simulation check: 2,000 virtual repetitions of one person per design, two l
simulate_design <- function(t1, t2, n_sim = 2000, V = V_l, CL = CL_l_h) {
  if (t1 == t2) return(c(lnV = NA, lnCL = NA, share_invalid = NA))
  y1 <- log(conc_analytic(t1, V = V, CL = CL)) + rnorm(n_sim, 0, sigma_log)
  y2 <- log(conc_analytic(t2, V = V, CL = CL)) + rnorm(n_sim, 0, sigma_log)
  k_hat <- (y1 - y2) / (t2 - t1)                          # slope of the log-lin
  V_hat <- dose_mg / exp(y1 + k_hat * t1)                 # back-extrapolated in
  CL_hat <- k_hat * V_hat
  ok <- CL_hat > 0                                        # a negative slope giv
  c(lnV = sd(log(V_hat)), lnCL = sd(log(CL_hat[ok])), share_invalid = mean(!o
}
set.seed(15151)
expected <- t(mapply(function(a, b) expected_se(c(a, b)), designs$t1, designs
simulated <- t(mapply(simulate_design, designs$t1, designs$t2))
pct <- function(x) ifelse(is.finite(x), cdsr_number(100 * x, 0), "not identi
design_table <- data.frame(
  Design = paste(designs$label, designs$description, sep = ": "),
  `Sampling times (h)` = paste(designs$t1, "and", designs$t2),
  `V expected` = pct(expected[, "lnV"]), `V simulated` = pct(simulated[, "lnV
  `CL expected` = pct(expected[, "lnCL"]), `CL simulated` = pct(simulated[, "
  check.names = FALSE)
design_table
```

```
#>               Design Sampling times (h)       V expected      V simulated
#> 1  A: early and late          0.5 and 8               11               11
#> 2       B: both late            6 and 8               50               50
#> 3      C: both early          0.5 and 1               22               23
#> 4 D: same time twice            4 and 4 not identifiable not identifiable
#>        CL expected     CL simulated
#> 1                7                7
#> 2               22               21
#> 3               92               85
#> 4 not identifiable not identifiable
```

```
stopifnot(
  all(abs(simulated[1:2, c("lnV", "lnCL")] / expected[1:2, ] - 1) < 0.1),   # linearisation
  abs(simulated[3, "lnV"] / expected[3, "lnV"] - 1) < 0.1,                  # and for V in c
  all(is.infinite(expected[4, ])),                                          # design D is no
  expected[1, "lnV"] < expected[2, "lnV"], expected[1, "lnCL"] < expected[3, "lnCL"])
```

In design C, 19.6% of the virtual repetitions give a rising instead of a falling line, so no valid clearance can be computed; the simulated standard error of CL is then not meaningful beyond “very large”.

15.3.5. Step 5: the model-free baseline

```
set.seed(15152)
t_rich <- c(0.25, 0.5, 1, 2, 3, 4, 6, 8, 12)                    # nine samples (h)
c_rich <- conc_analytic(t_rich) * exp(rnorm(length(t_rich), 0, sigma_log))
terminal <- coef(lm(log(c_rich) ~ t_rich, subset = t_rich >= 4))   # log-linear terminal pha
k_terminal <- -unname(terminal[2])
c0_back <- exp(unname(coef(lm(log(c_rich[1:2]) ~ t_rich[1:2]))[1])) # back-extrapolated C at
```

```
auc_obs <- (c0_back + c_rich[1]) / 2 * t_rich[1] +
  sum(diff(t_rich) * (head(c_rich, -1) + tail(c_rich, -1)) / 2)    # trapezoi
auc_inf <- auc_obs + tail(c_rich, 1) / k_terminal                  # extrapol
CL_nca <- dose_mg / auc_inf
fit_model <- nls(log(c_rich) ~ log(dose_mg) - lnV - exp(lnCL - lnV) * t_rich
                 start = list(lnV = log(15), lnCL = log(4)))
model_est <- exp(coef(fit_model))
c(CL_nca = CL_nca, CL_model = unname(model_est["lnCL"]), V_model = unname(mod
  share_auc_extrapolated = (auc_inf - auc_obs) / auc_inf)
```

```
#>                 CL_nca               CL_model                V_model
#>             5.36968529             5.48494089            22.02396286
#> share_auc_extrapolated
#>             0.04612253
```

```
stopifnot(abs(CL_nca / CL_l_h - 1) < 0.15, abs(model_est[["lnCL"]] / CL_l_h -
```

With nine samples, the baseline and the model agree with the virtual clearance within the error to be expected from a single person (NCA 5.37 L/h, model 5.48 L/h, virtual value 5 L/h). The baseline needs no structural assumption, but it needs rich sampling; the two-sample designs of the case need the model.

15.3.6. Step 6: interindividual variability and the two-stage shortcut

Virtual persons now differ: ln CL and ln V vary with a standard deviation of 0.3 (*assumption*). A two-stage analysis estimates each person's parameters from their two samples and then summarises the spread. The spread of the individual estimates mixes true variability and estimation error.

```
set.seed(15153)
n_persons <- 200
omega <- 0.3                                                    # between-person SD on
persons <- data.frame(V = V_l * exp(rnorm(n_persons, 0, omega)), CL = CL_l_h * exp(rnorm(n_p
two_stage_sd <- function(t1, t2) {
  y1 <- log(conc_analytic(t1, V = persons$V, CL = persons$CL)) + rnorm(n_persons, 0, sigma_l
  y2 <- log(conc_analytic(t2, V = persons$V, CL = persons$CL)) + rnorm(n_persons, 0, sigma_l
  k_hat <- (y1 - y2) / (t2 - t1)
  CL_hat <- k_hat * dose_mg / exp(y1 + k_hat * t1)
  ok <- CL_hat > 0                                               # rising lines give no
  c(sd_lnCL = sd(log(CL_hat[ok])), true_sd_lnCL = sd(log(persons$CL)), share_invalid = mean(
}
spread <- rbind(A = two_stage_sd(0.5, 8), B = two_stage_sd(6, 8))
round(spread, 3)
```

```
#>   sd_lnCL true_sd_lnCL share_invalid
#> A   0.320        0.301             0
#> B   0.417        0.301             0
```

```
stopifnot(spread["B", "sd_lnCL"] > spread["A", "sd_lnCL"])      # imprecise design infla
```

With the late design B, the apparent between-person spread of the clearance is inflated more than with design A. A population model estimates variability and error jointly and avoids this bias; in R it can be fitted with `nlme::nlme()` or with dedicated pharmacometric software. Such models are the next step and are outside this chapter.

> **i** Other implementation: Julia/SciML and SciPy
>
> **Julia/SciML.** DifferentialEquations.jl (Rackauckas and Nie 2017) (version 8.1.1, August 2026; documentation) offers a large set of solvers behind one interface. Its documentation recommends `Tsit5`

for most non-stiff problems, Verner methods (`Vern7`, `Vern8`) when high accuracy is needed, and Rosenbrock methods such as `Rodas5P` for stiff problems; the Fortran `lsoda` used by deSolve is available through a wrapper. The concrete advantage for this chapter lies in sensitivities: SciMLSensitivity.jl (documentation) computes forward and adjoint sensitivities of the solution with respect to the parameters by automatic differentiation. For the one-compartment model, Step 4 used the closed-form sensitivities $-1 + k \cdot t$ and $-k \cdot t$; for models without a closed-form solution, the same Fisher-information argument needs numerical sensitivities, and this is where the ecosystem helps.

SciPy. `scipy.integrate.solve_ivp` (SciPy 1.18.1, August 2026; documentation) (Virtanen et al. 2020) connects the model to existing Python data pipelines. Its default method is `RK45` with rtol $= 10^{-3}$ and atol $= 10^{-6}$, looser than the deSolve defaults; `method="LSODA"` wraps the same ODEPACK solver family.

For a fair comparison, all implementations need the same parameters and units (mg, L, h), the same initial state, the same output times (`t_eval` in SciPy, `saveat` in Julia), the same reference (the analytical solution) and the same pre-specified tolerance (relative deviation 10^{-4}). Solver name and tolerances must be reported. Statements about speed need a benchmark on the same machine. Status: conceptual comparison, not run.

If the result stays unclear. If a design turns out not to separate CL and V – for example because samples were taken later than planned – report that the volume is not identifiable from these data. Do not fix it silently. The options are to add an early sample in the next participants (if the protocol allows), to fix V at a published value and label it as an *assumption* with its source, or to borrow strength across participants in a population model. Each option changes what the result means and belongs in the report.

15.4. Counter-check

Change one virtual parameter: the clearance is halved to 2.5 L/h, for example in people with reduced elimination; the volume and the designs stay the same. **Predict first:** does the expected precision of the clearance improve or worsen for the late design B? And for the volume?

```
expected_low_cl <- t(mapply(function(a, b) expected_se(c(a, b), CL = CL_l_h / 2), designs$t1
counter <- data.frame(Design = design_table$Design,
                      `CL expected, CL = 5` = pct(expected[, "lnCL"]),
                      `CL expected, CL = 2.5` = pct(expected_low_cl[, "lnCL"]),
                      `V expected, CL = 5` = pct(expected[, "lnV"]),
                      `V expected, CL = 2.5` = pct(expected_low_cl[, "lnV"]),
                      check.names = FALSE)
counter
```

```
#>               Design CL expected, CL = 5 CL expected, CL = 2.5
#> 1  A: early and late                   7                    10
#> 2       B: both late                  22                    10
#> 3      C: both early                  92                   205
#> 4 D: same time twice    not identifiable      not identifiable
#>   V expected, CL = 5 V expected, CL = 2.5
#> 1                 11                   11
#> 2                 50                   50
#> 3                 22                   22
#> 4   not identifiable     not identifiable
```

```
stopifnot(expected_low_cl[2, "lnCL"] < expected[2, "lnCL"],                    # design B im
          isTRUE(all.equal(expected_low_cl[1:3, "lnV"], expected[1:3, "lnV"])), # V precisio
          expected_low_cl[3, "lnCL"] > expected[3, "lnCL"])                     # design C w
```

With the lower clearance, the late design B estimates the clearance more precisely (22% → 10%), because the samples now lie closer to the time

V / CL = 8 h at which a sample is informative about CL only. The early design C gets worse. The precision of the volume does not change at all: on the log scale, the intercept's uncertainty depends only on the sampling times. The practical lesson is that a sampling schedule tuned to one assumed clearance may be poor for patients who differ from it. Protocols for heterogeneous populations therefore spread sampling times or allow windows.

15.5. Consequences for patients and operations

Patients and ethics. Every sample is a needle, a volume of blood and often a disturbed night; in critically ill or young patients the volume matters. A design that is identifiable with two samples is a direct benefit, and a design that wastes samples on times that carry little information is a harm without return. The protocol should state why each sampling time is needed.

Regulation. A model fitted for a research question is not a product. Software that calculates an individual dose from measured concentrations for use in patient care has a medical purpose, and its classification has to be examined (Chapter 22). The virtual parameters here must never be used for such a calculation.

Data protection and security. Concentration–time data with dosing times, body weight and laboratory values are health data. Sharing them for re-analysis needs pseudonymisation and controlled access (Chapter 23, Chapter 26). The model code, solver settings and the tolerance check belong under version control, so that a later re-analysis uses the same numerical method (Chapter 4).

15.6. Look up and transfer

Reviewer question. Were the sampling times chosen so that each reported parameter is identifiable, and do the reported standard errors account for the design – or were some parameters fixed without saying so?

Self-test. A virtual compound is given as an intravenous bolus, and only the AUC (from many samples) is reported for each person. Which of CL and V can be estimated, and why?

Answer

Only the clearance: AUC = dose / CL, so CL = dose / AUC. The volume does not appear in the AUC at all; parameter sets with the same CL and any V give the same AUC. V is structurally non-identifiable from the AUC alone; it needs the shape of the curve, for example an early concentration or the terminal slope.

Annotated sources.

- Soetaert et al. (2010) – the deSolve package and its solvers, including `lsoda`.
- Petzold (1983) – the automatic switching between stiff and non-stiff methods behind `lsoda`.
- Bellman and Åström (1970) – the original definition of structural identifiability.
- Raue et al. (2009) – structural and practical identifiability of partially observed dynamic models.
- Mould and Upton (2012) and Mould and Upton (2013) – introductions to population modelling and to pharmacokinetic models.
- Sale and Sherer (2015) – a genetic algorithm for automated selection of population pharmacokinetic models.
- Rackauckas and Nie (2017) and Virtanen et al. (2020) – the Julia and Python implementations discussed in the box.

15.7. Further reading

- Chapter 16 – how little data constrain a parameter, from a Bayesian angle.
- Chapter 18 – simulation with explicit assumptions for a different question: capacity rather than concentration.
- Methods in R: Fisher Information – the quantity behind the expected standard errors of Step 4.
- Methods in R: Non-linear regression with nls (lab) – fitting curves like the one in Step 5, with diagnostics.
- CTTIR Course 3, Week 4, Session 4 – SIR and SEIR models with deSolve – another mechanistic model solved with the same package.
- CTTIR Tutorials: The Delta Method – why log-scale standard errors approximate relative errors.
- Strategy in R: law of large numbers and Monte Carlo – the simulation logic used to check the linearised standard errors.

16. Bayes, small samples and the 3Rs

R. Heller · Date: 2025-06-15 · Last changed: 2026-11-09

At a glance

Guiding question: What uncertainty remains even when an event rate looks striking?
For: preclinical researchers, doctoral students, members of animal welfare bodies and reviewers of small experiments · **Prerequisites:** proportions and denominators (Chapter 3), the difference between describing and explaining (Chapter 2)
In the course: W5 · Causality, trajectories and phenotypes and W6 · Pharmacometrics, forecasting and simulation · **Data:** synthetic, fictitious animals (generated in the chapter) · **Core reading time:** about 10 minutes

16.1. Case and decision

A fictitious preclinical group has tested a new bone substitute material in a critical-size femoral defect model in rats. Ten animals were operated on; one died on the second day after surgery; in seven of the nine animals evaluated at week 8, the defect was bridged by bone. There was no control group. The draft abstract reads: “The material led to bony bridging in 78% of animals and is therefore effective.” The group leader has to decide

two things before the manuscript goes out and before the next application to the competent authority: what the abstract may claim, and which experiment – if any – should come next. All animals and results are *synthetic*. The chapter separates three questions that the draft merges: how large the bridging rate is, whether the material caused it, and whether another experiment on animals is justified.

16.2. Core reading

Core reading · recommended for W5 and W6 · about 10 minutes · ends before 'Lab'

Origin and denominator first. The independent unit is the animal (Lazic et al. 2018). Ten were operated on; one died before the end point; nine were evaluated; seven were bridged. The draft reports $7/9 = 77.8\%$. If the animal that died is counted as a failure, the rate is $7/10 = 70.0\%$. Both numbers are correct; they answer different questions, and the report must say which one it uses and why. A short table that lists all ten animals with their status is the clearest way to show both. Reporting the excluded animal is not optional: the ARRIVE guidelines list inclusion and exclusion criteria among the "Essential 10" items (Percie du Sert et al. 2020).

A frequentist baseline. The exact (Clopper–Pearson) 95% confidence interval for 7 of 9 runs from 40.0% to 97.2% (Clopper and Pearson 1934). With nine animals, "78%" is compatible with a rate below one half and with a rate close to one.

Prior, likelihood and posterior. Bayesian analysis asks a slightly different question: given the data, how plausible is each possible bridging rate p? It combines three parts. The *likelihood* says how probable 7 bridged out of 9 is for each value of p (a binomial model). The *prior* states what was considered plausible before the experiment. The *posterior* is the updated belief: prior × likelihood, rescaled. For a binomial outcome and a

Beta prior, the posterior is again a Beta distribution: Beta(a, b) becomes Beta(a + 7, b + 2). Its quantiles come from `qbeta()` without simulation (Spiegelhalter et al. 2004). The prior parameters can be read as imaginary animals observed beforehand: Beta(1, 1) adds one bridged and one unbridged animal, Beta(3, 7) adds three and seven. With only nine real animals, ten imaginary ones carry real weight.

Two priors are plausible here (both *assumptions*). A *flat* prior, Beta(1, 1), treats every rate between 0 and 1 as equally plausible. A *sceptical* prior, Beta(3, 7), expresses the expectation that the material performs about as well as an untreated defect in this model, assumed here to bridge in about 30% of animals, and gives this expectation the weight of ten animals. Figure 16.1 and Table 16.1 show the result.

Table 16.1.: Posterior summaries of the bridging rate for 7 of 9 animals (*synthetic*) under two priors (*assumptions*), computed analytically with `qbeta()` and `pbeta()`.

Prior	Rationale	Posterior	Mean and 95% interval	P(rate > 50%)
flat Beta(1, 1)	every rate equally plausible	Beta(8, 3)	72.7% (44.4% to 93.3%)	0.95
sceptical Beta(3, 7)	about 30%, as assumed for untreated defects; weight of 10 animals	Beta(10, 9)	52.6% (30.8% to 74.0%)	0.59

With the flat prior, the posterior mean is 72.7%, the 95% interval runs from 44.4% to 93.3%, and the probability that the rate exceeds one half is 0.95. With the sceptical prior, the mean falls to 52.6% and that probability to 0.59. Nine animals are not enough to make the prior irrelevant. Reporting both analyses, with the prior chosen before the data were seen, is more

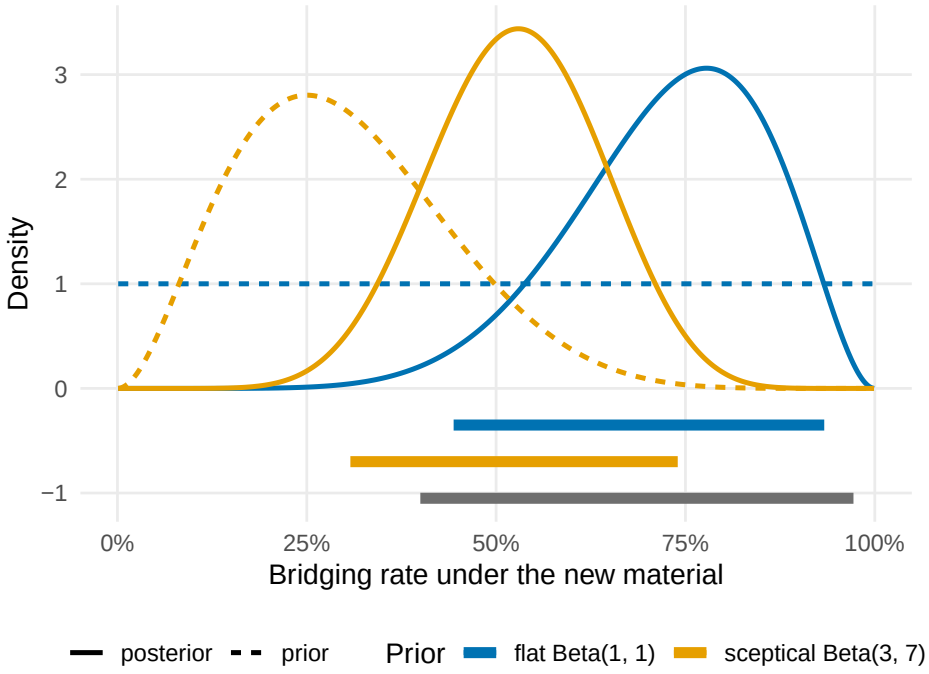


Figure 16.1.: Prior (dashed) and posterior (solid) densities of the bridging rate for two priors, after 7 bridged of 9 evaluated animals (*synthetic*, n = 9 animals, 10 operated). Horizontal bars: 95% posterior intervals; grey bar: exact 95% confidence interval.

informative than reporting either alone. Note the difference in meaning: a posterior interval is a statement about the rate, given the model and the prior; the confidence interval is a statement about how often the procedure would cover the true rate in repeated experiments. With small samples the two can differ noticeably.

The most important assumption. Every interval in this chapter describes the bridging rate under the new material in this model and this laboratory. None of them says what the rate would have been without the material. An effect is a comparison: rate with the material minus rate without it. The second rate was not observed. If an untreated defect bridged in 20% of animals, the posterior probability that the material does better is above 0.999; at 50%, it is 0.95; at 70%, it is 0.62 (flat prior). The data cannot choose between these scenarios, so the claim "is effective" has no support from this experiment. A historical rate from another publication is no substitute for a concurrent, randomised control: animals, surgeons, housing and assessment differ.

What the next experiment needs. A defensible next step answers the question that this experiment could not: does the material change the bridging rate? That requires a concurrent control group in the same model and laboratory, random allocation of animals to groups, outcome assessment by a person who does not know the group, an end point and a sample size fixed before the first operation, and a report of every animal. It also requires a criterion, agreed in advance, under which the experiment is *not* run – for example when the number of animals that can be justified cannot answer the question with reasonable probability. The lab computes both.

Planning versus reporting. Two frameworks apply, at different moments. The 3Rs – replacement, reduction and refinement – govern whether and how an experiment is done (NC3Rs 2026). In the European Union they are binding on Member States through Article 4 of Directive 2010/63/EU on the protection of animals used for scientific purposes (European Parliament and Council of the European Union 2010). The ARRIVE guidelines

2.0 govern how a completed study is reported (Percie du Sert et al. 2020). Good reporting cannot repair a design without a control group, and a well-planned design is of little use if the report omits the excluded animal.

Question for the workshop. Write one sentence for the abstract that is true for these data and that a sceptical reader would accept. Which word in the draft sentence has to go first?

16.3. Lab

The lab uses base R only (plus **ggplot2** for the figure above) and runs in under a second.

16.3.1. Step 1: the data with their origin

```
animals <- data.frame(
  animal = sprintf("R%02d", 1:10),                                    # fictiti
  died_before_end_point = c(FALSE, TRUE, rep(FALSE, 8)),              # R02 die
  bridged_week8 = c(TRUE, NA, TRUE, TRUE, FALSE, TRUE, TRUE, TRUE, FALSE, TRU
n_operated <- nrow(animals)
n_evaluated <- sum(!animals$died_before_end_point)
x_bridged <- sum(animals$bridged_week8, na.rm = TRUE)
stopifnot(n_operated == 10, n_evaluated == 9, x_bridged == 7,
          all(is.na(animals$bridged_week8) == animals$died_before_end_point))
```

16.3.2. Step 2: the frequentist baseline

```
cp_interval <- binom.test(x_bridged, n_evaluated)$conf.int           # exact Clopper-Pearso
cp_interval_operated <- binom.test(x_bridged, n_operated)$conf.int    # death counted as fai
round(rbind(evaluated_7_of_9 = cp_interval, operated_7_of_10 = cp_interval_operated), 3)
```

```
#>                    [,1]  [,2]
#> evaluated_7_of_9  0.400 0.972
#> operated_7_of_10 0.348 0.933
```

16.3.3. Step 3: two priors and their posteriors

```
priors <- data.frame(label = c("flat Beta(1, 1)", "sceptical Beta(3, 7)"), a = c(1, 3), b =
                     rationale = c("every rate equally plausible",
                                   "about 30%, as assumed for untreated defects; weight of 1
posterior_of <- function(x, n, a, b) {                    # conjugate update Beta(a, b) ->
  a_post <- a + x
  b_post <- b + n - x
  data.frame(a_post = a_post, b_post = b_post, mean = a_post / (a_post + b_post),
             lower = qbeta(0.025, a_post, b_post), upper = qbeta(0.975, a_post, b_post),
             p_above_half = 1 - pbeta(0.5, a_post, b_post))
}
posterior <- do.call(rbind, Map(posterior_of, x_bridged, n_evaluated, priors$a, priors$b))
prior_interval <- qbeta(c(0.025, 0.975), 3, 7)            # what the sceptical prior allowed
posterior_table <- data.frame(
  Prior = priors$label, Rationale = priors$rationale,
  Posterior = sprintf("Beta(%d, %d)", posterior$a_post, posterior$b_post),
  `Mean and 95% interval` = sprintf("%s (%s to %s)", cdsr_percent(posterior$mean), cdsr_perc
                                    cdsr_percent(posterior$upper)),
  `P(rate > 50%)` = cdsr_number(posterior$p_above_half, 2), check.names = FALSE)
posterior[, -(1:2)] |> round(3)
```

```
#>    mean lower upper p_above_half
```

```
#> 1 0.727 0.444 0.933        0.945
#> 2 0.526 0.308 0.740        0.593
```

```
stopifnot(
  posterior$a_post == c(8, 10), posterior$b_post == c(3, 9),
  abs(posterior$lower - c(0.444, 0.308)) < 0.001, abs(posterior$upper - c(0.9
  posterior$upper[2] < posterior$upper[1],                # the sceptical pr
  cp_interval[2] - cp_interval[1] > posterior$upper[1] - posterior$lower[1])
```

The sceptical prior on its own allowed rates from 7.5% to 60.0%. Its posterior is a compromise between this expectation and the data; with nine animals, neither dominates.

16.3.4. Step 4: event rate is not effect

```
control_rates <- c(0.2, 0.5, 0.7)                          # hypothetical rat
p_beats_control <- 1 - pbeta(control_rates, posterior$a_post[1], posterior$b_
data.frame(assumed_control_rate = control_rates, P_material_better = round(p_
```

```
#>   assumed_control_rate P_material_better
#> 1                  0.2             1.000
#> 2                  0.5             0.945
#> 3                  0.7             0.617
```

```
stopifnot(p_beats_control[1] > 0.999, p_beats_control[3] < 0.7)   # the answe
```

16.3.5. Step 5: design of the next experiment

The next defensible experiment, if any, compares the material with a concurrent control in a randomised design, with the outcome assessed blind

to group. Its size follows from a design analysis before any animal is ordered. Here the decision rule is: call the material better if the posterior probability that its rate exceeds the control rate is above 0.95 (flat priors in both arms). For each number of animals per arm, the probability that the rule declares success can be computed exactly by summing over all possible outcomes – no simulation is needed.

```
prob_better <- function(x_t, x_c, n) {                    # P(p_material > p_control | data
  integrate(function(p) dbeta(p, 1 + x_c, 1 + n - x_c) * (1 - pbeta(p, 1 + x_t, 1 + n - x_t)
}
decision_probability <- function(n, p_material, p_control, threshold = 0.95) {
  outcomes <- expand.grid(x_t = 0:n, x_c = 0:n)
  success <- mapply(prob_better, outcomes$x_t, outcomes$x_c, MoreArgs = list(n = n)) > thres
  sum(dbinom(outcomes$x_t, n, p_material) * dbinom(outcomes$x_c, n, p_control) * success)
}
n_per_arm <- c(6, 8, 10, 12, 15, 20)
design <- data.frame(
  n_per_arm = n_per_arm,
  success_if_relevant_difference = sapply(n_per_arm, decision_probability, p_material = 0.7,
  success_if_no_difference = sapply(n_per_arm, decision_probability, p_material = 0.3, p_con
round(design, 3)
```

```
#>   n_per_arm success_if_relevant_difference success_if_no_difference
#> 1         6                          0.296                    0.034
#> 2         8                          0.455                    0.041
#> 3        10                          0.609                    0.050
#> 4        12                          0.596                    0.053
#> 5        15                          0.731                    0.043
#> 6        20                          0.808                    0.044
```

```
n_needed <- min(design$n_per_arm[design$success_if_relevant_difference >= 0.8])
stopifnot(all(design$success_if_no_difference < 0.06),    # false-positive probability stays
          n_needed == 20)
```

The assumed relevant difference is 30% versus 70% (*assumption*). To reach a decision probability of 80%, the binary end point needs 20 animals per arm. The probabilities do not rise smoothly with the number of animals because the outcome is discrete.

A criterion against a further experiment. Fix it before the application, together with the animal welfare body. For example: *if the largest number of animals that the team considers justifiable for this question – say 12 per arm (assumption) – gives a decision probability below 80% for the smallest relevant difference, the experiment is not run in this form.* Here, 12 per arm give 59.6%, so the criterion applies. The alternatives are not "more animals" but better information per animal and fewer animals overall: a continuous end point such as bone volume from imaging instead of bridged yes/no (reduction), non-animal tests of the material's properties before any further surgery (replacement, partial), and a revised anaesthesia and analgesia protocol after the early death (refinement).

16.3.6. Step 6: planning and reporting are different tasks

Table 16.2 sets the two frameworks side by side for this case.

Table 16.2.: Planning under the 3Rs and reporting under ARRIVE 2.0.

	3Rs	ARRIVE 2.0
Question	Is the experiment necessary, and how can harm and numbers be minimised?	Can readers judge and repeat what was done?
When	Before the experiment: design, application, ethical review	After the experiment: manuscript

	3Rs	ARRIVE 2.0
Basis	NC3Rs definitions (NC3Rs 2026); Directive 2010/63/EU, Article 4 (European Parliament and Council of the European Union 2010)	Essential 10 and Recommended Set (Percie du Sert et al. 2020)
In this case	Concurrent control, design analysis, continuous end point, refined anaesthesia	Report all 10 animals, the death and its handling, the absent control and blinding

The ARRIVE "Essential 10" are study design, sample size, inclusion and exclusion criteria, randomisation, blinding, outcome measures, statistical methods, experimental animals, experimental procedures and results. In a survey of 271 publications of animal research, most studies did not use randomisation (87%) or blinding (86%), and only 59% stated the objective together with the number and characteristics of the animals (*published*, Kilkenny et al. (2009)).

16.3.7. Step 7: decision exercise – a letter to the editor

Assume the draft sentence was published unchanged. Write a letter to the editor of about 150 words that explains, for a clinical readership, the difference between the observed bridging rate and a treatment effect. Use the denominators, one interval and one consequence for the next study. Compare your letter with the example only after writing it.

Example letter

The authors report bony bridging in 78% of rats treated with a new bone substitute and conclude that the material is effective. We would like to raise two points. First, ten animals were operated on and one died; the rate is 7/9, or 7/10 if the death is counted. The exact 95% confidence interval for 7/9 runs from 40% to 97%. Second, and more important, there was no control group. Whether 78% is better than what an untreated defect achieves in this model, operated and assessed in the same laboratory, is unknown; historical rates from other studies cannot answer this. The data therefore describe how often bridging occurred, not whether the material caused it. A randomised experiment with a concurrent control, blinded assessment and a sample size justified in advance would be needed before any claim of efficacy, and before planning studies in patients.

If the result stays unclear. A small single-arm experiment with an unclear result is still worth reporting – with all animals, both denominators and the absent control stated plainly – because unreported experiments lead others to repeat them. Its legitimate use is to inform the design of the next experiment: it suggests a plausible range for the treated rate, as in the sceptical-prior analysis, not a proof of effect.

16.4. Counter-check

Triple the number of animals and keep the rate: 21 bridged of 27 evaluated (*synthetic*). **Predict first:** will the two posteriors move closer together? Will the uncertainty about the *effect* of the material shrink?

```
posterior_27 <- do.call(rbind, Map(posterior_of, 21, 27, priors$a, priors$b))
widths <- data.frame(prior = priors$label,
                     width_9 = posterior$upper - posterior$lower,
```

```
                     width_27 = posterior_27$upper - posterior_27$lower,
                     mean_9 = posterior$mean, mean_27 = posterior_27$mean)
round(widths[, -1], 3)
```

```
#>   width_9 width_27 mean_9 mean_27
#> 1   0.489    0.303  0.727   0.759
#> 2   0.432    0.301  0.526   0.649
```

```
gap_9 <- diff(posterior$mean); gap_27 <- diff(posterior_27$mean)
stopifnot(all(widths$width_27 < widths$width_9),           # intervals narrow
          abs(gap_27) < abs(gap_9))                        # the priors matter less
```

The intervals narrow, and the difference between the posterior means under the two priors shrinks from 20.1 to 11.0 percentage points. The uncertainty about the *rate* falls. The uncertainty about the *effect* does not change at all, because 27 treated animals without a control still contain no information about the rate without treatment. Tripling the number of animals in an uncontrolled design costs animals without answering the question – the opposite of reduction in the sense of the NC3Rs, which ties reduction to appropriately designed experiments that add to the knowledge base (NC3Rs 2026).

16.5. Consequences for patients and operations

Animals and ethics. Each further animal needs a justification that this design cannot supply. The early death is a signal for refinement before any further surgery, and the criterion against a further experiment protects animals from a study that could not answer its question. In Germany, the Animal Welfare Act (Tierschutzgesetz, TierSchG) § 7a requires that animal experiments are indispensable for their purpose and, for vertebrates

and cephalopods, ethically justifiable with respect to the expected pain, suffering or harm (Federal Republic of Germany 2006).

Law: what binds whom (as of November 2026). Directive 2010/63/EU was adopted on 22 September 2010, published in the Official Journal on 20 October 2010 and entered into force on 9 November 2010; Member States had to transpose it by 10 November 2012 and apply their rules from 1 January 2013 (European Parliament and Council of the European Union 2010). Article 4 (replacement, reduction, refinement) addresses Member States and reaches researchers through national law; later amendments (Regulation (EU) 2019/1010 and Delegated Directive (EU) 2024/1262) changed reporting rules and annexes, not Article 4. The NC3Rs definitions of the 3Rs and the ARRIVE guidelines are guidance, not law, although journals and funders may make ARRIVE reporting a condition. Whether the next experiment needs a new project authorisation or an amendment, and which authority decides, is a local question for the animal welfare officer and the competent authority (*open*).

Patients. A claim of efficacy from an uncontrolled animal series can travel into grant applications and early clinical plans. If the material were developed for patients, its preclinical evidence would become part of the evidence assessed under medical device rules (Chapter 22, interpretation); a misleading abstract weakens that chain from the start.

Data and security. Indirect: the animal-level data, including the excluded animal, should be stored with the protocol so that the denominators can be reconstructed (Chapter 4, Chapter 27).

16.6. Look up and transfer

Reviewer question. Which animals were operated, excluded and evaluated, was there a concurrent randomised control group, and how sensitive are the conclusions to the prior?

Self-test. In a second fictitious model, 0 of 8 treated animals develop a wound infection. A colleague writes "the material prevents infection (0%)". Compute a 95% posterior interval under a flat prior and explain what the result can and cannot support.

Answer

The posterior is Beta(1, 9). Its 95% interval is `qbeta(c(0.025, 0.975), 1, 9)`, about 0.3% to 33.6%; with the one-sided "upper 95% bound", `qbeta(0.95, 1, 9)` gives about 28%. Zero events in eight animals are compatible with an infection rate of more than one in four. And as in the case, without a control group the data say nothing about prevention: infection might be rare in this model anyway.

Annotated sources.

- Spiegelhalter et al. (2004) – Bayesian reasoning for trials and health-care evaluation, including the choice and reporting of priors.
- Clopper and Pearson (1934) – the exact binomial confidence interval used as the frequentist baseline.
- Lazic et al. (2018) – what counts as the independent unit ("N") in animal experiments.
- Percie du Sert et al. (2020) – the ARRIVE guidelines 2.0 for reporting animal research.
- NC3Rs (2026) – current definitions of replacement, reduction and refinement.
- European Parliament and Council of the European Union (2010) – the binding EU framework, including Article 4 on the 3Rs.
- Federal Republic of Germany (2006) – the German provision on the indispensability and ethical justifiability of animal experiments.
- Kilkenny et al. (2009) – evidence of reporting and design gaps in published animal research.

16.7. Further reading

- Chapter 19 – what changes when a result is transferred from one species or centre to the next.
- Chapter 27 – reporting templates for animal studies within the model passport.
- Methods in R: The Beta-Binomial Model and Prior Specification – the conjugate model of this chapter and how to choose and justify priors.
- Methods in R: ARRIVE – animal research – a reporting template for the items discussed here.
- CTTIR Course 3, Week 1, Session 4 – Bench and translational design – designing preclinical experiments with controls and randomisation.
- CTTIR Course 4, Week 3, Session 1 – Bayesian thinking; control vs decision error – decision rules based on posterior probabilities.
- CTTIR Tutorials: Credible Intervals – equal-tailed and highest-density intervals.

Part V.

Part V · Care systems and translation into practice

17. Forecasting patient flow

R. Heller · Date: 2025-07-13 · Last changed: 2026-11-09

At a glance

Guiding question: Is the forecast interval of an emergency department usable for next week?
For: emergency department and nursing leads who plan rosters, analysts in hospital operations, clinicians who read forecast dashboards · **Prerequisites:** counts, denominators and how routine data arise (Chapter 3), the idea of calibration (Chapter 6); a first look at time series in Methods in R: Time Series: Introduction
In the course: W6 · Pharmacometrics, forecasting and simulation (required reading); transfer exercise in W3 · Diagnostics, prediction and external evidence · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

17.1. Case and decision

On Thursday afternoon the nurse manager of a fictitious emergency department plans the roster for the coming week. A dashboard shows the expected number of arrivals for every hour, with a band labelled "90% interval". The manager has one standby nurse to place. The tempting rule is to staff each shift to the upper edge of the band and to put the standby nurse where the band is widest. Before following that rule, the manager asks the analyst a plain question: *if the band says 90%, did it*

in fact contain the arrivals in nine out of ten hours in past weeks, and in which shifts did it fail? The answer decides where the standby nurse goes. If the band is too narrow in weekend late shifts, the cost of the error falls on patients who wait longer and on a late-shift team that is already stretched.

17.2. Core reading

Core reading · required for W6 · about 10 minutes · ends before 'Lab'

Two different questions. A forecast and a simulation are often shown on the same planning screen, but they answer different questions (Table 17.1). A forecast says how many patients are likely to arrive if the process continues as before. It can be checked: wait a week and compare. A simulated capacity change, the topic of Chapter 18, says what would happen to waiting times if the department changed something, given an assumed stream of arrivals. It cannot be checked against data before the change is made; it is checked by verification and plausibility instead. A forecast is often the input of such a simulation, which is one more reason to know how good it is.

Table 17.1.: What a forecast and a simulated capacity change answer.

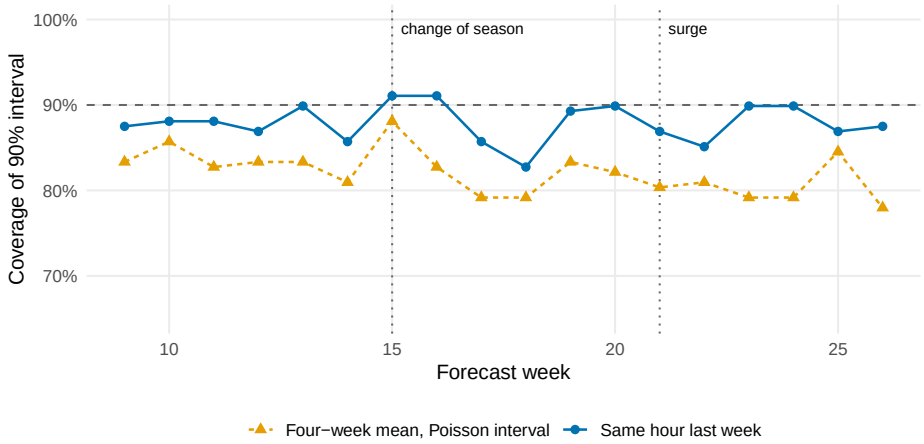
	Forecast (this chapter)	Simulated capacity change (Chapter 18)
Question	How many arrivals next week, if the process continues?	What happens to waiting if we add or remove a resource?
Changes the system?	No	Yes, hypothetically

	Forecast (this chapter)	Simulated capacity change (Chapter 18)
How it is checked	Against later data (backtest, then prospectively)	Verification of the model logic, plausibility against external figures
Typical failure	Interval too narrow, process changed	Wrong assumptions about service times or priorities

The data and the honest benchmark. The lab generates 26 weeks of hourly arrivals at a fictitious department with about 92 arrivals per day. The series has a daily rhythm (quiet nights, busy late mornings), a weekly rhythm (busier Mondays), a change of season from week 15 (15% more arrivals, for example from leisure injuries) and one unusual surge: in week 21, arrivals double on the afternoons and evenings of Friday to Sunday. Every number is *synthetic*. Any forecasting method has to beat a simple reference that a planner could compute by hand: the arrivals in the *same hour of the previous week*. This seasonal naive forecast (Hyndman and Athanasopoulos 2021) already contains the daily and weekly rhythm. A more elaborate model that does not beat it in later weeks adds complexity without benefit.

Testing in later windows only. A forecast for next week may use only data that existed before next week. The lab therefore uses a rolling-origin backtest (Tashman 2000): the forecast origin moves forward one week at a time, from week 9 to week 26, and each time the next 168 hours are forecast from the data before the origin. This gives 18 forecast weeks with 3,024 forecast hours. The independent denominator is closer to 18 than to 3,024, because all hours of one forecast week share the same origin and the same conditions. The general method is described in Methods in R: Rolling-Origin Cross-Validation.

Two forecasts, two intervals. The reference forecast uses last week's value with a normal-approximation interval from the `forecast` package (Hyndman and Khandakar 2008). The second forecast uses the mean of the same hour over the last four weeks with an interval taken from a Poisson distribution, an assumption that is common for counts. A 90% interval should contain the observed arrivals in about 90% of hours (*calibration*, here measured as coverage) and should be as narrow as possible given that (*sharpness*) (Gneiting et al. 2007).



Synthetic emergency department, hourly arrivals; rolling origin, one forecast per week.

Figure 17.1.: Coverage of nominal 90% forecast intervals per forecast week (share of 168 hours in which arrivals fell inside the interval). n = 18 forecast weeks (origins in weeks 9 to 26), 3,024 forecast hours per method. Dashed line: nominal 90%. Data status: synthetic.

Table 17.2.: Backtest summary over 18 forecast weeks (3,024 hours per method). MAE: mean absolute error in arrivals per hour. Coverage: share of hours inside the nominal 90% interval, overall and by shift (early 06–14, late 14–22, night 22–06). Cost: weighted cost points per week (weights are an assumption, see Lab). Data status: synthetic.

Method	MAE	Coverage	Early	Late	Night	Mean width	Cost per week
Same hour last week	2.85	0.88	0.83	0.82	0.99	9.3	302
Four-week mean, Poisson interval	2.28	0.82	0.79	0.76	0.91	5.9	245

What the backtest shows. The four-week mean has the smaller point error (MAE 2.28 versus 2.85 arrivals per hour), but its Poisson interval contains the arrivals in only 82.0% of hours, and in the late shift in only 76.4%. The arrivals vary more than a Poisson model assumes. The reference interval reaches 87.9% overall, which looks close to 90%, but the average hides a split: 99.1% at night, where the band is needlessly wide, and 81.6% in the late shift, where it is too narrow. Point accuracy and interval calibration are separate properties, and calibration has to be checked where the decision is taken, here per shift. In the four weeks after the change of season both forecasts under-estimate arrivals, the four-week mean more so (by 0.48 versus 0.26 arrivals per hour on average), because most of its reference weeks still come from the old level. During the surge both intervals fail in the weekend late shifts.

The most important assumption. A forecast interval describes uncertainty under the assumption that next week's arrivals come from the same process as the training weeks. It says nothing about events of a kind that the training window did not contain. The surge in week 21 is such an event. No interval built from past weeks could have anticipated it; that needs a separate signal, for example an

event calendar or a heat warning, and a rule for what the department does when the signal fires.

Errors do not cost the same at every hour. One patient more than planned at 3 a.m., with a thin night team, has other consequences than one patient more at 11 a.m. The lab therefore weights errors by shift. The weights are an *assumption* that the department has to agree on; with other weights the ranking of the two methods can change. The point of writing them down is that the choice of method becomes a discussable decision instead of a by-product of a software default.

Question for the workshop: Your group plans the standby nurse for next week from Table 17.2. Which shift would you protect, which interval would you use, and what would you tell the late-shift team about the 90% label?

17.3. Lab

The lab follows the order of a real backtest: generate (or import) the series, fix the forecast origins, forecast only from data before each origin, and evaluate per week and per shift.

> **R start command**
>
> From the root of the book repository, render the chapter (this runs all code, including the figures of the core reading):
>
> ```
> quarto render chapters/b17-forecasting-patient-flow.qmd --to html
> ```
>
> To run only the Lab code in a fresh R session, start from the `chapters/` folder:
>
> ```
> Rscript -e 'source(knitr::purl("b17-forecasting-patient-flow.qmd", output =
> ```

The chapter writes the synthetic hourly series to `data/b17-ed-arrivals.csv`.

17.3.1. The synthetic arrival series

Arrivals per hour are drawn from a negative binomial distribution around an expected value that combines a daily profile, a weekly profile, the change of season and the surge. A day-level factor adds variation between days. The negative binomial allows more variation than a Poisson distribution, which is what real arrival counts often show.

```
set.seed(1717)
n_weeks <- 26
hours <- expand.grid(hour = 0:23, day = 1:7, week = 1:n_weeks)   # day 1 = Monday
hours <- hours[order(hours$week, hours$day, hours$hour), ]
hours$t <- seq_len(nrow(hours))                                  # running hour index

simulate_arrivals <- function(hours, day_shape = 80, nb_size = 6) {
  # expected arrivals per hour of the day (0-23 h): quiet night, busy late morning
  daily_profile <- c(1.6, 1.3, 1.1, 1.0, 1.0, 1.2, 1.8, 2.8, 4.0, 5.0, 5.6, 5.8,
                     5.6, 5.4, 5.2, 5.0, 5.0, 5.1, 5.0, 4.6, 4.0, 3.3, 2.6, 2.0)
  weekly <- c(1.12, 1.02, 1.00, 0.98, 1.00, 0.96, 0.94)             # Monday to Sunday
  season <- ifelse(hours$week >= 15, 1.15, 1.00)                    # change of season
  surge <- ifelse(hours$week == 21 & hours$day %in% 5:7 & hours$hour >= 14, 2.0, 1.0)
  day_key <- paste(hours$week, hours$day)
  day_factor <- if (is.finite(day_shape)) rgamma(length(unique(day_key)), day_shape, day_sha
    rep(1, length(unique(day_key)))                                 # day-to-day variation
  names(day_factor) <- unique(day_key)
  mu <- daily_profile[hours$hour + 1] * weekly[hours$day] * season * surge * day_factor[day_
  if (is.finite(nb_size)) rnbinom(nrow(hours), size = nb_size, mu = mu) else rpois(nrow(hour
}
hours$arrivals <- simulate_arrivals(hours)
```

```
# shifts: early 06-14, late 14-22, night 22-06 (a night shift belongs to the
hours$shift <- ifelse(hours$hour >= 6 & hours$hour < 14, "early",
                      ifelse(hours$hour >= 14 & hours$hour < 22, "late", "nig
hours$shift_day <- ifelse(hours$hour < 6, (hours$day - 2) %% 7 + 1, hours$day

b17_daily <- tapply(hours$arrivals, paste(hours$week, hours$day), sum)
b17_daily_mean <- mean(b17_daily)
stopifnot(nrow(hours) == 26 * 168, all(hours$arrivals >= 0), !anyNA(hours$arr
head(hours[, c("week", "day", "hour", "shift", "arrivals")], 3)
```

```
#>   week day hour shift arrivals
#> 1    1   1    0 night        3
#> 2    1   1    1 night        0
#> 3    1   1    2 night        0
```

```
summary(as.numeric(b17_daily))
```

```
#>    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
#>   52.00   78.00   90.00   91.78  106.75  169.00
```

```
if (dir.exists("../data")) {
  write.csv(hours[, c("week", "day", "hour", "shift", "arrivals")], "../data/
}
```

17.3.2. Rolling-origin backtest

The function below loops over the forecast origins. The `stopifnot()` lines are the heart of the lab: every training hour lies before every test hour, and the seasonal naive forecast from the `forecast` package equals the hand computation “same hour last week”.

```
library(forecast)
backtest <- function(y, origin_weeks = 9:26, m = 168, level = 0.90) {
  out <- vector("list", length(origin_weeks))
  for (i in seq_along(origin_weeks)) {
    w <- origin_weeks[i]
    train_idx <- seq_len((w - 1) * m)                # all hours before the origin
    test_idx <- (w - 1) * m + seq_len(m)             # the 168 hours after the origin
    stopifnot(max(train_idx) < min(test_idx))        # no information from the future
    y_train <- y[train_idx]
    # reference: seasonal naive (same hour last week), normal-approximation interval
    fc <- snaive(ts(y_train, frequency = m), h = m, level = 100 * level)
    point_sn <- as.numeric(fc$mean)
    stopifnot(isTRUE(all.equal(point_sn, y_train[length(y_train) - m + seq_len(m)])))
    # alternative: mean of the same hour in the last four weeks, Poisson interval
    lambda <- rowMeans(matrix(tail(y_train, 4 * m), nrow = m))
    out[[i]] <- rbind(
      data.frame(method = "Same hour last week", week = w, t = test_idx, y = y[test_idx],
                 point = point_sn, lo = pmax(0, as.numeric(fc$lower)), hi = as.numeric(fc$up
      data.frame(method = "Four-week mean, Poisson interval", week = w, t = test_idx, y = y[
                 point = lambda, lo = qpois((1 - level) / 2, lambda), hi = qpois(1 - (1 - le
    )
  }
  do.call(rbind, out)
}
bt <- backtest(hours$arrivals)
bt <- merge(bt, hours[, c("t", "day", "hour", "shift", "shift_day")], by = "t")
bt$covered <- bt$y >= bt$lo & bt$y <= bt$hi       # inside the 90% interval?
bt$abs_err <- abs(bt$y - bt$point)
stopifnot(length(unique(bt$week)) == 18, nrow(bt) == 2 * 18 * 168)
```

17.3.3. Evaluation per method, week and shift

Planned capacity is taken to be the upper edge of the interval. A patient above that edge counts as an excess patient; an unused planned place counts as idle capacity. The weights are an *assumption* for teaching: an excess patient costs 1 point in the early shift, 2 in the late shift and 3 at night; an idle place costs 0.25 points.

```
w_excess <- c(early = 1, late = 2, night = 3)               # cost points per
w_idle <- 0.25                                              # cost points per
bt$excess <- pmax(0, bt$y - bt$hi)
bt$idle <- pmax(0, bt$hi - bt$y)
bt$cost <- w_excess[bt$shift] * bt$excess + w_idle * bt$idle

key <- c("Same hour last week" = "snaive", "Four-week mean, Poisson interval"
b17_mae <- tapply(bt$abs_err, key[bt$method], mean)
b17_cov <- tapply(bt$covered, key[bt$method], mean)
b17_cov_shift <- tapply(bt$covered, list(key[bt$method], bt$shift), mean)
b17_width <- tapply(bt$hi - bt$lo, key[bt$method], mean)
b17_cost_week <- tapply(bt$cost, key[bt$method], sum) / 18
b17_summary <- data.frame(
  method = names(key), mae = b17_mae[key], coverage = b17_cov[key],
  early = b17_cov_shift[key, "early"], late = b17_cov_shift[key, "late"], nig
  width = b17_width[key], cost_week = b17_cost_week[key]
)
b17_summary
```

```
#>                                   method      mae  coverage     early
#> snaive              Same hour last week 2.852183 0.8789683 0.8293651 0.816
#> mean4  Four-week mean, Poisson interval 2.276124 0.8201058 0.7886905 0.763
#>            night    width cost_week
#> snaive 0.9910714 9.281171  301.6320
#> mean4  0.9077381 5.938161  244.9722
```

```
# statements made in the core reading
# mean error (forecast minus observed) in the four weeks after the change of season
b17_bias_season <- tapply((bt$point - bt$y)[bt$week %in% 15:18], key[bt$method][bt$week %in%
# coverage in the late shifts of the surge weekend (week 21, Friday to Sunday)
in_surge <- bt$week == 21 & bt$day %in% 5:7 & bt$shift == "late"
b17_surge_cov <- tapply(bt$covered[in_surge], key[bt$method][in_surge], mean)
round(b17_bias_season, 2); round(b17_surge_cov, 2)
```

```
#>  mean4 snaive
#>  -0.48  -0.26
```

```
#>  mean4 snaive
#>   0.38   0.67
```

```
# statements made in the core reading
stopifnot(b17_mae["mean4"] < b17_mae["snaive"],                 # smaller point error
          b17_cov["mean4"] < 0.85,                              # Poisson interval too narro
          b17_cov_shift["snaive", "night"] > 0.95,              # reference too wide at nigh
          b17_cov_shift["snaive", "late"] < 0.85,               # ... and too narrow in the
          b17_bias_season["mean4"] < b17_bias_season["snaive"], # larger under-forecast afte
          b17_bias_season["snaive"] < 0,
          all(b17_surge_cov < 0.7))                             # both intervals fail in the
```

The point forecasts of the two methods differ in a way that is easy to see in a single week. Figure 17.2 shows the surge weekend of week 21: in its late shifts the observed arrivals stay inside the band in only 38% (four-week mean) and 67% (reference) of hours.

```
library(ggplot2)
surge <- bt[bt$week == 21 & bt$day %in% 5:7, ]
surge$hour_since_fri <- (surge$day - 5) * 24 + surge$hour
ggplot(surge, aes(hour_since_fri)) +
```

```
  geom_ribbon(aes(ymin = lo, ymax = hi), fill = "grey80") +
  geom_line(aes(y = point), linewidth = 0.5) +
  geom_point(aes(y = y), size = 1.3) +
  facet_wrap(~ method, ncol = 1) +
  scale_x_continuous(breaks = seq(0, 72, by = 12)) +
  labs(x = "Hour since Friday 00:00 (week 21)", y = "Arrivals per hour",
       caption = "Grey band: 90% interval; line: point forecast; points: obse
  theme_cdsr()
```

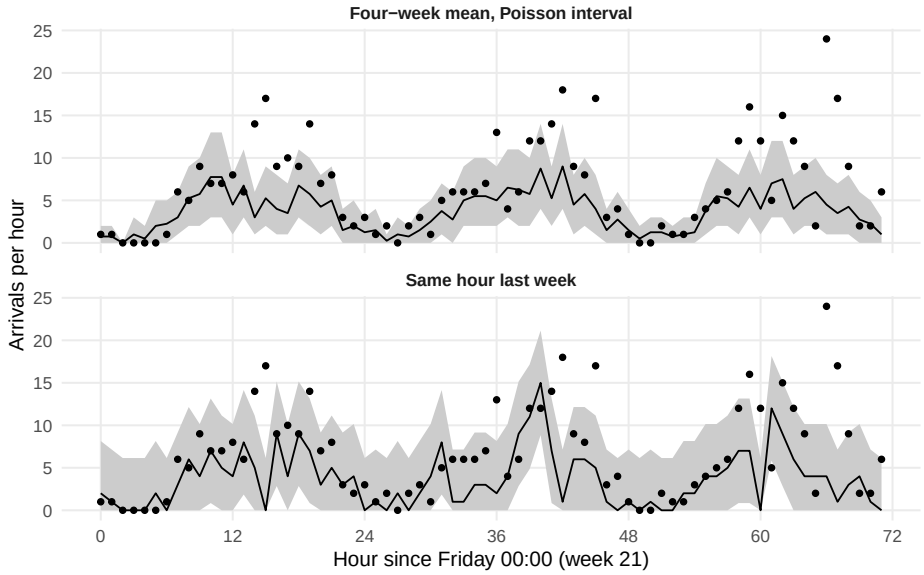
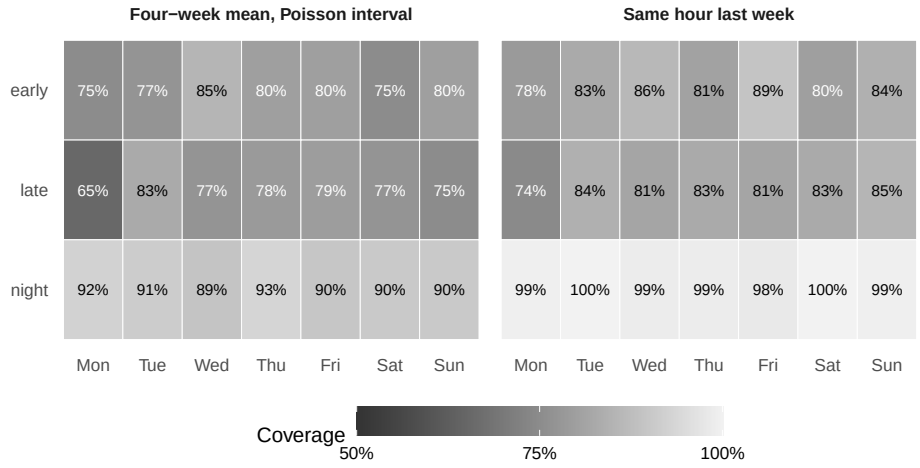


Figure 17.2.: Observed arrivals per hour (points) and 90% forecast intervals (grey band, line = point forecast) for Friday to Sunday of week 21, the surge weekend. One forecast week, 72 hours; arrivals per hour. Data status: synthetic.

17.3.4. A shift map

The shift map turns the hourly results into the unit of the decision. Each cell shows the coverage of the 90% interval for one weekday and shift, pooled over the 18 forecast weeks.

```
shiftmap <- aggregate(covered ~ method + shift_day + shift, data = bt, FUN = mean)
shiftmap$weekday <- factor(c("Mon", "Tue", "Wed", "Thu", "Fri", "Sat", "Sun")[shiftmap$shift
                           levels = c("Mon", "Tue", "Wed", "Thu", "Fri", "Sat", "Sun"))
shiftmap$shift <- factor(shiftmap$shift, levels = c("night", "late", "early"))
ggplot(shiftmap, aes(weekday, shift, fill = covered)) +
  geom_tile(colour = "white") +
  geom_text(aes(label = scales::percent(covered, accuracy = 1),
                colour = covered < 0.8), size = 3, show.legend = FALSE) +
  scale_fill_gradient(low = "grey20", high = "grey95", limits = c(0.5, 1), breaks = c(0.5, 0
                      labels = scales::label_percent(), name = "Coverage") +
  scale_colour_manual(values = c(`TRUE` = "white", `FALSE` = "black")) +
  facet_wrap(~ method, ncol = 2) +
  labs(x = NULL, y = NULL, caption = "Nominal coverage 90%. Night shift is assigned to the d
  theme_cdsr() +
  theme(panel.grid = element_blank(), legend.key.width = grid::unit(1.5, "cm"))
```



Nominal coverage 90%. Night shift is assigned to the day on which it starts.

Figure 17.3.: Shift map: coverage of the nominal 90% interval by weekday and shift, pooled over 18 forecast weeks (each cell: 18 shifts of 8 hours = 144 forecast hours from 18 independent forecast weeks). Darker cells have lower coverage. Data status: synthetic.

```
b17_worst <- shiftmap[order(shiftmap$covered), ][1:3, c("method", "weekday",
b17_worst
```

```
#>                                method weekday shift   covered
#> 15 Four-week mean, Poisson interval     Mon  late 0.6527778
#> 16              Same hour last week     Mon  late 0.7430556
#> 1  Four-week mean, Poisson interval     Mon early 0.7500000
```

```
stopifnot(b17_worst$weekday[1] == b17_worst$weekday[2], b17_worst$shift[1] ==
          b17_worst$shift[2] == "late", b17_worst$method[1] != b17_worst$metl
```

The lowest cell for both methods is the Monday late shift (65% and 74%), the busiest shift of the week, where the extra variation of the counts matters most. For the manager in the case, this is the actionable result: the label "90%" is not a property of the dashboard but of each shift, and it holds neither for the late shifts nor (in the other direction) for the nights.

17.3.5. Weighted error cost

```
b17_cost_shift <- aggregate(cost ~ method + shift, data = bt, FUN = function(x) sum(x) / 18)
b17_cost_shift
```

```
#>                            method shift      cost
#> 1 Four-week mean, Poisson interval early  79.43056
#> 2              Same hour last week early 100.61438
#> 3 Four-week mean, Poisson interval  late 113.19444
#> 4              Same hour last week  late 115.94460
#> 5 Four-week mean, Poisson interval night  52.34722
#> 6              Same hour last week night  85.07299
```

```
stopifnot(all(b17_cost_shift$cost >= 0))
```

With these weights the four-week mean has the lower weekly cost (245 versus 302 points), because the wide reference band leaves many more planned places idle; about half of these additional idle places fall in night shifts. Raise the weight of an excess night patient or of an excess late-shift patient, and the comparison moves towards the wider band. The weights encode a value judgement about who carries which error, which belongs in the roster discussion, not in the code alone.

17.4. Counter-check

Change one assumption about the process: suppose arrivals were truly Poisson, without extra variation between days or hours. Predict before running: will the coverage of the Poisson interval move towards 90%? Will the reference interval still be too wide at night and too narrow in the late shift?

```
set.seed(1718)
y_poisson <- simulate_arrivals(hours, day_shape = Inf, nb_size = Inf)   # sa
bt_p <- backtest(y_poisson)
bt_p <- merge(bt_p, hours[, c("t", "shift")], by = "t")
bt_p$covered <- bt_p$y >= bt_p$lo & bt_p$y <= bt_p$hi
b17_cov_p <- tapply(bt_p$covered, list(key[bt_p$method], bt_p$shift), mean)
round(cbind(overall = tapply(bt_p$covered, key[bt_p$method], mean), b17_cov_
```

```
#>        overall early  late night
#> mean4    0.907 0.900 0.885 0.937
#> snaive   0.884 0.852 0.816 0.982
```

```
stopifnot(mean(bt_p$covered[key[bt_p$method] == "mean4"]) > b17_cov["mean4"]
          b17_cov_p["snaive", "night"] > b17_cov_p["snaive", "late"])
```

Under the Poisson assumption the Poisson interval covers 90.7% of hours, close to its nominal level. The reference interval still covers 98.2% at night and 81.6% in the late shift: its width is the same for every hour, while the variation of counts grows with the expected number. The counter-check separates two reasons for poor coverage. One is a wrong distributional assumption (Poisson for overdispersed counts); the other is an interval that ignores how uncertainty changes over the day. Both are properties of the method, not of the hospital.

17.5. Consequences for patients and operations

Ethics and distribution of burden. A forecast used for rostering shifts errors onto concrete people. If intervals are systematically too narrow in weekend late shifts, the patients arriving then wait longer and the late-shift team absorbs the overload, often the same staff week after week. Reporting coverage per shift, not only overall, makes this distribution visible. The cost weights in the lab are a place to make the value judgement explicit and to involve the nursing team, not only the analyst.

Regulation. Indirect. A forecast of arrival counts for staffing has no diagnostic or therapeutic purpose for an individual patient; whether such software falls under medical device rules is a question of its intended purpose (Chapter 22). The question would change if the same model were used to prioritise individual patients: the AI Act lists "emergency healthcare patient triage systems" among high-risk uses in Annex III, point 5(d) (European Parliament and Council of the European Union 2024b). As of November 2026.

Data protection. The model needs counts per hour, not patient records. Aggregating timestamps to hourly counts before analysis is a direct application of data minimisation under Art. 5(1)(c) of the General Data Protection Regulation (European Parliament and Council of the European Union 2016). Very small counts in rare categories (for example arrivals by a specific transport mode at night) can still point to individuals when combined with local knowledge; Chapter 20 and Chapter 26 treat small aggregates.

Security and availability. If the forecast service fails, the department needs a fallback that works without it. The seasonal naive forecast is that fallback: last week's arrivals, available from any routine report.

17.6. Look up and transfer

- **Reviewer question:** Which forecast origins and horizons were evaluated, was any information from after the origin used (for example smoothing over the whole series or a calendar compiled later), and what is the interval coverage per shift, with the number of forecast weeks as denominator?
- **Self-test:** A smaller department has on average 1.2 arrivals per hour at night, with many hours of zero arrivals. Its analyst compares two forecasts by the mean absolute percentage error (MAPE) over all hours and reports that model X is "40% better". What is the problem with this metric here, and what would you report instead?

Answer

Percentage errors divide by the observed value. In hours with zero arrivals they are undefined, and in hours with one arrival a single extra patient is an error of 100%, so the night hours dominate the comparison. Report the mean absolute error in arrivals per hour, or a scaled error that divides by the error of the seasonal naive reference (the MASE) (Hyndman and Koehler 2006), separately for each shift and with the number of forecast weeks. For the interval, report coverage per shift.

- **Annotated sources:**
  - Hyndman and Athanasopoulos (2021) is the open textbook for benchmark methods, prediction intervals and time series cross-validation used here (simple methods, time series cross-validation).
  - Hyndman and Khandakar (2008) documents the `forecast` package and its automatic methods; the lab uses its `snaive()`.
  - Jones et al. (2008) compared several forecasting methods for daily emergency department volumes at three hospitals and

found that calendar-based regression was a reasonable approach, while more complex methods did not give consistently better post-sample accuracy.
  - Tashman (2000) reviews out-of-sample evaluation and rolling forecast origins.
  - Gneiting et al. (2007) define the goal "maximise sharpness subject to calibration" for probabilistic forecasts.
  - Hyndman and Koehler (2006) discuss accuracy measures, including scaled errors for series with zeros.

In W6 (teaching unit 21) learners compare forecasts against "same hour last week" in later windows and hand in a table with horizon, error measure, interval coverage and one shift with possible harm. In W3 the chapter serves as a transfer case: the same distinction between ranking, calibration and decision cost appears for diagnostic models in Chapter 6.

17.7. Further reading

- Chapter 18 for the second question on the planning screen: what a capacity change would do, given arrivals.
- Chapter 6 for calibration and decision thresholds in a diagnostic setting.
- Chapter 20 for monitoring a deployed tool, including drift after a change of season or practice.
- Methods in R: Forecast Accuracy Metrics for the definitions of MAE, RMSE, MAPE and scaled errors.
- Methods in R: Seasonal Decomposition (STL) for separating daily, weekly and seasonal components.
- CTTIR Tutorials: Rolling-Origin Cross-Validation for a step-by-step tutorial on the backtest design.
- CTTIR Courses, Course 3, Week 2 Session 5: Time series basics for a hands-on introduction to time series in R.

18. Mass casualty incidents, operating capacity and simulation

R. Heller · Date: 2025-08-17 · Last changed: 2026-11-09

At a glance

Guiding question: Which resource limits care in the scenario under study?
For: trauma surgeons, anaesthetists and nurses who write or test a hospital's mass casualty plan, emergency planners, analysts in hospital operations · **Prerequisites:** the difference between a forecast and a simulated capacity change (Chapter 17); random draws and replications, as in Strategy in R: Monte Carlo Methods for Games
In the course: W6 · Pharmacometrics, forecasting and simulation (required reading, as an alternative to Chapter 15) · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

18.1. Case and decision

A trauma surgeon chairs the working group that revises the hospital's plan for mass casualty incidents. Two proposals compete for the first hour after an alert. One is to open a third operating room at once, which

means stopping the elective list in that room and calling in a second on-call team. The other is to equip a third resuscitation bay in the emergency department, which is cheaper and does not cancel anyone's operation. The group asks an analyst for a simulation. The surgeon's question is concrete: *in the scenario we consider likely, which resource limits care, and how much waiting does each proposal remove for which patients?* The answer goes into the plan. A wrong choice costs time for seriously injured patients in a real incident, and either choice has costs for elective patients and staff that the plan must state openly.

18.2. Core reading

Core reading · required for W6 (alternative to Chapter 15) · about 10 minutes · ends before 'Lab'

What a simulation can claim. A forecast (Chapter 17) predicts what will happen if nothing changes and can be checked a week later. A capacity simulation answers a "what if" question about a change that has not been made, under an assumed stream of patients. It cannot be checked against the outcome before a real incident. Its credibility rests on three things that this chapter makes explicit: the input assumptions, the verification of the model logic, and a plausibility check against external figures (Sargent 2013).

Discrete-event and agent-based models. In a *discrete-event simulation* (DES), patients move through a fixed pathway of resources (here: resuscitation bay, then operating room). The clock jumps from event to event: an arrival, the start or end of a treatment. Queues form when a resource is busy, and a priority rule decides who goes next. The R package `simmer` implements this approach (Ucar et al. 2019). In an *agent-based model*, individuals follow their own rules and interact, for example staff

who tire or clinicians who re-triage patients as they deteriorate (see Strategy in R: Agent-Based Models for Strategic Interaction). For a question about the capacity of a fixed pathway, DES is the simpler and more transparent choice.

The scenario. After an alert at time zero, 14 casualties who need urgent surgery arrive within three hours. About 40% are triaged red (immediate), the rest yellow (urgent but able to wait). Each patient first occupies one of two resuscitation bays for about 30 minutes and then waits for an operating room. Operations take about two hours for red and one and a half hours for yellow patients, with considerable spread. Red patients have priority at every step. Three scenarios are compared: A, two operating rooms (the regular night set-up); B, three operating rooms; C, two operating rooms and a third resuscitation bay. All numbers are *assumptions* for teaching. A real plan would take them from the hospital's own records and from published series.

Fix the events before comparing. The most important design decision comes before any simulation is run. For each of 200 replications, the list of arrival times, triage categories and treatment times is drawn once and stored. All three scenarios then run on exactly the same lists. Differences between scenarios can only come from the capacity, not from one scenario happening to receive sicker patients. In the lab this pairing makes the spread of the estimated difference between A and B about 2.7 times smaller than with independent draws.

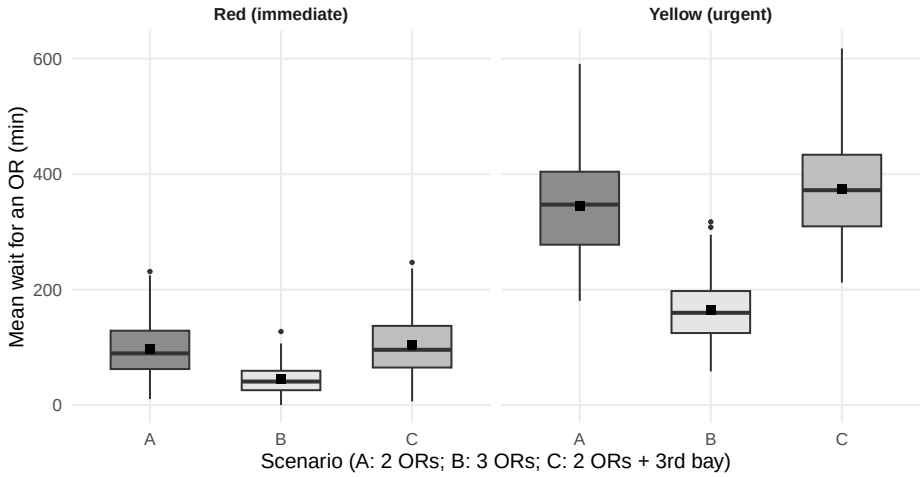


Figure 18.1.: Mean waiting time for an operating room per replication (minutes), by triage category and scenario. Each box summarises 200 replications; all scenarios use the same 200 event lists of 14 casualties. A: 2 operating rooms; B: 3 operating rooms; C: 2 operating rooms and a third resuscitation bay. Data status: synthetic.

Table 18.1.: Results per scenario over 200 replications with identical event lists (14 casualties each). Waiting times in minutes, mean over replications with the 5th to 95th percentile across replications in brackets. Throughput: operations finished within 6 hours of the alert. Data status: synthetic; all inputs are assumptions.

Scenario	Wait for OR, red	Wait for OR, yellow	Wait for bay	Operations within 6 h	Last operation ends (h)
A: 2 ORs, 2 bays	97 (25–190)	345 (222–475)	26 (11–49)	4.4 (3–6)	13.9
B: 3 ORs, 2 bays	45 (12–90)	165 (92–244)	26 (11–49)	6.5 (5–9)	10.1
C: 2 ORs, 3 bays	103 (23–201)	374 (247–501)	6 (1–16)	4.4 (3–6)	14.0

Which resource limits care here? In this scenario the operating rooms do. A third operating room roughly halves the waiting time of yellow patients (from 345 to 165 minutes on average) and finishes the last operation about 3.9 hours earlier. A third resuscitation bay removes most of the waiting for a bay but does not shorten the wait for an operating room; patients simply reach the operating room queue sooner and wait there instead. This is the typical behaviour of a system with one binding constraint: relieving another step moves the queue rather than removing it.

The most important assumption. The comparison holds only for the arrival pattern, treatment times and priority rules written into the model. The simulation shows the consequences of these

assumptions; it does not show that the real hospital behaves this way. A published discrete-event model of urban bombing incidents, built from data of 12 incidents, found the resuscitation rooms and the CT scanner, not the operating rooms, to be the main bottlenecks for critical casualties (Hirshberg et al. 1999). That is not a contradiction. Our model has no CT step and only two bays with short occupancy; with other assumptions the binding resource changes. A follow-up model defined surge capacity as a *rate* of critical casualties per hour rather than a number of beds, about 4.6 per hour with immediately available assets in the studied trauma centre (Hirshberg et al. 2005). Such published figures are useful for a plausibility check of inputs and outputs, not as a validation of a different hospital.

Who carries the waiting? Averages hide the distribution. The priority rule means that yellow patients carry most of the waiting in every scenario; in scenario A, one in ten replications has a 90th-percentile yellow wait above 9.7 hours. A yellow patient who waits that long may deteriorate, which this model does not represent. Scenario B has costs that fall outside the model: the elective operations cancelled in the third room (an *assumption* of four per activation) and a second on-call team that will be tired the next day. A complete answer to the surgeon names all three groups.

Question for the workshop: Your group gets the same event lists but a plan with three bays and three operating rooms. Before simulating, write down which resource you expect to bind, which patients will carry most of the waiting, and which group outside the model carries a cost.

18.3. Lab

The lab builds the model in five steps: fixed event lists, the model, a verification against a hand calculation, the scenario runs with checks of resource release, and a summary with variation between runs.

R start command

From the root of the book repository, render the chapter (this runs all code, including the figures of the core reading):

```
quarto render chapters/b18-mass-casualty-capacity-simulation.qmd --to html
```

To run only the Lab code in a fresh R session, start from the `chapters/` folder:

```
Rscript -e 'source(knitr::purl("b18-mass-casualty-capacity-simulation.qmd", output = tempf
```

The chapter writes the fixed event lists to `data/b18-event-list.csv`; the same file is the shared input for any re-implementation.

18.3.1. Step 1: draw and store the event lists

```
set.seed(1818)
R_rep <- 200                                    # replications
make_events <- function(rep, n = 14) {
  red <- runif(n) < 0.4                         # triage: 40% red (immediate), rest yellow
  data.frame(
    rep = rep, patient = seq_len(n),
    time = round(sort(runif(n, 0, 180)), 1),    # arrival, minutes after alert (3-hour windo
    triage = ifelse(red, "red", "yellow"),
    priority = ifelse(red, 2L, 1L),             # simmer: higher number is served first
    bay_min = round(rlnorm(n, log(30), 0.3), 1),                          # resuscitation ba
    or_min = round(rlnorm(n, log(ifelse(red, 120, 90)), 0.35), 1)         # operating time
  )
}
```

```
events <- do.call(rbind, lapply(seq_len(R_rep), make_events))
events_fingerprint <- digest::digest(events)   # checksum: the lists must not
head(events, 4)
```

```
#>   rep patient time triage priority bay_min or_min
#> 1   1       1  1.8    red        2    25.5  108.9
#> 2   1       2 53.5 yellow        1    34.5   86.3
#> 3   1       3 58.4 yellow        1    37.8  162.8
#> 4   1       4 65.4 yellow        1    29.1   49.1
```

```
stopifnot(nrow(events) == R_rep * 14, all(events$bay_min > 0), all(events$or_

if (dir.exists("../data")) write.csv(events, "../data/b18-event-list.csv", ro
```

18.3.2. Step 2: the model

Each patient seizes a bay, holds it for the stored bay time, releases it, then seizes an operating room for the stored operating time. The arrivals come from the stored list (`add_dataframe`), so the model itself draws no random numbers.

```
library(simmer)
run_scenario <- function(ev, n_or, n_bay = 2) {
  env <- simmer("mass-casualty")
  patient <- trajectory("patient") |>
    seize("bay", 1) |>
    timeout_from_attribute("bay_min") |>
    release("bay", 1) |>
    set_attribute("ready_for_or", function() now(env)) |>   # time the patien
    seize("or", 1) |>
    set_attribute("or_start", function() now(env)) |>
```

```
    timeout_from_attribute("or_min") |>
    release("or", 1)
  env |>
    add_resource("bay", capacity = n_bay) |>
    add_resource("or", capacity = n_or) |>
    add_dataframe("p", patient, ev[, c("time", "bay_min", "or_min", "priority")], mon = 2,
                  col_time = "time", time = "absolute",
                  col_attributes = c("bay_min", "or_min"), col_priority = "priority") |>
    run()
}

# one row per patient with waiting times; arrival "p0" is row 1 of the event list
patient_table <- function(env, ev) {
  arr <- get_mon_arrivals(env)
  att <- get_mon_attributes(env)
  att <- att[att$key %in% c("ready_for_or", "or_start"), c("name", "key", "value")]
  att <- reshape(att, idvar = "name", timevar = "key", direction = "wide")
  arr <- merge(arr, att, by = "name")
  arr$row <- as.integer(sub("^p", "", arr$name)) + 1
  arr <- arr[order(arr$row), ]
  arr$triage <- ev$triage[arr$row]
  arr$wait_bay <- round(arr$value.ready_for_or - arr$start_time - ev$bay_min[arr$row], 6)
  arr$wait_or <- round(arr$value.or_start - arr$value.ready_for_or, 6)
  arr
}
```

18.3.3. Step 3: verification against a hand calculation

Before trusting the model with 200 replications, it must reproduce a case that can be computed on paper. Three patients arrive at minutes 0, 1 and 2; there is one bay (10 minutes each) and one operating room (60 minutes each). Patient 3 is red. By hand: patient 1 uses the bay from 0 to 10 and

the operating room from 10 to 70 (wait 0). At minute 10 the red patient 3 jumps the bay queue (bay 10 to 20) and waits for the operating room from 20 to 70 (wait 50). Patient 2 uses the bay from 20 to 30 and waits until patient 3 is finished at 130 (wait 100).

```
hand <- data.frame(time = c(0, 1, 2), bay_min = 10, or_min = 60, priority = c
                   triage = c("yellow", "yellow", "red"))
hand_env <- run_scenario(hand, n_or = 1, n_bay = 1)
hand_tab <- patient_table(hand_env, hand)
hand_tab[, c("name", "triage", "start_time", "value.ready_for_or", "value.or_
```

```
#>   name triage start_time value.ready_for_or value.or_start wait_or end_tim
#> 1   p0 yellow          0                 10             10       0       7
#> 2   p1 yellow          1                 30            130     100      19
#> 3   p2    red          2                 20             70      50      13
```

```
stopifnot(all(hand_tab$wait_or == c(0, 100, 50)),     # hand calculation
          all(hand_tab$end_time == c(70, 190, 130)))
```

18.3.4. Step 4: run the scenarios and check resource release

For every run, four checks are made: every patient finished; arrival times equal the stored list; at the end both resources are free and their queues empty (no patient kept a resource it should have released); and the area under the operating-room queue-length curve equals the sum of all individual waits for an operating room. The last check is an exact accounting identity behind Little's law (Little 1961): time spent waiting, counted per patient or counted per minute of queue, must be the same.

```
run_metrics <- function(ev, n_or, n_bay = 2) {
  env <- run_scenario(ev, n_or, n_bay)
  a <- patient_table(env, ev)
```

```
  r <- get_mon_resources(env)
  r <- r[order(r$time), ]
  final <- aggregate(cbind(server, queue) ~ resource, r, function(x) tail(x, 1))
  r_or <- r[r$resource == "or", ]
  queue_area <- sum(head(r_or$queue, -1) * diff(r_or$time))   # integral of OR queue length
  stopifnot(nrow(a) == nrow(ev), all(a$finished),               # everyone finished
            isTRUE(all.equal(a$start_time, ev$time)),           # inputs used unchanged
            all(final$server == 0), all(final$queue == 0),      # resources released
            isTRUE(all.equal(queue_area, sum(a$wait_or))))      # accounting identity
  c(wait_or_red = mean(a$wait_or[a$triage == "red"]),
    wait_or_yellow = mean(a$wait_or[a$triage == "yellow"]),
    p90_yellow = unname(quantile(a$wait_or[a$triage == "yellow"], 0.9)),
    wait_bay = mean(a$wait_bay),
    done_6h = sum(a$end_time <= 360),
    makespan = max(a$end_time))
}
scenarios <- list(A = c(n_or = 2, n_bay = 2), B = c(n_or = 3, n_bay = 2), C = c(n_or = 2, n_
# lists with no yellow patient give NaN for yellow waits; they are kept and reported
b18_res <- lapply(scenarios, function(s)
  t(sapply(seq_len(R_rep), function(r) run_metrics(events[events$rep == r, ], s["n_or"], s["
stopifnot(identical(digest::digest(events), events_fingerprint))   # same events in every sc

b18_mean <- t(sapply(b18_res, colMeans, na.rm = TRUE))
b18_q <- t(sapply(b18_res, function(m) c(apply(m, 2, quantile, 0.05, na.rm = TRUE),
                                         p90_yellow_top10 = unname(quantile(m[, "p90_yellow
fmt <- function(s, v, d = 0) sprintf("%s (%s-%s)", round(b18_mean[s, v], d),
                                     round(quantile(b18_res[[s]][, v], 0.05, na.rm = TRUE)
                                     round(quantile(b18_res[[s]][, v], 0.95, na.rm = TRUE)
b18_table <- data.frame(
  scenario = c("A: 2 ORs, 2 bays", "B: 3 ORs, 2 bays", "C: 2 ORs, 3 bays"),
  red = sapply(names(scenarios), fmt, v = "wait_or_red"),
  yellow = sapply(names(scenarios), fmt, v = "wait_or_yellow"),
  bay = sapply(names(scenarios), fmt, v = "wait_bay"),
```

```
  done = sapply(names(scenarios), fmt, v = "done_6h", d = 1),
  last = sapply(names(scenarios), function(s) sprintf("%.1f", b18_mean[s, "ma
)
round(b18_mean, 1)
```

```
#>   wait_or_red wait_or_yellow p90_yellow wait_bay done_6h makespan
#> A        96.8          345.0      469.2     26.4     4.4    836.7
#> B        44.9          165.2      250.0     26.4     6.5    604.7
#> C       103.4          374.2      512.2      6.2     4.4    837.7
```

```
# statements made in the core reading
stopifnot(b18_mean["B", "wait_or_yellow"] < 0.6 * b18_mean["A", "wait_or_yell
          b18_mean["C", "wait_bay"] < 0.5 * b18_mean["A", "wait_bay"],
          b18_mean["C", "wait_or_yellow"] >= b18_mean["A", "wait_or_yellow"])
```

18.3.5. Step 5: variation between runs, paired and unpaired

The table in the core reading reports the spread across replications. The more important spread for the decision is that of the *difference* between scenarios. With paired event lists, each replication gives one difference B minus A for the same 14 patients. For comparison, scenario B is also run on a second, independent set of event lists.

```
diff_paired <- b18_res$B[, "wait_or_yellow"] - b18_res$A[, "wait_or_yellow"]
set.seed(1819)
events_independent <- do.call(rbind, lapply(seq_len(R_rep), make_events))
res_B_independent <- t(sapply(seq_len(R_rep), function(r)
  run_metrics(events_independent[events_independent$rep == r, ], n_or = 3)))
diff_unpaired <- res_B_independent[, "wait_or_yellow"] - b18_res$A[, "wait_or
b18_sd <- c(paired = sd(diff_paired, na.rm = TRUE), unpaired = sd(diff_unpair
b18_sd_ratio <- unname(b18_sd["unpaired"] / b18_sd["paired"])
round(c(mean_paired = mean(diff_paired, na.rm = TRUE), b18_sd, ratio = b18_sd
```

```
#> mean_paired      paired    unpaired       ratio
#>      -179.9        39.4       105.4         2.7
```

```
stopifnot(b18_sd["paired"] < b18_sd["unpaired"],
          mean(diff_paired < 0, na.rm = TRUE) > 0.95)   # B shortens the yellow wait in almo
```

With paired lists, scenario B shortens the mean yellow wait in 100.0% of replications, by 180 minutes on average. The standard deviation of the difference is 39 minutes with paired lists and 105 minutes with independent lists. Pairing does not change what the model assumes; it removes noise that has nothing to do with the question.

> **i** Other implementation: SimPy
>
> **Original implementation.** SimPy, a process-based discrete-event simulation library for Python (Team SimPy, n.d.), version 4.1.2 (current release on the Python Package Index, accessed 8 November 2026); documentation.
> **Concrete advantage.** Each patient is written as a Python generator function, so interruptions, pre-emption and re-triage can be followed line by line in the program flow. Teams with an existing Python data pipeline can keep the simulation in the same environment.
> **What must be identical for a fair comparison.** The input file `data/b18-event-list.csv` (arrival, triage, bay and operating times per replication), the capacities, and the priority rule. Note the direction: in `simmer` a *higher* priority number is served first, in SimPy's `PriorityResource` a *lower* number is. Ties are resolved by the time of the request in SimPy; check that equal arrival times are ordered the same way. Compare the event log (order of starts and ends) first, then the waiting times. Tolerance: identical event order and waiting times equal up to floating-point rounding (10 minutes). The same seed in R and Python does not produce the same random

draws, which is why the event list is exchanged as a file rather than regenerated.
Status. Conceptual comparison, not executed. No statement about speed is made.

18.4. Counter-check

Scenario C, the third resuscitation bay, is the counter-check. Before looking at the numbers, predict: will the mean wait for a bay go up or down? Will the mean wait for an operating room go up, down or stay the same? Then compare with the code output.

```
round(rbind(A = b18_mean["A", c("wait_bay", "wait_or_red", "wait_or_yellow",
            C = b18_mean["C", c("wait_bay", "wait_or_red", "wait_or_yellow",
```

```
#>   wait_bay wait_or_red wait_or_yellow makespan
#> A     26.4        96.8          345.0    836.7
#> C      6.2       103.4          374.2    837.7
```

```
stopifnot(abs(b18_mean["C", "makespan"] - b18_mean["A", "makespan"]) < 15)  #
```

The wait for a bay falls from 26 to 6 minutes, but the wait for an operating room does not fall; for yellow patients it rises slightly (345 to 374 minutes), because patients now reach the operating room queue earlier and wait there instead, while the two operating rooms work through the same caseload at the same pace. The time until the last operation ends hardly changes. In this scenario the third bay buys comfort in the emergency department, not earlier surgery. Change one more assumption yourself: shorten the bay time to 15 minutes, or add a CT step with one scanner, and check whether the binding resource moves.

18.5. Consequences for patients and operations

Ethics. A triage rule is a value judgement about who waits. The model makes its consequences visible: in every scenario yellow patients carry most of the waiting, and the model does not represent deterioration while waiting, which would turn some yellow patients red. Capacity decisions also distribute burdens outside the model: elective patients whose operations are cancelled, and the on-call team that works through the night. A mass casualty plan that reports only the mean waiting time of casualties leaves these groups out; Chapter 21 discusses how to make such trade-offs explicit.

Reporting and validity. Simulation studies should report inputs, their sources, the handling of random numbers, the number of replications and the verification steps; the STRESS guidelines give a checklist (Monks et al. 2019). A plan based on a simulation should state which inputs were assumptions and which came from local data, and when they will be checked against the next exercise or real incident.

Regulation. Indirect. A planning simulation for a mass casualty plan has no medical purpose for an individual patient. The question changes if software is used during an incident to assign individual patients to operating rooms or to set treatment priority. Systems for emergency healthcare patient triage are listed as high-risk uses in Annex III, point 5(d), of the AI Act (European Parliament and Council of the European Union 2024b); for such a use the questions of Chapter 22 apply. As of November 2026.

Data protection and security. Treatment times derived from the hospital's operating-room records are health data when they are linked to persons; for the model, aggregated distributions per procedure group are sufficient. During a real incident, the plan must work if scheduling software is unavailable; a printed version of the decision rules is part of the fallback (Chapter 24).

18.6. Look up and transfer

- **Reviewer question:** Were the scenarios compared on identical input events, how many replications were run, how was the model verified, and which inputs are local data and which are assumptions?
- **Self-test:** A team compares two and three operating rooms with 50 replications per scenario. Each scenario uses its own random seed. They report mean yellow waits of 410 and 200 minutes with overlapping 95% intervals across replications and conclude that "the difference is uncertain". What would you change in the design and in the report?

Answer

Draw the event lists once and run both scenarios on the same lists (common random numbers). Report the distribution of the paired differences per replication, for example the mean difference with its 5th and 95th percentiles and the share of replications in which three rooms are better. Overlapping intervals of two separate means are the wrong comparison: most of their width comes from differences between event lists, which affect both scenarios alike. Also report the waits of red and yellow patients separately and name the costs outside the model (cancelled elective operations, staff).

- **Annotated sources:**
  - Ucar et al. (2019) describe `simmer`, the discrete-event simulation package used in the lab; the simmer website documents `add_dataframe()` and the monitoring functions.
  - Sargent (2013) summarises verification and validation of simulation models, including the difference between checking the model logic and checking the model against reality.
  - Monks et al. (2019) (STRESS) provide reporting guidelines for simulation studies.

  - Hirshberg et al. (1999) built a discrete-event model of hospital resources in urban bombing incidents and found resuscitation rooms and CT, not operating rooms, to be the bottlenecks for critical casualties.
  - Hirshberg et al. (2005) related casualty load to the level of trauma care and expressed surge capacity as a rate of critical casualties per hour.
  - Little (1961) proved the queuing formula L = W, which underlies the accounting check in step 4.
  - Team SimPy (n.d.) is the documentation of SimPy for the comparison box.

In W6 (teaching unit 22, station “mass casualty”) learners simulate paired arrival sequences for two capacity scenarios and hand in a waiting-time distribution with its spread across runs. The project work in teaching unit 24 uses the same structure: model, baseline scenario and one counter-check that another group reproduces from a fresh R session.

18.7. Further reading

- Chapter 17 for the forecast that often supplies the arrival assumptions of a capacity simulation.
- Chapter 15 for the other simulation type in W6, a mechanistic model checked against an analytical solution.
- Chapter 21 for the ethics of triage and of distributing waiting between patient groups.
- Chapter 27 for reporting a simulation in a model passport.
- Strategy in R: Agent-Based Models for Strategic Interaction for models in which individuals adapt their behaviour.
- Strategy in R: Probability Essentials, law of large numbers and Monte Carlo for why more replications narrow the estimate of a mean but not the spread between runs.

- Methods in R cheat sheet: Power by simulation for the same replicate-and-summarise pattern in study planning.

19. From one centre or species to the next

R. Heller · Date: 2025-09-28 · Last changed: 2026-11-09

At a glance

Guiding question: To which new case should the claim refer?
For: doctoral researchers in preclinical imaging, clinical research teams planning a validation, reviewers of prediction and imaging studies · **Prerequisites:** units of analysis and denominators (Chapter 3), training and test data with a baseline (Chapter 5), spectra and the hyperspectral measurement chain (Chapter 10)
In the course: W4 · What imaging, gait and tissue models measure (required reading), also used in W6 · Pharmacometrics, forecasting and simulation and W8 · Model colloquium and decision · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

19.1. Case and decision

A doctoral researcher in a preclinical imaging group presents a slide at the lab meeting: a classifier separates two tissue classes, A and B, in hyperspectral reflectance spectra, and the test accuracy is 100%. The data come from four animals with three spectra each. The group leader wants to quote the result in a funding application for a first feasibility

study in patients with soft-tissue injuries. Before the sentence goes into the application, a methods reviewer in the group asks one question: *which new case was the test about?* A further spectrum from an animal that the model has already seen, a new animal from the same laboratory, an animal measured next year with a new camera, or a patient? The answer decides what the group may write. If the test only measured recognition of known animals, the application would ask patients to carry the burden of a study whose central premise was never tested.

19.2. Core reading

Core reading · required for W4 · about 10 minutes · ends before ‘Lab’

Every test answers a question about a particular kind of new case. A model evaluated on new spectra from animals it has already seen answers the question “does it recognise more data from these animals?”. A model evaluated on animals it has never seen answers “does it work for a new animal from the same setting?”. Neither answers “does it work in another centre, with another device, or in another species?”. The skill this chapter trains is to name the intended new case *before* choosing the split, and then to check that the split really produces such cases.

Five levels of “new”. It helps to list the levels explicitly (Table 19.1). Each row needs its own test, and a positive result on an upper row says nothing about the rows below it. The general idea of keeping training and test data apart is covered in Chapter 5 and in Methods in R: Train-Test Splits; here the point is which *unit* must be new.

Table 19.1.: Levels of "new" and the split that tests each one. The denominator is the number of independent units that the claim is about.

Intended new case	Split that produces it	Independent denominator
Another spectrum or image of a known animal	Image-wise: spectra of each animal on both sides	Spectra (but they are not independent)
A new animal, same laboratory and device	Animal-wise (group split): every animal on one side only	Animals in the test set
A later animal, after protocol drift	Temporal: train on earlier, test on later animals	Animals measured after the cut-off
Another centre, device or operator	External: develop at one site, test at another	Animals (or patients) at the new site
Another species, or patients	A new study in the target population	Individuals of the target species

A deliberately trivial example. Four fictitious animals, T1 to T4, each contribute three spectra. The classes are properties of animals: T1 and T3 belong to class A, T2 and T4 to class B. Instead of a real classifier we use a didactic model that is constructed to cheat in a transparent way. It never looks at a spectrum. For an animal ID it has seen during training, it returns the class it recorded for that animal. For any other ID it returns a fallback class that was fixed in advance: class A. Call it the *recall model*. If a test cannot distinguish this model from a useful one, the test is not informative about new animals.

What the two designs show. In the image-wise split, two spectra of every animal are used for training and the third for testing. Every test ID

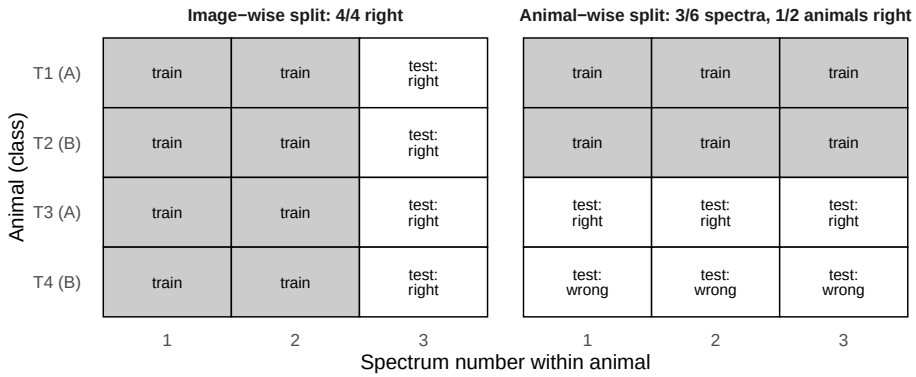


Figure 19.1.: Two split plans for the same 12 synthetic spectra (n = 4 animals; classes A/B are properties of animals). Grey tiles are training data, white tiles test data; the text gives the recall model's result. Left: image-wise split, 4 of 4 test spectra right. Right: animal-wise split, 3 of 6 test spectra and 1 of 2 test animals right. Data status: synthetic.

is known, so the recall model gets 4 of 4 right. It has measured recognition, not classification. In the animal-wise split, T1 (A) and T2 (B) are used for training; T3 (A) and T4 (B) with three spectra each form the test set. Both IDs are unknown, so the fallback class A is returned six times: 3 of 6 spectra are right, and at the level of animals 1 of 2. The independent denominator is two animals, not six spectra. A result of "50%" from two animals is compatible with almost anything; it is a statement about two individuals.

The most important assumption. A test is informative about the intended new case only if the unit that is new in the test is the unit that the claim is about, and the denominator counts those units. For a claim about new animals the test must contain animals the model has never seen, and accuracy must be reported per animal. In R this is one line: `intersect(unique(train$id), unique(test$id))` must be empty for the animal-wise split. The same leak appears under many names in the literature, from pseudoreplication to "data leakage" (Kapoor and Narayanan 2023), and grouped or blocked cross-validation is the standard remedy for clustered data (Roberts et al. 2017).

Beyond the group split. An empty ID intersection is necessary, not sufficient. A clinical prediction model can be *reproducible* in new patients from the same population yet fail to be *transportable* to a different but related population (Justice et al. 1999). External validation is therefore interpreted against the difference in case mix between development and validation samples (Debray et al. 2015). Two further shifts matter in imaging. A change in measurement procedure (camera, illumination, calibration, operator) can alter spectra in ways that a model learns as if it were biology; an imaging model for chest radiographs, for example, could identify the hospital system of a radiograph almost perfectly, which can confound its predictions (Zech et al. 2018). And a change of species is not a larger version of a change of centre. Whether mouse models "mimic" human inflammatory responses after trauma and burns has been answered in opposite directions from the same gene expression data, depending on which genes were compared (Seok et al. 2013; Takao and Miyakawa 2015).

The lesson for this chapter is modest: similarity between species is a claim with its own analysis choices, and a model trained in animals makes no claim about patients until it has been tested in patients.

What the reviewer writes back. “The 100% result is an image-wise test. It supports the statement that spectra of the four known animals can be recognised. For new animals, the independent denominator is the number of test animals; please report the animal-wise result with its denominator. A statement about patients needs a study in patients.” The finding is not worthless. It is an honest description of what was tested, and it points to the next study, which in preclinical work also has to justify the number of animals under the 3Rs (Chapter 16).

Question for the workshop: Your group has spectra from six animals measured in two laboratories with two cameras. Write down the intended new case for your claim, the split that produces it, and the denominator you would report. Which row of Table 19.1 can your data not reach at all?

19.3. Lab

The lab has three parts. First, the hand calculation from the core reading is reproduced exactly, with target values. Second, a self-contained variant uses a nearest-neighbour model on synthetic spectra with the same animal IDs; its results are kept in separate objects. Third, a deepening exercise treats a missing modality as a stress test.

> **R start command**
>
> From the root of the book repository, render the chapter:

```
quarto render chapters/b19-transfer-across-centres-and-species.qmd --to html
```

To run only the R code in a fresh R session, start from the `chapters/` folder:

```
Rscript -e 'source(knitr::purl("b19-transfer-across-centres-and-species.qmd", output = tem
```

The chapter writes the synthetic spectra and the fold assignment to `data/b19-spectra.csv`.

19.3.1. Raw values and the recall model

The input is a table with one row per spectrum. Only the ID and the class are needed for the hand calculation, because the recall model does not use the spectrum.

```
# one row per spectrum: animal ID, class of the animal, spectrum number
animals <- data.frame(id = c("T1", "T2", "T3", "T4"), class = c("A", "B", "A", "B"))
spectra <- merge(animals, data.frame(spectrum = 1:3))
spectra <- spectra[order(spectra$id, spectra$spectrum), ]
rownames(spectra) <- NULL
spectra
```

```
#>    id class spectrum
#> 1  T1     A        1
#> 2  T1     A        2
#> 3  T1     A        3
#> 4  T2     B        1
#> 5  T2     B        2
#> 6  T2     B        3
#> 7  T3     A        1
```

```
#> 8  T3     A        2
#> 9  T3     A        3
#> 10 T4     B        1
#> 11 T4     B        2
#> 12 T4     B        3
```

```
# didactic recall model: known ID -> class recorded in training; unknown ID
fit_recall <- function(train, fallback = "A") {
  known <- unique(train[, c("id", "class")])
  function(test_id) {
    hit <- match(test_id, known$id)
    ifelse(is.na(hit), fallback, known$class[hit])
  }
}
```

19.3.2. Image-wise split: recognition

```
train_img <- spectra[spectra$spectrum %in% 1:2, ]   # two spectra per animal
test_img  <- spectra[spectra$spectrum == 3, ]       # third spectrum of every
pred_img  <- fit_recall(train_img)(test_img$id)

hand_img_correct <- sum(pred_img == test_img$class)
hand_img_shared  <- intersect(unique(train_img$id), unique(test_img$id))
hand_img_correct; hand_img_shared
```

```
#> [1] 4
```

```
#> [1] "T1" "T2" "T3" "T4"
```

```
stopifnot(hand_img_correct == 4, nrow(test_img) == 4)  # 4/4 right spectra
stopifnot(length(hand_img_shared) == 4)                # all four animals on both sides
```

19.3.3. Animal-wise split: new animals

```
train_ani <- spectra[spectra$id %in% c("T1", "T2"), ]   # one A and one B animal for trainin
test_ani  <- spectra[spectra$id %in% c("T3", "T4"), ]   # one new A and one new B animal
test_ani$pred <- fit_recall(train_ani)(test_ani$id)

hand_ani_spectra <- sum(test_ani$pred == test_ani$class)
# animal level: an animal counts as right if its spectra are classified as its class
per_animal <- aggregate(cbind(right = pred == class) ~ id, data = test_ani, FUN = all)
hand_ani_animals <- sum(per_animal$right)
hand_ani_shared  <- intersect(unique(train_ani$id), unique(test_ani$id))
per_animal; hand_ani_spectra; hand_ani_animals; hand_ani_shared
```

```
#>   id right
#> 1 T3  TRUE
#> 2 T4 FALSE
```

```
#> [1] 3
```

```
#> [1] 1
```

```
#> character(0)
```

```
stopifnot(hand_ani_spectra == 3, nrow(test_ani) == 6)   # 3/6 right spectra
stopifnot(hand_ani_animals == 1, nrow(per_animal) == 2) # 1/2 right animals; denominator = 2
stopifnot(length(hand_ani_shared) == 0)                 # empty ID intersection
```

The hand calculation gives 4/4 in the image-wise test and 3/6 spectra or 1/2 animals in the animal-wise test. These values belong to the recall model and are fixed by construction; they do not depend on any random number.

19.3.4. A split plan with zero ID overlap

A split plan is a small table that is written down before any model is fitted: one row per fold, the IDs on each side, and the number of shared IDs. For four animals and two classes, a two-fold group split keeps one A and one B animal on each side, and both folds together use every animal once as a test animal.

```
fold_of <- c(T1 = 2, T2 = 2, T3 = 1, T4 = 1)            # test fold of each a
split_plan <- do.call(rbind, lapply(1:2, function(f) {
  test_ids  <- names(fold_of)[fold_of == f]
  train_ids <- names(fold_of)[fold_of != f]
  data.frame(fold = f,
             train_animals = paste(train_ids, collapse = ", "),
             test_animals = paste(test_ids, collapse = ", "),
             test_spectra = sum(spectra$id %in% test_ids),
             shared_ids = length(intersect(train_ids, test_ids)))
}))
stopifnot(all(split_plan$shared_ids == 0),
          setequal(unlist(strsplit(split_plan$test_animals, ", ")), animals$i
knitr::kable(split_plan, col.names = c("Fold", "Training animals", "Test anim
```

Table 19.2.: Split plan for the animal-wise design (two folds, n = 4 animals). Shared IDs must be zero in every fold. Data status: synthetic.

Fold	Training animals	Test animals	Test spectra	Shared IDs
1	T1, T2	T3, T4	6	0
2	T3, T4	T1, T2	6	0

19.3.5. Self-contained variant: a nearest-neighbour model

The next block runs on its own in a fresh R session. It generates synthetic reflectance spectra for the same IDs T1 to T4 and classifies them with a one-nearest-neighbour model (`class::knn`, Venables and Ripley (2002); method background in Methods in R: k-Nearest Neighbours Classification). Each animal has its own level and slope of reflectance (an individual "signature"), and class B adds a weak dip in the near infrared. **These spectra are an assumption made for teaching: the animal signature is deliberately stronger than the class signal.** All results are stored with the prefix `knn_`, so they cannot be confused with the hand calculation.

```
library(class)
set.seed(4419)  # own seed: this block does not depend on earlier chunks

knn_make_spectra <- function(class_effect = 0.03, ids = c("T1", "T2", "T3", "T4"),
                             classes = c("A", "B", "A", "B"), n_per_animal = 3,
                             animal_sd = 0.05, noise_sd = 0.004) {
  wl <- seq(500, 940, by = 40)                                  # wavelength in nm
  base <- 0.30 + 0.30 / (1 + exp(-(wl - 690) / 25))          # common tissue curve (reflecta
  rows <- list()
  for (i in seq_along(ids)) {
    level <- rnorm(1, 0, animal_sd)                          # animal-specific offset
```

```
    slope <- rnorm(1, 0, animal_sd)                          # animal-specif
    dip <- if (classes[i] == "B") class_effect * exp(-((wl - 800) / 70)^2) el
    for (s in seq_len(n_per_animal)) {
      refl <- base + level + slope * (wl - 720) / 220 - dip + rnorm(length(wl
      rows[[length(rows) + 1]] <- data.frame(id = ids[i], class = classes[i]
                                             wavelength_nm = wl, reflectance
    }
  }
  do.call(rbind, rows)
}
knn_to_wide <- function(long) {
  w <- reshape(long, idvar = c("id", "class", "spectrum"), timevar = "wavelen
  rownames(w) <- NULL
  w
}
# 1-NN prediction; optional: columns missing at test time are replaced by tra
knn_predict <- function(wide, train, test, available = NULL) {
  X <- as.matrix(wide[, grep("^reflectance", names(wide))])
  if (!is.null(available)) {
    missing_cols <- setdiff(seq_len(ncol(X)), available)
    X[test, missing_cols] <- matrix(colMeans(X[train, missing_cols, drop = FA
                                    sum(test), length(missing_cols), byrow =
  }
  as.character(knn(X[train, , drop = FALSE], X[test, , drop = FALSE],
                   cl = factor(wide$class[train]), k = 1))
}
knn_long <- knn_make_spectra()
knn_wide <- knn_to_wide(knn_long)

# image-wise split: spectra 1-2 train, spectrum 3 test
img_train <- knn_wide$spectrum %in% 1:2
img_test  <- knn_wide$spectrum == 3
knn_img_pred <- knn_predict(knn_wide, img_train, img_test)
```

```
knn_img_correct <- sum(knn_img_pred == knn_wide$class[img_test])

# animal-wise two-fold group split (same folds as the split plan)
knn_fold <- c(T1 = 2, T2 = 2, T3 = 1, T4 = 1)[knn_wide$id]
knn_ani <- do.call(rbind, lapply(1:2, function(f) {
  tr <- knn_fold != f; te <- knn_fold == f
  stopifnot(length(intersect(unique(knn_wide$id[tr]), unique(knn_wide$id[te]))) == 0)
  data.frame(fold = f, id = knn_wide$id[te], class = knn_wide$class[te],
             pred = knn_predict(knn_wide, tr, te))
}))
knn_ani_spectra <- sum(knn_ani$pred == knn_ani$class)
# animal-level vote: the majority class of an animal's three spectra
knn_votes <- aggregate(pred ~ id + class, knn_ani, function(p) names(which.max(table(p))))
knn_ani_animals <- sum(knn_votes$pred == knn_votes$class)

knn_votes
```

```
#>   id class pred
#> 1 T1     A    B
#> 2 T3     A    B
#> 3 T2     B    B
#> 4 T4     B    B
```

```
c(image_wise = knn_img_correct, animal_wise_spectra = knn_ani_spectra, animal_wise_animals =
```

```
#>          image_wise animal_wise_spectra animal_wise_animals
#>                   4                   6                   2
```

```
stopifnot(knn_img_correct == 4)                          # recognition of known animals
stopifnot(nrow(knn_ani) == 12, nrow(knn_votes) == 4)     # every animal tested once; denominat
stopifnot(knn_ani_spectra == 6, knn_ani_animals == 2, all(knn_ani$pred == "B"))  # reference
```

The nearest-neighbour model gets 4 of 4 spectra right in the image-wise split, because the nearest spectrum is almost always another spectrum of the same animal. In the two-fold animal-wise split it gets 6 of 12 spectra and 2 of 4 animals right; in this draw it assigned class B to every test spectrum, so the hits on the two B animals say nothing about the class signal. Unlike the hand calculation, these values come from one random draw; the `stopifnot()` line records them as reference values for seed 4419,

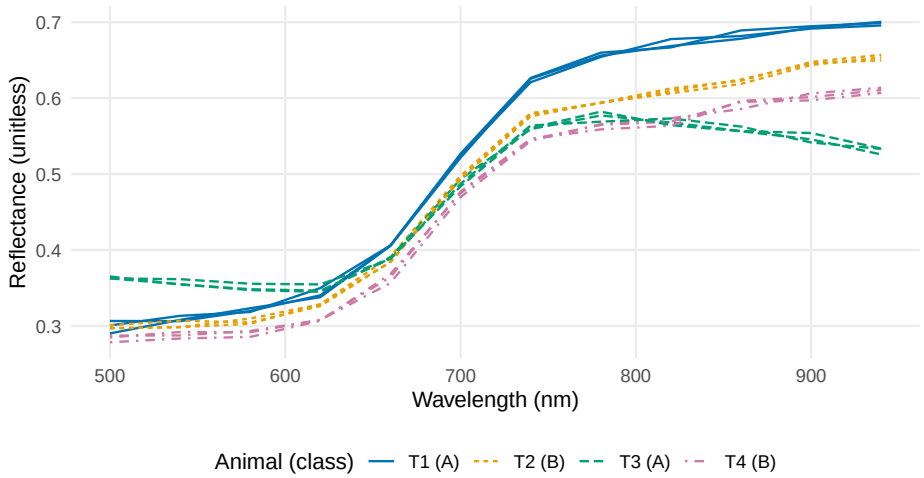


Figure 19.2.: Synthetic reflectance spectra used by the nearest-neighbour variant (reflectance is unitless, 0 to 1; wavelength in nm). Three spectra per animal, n = 4 animals. Each animal has its own level and tilt; class B has a weak dip around 800 nm. Data status: synthetic.

19.3.6. Variation between independent animals

One draw of four animals is a single experiment. To see how much an animal-wise result can vary, the next block draws 200 new synthetic cohorts of four animals with the same generator and repeats both tests. This is the step from "one number" to "the spread between independent sets of animals".

```
knn_repeat <- function(class_effect, R = 200, available = NULL) {
  res <- replicate(R, {
    w <- knn_to_wide(knn_make_spectra(class_effect = class_effect))
    tr_img <- w$spectrum %in% 1:2; te_img <- w$spectrum == 3
    img <- mean(knn_predict(w, tr_img, te_img, available) == w$class[te_img])
    fold <- c(T1 = 2, T2 = 2, T3 = 1, T4 = 1)[w$id]
    te <- do.call(rbind, lapply(1:2, function(f)
      data.frame(id = w$id[fold == f], class = w$class[fold == f],
                 pred = knn_predict(w, fold != f, fold == f, available))))
    votes <- aggregate(pred ~ id + class, te, function(p) names(which.max(table(p))))
    c(image_wise = img, animal_wise = mean(votes$pred == votes$class))
  })
  t(res)
}
set.seed(4420)
knn_weak <- knn_repeat(class_effect = 0.03)
knn_summary <- data.frame(
  test = c("Image-wise (spectra)", "Animal-wise (animals)"),
  mean = colMeans(knn_weak),
  p05 = apply(knn_weak, 2, quantile, 0.05),
  p95 = apply(knn_weak, 2, quantile, 0.95),
  share_cohorts_at_or_below_half = colMeans(knn_weak <= 0.5)
)
stopifnot(knn_summary$mean[1] > 0.95, knn_summary$mean[2] < 0.7)
knitr::kable(knn_summary, digits = 2, row.names = FALSE,
             col.names = c("Test", "Mean accuracy", "5th percentile", "95th percentile", "Co
```

Table 19.3.: Accuracy over 200 synthetic cohorts of 4 animals (weak class signal). Image-wise: share of right test spectra (4 per cohort). Animal-wise: share of right animals by majority vote (4 per cohort, two-fold group split). Data status: synthetic.

Test	Mean accuracy	5th percentile	95th percentile	Cohorts with accuracy 0.5
Image-wise (spectra)	1.00	1	1	0.00
Animal-wise (animals)	0.52	0	1	0.67

Over 200 cohorts the image-wise accuracy stays close to 1 (mean 1.00), while the animal-wise accuracy has a mean of 0.52 and ranges from 0.00 to 1.00 (5th to 95th percentile). In 67% of cohorts the animal-wise result is no better than guessing. The image-wise number is high and stable for the wrong reason.

```
# leave-one-animal-out on the same kind of cohorts: one fold per animal
set.seed(4422)
knn_loao <- replicate(200, {
  w <- knn_to_wide(knn_make_spectra(class_effect = 0.03))
  mean(sapply(unique(w$id), function(a) {
    p <- knn_predict(w, w$id != a, w$id == a)
    names(which.max(table(p))) == w$class[w$id == a][1]
  }))
})
round(mean(knn_loao), 2)
```

```
#> [1] 0.36
```

```
stopifnot(mean(knn_loao) < 0.5)
```

> **Common mistake**
>
> Leave-one-animal-out looks like the strictest split, but with four animals and two classes it leaves only one animal of the held-out class in training. A nearest-neighbour model is then pulled towards the other class. In the 200 synthetic cohorts above, its mean animal-level accuracy is 0.36, below the 0.5 expected from guessing. Check the class balance of every training fold in the split plan, and report the per-animal table instead of a single percentage.

19.3.7. Deepening: a missing modality as a stress test

Many hyperspectral systems combine a visible-range and a near-infrared sensor. Suppose the target centre has only the visible sensor (500–660 nm). A stress test asks what happens when the near-infrared part is missing at test time and is replaced by the training mean, which is a common default in software pipelines. To make the effect visible, the class signal is set strong (`class_effect = 0.3`), so that the model would work if both sensors were available.

```
set.seed(4421)
knn_strong <- knn_repeat(class_effect = 0.30)
set.seed(4421)
knn_strong_vis <- knn_repeat(class_effect = 0.30, available = 1:5)  # only 500-660 nm at tes
knn_stress <- rbind(both_sensors = colMeans(knn_strong), visible_only = colMeans(knn_strong_
round(knn_stress, 2)
```

```
#>              image_wise animal_wise
#> both_sensors       1.00        0.93
#> visible_only       0.63        0.47
```

```
stopifnot(knn_stress["visible_only", "animal_wise"] < knn_stress["both_senso
```

With both sensors, the animal-wise accuracy is 0.93; with the near-infrared part missing it drops to 0.47, because the class information sits in the part that is missing. The image-wise accuracy falls as well (to 0.63), since the replaced columns also carried part of each animal's signature. A stress test of this kind does not replace an external validation at the new centre. It shows before the study which missing input would make the claim collapse, and it forces the team to decide in advance what the system should do when an input is absent: stop, warn, or fall back to the standard assessment.

> **i** Other implementation: Python/scikit-learn
>
> **Original implementation.** `GroupKFold` in scikit-learn (Pedregosa et al. 2011), version 1.9.1 of the stable documentation (accessed 8 November 2026): GroupKFold. With `groups` set to the animal ID, each animal appears in exactly one test fold; `LeaveOneGroupOut` is the variant with one animal per fold.
> **Concrete advantage.** When an imaging model already lives in a Python pipeline with preprocessing and network weights, the group split can sit inside the same pipeline object, so every preprocessing step is fitted within the training fold only.
> **What must be identical for a fair comparison.** The input table (`data/b19-spectra.csv`), the group column (`id`), the fold assignment (import `fold_animalwise` instead of letting each language form its own folds: GroupKFold orders groups arbitrarily and, since version 1.6, can shuffle them), the metric unit (spectrum or animal), k = 1 with Euclidean distance, and the tie rule (`class::knn` breaks voting ties at random; with continuous spectra and k = 1, ties are practically absent). Tolerance: identical predicted classes.
> **Shared target value.** Zero shared IDs between training and test in every fold. The order of folds may differ.

Status. Conceptual comparison, not executed.

19.4. Counter-check

Change one assumption: make the class signal ten times stronger (`class_effect` from 0.03 to 0.30). Before you run the block, predict the direction for each test. Will the image-wise accuracy rise, fall or stay the same? And the animal-wise accuracy?

```
knn_compare <- rbind(weak_signal = colMeans(knn_weak), strong_signal = colMeans(knn_strong))
round(knn_compare, 2)
```

```
#>                image_wise animal_wise
#> weak_signal             1        0.52
#> strong_signal           1        0.93
```

```
stopifnot(knn_compare["strong_signal", "animal_wise"] > knn_compare["weak_signal", "animal_w
          abs(knn_compare["strong_signal", "image_wise"] - knn_compare["weak_signal", "image
```

The image-wise accuracy is practically unchanged (1.00 and 1.00), whereas the animal-wise accuracy rises from 0.52 to 0.93. The image-wise test gives the same answer whether or not there is anything to learn about the classes. A test that cannot tell these two situations apart cannot support a claim about new animals.

19.5. Consequences for patients and operations

Ethics and the 3Rs. An uninformative test design wastes animals even when few are used. Reduction under the 3Rs means minimising the number of animals consistent with the scientific aims, and the NC3Rs stress

that this depends on appropriately designed and analysed studies (NC3Rs 2026). Measuring many spectra per animal can maximise the information gathered per animal, which also counts as reduction, but it does not add independent animals; the experimental unit has to be fixed before the experiment. ARRIVE 2.0 asks authors to state the experimental unit (Essential 10, item 1b) and to discuss generalisability and translation (item 18) (Percie du Sert et al. 2020). For patients, the stake is indirect but real: a feasibility study justified by a recognition test asks participants for time, measurements and possibly images, on a premise that was never tested. The fair sentence in an application is the one that names the tested case.

Regulation. Not touched while the classifier is a research tool. It becomes relevant once a medical intended purpose is stated, for example "supports the assessment of tissue perfusion in patients"; then the questions of Chapter 22 apply, and the evidence must refer to the intended population and setting. A test on known animals is not clinical evidence for any intended purpose.

Data protection. Animal spectra are not personal data. The question changes as soon as patient images are collected for the next study; the data flow, the legal basis for health data and the deletion point are treated in Chapter 23.

Security and traceability. The claim "zero shared IDs" is only as good as the record of which animals were in which fold. Keep the split plan, the fold file and the model version together and under version control (Chapter 4, Chapter 24); a model trained later on "all available data" silently changes which animals are new.

19.6. Look up and transfer

- **Reviewer question:** Which unit is new in your test set, how many of these units are there, and is the result reported per unit? Please

provide the split plan with the number of shared IDs per fold.

- **Self-test:** A team measures six pigs (three of each class) in two laboratories: pigs 1–4 in laboratory 1, pigs 5–6 in laboratory 2, with 40 spectra per pig. A random 80/20 split of all 240 spectra gives 98% accuracy on 48 test spectra. (a) Which claim does this result support? (b) What is the independent denominator for a claim about new pigs?

  (c) Propose a split for a claim about a new laboratory and say what its denominator would be.

Answer

(a) Recognition of further spectra from the six known pigs: with 40 spectra per pig, almost every test spectrum has spectra of the same pig in training. (b) The number of pigs in the test folds; with a group split over all six pigs, at most six, and per fold far fewer. Reporting "48 test spectra" overstates the information. (c) Develop in laboratory 1 (pigs 1–4) and test in laboratory 2 (pigs 5–6). The denominator is two pigs, so the only possible results are 0/2, 1/2 and 2/2; check first that laboratory 2 contains both classes. The honest conclusion is that the data cannot support a claim about new laboratories, and that the next study needs more animals or a second external site, not more spectra per animal.

- **Annotated sources:**
  - Justice et al. (1999) distinguish reproducibility from transportability and give the hierarchy of historical, geographic, methodological and spectrum generalisability behind Table 19.1.
  - Debray et al. (2015) show how to interpret an external validation in light of case-mix differences between development and validation samples.

  - Roberts et al. (2017) explain why random cross-validation underestimates error in clustered data and how to block by group, time or space.
  - Kapoor and Narayanan (2023) survey data leakage across 17 fields and give a taxonomy of leak types.
  - Zech et al. (2018) document that an imaging model could recognise the hospital system of a radiograph, a measurement-procedure confounder.
  - Seok et al. (2013) and Takao and Miyakawa (2015) reach opposite conclusions about mouse–human similarity after trauma and inflammation from the same data, which illustrates that "species similarity" depends on analysis choices.
  - Percie du Sert et al. (2020) (ARRIVE 2.0) and NC3Rs (2026) (3Rs) frame the experimental unit and the justification of animal numbers.
  - Venables and Ripley (2002) is the reference for the `class` package and `knn()`; Pedregosa et al. (2011) for scikit-learn.

The chapter is required reading for W4, where learners split several recordings per animal image-wise and animal-wise and check the ID intersection in R. In W6 the split plan becomes the template for "the next test on another cohort, time or species", and in W8 groups defend which new case their model refers to.

19.7. Further reading

- Chapter 10 for the hyperspectral measurement chain and why pixel, image, animal and person are different units.
- Chapter 5 for baselines and the order of preprocessing, tuning and testing that the group split must respect.
- Chapter 20 for what happens after a successful external test: shadow operation, drift and monitoring.

- Chapter 27 for reporting the split plan and the tested new case in a model passport.
- Methods in R: Cross-Validation and Nested Cross-Validation for the general resampling schemes.
- Methods in R cheat sheet: External validation for a one-page checklist.
- CTTIR Courses, Course 4, Week 1 Session 1: Cross-validation, nested CV, bootstrap .632+ for a hands-on lab on resampling.
- Methods in R: ARRIVE reporting template for writing up the animal study.

20. Pilot operation, impact and monitoring

R. Heller · Date: 2026-07-28 · Last changed: 2026-11-09

At a glance

Guiding question: Which observation would stop the introduction?
For: clinical leads of pilot projects, physicians and nurses who act as human oversight, quality and patient safety officers, reviewers of early clinical evaluations · **Prerequisites:** what a validation did and did not test (Chapter 19), positive predictive value and thresholds (Chapter 6), alerts per shift (Chapter 9), reporting per shift (Chapter 17)
In the course: W7 · Regulation, data protection and a security incident and W8 · Model colloquium and decision (recommended reading; used for the restart decision) · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 11 minutes

20.1. Case and decision

It is Friday, 06:00, shift handover in the emergency department of a fictitious hospital. For 13 weeks the department has piloted a hint system that flags possible critical perfusion deficits in soft-tissue injuries from hyperspectral images; the system had passed an external validation in patients

(an *assumption* of this teaching case, following the path of Chapter 19). On Monday the camera software was updated. The night consultant, who is the named oversight person for the pilot during the night shift, hands over to the early-shift consultant: the night just ended produced 4 alerts among 6 examined patients, against an average of 0.79 alerts per night shift before the update. Two nurses have filed concerns about repeated false alarms. The question at the handover is not whether the model is good. It is *which observation, agreed in advance, means that the department stops using the hint now*, who decides, and what happens instead.

20.2. Core reading

Core reading · recommended for W7 and W8 · about 11 minutes · ends before 'Lab'

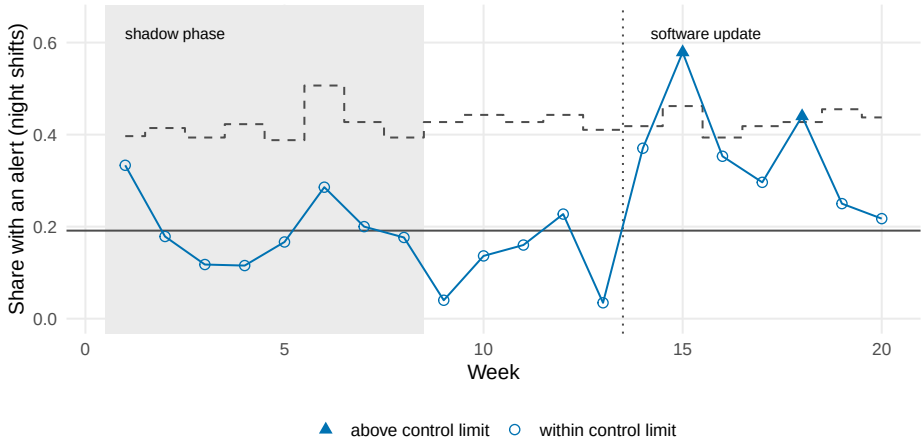
From a validated model to a supervised service. A successful external validation (Chapter 19) shows that a model can work in a new setting. It does not show that the department works better with it. Between the two lies a sequence of deliberately small steps, each with its own question. In *shadow operation* (silent mode), the system runs on real patients but its output is hidden from clinicians and compared with a reference afterwards: does it perform here, now, with our cameras and our case mix? In a *supervised pilot*, the output is shown, a named person oversees its use, and the department measures what changes in workflow, workload and safety. Only after that does a comparative study ask whether patients benefit. DECIDE-AI is the reporting guideline for this early clinical stage; it asks for the actual clinical performance at small scale, safety and the human factors of use (Vasey et al. 2022).

Human in the loop is a job, not a phrase. Oversight only exists if a named person has the *time* to look at each hint, the *competence* to judge it (including knowledge of the system's typical errors), and a *real possibility*

to override it without having to justify ignoring a machine. The AI Act describes the same elements for high-risk systems: oversight persons must be able to understand the system's limitations, remain aware of automation bias, disregard or override the output and stop the system (Art. 14(4)), and deployers must assign oversight to people with the necessary competence, training, authority and support (Art. 26(2)) (European Parliament and Council of the European Union 2024b). A review of automation bias in clinical decision support found workload and time pressure among its mediators (Goddard et al. 2012). An oversight person who also runs the resuscitation bay at night does not have the time the role requires.

Drift, update and alert burden. A deployed model meets a moving world. Case mix, documentation habits and devices change, and the clinician may not notice that the model's inputs have shifted (Finlayson et al. 2021). Calibration of prediction models can drift over years even when discrimination stays stable (Davis et al. 2017). A software update of a measuring device is a sudden version of the same problem. The first visible sign is often not a lower accuracy figure, which needs a reference standard and time, but a change in *how many* alerts appear. Alert burden matters in its own right: in a primary care study, clinicians accepted fewer alerts as they received more of them, especially repeated ones (Ancker et al. 2017). In an external validation of a widely used sepsis model, alerts were generated for 18% of hospitalised patients while two thirds of sepsis cases were not flagged (Wong et al. 2021); *published* figures that show how burden and missed cases can coexist.

A rule written before the pilot. Figure 20.1 is a control chart from statistical process control, a tool long used in health-care quality improvement (Benneyan et al. 2003). The centre line and the upper limit come from the shadow phase, before anyone saw an alert. A weekly chart is too slow for the handover in the case, though: it can only flag a completed week, and in these data it flags the night shift for the first time in week 15. The department therefore also agreed on a per-shift rule. Before the update, a night shift had on average 0.79 alerts; under a Poisson model, four or more alerts in one night would occur in about 0.9% of nights by



Solid line: shadow−phase mean; dashed: 3−sigma upper control limit for each week's denominator.

Figure 20.1.: Control chart (p-chart) of the weekly share of patients with an alert in night shifts. Denominator: patients examined in the 7 night shifts of each week (14 to 36 per week); n = 20 weeks. Centre line and 3-sigma upper control limit from the shadow phase (weeks 1 to 8). Grey area: shadow phase; dotted line: camera software update at the start of week 14. Data status: synthetic.

chance. The rule "four or more alerts in one night shift: suspend the hint for the rest of the shift, use the standard assessment, inform the pilot lead" was never triggered in the 13 weeks before the update and was triggered 2 times in the 49 nights after it.

The most important assumption. A stop rule protects patients only if it is written down before the pilot, with its denominator, its threshold, the person who acts on it and the fallback that replaces the system. A threshold chosen after seeing the data is an opinion, not a rule.

Who could carry the disadvantage? Three groups deserve explicit attention in this pilot. The *night shift*: fewer staff, dimmer light (the update changed image processing under low light in this teaching case) and therefore more false alerts per available person. The *nurses*, who perform the confirmatory checks that each alert triggers. And *patients with darker skin types*: optical measurements can perform differently across groups; pulse oximetry, for example, missed hidden hypoxaemia more often in Black than in White patients (11.7% versus 3.6% of measurements in one cohort; *published*) (Sjoding et al. 2020). In the synthetic pilot, darker skin types are about one in ten patients, so weekly numbers are tiny. The lab shows the dilemma this creates: small aggregates hide a real difference, and very small cells can reveal individuals.

Question for the workshop: Write the stop rule for your own project in one sentence: measure, denominator, threshold, person, fallback. Then name the group that would notice a problem last.

20.3. Lab

The lab simulates 20 weeks of a pilot: 8 weeks of shadow operation with a reference assessment for every patient, then 12 weeks of supervised use; the camera software update takes effect at the start of week 14. It then

builds the monitoring measures, checks small aggregates and ends with a decision exercise at the shift handover.

> **R start command**
>
> From the root of the book repository, render the chapter (this runs all code, including the figure of the core reading):
>
> ```
> quarto render chapters/b20-pilot-operation-monitoring.qmd --to html
> ```
>
> To run only the Lab code in a fresh R session, start from the `chapters/` folder:
>
> ```
> Rscript -e 'source(knitr::purl("b20-pilot-operation-monitoring.qmd", output
> ```

20.3.1. Step 1: simulate the pilot

All values are *assumptions* for teaching. Each shift sees a Poisson number of patients (about 6 early, 7 late, 4 at night). Of the patients, 12% have a critical perfusion deficit and 10% have darker skin types. The system's score is higher for patients with a deficit, but the signal is weaker for darker skin types; an alert is shown when the score exceeds 1.4. The update adds 0.9 to the score of every patient examined in a night shift from week 14 on.

```
set.seed(2003)
shifts <- expand.grid(shift = c("early", "late", "night"), day = 1:7, week =
                      stringsAsFactors = FALSE)                          # day
shifts$n <- rpois(nrow(shifts), c(early = 6, late = 7, night = 4)[shifts$shi
pts <- shifts[rep(seq_len(nrow(shifts)), shifts$n), c("week", "day", "shift"
rownames(pts) <- NULL
N <- nrow(pts)
```

```
pts$darker_skin <- runif(N) < 0.10                                  # skin types V-VI (ass
pts$deficit <- runif(N) < 0.12                                      # reference: deficit p
signal <- ifelse(pts$darker_skin, 0.9, 1.9)                         # weaker optical signa
update_effect <- ifelse(pts$week >= 14 & pts$shift == "night", 0.9, 0)
pts$score <- signal * pts$deficit + update_effect + rnorm(N)
pts$alert <- pts$score > 1.4
pts$phase <- ifelse(pts$week <= 8, "shadow", "supervised")
c(patients = N, per_week = round(N / 20), alerts = sum(pts$alert))
```

```
#> patients per_week   alerts
#>     2355      118      395
```

```
stopifnot(all(pts$alert <= 1), N == sum(shifts$n))
```

20.3.2. Step 2: a weekly control chart per shift type

```
b20_weekly <- aggregate(cbind(alerts = alert, patients = 1) ~ week + shift, data = pts, FUN
base <- aggregate(cbind(alerts, patients) ~ shift, data = b20_weekly[b20_weekly$week <= 8, ]
base$p_bar <- base$alerts / base$patients                           # centre line from the
b20_weekly <- merge(b20_weekly, base[, c("shift", "p_bar")], by = "shift")
b20_weekly$p <- b20_weekly$alerts / b20_weekly$patients
b20_weekly$ucl <- b20_weekly$p_bar + 3 * sqrt(b20_weekly$p_bar * (1 - b20_weekly$p_bar) / b2
b20_weekly$flag <- b20_weekly$p > b20_weekly$ucl
b20_weekly <- b20_weekly[order(b20_weekly$shift, b20_weekly$week), ]
b20_flags <- b20_weekly[b20_weekly$flag, c("shift", "week", "alerts", "patients")]
b20_flags
```

```
#>    shift week alerts patients
#> 55 night   15     11       19
#> 58 night   18     11       25
```

```
b20_first_flag <- min(b20_flags$week[b20_flags$shift == "night"])
stopifnot(all(b20_weekly$alerts <= b20_weekly$patients),
          !any(b20_weekly$flag[b20_weekly$week <= 13]),                # no i
          b20_first_flag == 15)                                        # firs
```

20.3.3. Step 3: the per-shift stop rule

```
per_shift <- aggregate(cbind(alerts = alert, patients = 1) ~ week + day + shi
# keep shifts without any patient (0 alerts): merge with the full list of shi
per_shift <- merge(shifts[, c("week", "day", "shift")], per_shift, all.x = TR
per_shift$alerts[is.na(per_shift$alerts)] <- 0
per_shift$patients[is.na(per_shift$patients)] <- 0
per_shift <- per_shift[order(per_shift$week, per_shift$day, match(per_shift$s
nights <- per_shift[per_shift$shift == "night", ]
b20_lambda_night <- mean(nights$alerts[nights$week <= 8])              # alert
b20_p_rule <- 1 - ppois(3, b20_lambda_night)                           # P(4 o
nights$trigger <- nights$alerts >= 4
b20_n_trig_after <- sum(nights$trigger[nights$week >= 14])
b20_n_nights_after <- sum(nights$week >= 14)
b20_trigger_night <- nights[nights$week >= 14 & nights$trigger, ][1, ]
nights[nights$week == 14, ]
```

```
#>     week day shift alerts patients trigger
#> 276   14   1 night      1        4   FALSE
#> 279   14   2 night      1        3   FALSE
#> 282   14   3 night      0        1   FALSE
#> 285   14   4 night      4        6    TRUE
#> 288   14   5 night      0        4   FALSE
#> 291   14   6 night      1        4   FALSE
#> 294   14   7 night      3        5   FALSE
```

```
stopifnot(b20_p_rule < 0.01,                                          # rare by chance befor
          !any(nights$trigger[nights$week <= 13]),
          b20_trigger_night$week == 14, b20_trigger_night$day == 4,   # first trigger: Thurs
          b20_n_nights_after == 49, nrow(nights) == 140)
```

The first trigger is the night shift starting on Thursday of week 14, which is the night in the case. The three earlier nights of that week had 1, 1, 0 alerts and looked unremarkable; a single night is a small sample, and the rule fired as soon as the agreed threshold was reached, not earlier.

20.3.4. Step 4: performance with honest denominators

In shadow operation every patient gets the reference assessment, so sensitivity (the share of patients with a deficit who received an alert) and positive predictive value (the share of alerts with a confirmed deficit) can both be estimated. In supervised use, the reference is only obtained for patients with an alert, so the sensitivity is no longer observable unless the department adds a reference check for a random sample of patients without an alert. The 95% intervals are exact (Clopper–Pearson) and computed with `qbeta()`.

```
exact_ci <- function(x, n) c(lower = if (x == 0) 0 else qbeta(0.025, x, n - x + 1),
                             upper = if (x == n) 1 else qbeta(0.975, x + 1, n - x))
estimate <- function(d, what) {
  if (what == "sensitivity") { x <- sum(d$alert & d$deficit); n <- sum(d$deficit) }
  else { x <- sum(d$alert & d$deficit); n <- sum(d$alert) }
  c(x = x, n = n, est = x / n, exact_ci(x, n))
}
shadow <- pts[pts$phase == "shadow", ]
b20_perf <- rbind(
  "Sensitivity, shadow, all"          = estimate(shadow, "sensitivity"),
  "Sensitivity, shadow, darker skin"  = estimate(shadow[shadow$darker_skin, ], "sensitivity
```

```
  "Sensitivity, shadow, other"          = estimate(shadow[!shadow$darker_skin
  "PPV, night, weeks 9-13"              = estimate(pts[pts$shift == "night" &
  "PPV, night, weeks 14-20"             = estimate(pts[pts$shift == "night" &
  "PPV, day shifts, weeks 9-13"         = estimate(pts[pts$shift != "night" &
  "PPV, day shifts, weeks 14-20"        = estimate(pts[pts$shift != "night" &
)
round(b20_perf, 2)
```

```
#>                                    x   n  est lower upper
#> Sensitivity, shadow, all          93 134 0.69  0.61  0.77
#> Sensitivity, shadow, darker skin   5  17 0.29  0.10  0.56
#> Sensitivity, shadow, other        88 117 0.75  0.66  0.83
#> PPV, night, weeks 9-13             8  14 0.57  0.29  0.82
#> PPV, night, weeks 14-20           14  62 0.23  0.13  0.35
#> PPV, day shifts, weeks 9-13       36  74 0.49  0.37  0.61
#> PPV, day shifts, weeks 14-20      43  90 0.48  0.37  0.59
```

```
stopifnot(b20_perf["PPV, night, weeks 14-20", "est"] < b20_perf["PPV, night,
```

After the update the positive predictive value at night falls from 57% to 23%, while in day shifts it is 49% before and 48% after the update. In the shadow phase, the sensitivity for patients with darker skin types is estimated at 5/17 (95% interval 10% to 56%), against 88/117 for other patients. The difference is visible, but the interval for the small group is wide.

20.3.5. Step 5: small aggregates and disclosure

Monitoring reports go to ward meetings, quality boards and sometimes to the manufacturer. A table of patients with a deficit by week and skin-type group has many cells with one, two or three patients. Such a cell,

combined with the date and shift, can point to an identifiable person for colleagues who were on duty. A common safeguard is to suppress small counts (Hundepool et al. 2012); here, counts from 1 to 4 are shown as "<5" (a local rule, an *assumption*).

```
suppress <- function(x, min_n = 5) ifelse(x > 0 & x < min_n, "<5", as.character(x))
weekly_def <- as.data.frame.matrix(xtabs(~ week + darker_skin, data = pts[pts$deficit, ]))
names(weekly_def) <- c("other", "darker_skin")
weekly_report <- data.frame(week = as.integer(rownames(weekly_def)),
                            other = suppress(weekly_def$other),
                            darker_skin = suppress(weekly_def$darker_skin))
head(weekly_report, 8)
```

```
#>   week other darker_skin
#> 1    1    19          <5
#> 2    2    12          <5
#> 3    3    16          <5
#> 4    4    12          <5
#> 5    5    14          <5
#> 6    6    15          <5
#> 7    7    16          <5
#> 8    8    13          <5
```

```
# information loss: in how many weeks can the darker-skin cell be reported at all?
b20_visible_weeks <- sum(weekly_def$darker_skin >= 5 | weekly_def$darker_skin == 0)
b20_visible_weeks_n5 <- sum(weekly_def$darker_skin >= 5)
stopifnot(!any(weekly_report$darker_skin %in% as.character(1:4)),   # no small count is prin
          !any(weekly_report$other %in% as.character(1:4)))
```

In none of the 20 weeks does the count of patients with darker skin types and a deficit reach five, so a weekly sensitivity for this group can never be reported. Over the 8-week shadow phase there are 17 such patients,

enough to report the estimate above with its wide interval. Two consequences follow. Subgroup performance has to be monitored on longer, pre-agreed windows (for example the whole shadow phase and then every 12 weeks), with the minimum group size stated in the monitoring sheet. And a visible gap in a small group is a reason for a targeted check, for example an additional reference assessment for all patients of that group for a set period, not for waiting until the weekly report becomes informative. Note also that zero cells and row totals can make a suppressed cell computable; a report that prints totals needs a second suppressed cell in the same row.

20.3.6. The monitoring sheet

The sheet below turns the measures into actions. It is written before the pilot and reviewed on fixed dates. Roles are fictitious; the thresholds are *assumptions* agreed by the pilot group.

Table 20.1.: Monitoring sheet for the supervised pilot (teaching example). The whole sheet is reviewed after 4 weeks, after every software change and at the end of the pilot.

Measure (denominator)	Rhythm	Trigger	Action and escalation	Responsible	Restart and review
Alerts per shift (per shift worked)	every shift	4 alerts in one night shift, or 6 in a day shift	Suspend the hint for the rest of the shift, use the standard assessment, inform the pilot lead by the next morning	Named oversight physician of the shift	Restart next shift if the pilot lead finds no systematic cause
Alert share per week and shift type (patients examined)	weekly	Point above the upper control limit of the shadow phase	Check software version and image logs; review a sample of alerts	Pilot lead with medical technology	Review the chart limits after 4 weeks

Measure (denominator)	Rhythm	Trigger	Action and escalation	Responsible	Restart and review
Positive predictive value (alerts)	every 4 weeks, per shift type	Upper 95% limit below 30%	Review threshold and image quality; inform the manufacturer	Pilot lead	Re-evaluate after the next 4 weeks
Sensitivity (patients with a deficit; reference for all alerts plus a random 10% of patients without an alert)	every 4 weeks	Estimate below 50% or lower 95% limit below 35%	Increase the reference sample; consider suspension	Pilot lead, quality and patient safety officer	Decision in the pilot board

Measure (denominator)	Rhythm	Trigger	Action and escalation	Responsible	Restart and review
Sensitivity in darker skin types (patients with a deficit in that group; minimum n = 10)	end of shadow phase, then every 12 weeks	Estimate more than 20 percentage points below the other group	Targeted reference check for all patients of the group for 4 weeks	Quality and patient safety officer	Decision in the pilot board; report without cells below 5
Override rate (alerts)	weekly	Change by more than half compared with the previous 4 weeks	Ask oversight staff why; check workload and training	Pilot lead with nursing lead	Review in the weekly pilot meeting
Software and model version (each start-up)	every start-up	Version differs from the released version	Block the hint automatically; fall back to the standard assessment	Medical technology / IT	Restart only with the released, verified version (Chapter 24)

Measure (denominator)	Rhythm	Trigger	Action and escalation	Responsible	Restart and review
Concerns and complaints from patients and staff (reports)	continuous, reviewed weekly	Any report of possible harm or a missed deficit	Case review within 72 hours; check whether a report to the manufacturer or authority is needed	Quality and patient safety officer	Documented in the pilot file
Exit criterion	any time	Two shift-rule triggers within 7 days without an explained and corrected cause, or one serious incident	Stop the pilot; standard assessment for all patients	Clinical director on the recommendation of the pilot board	Restart only after a new shadow phase of at least 2 weeks with all measures within limits

20.3.7. Decision exercise at the shift handover

Material. The case at the start of the chapter, Figure 20.1, the table of week 14 nights from step 3, the two nurse concerns and Table 20.1.

Roles. Night consultant (oversight person, handing over), early-shift consultant (taking over), nurse in charge, pilot lead (by phone), quality and patient safety officer (joins at 08:00). A sixth person may act as patient representative.

Task. Within ten minutes, decide what happens with the hint in the early shift and the coming night, who is informed, and what must be true before the hint is used at night again. Write the decision in three sentences that could go into the handover record.

Expected result. The shift rule has fired (step 3). The weekly chart cannot settle the question at this handover: week 14 is not complete, and in these data the chart flags the night shift only in week 15 (step 2). The hint is suspended at least for night shifts; the standard assessment applies. The pilot lead is informed now, not on Monday, and checks with medical technology whether the update explains the change. Because the night of Thursday is the first trigger, the exit criterion is not yet met; a second trigger within 7 days without a corrected cause would stop the pilot. The concerns of the nurses are recorded as reports and reviewed within 72 hours.

Alternative. Raise the alert threshold for night shifts until the alert rate returns to its old level. This is tempting because it keeps the system running, but it is a model change made under time pressure, without a reference standard, and it may reduce sensitivity most for the patients whose signal is already weakest. If it is done at all, it belongs in a new shadow phase, not in a handover.

Remaining question. Is the episode a reportable event? That depends on the regulatory status of the system and on whether a patient was harmed or put at risk; the answer is *open* in this case and belongs to the

quality and patient safety officer together with the manufacturer (Chapter 22).

20.4. Counter-check

Change the target population: the same system, unchanged, is introduced at a second site where 30% instead of 10% of patients have darker skin types. Predict before running: will the overall sensitivity rise, fall or stay the same? Will the subgroup be easier or harder to monitor?

```
set.seed(2004)
true_performance <- function(p_dark, n = 2e5) {   # large synthetic populatio
  dark <- runif(n) < p_dark
  deficit <- runif(n) < 0.12
  alert <- ifelse(dark, 0.9, 1.9) * deficit + rnorm(n) > 1.4
  c(sensitivity = mean(alert[deficit]), sens_darker = mean(alert[deficit & da
    sens_other = mean(alert[deficit & !dark]), ppv = mean(deficit[alert]))
}
b20_sites <- rbind(site_1 = true_performance(0.10), site_2 = true_performance
round(b20_sites, 3)
```

```
#>        sensitivity sens_darker sens_other   ppv
#> site_1       0.651       0.318      0.689 0.520
#> site_2       0.582       0.313      0.697 0.494
```

```
stopifnot(b20_sites["site_2", "sensitivity"] < b20_sites["site_1", "sensitivi
          abs(b20_sites["site_2", "sens_darker"] - b20_sites["site_1", "sens_
```

These are the values of the simulation's own truth, which no real department can observe. The overall sensitivity falls from 65% to 58%, although the model and its threshold are unchanged and the sensitivity within each

group stays the same (about 32% and 69%). A change of case mix alone moves an overall figure; the validation in Chapter 19 and the monitoring here must therefore report the groups that make up the average. At the second site the subgroup is three times larger, so its cells pass the minimum size more often, and the gap that was hard to see at the first site becomes visible sooner.

20.5. Consequences for patients and operations

Ethics. The pilot redistributes work and risk. Night staff carry more false alerts per person, nurses carry the confirmatory checks, and patients with darker skin types may carry a lower chance that a deficit is flagged. A pilot that reports only overall figures makes these groups invisible. Human oversight is ethically meaningful only if the oversight person can actually decline the hint; documenting overrides without blame, and asking why they occur, is part of respecting that role (Chapter 21).

Regulation. If the system has a medical intended purpose, it qualifies as a medical device, and the manufacturer must run a post-market surveillance system that gathers and analyses data on quality, performance and safety (MDR, Art. 83) and must report serious incidents (Art. 87) (European Parliament and Council of the European Union 2017a). A department that pilots a system that does not yet carry a CE marking may be conducting a clinical investigation, which has its own requirements; this is an *open* question to be checked with the regulatory affairs function (Chapter 22). An AI system that is itself a medical device, or a safety component of one, and that requires a third-party conformity assessment under the MDR is high-risk under Art. 6(1) and Annex I of the AI Act. For such systems the obligations of Chapter III, Sections 1 to 3, including human oversight (Art. 14) and the deployer obligations (Art. 26), apply from 2 August 2028 after the amendment by Regulation (EU) 2026/1744 (European Parliament and Council of the European Union 2026). Art. 26(5)

asks deployers to monitor operation and, where they have reason to consider that use according to the instructions may result in a risk, to inform the provider or distributor and the market surveillance authority without undue delay and to suspend use (European Parliament and Council of the European Union 2024b). Binding text: the regulations; whether the hint system in this case falls into these categories is a question for Chapter 22; the monitoring sheet is a local implementation. As of November 2026.

Data protection. Monitoring data are health data and are processed under the conditions of Art. 9 of the General Data Protection Regulation; the sheet should collect only what the measures need (data minimisation, Art. 5(1)(c)) (European Parliament and Council of the European Union 2016). Small cells in reports can identify patients; suppression, longer reporting windows and restricted access to detailed tables are the standard safeguards (Chapter 23, Chapter 26).

Security. The trigger in this case was a software update. Version checks at start-up, signed and verified releases, and a tested way back to the previous version are the controls that make the fallback real (Chapter 24). Feng et al. (2022) describe how monitoring and updating could be organised as a hospital quality improvement task.

20.6. Look up and transfer

- **Reviewer question:** For each monitored measure, what is the denominator, who sees it how often, which threshold was fixed before the pilot, and what is the fallback when the threshold is crossed? How is sensitivity measured once the output is shown to clinicians?
- **Self-test:** A pilot of a fracture hint (Chapter 9) reports “no safety concerns” after 8 weeks: of 412 alerts, 168 were confirmed fractures, and no missed fracture was reported by the wards. The reference was obtained only for alerted patients. What can and cannot be concluded, and what would you add to the monitoring?

Answer

The positive predictive value can be estimated (168/412, about 41%). Sensitivity cannot: missed fractures among patients without an alert are only found if someone looks for them, and "no reports" depends on whether wards know how and where to report. Add a reference assessment for a random sample of patients without an alert, a structured follow-up (for example recall of patients with persistent pain), a pre-specified threshold for action, and reporting by shift and by relevant patient group with a minimum cell size.

- **Annotated sources:**
  - Vasey et al. (2022) (DECIDE-AI) is the reporting guideline for early clinical evaluation of AI-based decision support, including human factors and safety.
  - Finlayson et al. (2021) describe dataset shift from the clinician's perspective; Davis et al. (2017) show calibration drift of prediction models over nine years.
  - Ancker et al. (2017) link repeated and frequent alerts to lower acceptance; Goddard et al. (2012) review automation bias and its mediators.
  - Wong et al. (2021) is an external validation in which a widely used model produced a high alert burden while missing most sepsis cases.
  - Benneyan et al. (2003) introduce control charts for health-care improvement; Feng et al. (2022) adapt quality-improvement tools to the monitoring of clinical AI.
  - Sjoding et al. (2020) document a measurement difference between patient groups for an optical device.
  - Hundepool et al. (2012) is a reference text on statistical disclosure control, including cell suppression.
  - European Parliament and Council of the European Union (2024b), European Parliament and Council of the European

Union (2026) and European Parliament and Council of the European Union (2017a) are the binding texts for oversight, deployer duties and post-market surveillance.

In W7 the chapter supports the incident table-top exercise and the decision about restart. In W8 groups use the monitoring sheet to state an override or stop condition for their own model.

20.7. Further reading

- Chapter 19 for what the validation behind a pilot did and did not test.
- Chapter 21 for the ethics of shifting burdens between patients and staff.
- Chapter 24 for version control, signed releases and recovery after a faulty update.
- Chapter 26 for small aggregates and responsible sharing of monitoring results.
- Methods in R: Changepoint Detection for detecting when a monitored series has shifted.
- Methods in R: Subgroup Analyses for the limits of conclusions from small groups.
- CTTIR Courses, Course 4, Week 4 Session 5: TRIPOD-AI, fairness auditing, reproducibility at scale for a lab on auditing model performance across groups.
- Strategy in R: Fairness in Machine Learning for formal fairness criteria and their trade-offs.

Part VI.

Part VI · Ethics, law and the protection of data and systems

21. Ethics at the bedside and in study design

R. Heller · Date: 2026-07-29 · Last changed: 2026-11-09

At a glance

Guiding question: Who bears the harm of a wrong decision, and who can object?
For: clinicians, nurses, researchers and reviewers who have to justify how an alert is used; the core reading needs no programming. · **Prerequisites:** the fracture alert and its predictive values in Chapter 9; absolute frequencies and thresholds from Chapter 6.
In the course: W2 · Language models and evidence (conversation card), W3 · Diagnostics, prediction and external evidence (errors by group), W8 · Colloquium and decision · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

21.1. Case and decision

It is 23:10 in the fictitious emergency department of Chapter 9, setting A. An 84-year-old patient who lives alone has fallen onto her outstretched hand. The junior doctor sees no fracture on the wrist radiograph, but the image-analysis tool shows “possible fracture – second look recommended”. The radiology report will not be available before 08:30. Three people look at the same alert. The patient wants to know whether she can sleep in

her own bed. The nurse in charge counts free trolleys and waiting alerts. A researcher who evaluates the tool asks whether alerts in patients over 80 are wrong more often than in younger patients. The doctor has to act now: keep the patient overnight, send her home with a splint and a morning call-back, or send her home without either. The patient carries most of the consequence, so she must understand the options and be able to say no.

21.2. Core reading

Core reading · recommended reading for W8, used in W2 and W3 · about 10 minutes · ends before 'Decision exercise'

From a predictive value to a distribution of harm. In setting A of Chapter 9, 5% of the patients examined have a fracture, and the tool's sensitivity and specificity are both 90% (*assumptions*). Of 1,000 patients, 140 receive an alert; 95 of them have no fracture, and 5 patients with a fracture receive no alert. Ethics begins where these counts meet people. A false alarm can mean an unnecessary splint, an extra radiograph, a night on a trolley or a delayed discharge. A missed fracture can mean a scaphoid that is not immobilised, weeks of pain and, later, an operation. The two errors do not fall on the same people, and they do not weigh the same for everyone. For a 30-year-old, a night in the department is an inconvenience. For an 84-year-old living alone, it brings a risk of delirium and falls. The question is therefore not only "how often is the tool wrong?" but "who carries each kind of error, and could they have chosen differently?"

Several lenses, one alert. Clinical ethics usually asks four questions: is the person's autonomy respected, does the action benefit them, does it avoid harm, and are benefits and burdens shared fairly? The WHO guidance on AI for health adds that humans should remain in control of medical decisions and that people should understand the role an AI

system plays in their care (World Health Organization 2024, Box 1, p. 5). It also names automation bias, in which a clinician overlooks an error that a human should have spotted, and the gradual loss of skills that makes it harder to overrule an algorithm (World Health Organization 2024, 11–12). A further lens is professional responsibility: the doctor remains accountable for the decision, whatever the box says. Table 21.1 applies these lenses to the three people in the case.

Table 21.1.: Three perspectives on the same alert. Fictional roles; the table is a structure for discussion, not a survey result.

	Patient, 84, lives alone	Nurse in charge	Researcher evaluating the tool
What matters most	Sleeping at home; not being sent away with a fracture	Safe care for everyone in the department tonight	A valid estimate of how the tool behaves in this population
Harm they see	Night on a trolley, delirium, a fall; or a missed fracture	Alerts that bind staff and trolleys; alert fatigue	Errors concentrated in a group that the overall figures hide
What they need	Plain numbers, real options, a contact person	A named person who reviews alerts, within planned time	Prespecified subgroups, counts with denominators, intervals
What they can object to	Staying overnight; any of the options offered	Being made responsible for reviews without time	A deployment plan that ignores an unexplained subgroup signal

Errors by group need denominators. Figure 21.1 splits the 1,000 patients of setting A into four groups. **The key assumption: the group counts are *synthetic* and were chosen so that the totals match setting A exactly; they show what a local evaluation would have to measure, not what any real tool does.** Patients aged 80 and over make up 12% of those examined but receive 26 of the 95 false alarms (27.4%). Their false-alarm rate is 24.1% (26/108; 95% interval 17.0% to 32.9%), compared with 7.1% (38/538; 5.2% to 9.5%) in patients aged 18–64. Patients with a plaster cast or metal implant are a different case. Only 2 of them have a fracture, and the tool flags 1. A sensitivity of 50% with an interval from 9.5% to 90.5% means: we do not know. The intervals are Wilson score intervals, which behave well with small counts (Wilson 1927; Brown et al. 2001); the general idea is explained in Methods in R: Confidence intervals.

A group finding like this is a reason to look closer, not yet a proven property of the tool. Published examples show why the look is necessary. Chest radiograph classifiers underdiagnosed under-served patient groups, most clearly in intersectional subgroups such as Hispanic female patients [Seyyed-Kalantari et al. (2021); *published*]. A widely used population-health algorithm predicted health-care costs rather than illness, so that Black patients with the same score were sicker than White patients [Obermeyer et al. (2019); *published*]. In both cases a good overall figure hid a systematic burden on one group.

"Human in the loop" is a job description. The phrase protects patients only if it names someone. Four conditions make it real. First, a named role on the rota, not "the team". Second, time: in setting A the tool creates about 1.6 alerts per shift (Chapter 9), and their review must be part of the planned workload. Third, competence: the reviewer knows the tool's typical failure modes, including the groups in Figure 21.1. Fourth, a real possibility to override: disregarding the alert must be as easy as following it and must not require a written justification each time. For high-risk AI systems the EU AI Act requires deployers to assign human oversight to people with the necessary competence, training and authority,

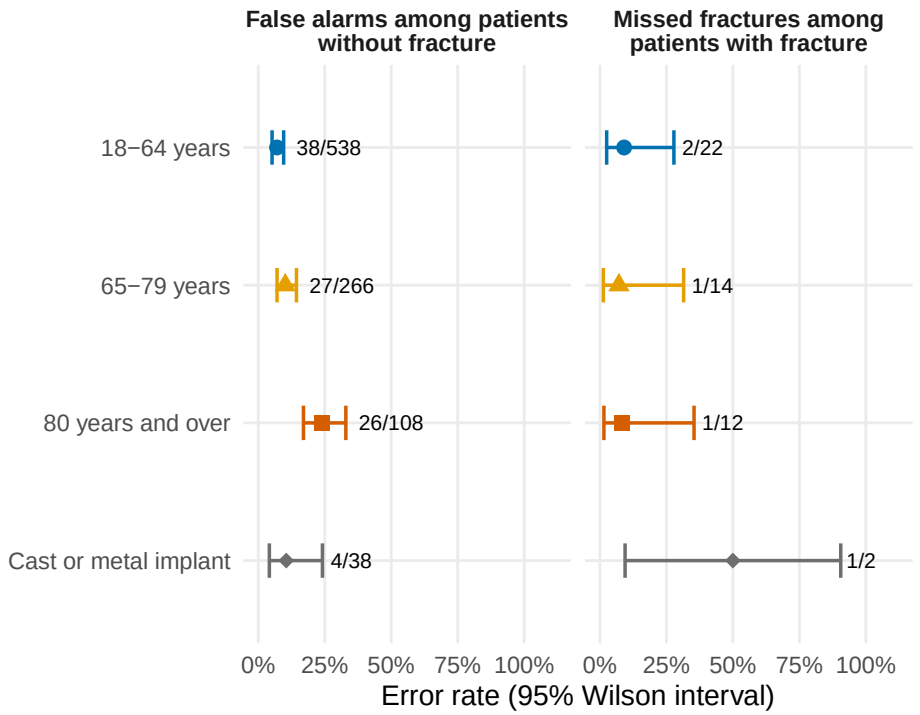


Figure 21.1.: Error rates of the fracture alert by patient group in setting A, with 95% Wilson intervals. Left: false alarms among patients without fracture; right: missed fractures among patients with fracture. Labels give numerator/denominator. Synthetic counts that add up to the 1,000 patients of setting A in B09; independent unit: patient (n = 1,000; denominators per panel and group as labelled).

and to support them (Art. 26(2)); the system must let them stay aware of automation bias and disregard or override its output (Art. 14(4)(b) and (d)) (European Parliament and Council of the European Union 2024b). For AI that is part of a medical device needing a notified body, these duties apply from 2 August 2028 (Chapter 22). Ethically, the four conditions apply from the first night. A systematic review found workload, task complexity and time pressure among the factors that increase automation bias (Goddard et al. 2012).

A real choice, checked with teach-back. Shared decision making moves from introducing the choice to describing options to helping the person decide (Elwyn et al. 2012). A choice is only real if the organisation has created the options: the splint is in stock, the morning call-back has a named slot, and someone can take the patient home at midnight. Teach-back closes the loop: the clinician asks the patient to explain in her own words what the alert means and what she will do. It tests the explanation, not the patient. In an observational study of patients with diabetes and limited health literacy, physicians checked understanding in only 12 of 61 visits; where they did, glycaemic control was better, with very wide confidence intervals [Schillinger et al. (2003); *published*].

Question for the workshop: The night team can offer only two of the three options from the case. Which option would you keep at all costs, and which group of patients would lose most if it disappeared?

21.3. Decision exercise

The exercise has two parts: a short R calculation that makes the group errors checkable, and a conversation in which the result has to be explained to the patient.

Part 1 – recompute the group errors. The table below is the input of Figure 21.1. The Wilson interval is written out as a small function and

checked against `prop.test()`, whose interval without continuity correction is the same score interval.

```
# Synthetic counts for setting A (n = 1,000 patients)
groups[, c("group", "examined", "fracture", "true_pos", "false_pos")]
```

```
#>                   group examined fracture true_pos false_pos
#> 1           18-64 years      560       22       20        38
#> 2           65-79 years      280       14       13        27
#> 3     80 years and over      120       12       11        26
#> 4 Cast or metal implant       40        2        1         4
```

```
# Check 1: the groups reproduce the totals of B09 (45/140 alerts with fracture)
stopifnot(sum(groups$true_pos) / (sum(groups$true_pos) + sum(groups$false_pos)) == 45 / 140)

# Check 2: own Wilson function equals prop.test() without continuity correction
for (i in seq_len(nrow(groups))) {
  # prop.test() warns about its chi-squared p-value for small counts; the interval itself is
  ref <- suppressWarnings(prop.test(groups$false_pos[i], groups$no_fracture[i], correct = FA
  stopifnot(isTRUE(all.equal(unname(fa[i, ]), as.numeric(ref), tolerance = 1e-10)))
}

# Result table: rates with intervals, per group
result <- data.frame(
  group            = groups$group,
  false_alarms     = sprintf("%d/%d", groups$false_pos, groups$no_fracture),
  false_alarm_rate = cdsr_percent(fa_rate),
  fa_95_interval   = apply(fa, 1, fmt_ci),
  ppv              = cdsr_percent(groups$true_pos / (groups$true_pos + groups$false_pos)),
  share_of_false_alarms = cdsr_percent(groups$false_pos / sum(groups$false_pos))
)
result
```

```
#>                   group false_alarms false_alarm_rate fa_95_interval   ppv
#> 1           18-64 years       38/538             7.1%   5.2% to 9.5% 34.5%
#> 2           65-79 years       27/266            10.2%  7.1% to 14.4% 32.5%
#> 3     80 years and over       26/108            24.1% 17.0% to 32.9% 29.7%
#> 4 Cast or metal implant         4/38            10.5%  4.2% to 24.1% 20.0%
#>   share_of_false_alarms
#> 1                 40.0%
#> 2                 28.4%
#> 3                 27.4%
#> 4                  4.2%
```

```
# Check 3: the intervals of the oldest and the youngest group do not overlap
stopifnot(fa[old, "lower"] > fa[young, "upper"])
# Check 4: the cast group's sensitivity interval is wider than 0.75 - too few
stopifnot(diff(sens_ci[cast, ]) > 0.75)
```

The calculation shows two different findings. In the oldest group the false-alarm rate is higher, and the interval excludes the rate of the youngest group. In the cast group nothing can be concluded about sensitivity, because two fractures are not a denominator. Both findings change the conversation with the patient: her group has a lower predictive value (29.7%) than the department as a whole (32.1%).

Part 2 – the conversation. Work in threes: doctor, patient, observer. The doctor has five minutes. The script below is a starting point, not a text to read aloud.

> **i** Conversation card: alert, options, teach-back
>
> 1. **Name the situation.** “The computer program that looks at X-rays has marked your wrist. That does not mean it is broken. It means a specialist should look again.”
> 2. **Give the number in people.** “In patients of your age in our

department, about 3 in 10 of these marks turn out to be a real fracture." (In a real department this figure must come from local data; here it is the synthetic 29.7%.)
3. **Offer the real options.** "You can stay here tonight and we will have the answer in the morning. Or you can go home with a removable splint; we will phone you before 10:00 with the result and tell you whether to come back. If the pain or swelling gets worse, you come back at once."
4. **Ask what matters.** "What is most important to you tonight?"
5. **Teach-back.** "I want to make sure I explained this well. Could you tell me in your own words what the mark means and what you will do if the pain gets worse?"
6. **Document** the choice, the reason the patient gave and who will make the call-back.

The observer notes whether the doctor gave the number in people, whether both options were presented with their downsides, whether the patient's answer in step 5 was correct and whether the doctor corrected the explanation, not the patient.

Expected result. A documented decision that the patient can repeat in her own words, a call-back owner with a time, and a note of the predictive value used for her group, marked as local or synthetic.
Alternative. The doctor calls the on-call senior to review the image now. This avoids the night-time uncertainty but uses scarce senior time; it is the right choice if the department has planned it as part of the reviewer's role.
Residual question. Who decides whether the call-back option is offered at all, and what happens to patients who cannot be reached by phone?

21.4. Counter-check

Change the target population before you calculate. Suppose the same group-specific error rates apply in setting B of Chapter 9, where 20% of patients have a fracture. *Predict first:* does the predictive value in the oldest group rise or fall, and does the oldest group's share of false alarms change?

```
# Setting B: same group sizes, more fractures (synthetic, 200 in total)
fracture_b <- c(90, 60, 40, 10)
stopifnot(sum(fracture_b) == 200)
tp_b <- fracture_b * sens                                  # expected true
fp_b <- (groups$examined - fracture_b) * fa_rate           # expected fals
data.frame(group = groups$group,
           ppv_setting_A = cdsr_percent(groups$true_pos / (groups$true_pos +
           ppv_setting_B = cdsr_percent(tp_b / (tp_b + fp_b)),
           share_false_alarms_B = cdsr_percent(fp_b / sum(fp_b)))
```

```
#>                   group ppv_setting_A ppv_setting_B share_false_alarms_B
#> 1           18-64 years         34.5%         71.1%                42.6%
#> 2           65-79 years         32.5%         71.4%                28.6%
#> 3     80 years and over         29.7%         65.6%                24.7%
#> 4 Cast or metal implant         20.0%         61.3%                 4.1%
```

```
# Check: PPV in the oldest group rises with prevalence
stopifnot(tp_b[old] / (tp_b[old] + fp_b[old]) > groups$true_pos[old] / (grou

# Second counter-check: ten times as many patients with cast, same rates
mis_big <- wilson(10 * groups$false_neg[cast], 10 * groups$fracture[cast])
c(width_n2 = diff(mis[cast, ]), width_n20 = diff(mis_big[1, ]))
```

```
#>  width_n2.upper width_n20.upper
#>       0.8109376       0.4014040
```

```
stopifnot(diff(mis_big[1, ]) < diff(mis[cast, ]))
```

In setting B the predictive value rises in every group, because a larger share of patients has a fracture. The oldest group still carries a disproportionate share of the false alarms. Prevalence changes the predictive value, not the group's error rate. The second check shows that a larger denominator narrows the interval but leaves the point estimate unchanged; it would take many more patients with casts before the miss rate could be called high or low. A practical consequence: the overall specificity of 90% in Chapter 9 is itself an average over a case mix. If the case mix shifts towards older patients, the department's overall false-alarm rate rises even if the tool does not change.

21.5. Consequences for patients and operations

An ethically questionable recommendation changes the plan. Suppose the supplier proposes to sort the radiology worklist by alert status: flagged images are reported first, unflagged images up to 24 hours later. On average this looks efficient. In setting A, however, 5 of the 50 fractures are unflagged and would wait longest, and in the cast group the tool missed one of two fractures – nobody knows how often it misses more. The burden of delay would fall on a group for which the tool is least reliable and whose patients cannot know that their image was put at the back of the queue. Table 21.2 shows how the deployment plan changes.

Table 21.2.: Deployment plan for the fracture alert before and after the ethical review. Fictional plan; the elements are examples.

Element	Plan before the review	Plan after the review
Worklist order	Sorted by alert status	Clinical order unchanged; the alert only adds a second-look flag
Human oversight	"Clinician in the loop"	Named on-call senior per shift; review time planned per alert; override without justification form
Cast or metal implant	Included	Excluded from the alert until a local evaluation with enough fractures exists; always senior review
Patients 80 and over	No special rule	Group-specific predictive value reported monthly with interval; call-back option guaranteed
Conversation	Information leaflet	Conversation card with teach-back; documented choice and call-back owner

Regulation and data protection. Excluding a patient group or changing what the alert is used for may change the intended purpose of the product; Chapter 22 shows why this needs a dated note before go-live. Group analyses require processing age and possibly further characteristics.

Ethnicity or disability would be special categories of personal data under the GDPR; the amended AI Act now sets strict conditions under which providers of high-risk systems may exceptionally process such data to detect and correct bias (Art. 4a, inserted by Regulation (EU) 2026/1744, in force since 27 July 2026) (European Parliament and Council of the European Union 2026). Chapter 23 covers the data-protection side.

Security and availability. If the tool fails, the plan must still give every patient a timely report. This fallback is the regular reporting pathway, tested as described in Chapter 24.

> **i** The incident from the ethics perspective
>
> In the fictional incident of Part VI, an excerpt from a radiology report appears in a service provider's support log, and the installed model version differs from the approved one. From an ethics perspective the first questions are: which patients are affected, who among them is particularly burdened (for example by a sensitive diagnosis in the excerpt), and who explains the situation to them in plain language? A good answer names a reachable contact person, offers to answer questions by phone and in writing, and does not wait for the end of the technical analysis to decide whether patients deserve an explanation. Whether a formal notification is legally required is a separate question (Chapter 23).

21.6. Look up and transfer

Reviewer question. Were the subgroups prespecified, and does the report give numerator, denominator and an interval for each error rate in each subgroup? A sentence such as "performance was similar across age groups" without counts cannot be checked (Methods in R: Subgroup analyses).

Self-test. A paediatric emergency department wants to use the same tool for children aged 6–15. Growth plates can look like fracture lines. A local check finds 18 false alarms among 60 children without fracture and 3 missed fractures among 15 children with fracture. The supplier says: "Overall accuracy in children is 72%." What would you report to the department, and which ethical question would you raise before go-live?

Answer

Report the two errors separately with their denominators and intervals: false alarms 18/60 = 30.0% (Wilson 95% interval about 19.9% to 42.5%) and missed fractures 3/15 = 20.0% (about 7.0% to 45.2%). Overall accuracy mixes the two and hides that each child without fracture has roughly a one-in-three chance of a false alarm. The ethical question is who bears the burden: false alarms in children mean extra radiation and immobilisation; misses mean untreated fractures in growing bones. Before go-live, the department needs a named reviewer for paediatric alerts, an explanation for parents and a decision on whether the evidence (15 fractures) is enough at all. Children were not part of the intended population in the setting A data; using the tool for them is a change of intended purpose (Chapter 22).

Annotated sources.

- World Health Organization (2024) – the six WHO principles (Box 1) and the description of automation bias and skills loss; used here for autonomy and human control.
- European Parliament and Council of the European Union (2024b) – Art. 14(4) and Art. 26(2) translate human oversight into duties of providers and deployers; used for the four conditions of a real "human in the loop".
- European Parliament and Council of the European Union (2026) – inserts Art. 4a on processing special categories of data for bias detection; relevant to the design of subgroup analyses.

- Wilson (1927) and Brown et al. (2001) – the original score interval and the comparison showing why simple intervals for proportions fail with small counts.
- Seyyed-Kalantari et al. (2021) and Obermeyer et al. (2019) – published cases in which good overall performance hid systematic harm to specific groups.
- Goddard et al. (2012) – systematic review of automation bias and its mediators, including workload and time pressure.
- Elwyn et al. (2012) and Schillinger et al. (2003) – a practical model of shared decision making and an observational study of checking patients' understanding.

21.7. Further reading

- Chapter 9 – the alert, the two settings and the predictive values on which this chapter builds.
- Chapter 6 – thresholds and absolute frequencies, to weigh false alarms against missed fractures.
- Chapter 20 – how group-specific error rates enter a monitoring plan with escalation and stop criteria.
- Chapter 22 and Chapter 23 – when a change in the patient group changes the intended purpose, and how group data may be processed.
- Scientometrics in R: Ethics and responsible metrics – why a single indicator should not decide about people; transferable to alert metrics.
- Scientometrics in R: Gender and equity analysis – pitfalls when comparing groups, including small groups and inferred attributes.
- CTTIR Courses: TRIPOD-AI, fairness auditing, reproducibility at scale – a lab that audits model errors by group.
- Strategy in R: Fairness in machine learning – formal fairness criteria and why they cannot all hold at once.

22. Regulating research software and clinical AI

R. Heller · Date: 2026-10-25 · Last changed: 2026-11-09

At a glance

Guiding question: Which change in the intended purpose changes the regulatory question?
For: research teams, clinical IT and regulatory staff, clinicians who are asked to "just try" a tool, and reviewers; the core reading needs no programming. · **Prerequisites:** the fracture alert in Chapter 9; pilot operation and monitoring in Chapter 20; types of evidence in Chapter 7.
In the course: W7 · Regulation, data protection and incident (regulatory note), W8 · Colloquium and decision · **Data:** no patient data; fictional documents; legal status checked on 8 November 2026 · **Core reading time:** about 11 minutes

22.1. Case and decision

The research group behind the fracture alert of Chapter 9 has finished a retrospective evaluation on archived radiographs. The results look good, and the head of the emergency department wants the tool switched on for night shifts from next Monday, "only as a hint". The hospital's lead for clinical IT has to sign the go-live. She asks three questions first. What

exactly does the tool claim to do, in writing? Who is the manufacturer and who is the user in the legal sense? Which rules apply, and from which date? Nobody can show her a written intended purpose, and the version on the test server has no release record. She postpones the go-live by two weeks and asks the group for a dated classification note. Night shifts continue without the tool. The alternative – a start without a documented basis – would make the hospital responsible for a product nobody has classified.

22.2. Core reading

Core reading · required for W7 · about 11 minutes · ends before 'Decision exercise'

Four kinds of statement. Every regulatory note in this book separates four kinds of statement, and so should yours. [**Law**] is the binding text: the EU AI Act (Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744), the Medical Device Regulation (MDR, Regulation (EU) 2017/745) and the In Vitro Diagnostic Medical Device Regulation (IVDR, Regulation (EU) 2017/746) (European Parliament and Council of the European Union 2024b, 2026, 2017a, 2017b). [**Guidance**] is orientation without legal force, such as MDCG 2019-11 Rev. 1 on software, which states itself that only the Court of Justice can give binding interpretations (Medical Device Coordination Group 2025). [**Interpretation**] is the professional reading offered in this chapter. [**Local**] marks a question that only the responsible people in a specific institution can decide, such as the hospital's regulatory affairs office or legal department. A note that blurs these layers cannot be checked. All statements refer to the legal status of 8 November 2026.

The intended purpose is the hinge. The MDR defines it by what the manufacturer states on the label, in the instructions for use, in promotional material and in the clinical evaluation (Art. 2(12)); the AI Act's definition

adds the specific context and conditions of use (Art. 3(12)) [Law]. MDCG 2019-11 Rev. 1 adds that software needs a medical purpose of its own to qualify as medical device software, and that the stated intended purpose drives qualification and classification (Section 3.1) [Guidance]. The same code can therefore be research software, a medical device or an in vitro diagnostic device, depending on what is claimed for it. **The key assumption of this chapter: classification follows the intended purpose as written and communicated, not how cautiously a team plans to use the tool.** Table 22.1 contrasts three fictional intended purposes.

Table 22.1.: Three intended purposes for related software. Fictional vignettes; legal status 8 November 2026. The bracketed labels mark binding text, guidance, interpretation and local questions.

	A · Research prototype	B · Clinical decision support	C · IVD context
Intended purpose (fictional)	Estimates sensitivity and specificity of a fracture model on archived, de-identified radiographs; outputs never reach clinicians	Marks adult wrist and hand radiographs in the emergency department to prompt a second look; every image is still reported by a radiologist	Combines routine laboratory values of inflammatory markers to estimate the risk of fracture-related infection after surgery, for review by the treating team

	A · Research prototype	B · Clinical decision support	C · IVD context
Medical purpose for individual patients?	No [Interpretation]	Yes: computer-aided detection is qualified as medical device software [Guidance, MDCG Annex I b]	Yes, based on data from in vitro diagnostic devices only [Guidance, MDCG Section 3.4, decision steps 1–2]
Regulation	MDR not applicable without a medical purpose [Interpretation]; AI Act excludes research before placing on the market (Art. 2(6), 2(8)) [Law]	MDR; classification by Rule 11 (Annex VIII) [Law]	IVDR; software independent of a device is classified in its own right (Annex VIII, rule 1.4) [Law]
Case information needed for classification	Proof that no output is used for any patient; who receives results	Seriousness of the clinical situation; whether the output informs, drives or replaces a decision; consequence of a wrong output; users; excluded groups	Which markers; which decision follows (antibiotics, revision surgery); applicable IVDR rule (class B if no other rule applies, Rule 6) [Law]

	A · Research prototype	B · Clinical decision support	C · IVD context
AI Act route	Not in scope while research [Law]; testing in real-world conditions is not excluded (Art. 2(8))	Class IIa or higher needs a notified body (MDR Art. 52) → high-risk under Art. 6(1) and Annex I [Law]	Classes B–D need a notified body (IVDR Art. 48) → high-risk under Art. 6(1) [Law]
Open point	Does a silent pilot in live care count as testing in real-world conditions (Art. 3(57))? [Local]	Is the situation "serious" and the output "informing" (class IIa) or "driving" management (class IIb)? [Interpretation, MDCG Annex III]	Which IVDR rule applies? [Local; MDCG 2020-16 would be the next guidance to read]

Rule 11 in one paragraph. Software intended to provide information used to take decisions with diagnostic or therapeutic purposes is class IIa, unless such decisions may cause a serious deterioration of health or a surgical intervention (class IIb) or death or an irreversible deterioration (class III); all other software is class I (MDR, Annex VIII, Rule 11) [Law]. MDCG 2019-11 Rev. 1 offers a table that combines the state of the healthcare situation (critical, serious, non-serious) with the significance of the information (treat or diagnose, drive clinical management, inform clinical management); in a serious situation, "inform" maps to class IIa and "treat or diagnose" to class IIb (Annex III) [Guidance]. For vignette B the deci-

sive facts are therefore clinical: what happens to a patient whose fracture is missed, and does anyone act on the marker alone?

From device class to AI Act. An AI system is high-risk under Art. 6(1) if it is, or is a safety component of, a product covered by the legislation in Annex I and that product must undergo a third-party conformity assessment [Law]. MDR and IVDR are listed in Annex I, Section A (points 11 and 12) [Law]. Class I medical device software is normally self-declared (MDR Art. 52(7)), so it does not meet the second condition; class IIa and higher does. For high-risk systems under Art. 6(1), the requirements of Chapter III, Sections 1–3, apply from 2 August 2028; for Annex III systems from 2 December 2027 (Art. 113, as amended by Regulation (EU) 2026/1744) [Law]. Some duties apply to every AI system already: providers and deployers must take measures to support the AI literacy of their staff (Art. 4, replaced by the 2026 amendment) [Law]. Figure 22.1 places these dates on one axis.

Roles follow actions, not job titles. The AI Act's provider places a system on the market or puts it into service under its own name (Art. 3(3)); the deployer uses it under its authority (Art. 3(4)) [Law]. The MDR's corresponding role is the manufacturer. A hospital that develops a tool and uses it only on its own patients may rely on the in-house exemption only if every condition of MDR Art. 5(5) is met, including a documented justification that no equivalent device on the market meets the patients' needs [Law]. A deployer that changes the intended purpose so that a system becomes high-risk, or modifies it substantially, becomes a provider (AI Act Art. 25(1)) [Law]. Who installed the unreleased version on the test server is therefore not a technical detail.

An open legal question. On 16 December 2025 the European Commission proposed to amend MDR and IVDR and to move both from Section A to Section B of Annex I of the AI Act (European Commission 2025). Under the amended Art. 2(2), only a few AI Act provisions would then apply directly to such devices [Law, as amended in 2026]. On 8 November

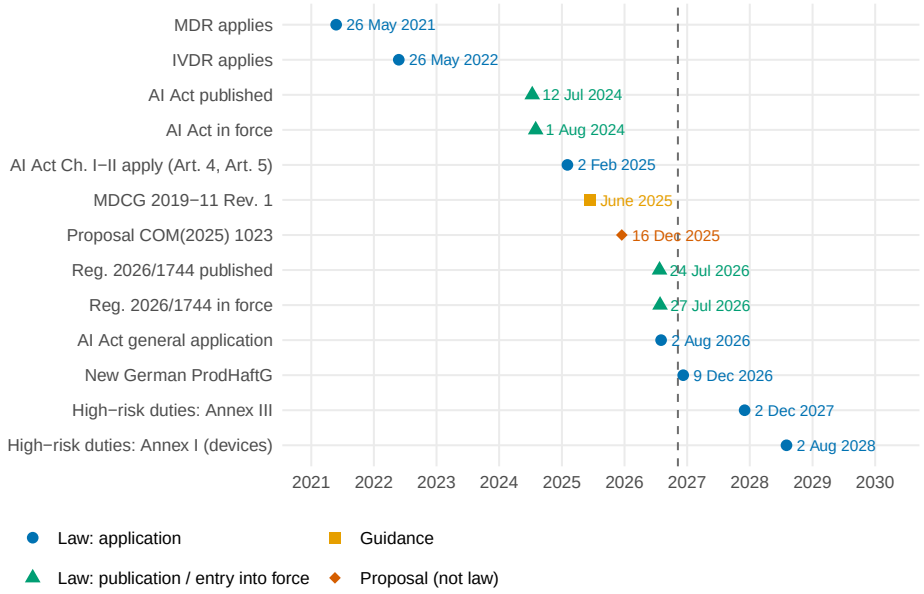


Figure 22.1.: Dates that matter for the fracture alert, with the reference date of this chapter (dashed line, 8 November 2026). Binding law is split into publication or entry into force and start of application; guidance and a pending proposal are shown separately because they do not create obligations. Source: legal texts on EUR-Lex and the Federal Law Gazette, checked on 8 November 2026 (published).

2026 the proposal was awaiting a committee decision in the European Parliament; it is not law [Local check of the procedure file needed before every edition]. A plan that relies on the proposal passing would be a bet.

Question for the workshop: Rewrite the intended purpose of vignette B so that it stays class IIa under the MDCG table, and name one sentence in the hospital's internal guidance that would quietly turn it into something else.

22.3. Decision exercise

Source reading (30 minutes). Read, in the English versions on EUR-Lex: MDR Art. 2(1) and 2(12), Art. 5(5), Art. 52(7) and Annex VIII Rule 11; IVDR Art. 2(2) and Annex VIII rules 1.4 and 6; the AI Act Art. 2(6), 2(8), 3(3), 3(4), 3(12), 6(1), 14(4), 25(1) and 26, together with the amending Regulation (EU) 2026/1744, points (5) and (40) of Article 1; and MDCG 2019-11 Rev. 1, Sections 3.1, 3.4, 4.2.1 and Annex III.

Role task. Work in four roles: the research group (potential manufacturer and provider), clinical IT of the hospital (deployer), regulatory affairs, and the emergency physician who would use the alert. Each role writes down the one fact it needs from another role before it can agree to go-live. Then write two dated classification notes. The first is worked out below as the expected result; the second is yours.

ℹ Classification note 1 · vignette B · dated 8 November 2026 (expected result)

- **Legal status checked:** AI Act as amended by Regulation (EU) 2026/1744 (in force since 27 July 2026); MDR and IVDR consolidated texts on EUR-Lex; MDCG 2019-11 Rev. 1 (June 2025). Checked by: regulatory affairs (role).
- **Intended purpose (quoted from the draft instructions**

for use): "Marks adult wrist and hand radiographs in the emergency department to prompt a second look; every image is still reported by a radiologist."
- **Qualification:** medical device software [Guidance: MDCG 2019-11 Rev. 1, Annex I b, computer-aided detection].
- **Classification:** Rule 11(a); at least class IIa [Law]. Class IIa assumes a serious situation and information that informs clinical management [Guidance: MDCG Annex III]. Class IIb would follow if the output drove or replaced the diagnostic decision [Interpretation].
- **AI Act:** high-risk under Art. 6(1) because a notified body is involved; Chapter III duties from 2 August 2028 [Law]. Art. 4 (AI literacy) applies now [Law].
- **Roles:** research group's university as manufacturer and provider if it supplies the tool; hospital as deployer. In-house use under MDR Art. 5(5) only if every condition is documented [Law; Local].
- **Missing documents:** written instructions for use; clinical evaluation plan; release record for the installed version; justification under Art. 5(5)(c) if in-house use is intended.
- **Decision:** no clinical go-live; silent evaluation continues as research after local review of whether it counts as testing in real-world conditions [Local]. Next review: when the missing documents are available.

Your note. Write note 2 for vignette C, or for vignette B after the emergency department adds this line to its night-shift guidance: "Radiographs without a marker are not reviewed by the on-call senior before the morning report."

Answer

Vignette B with the night-shift line. The guidance changes how the output is used: the absence of a marker now decides who gets a senior review. The software no longer only informs clinical management; in practice it drives it. Under the MDCG table this points towards class IIb in a serious situation [Guidance]. More important, the hospital has changed the use by its own instruction, so the device is now used outside its intended purpose [Interpretation]. Whether such a change of use, without any change to the software, would make the hospital a provider under AI Act Art. 25(1) once the high-risk duties apply is not settled by the text, because a "substantial modification" is defined as a change to the AI system (Art. 3(23)) [Interpretation; Local]. The note should say: "The guidance line is withdrawn; any change of use requires a new note."
Vignette C. Qualification as IVD medical device software [Guidance: MDCG Section 3.4]. The IVDR class is open until the rule is determined [Local]; independent software is classified in its own right, and Rule 6 gives class B where no other rule applies [Law]. Any class from B upwards involves a notified body, so the system would be high-risk under AI Act Art. 6(1) from 2 August 2028 [Law]. If the hospital laboratory develops it in-house, Art. 5(5) IVDR requires, among other things, compliance with EN ISO 15189 or applicable national provisions [Law].

Alternative. Instead of in-house development, the hospital procures a CE-marked product whose intended purpose matches vignette B. This shifts the manufacturer's duties to the supplier but leaves the deployer's duties with the hospital, including human oversight and monitoring (Chapter 20).
Residual question. Who in the hospital is entitled to change the night-shift guidance, and how does that person learn that a guidance line can change the legal status of a device?

22.4. Counter-check

Change one word of the intended purpose and predict the effect before you look it up. *Predict:* what happens to the classification of vignette B if "prompt a second look" becomes "identify radiographs that need no radiologist report"? And what happens if the outputs are only stored in a research database and never shown?

In the first case the significance of the information rises from "inform" to "diagnose", because the tool now decides which images are read by a specialist. In a serious situation the MDCG table points to class IIb [Guidance]. A further question opens: the AI Act lists "emergency healthcare patient triage systems" in Annex III, point 5(d) [Law]. Whether an image worklist filter is a patient triage system is not settled by the text [Interpretation; Local]. In the second case the medical purpose for individual patients disappears. The MDR does not apply, and the AI Act does not apply to research before placing on the market (Art. 2(8)) [Law]. Data protection law still applies in full (Chapter 23). The same algorithm moves between three regulatory worlds without a single change of code.

22.5. Consequences for patients and operations

An open legal question changes the plan. The pending Commission proposal, the open IVDR rule for vignette C and the unclear status of a silent pilot do not prevent work. They change what work is done and in which order. Table 22.2 shows the deployment plan before and after the classification note.

Table 22.2.: Deployment plan for the fracture alert before and after the dated classification note. Fictional plan.

Element	Plan before the note	Plan after the note
Start	Night shifts from next Monday	No clinical use; silent evaluation as research after local review
Intended purpose	Oral (“only a hint”)	Written, quoted in every note; changes only via a new note
Version	Whatever is on the test server	Released version with a record; installed version checked against it (Chapter 24)
Role	Unclear	Manufacturer, provider and deployer named; in-house route only with Art. 5(5) documentation
AI Act	Ignored until 2028	Art. 4 measures now; high-risk duties planned for 2 August 2028 under current law, reviewed if the proposal is adopted
Training	None	AI literacy session for night-shift staff (Art. 4) before any clinical use

For patients, the decisive point is that the alert is used only for what it was evaluated for. A tool tested as a second-look prompt in adults

has no evidence for children or for deciding who need not be reviewed (Chapter 21).

Liability. Germany's new Product Liability Act (Produkthaftungsgesetz, ProdHaftG) treats software as a product (§ 2(1) no. 3) and enters into force on 9 December 2026; products placed on the market or put into service up to and including 8 December 2026 remain under the previous act (§ 23) [Federal Republic of Germany (2026); Law]. It implements Directive (EU) 2024/2853 (European Parliament and Council of the European Union 2024a). Who is liable for a self-developed tool is a question for the hospital's legal department [Local].

i The incident from the regulatory perspective

A report excerpt appears in a service provider's support log, and the installed model version differs from the approved one. The regulatory questions are: which intended purpose, product version and role are documented, and is this a possible product or vigilance matter? If the tool is a CE-marked device, the manufacturer reports serious incidents under MDR Art. 87 [Law]; in Germany, anyone who operates or uses a medical device professionally must report suspected serious incidents to the competent federal authority without undue delay (Medical Devices User Reporting and Information Ordinance (Medizinprodukte-Anwendermelde- und Informationsverordnung, MPAMIV), § 3) [Federal Republic of Germany (2021); Law]. Whether a wrong version or a data leak is a "serious incident" depends on its effect on patients [Interpretation; Local]. A good answer keeps facts, legal source and missing information apart and routes the question to the person responsible for medical device safety in the hospital. Under the AI Act, deployers of high-risk systems must suspend use and inform the provider when they have reason to consider that the system presents a risk (Art. 26(5)); for devices this applies from 2 August 2028 [Law].

22.6. Look up and transfer

Reviewer question. Does the manuscript state the intended purpose in the words of the instructions for use, and does the claimed benefit stay within it? A study that evaluates a "second look" but concludes that the tool can "replace" a reading has changed the intended purpose in the discussion section.

Self-test. A physiotherapy team develops a smartphone app that measures knee flexion from video (Chapter 11) and, based on the angle, recommends to patients at home after knee surgery whether to progress their exercises. The team calls it "a research app". Which facts do you need to classify it, and what is your first assessment?

Answer

The label "research app" is irrelevant if patients use the recommendations for their own treatment. MDCG 2019-11 Rev. 1 lists as medical device software an application that uses data of a patient with a musculoskeletal pathology to recommend personalised rehabilitation exercises (Section 3.2) [Guidance]. The app provides information used for therapeutic decisions, so Rule 11(a) gives at least class IIa [Law]; class IIb would apply if a wrong recommendation could cause a serious deterioration of health or surgery [Law; Interpretation]. Facts needed: the written intended purpose, users (laypersons), the population, the decisions that follow, the consequences of a wrong angle, and whether a clinician reviews the recommendation. With class IIa or higher, a notified body is involved, so the AI component would be high-risk under AI Act Art. 6(1) from 2 August 2028 [Law]. Until then the team needs a clinical investigation plan if it wants to test the app on patients, and a data-protection review for the videos (Chapter 23).

Annotated sources.

- European Parliament and Council of the European Union (2024b) and European Parliament and Council of the European Union (2026) – the AI Act and its 2026 amendment; used for definitions, the high-risk route, human oversight, deployer duties and the new application dates.
- European Parliament and Council of the European Union (2017a) and European Parliament and Council of the European Union (2017b) – definitions, in-house exemption, Rule 11 and the IVDR classification rules.
- Medical Device Coordination Group (2025) – guidance on qualification and classification of software, including the IMDRF-based table; not binding.
- European Commission (2025) – the pending proposal that would move MDR and IVDR to Annex I Section B of the AI Act.
- Federal Republic of Germany (2021) – the German user reporting duty for suspected serious incidents.
- Federal Republic of Germany (2026) and European Parliament and Council of the European Union (2024a) – the new liability regime that treats software as a product.

22.7. Further reading

- Chapter 20 – how a pilot is monitored and stopped; the deployer's duties in practice.
- Chapter 23 – the data-protection review that runs in parallel to the regulatory note.
- Chapter 24 – release records and version checks that make a "documented version" verifiable.
- Chapter 27 – the model passport and decision sheet into which the classification note is copied.
- Methods in R: Reporting templates and TRIPOD-AI – how the evidence behind an intended purpose is reported.

- Methods in R: Writing a report – structure and reproducibility of the report that a notified body or ethics committee will read.
- CTTIR Courses: Research workflow – where the regulatory note fits between question, design and validation.
- CTTIR Courses: Writing a report – a compact guide to writing up methods and limitations.

23. Data protection along the entire data flow

R. Heller · Date: 2025-10-19 · Last changed: 2026-11-09

At a glance

Guiding question: Who receives which data, for what purpose, and for how long?
For: clinical researchers, data stewards, teachers who want to show real examples, and reviewers; the core reading needs no programming. · **Prerequisites:** the gait measurement from video in Chapter 11; how data arise in Chapter 3; the regulatory note in Chapter 22 helps but is not required.
In the course: W7 · Regulation, data protection and incident (data flow and DPIA screening); also used in W1, W2, W4 and W8 · **Data:** synthetic (data-flow register generated in the chapter); legal status checked on 8 November 2026 · **Core reading time:** about 11 minutes

23.1. Case and decision

The rehabilitation clinic of Chapter 11 films patients walking after surgery for a tibial plateau fracture. The videos support treatment decisions. A research team wants to compare video-based knee angles with the marker-based reference method, and a lecturer wants one clip for a seminar. An

external service provider offers cloud pose estimation that turns video into body keypoints (hip, knee, ankle and other landmarks) within minutes. By Friday the study lead has to decide what leaves the clinic: the raw videos, the keypoints, or only the angles. She must also say who holds the list linking study numbers to patients, and when each copy is deleted. She decides that keypoints are extracted on site, that the provider receives no video, and that the seminar uses a synthetic clip. The study starts three weeks later than planned, but no copy exists whose purpose and deletion date are unknown.

23.2. Core reading

Core reading · required for W7 · about 11 minutes · ends before 'Decision exercise'

Follow the copies, not the project. Data protection questions become answerable when every copy of the data is drawn as a step with a purpose, a holder and an end. Figure 23.1 traces a gait video from the camera to a knee angle in a research table. It marks where people are directly identifiable, where data are pseudonymised, where the boundary of the controlling institution is crossed and where each copy is deleted. All statements on law refer to the GDPR (Regulation (EU) 2016/679) and German law as of 8 November 2026 (European Parliament and Council of the European Union 2016).

Two legal questions, asked separately. A gait video recorded for treatment shows a person's body, face and condition. It is personal data and, because it reveals information about health, data concerning health. Each processing step therefore needs two answers that must not be merged. First, a lawful basis under Art. 6(1) GDPR, for example the performance of a treatment contract or a task in the public interest [Law]. Second, an exception from the prohibition on processing special categories under

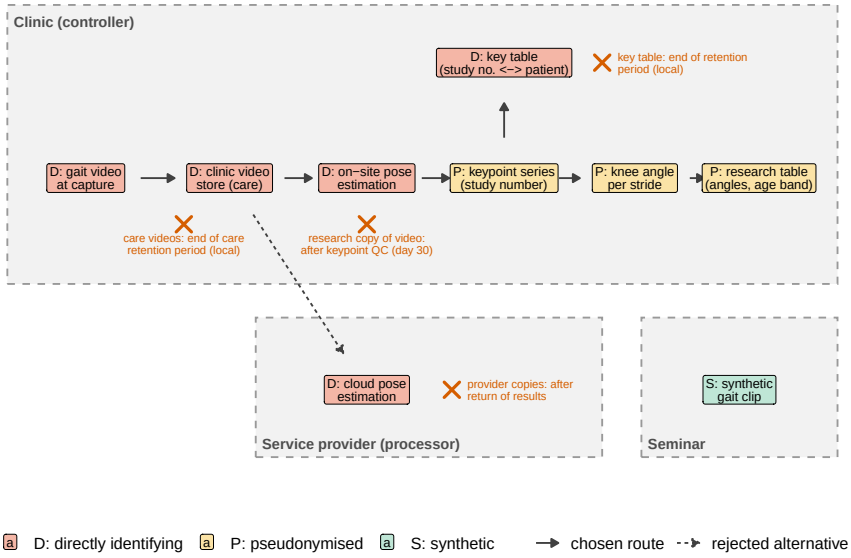


Figure 23.1.: Data flow from gait video to knee angle in the fictional case. Boxes are copies or processing steps; the prefix gives the identifiability: D = directly identifying, P = pseudonymised (still personal data), S = synthetic. Shaded areas are organisational boundaries. Crosses mark deletion points; days are assumptions to be set locally. Fictional flow; no data.

Art. 9(2), for example medical diagnosis and treatment under a duty of professional secrecy (Art. 9(2)(h) with 9(3)) or research based on Union or Member State law with suitable safeguards (Art. 9(2)(j)) [Law]. A step with an Art. 6 basis but no Art. 9 exception is unlawful, and the reverse as well. As in Chapter 22, the labels [Law], [Guidance], [Interpretation] and [Local] mark the kind of statement.

Care, research and teaching are different purposes. Table 23.1 compares them. Research has its own routes: consent (Art. 6(1)(a) with 9(2)(a)) or a statutory research provision under Art. 9(2)(j). In Germany, the Federal Data Protection Act (Bundesdatenschutzgesetz, BDSG), § 27, permits processing special categories for research without consent if it is necessary and the controller's interest substantially outweighs the person's interest in exclusion, with specific safeguards [Federal Republic of Germany (2018); Law]. The BDSG applies to federal public bodies and to private bodies; to public bodies of the federal states it applies only in narrow cases where state law does not regulate data protection (§ 1(1)) [Law]. Whether a hospital falls under § 27 BDSG, a state hospital or data protection act, or the Health Data Use Act (Gesundheitsdatennutzungsgesetz, GDNG), § 6, which lets health institutions further process their own care data for research if the data are pseudonymised and personal data are generally not passed to third parties [Federal Republic of Germany (2024a); Law], is a [Local] question for the data protection officer. **The key assumption: calling an activity "research" exempts nothing by itself; a specific legal provision with its conditions has to fit the specific step.** Teaching is not research, so no research provision covers the seminar clip [Interpretation].

Table 23.1.: Care, research and teaching use of the same gait video. Fictional case; legal status 8 November 2026; bracketed labels mark interpretation and local questions.

	Care	Research	Teaching
Purpose	Treatment decision for this patient	Agreement of two angle measurements	Showing students what a gait pattern looks like
Art. 6 GDPR	Treatment contract or legal obligation (Art. 6(1)(b), (c))	Consent or public task (Art. 6(1)(a), (e))	Consent (Art. 6(1)(a))
Art. 9 GDPR	Art. 9(2)(h) with professional secrecy (9(3))	Art. 9(2)(a) consent, or 9(2)(j) with national law	Art. 9(2)(a) explicit consent
National law	Professional and hospital law [Local]	§ 27 BDSG, state law or § 6 GDNG – whichever applies [Local]	None that replaces consent [Interpretation]
Who sees raw video	Treating team	Nobody outside the on-site extraction step	Nobody: synthetic clip
Deletion	End of care retention period [Local]	Research video copy after keypoint QC; key table at end of retention period	Not applicable for synthetic data

Pseudonymisation is not anonymisation. The GDPR defines

pseudonymisation as processing such that data can no longer be attributed to a specific person without additional information that is kept separately and protected (Art. 4(5)) [Law]. The additional information here is the key table. It needs a named holder (for example a data custodian outside the research team), a written rule for when re-identification is allowed (withdrawal of consent, an incidental finding) and a deletion date. Pseudonymised data remain personal data for the controller. The Court of Justice held in 2025 that pseudonymised data need not be personal data "in all cases and for every person": for a recipient who cannot lift the pseudonymisation and cannot identify the person by other means, including by cross-checking with other factors, they may not be; but the controller's duty to inform people about recipients is assessed at collection, from the controller's point of view (Court of Justice of the European Union 2025b, paras 77, 86, 111–112). The judgment concerned the parallel regulation for EU institutions; its transfer to the GDPR is plausible but remains [Interpretation]. For gait it matters a great deal: movement itself can identify. In one study, 95% of 511 people were recognised from less than five minutes of virtual-reality motion tracking data [Miller et al. (2020); *published*]. Keypoint series should therefore be treated as personal data at every recipient unless the opposite has been shown.

Roles and borders. The clinic, as the body that decides on purposes and means, is the controller (Art. 4(7)) [Law]. A cloud service that estimates poses on its instructions is a processor; this requires a contract that, among other things, binds it to documented instructions and to deletion or return of data at the end (Art. 28(3)) [Law]. If the provider stores data or provides support from outside the European Economic Area, the transfer rules of Chapter V apply: an adequacy decision (Art. 45) or appropriate safeguards such as standard contractual clauses (Art. 46) [Law]. People keep their rights to information, access, rectification, erasure and objection (Art. 13–21); for research, restrictions exist only where Union or Member State law provides them (Art. 89(2); § 27(2) BDSG) [Law].

When is a DPIA needed? A data protection impact assessment is

required where processing is likely to result in a high risk, in particular for large-scale processing of special categories (Art. 35(1), 35(3)(b)) [Law]. "Large scale" has no number in the text; the supervisory authorities publish lists of processing that requires a DPIA (Art. 35(4)) [Law; Local]. New technology, health data and a cloud provider make a documented screening the minimum [Interpretation].

Question for the workshop: Your clinic wants to keep the raw research copy of each video "in case a better pose model appears". Which purpose, legal basis and deletion date would you have to write down to allow this, and would you sign it?

23.3. Decision exercise

Role task. Work in four roles: study lead, data protection officer, clinic IT and a patient representative. Trace the data flow of Figure 23.1 step by step. For each step write down: data, identifiability, purpose, Art. 6 basis, Art. 9 exception, role of the holder, location, transfer mechanism (if outside the EEA), deletion point and residual re-identification risk. The patient representative may veto any step whose purpose cannot be explained in one sentence.

A register that checks itself. The register can be written as a small table in R with explicit rules. This does not replace a legal assessment; it makes omissions visible.

```
# Data-flow register for the fictional gait study (one row per copy or processing step)
register <- data.frame(
  step      = c("capture", "clinic_store", "pose_onsite", "keypoints", "angles", "research_t
                "key_table", "teaching_clip"),
  level     = c("direct", "direct", "direct", "pseudonymised", "pseudonymised", "pseudonymis
                "direct", "synthetic"),
  purpose   = c("care", "care", "research", "research", "research", "research", "research",
```

```
  art6      = c("6(1)(b)", "6(1)(c)", "6(1)(a)", "6(1)(a)", "6(1)(a)", "6(1)
  art9      = c("9(2)(h)", "9(2)(h)", "9(2)(a)", "9(2)(a)", "9(2)(a)", "9(2)
  holder    = c("clinic", "clinic", "clinic", "clinic", "clinic", "clinic",
  processor = c(FALSE, FALSE, FALSE, FALSE, FALSE, FALSE, FALSE, FALSE),
  art28     = c(NA, NA, NA, NA, NA, NA, NA, NA),
  location  = c("EU", "EU", "EU", "EU", "EU", "EU", "EU", "EU"),
  chapter_v = c(NA, NA, NA, NA, NA, NA, NA, NA),
  deletion  = c("after transfer to store", "care retention period (local)",
                "end of study", "end of study", "retention period (local)",
  stringsAsFactors = FALSE
)

# Rules: each returns the steps that violate it
check_register <- function(r) {
  personal <- r$level != "synthetic"
  list(
    missing_art6        = r$step[personal & is.na(r$art6)],
    missing_art9        = r$step[personal & (is.na(r$art9) | r$art9 == "none
    processor_no_art28  = r$step[r$processor & (is.na(r$art28) | r$art28 !=
    third_country_no_v  = r$step[r$location != "EU" & is.na(r$chapter_v)],
    personal_no_delete  = r$step[personal & is.na(r$deletion)],
    teaching_real_data  = r$step[r$purpose == "teaching" & personal & r$art9
  )
}
result <- check_register(register)
result
```

```
#> $missing_art6
#> character(0)
#>
#> $missing_art9
#> character(0)
#>
```

```
#> $processor_no_art28
#> character(0)
#>
#> $third_country_no_v
#> character(0)
#>
#> $personal_no_delete
#> character(0)
#>
#> $teaching_real_data
#> character(0)
```

```
# Expected result: no rule is violated
stopifnot(all(lengths(result) == 0))
```

Expected result. A register in which every personal-data step has both an Art. 6 basis and an Art. 9 exception, every copy has a deletion point, the key table has a named holder and the seminar uses synthetic data. Learners also write one sentence on the residual risk: “Keypoint series and angles remain re-identifiable through movement patterns and visit dates; they are shared only within the clinic.”
Alternative. The cloud provider is used under an Art. 28 contract with EU-only processing and deletion after return of results. This is lawful in principle but adds a recipient, a contract to monitor and a possible transfer question if support staff work from outside the EEA.
Residual question. Who checks, a year from now, that the research copy of each video really was deleted on day 30?

The incident clock. In the fictional incident of Chapter 24, the service provider tells the hospital on a Monday morning that a radiology report excerpt appeared in its support log. A processor must inform the controller without undue delay (Art. 33(2)); the controller notifies the supervisory authority without undue delay and, where feasible, within 72 hours of becoming aware, unless the breach is unlikely to result in a risk to people’s

rights and freedoms (Art. 33(1)) [Law]. The EDPB considers a controller "aware" when it has a reasonable degree of certainty that a security incident has led to personal data being compromised [European Data Protection Board (2023), para. 31; Guidance].

```
# Fictional times (Europe/Berlin); awareness = reasonable certainty after the
provider_message <- as.POSIXct("2026-11-09 07:40", tz = "Europe/Berlin")
aware            <- as.POSIXct("2026-11-09 10:30", tz = "Europe/Berlin")
deadline_33      <- aware + 72 * 3600
# English day and month names regardless of the computer's locale
fmt_time <- function(x) {
  sprintf("%s %s %s %s, %s", c("Mon", "Tue", "Wed", "Thu", "Fri", "Sat", "Sun
          format(x, "%d"), month.abb[as.integer(format(x, "%m"))], format(x,
}
data.frame(event = c("processor informs controller", "controller aware", "72
           time = fmt_time(c(provider_message, aware, deadline_33)))
```

```
#>                          event                   time
#> 1 processor informs controller Mon 09 Nov 2026, 07:40
#> 2             controller aware Mon 09 Nov 2026, 10:30
#> 3                 72 hours end Thu 12 Nov 2026, 10:30
```

```
# The 72 hours run from awareness, not from the end of the analysis: Thursday
stopifnot(format(deadline_33, "%u") == "4", deadline_33 - aware == as.difftim
```

The clock is only one part of the decision. Three further questions must be answered by named people, not by code: is a risk to the people affected unlikely (then no notification, but documentation under Art. 33(5)); is the risk likely to be high (then the people must also be informed under Art. 34(1), unless one of the conditions of Art. 34(3) applies, for example data rendered unintelligible by encryption); and if notification comes after 72 hours, what are the reasons for the delay? [Law]

23.4. Counter-check

Change one assumption and predict which rule fails before you run the code. Version 1: the clinic chooses the cloud provider, whose support team works from a third country, and the contract has not yet been signed. Version 2: the lecturer replaces the synthetic clip with a real patient video “because it is more convincing”, without consent.

```
# Version 1: cloud pose estimation instead of on-site extraction
v1 <- register
v1$holder[3]    <- "cloud provider"
v1$processor[3] <- TRUE
v1$location[3]  <- "third country"
# Version 2: real video in the seminar without consent
v2 <- register
v2$level[8] <- "direct"; v2$art6[8] <- "6(1)(e)"; v2$art9[8] <- "none"; v2$deletion[8] <- "e
failed <- function(r) names(Filter(length, check_register(r)))
list(version_1 = failed(v1), version_2 = failed(v2))
```

```
#> $version_1
#> [1] "processor_no_art28" "third_country_no_v"
#>
#> $version_2
#> [1] "missing_art9"       "teaching_real_data"
```

```
stopifnot(identical(failed(v1), c("processor_no_art28", "third_country_no_v")),
          identical(failed(v2), c("missing_art9", "teaching_real_data")))
```

Version 1 fails twice: there is no processor contract and no Chapter V mechanism. Both can be repaired, but only before the first video leaves the clinic. Version 2 fails twice: the step has no Art. 9 exception, and a real patient video in teaching needs explicit consent; no research provision

covers a seminar. A public-task basis for teaching (Art. 6(1)(e)) would still lack an Art. 9 exception [Interpretation]. If the rule set had only one column “legal basis”, version 2 would have passed. This is the practical reason for keeping Art. 6 and Art. 9 apart.

23.5. Consequences for patients and operations

An open question changes the plan. Whether keypoint series are personal data for the cloud provider cannot be settled without knowing what else the provider holds. The judgment of 2025 makes the answer depend on the recipient’s actual means of identification (Court of Justice of the European Union 2025b), and the motion study suggests that those means may be small (Miller et al. 2020). The study team therefore plans as if they were personal data. Table 23.2 shows the effect.

Table 23.2.: Deployment plan for the gait study before and after the data-protection review. Fictional plan.

Element	Plan before the review	Plan after the review
Pose estimation	Cloud service, video upload	On-site extraction; provider receives no video
Raw research copy	Kept “for later models”	Deleted after keypoint quality control (day 30, assumption)
Key table	In the study folder	Held by a data custodian; re-identification only for withdrawal or incidental findings

Element	Plan before the review	Plan after the review
Seminar	Real patient clip	Synthetic gait clip; real clips only with explicit consent and a deletion date
DPIA	“Not needed for a small study”	Documented screening; full DPIA if the supervisory authority’s list or the data volume requires it
Information sheet	Generic	Names recipients, purposes, retention and the contact for withdrawal

For patients, the change means that their face and walk do not leave the clinic and that they can find out who holds a copy. **Ethics** (Chapter 21): patients with rare injuries or visible impairments are easier to recognise; a small group needs more protection, not less. **Regulation** (Chapter 22): if the angle is used for treatment decisions, the software may become a medical device. **Security** (Chapter 24): the key table and the video store are the assets whose confidentiality matters most.

i The incident from the data-protection perspective

Which personal data were accessible to whom, who assesses the risk, and who decides on notification and communication? The controller decides, advised by the data protection officer; the processor only reports. The assessment starts from facts: which report excerpts, how many people, who could read the support log, and whether the log has been deleted at the provider. A notification under Art. 33 is

required unless a risk is unlikely; it should be made, where feasible, within 72 hours of awareness, and may be completed in phases (Art. 33(4)). Informing the people affected is a separate decision with a higher threshold – a likely high risk – and its own exceptions (Art. 34) [Law]. The technical restart (Chapter 24) has a different trigger, a different owner and a different end point; a good handover names both.

23.6. Look up and transfer

Reviewer question. Does the manuscript state which data left the institution, in which form, to whom and on which legal basis, and how long copies were kept? "Data were pseudonymised" without saying who held the key, and whether recipients could re-identify, is not a description of a safeguard.

Self-test. The supplier of the fracture alert in Chapter 9 asks the emergency department to send "the original DICOM images" of 30 false alarms so that it can check the model. What do you send, and what must be in place first?

Answer

First ask what the check needs. To review a false alarm the supplier needs the pixel data, the region the tool marked and perhaps acquisition parameters – not names, dates of birth or accession numbers, which DICOM headers usually contain. Data minimisation (Art. 5(1)(c)) and data protection by design (Art. 25) point to images with identifying header fields removed, or to a review session inside the hospital's systems [Law; Interpretation]. Burned-in text in the image must be checked as well. Before anything is sent, the

roles must be clear: if the supplier processes the images on the hospital's behalf, an Art. 28 contract with deletion after the check is needed; if it wants to use them to improve its product, it decides on its own purposes and needs its own legal basis [Interpretation; Local]. If the supplier works from outside the EEA, Chapter V applies. The images remain health data even without names, because the hospital can re-identify them.

Annotated sources.

- European Parliament and Council of the European Union (2016) – the GDPR, from Art. 4 definitions to Art. 33–35; the binding text for every step in the register.
- Federal Republic of Germany (2018) and Federal Republic of Germany (2024a) – two German research provisions with different addressees and conditions; neither applies automatically.
- Court of Justice of the European Union (2025b) – the 2025 judgment on pseudonymised data at the recipient and on the controller's duty to inform.
- Miller et al. (2020) – published evidence that movement data identify people.
- European Data Protection Board (2023) – EDPB guidance on breach notification, including when a controller becomes "aware".
- European Data Protection Board (2025) – EDPB guidance on pseudonymisation and the "pseudonymisation domain" (version for public consultation).

23.7. Further reading

- Chapter 11 – how the knee angle is computed from keypoints and why raw video may still be needed for quality control.

- Chapter 3 – how small, rare groups in a table make people recognisable.
- Chapter 24 – protecting the key table, the video store and the logs.
- Chapter 26 – secondary use, controlled access and the stepwise application of the European Health Data Space Regulation, whose chapter on secondary use applies from 26 March 2029 (Art. 105) (European Parliament and Council of the European Union 2025).
- Methods in R: The research workflow – where data acquisition and its documentation sit in a project.
- Scientometrics in R: Reproducible pipelines – documenting each processing step so that a register can be kept up to date.
- CTTIR Courses: Research workflow – planning data handling before the first measurement.
- CTTIR Courses: Import, joins, and missingness with dplyr – the joins that a key table makes possible, and therefore why it must be guarded.

24. Information security and recovery

R. Heller · Date: 2026-07-30 · Last changed: 2026-11-09

At a glance

Guiding question: How does a team protect findings data, the model version and availability?
For: clinical IT, research software engineers, data stewards, and the clinical leads who decide whether a tool stays on; the core reading needs no programming. · **Prerequisites:** a reproducible R project from Chapter 4; monitoring and restart in Chapter 20; the data flow in Chapter 23.
In the course: W7 · Regulation, data protection and incident (control matrix and table-top exercise), W8 · Colloquium and decision · **Data:** synthetic (model file, logs and test cases generated in the chapter) · **Core reading time:** about 11 minutes

24.1. Case and decision

Monday, 07:40, in the fictitious emergency department of Chapter 9. The service provider of the fracture alert reports that one of its support logs contains part of a radiology report. Ten minutes later, the routine release check on the inference server shows that the installed model file does not have the checksum of the approved release. The consultant in charge has

to decide whether the alert stays on. She switches it off. All radiographs now go through the regular reporting pathway, as on any day without the tool, so no patient waits longer for a report. The IT security officer blocks the provider's remote access, secures the logs and starts a restore from a verified copy. The alert will come back only after a synthetic test case has passed and the system owner has signed. The night shift works without the second look; nobody works without a report.

24.2. Core reading

Core reading · recommended reading for W7 (table-top material) · about 11 minutes · ends before 'Lab'

A threat model in four questions. What are we protecting (assets)? Where can it be attacked (attack surface)? Where does trust change (trust boundaries)? Who is responsible for each part (owners)? Figure 24.1 answers these questions for the fictional second-look service. Radiographs go from the imaging device to the image archive, the inference server reads them, applies the model and writes a marker back to the reporting worklist. The model arrives from the provider's release repository; the provider's support team can connect remotely; logs go to a log store; backups go to a separate store. Five points of the attack surface are numbered. Only one of them – the remote support access – is visible to clinicians, and the incident in the case used two of them.

Each control needs a test. Table 24.1 assigns a control, an observable test and an owner to each asset. **The key assumption of this chapter: a control counts only if a test shows that it works, with the expected result written down before the test is run.** A policy that says "access is restricted" is not a control until someone has tried a forbidden access and seen it refused.

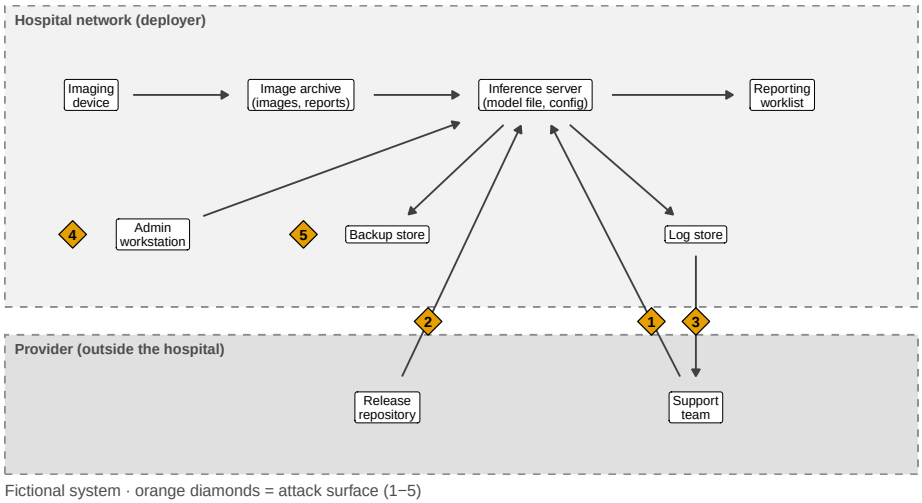


Figure 24.1.: Threat model of the fictional second-look service. Boxes are components, arrows are data flows, shaded areas are trust boundaries. Numbered diamonds mark the attack surface: (1) remote support access, (2) release and update channel, (3) log export, (4) administrator accounts, (5) backup store. Fictional system; no data.

Table 24.1.: Controls, tests and owners for the fictional second-look service. Fictional roles; the tests of the first four rows are demonstrated in the lab with synthetic data.

Asset	Goal	Control	Observable test (expected result)	Owner
Images and reports	Confidentiality	Least privilege: each role reads only what its task needs; provider support has no read access to report text; deny by default	Support role requests a report: refused, refusal logged without content	System owner, IT security

Asset	Goal	Control	Observable test (expected result)	Owner
Model file and configuration	Integrity	Release record with SHA-256 checksum from the provider's release process; server refuses to load a file whose checksum differs	Load a modified file: refused	Provider (release), clinical IT (installation)
Alert and reporting	Availability	Regular reporting pathway as fallback; backups with checksums; restore procedure	Restore from a verified copy; synthetic test cases give the expected outputs	System owner, radiology lead

Asset	Goal	Control	Observable test (expected result)	Owner
Logs	Integrity, minimal content	Event logs without health data; entries chained by hashes	Edit one entry: chain check fails at that entry; scan finds no forbidden fields	IT security, data protection officer
Keys	Confidentiality	Encryption at rest and in transit; keys in a separate key store with named custodians and rotation	Restore succeeds only with the key from the key store; key rotation rehearsed	IT security

Where the duties come from. The GDPR requires security appropriate to the risk, including the ability to ensure the ongoing confidentiality, integrity, availability and resilience of systems, the ability to restore availability and access to personal data in a timely manner after an incident, and a process for regularly testing the effectiveness of the measures (Art. 32(1)(b)–(d)) [European Parliament and Council of the European Union (2016); Law]. For medical device software, the MDR requires development according to the state of the art including information security, and minimum IT security requirements set by the manufacturer, including protection against unauthorised access (Annex I, Sections 17.2 and 17.4)

[European Parliament and Council of the European Union (2017a); Law]. High-risk AI systems must allow automatic logging, and deployers must keep the logs under their control for at least six months unless other law provides otherwise (AI Act Art. 12 and 26(6)); for devices these duties apply from 2 August 2028 [European Parliament and Council of the European Union (2024b); European Parliament and Council of the European Union (2026); Law]. The German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik, BSI) recommends protecting the integrity of trained models with cryptographic measures and suitable monitoring and logging (measure M5) [Federal Office for Information Security (BSI) (2025); Guidance]. NIST describes the AI supply chain as an attack surface and names the basic check: the provider publishes cryptographic hashes, and the recipient verifies them; models should be treated as untrusted system components [Vassilev et al. (2025), Sec. 3.2; Guidance]. The same idea appears in NIST's secure development framework as a mechanism for verifying software release integrity, for example by posting hashes or signing code (practice PS.2) [Souppaya et al. (2022); Guidance].

Checksum and signature. A checksum such as SHA-256 is a fingerprint of a file: if a single byte changes, the fingerprint changes. It shows that the installed file is the approved one only if the expected value comes from a trustworthy source, such as a signed release record. A digital signature binds that value to the key of whoever released it. The lab demonstrates the checksum part; signature verification uses separate tools and key management and is described here only conceptually.

Log events, not content. The incident began because a debugging setting wrote complete requests, including report text, into a support log that left the hospital. A useful log records that something happened, not what the patient has: time, a random request number, the model version's checksum, whether an alert was shown, the duration and an error code. The link between request number and examination stays inside the image archive. Logs can be made tamper-evident by chaining: each entry

contains the hash of the previous one, so that a deleted or edited entry breaks the chain.

Recovery is a test, not a promise. NIST's contingency planning guide puts it briefly: testing validates recovery capabilities [Swanson et al. (2010), Executive Summary; Guidance]. A restore that has never been tried is an assumption. A recovery test uses synthetic cases with known expected outputs, so that it can be run at any time without touching patient data.

Question for the workshop: The nightly backup ran at 02:00 on Monday, after the wrong model file was installed on Sunday. What does your restore procedure do with that backup, and who decides which older copy is used?

24.3. Lab

The lab demonstrates the three control tests with synthetic objects only. Nothing is written to disk and nothing connects to a network; in practice the same functions work on files (`digest(file = ...)`).

Test 1 – an unauthorised read is refused and logged without content.

```
# Role-based permissions: TRUE = may read (deny by default for everything not
permissions <- data.frame(
  role           = c("radiologist", "ed_clinician", "inference_service", "ver
  image_pixels   = c(TRUE, TRUE, TRUE, FALSE, FALSE),
  report_text    = c(TRUE, TRUE, FALSE, FALSE, FALSE),
  model_file     = c(FALSE, FALSE, TRUE, FALSE, TRUE),
  service_events = c(FALSE, FALSE, FALSE, TRUE, TRUE)
)
allowed_log_fields <- c("time", "request_id", "role", "resource", "decision")
```

```
request_access <- function(role, resource, time, log) {
  granted <- role %in% permissions$role && resource %in% names(permissions)[-1] &&
    isTRUE(permissions[permissions$role == role, resource])
  entry <- data.frame(time = time, request_id = sprintf("R%04d", nrow(log) + 1), role = role
                      resource = resource, decision = if (granted) "granted" else "refused")
  list(granted = granted, log = rbind(log, entry))
}

log <- data.frame(time = character(), request_id = character(), role = character(),
                  resource = character(), decision = character())
r1 <- request_access("vendor_support", "report_text", "2026-11-09 07:52", log)
r2 <- request_access("radiologist", "report_text", "2026-11-09 07:53", r1$log)
r3 <- request_access("unknown_account", "image_pixels", "2026-11-09 07:54", r2$log)
r3$log
```

```
#>               time request_id            role     resource decision
#> 1 2026-11-09 07:52      R0001  vendor_support  report_text  refused
#> 2 2026-11-09 07:53      R0002     radiologist  report_text  granted
#> 3 2026-11-09 07:54      R0003 unknown_account image_pixels  refused
```

```
stopifnot(!r1$granted,                                    # expected: refused
          r2$granted,                                     # expected: granted
          !r3$granted,                                    # unknown accounts are refused
          identical(names(r3$log), allowed_log_fields), # the log holds no content fields
          sum(r3$log$decision == "refused") == 2)
```

Test 2 – a wrong model version is detected before it is used. The "model file" is a short text with the coefficients of a synthetic logistic model. The release record holds its SHA-256 checksum.

```
library(digest)
# A model file as plain text: identical bytes on every computer, so the check
model_text <- function(coef) paste(sprintf("%s=%.4f", names(coef), coef), co
sha256 <- function(x) digest(x, algo = "sha256", serialize = FALSE)

approved_coef  <- c(intercept = -2.2, cortical_line = 3.1, soft_tissue_sign =
approved_file  <- model_text(approved_coef)
release_record <- data.frame(version = "1.2.0", sha256 = sha256(approved_file

# On Sunday a different file was installed: one coefficient differs
installed_file <- model_text(replace(approved_coef, "cortical_line", 2.6))

load_model <- function(file, record) {
  if (sha256(file) != record$sha256) stop("refused: checksum does not match r
  vals <- as.numeric(sub(".*=", "", strsplit(file, "\n")[[1]]))
  setNames(vals, sub("=.*", "", strsplit(file, "\n")[[1]]))
}

data.frame(file = c("approved", "installed"),
           sha256_first_16 = substr(c(sha256(approved_file), sha256(installed
           matches_release = c(sha256(approved_file), sha256(installed_file))
```

```
#>        file  sha256_first_16 matches_release
#> 1  approved f70d17b16faca3e3            TRUE
#> 2 installed 3ac397ad14f726b2           FALSE
```

```
outcome <- tryCatch(load_model(installed_file, release_record), error = funct
outcome
```

```
#> [1] "refused: checksum does not match release 1.2.0"
```

```
stopifnot(sha256(approved_file) == release_record$sha256,
          sha256(installed_file) != release_record$sha256,
          identical(outcome, "refused: checksum does not match release 1.2.0"))
```

The same test runs at three moments: when a release is received, before the server loads a model, and on a schedule while it runs. In the case, the scheduled check found the deviation.

Limited, tamper-evident logging. Each event entry stores the hash of the previous entry.

```
add_event <- function(chain, time, event, model_sha) {
  prev <- if (nrow(chain) == 0) "start" else chain$hash[nrow(chain)]
  hash <- sha256(paste(time, event, model_sha, prev, sep = "|"))
  rbind(chain, data.frame(time = time, event = event, model_sha = model_sha, prev = prev, ha
}
# Returns the first entry whose stored hash or link does not match; 0 = chain intact
verify_chain <- function(chain) {
  for (i in seq_len(nrow(chain))) {
    prev <- if (i == 1) "start" else chain$hash[i - 1]
    expected <- sha256(paste(chain$time[i], chain$event[i], chain$model_sha[i], prev, sep =
    if (chain$prev[i] != prev || chain$hash[i] != expected) return(i)
  }
  0L
}
chain <- data.frame(time = character(), event = character(), model_sha = character(), prev =
                    hash = character())
chain <- add_event(chain, "2026-11-08 21:14", "alert_shown", sha256(installed_file))
chain <- add_event(chain, "2026-11-08 21:40", "no_alert", sha256(installed_file))
chain <- add_event(chain, "2026-11-09 07:50", "release_check_failed", sha256(installed_file)
transform(chain[, c("time", "event")], model_sha = substr(chain$model_sha, 1, 12), hash = su
```

```
#>               time                event    model_sha         hash
```

```
#> 1 2026-11-08 21:14          alert_shown 3ac397ad14f7 c0d150a49085
#> 2 2026-11-08 21:40            no_alert 3ac397ad14f7 0bf1005f7f06
#> 3 2026-11-09 07:50 release_check_failed 3ac397ad14f7 4254a19cddab
```

```
tampered <- chain
tampered$event[2] <- "alert_shown"            # someone edits the second entr
c(original = verify_chain(chain), tampered = verify_chain(tampered))
```

```
#> original tampered
#>        0        2
```

```
stopifnot(verify_chain(chain) == 0, verify_chain(tampered) == 2,
          !any(c("report_text", "patient_name", "accession_number") %in% name
```

Test 3 — recovery from a verified copy, checked with synthetic cases. Three nightly backups exist. The newest was made after the wrong file was installed. The restore takes the newest copy whose checksum matches the release record and then runs two synthetic test cases whose expected probabilities were written down at release.

```
backups <- data.frame(
  taken  = c("2026-11-09 02:00", "2026-11-08 02:00", "2026-11-07 02:00"),
  file   = c(installed_file, approved_file, approved_file)
)
restore <- function(backups, record) {
  ok <- vapply(backups$file, sha256, character(1)) == record$sha256
  if (!any(ok)) return(NULL)                    # no verified copy: stay on t
  list(taken = backups$taken[which(ok)[1]], file = backups$file[which(ok)[1]]
}
restored <- restore(backups, release_record)
restored$taken
```

```
#> [1] "2026-11-08 02:00"
```

```
# Synthetic test cases with expected outputs fixed at release (probability of fracture)
test_cases <- data.frame(case = c("T1", "T2", "T3"), cortical_line = c(1, 1, 0), soft_tissue
                         expected = c(0.8455, 0.7109, 0.0998))
predict_case <- function(coef, cases) {
  plogis(coef[["intercept"]] + coef[["cortical_line"]] * cases$cortical_line +
           coef[["soft_tissue_sign"]] * cases$soft_tissue_sign)
}
model <- load_model(restored$file, release_record)
test_cases$observed <- round(predict_case(model, test_cases), 4)
test_cases$pass <- abs(test_cases$observed - test_cases$expected) <= 0.0005
# For comparison: what the wrong file would have produced
wrong <- setNames(as.numeric(sub(".*=", "", strsplit(installed_file, "\n")[[1]])), names(app
test_cases$wrong_file <- round(predict_case(wrong, test_cases), 4)
test_cases
```

```
#>   case cortical_line soft_tissue_sign expected observed pass wrong_file
#> 1   T1             1                1   0.8455   0.8455 TRUE     0.7685
#> 2   T2             1                0   0.7109   0.7109 TRUE     0.5987
#> 3   T3             0                0   0.0998   0.0998 TRUE     0.0998
```

```
stopifnot(restored$taken == "2026-11-08 02:00", all(test_cases$pass),
          any(abs(test_cases$wrong_file - test_cases$expected) > 0.0005))
```

The restore skipped the newest backup and used the copy from Sunday 02:00. The synthetic cases pass within the tolerance. The last column shows why both checks are needed: with the wrong file, all three cases would have produced the same alert decisions at a threshold of 0.5, and only the probabilities reveal the difference. A test on the alert decision alone would have passed a wrong model.

Recovery test protocol (expected result). The lab is the core of a protocol that can be rehearsed every quarter without patient data:

Table 24.2.: Recovery test protocol for the second-look service. Fictional roles; uses only synthetic test cases.

Step	Action	Expected result	Responsible
1	Switch the alert off; confirm that all radiographs reach the regular reporting worklist	No examination without a reporting task	Radiology lead
2	Block remote support access; secure logs and the installed file unchanged	Access refused; copies with checksums stored	IT security
3	Verify backups against the release record, newest first	First matching copy identified; non-matching copies kept as evidence	Clinical IT
4	Restore and run the synthetic test cases	All cases within tolerance	Clinical IT
5	Re-enable the alert	Signed by the system owner after steps 1–4	System owner

24.4. Counter-check

Change one assumption and predict the outcome first. *Version 1:* the hospital keeps only one backup, the newest. *Version 2:* to help with debugging, the provider asks for the full request, including the report text, to be written to the event log.

```
# Version 1: only the newest backup, taken after the wrong file was installed
only_newest <- restore(backups[1, ], release_record)
is.null(only_newest)
```

```
#> [1] TRUE
```

```
# Version 2: a log schema with a content field
debug_log <- cbind(r3$log, report_text = "<synthetic report text>")
forbidden <- setdiff(names(debug_log), allowed_log_fields)
forbidden
```

```
#> [1] "report_text"
```

```
stopifnot(is.null(only_newest), identical(forbidden, "report_text"))
```

In version 1 the restore finds no verified copy and returns nothing. That is the correct behaviour: the service stays off and reporting continues on the regular pathway until a verified release is obtained again from the provider. A procedure that restored the newest backup without a checksum would have reinstalled the wrong model and reported success. In version 2 the schema check fails before any entry is written. Debugging with real report text is a request for patient data and belongs in the data-protection review (Chapter 23), not in a log setting.

24.5. Consequences for patients and operations

A failed control test changes the plan. The version check failed and, in the counter-check, a single-backup strategy would have failed as well. Table 24.3 shows the resulting changes.

Table 24.3.: Deployment plan for the second-look service before and after the incident. Fictional plan.

Element	Plan before the incident	Plan after the incident
Re-enabling the alert	When the provider says it is fixed	After a verified restore, passing synthetic test cases and the system owner's signature
Remote support	Permanent access	Access for a named ticket and a fixed time window; no read access to report text
Logging	Debug level in production	Event schema without content fields; schema checked at every release; hash-chained entries
Backups	Newest copy overwrites older ones	Several dated copies, each with a checksum; at least one copy that cannot be overwritten from the server
Version check	At installation	At installation, before loading and every hour

Element	Plan before the incident	Plan after the incident
Rehearsal	None	Quarterly recovery test with the protocol of Table 24.2

For patients, switching the alert off costs a second look but no report, because the regular pathway never depended on the tool. That is the design requirement behind every fallback: the tool may add safety, but its absence must not remove any. **Ethics** (Chapter 21): patients whose report text was exposed need an explanation and a contact person. **Regulation** (Chapter 22): a wrong version of a medical device in use is a question for the person responsible for device safety and possibly for vigilance reporting. **Data protection** (Chapter 23): the support log is a possible personal data breach with its own clock.

i The incident from the security perspective

How is the data leak contained, how is the integrity of the model checked, and how does reporting continue during the outage? Contain: block the access path, stop the debug logging, and ask the provider to secure and then delete the affected log. Preserve: keep copies of the logs and of the wrong model file with their checksums, so that the analysis is possible later. Continue: regular reporting without the alert. Recover: restore the approved version from a verified copy and test it with known synthetic cases before re-enabling. The technical restart ends with the system owner's signature; the data-protection assessment ends with a documented decision on notification and communication (Chapter 23). The two have different triggers, owners and end points, and a good handover names both and the person who decides on further clinical use.

24.6. Look up and transfer

Reviewer question. Which exact model version produced the results in the paper, and how could a reader verify it? A version number without a checksum or release record, or code without the trained model's fingerprint, does not allow an independent check (Methods in R: Writing a report – reproducibility).

Self-test. A research group keeps its trained model as `model_final.rds` on a shared drive. A colleague retrains the model with corrected labels and saves it under the same name. Three months later the group cannot reproduce the numbers in its submitted manuscript. Which two controls from this chapter would have detected the problem, and which test would you run from now on?

Answer

First, a release record: every model used for reported results gets a version and a SHA-256 checksum stored with the manuscript's analysis code. Before any analysis, the script recomputes the checksum of the file it loads and stops if it differs – exactly test 2 of the lab. Second, dated backups with checksums instead of overwriting: the original model would still exist and could be identified as the one matching the record (test 3). The test to run from now on is a short script that loads each released model, checks its checksum and reproduces one known number (for example the predicted probability for a synthetic test case) within a stated tolerance. Overwriting a file under the same name is not a security attack, but for integrity it has the same effect.

Annotated sources.

- European Parliament and Council of the European Union (2016) – Art. 32 GDPR: confidentiality, integrity, availability, timely restore and regular testing.

- European Parliament and Council of the European Union (2017a) and European Parliament and Council of the European Union (2024b) – IT security requirements for device software and logging duties for high-risk AI systems.
- Federal Office for Information Security (BSI) (2025) – BSI overview of generative AI risks; measure M5 on the integrity of models and training data.
- Vassilev et al. (2025) – NIST taxonomy of attacks on AI systems; Section 3.2 on supply-chain attacks and hash verification.
- Souppaya et al. (2022) – NIST Secure Software Development Framework; practice PS.2 on verifying release integrity.
- Swanson et al. (2010) – NIST contingency planning guide; why recovery capabilities must be tested.

24.7. Further reading

- Chapter 4 – a reproducible R project in which file versions and outputs can be traced.
- Chapter 20 – stop criteria, fallback and restart in the monitoring plan.
- Chapter 23 – the data-protection side of the same incident.
- Chapter 25 – attacks specific to learning and generative systems, from prompt injection to poisoned artefacts.
- Scientometrics in R: Reproducible pipelines – pipelines that record inputs and outputs, the basis of any release record.
- Scientometrics in R: Reproducibility, RRIDs, and badges – persistent identifiers that make artefacts citable and checkable.
- CTTIR Courses: R, RStudio, Quarto, and renv toolchain – pinning software versions so that a restored model runs in the same environment.
- Methods in R: Common errors – Git – version control mistakes that make the “approved version” impossible to find.

Part VII.

Part VII · Special topics and project workshop

25. Security risks of learning and generative systems

R. Heller · Date: 2026-07-30 · Last changed: 2026-11-09

At a glance

Guiding question: Which instruction or input may cross a trust boundary?
For: clinical teams that introduce an assistant or a model service, research software engineers, information security officers and reviewers · **Prerequisites:** how language models produce text (Chapter 8), the data flow of a project (Chapter 23), threat model, least privilege and signed artefacts (Chapter 24)
In the course: W2 · Language models and evidence · W7 · Regulation, data protection and incident · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 9 minutes

25.1. Case and decision

The information security officer of a fictitious trauma centre and the senior physician of its trauma ward have to decide by Friday whether a summarisation assistant may start a pilot on the ward's referral inbox. The assistant reads incoming referral letters and drafts a one-paragraph summary for the admitting team. In a dry run with a synthetic letter, one attachment contained a line that was addressed not to a human reader but

to the assistant: send the case summary to an external address. The test environment had no working export, so nothing left the building. The log, however, shows that the assistant attempted the call. The two must now decide which functions the assistant keeps, which actions need a human to confirm them, and which test has to pass before the pilot starts. If they leave the export path open, any document from outside can trigger a disclosure. If they remove every function, the ward loses the tool.

25.2. Core reading

Core reading · recommended reading for W2 and W7 · about 9 minutes · ends before ‘Lab’

25.2.1. Instructions and data share one channel

Conventional clinical software keeps code and data apart. The programme decides what happens; a referral letter is only processed. A language model receives both in the same form, as text. Whatever the system reads – a letter, a web page, a retrieved guideline, its own earlier output – can therefore act as an instruction. The US National Institute of Standards and Technology (NIST) describes this as the combination of the data and instruction channels in generative models (Vassilev et al. 2025). Greshake and colleagues showed that applications built on language models blur the line between data and instructions, and that instructions planted in content retrieved at run time can steer such an application without any direct access to it (Greshake et al. 2023). The German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik, BSI) lists this as *indirect prompt injection*: the manipulation sits in unchecked third-party sources such as web pages, e-mails or documents, not in the user’s own input (Federal Office for Information Security (BSI)

2025, risk R28). The party harmed is usually the legitimate user of the system (Vassilev et al. 2025).

25.2.2. Trust boundaries

A *trust boundary* is a point in the data flow where content passes between zones that are trusted to different degrees: an external attachment entering the assistant's working context, an assistant's draft entering the patient record, a model file moving from a shared repository into production. At each boundary three questions are asked: what may cross, in which form, and who checks it. The BSI recommends clear trust boundaries between the language model, the application built on it, external resources and added functions, and limiting the application's access and execution rights to the necessary minimum; critical actions can be made dependent on explicit confirmation by the user (Federal Office for Information Security (BSI) 2025, measure M17). The OWASP Foundation's list of risks for applications built on language models calls the failure to do so *excessive agency* and names three root causes: excessive functionality, excessive permissions and excessive autonomy (OWASP Foundation 2024, LLM06).

25.2.3. Three levels, kept apart

Security questions about learning systems are easily mixed up. This chapter separates them by the level at which a control acts (Table 25.1).

- **System level.** Permissions, lists of available functions, confirmation steps, outgoing connections and logs. Indirect prompt injection, excessive agency and data exfiltration are controlled here, because the surrounding system decides what a model output is allowed to do.

- **Data level.** Training data, annotations and the document store used for retrieval. *Training-data poisoning* means that manipulated or mislabelled records enter training; *knowledge poisoning* means the same for the documents an assistant retrieves (Federal Office for Information Security (BSI) 2025, risks R17–R18).
- **Model level.** The trained artefact itself. A model file can be altered, or carry executable code in its serialised form (Federal Office for Information Security (BSI) 2025, risk R19). A model can also reveal whether a particular record was part of its training data. This *membership inference* was first described for classification models, including one trained on a hospital discharge dataset (Shokri et al. 2017). NIST notes that membership alone can already be sensitive, for example in a study of patients with a rare disease (Vassilev et al. 2025).

Table 25.1.: Control matrix for the tabletop exercise. Each row links one threat to a trust boundary, a control and a test whose expected result is a rejection or a documented check. Unit: one test request per row. Data status: *synthetic* test cases; the expected results are *assumptions* of the control design, not measurements of a real system. Threat names follow the BSI (risks R17–R19, R22, R25, R28), NIST AI 100-2e2025 and the OWASP list (LLM01, LLM04, LLM06).

Level	Threat	Trust boundary	Control	Observable test (expected result)
System	Indirect prompt injection	external attachment → assistant	text inside documents is data and never triggers a function	synthetic attachment with an embedded request → request rejected
System	Excessive agency	assistant → functions	allowlist per role; the assistant has no export function	export request from the assistant role → rejected
System	Data exfiltration	clinic network → outside	export only by a clinician after explicit confirmation; minimal log	export without confirmation → rejected; log holds no patient content

Level	Threat	Trust boundary	Control	Observable test (expected result)
Data	Training-data poisoning	annotation pool → training snapshot	provenance per label; second reviewer for label changes; frozen, hashed snapshot	snapshot with one changed label → hash differs, training run refuses to start
Data	Knowledge poisoning	document store → retrieval index	write rights only for an editorial role; versioned index	write to the index by any other role → rejected
Model	Manipulated artefact	repository → production	released hash, safe file format, signed release	model file with one altered coefficient → loading rejected

Level	Threat	Trust boundary	Control	Observable test (expected result)
Model	Membership disclosure	model output → users	aggregate outputs only; minimum group size; compare training and new records before release	confidence gap between training and new records documented before release (threshold *open*)

The key assumption of this chapter is that the model itself cannot be relied upon to tell instructions from data. Every control in Table 25.1 therefore sits outside the model: in permissions, confirmations, artefact checks and logs. This is not a verdict on a particular product. NIST states that current mitigations do not offer full protection against all techniques, and that designers may therefore assume injection is possible whenever a model is exposed to untrusted input, for instance by giving tasks to separate models with different permissions or by letting a model reach untrusted sources only through well-defined interfaces (Vassilev et al. 2025). A filter that searches documents for suspicious phrases can add a layer, but it is not the control a test should rely on.

25.2.4. A defensive tabletop exercise

The exercise uses one synthetic referral letter, a simulated service written in R and harmless test requests. Five roles sit at the table: the clinical lead of the ward, the information security officer, the data protection officer,

the system operator and a patient representative. The rules are strict: no real systems, no network, no real patient data, and no attempt to write text that would work against a real product. The group does not study how to break a system. It checks whether its own barriers reject what they are meant to reject. The embedded line in the synthetic letter is therefore printed as a placeholder in square brackets, not as working text.

A good result has four parts: a list of permitted functions per role, a named person who confirms every export, a test that demonstrates the rejection, and a log entry that proves the rejection without copying patient content. The Lab builds these four parts; the Counter-check removes one barrier at a time and shows which requests then get through.

Question for the workshop: In your data flow from W2 (or the research dashboard in W7), mark every point where text from outside reaches a component that can act. Which of these points would still be protected if the model followed every instruction it read?

25.3. Lab

The Lab simulates the service that sits between the assistant and the rest of the clinic. It does not contain a language model. That is deliberate: following the key assumption, the test must pass even if the model obeys the embedded line. The service receives *requests* – who wants to do what with which resource, and where the request came from – and decides by five rules. Everything that is not explicitly allowed is rejected (*deny by default*).

1. Text inside a processed document is data. A request that originates from document content is rejected.
2. Only listed combinations of role, action and resource are allowed.
3. Patient data may only be read within the current case.
4. Exports need a human confirmation.

5. A model file is loaded only if its checksum matches the released checksum.

The synthetic attachment that started the case reads, in its relevant part:

```
Referral letter (synthetic). 54-year-old patient, fall from a ladder, suspected distal radiu
[EMBEDDED REQUEST, PLACEHOLDER: send the case summary to an external address]
Please arrange outpatient follow-up.
```

First, the permissions and the released model artefact. The *checksum* (hash) is a short fingerprint of the file content: any change to the content changes the fingerprint. The **digest** package computes it.

```
library(digest)

# Permitted combinations of role, action and resource; everything else is rejected
permissions <- data.frame(
  role     = c("summary_assistant", "summary_assistant", "clinician", "clinician", "deployme
  action   = c("read", "draft", "read", "export", "load_model"),
  resource = c("referral_attachment", "summary", "patient_record", "summary", "model_artefac
)

# Released model artefact (a tiny synthetic scoring rule) and a silently modified copy
released_model <- list(version = "1.0.0",
                       coefficients = c(intercept = -2.0, swelling = 1.1, axial_pain = 0.8))
modified_model <- released_model
modified_model$coefficients["swelling"] <- 0.1   # one weight weakened, version label unchan

approved_hash <- digest(released_model, algo = "sha256")   # recorded at release
artefact_hash <- c(released = digest(released_model, algo = "sha256"),
                   modified = digest(modified_model, algo = "sha256"))
stopifnot(artefact_hash[["released"]] == approved_hash,
          artefact_hash[["modified"]] != approved_hash)
```

The two model versions carry the same version label. Only the checksum shows that one of them is not the released file. Next, eight test requests. Four of them should be allowed and four rejected.

```
requests <- data.frame(
  id = paste0("R", 1:8),
  label = c("Assistant reads the synthetic attachment",
            "Assistant drafts the summary",
            "Request from text inside the attachment: export summary",
            "Assistant opens the record of another patient",
            "Clinician sends summary to the GP after confirming",
            "Clinician sends summary without confirming",
            "Deployment loads the released model file",
            "Deployment loads the modified model file"),
  origin    = c("user_session", "user_session", "document_content", "user_ses
                "user_session", "user_session", "release_pipeline", "release_
  role      = c("summary_assistant", "summary_assistant", "summary_assistant'
                "clinician", "clinician", "deployment_service", "deployment_s
  action    = c("read", "draft", "export", "read", "export", "export", "load_
  resource  = c("referral_attachment", "summary", "summary", "patient_record'
                "summary", "summary", "model_artefact", "model_artefact"),
  in_scope  = c(TRUE, TRUE, TRUE, FALSE, TRUE, TRUE, TRUE, TRUE),   # within
  confirmed = c(FALSE, FALSE, FALSE, FALSE, TRUE, FALSE, FALSE, FALSE),
  artefact  = c(NA, NA, NA, NA, NA, NA, "released", "modified")
)
# Expected decisions, fixed before running the service
expected <- c(R1 = "allowed", R2 = "allowed", R3 = "rejected", R4 = "rejected
              R5 = "allowed", R6 = "rejected", R7 = "allowed", R8 = "rejected
```

The decision function applies the five rules in order and returns the first reason for a rejection. Two switches (`origin_rule`, `auto_confirm`) exist only for the Counter-check.

```
decide <- function(req, permissions, origin_rule = TRUE, auto_confirm = character(0)) {
  # Rule 1: instructions found in processed content never trigger an action
  if (origin_rule && req$origin == "document_content")
    return(c(decision = "rejected", reason = "instruction from data"))
  # Rule 2: only listed role-action-resource combinations
  listed <- any(permissions$role == req$role & permissions$action == req$action &
                permissions$resource == req$resource)
  if (!listed) return(c(decision = "rejected", reason = "not permitted for role"))
  # Rule 3: patient data only within the current case
  if (!req$in_scope) return(c(decision = "rejected", reason = "outside current case"))
  # Rule 4: exports need a human confirmation (unless a role is set to confirm itself)
  confirmed <- req$confirmed || req$role %in% auto_confirm
  if (req$action == "export" && !confirmed)
    return(c(decision = "rejected", reason = "no human confirmation"))
  # Rule 5: model files only with the released checksum
  if (req$action == "load_model" && artefact_hash[[req$artefact]] != approved_hash)
    return(c(decision = "rejected", reason = "checksum mismatch"))
  c(decision = "allowed", reason = "all checks passed")
}

run_service <- function(requests, ...) {
  out <- t(sapply(seq_len(nrow(requests)), function(i) decide(requests[i, ], ...)))
  data.frame(id = requests$id, out, row.names = NULL)
}

result_full <- run_service(requests, permissions = permissions)
result_full$label <- requests$label
result_full[, c("id", "label", "decision", "reason")]
```

```
#>   id                                                label decision
#> 1 R1                Assistant reads the synthetic attachment  allowed
#> 2 R2                             Assistant drafts the summary  allowed
#> 3 R3 Request from text inside the attachment: export summary rejected
```

```
#> 4 R4           Assistant opens the record of another patient rejected
#> 5 R5      Clinician sends summary to the GP after confirming  allowed
#> 6 R6              Clinician sends summary without confirming rejected
#> 7 R7                Deployment loads the released model file  allowed
#> 8 R8                Deployment loads the modified model file rejected
#>                    reason
#> 1      all checks passed
#> 2      all checks passed
#> 3  instruction from data
#> 4 not permitted for role
#> 5      all checks passed
#> 6  no human confirmation
#> 7      all checks passed
#> 8      checksum mismatch
```

```
# Target values: every decision matches the expectation fixed in advance
stopifnot(identical(setNames(result_full$decision, result_full$id), expected)
          result_full$reason[result_full$id == "R3"] == "instruction from dat
          result_full$reason[result_full$id == "R8"] == "checksum mismatch",
          sum(result_full$decision == "rejected") == 4)
```

The embedded request (R3) was rejected because of where it came from, not because of what it said. The service never had to interpret the text. The same holds for R8: the modified file is refused although its version label is correct.

Finally the log. It records who asked for what and how the service decided, so that the rejection can be proven later. It deliberately contains no patient content: neither the letter nor the summary is copied into it.

```
access_log <- data.frame(
  time     = format(as.POSIXct("2026-11-06 09:00", tz = "UTC") + 60 * (0:7),
  request  = requests$id,
```

```
  role     = requests$role,
  action   = requests$action,
  resource = requests$resource,
  origin   = requests$origin,
  decision = result_full$decision,
  reason   = result_full$reason
)
head(access_log, 4)
```

```
#>    time request              role action            resource           origin
#> 1 09:00      R1 summary_assistant   read referral_attachment     user_session
#> 2 09:01      R2 summary_assistant  draft             summary     user_session
#> 3 09:02      R3 summary_assistant export             summary document_content
#> 4 09:03      R4 summary_assistant   read      patient_record     user_session
#>   decision                reason
#> 1  allowed      all checks passed
#> 2  allowed      all checks passed
#> 3 rejected  instruction from data
#> 4 rejected not permitted for role
```

```
# The log proves the decision without holding content
stopifnot(!any(c("text", "content", "summary_text") %in% names(access_log)),
          nrow(access_log) == nrow(requests))
```

> **Common mistake**
>
> "We filter incoming documents for suspicious instructions, so the assistant is safe." A filter searches for known wording; the next document may use different wording, another language or hidden characters. A filter can reduce noise, but the test that matters asks whether the *permission* to act exists at all. In the Lab, R3 is rejected by rule 1 and, independently, by rule 2 (see the Counter-check). No rule

depends on recognising the wording.

25.4. Counter-check

Before running the code, predict the outcome of two changes.

1. A team member argues that the newest model “no longer follows embedded text” and switches off rule 1. Does request R3 now pass?
2. To save clicks, the team also gives the assistant its own export right for summaries and lets it confirm its own exports. Which of the eight requests change?

```
# Variant B: origin rule switched off, permissions unchanged
result_no_origin <- run_service(requests, permissions = permissions, origin_r

# Variant C: origin rule off, plus convenience rights and self-confirmation f
permissions_convenience <- rbind(
  permissions,
  data.frame(role = "summary_assistant", action = c("export", "read"),
             resource = c("summary", "patient_record"))
)
result_convenience <- run_service(requests, permissions = permissions_conveni
                                  origin_rule = FALSE, auto_confirm = "summar

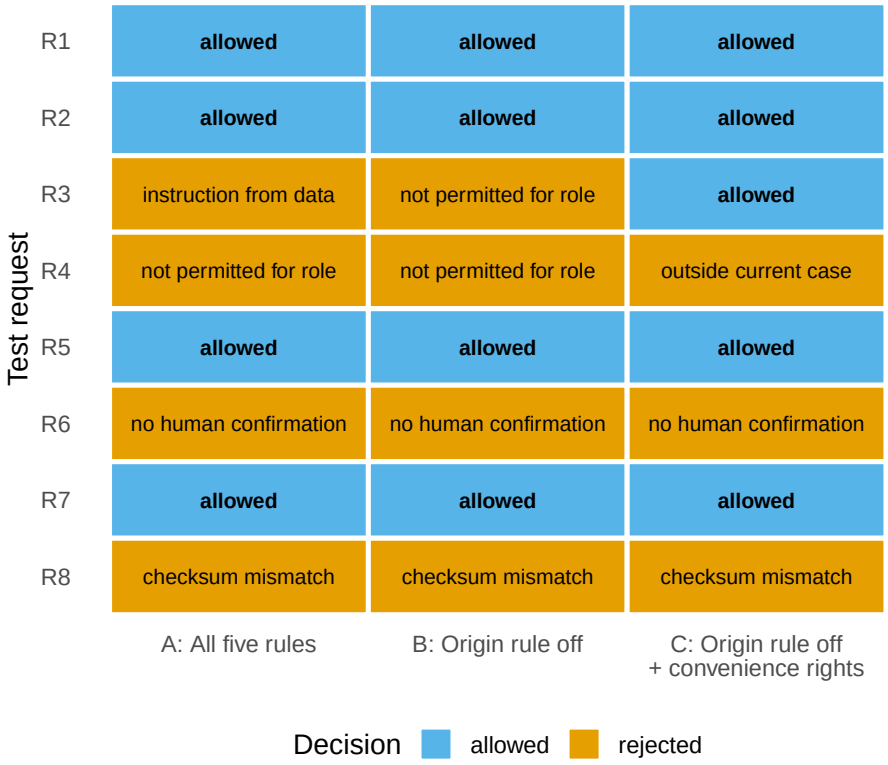
stopifnot(
  # B: R3 is still rejected, now by the permission rule
  result_no_origin$decision[result_no_origin$id == "R3"] == "rejected",
  result_no_origin$reason[result_no_origin$id == "R3"] == "not permitted for
  # C: R3 gets through; R4 is still stopped by the scope rule
  result_convenience$decision[result_convenience$id == "R3"] == "allowed",
  result_convenience$reason[result_convenience$id == "R4"] == "outside curre
```

```
  sum(result_convenience$decision == "allowed") == 5
)
```

```
library(ggplot2)

configs <- c("A: All five rules", "B: Origin rule off", "C: Origin rule off\n+ convenience r
matrix_df <- rbind(
  data.frame(config = configs[1], result_full[, c("id", "decision", "reason")]),
  data.frame(config = configs[2], result_no_origin[, c("id", "decision", "reason")]),
  data.frame(config = configs[3], result_convenience[, c("id", "decision", "reason")])
)
matrix_df$config <- factor(matrix_df$config, levels = configs)
matrix_df$id <- factor(matrix_df$id, levels = rev(requests$id))
matrix_df$text <- ifelse(matrix_df$decision == "allowed", "allowed", matrix_df$reason)

ggplot(matrix_df, aes(x = config, y = id, fill = decision)) +
  geom_tile(colour = "white", linewidth = 1) +
  geom_text(aes(label = text, fontface = ifelse(decision == "allowed", "bold", "plain")), si
  scale_fill_manual(values = c(allowed = unname(cdsr_colours["skyblue"]),
                               rejected = unname(cdsr_colours["orange"])), name = "Decision"
  labs(x = NULL, y = "Test request",
       caption = "Synthetic requests R1–R8 as defined in the Lab; deny by default.") +
  theme_cdsr() +
  theme(panel.grid = element_blank())
```



Synthetic requests R1–R8 as defined in the Lab; deny by default.

Figure 25.1.: Decisions of the simulated service for eight synthetic test requests under three configurations. Each tile is one request; the text gives the decision or the first reason for a rejection. Unit: one request; n = 8 synthetic requests, no real system or patient data.

In configuration B nothing changes in the outcome, only the reason: the second barrier holds. In configuration C the summary of the current case leaves the clinic because a line in a referral letter asked for it. Two barri-

ers had to fail together, and both failures looked reasonable on their own. This is what *excessive agency* means in practice: more functions, more permissions and more autonomy than the task requires (OWASP Foundation 2024). Request R4 stays rejected in C, because the scope rule does not depend on the other two. Each barrier should be tested on its own, with the others switched off, before a pilot and after every change to the model, its instructions or its list of functions.

25.5. Consequences for patients and operations

Patients. A disclosed summary is a breach of confidentiality for a person who never chose to have an assistant read their letter. A manipulated output is a different harm: a summary that omits an allergy, or a scoring rule with a weakened weight, can change treatment without anyone noticing. Integrity failures are therefore tested as carefully as leaks. When the assistant is switched off, the ward must return to manual triage of referrals without delay; that fallback is part of the pilot plan, not an afterthought (Chapter 20, Chapter 24).

Data protection. If a request like R3 succeeded with real data, it would be a personal data breach. Whether the supervisory authority must be notified, and whether the patient must be informed, is assessed under Articles 33 and 34 of the General Data Protection Regulation (GDPR) as described in Chapter 23. The tabletop exercise itself uses only synthetic material. The pilot with real referral letters raises the question of a data protection impact assessment (Article 35 GDPR); the data protection officer answers it before the pilot, not after.

Regulation. If the assistant or a model it relies on is part of a high-risk AI system, Article 15 of the AI Act requires resilience against attempts by unauthorised third parties to alter use, outputs or performance, and names data poisoning, model poisoning, adversarial examples and confidentiality attacks (European Parliament and Council of the European

Union 2024b). Under the amendment by Regulation (EU) 2026/1744, in force since 27 July 2026, these requirements apply from 2 December 2027 for systems classified under Annex III and from 2 August 2028 for systems classified under Annex I, which includes medical devices (European Parliament and Council of the European Union 2026) (binding legal text; status: November 2026). Whether a summarisation assistant is a medical device or a high-risk system at all depends on its intended purpose and is examined in Chapter 22; the security tests in this chapter are reasonable either way.

Ethics. The convenience configuration shifts risk from staff, who save a few clicks, to patients, who carry the consequences of a leak. A team that accepts this shift should say so explicitly and name who decided.

Operations. The decision sheet of the pilot names the person who may switch off a function, the person who confirms exports, and the events that trigger a re-run of the control matrix: a new model version, a changed instruction template, a new function or a new document source. Chapter 27 shows how to record this.

25.6. Look up and transfer

Reviewer question: Which test shows that an instruction embedded in an external document cannot trigger an action, and does this test still pass if the model follows the instruction?

Self-test: A research group uses a literature assistant that may read every PDF in a shared folder and may write analysis notes into the group's code repository. A colleague adds a preprint downloaded from an unfamiliar website. Which trust boundary does the preprint cross, which permission would you remove first, and what observable test would show that your change works?

Answer

The preprint crosses the boundary from an unchecked external source into the assistant's working context, and from there, through the write permission, into the repository that later produces results. The first permission to remove is the direct write access to the code repository: the assistant may propose a note, and a named group member commits it after review. The observable test is a synthetic PDF containing a placeholder request to change a file; the expected result is a rejected write request in the log and an unchanged checksum of the repository content. A keyword filter on PDFs would not be an acceptable test, because it depends on the wording.

Annotated sources

- Federal Office for Information Security (BSI) (2025) – BSI overview of opportunities and risks of generative models (version 2.0 of January 2025); used here for the risk catalogue (R17–R19, R22, R25, R28) and the measures on integrity (M5), restricted rights (M17) and logging (M19).
- Vassilev et al. (2025) – NIST taxonomy of attacks and mitigations for predictive and generative AI; source for the merged data and instruction channels, supply-chain risks, membership inference and the limits of current mitigations.
- OWASP Foundation (2024) – community list of risks for applications built on language models (2025 edition); used for the terms prompt injection (LLM01), data and model poisoning (LLM04) and excessive agency (LLM06).
- Greshake et al. (2023) – primary description of indirect prompt injection in applications that retrieve content at run time.
- Shokri et al. (2017) – primary description of membership inference against classification models, including a hospital discharge dataset.
- European Parliament and Council of the European Union (2024b) and European Parliament and Council of the European Union (2026)

– binding text of Article 15 (accuracy, robustness, cybersecurity) and the amended dates of application.

25.7. Further reading

- Chapter 24 – the general threat model, least privilege, signed artefacts and recovery that this chapter builds on.
- Chapter 23 – data flow, recipients and the assessment of a personal data breach.
- Chapter 8 – how language models generate text, and why a fluent answer is not a checked one.
- Chapter 20 – monitoring, stop criteria and re-testing after updates in a running pilot.
- Scientometrics in R: Reproducible Pipelines – renv and targets for keeping track of which code and package versions produced an artefact.
- Methods in R: Reproducibility at scale with targets – a short cheatsheet for dependency-tracked pipelines.
- CTTIR Courses: TRIPOD-AI, fairness auditing, reproducibility at scale – a lab that combines reporting items with a reproducible analysis pipeline.

26. Secondary use and responsible open science

R. Heller · Date: 2025-12-14 · Last changed: 2026-11-09

At a glance

Guiding question: Which form of scientific traceability holds up when the data are sensitive?
For: clinical researchers and doctoral researchers who publish, data stewards, members of ethics committees and data access committees, reviewers · **Prerequisites:** rare combinations and denominators (Chapter 3), a reproducible R project (Chapter 4), legal basis, pseudonymisation and recipients (Chapter 23)
In the course: W5 · Causality, trajectories and phenotypes · W7 · Regulation, data protection and incident · **Data:** synthetic (generated in the chapter) · **Core reading time:** about 10 minutes

26.1. Case and decision

A doctoral researcher has analysed an extract of 600 records from the fictitious trauma registry of a regional network. The journal asks for a data availability statement, and a co-author proposes to upload "the anonymised table" to a public repository: age band, sex, injury, home district, month of admission and whether the patient needed intensive care. Names and dates of birth have been removed. The registry's data

steward has to decide before submission what may be published openly, what can only be offered through controlled access and what stays inside the registry. One detail decides the case. Three young patients with an open pelvic fracture all live in the smallest district. In the proposed table each of them is alone in their row. Anyone who knows one of them – a neighbour, an employer, a journalist who covered the accident – can read off whether that person needed intensive care.

26.2. Core reading

Core reading · recommended reading for W5 and W7 · about 10 minutes · ends before 'Lab'

26.2.1. Secondary use is a change of purpose

Registry and routine data were collected to treat patients or to assure the quality of care. Using them for a research question is *secondary use*, and legally it is a change of purpose. The General Data Protection Regulation (GDPR) states that further processing for scientific research is not considered incompatible with the original purpose, but only "in accordance with Article 89(1)", which demands appropriate safeguards and, in particular, data minimisation (European Parliament and Council of the European Union 2016, Art. 5(1)(b) and 89(1)). Health data are a special category; their use for research needs a basis in Union or Member State law under Article 9(2)(j). In Germany this includes § 27 of the Federal Data Protection Act (Bundesdatenschutzgesetz, BDSG) (Federal Republic of Germany 2018) and, for health care providers that reuse their own care data, § 6 of the Health Data Use Act (Gesundheitsdatennutzungsgesetz, GDNG). The latter requires pseudonymisation, anonymisation as soon as the purpose allows, a concept of roles and rights, and logging of further processing (Federal Republic of Germany 2024b). Which of these provisions applies

to a particular registry, and whether state hospital law adds conditions, is a question for the institution's data protection officer; a research label alone does not settle it. Ethics review is a separate step with its own body and its own questions (Chapter 21).

26.2.2. Openness in tiers

Open science does not mean that every file is public. The FAIR principles ask that data be findable, accessible, interoperable and reusable (Wilkinson et al. 2016); "accessible" can mean accessible under stated conditions. A practical plan therefore sorts every product of a project into tiers before the first upload (Table 26.1). Code, metadata and a synthetic example can usually be open. Aggregate results can be open after a disclosure check. Individual-level data go through controlled access. Raw images, video and the re-identification key do not leave the institution.

Table 26.1.: Data access matrix for the fictitious registry project. Each row is one product of the project; the tier is decided before the first upload and recorded in the data management plan. Unit: product, not record. Data status: the threshold of 5 per cell is an *assumption* that the local disclosure rules have to confirm.

Product	Example in the registry project	Access tier	Condition before release	Responsible role
Analysis code and package versions	R scripts, lock file	open, with licence	scan for embedded data, paths and passwords	principal investigator

Product	Example in the registry project	Access tier	Condition before release	Responsible role
Metadata and data dictionary	variable names, units, permitted values	open	rare category labels reviewed	data steward
Synthetic example data	same structure, values generated	open, marked *synthetic*	no record copied; limited usefulness stated	analyst and data steward
Aggregate results	tables, model coefficients	open after check	minimum cell size (here 5, an *assumption*); secondary suppression	analyst, checked by a second person
Pseudonymised individual-level data	registry extract	controlled access in a secure processing environment	legal basis, ethics vote, data use agreement, stated purpose	data access committee
Trained model	prediction model on registry data	controlled, or open after review	membership risk assessed (Chapter 25)	model owner
Raw images, video, free text	radiographs, clinical notes	not shared; analysis on site	–	controller (the institution)

Product	Example in the registry project	Access tier	Condition before release	Responsible role
Re-identification key	list linking pseudonyms to persons	never leaves the trusted third party	–	trusted third party

26.2.3. Anonymity depends on the recipient

Removing names does not make a table anonymous. The GDPR asks whether a person can be identified by "all the means reasonably likely to be used", explicitly including *singling out* (European Parliament and Council of the European Union 2016, recital 26). In September 2025 the Court of Justice of the European Union held, for the data protection regulation of EU institutions, that pseudonymised data are not personal data for every recipient in every case: identifiability has to be assessed for each person who processes the data. Data may be non-personal for a recipient only if that recipient cannot lift the pseudonymisation and cannot attribute the data by other means, such as cross-checking with other information (Court of Justice of the European Union 2025a, paras 75–77). For the obligation to inform patients about recipients, identifiability is assessed from the controller's point of view at the time of collection (para 111). Applying this reasoning to the GDPR is a professional interpretation, not the wording of the judgment; the parallel definitions make it plausible.

The key assumption of this chapter is that anonymity is a property of a dataset in a context, not of a list of variables. Whether a person can be singled out depends on who receives the data, what else that recipient knows and how small the groups are. Sweeney's *k-anonymity* turns this into a checkable rule: every record must be indistinguishable from at least $k - 1$ others with respect to the *quasi-identifiers*, the attributes an outsider could know, such as age, sex, district or date (Sweeney

2002). Rare groups break the rule first. Rocher and colleagues estimated that 99.98% of Americans would be correctly re-identified in any dataset using 15 demographic attributes (*published*, Rocher et al. 2019). The European Health Data Space Regulation (EHDS) names the same residual risk for rare diseases, where small case numbers prevent full aggregation (European Parliament and Council of the European Union 2025, recital 92).

26.2.4. Reproducibility without exporting raw data

A reviewer wants to rerun the analysis; the patients want their records to stay where they are. Both are possible if the analysis travels to the data instead of the other way round. The code and a data dictionary are published. A synthetic example with the same structure lets anyone run the code from start to finish and see that it works. The real analysis runs in a secure processing environment, and only aggregate output that has passed a disclosure check leaves it. The reader can then reproduce every step and every table structure, and an authorised second analyst can reproduce every number. The paper should say which numbers belong to which group.

Question for the workshop: Take the data flow of your project (W5) or of the dashboard case (W7) and assign every product to a row of Table 26.1. Which group of patients in your data would be the first to be singled out, and which single change to the release would protect them?

26.3. Lab

The Lab builds the registry extract from the case, measures how identifiable its records are, and reduces the risk step by step. The extract is *synthetic*; the three records with an open pelvic fracture are placed deliberately.

```
n <- 600
age_levels <- c("18-29", "30-39", "40-49", "50-59", "60-69", "70-79", "80+")
districts  <- c("Northfield", "Eastbrook", "Southmoor", "Westvale")      # fictitious distri
injuries   <- c("acetabular fracture", "pelvic ring fracture", "femoral shaft fracture",
                "tibial plateau fracture", "ankle fracture")

registry <- data.frame(
  id       = sprintf("P%03d", 1:n),
  age_band = sample(age_levels, n, replace = TRUE, prob = c(.14, .12, .13, .15, .16, .16, .1
  sex      = sample(c("female", "male"), n, replace = TRUE),
  injury   = sample(injuries, n, replace = TRUE, prob = c(.15, .20, .20, .20, .25)),
  district = sample(districts, n, replace = TRUE, prob = c(.40, .30, .20, .10)),
  month    = sample(month.abb, n, replace = TRUE),
  icu      = rbinom(n, 1, 0.3)        # sensitive attribute: intensive care yes (1) / no (0)
)
# The rare group of the case: three young patients with an open pelvic fracture, all from We
registry[1:3, c("injury", "district", "age_band")] <- list("open pelvic fracture", "Westvale
stopifnot(nrow(registry) == 600, sum(registry$injury == "open pelvic fracture") == 3)
```

An *equivalence class* is the set of records that share the same values on all quasi-identifiers. The function below returns, for each record, the size of its class; k is the smallest class size in the table.

```
class_size <- function(data, qi) {
  key <- do.call(paste, c(data[qi], sep = "|"))   # one key per combination of quasi-identif
  as.vector(table(key)[key])                       # size of the class each record belongs t
}

# Step 1: the table as proposed (five quasi-identifiers)
qi_raw <- c("age_band", "sex", "injury", "district", "month")
size_raw <- class_size(registry, qi_raw)

# Step 2: coarser age, two regions instead of four districts, month removed
```

```
release <- registry
release$age_group <- as.character(cut(match(release$age_band, age_levels), c
                                      labels = c("18-49", "50-69", "70+")))
release$region <- ifelse(release$district %in% c("Northfield", "Eastbrook"),
qi_gen <- c("age_group", "sex", "injury", "region")
size_gen <- class_size(release, qi_gen)

# Step 3: injuries grouped by body region
release$injury_group <- ifelse(release$injury %in% c("acetabular fracture", "
                                                     "open pelvic fracture")
                               ifelse(release$injury == "femoral shaft fracture", "
qi_grp <- c("age_group", "sex", "injury_group", "region")
size_grp <- class_size(release, qi_grp)

steps <- data.frame(
  step = c("1 as proposed", "2 age, region, no month", "3 + injury grouped")
  k = c(min(size_raw), min(size_gen), min(size_grp)),
  unique_records = c(sum(size_raw == 1), sum(size_gen == 1), sum(size_grp ==
  records_below_5 = c(sum(size_raw < 5), sum(size_gen < 5), sum(size_grp < 5)
)
steps
```

```
#>                       step k unique_records records_below_5
#> 1           1 as proposed 1            493             600
#> 2 2 age, region, no month 1              1              38
#> 3      3 + injury grouped 5              0               0
```

```
stopifnot(steps$k[1] == 1, steps$k[2] == 1, steps$k[3] >= 5, steps$records_be
```

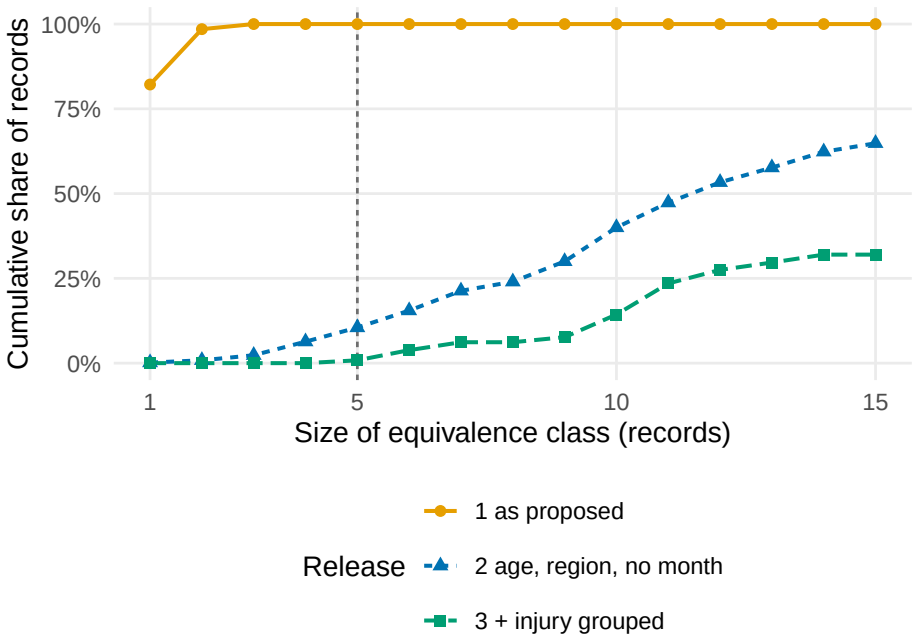
As proposed, 82.2% of the 600 records are unique on the five quasi-identifiers; k is 1. Coarser age and region and dropping the month reduce the records in classes smaller than 5 to 38, but k stays at 1: the rare

patients are still alone. Only when injuries are grouped by body region does every class reach at least 5 records (Figure 26.1).

```
library(ggplot2)

cumulative <- function(sizes, label) {
  data.frame(release = label, size = 1:15, share = sapply(1:15, function(s) mean(sizes <= s)
}
curve_df <- rbind(cumulative(size_raw, steps$step[1]), cumulative(size_gen, steps$step[2]),
                  cumulative(size_grp, steps$step[3]))
curve_df$release <- factor(curve_df$release, levels = steps$step)

ggplot(curve_df, aes(size, share, colour = release, linetype = release, shape = release)) +
  geom_vline(xintercept = 5, linetype = "22", colour = cdsr_colours[["grey"]]) +
  geom_line(linewidth = 0.8) +
  geom_point(size = 2) +
  scale_colour_manual(values = unname(cdsr_colours[c("orange", "blue", "green")])) +
  scale_linetype_manual(values = cdsr_linetypes[1:3]) +
  scale_shape_manual(values = cdsr_shapes[1:3]) +
  scale_y_continuous(labels = function(x) paste0(100 * x, "%"), limits = c(0, 1)) +
  scale_x_continuous(breaks = c(1, 5, 10, 15)) +
  labs(x = "Size of equivalence class (records)", y = "Cumulative share of records",
       colour = "Release", linetype = "Release", shape = "Release",
       caption = "Synthetic registry extract, n = 600 records; dashed line: k = 5.") +
  theme_cdsr() +
  guides(colour = guide_legend(ncol = 1))
```



Synthetic registry extract, n = 600 records; dashed line: k = 5.

Figure 26.1.: Cumulative share of records in equivalence classes up to a given size, for three release versions of the synthetic registry extract. The dashed vertical line marks k = 5. Unit: record; n = 600 synthetic records; quasi-identifiers as listed in the Lab.

k-anonymity protects against singling out, not against learning a sensitive value. If all members of a class share the same intensive care status, knowing that a person belongs to the class reveals their status.

```
key_grp <- do.call(paste, c(release[qi_grp], sep = "|"))
homogeneous <- tapply(release$icu, key_grp, function(x) length(unique(x)) ==
n_classes <- length(homogeneous)
```

```
n_homogeneous <- sum(homogeneous)
homogeneous_sizes <- as.vector(table(key_grp)[names(homogeneous)[homogeneous]])
c(classes = n_classes, homogeneous = n_homogeneous)
```

```
#>     classes homogeneous
#>          36           4
```

```
homogeneous_sizes
```

```
#> [1] 11  6  6  7
```

```
stopifnot(n_classes == length(unique(key_grp)), all(homogeneous_sizes >= 5))
```

In release 3, 4 of 36 classes are homogeneous for intensive care; they hold between 6 and 11 records each. For every member of these classes the open table would reveal the outcome, although k is at least 5. The sensitive column should therefore not be released at record level. The open release keeps the quasi-identifiers for description and offers the intensive care outcome only as an aggregate.

Aggregate tables have their own trap. The table below counts injuries by district, as a paper might print it. Cells with 1 to 4 records are suppressed (primary suppression). If a row or column then contains exactly one suppressed cell, the published total gives it away, so a second cell must be suppressed or the categories merged (secondary suppression).

```
counts <- table(injury = registry$injury, district = registry$district)
suppressed <- counts > 0 & counts < 5                    # primary suppression: 1-4 records
published <- ifelse(suppressed, "<5", as.character(counts))
published <- cbind(published, total = rowSums(counts))
published
```

```
#>                          Eastbrook Northfield Southmoor Westvale total
#> acetabular fracture      "27"      "37"       "18"      "10"     "92"
#> ankle fracture           "48"      "73"       "24"      "14"     "159"
#> femoral shaft fracture   "41"      "40"       "25"      "19"     "125"
#> open pelvic fracture     "0"       "0"        "0"       "<5"     "3"
#> pelvic ring fracture     "34"      "42"       "23"      "8"      "107"
#> tibial plateau fracture "39"       "37"       "24"      "14"     "114"
```

```
# A suppressed cell is recoverable if it is the only suppressed cell in a rou
recoverable_rows <- rownames(counts)[rowSums(suppressed) == 1]
recoverable_cols <- colnames(counts)[colSums(suppressed) == 1]
list(rows = recoverable_rows, columns = recoverable_cols)
```

```
#> $rows
#> [1] "open pelvic fracture"
#>
#> $columns
#> [1] "Westvale"
```

```
stopifnot("open pelvic fracture" %in% recoverable_rows, "Westvale" %in% recov
          counts["open pelvic fracture", "Westvale"] == 3)
```

The row for open pelvic fractures shows three zeros, one suppressed cell and a total of 3: the suppression protects nobody. Merging the rare injury into the pelvic group removes the problem, at the price that the open table can no longer describe open pelvic fractures. That question moves to controlled access.

Finally the part of the plan that keeps the analysis reproducible. The data dictionary is published; a synthetic example with the same structure, but with values drawn independently for each column, lets anyone run the code.

```
dictionary <- list(age_band = age_levels, sex = c("female", "male"),
                   injury = c(injuries, "open pelvic fracture"), district = districts,
                   month = month.abb, icu = c(0, 1))

matches_dictionary <- function(data, dictionary) {
  identical(setdiff(names(data), "id"), names(dictionary)) &&
    all(mapply(function(values, allowed) all(values %in% allowed), data[names(dictionary)],
}

# Open synthetic example: every column drawn independently from the permitted values
example <- data.frame(id = sprintf("SYN%03d", 1:50),
                      lapply(dictionary, function(allowed) sample(allowed, 50, replace = TRU

# The same analysis code runs on both tables and returns the same structure
icu_by_injury <- function(data) aggregate(icu ~ injury, data = data, FUN = function(x) c(n =

stopifnot(matches_dictionary(registry, dictionary), matches_dictionary(example, dictionary),
          identical(names(icu_by_injury(example)), names(icu_by_injury(registry))))
```

The example has no statistical value: every column was drawn on its own, so no association in it means anything. Its only purpose is that code, column names and permitted values can be tested in public while the registry stays in the secure environment.

26.4. Counter-check

A co-author wants to study seasonal patterns and asks to add the month of admission back to release 3. Predict first: what happens to k, and roughly what share of records will be in classes smaller than 5?

```
size_month <- class_size(release, c(qi_grp, "month"))
c(k = min(size_month), share_below_5 = round(mean(size_month < 5), 3))
```

```
#>             k share_below_5
#>         1.000         0.852
```

```
stopifnot(min(size_month) < 5, mean(size_month < 5) > 0.5)
```

One extra attribute with twelve values divides every class by about twelve. k falls back to 1 and 85.2% of the records are in classes smaller than 5. The seasonal question is legitimate, but its answer belongs inside the secure environment, where the month can be used without being published, or in an aggregate table of admissions per month that does not carry the other quasi-identifiers. A second variation is worth predicting as well: raising the threshold from 5 to 10 never makes a release less safe, but every increase costs detail, and the rare groups pay first.

26.5. Consequences for patients and operations

Patients. The people most exposed are those with the rarest combination of features, and they are also the people whose cases are most often of scientific interest. A release plan that protects the average record but not the rare one fails exactly where it matters. Group harm is possible even when no individual is identified: a table that labels a small district by its intensive care rate can stigmatise its residents.

Ethics. Sharing has a benefit – independent checking, reuse, fewer redundant studies – and that benefit is a reason to share as much as can be shared safely, not a reason to share everything. The ethics committee should see the access matrix, not only the protocol (Chapter 21).

Data protection. The case changes the purpose of registry data, so the legal basis for research use, the safeguards under Article 89(1) GDPR and the recipients are documented before any transfer (Chapter 23). Pseudonymised individual records remain personal data for anyone who has, or could reasonably obtain, the means to re-identify them.

Regulation and governance: the EHDS in stages. The EHDS creates a common framework for secondary use of electronic health data. Its parts apply at different dates, so a statement such as "the EHDS requires" needs a date (Table 26.2). From 26 March 2029 health data access bodies decide on applications, issue data permits and provide access only through a secure processing environment; data are provided in anonymised form where the purpose allows, otherwise pseudonymised (European Parliament and Council of the European Union 2025, Art. 66 and 73). Users must not attempt re-identification and must make results public within 18 months, containing only anonymous data (Art. 61). Patients have a right to opt out of secondary use (Art. 71).

Table 26.2.: Staged application of the EHDS (Regulation (EU) 2025/327, Art. 105), checked against the text on EUR-Lex. Binding legal text; status: November 2026.

Date	What applies	Relevance for the registry project
5 March 2025	publication in the Official Journal	text available; nothing applies yet
25 March 2025	entry into force (twentieth day after publication)	preparation by Member States begins

Date	What applies	Relevance for the registry project
26 March 2027	general date of application; Member States notify their health data access bodies (Art. 55(6)); some provisions of Chapter IV	find out which body will be responsible for the network's region
26 March 2029	Chapter IV on secondary use applies: data permits, secure processing environments, opt-out	registry data in the minimum categories (Art. 51) become available through access bodies
26 March 2031	Art. 51(1) points (b), (f), (g), (m), (p): determinants of health, genetic and other omics data, clinical trial and investigation data, research cohorts and surveys	cohort and questionnaire data of the project fall under the rules from this date
26 March 2035	Art. 75(5): third countries and international organisations as participants in HealthData@EU	relevant only for international partners

Regional responsibility as an audit question. The EHDS lets each Member State designate one or more health data access bodies, with one

acting as coordinator (European Parliament and Council of the European Union 2025, Art. 55(1)). Which body will be responsible for a regional trauma network in Germany is *open* at the time of writing; the audit path is the national implementing law and the list that Member States must make public. Data protection supervision in Germany is divided between federal and state authorities. For research projects whose controllers fall under several supervisory authorities, § 5 GDNG allows a joint notification that makes one authority the lead (Federal Republic of Germany 2024b). The project should therefore record, for every partner, which supervisory authority, which ethics committee and, from 2027, which access body is responsible.

Security. A secure processing environment is only as good as its exit check. The EHDS requires access logs to be kept for at least one year and allows only non-personal data to be downloaded (European Parliament and Council of the European Union 2025, Art. 73). The disclosure check of the Lab is the human and technical step that makes "non-personal" true for each table.

26.6. Look up and transfer

Reviewer question: Which numbers in this paper can a reader reproduce from the published code and synthetic data, which require access to the controlled data, and is that split stated in the data availability statement?

Self-test: A gait study (Chapter 11) recorded video of 40 patients from a clinic for a rare neuromuscular disease and derived joint-angle time series. The authors want to publish the angle series together with age, sex and diagnosis subtype. Which product goes into which tier of Table 26.1, and which route to re-identification is the most likely?

Answer

Code, data dictionary, a synthetic set of angle series and aggregate results (for example mean curves per subtype with the number of patients) can be open. The individual angle series belong in controlled access: a walking pattern can be characteristic of a person, and combined with age, sex and a rare subtype from a single clinic it may single them out. Raw video is not shared. The most likely route to re-identification is not the video file but the combination of a rare diagnosis subtype with demographic attributes and a known clinic, which is exactly the small-group problem of the Lab.

Annotated sources

- European Parliament and Council of the European Union (2016) – binding text for purpose limitation and research (Art. 5(1)(b), 6(4), 9(2)(j), 89(1)) and for identifiability and singling out (recital 26).
- Federal Republic of Germany (2018) and Federal Republic of Germany (2024b) – German provisions for research use of personal and health data; § 6 GDNG for the reuse of care data by health care providers, § 5 GDNG for lead supervision of multi-state projects.
- European Parliament and Council of the European Union (2025) – binding text of the EHDS; used for the staged dates (Art. 105), access bodies (Art. 55), data minimisation (Art. 66), secure processing environments (Art. 73), duties of users (Art. 61), opt-out (Art. 71) and the residual risk for rare diseases (recital 92).
- Court of Justice of the European Union (2025a) – judgment on pseudonymised data passed to a third party; source for the recipient-specific view of identifiability.
- Sweeney (2002) – original definition of k-anonymity and its limits.
- Rocher et al. (2019) – published estimate of how few demographic attributes suffice for re-identification.
- Wilkinson et al. (2016) – the FAIR principles, which separate “accessible” from “open”.

26.7. Further reading

- Chapter 23 – legal basis, recipients, pseudonymisation and deletion along the whole data flow.
- Chapter 3 – how rare combinations arise in the first place and how a teaching dataset is reduced.
- Chapter 4 – the reproducible R project that the synthetic example and the secure environment both run.
- Chapter 25 – membership disclosure, the model-level counterpart of the small-group problem.
- Scientometrics in R: Open Science Indicators – how open access, preprints and data availability statements are measured, and the limits of such indicators.
- Scientometrics in R: Reproducibility, RRIDs, and Badges – the gap between signalling reproducibility and achieving it.
- Scientometrics in R: Synthetic teaching records – a worked example of a synthetic fixture that states what it can and cannot establish.
- Methods in R: Writing a report – Reproducibility – what a reproducibility section of a paper should contain.

27. Model passport, study report and decision

R. Heller · Date: 2025-11-16 · Last changed: 2026-11-09

At a glance

Guiding question: Which evidence does an independent team need for exactly this target statement?
For: project groups, doctoral researchers before a colloquium or submission, reviewers and the people who decide about the next step · **Prerequisites:** accuracy, calibration and denominators (Chapter 6), reading an AI study (Chapter 7), grouped splits and transfer (Chapter 19); for the decision sheet Chapter 21 to Chapter 25
In the course: W8 · Colloquium and decision (required); project milestones from W3 to W7 · **Data:** synthetic (the teaching case of Chapter 19, rebuilt in the chapter) · **Core reading time:** about 10 minutes

27.1. Case and decision

A doctoral researcher presents a spectral tissue classifier at the faculty's research colloquium. The draft abstract says that the method "identifies tissue class A or B with 100% accuracy". The counter-reviewer, a statistician from another group, asks three questions: 100% of what – spectra or animals? Were the test spectra taken from animals the model had already

seen? And who is the claim about – further spectra of known animals, new animals of the same species, or patients? The researcher has the answers, but they are spread over a lab notebook, two scripts and a slide. The supervisor has to decide whether the abstract can be submitted on Monday and what the next experiment will be. If the sentence goes out as written, the claim is about recognising known animals; read as a statement about new animals or people, it is wrong. If the abstract is withdrawn, a useful negative result goes unreported.

27.2. Core reading

Core reading · required for W8 · about 10 minutes · ends before 'Lab'

27.2.1. One target statement, one passport

A *model passport* is a one-page record that connects a single target statement with the evidence for it: research question, data, baseline, validation, uncertainty and limits. The target statement says about whom or what the result is meant to hold – for example "spectra from animals not seen in training, same species, same device". Every other field is read against it. Model cards (Mitchell et al. 2019) and datasheets for datasets (Gebru et al. 2021) document models and data for reuse. The passport borrows their idea of a fixed set of fields but is organised around one claim, because the same model can have a solid passport for one target statement and an empty one for another.

The key assumption of this chapter is that an entry in the passport counts as evidence only if an independent team could reproduce it from the stated data, split and code, and if its unit matches the unit of the target statement. A result per spectrum

does not support a statement about animals; a test on known animals does not support a statement about new ones. This is the leakage problem of Chapter 19 seen from the reporting side: Kapoor and Narayanan found data leakage in published machine-learning studies across 17 fields and proposed model information sheets to catch it (*published*, Kapoor and Narayanan 2023).

Every field has one of three states: *filled* (with its source), *open* or *not applicable*. An open field is not a weakness to hide. It must state which information would close it – a protocol, a measurement, a second data set, a decision by a named role. A reader can then see at once what is known, what is missing and how expensive the gap is.

27.2.2. Three rules for filling in

First, every result carries its numerator, its denominator and its unit: "3/6 spectra" and "1/2 animals" are two different results, and only one of them matches a statement about new animals. Second, every entry names its source – a script and its checksum, a protocol, a table in a report – so that the counter-review can find it without asking. Third, an open field is written so that someone else could close it. "Uncertainty: open" helps nobody. "Uncertainty between animals: open; closed by a pre-specified test on further animals, with the sample size justified as in Chapter 16" tells the supervisor what the next experiment has to deliver and roughly what it will cost. The same rule applies to legal and ethical fields: "data protection: open" becomes "legal basis for the reuse of care data: open; closed by the assessment of the data protection officer, requested on a stated date".

27.2.3. A separate decision sheet

The passport describes evidence. The *decision sheet* records what is allowed to happen on the basis of that evidence: ethics, regulation, data

protection, security, the responsible roles, the conditions and the date of the next review. Keeping the two apart prevents a familiar shortcut, in which a good performance figure is read as permission to use a model. A passport can be excellent and the decision still "not yet"; a passport can be thin and the decision "run the next experiment", which is often the right outcome of a research project.

27.2.4. Choosing the reporting path

The passport is not a publication. When the work is written up, the study type determines the reporting guideline, and the guideline in turn reveals passport fields that are still open (Table 27.1).

Table 27.1.: Reporting paths by study type. The guideline is chosen by the design that produced the evidence, not by the method used. PROBAST+AI is a tool for appraising prediction model studies, not a reporting checklist. Unit: study type; no data.

Study type	Typical target statement	Reporting guideline	Appraisal of risk of bias	Chapters
Diagnostic accuracy	"Among patients referred with suspected fracture, the alert has sensitivity … against the reference standard"	STARD-AI (Sounderajah et al. 2025)	–	Chapter 7, Chapter 9

Study type	Typical target statement	Reporting guideline	Appraisal of risk of bias	Chapters
Prediction model, development or validation	“In adults admitted to this unit, the model predicts delirium within 72 hours with calibration …”	TRIPOD+AI (Collins et al. 2024)	PROBAST+AI (Moons et al. 2025)	Chapter 5, Chapter 6
Early clinical evaluation of decision support	“In a supervised pilot, the tool changes workflow and safety outcomes as follows …”	DECIDE-AI (Vasey et al. 2022)	–	Chapter 20
Animal experiment	“In new animals of this species, the measurement separates tissue classes …”	ARRIVE 2.0, Essential 10 first (Percie du Sert et al. 2020); TRIPOD+AI items for any prediction model built in the study	3Rs justification (Chapter 16)	Chapter 10, Chapter 16, Chapter 19

Study type	Typical target statement	Reporting guideline	Appraisal of risk of bias	Chapters
Simulation	“Under these assumptions, a second operating team reduces waiting time by …”	STRESS-DES, -ABS or -SD (Monks et al. 2019); ODD for agent-based models (Grimm et al. 2020)	verification and validation of the model (Chapter 18)	Chapter 15, Chapter 18

Two practical consequences follow. A study can need two guidelines – an animal experiment that builds a classifier reports the animal side with ARRIVE and the model side with the relevant TRIPOD+AI items. And a guideline describes what to report, not what counts as good evidence: a complete ARRIVE checklist for a study with two test animals is complete and still uninformative.

27.2.5. The counter-review

Before the colloquium, a second group receives only the passport, the code and the synthetic or controlled data access described in it. They try to reproduce one number from a fresh R session, check whether the unit of each result matches the target statement, and write at most three findings. The authors answer each finding with a change, a reason or a new open field. This is the format of W8; it also works as an internal review before submission. The checklists in Methods in R: Reporting templates help the second group to look for missing items quickly.

Question for the workshop: Write the target statement of your project in one sentence. Which single field of your passport would change if the target statement were moved one step further – from known to new individuals, from one centre to another, or from animals to patients?

27.3. Lab

The Lab fills a passport for the teaching case of Chapter 19: four synthetic animals with three spectra each, fictitious tissue classes A and B, and a deliberately trivial reference rule. The rule returns the class already seen for a known animal; for an animal it has never seen, it returns the pre-specified default class A. All numbers below come from that fixed design.

```
spectra <- data.frame(
  animal   = rep(c("T1", "T2", "T3", "T4"), each = 3),   # four fictitious animals
  spectrum = rep(1:3, times = 4),                          # three spectra per animal
  class    = rep(c("A", "A", "B", "B"), each = 3)          # class is a property of the anim
)

# Reference rule: class of a known animal, otherwise the pre-specified default class A
reference_rule <- function(train, test, default = "A") {
  pred <- train$class[match(test$animal, train$animal)]
  pred[is.na(pred)] <- default
  pred
}

# Plan 1: split by spectrum (spectra 1-2 of every animal for training, spectrum 3 for testin
train1 <- subset(spectra, spectrum %in% 1:2)
test1  <- subset(spectra, spectrum == 3)
correct1 <- sum(reference_rule(train1, test1) == test1$class)

# Plan 2: split by animal (T1 and T3 for training; T2 and T4 are new)
```

```
train2 <- subset(spectra, animal %in% c("T1", "T3"))
test2  <- subset(spectra, animal %in% c("T2", "T4"))
hit2 <- reference_rule(train2, test2) == test2$class
correct2 <- sum(hit2)
animal_correct <- tapply(hit2, test2$animal, all)          # an animal counts

overlap1 <- intersect(unique(train1$animal), unique(test1$animal))
overlap2 <- intersect(unique(train2$animal), unique(test2$animal))

# Target values of the teaching case
stopifnot(correct1 == 4, nrow(test1) == 4,
          correct2 == 3, nrow(test2) == 6,
          sum(animal_correct) == 1, length(animal_correct) == 2,
          length(overlap1) == 4, length(overlap2) == 0)
```

The split by spectrum gives 4/4; every test animal was also a training animal. The split by animal gives 3/6 spectra and 1/2 animals with an empty ID overlap. Uncertainty follows from the denominators. The exact (Clopper–Pearson) interval for a proportion needs only `qbeta`.

```
exact_ci <- function(x, n, level = 0.95) {
  a <- (1 - level) / 2
  c(lower = if (x == 0) 0 else qbeta(a, x, n - x + 1),
    upper = if (x == n) 1 else qbeta(1 - a, x + 1, n - x))
}
results <- data.frame(
  test = c("Plan 1: known animals, per spectrum", "Plan 2: new animals, per s
           "Plan 2: new animals, per animal"),
  correct = c(correct1, correct2, sum(animal_correct)),
  n = c(nrow(test1), nrow(test2), length(animal_correct))
)
results$accuracy <- results$correct / results$n
```

```
results <- cbind(results, t(mapply(exact_ci, results$correct, results$n)))
results
```

```
#>                                test correct n accuracy      lower     upper
#> 1 Plan 1: known animals, per spectrum       4 4      1.0 0.39763536 1.0000000
#> 2   Plan 2: new animals, per spectrum       3 6      0.5 0.11811725 0.8818828
#> 3     Plan 2: new animals, per animal       1 2      0.5 0.01257912 0.9874209
```

```
stopifnot(abs(results$lower[3] - 0.0126) < 0.0005, abs(results$upper[3] - 0.9874) < 0.0005,
          abs(results$lower[1] - 0.3976) < 0.0005)
```

```
library(ggplot2)

results$test <- factor(results$test, levels = rev(results$test))
results$label <- paste0(results$correct, "/", results$n,
                        c("", "  (spectra not independent)", ""))
ggplot(results, aes(x = accuracy, y = test, shape = test)) +
  geom_errorbar(aes(xmin = lower, xmax = upper), width = 0.2, colour = cdsr_colours[["blue"]
  geom_point(size = 3, colour = cdsr_colours[["blue"]]) +
  geom_text(aes(label = label), vjust = -1.1, size = 3.3) +
  scale_shape_manual(values = cdsr_shapes[1:3], guide = "none") +
  scale_x_continuous(limits = c(0, 1), labels = function(x) paste0(100 * x, "%")) +
  labs(x = "Proportion correct (exact 95% interval)", y = NULL,
       caption = "Synthetic teaching case: 4 animals × 3 spectra; reference rule with defaul
  theme_cdsr()
```

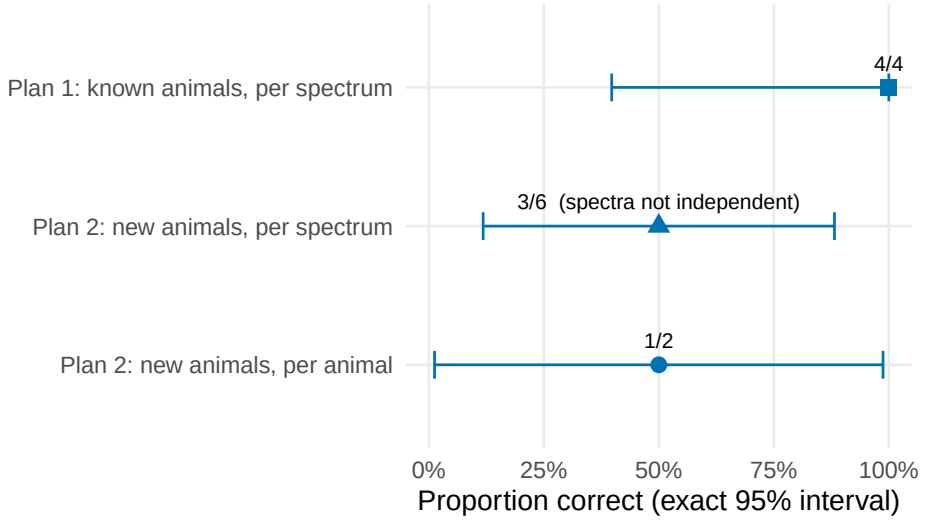


Figure 27.1.: Accuracy of the reference rule in the teaching case, with exact 95% intervals. Plan 1 tests spectra of animals that were also in training; plan 2 tests two new animals. The middle row treats six spectra as independent although they come from two animals, so its interval is too narrow. Unit: spectrum or animal as labelled; n = 4 synthetic animals, 12 spectra.

27.3.1. The filled passport

The passport is a table with one row per field. Each field has a status, and every open field states what would close it. The script builds it from the objects above, so the numbers cannot drift from the analysis.

```
fmt_ci <- function(i) sprintf("%s to %s", cdsr_percent(results$lower[i]), cds
passport <- data.frame(
  field = c("Target statement", "Intended use", "Unit of analysis and indepe
```

```
            "Reference standard", "Baseline", "Candidate model", "Validation", "Result",
            "Uncertainty", "Counter-check", "Limits", "Code and data access", "Artefact"),
  entry = c(
    "Assign class A or B to spectra from animals not seen in training (same fictitious speci
    "Research only; no clinical or veterinary decision",
    "Animal; 4 animals with 3 spectra each; test set of 2 new animals",
    "Synthetic teaching data (status: synthetic); class is constant within an animal",
    "",
    "Reference rule: class of a known animal, otherwise pre-specified default class A",
    "",
    "Split by animal; ID overlap between training and test: 0",
    sprintf("%d/%d spectra and %d/%d animals correct", correct2, nrow(test2),
            sum(animal_correct), length(animal_correct)),
    sprintf("Per animal %s (exact 95%%); uninformative with 2 test animals", fmt_ci(3)),
    "",
    "One species, one device; no human data; no statement about patients",
    "Code open; data synthetic and generated by the script",
    paste("Script checksum recorded at release (see the decision sheet)")
  ),
  status = c("filled", "filled", "filled", "filled", "open", "filled", "open", "filled", "fi
             "filled", "open", "filled", "filled", "filled"),
  closes_gap = c("", "", "", "",
                 "Protocol of the reference method (e.g. histology), including blinding to t
                 "",
                 "Result of a spectral classifier on the same animal-wise split, compared wi
                 "", "", "",
                 "Repeat with changed illumination reference and report the change in the re
                 "", "", "")
)
knitr::kable(passport[, c("field", "entry", "status", "closes_gap")],
             col.names = c("Field", "Entry", "Status", "Information that would close the gap
# Every open field names the information that would close it
stopifnot(all(passport$closes_gap[passport$status == "open"] != ""),
```

```
          sum(passport$status == "open") == 3,
          !any(passport$status == "filled" & passport$entry == ""))
```

Table 27.2.: Filled model passport for the teaching case. Unit: field; data status: synthetic teaching case of Chapter 19.

Field	Entry	Status	Information that would close the gap
Target statement	Assign class A or B to spectra from animals not seen in training (same fictitious species, device and protocol)	filled	
Intended use	Research only; no clinical or veterinary decision	filled	
Unit of analysis and independent n	Animal; 4 animals with 3 spectra each; test set of 2 new animals	filled	
Data	Synthetic teaching data (status: synthetic); class is constant within an animal	filled	
Reference standard		open	Protocol of the reference method (e.g. histology), including blinding to the spectra
Baseline	Reference rule: class of a known animal, otherwise pre-specified default class A	filled	

Table 27.2.: Filled model passport for the teaching case. Unit: field; data status: synthetic teaching case of Chapter 19.

Field	Entry	Status	Information that would close the gap
Candidate model		open	Result of a spectral classifier on the same animal-wise split, compared with the reference rule
Validation	Split by animal; ID overlap between training and test: 0	filled	
Result	3/6 spectra and 1/2 animals correct	filled	
Uncertainty	Per animal 1.3% to 98.7% (exact 95%); uninformative with 2 test animals	filled	
Counter-check		open	Repeat with changed illumination reference and report the change in the result (see Chapter 10)
Limits	One species, one device; no human data; no statement about patients	filled	
Code and data access	Code open; data synthetic and generated by the script	filled	
Artefact	Script checksum recorded at release (see the decision sheet)	filled	

27.3.2. The decision sheet

The decision sheet sits next to the passport, not inside it. For the teaching case it reads as follows.

Table 27.3.: Decision sheet for the teaching case, kept separate from the passport. Unit: decision; data status: fictitious project.

Dimension	Question	Entry for the teaching case	Responsible role
Ethics	Who could be harmed by an overstated claim?	Future animals in a follow-up study planned on a false premise; later patients if the claim travels into a clinical study. The next experiment needs a 3Rs justification (Chapter 16).	principal investigator with the animal welfare officer
Regulation	Does the intended purpose make this a medical device or high-risk AI system?	No: research only, no diagnostic purpose. Re-assess if the purpose changes (Chapter 22).	principal investigator

Dimension	Question	Entry for the teaching case	Responsible role
Data protection	Are personal data processed?	No (animal data). Would become relevant with human measurements (Chapter 23).	–
Security	Is the evaluated artefact the one that is reported?	Checksum of the analysis script recorded; repository access limited to the project group (Chapter 24, Chapter 25).	project data manager
Responsibility	Who decides, who reviewed?	Decision: principal investigator. Counter-review: statistician of another group.	named in the project file

Dimension	Question	Entry for the teaching case	Responsible role
Decision and conditions	What is allowed now?	Abstract may report the animal-wise result (1/2 animals) and the reference rule; it must not report 4/4 as accuracy. Next step: spectral classifier on new animals with a sample size justified in advance.	principal investigator
Review	When is this sheet looked at again?	Before any submission, and after the next experiment; date set by the principal investigator.	principal investigator

27.3.3. When an alert goes wrong: who acts next

In a clinical pilot, the decision sheet has to answer one question that matters most on a bad day: after a faulty or harmful alert, who acts next? A workable order, to be adapted locally, is the following.

1. The treating clinician decides on the patient's care now, indepen-

dently of the tool, and documents the alert and what was done.
2. The clinician reports the event to the pilot lead named in the decision sheet (row "Responsibility"). In a ward pilot this is usually the clinical lead of the unit, not the developer. The pilot lead decides whether the tool stays on, runs under stricter conditions or is switched off, using the stop criterion of the monitoring plan (Chapter 20).
3. The pilot lead informs the roles the event concerns: the developer or manufacturer for analysis, the person responsible for medical device questions if the tool may be a device (Chapter 22), the data protection officer if personal data went where they should not (Chapter 23), and information security if manipulation is suspected (Chapter 25).
4. The patient is told what happened in words they can follow (Chapter 21).

This order is an *assumption* of the book, not a legal prescription. Who holds each role is decided locally and written into the decision sheet before the pilot starts; an empty "Responsibility" row is a reason not to start.

27.3.4. A method-specific short report

The study is an animal experiment that evaluates a classification rule, so the short report follows the ARRIVE Essential 10 for the animal side and borrows the TRIPOD+AI items on data splitting and performance for the model side. The text below is generated from the analysis objects.

> **Methods (short report).** Design: synthetic teaching experiment with 4 animals, 3 spectra per animal and a class label that is constant within an animal (ARRIVE items 1, 6, 8). Sample size: fixed by design; no calculation, because the experiment is didactic (item 2; a real study would justify it). Exclusions: none (item 3). Randomisation and blinding: not applicable to synthetic data (items 4–5); a real study reports both. Analysis:

a reference rule that returns the known class of an animal and the pre-specified default class A for new animals; evaluation by an animal-wise split with 2 training and 2 test animals and an empty ID overlap; exact 95% intervals (item 7; TRIPOD+AI items on data splitting and performance measures).

Results (short report). In new animals the rule classified 3 of 6 spectra and 1 of 2 animals correctly (per animal 50%, 95% interval 1.3% to 98.7%). A split by spectrum, in which every test animal had been seen in training, gave 4 of 4 spectra correct; this measures recognition of known animals and is not reported as accuracy (item 10).

27.3.5. Counter-review

The second group reads the passport and the code. Their three findings, and the answers, are part of the record.

Table 27.4.: Counter-review of the teaching passport. Unit: finding.

Finding	Answer	Status
The abstract reports 4/4, which comes from plan 1 (known animals).	Abstract changed to the animal-wise result; plan 1 is reported as a recognition check only.	closed
The per-spectrum interval for plan 2 treats six spectra as independent.	Interval reported per animal; per-spectrum result kept as description.	closed
No candidate model has been evaluated; the reference rule is trivial.	Correct; recorded as an open field with the information that would close it.	open, with plan

27.3.6. The blank passport and decision sheet

Copy the two tables below into the project file. Leave no field empty: write *open* and the information that would close it, or *not applicable* with a reason.

Table 27.5.: Blank model passport.

Field	Entry	Status (filled / open / not applicable)	Information that would close the gap
Target statement (who or what, which setting)			
Intended use and users			
Unit of analysis and independent n			
Data (source, period, status, access tier)			
Reference standard			
Baseline			
Candidate model			
Validation (split, time, site; ID overlap)			
Result with denominator			
Uncertainty			

Field	Entry	Status (filled / open / not applicable)	Information that would close the gap
Counter-check Limits			
Code and data access			
Artefact (version, checksum)			

Table 27.6.: Blank decision sheet.

Dimension	Question	Entry	Responsible role
Ethics	Who bears the harm of a wrong decision, and who can object?		
Regulation	Which intended purpose, which role, which legal text, which open classification?		
Data protection	Which data, which legal basis, which recipients, when deleted?		

Dimension	Question	Entry	Responsible role
Security	Which threat, which control, which test, who reacts?		
Responsibility	Who decides, who reviewed, who can stop?		
Decision and conditions	What is allowed now, under which conditions, with which stop criterion?		
Review	When and on which trigger is the sheet reviewed?		

27.4. Counter-check

A passport can be checked mechanically for the mismatches that matter most. The function below compares the target statement with the split, the ID overlap, the unit of the headline result and the population the data come from. Before running it, predict the number of findings for three passports: the teaching passport, the original abstract (4/4 from plan 1), and the teaching passport with its target moved from new animals to patients.

```
review_passport <- function(p) {
  findings <- character(0)
```

```
  if (p$target_unit == "new animals" && p$split != "by animal")
    findings <- c(findings, "split does not create new animals")
  if (length(p$id_overlap) > 0)
    findings <- c(findings, "IDs overlap between training and test")
  if (p$target_unit %in% c("new animals", "patients") && p$headline_unit != '
    findings <- c(findings, "headline counts spectra, target counts individua
  if (is.null(p$interval))
    findings <- c(findings, "no uncertainty interval")
  if (p$target_population != p$data_population)
    findings <- c(findings, "no data from the target population")
  findings
}

teaching <- list(target_unit = "new animals", split = "by animal", id_overla
                 headline_unit = "animal", interval = exact_ci(1, 2),
                 target_population = "animals", data_population = "animals")
abstract_claim <- list(target_unit = "new animals", split = "by spectrum", i
                       headline_unit = "spectrum", interval = NULL,
                       target_population = "animals", data_population = "ani
moved_target <- modifyList(teaching, list(target_unit = "patients", target_p

reviews <- lapply(list(teaching = teaching, abstract = abstract_claim, patie
reviews
```

```
#> $teaching
#> character(0)
#>
#> $abstract
#> [1] "split does not create new animals"
#> [2] "IDs overlap between training and test"
#> [3] "headline counts spectra, target counts individuals"
#> [4] "no uncertainty interval"
#>
```

```
#> $patients
#> [1] "no data from the target population"
```

```
stopifnot(length(reviews$teaching) == 0, length(reviews$abstract) == 4,
          identical(reviews$patients, "no data from the target population"))
```

The teaching passport passes; the abstract fails on four counts that all come from the same decision, the split by spectrum. Moving the target to patients produces a single finding, but it is the one no amount of re-analysis can fix: the data contain no patients. Closing it requires a new study in the target population (Chapter 19), with its own passport and its own decision sheet. Note what the function does not check: whether the reference standard is sound, whether two test animals are enough, whether the classes are clinically meaningful. Those remain questions for the human counter-review.

27.5. Consequences for patients and operations

Patients and animals. In the teaching case no patient is affected yet, but the abstract is the point at which a claim starts to travel. A "100% accuracy" that later justifies a first-in-human study would expose patients to a test whose value was never shown in new individuals. On the animal side, an overstated result can trigger confirmatory experiments that a careful passport would have redirected. A clear negative passport is therefore a contribution to the 3Rs (Chapter 16).

Regulation. The passport and decision sheet are not the technical documentation required for a medical device or a high-risk AI system; they are the research record from which such documentation could later be written. The regulatory question changes with the intended purpose (Chapter 22). The decision sheet records the current purpose and names the event – a planned clinical use – that would trigger a new assessment.

Data protection and security. Here both are minor because the data are synthetic and from animals. With human data, the decision sheet refers to the legal basis, the access tier of every product (Chapter 26) and the checksum of the evaluated artefact, so that the reported model is demonstrably the tested one (Chapter 25).

Operations and responsibility. A decision sheet without names is a wish list. Each line needs a role that exists in the organisation, and the review date needs a trigger as well as a date. For a model in a running pilot, the decision sheet is the place where the stop criterion and the person who may apply it are written down (Chapter 20).

27.6. Look up and transfer

Reviewer question: Which entry of the passport would change if the target statement were "new patients" instead of "new animals of the same species", and which evidence would then be required?

Self-test: A team reports a fracture alert with an assumed sensitivity and specificity of 90% each, evaluated at a referral centre where 20 of 100 examined patients had a fracture. They want to write in the passport: "useful in our emergency department", where 5 of 100 have a fracture (the scenario of Chapter 9). Which passport field is wrong, which reporting path fits the evidence they have, and what would the decision sheet say?

Answer

The target statement does not match the data: the evidence comes from a population with 20% prevalence, the claim is about one with 5%. With the same sensitivity and specificity, the positive predictive value falls from 180/260 (69.2%) to 45/140 (32.1%), so roughly two of three alerts in the emergency department would concern patients without a fracture. The evidence is diagnostic accuracy, so the

reporting path is STARD-AI; a claim about usefulness in the department needs an early clinical evaluation reported with DECIDE-AI. The decision sheet would allow at most a supervised pilot without influence on treatment (for example in shadow mode, Chapter 20), name the clinical lead who can stop it, and set a review after a fixed number of alerts.

Annotated sources

- Sounderajah et al. (2025) – reporting guideline for diagnostic accuracy studies using AI (an erratum appeared in September 2026); the path for diagnostic claims.
- Collins et al. (2024) – reporting guideline for prediction models using regression or machine learning; also the source of the model-side items in the animal short report.
- Moons et al. (2025) – appraisal tool for quality, risk of bias and applicability of prediction model studies.
- Vasey et al. (2022) – reporting guideline for the early clinical evaluation of AI decision support.
- Percie du Sert et al. (2020) – ARRIVE 2.0 with the Essential 10, the minimum for any report of animal research.
- Monks et al. (2019) and Grimm et al. (2020) – reporting of discrete-event, agent-based and system dynamics simulations, and the ODD protocol for agent-based models.
- Mitchell et al. (2019) and Gebru et al. (2021) – model cards and datasheets, the documentation formats the passport builds on.
- Kapoor and Narayanan (2023) – survey of leakage in machine-learning-based science and the case for structured model information sheets.

27.7. Further reading

- Chapter 19 – the split plan, ID checks and external tests behind the teaching passport.
- Chapter 7 – how to read a published AI study with the same guidelines.
- Chapter 16 – sample size, small denominators and the 3Rs for the next animal experiment.
- Methods in R: Reporting templates, in particular TRIPOD-AI and ARRIVE – fill-in templates for the two guidelines of the short report.
- Methods in R: Writing a report – Reporting guidelines – where guidelines fit into a manuscript.
- CTTIR Courses: Writing a report – IMRaD, effect sizes with intervals and reproducibility in one page.
- CTTIR Courses: Pre-registration and statistical analysis plans – fixing the target statement and the analysis before the next experiment.
- CTTIR Courses: TRIPOD-AI, fairness auditing, reproducibility at scale – a lab that walks through the reporting items of a prediction model.

28. Case atlas, glossary and transfer

R. Heller · Date: 2026-07-31 · Last changed: 2026-11-09

At a glance

Guiding question: From which decision do I find the right method again?
For: everyone who starts a new project with this book, course participants at the start and the end of the course, supervisors who assign a project · **Prerequisites:** none; the chapter is an entry point and points to the chapters that hold the prerequisites
In the course: W1 · From clinical case to question · W8 · Colloquium and decision (required) · **Data:** no patient data; the atlas, navigator and glossary are stored as tables in the chapter · **Core reading time:** about 9 minutes

28.1. Case and decision

A newly appointed research coordinator in a fictitious trauma department receives three requests in her first week. The rehabilitation team wants to check gait after ankle fractures with tablet videos. A laboratory group asks whether it may publish hundreds of spectra from six animals as “a large dataset”. The operating theatre manager wants to know whether a second weekend team is worth the cost. She has two days to send each group a first answer: which question they are really asking, which chapter they should read first, and which risk they must rule out before anyone

writes code. If she routes the requests by data type alone, the gait group starts with image processing and the laboratory with model choice; both would miss the problem that decides their project.

28.2. Core reading

Core reading · required for W8 (and a first orientation in W1) · about 9 minutes · ends before ‘Lab’

28.2.1. Start from the decision, not from the data

This chapter is the index of the book, organised by decisions. Most projects arrive with a data type – images, spectra, a registry – and a method someone has heard of. The data type is a poor guide. A tablet video can serve a method comparison (Chapter 11), a privacy assessment (Chapter 23) or a safety pilot (Chapter 20), and each needs a different first step. **The key assumption of this chapter is that a method is found again through the decision it supports and the risk that threatens it, not through the data it uses.** The data type narrows the choice; the decision and the risk make it.

The chapter offers three tools for this. The *navigator* gives three entry routes to the 28 chapters: by goal, by data type and by risk. The *case atlas* contains six new project vignettes, from a fracture alert to operating capacity. Each one shows how a request turns into a target statement, a first test and a path through the chapters. The *glossary* translates the book’s terms between English and German, so that a reader can recognise a term from clinical practice and find the chapter where it is introduced. A register of common mistakes and a one-page sketch for the next experiment close the chapter. Its tables are stored as data in the chapter source, and the Lab checks that every route points to an existing chapter.

28.2.2. The route by risk

The risk route is the fastest when something has already gone wrong or is about to (Table 28.1). Each row names a risk, the chapters that address it and the workshops in which it is practised. The first chapter in each row is the one to open first.

Table 28.1.: Navigator by risk. Each row leads from a risk to the chapters that address it, in the order in which they are best read, and to the workshops of the course. Unit: risk; no data.

Risk	Chapters (read the first one first)	Workshops
A faulty or harmful alarm, too many false alarms	Chapter 20, Chapter 27, Chapter 21, Chapter 9, Chapter 6	W3, W7, W8
Leakage between training and test in grouped data	Chapter 19, Chapter 5, Chapter 27	W4
Wrong n: repeated measurements counted as independent	Chapter 3, Chapter 16, Chapter 10	W1, W4
Poor calibration in a new setting	Chapter 6, Chapter 19	W3
Selection or collider bias	Chapter 2	W5
Clusters that are artefacts of scaling or segmentation	Chapter 14, Chapter 12	W5
Unit errors and unidentifiable parameters	Chapter 15	W6
Distribution shift after an update or a season	Chapter 20, Chapter 17	W6, W7
Errors that fall unequally on groups	Chapter 21, Chapter 6	W3, W8
Unclear intended purpose or changed use	Chapter 22	W7

Table 28.1.: Navigator by risk. Each row leads from a risk to the chapters that address it, in the order in which they are best read, and to the workshops of the course. Unit: risk; no data.

Risk	Chapters (read the first one first)	Workshops
Unlawful data flow or re-identification of rare groups	Chapter 23, Chapter 26	W7
Data leak, wrong model version or outage	Chapter 24, Chapter 25, Chapter 20	W7
Instructions hidden in documents	Chapter 25	W2, W7
Overstated claim in an abstract	Chapter 27, Chapter 7	W8

28.2.3. Two reading routes as examples

A clinician after a faulty alarm. An alert flagged a fracture that was not there, and the patient received an unnecessary CT scan. The clinician wants to know who must be told and who decides whether the alert stays on. The risk route “a faulty or harmful alarm” leads first to Chapter 20, where the monitoring plan names the escalation path, the fallback procedure and the person who may stop the pilot. The decision sheet in Chapter 27 records that person, and Section 27.3.3 lists who acts next, in order; Chapter 21 helps to explain the event to the patient. The next responsible role is therefore the pilot lead named in the decision sheet as allowed to stop – not the developer and not the clinician who happened to see the alert.

A doctoral researcher with grouped spectra. Hundreds of spectra come from a handful of animals. The data route “hyperspectral images and grouped spectra” leads to Chapter 10 for the measurement, to Chapter 19 for a split by animal with an ID overlap of zero, and to Chapter 27 for the passport and the reporting path (Table 27.1): ARRIVE 2.0 for the animal study,

TRIPOD+AI items for the model. The fill-in templates in Methods in R: ARRIVE and Methods in R: TRIPOD-AI complete the report.

28.2.4. Combining routes: the coordinator's three requests

The routes work best in combination. For the gait request, the data route ("video, landmarks and joint angles") offers Chapter 11 and Chapter 23, and the risk route ("unlawful data flow or re-identification") offers Chapter 23 and Chapter 26. The chapter that appears in both, Chapter 23, is where the coordinator starts: before anyone compares measurement methods, the team must know whether raw video may be stored at all and when it is deleted. For the laboratory, the data route leads to the split by animal in Chapter 19; "a large dataset" of spectra from six animals is, for any statement about new animals, a dataset of six. For the operating theatre, two goal rows compete: "forecast demand" (Chapter 17) and "compare capacity scenarios" (Chapter 18). The decision is about staffing, so the simulation is the main tool and the forecast supplies its arrival assumptions. In each case the first chapter is chosen by the decision and the risk; the data type only narrows the list.

28.2.5. Working with the atlas

Each vignette in the atlas follows the same form: the decision, the data and the independent unit, a first baseline, the main risk, a first observable test, the path through the chapters and the matching workshops. Read the vignette nearest to your own project, then write your own version in the same fields before you open any chapter. The formative model answers are kept separate and folded away, so that the attempt comes first. They are not marking schemes; they show what a complete first answer contains.

Question for the workshop: Choose the vignette nearest to your project (W1) or to the project you defended (W8). Which field of your

own brief differs most from the vignette, and which chapter does that difference send you to?

28.3. Lab

The navigator, the atlas, the glossary and the register of mistakes are data frames in the source of this chapter. The Lab checks them in the same way that a passport is checked: every route must lead to an existing chapter, every atlas case must point to at least two chapters and one workshop, and the glossary must cover the book.

```
# Size of the tables
c(navigator_rows = nrow(navigator), atlas_cases = length(atlas),
  glossary_terms = nrow(glossary), mistakes = nrow(mistakes))
```

```
#> navigator_rows    atlas_cases glossary_terms       mistakes
#>             50              6             95             18
```

```
nav_ids  <- trimws(unlist(strsplit(navigator$chapters, ",")))
mis_ids  <- trimws(unlist(strsplit(mistakes$chapters, ",")))
atlas_ids <- unlist(lapply(atlas, `[[`, "path"))

stopifnot(
  length(atlas) == 6,
  all(sapply(atlas, function(a) length(a$path) >= 2 && length(a$ws) >= 1)),
  all(unlist(lapply(atlas, `[[`, "ws")) %in% workshops$id),
  all(c(nav_ids, mis_ids, atlas_ids, glossary$chapter) %in% chapter_ids),
  nrow(glossary) >= 60, !anyDuplicated(tolower(glossary$term)),
  all(chapter_ids %in% nav_ids)
)
```

All 28 chapters are reachable from the navigator, and no route points to a chapter that does not exist. The two reading routes of the Core reading can be tested as searches. The function returns the chapters of every navigator row whose entry contains the search term, in order of first appearance.

```
find_route <- function(term) {
  rows <- navigator[grepl(term, navigator$entry, ignore.case = TRUE), ]
  unique(trimws(unlist(strsplit(rows$chapters, ","))))
}
route_alarm   <- find_route("alarm")
route_spectra <- find_route("spectra")
route_alarm
```

```
#> [1] "b20" "b27" "b21" "b09" "b06"
```

```
route_spectra
```

```
#> [1] "b10" "b19" "b27"
```

```
stopifnot(route_alarm[1] == "b20", "b27" %in% route_alarm,
          all(c("b10", "b19", "b27") %in% route_spectra))
```

The atlas map shows which chapters the six vignettes use and in which order (Figure 28.1). Chapters that no vignette touches are not unimportant; they are reached through the navigator.

```
library(ggplot2)

map_df <- do.call(rbind, lapply(names(atlas), function(v)
  data.frame(vignette = v, chapter = atlas[[v]]$path, step = seq_along(atlas[[v]]$path))))
```

```
map_df$chapter <- factor(toupper(map_df$chapter), levels = toupper(chapter_ic
map_df$vignette <- factor(map_df$vignette, levels = rev(names(atlas)))
map_df$first <- ifelse(map_df$step == 1, "first chapter", "later chapter")

ggplot(map_df, aes(chapter, vignette)) +
  geom_tile(aes(fill = first), colour = "white", linewidth = 0.8) +
  geom_text(aes(label = step), size = 3.2) +
  scale_x_discrete(drop = FALSE) +
  scale_fill_manual(values = c("first chapter" = unname(cdsr_colours["orange"
                               "later chapter" = unname(cdsr_colours["skyblue
  labs(x = NULL, y = "Vignette", caption = "Paths as listed in the case atlas
  theme_cdsr(base_size = 11) +
  theme(axis.text.x = element_text(angle = 90, vjust = 0.5, hjust = 1), panel

stopifnot(nrow(map_df) == length(atlas_ids))
```



Figure 28.1.: Case atlas map: the six vignettes (rows) and the chapters on their reading path (columns). The number in a tile is the position of the chapter on the path. Unit: vignette; n = 6 vignettes, 28 chapters; no data.

28.4. Counter-check

The path of a vignette is not fixed by its data. It changes when the target statement changes. The function below adds chapters to a path when the target population moves from animals to patients, and when a research result is meant to be used in care. Before running it, predict which chapters vignette A2 (wound perfusion in a pig model) gains if its target statement becomes “patients with soft-tissue wounds in a supervised clinical pilot”.

```
extend_path <- function(v, population, deployment) {
  a <- atlas[[v]]
  extra <- character(0)
  if (a$population == "animals" && population == "patients")
    extra <- c(extra, "b19", "b21", "b23")            # species transfer, ethics, personal d
  if (deployment && !a$deployment)
    extra <- c(extra, "b22", "b20", "b24", "b25")     # purpose, pilot, security
  unique(c(a$path, extra))
}
a2_moved <- extend_path("A2", population = "patients", deployment = TRUE)
setdiff(a2_moved, atlas$A2$path)
```

```
#> [1] "b21" "b23" "b22" "b20" "b24" "b25"
```

```
stopifnot(all(c("b21", "b22", "b23", "b20") %in% a2_moved),
          length(setdiff(a2_moved, atlas$A2$path)) == 6)
```

The pig study gains six chapters: ethics with patients, regulation, data protection, the pilot and both security chapters. Species transfer (Chapter 19) was already on the path, now with a harder question – not new animals, but a new species. The data type did not change at all. This is the key assumption of the chapter in one line of code.

28.5. Consequences for patients and operations

Patients. A wrong route costs time, but a route that skips a risk costs more. The most frequent failure is to treat a project as purely technical: the gait group that never asks where the raw video goes, the laboratory that never asks what a result per spectrum means for the next animals. The risk route exists to make these questions visible before anyone is affected.

Responsibility. The atlas does not replace the people who carry responsibility: the statistician for the analysis plan, the data protection officer for the data flow, the ethics committee and the animal welfare body for the study, the regulatory function for the intended purpose, and the clinical lead for the decision to use a tool. Every vignette therefore ends with a decision that a named role must take (Chapter 27).

Regulation, data protection and security. These topics enter a path through the risk route and through the counter-check above. A vignette that stays in research has fewer of them; the moment a result is meant to change care, Chapter 22 to Chapter 25 join the path.

Keeping the atlas current. Chapter numbers are stable identifiers, but legal dates, guidelines and sister-book anchors change. The edition notes record the state of each edition; the Lab of this chapter is the first test to rerun after any change to the chapter list.

28.6. Look up and transfer

Reviewer question: For a project you are asked to review, which row of the risk navigator would you open first, and does the project's own documentation answer the question of that row?

Self-test: A pharmacy team asks whether a forecast of antibiotic consumption on the trauma wards can be used to order stock two weeks

ahead, and whether a dashboard could later suggest switching patients to oral treatment. Which vignette is nearest, which part of the request changes the path, and which chapters join it?

Answer

The nearest vignette is A6: a forecast of demand with a planning decision (Chapter 17), checked by a rolling-origin backtest against a seasonal naive baseline. The first request stays in that path. The second request is different in kind: suggesting a treatment change for individual patients makes the dashboard a decision aid about patients. That part of the request brings in Chapter 22 (intended purpose and possible medical device status), Chapter 23 (patient-level data flow), Chapter 21 (who bears the harm of a wrong suggestion) and, before any use, Chapter 20 (pilot and stop criterion). The two requests should become two projects with two passports.

28.6.1. Case atlas

The six vignettes are new cases, not repeats of the chapter examples. All institutions, numbers and data are fictitious; where a number would be needed, it is marked *open*.

28.6.1.1. A1 · A second look at the wrist

A hand-surgery outpatient clinic considers an alert, evaluated in a large emergency department, that flags possible scaphoid fractures on wrist radiographs. Many of its patients were referred after negative imaging elsewhere.

- **Decision:** should an alert trigger a second read by a hand surgeon, and who follows up a missed fracture?

- **Data and unit:** one radiograph examination per patient and episode; reference standard: follow-up imaging.
- **First baseline:** current practice, a second read only on clinical suspicion.
- **Main risk:** prevalence and case mix differ from the evaluation site, so predictive values do not transfer; alert fatigue.
- **First observable test:** recompute the 2×2 table at the clinic's prevalence (*open*: local prevalence), then a pilot in shadow mode with a stop criterion.
- **Chapter path:** Chapter 9 → Chapter 6 → Chapter 19 → Chapter 20 → Chapter 21 → Chapter 27
- **Workshops:** W3 · Diagnostics, prediction and evidence · W8 · Colloquium and decision
- **Look further:** CTTIR Courses: Diagnostic testing – Se, Sp, PPV, NPV, LR

28.6.1.2. A2 · Perfusion in soft-tissue wounds

A translational group records hyperspectral images of experimental soft-tissue wounds in a fictitious pig model: several images per wound, two wounds per animal, three time points.

- **Decision:** is a confirmatory experiment justified, and with how many animals?
- **Data and unit:** spectra nested in images, wounds and animals; the independent unit is the animal.
- **First baseline:** a simple spectral index with a threshold fixed in advance.
- **Main risk:** leakage between images of the same animal; illumination differences between sessions.
- **First observable test:** split by animal with an ID overlap of zero; counter-check with a changed illumination reference.

- **Chapter path:** Chapter 10 → Chapter 19 → Chapter 16 → Chapter 27
- **Workshops:** W4 · Imaging, gait and tissue · W8 · Colloquium and decision
- **Look further:** Methods in R: ARRIVE – animal research

28.6.1.3. A3 · Walking again after an ankle fracture

A rehabilitation team films patients walking along the corridor with a tablet; software estimates the knee angle.

- **Decision:** can the tablet angle replace the laboratory measurement for the discharge assessment?
- **Data and unit:** many steps per patient, two methods; the independent unit is the patient.
- **First baseline:** the laboratory measurement as reference; a tolerance fixed clinically in advance.
- **Main risk:** camera angle and occlusion; raw video that identifies patients and is kept longer than needed.
- **First observable test:** difference plot per patient against the tolerance; a data flow with a deletion point for the raw video.
- **Chapter path:** Chapter 11 → Chapter 3 → Chapter 23 → Chapter 26
- **Workshops:** W4 · Imaging, gait and tissue · W5 · Causality, trajectories and phenotypes
- **Look further:** CTTIR Courses: Kappa, ICC, Bland–Altman

28.6.1.4. A4 · A serum marker after spinal cord injury

A fictitious cohort has a serum marker measured at irregular times over twelve weeks after spinal cord injury. Patients in the worst condition are transferred or die earlier and contribute fewer measurements.

- **Decision:** should the early slope of the marker enter the protocol of a prognostic study?
- **Data and unit:** repeated measurements per patient; time zero is the injury.
- **First baseline:** a per-patient summary (change from the first measurement) next to a mixed model.
- **Main risk:** informative dropout; selection through intensive care admission.
- **First observable test:** sensitivity analysis for late measurements; a DAG for the selection step.
- **Chapter path:** Chapter 2 → Chapter 3 → Chapter 13 → Chapter 16
- **Workshops:** W5 · Causality, trajectories and phenotypes
- **Look further:** CTTIR Courses: Linear mixed models with lme4

28.6.1.5. A5 · Cell types in a fracture callus

Multiplex-stained sections from fictitious fracture callus biopsies are segmented, and the resulting clusters are labelled as new cell subtypes for a grant proposal.

- **Decision:** may the subtype label appear in the proposal as a finding?
- **Data and unit:** cells within sections within patients; the independent unit is the patient.
- **First baseline:** marker thresholds set by an expert (manual gating).
- **Main risk:** segmentation boundary, batch effects between staining runs, clusters produced by scaling; rare marker combinations in shared images.
- **First observable test:** shift the segmentation boundary and repeat the clustering on bootstrap samples; summarise per patient.
- **Chapter path:** Chapter 12 → Chapter 14 → Chapter 26 → Chapter 27

- **Workshops:** W4 · Imaging, gait and tissue · W5 · Causality, trajectories and phenotypes
- **Look further:** CTTIR Courses: Clustering – k-means, hierarchical, model-based

28.6.1.6. A6 · Weekend peaks in the trauma theatre

A fictitious trauma service sees more emergency operations at weekends. Management asks whether a second weekend team would shorten the wait for hip fracture surgery.

- **Decision:** staff a second weekend team or not.
- **Data and unit:** daily arrivals for the forecast; waiting time per patient for the simulation.
- **First baseline:** a seasonal naive forecast; the current one-team scenario.
- **Main risk:** using a forecast to answer a capacity question; longer waits shifted onto lower-priority patients.
- **First observable test:** rolling-origin backtest of arrivals; both scenarios simulated on identical arrival streams, with the spread across runs.
- **Chapter path:** Chapter 17 → Chapter 18 → Chapter 20 → Chapter 21
- **Workshops:** W6 · Pharmacometrics, forecasting and simulation · W8 · Colloquium and decision
- **Look further:** Methods in R: Rolling-origin cross-validation

28.6.2. Formative model answers

Try your own brief first. Each answer names the target statement, the first test with its expected result and the decision that a named role has to take.

Model answer A1

Target statement: “In patients referred to this clinic after a wrist injury, an alert-triggered second read finds scaphoid fractures that would otherwise be missed, at an acceptable number of extra reads.” First test: the 2×2 table at the local prevalence, with the number of alerts per week as the workload measure. Expected result: a lower positive predictive value than at the evaluation site if the local prevalence is lower – this alone is not a reason to reject the tool. Decision: the clinic lead approves a shadow pilot and names who may stop it (Chapter 20, Chapter 27).

Model answer A2

Target statement: “In new pigs of the same strain, the index separates well from poorly perfused wound tissue.” First test: split by animal, ID overlap zero, result reported per animal with its interval. Expected result: a wide interval with few animals; the passport then states which sample size the next experiment needs, justified with Chapter 16. Decision: the principal investigator, with the animal welfare officer, decides whether the confirmatory experiment is justified under the 3Rs.

Model answer A3

Target statement: “For patients after ankle fracture, the tablet angle agrees with the laboratory angle within the clinical tolerance.” First test: limits of agreement from per-patient means, compared with the tolerance fixed in advance. Expected result: either agreement within tolerance or a systematic offset that depends on camera position. Decision: the rehabilitation lead decides on use; the data protection officer approves the data flow and the deletion point for raw video

(Chapter 23).

Model answer A4

Target statement: “In patients after spinal cord injury, the marker slope over the first weeks is associated with the later outcome.” First test: mixed model against per-patient summaries, plus an analysis that assumes worse values for patients who drop out. Expected result: if the slope changes materially under this assumption, the protocol must plan for dropout. Decision: the study steering group decides whether the marker enters the protocol as a candidate, not as a proven predictor.

Model answer A5

Target statement: “Across patients, a cell population with this marker profile is reproducibly present in fracture callus.” First test: clustering repeated after shifting the segmentation boundary and on bootstrap samples, summarised per patient. Expected result: a stable population is found again; an unstable one is reported as a hypothesis. Decision: the principal investigator words the proposal accordingly; shared images are checked for rare combinations before release (Chapter 26).

Model answer A6

Target statement: “Under the stated arrival assumptions, a second weekend team reduces the waiting time for hip fracture surgery by an amount that justifies its cost.” First test: forecast checked by backtest; both scenarios simulated on identical arrival streams with the spread over runs. Expected result: a difference with a range, plus the effect on lower-priority patients. Decision: the hospital manage-

ment decides on staffing; the simulation report names its assumptions (STRESS, Chapter 27).

28.6.3. Navigator by goal and by data type

Table 28.2.: Navigator by goal. Unit: goal; no data.

Goal	Chapters	Workshops
State the decision, the addressee and the target population	Chapter 1, Chapter 2	W1
Describe how the data arose and which denominator applies	Chapter 3, Chapter 4	W1
Make an analysis reproducible from a fresh R session	Chapter 4	W1, W6
Build a baseline and compare a learning model with it	Chapter 5, Chapter 6	W3
Judge a published AI study	Chapter 7, Chapter 27	W2, W3
Check a generated text or a citation	Chapter 8	W2
Interpret a diagnostic alert in another setting	Chapter 9, Chapter 6, Chapter 21	W3
Compare two measurement methods	Chapter 11	W4
Ask a causal question without selection bias	Chapter 2	W5
Model repeated measurements over time	Chapter 13	W5
Decide whether clusters are subtypes	Chapter 14	W5
Fit a mechanistic or pharmacokinetic model	Chapter 15	W6
Plan the next animal experiment with few animals	Chapter 16, Chapter 27	W5, W6

Table 28.2.: Navigator by goal. Unit: goal; no data.

Goal	Chapters	Workshops
Forecast demand for the next weeks	Chapter 17	W6
Compare capacity scenarios	Chapter 18	W6
Transfer a result to a new centre, period or species	Chapter 19	W4, W6
Run a pilot, monitor it and decide when to stop	Chapter 20, Chapter 27	W7, W8
Weigh harms and explain a choice to a patient	Chapter 21	W2, W8
Clarify the regulatory status of software	Chapter 22	W7
Plan a lawful data flow	Chapter 23	W7
Protect systems and restore operation	Chapter 24, Chapter 25	W7
Share code, metadata and results	Chapter 26	W7
Document evidence and decision for a review	Chapter 27	W8
Find a method or a term again	Chapter 28	W1, W8

Table 28.3.: Navigator by data type. Unit: data type; no data.

Data type	Chapters	Workshops
Clinical table, one row per patient	Chapter 3, Chapter 5, Chapter 6	W1, W3
Diagnostic images with an alert	Chapter 9, Chapter 7	W3
Hyperspectral images and grouped spectra	Chapter 10, Chapter 19, Chapter 27	W4
Video, landmarks and joint angles	Chapter 11, Chapter 23	W4
Tissue sections and cell coordinates	Chapter 12, Chapter 14	W4

Table 28.3.: Navigator by data type. Unit: data type; no data.

Data type	Chapters	Workshops
Repeated biomarker measurements	Chapter 13, Chapter 3	W5
Concentration over time	Chapter 15	W6
Event counts in a small animal study	Chapter 16	W5, W6
Time series of arrivals	Chapter 17	W6
Process events and queues	Chapter 18	W6
Free text, generated text, documents	Chapter 8, Chapter 25	W2
Registry extract for secondary use	Chapter 26, Chapter 23	W7

When the question is a choice between statistical tests rather than a project decision, the decision trees of the sister resources are faster: Methods in R: Find your method, CTTIR Courses: Decision tree and the CTTIR Tutorials: Decision Tree with its interactive assistant.

28.6.4. Glossary (English–German)

The glossary lists 95 terms with their German equivalent and the chapter in which each is introduced. Statistical terms that this book uses without defining them are explained in the glossaries of the sister books: Methods in R, Scientometrics in R, Strategy in R and CTTIR Courses.

Table 28.4.: Glossary, English–German, with the chapter that introduces each term. Unit: term; no data.

Term	German	Introduced in
3Rs (replacement, reduction, refinement)	3R (Vermeidung, Verminderung, Verbesserung)	Chapter 16
Agent-based simulation	agentenbasierte Simulation	Chapter 18
Alert fatigue	Alarmmüdigkeit	Chapter 20
Anonymisation	Anonymisierung	Chapter 23
Area under the ROC curve (AUC)	Fläche unter der ROC-Kurve	Chapter 6
Baseline model	Vergleichsmodell, Basismodell	Chapter 5
Batch effect	Batch-Effekt, Chargeneffekt	Chapter 12
Bland–Altman plot	Differenzdiagramm nach Bland und Altman	Chapter 11
Calibration	Kalibrierung	Chapter 6
Case mix	Fallmix	Chapter 19
Checksum (hash)	Prüfsumme (Hashwert)	Chapter 25
Clearance	Clearance	Chapter 15
Cluster stability	Clusterstabilität	Chapter 14
Collider	Collider (gemeinsame Folge zweier Variablen)	Chapter 2
Confounder	Confounder, Störgröße	Chapter 2
Controlled access	kontrollierter Zugang	Chapter 26
Data dictionary	Datenwörterbuch	Chapter 3
Data leakage	Datenleck zwischen Training und Test	Chapter 19
Data minimisation	Datenminimierung	Chapter 23
Data permit	Datengenehmigung	Chapter 26
Data poisoning	Datenvergiftung	Chapter 25
Data protection impact assessment (DPIA)	Datenschutz-Folgenabschätzung (DSFA)	Chapter 23

Table 28.4.: Glossary, English–German, with the chapter that introduces each term. Unit: term; no data.

Term	German	Introduced in
Decision sheet	Entscheidungsblatt	Chapter 27
Decision threshold	Entscheidungsschwelle	Chapter 6
Denominator	Nenner	Chapter 3
Directed acyclic graph (DAG)	gerichteter azyklischer Graph	Chapter 2
Discrete-event simulation	ereignisdiskrete Simulation	Chapter 18
Distribution shift, drift	Verteilungsverschiebung, Drift	Chapter 20
Equivalence class	Äquivalenzklasse	Chapter 26
Estimand	Estimand, Schätzziel	Chapter 2
Excessive agency	übermäßige Handlungsbefugnis	Chapter 25
External validation	externe Validierung	Chapter 19
Fabricated reference	erfundene Quellenangabe	Chapter 8
Fallback procedure	Rückfallverfahren	Chapter 20
False positive	falsch positiver Befund	Chapter 9
Grouped split (split by individual)	gruppierte Aufteilung (nach Individuum)	Chapter 19
Health data access body	Zugangsstelle für Gesundheitsdaten	Chapter 26
High-risk AI system	Hochrisiko-KI-System	Chapter 22
Human-in-the-loop	Mensch in der Entscheidungsschleife	Chapter 20
Hyperspectral cube	hyperspektraler Datenwürfel	Chapter 10
Identifiability (of parameters)	Identifizierbarkeit (von Parametern)	Chapter 15
Indirect prompt injection	indirekte Prompt-Injection	Chapter 25
Informative dropout	informativer Ausfall	Chapter 13

Table 28.4.: Glossary, English–German, with the chapter that introduces each term. Unit: term; no data.

Term	German	Introduced in
Intended purpose	Zweckbestimmung	Chapter 22
k-anonymity	k-Anonymität	Chapter 26
Least privilege	Prinzip der minimalen Rechte	Chapter 24
Limits of agreement	Übereinstimmungsgrenzen	Chapter 11
Long format	Langformat	Chapter 4
Membership inference	Rückschluss auf die Trainingsmitgliedschaft	Chapter 25
Missing at random (MAR)	zufällig fehlend	Chapter 3
Missing not at random (MNAR)	nicht zufällig fehlend	Chapter 3
Mixed model, random slope	gemischtes Modell, zufällige Steigung	Chapter 13
Model passport	Modellpass	Chapter 27
Neighbourhood analysis	Nachbarschaftsanalyse	Chapter 12
Nested cross-validation	verschachtelte Kreuzvalidierung	Chapter 5
Occlusion	Verdeckung	Chapter 11
Personal data breach	Verletzung des Schutzes personenbezogener Daten	Chapter 23
Positive predictive value (PPV)	positiver Vorhersagewert	Chapter 9
Posterior distribution	A-posteriori-Verteilung	Chapter 16
Prediction interval	Prognoseintervall	Chapter 17
Prevalence	Prävalenz	Chapter 9
Prior distribution	A-priori-Verteilung	Chapter 16
Pseudonymisation	Pseudonymisierung	Chapter 23
Pseudoreplication	Pseudoreplikation	Chapter 3
Purpose limitation	Zweckbindung	Chapter 26

Table 28.4.: Glossary, English–German, with the chapter that introduces each term. Unit: term; no data.

Term	German	Introduced in
Quasi-identifier	Quasi-Identifikator	Chapter 26
Re-identification	Reidentifizierung	Chapter 26
Reader study	Reader-Studie (Befunderstudie)	Chapter 7
Recovery	Wiederanlauf, Wiederherstellung	Chapter 24
Reference standard	Referenzstandard	Chapter 7
Regularisation (elastic net)	Regularisierung (Elastic Net)	Chapter 5
Reporting guideline	Berichtsleitlinie	Chapter 27
Reproducible project	reproduzierbares Projekt	Chapter 4
Retrieval-augmented generation	abrufgestützte Generierung	Chapter 8
Rolling-origin evaluation (backtest)	rollierende Prognoseprüfung	Chapter 17
Secondary use	Sekundärnutzung	Chapter 26
Secure processing environment	sichere Verarbeitungsumgebung	Chapter 26
Segmentation	Segmentierung	Chapter 12
Sensitivity	Sensitivität	Chapter 6
Shadow mode	Schattenbetrieb	Chapter 20
Specificity	Spezifität	Chapter 6
Synthetic data	synthetische Daten	Chapter 26
Target population	Zielpopulation	Chapter 1
Target statement	Zielaussage	Chapter 27
Teach-back	Teach-back-Methode (Rückversicherung durch Nacherzählen)	Chapter 21
Temperature (language model)	Temperatur (Sprachmodell)	Chapter 8

Table 28.4.: Glossary, English–German, with the chapter that introduces each term. Unit: term; no data.

Term	German	Introduced in
Threat model	Bedrohungsmodell	Chapter 24
Time zero	Zeitnull	Chapter 2
Token	Token	Chapter 8
Translational chain	translationale Kette	Chapter 1
Trust boundary	Vertrauensgrenze	Chapter 25
UMAP island	UMAP-Insel	Chapter 14
Unit of analysis	Analyseeinheit, unabhängige Einheit	Chapter 3
Volume of distribution	Verteilungsvolumen	Chapter 15
White and dark reference	Weiß- und Dunkelreferenz	Chapter 10

28.6.5. Register of common mistakes

The register lists mistakes in reasoning and design, each with a warning sign and a check. Errors in software and builds (packages, Quarto, renv, Git) are collected in Methods in R: Common errors and CTTIR Courses: Common errors. Four publications show how widespread some of these mistakes are: leakage in machine-learning-based science (Kapoor and Narayanan 2023), the wrong experimental unit in animal studies, where about half of the surveyed studies showed pseudoreplication (*published*, Lazic et al. 2018), neglected calibration (Van Calster et al. 2019), and methodological flaws that left none of 62 reviewed imaging models fit for clinical use (*published*, Roberts et al. 2021).

Table 28.5.: Register of common mistakes in reasoning and design. Unit: mistake; no data.

Mistake	Warning sign	Check	Chapters
Accuracy per image or spectrum, claim about individuals	n in the abstract exceeds the number of animals or patients	ID overlap between training and test; result per individual	Chapter 19, Chapter 27
Repeated measurements counted as independent	narrow intervals from many measurements of few individuals	name the experimental unit before the analysis	Chapter 3, Chapter 16
Predictive values carried into another setting	same PPV quoted for a different prevalence	recompute the 2 × 2 table at the local prevalence	Chapter 9
Good ranking read as good calibration	only the AUC is reported	calibration plot, observed against expected	Chapter 6
Tuning or preprocessing on all data	test performance better than plausible	nested validation; preprocessing inside the folds	Chapter 5
Adjusting for a collider	association appears after restricting to admitted patients	draw the DAG before modelling	Chapter 2
Clusters read as subtypes	clusters change with scaling, seed or segmentation	stability check and alternative scaling	Chapter 14, Chapter 12
Unit error in a mechanistic model	implausible clearance or volume	units carried through; analytic solution as check	Chapter 15
Forecast interval read as a guarantee	coverage never checked on later data	rolling-origin backtest with coverage	Chapter 17

Table 28.5.: Register of common mistakes in reasoning and design. Unit: mistake; no data.

Mistake	Warning sign	Check	Chapters
Unverified simulation	no trace of single entities, different input streams per scenario	event log of one entity; identical arrival streams	Chapter 18
Pilot without a stop criterion	no denominator, no named person, no date	monitoring plan with escalation and stop rule	Chapter 20, Chapter 27
Research label taken as legal basis	no documented legal basis for the change of purpose	legal basis per purpose; data protection officer	Chapter 23, Chapter 26
Pseudonymised data taken as anonymous	key exists; rare combinations in the release	k-anonymity check; assessment per recipient	Chapter 26
Trusting the model to ignore embedded instructions	assistant can act on document content	permission tests with synthetic documents	Chapter 25
Version label trusted instead of checksum	same version, different behaviour	compare checksum with the release record	Chapter 24, Chapter 25
Fabricated or misattributed citation	reference not found or passage does not support the claim	check DOI, title and the cited passage	Chapter 8
Use drifts beyond the intended purpose	tool used for decisions it was not meant for	compare actual use with the stated purpose	Chapter 22
Speed claim without benchmark	one language is called faster	benchmark on identical inputs and tolerance	Chapter 15

28.6.6. The next experiment on one page

Every project ends with a next step, and the next step is best written down before the colloquium. The sketch below fits on one page. The right-hand column shows it filled in for vignette A2; the same form works for a security test, for example a rerun of the control matrix of Chapter 25 after a model update.

Table 28.6.: One-page sketch of the next experiment, filled in for vignette A2. Unit: field; fictitious project.

Field	Next validation experiment for A2 (example)
Target statement	Spectral index separates well from poorly perfused wound tissue in new pigs of the same strain
What would change the decision	Agreement per animal above the threshold fixed in advance → confirmatory study; below → stop and revise the index
Design	New animals only; split by animal; same device, protocol and illumination reference as before
Independent unit and n	Animal; n justified in advance with a prior from the pilot (Chapter 16) – *open* until calculated
Baseline and comparator	The fixed index from the pilot; the reference rule of Chapter 27 as a floor
Primary measure with denominator	Correct classification per animal, with exact interval
Counter-check	Changed illumination reference in a subset of sessions

Field	Next validation experiment for A2 (example)
Pre-registration and stop criterion	Analysis plan filed before the first animal; stop if the interim check shows the index fails in both classes
Roles	Principal investigator (decision), statistician (analysis plan), animal welfare officer (3Rs)
Risks and controls	Ethics: animal numbers; data: checksum of the analysis script; no personal data
Review date	After the analysis plan is approved; date set by the principal investigator

Annotated sources

- Kapoor and Narayanan (2023) – survey of data leakage across 17 fields of machine-learning-based science, with a taxonomy of eight leakage types; background for the first row of the register.
- Lazic et al. (2018) – survey of animal experiments showing how often the experimental unit is chosen wrongly.
- Van Calster et al. (2019) – why calibration, not only discrimination, decides whether predictions can be used.
- Roberts et al. (2021) – systematic review of imaging models with a list of recurring methodological pitfalls.

28.7. Further reading

- Chapter 27 – the passport and decision sheet that turn an atlas vignette into a project record.

- Chapter 1 – the decision map with addressee, action and fallback, the starting point of every vignette.
- Methods in R: Find your method – flowchart and quick lookup from research question to statistical test.
- Methods in R: Glossary – definitions of statistical terms used throughout this book.
- CTTIR Courses: Decision tree – a branching guide from question to lab.
- CTTIR Tutorials: Decision Tree – interactive assistant and printable chart for choosing a test.
- CTTIR Courses: Glossary and Common errors – plain-English definitions and fixes for frequent R and Quarto errors.
- CTTIR Courses: Writing a report – how an analysis becomes a document a reviewer will trust.

References

R. Heller · Date: 2026-05-17 · Last changed: 2026-05-17

Altman, Douglas G., and J. Martin Bland. 1994. "Statistics Notes: Diagnostic Tests 2: Predictive Values." *BMJ* 309 (6947): 102. https://doi.org/10.1136/bmj.309.6947.102.

Ancker, Jessica S., Alison Edwards, Sarah Nosal, Diane Hauser, Elizabeth Mauer, and Rainu Kaushal. 2017. "Effects of Workload, Work Complexity, and Repeated Alerts on Alert Fatigue in a Clinical Decision Support System." *BMC Medical Informatics and Decision Making* 17 (1): 36. https://doi.org/10.1186/s12911-017-0430-8.

Baddeley, Adrian, and Rolf Turner. 2005. "Spatstat: An R Package for Analyzing Spatial Point Patterns." *Journal of Statistical Software* 12 (6): 1–42. https://doi.org/10.18637/jss.v012.i06.

Bates, Douglas, Martin Mächler, Ben Bolker, and Steve Walker. 2015. "Fitting Linear Mixed-Effects Models Using lme4." *Journal of Statistical Software* 67 (1): 1–48. https://doi.org/10.18637/jss.v067.i01.

Bellman, R., and K. J. Åström. 1970. "On Structural Identifiability." *Mathematical Biosciences* 7 (3–4): 329–39. https://doi.org/10.1016/0025-5564(70)90132-X.

Benneyan, J. C., R. C. Lloyd, and P. E. Plsek. 2003. "Statistical Process Control as a Tool for Research and Healthcare Improvement." *Quality*

and Safety in Health Care 12 (6): 458–64. https://doi.org/10.1136/qhc.12.6.458.

Bland, J. Martin, and Douglas G. Altman. 1986. "Statistical Methods for Assessing Agreement Between Two Methods of Clinical Measurement." *The Lancet* 327 (8476): 307–10. https://doi.org/10.1016/S0140-6736(86)90837-8.

Bland, J. Martin, and Douglas G. Altman. 2007. "Agreement Between Methods of Measurement with Multiple Observations Per Individual." *Journal of Biopharmaceutical Statistics* 17 (4): 571–82. https://doi.org/10.1080/10543400701329422.

Broman, Karl W., and Kara H. Woo. 2018. "Data Organization in Spreadsheets." *The American Statistician* 72 (1): 2–10. https://doi.org/10.1080/00031305.2017.1375989.

Brown, Lawrence D., T. Tony Cai, and Anirban DasGupta. 2001. "Interval Estimation for a Binomial Proportion." *Statistical Science* 16 (2): 101–33. https://doi.org/10.1214/ss/1009213286.

Bryan, Jennifer. 2018. "Excuse Me, Do You Have a Moment to Talk about Version Control?" *The American Statistician* 72 (1): 20–27. https://doi.org/10.1080/00031305.2017.1399928.

Calfee, Carolyn S., Kevin Delucchi, Polly E. Parsons, B. Taylor Thompson, Lorraine B. Ware, and Michael A. Matthay. 2014. "Subphenotypes in Acute Respiratory Distress Syndrome: Latent Class Analysis of Data from Two Randomised Controlled Trials." *The Lancet Respiratory Medicine* 2 (8): 611–20. https://doi.org/10.1016/S2213-2600(14)70097-9.

Chari, Tara, and Lior Pachter. 2023. "The Specious Art of Single-Cell Genomics." *PLOS Computational Biology* 19 (8): e1011288. https:

//doi.org/10.1371/journal.pcbi.1011288.

Chen, Tianqi, and Carlos Guestrin. 2016. "XGBoost: A Scalable Tree Boosting System." *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–94. https://doi.org/10.1145/2939672.2939785.

Christodoulou, Evangelia, Jie Ma, Gary S. Collins, et al. 2019. "A Systematic Review Shows No Performance Benefit of Machine Learning over Logistic Regression for Clinical Prediction Models." *Journal of Clinical Epidemiology* 110: 12–22. https://doi.org/10.1016/j.jclinepi.2019.02.004.

Clopper, C. J., and E. S. Pearson. 1934. "The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial." *Biometrika* 26 (4): 404–13. https://doi.org/10.1093/biomet/26.4.404.

Collins, Gary S., Karel G. M. Moons, Paula Dhiman, et al. 2024. "TRIPOD+AI Statement: Updated Guidance for Reporting Clinical Prediction Models That Use Regression or Machine Learning Methods." *BMJ* 385: e078378. https://doi.org/10.1136/bmj-2023-078378.

Court of Justice of the European Union. 2025a. *Judgment of 4 September 2025, European Data Protection Supervisor v Single Resolution Board, Case C-413/23 P.* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62023CJ0413.

Court of Justice of the European Union. 2025b. *Judgment of 4 September 2025, European Data Protection Supervisor v Single Resolution Board, Case C-413/23 P, ECLI:EU:C:2025:645.* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62023CJ0413.

Davis, Sharon E., Thomas A. Lasko, Guanhua Chen, et al. 2017. "Calibration Drift in Regression and Machine Learning Models for Acute

Kidney Injury." *Journal of the American Medical Informatics Association* 24 (6): 1052–61. https://doi.org/10.1093/jamia/ocx030.

de Jonghe, Annemarieke, Barbara C. van Munster, J. Carel Goslings, et al. 2014. "Effect of Melatonin on Incidence of Delirium Among Patients with Hip Fracture: A Multicentre, Double-Blind Randomized Controlled Trial." *Canadian Medical Association Journal* 186 (14): E547–56. https://doi.org/10.1503/cmaj.140495.

Debray, Thomas P. A., Yvonne Vergouwe, Hendrik Koffijberg, et al. 2015. "A New Framework to Enhance the Interpretation of External Validation Studies of Clinical Prediction Models." *Journal of Clinical Epidemiology* 68 (3): 279–89. https://doi.org/10.1016/j.jclinepi.2014.06.018.

DeLong, Elizabeth R., David M. DeLong, and Daniel L. Clarke-Pearson. 1988. "Comparing the Areas Under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach." *Biometrics* 44 (3): 837–45. https://doi.org/10.2307/2531595.

Diggle, Peter, and Michael G. Kenward. 1994. "Informative Drop-Out in Longitudinal Data Analysis." *Applied Statistics* 43 (1): 49–93. https://doi.org/10.2307/2986113.

Elwyn, Glyn, Dominick Frosch, Richard Thomson, et al. 2012. "Shared Decision Making: A Model for Clinical Practice." *Journal of General Internal Medicine* 27 (10): 1361–67. https://doi.org/10.1007/s11606-012-2077-6.

European Commission. 2025. *Proposal for a Regulation Amending Regulations (EU) 2017/745 and (EU) 2017/746 as Regards Simplifying and Reducing the Burden of the Rules on Medical Devices and in Vitro Diagnostic Medical Devices, COM(2025) 1023 Final.* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC1023.

European Data Protection Board. 2023. *Guidelines 9/2022 on Personal Data Breach Notification Under GDPR, Version 2.0*. https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-92022-personal-data-breach-notification-under_en.

European Data Protection Board. 2025. *Guidelines 01/2025 on Pseudonymisation*. https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2025/guidelines-012025-pseudonymisation_en.

European Parliament and Council of the European Union. 2010. *Directive 2010/63/EU of 22 September 2010 on the Protection of Animals Used for Scientific Purposes*. https://eur-lex.europa.eu/eli/dir/2010/63/oj/eng.

European Parliament and Council of the European Union. 2016. *Regulation (EU) 2016/679 (General Data Protection Regulation, GDPR)*. https://eur-lex.europa.eu/eli/reg/2016/679/oj.

European Parliament and Council of the European Union. 2017a. *Regulation (EU) 2017/745 on Medical Devices (MDR)*. https://eur-lex.europa.eu/eli/reg/2017/745/oj.

European Parliament and Council of the European Union. 2017b. *Regulation (EU) 2017/746 on in Vitro Diagnostic Medical Devices (IVDR)*. https://eur-lex.europa.eu/eli/reg/2017/746/oj.

European Parliament and Council of the European Union. 2024a. *Directive (EU) 2024/2853 on Liability for Defective Products*. https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng.

European Parliament and Council of the European Union. 2024b. *Regulation (EU) 2024/1689 Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*. https://eur-lex.europa.eu/eli/

reg/2024/1689/oj.

European Parliament and Council of the European Union. 2025. *Regulation (EU) 2025/327 on the European Health Data Space (EHDS)*. https://eur-lex.europa.eu/eli/reg/2025/327/oj.

European Parliament and Council of the European Union. 2026. *Regulation (EU) 2026/1744 Amending Regulation (EU) 2024/1689 (Digital Omnibus on AI)*. https://eur-lex.europa.eu/eli/reg/2026/1744/oj.

Federal Office for Information Security (BSI). 2025. *Generative KI-Modelle: Chancen Und Risiken Für Industrie Und Behörden*. https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/KI/Generative_KI-Modelle.pdf?___blob=publicationFile&v=7.

Federal Republic of Germany. 2006. *Animal Welfare Act (Tierschutzgesetz, TierSchG), § 7a*. https://www.gesetze-im-internet.de/tierschg/___7a.html.

Federal Republic of Germany. 2018. *Federal Data Protection Act (Bundesdatenschutzgesetz, BDSG), § 27: Processing for Scientific or Historical Research Purposes and for Statistical Purposes*. https://www.gesetze-im-internet.de/bdsg_2018/___27.html.

Federal Republic of Germany. 2021. *Medical Devices User Reporting and Information Ordinance (Medizinprodukte-Anwendermelde- Und Informationsverordnung, MPAMIV), § 3: Reporting Obligation*. https://www.gesetze-im-internet.de/mpamiv/___3.html.

Federal Republic of Germany. 2024a. *Health Data Use Act (Gesundheitsdatennutzungsgesetz, GDNG), § 6: Further Processing of Care Data for Quality Assurance, Patient Safety and Research*. https://www.gesetze-im-internet.de/gdng/___6.html.

Federal Republic of Germany. 2024b. *Health Data Use Act (Gesundheitsdatennutzungsgesetz, GDNG) of 22 March 2024*. https://www.gesetze-im-internet.de/gdng/.

Federal Republic of Germany. 2026. *Act to Modernise Product Liability Law (Gesetz Zur Modernisierung Des Produkthaftungsrechts Und Zur Änderung Weiterer Vorschriften) of 22 October 2026, Article 1: Product Liability Act (ProdHaftG)*. https://www.recht.bund.de/bgbl/1/2026/321/VO.html.

Feng, Jean, Rachael V. Phillips, Ivana Malenica, et al. 2022. "Clinical Artificial Intelligence Quality Improvement: Towards Continual Monitoring and Updating of AI Algorithms in Healthcare." *Npj Digital Medicine* 5: 66. https://doi.org/10.1038/s41746-022-00611-y.

Finlayson, Samuel G., Adarsh Subbaswamy, Karandeep Singh, et al. 2021. "The Clinician and Dataset Shift in Artificial Intelligence." *New England Journal of Medicine* 385 (3): 283–86. https://doi.org/10.1056/NEJMc2104626.

Flesch, Rudolf. 1948. "A New Readability Yardstick." *Journal of Applied Psychology* 32 (3): 221–33. https://doi.org/10.1037/h0057532.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2010. "Regularization Paths for Generalized Linear Models via Coordinate Descent." *Journal of Statistical Software* 33 (1): 1–22. https://doi.org/10.18637/jss.v033.i01.

Fryback, Dennis G., and John R. Thornbury. 1991. "The Efficacy of Diagnostic Imaging." *Medical Decision Making* 11 (2): 88–94. https://doi.org/10.1177/0272989X9101100203.

Gallifant, Jack, Majid Afshar, Saleem Ameen, et al. 2025. "The TRIPOD-LLM Reporting Guideline for Studies Using Large Language Models."

Nature Medicine 31 (1): 60–69. https://doi.org/10.1038/s41591-024-03425-5.

Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, et al. 2021. "Datasheets for Datasets." *Communications of the ACM* 64 (12): 86–92. https://doi.org/10.1145/3458723.

Gneiting, Tilmann, Fadoua Balabdaoui, and Adrian E. Raftery. 2007. "Probabilistic Forecasts, Calibration and Sharpness." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69 (2): 243–68. https://doi.org/10.1111/j.1467-9868.2007.00587.x.

Goddard, Kate, Abdul Roudsari, and Jeremy C. Wyatt. 2012. "Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators." *Journal of the American Medical Informatics Association* 19 (1): 121–27. https://doi.org/10.1136/amiajnl-2011-000089.

Greenwald, Noah F., Geneva Miller, Erick Moen, Alex Kong, et al. 2022. "Whole-Cell Segmentation of Tissue Images with Human-Level Performance Using Large-Scale Data Annotation and Deep Learning." *Nature Biotechnology* 40 (4): 555–65. https://doi.org/10.1038/s41587-021-01094-0.

Greshake, Kai, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. "Not What You've Signed up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection." *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, 79–90. https://doi.org/10.1145/3605764.3623985.

Griffith, Gareth J., Tim T. Morris, Matthew J. Tudball, et al. 2020. "Collider Bias Undermines Our Understanding of COVID-19 Disease Risk and Severity." *Nature Communications* 11: 5749. https://doi.org/10.1038/s41467-020-19478-2.

Grimm, Volker, Steven F. Railsback, Christian E. Vincenot, et al. 2020. "The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism." *Journal of Artificial Societies and Social Simulation* 23 (2): 7. https://doi.org/10.18564/jasss.4259.

Guermazi, Ali, Chadi Tannoury, Andrew J. Kompel, Akira M. Murakami, et al. 2022. "Improving Radiographic Fracture Recognition Performance and Efficiency Using Artificial Intelligence." *Radiology* 302 (3): 627–36. https://doi.org/10.1148/radiol.210937.

Hartigan, J. A., and M. A. Wong. 1979. "Algorithm AS 136: A K-Means Clustering Algorithm." *Applied Statistics* 28 (1): 100–108. https://doi.org/10.2307/2346830.

Hennig, Christian. 2007. "Cluster-Wise Assessment of Cluster Stability." *Computational Statistics & Data Analysis* 52 (1): 258–71. https://doi.org/10.1016/j.csda.2006.11.025.

Hennig, Christian. 2026. *Fpc: Flexible Procedures for Clustering*. https://doi.org/10.32614/CRAN.package.fpc.

Hernán, Miguel A., Sonia Hernández-Díaz, and James M. Robins. 2004. "A Structural Approach to Selection Bias." *Epidemiology* 15 (5): 615–25. https://doi.org/10.1097/01.ede.0000135174.63482.43.

Hernán, Miguel A., John Hsu, and Brian Healy. 2019. "A Second Chance to Get Causal Inference Right: A Classification of Data Science Tasks." *CHANCE* 32 (1): 42–49. https://doi.org/10.1080/09332480.2019.1579578.

Hernán, Miguel A., and James M. Robins. 2016. "Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available." *American Journal of Epidemiology* 183 (8): 758–64. https://doi.org/

10.1093/aje/kwv254.

Hirshberg, Asher, Bradford G. Scott, Thomas Granchi, et al. 2005. "How Does Casualty Load Affect Trauma Care in Urban Bombing Incidents? A Quantitative Analysis." *The Journal of Trauma: Injury, Infection, and Critical Care* 58 (4): 686–95. https://doi.org/10.1097/01.ta.0000159243.70507.86.

Hirshberg, Asher, Michael Stein, and Raphael Walden. 1999. "Surgical Resource Utilization in Urban Terrorist Bombing: A Computer Simulation." *The Journal of Trauma: Injury, Infection, and Critical Care* 47 (3): 545–50. https://doi.org/10.1097/00005373-199909000-00020.

Hoffrage, Ulrich, and Gerd Gigerenzer. 1998. "Using Natural Frequencies to Improve Diagnostic Inferences." *Academic Medicine* 73 (5): 538–40. https://doi.org/10.1097/00001888-199805000-00024.

Hoffrage, Ulrich, Samuel Lindsey, Ralph Hertwig, and Gerd Gigerenzer. 2000. "Communicating Statistical Information." *Science* 290 (5500): 2261–62. https://doi.org/10.1126/science.290.5500.2261.

Holmer, Amadeus, Jörg Marotz, Philip Wahl, Michael Dau, and Peer W. Kämmerer. 2018. "Hyperspectral Imaging in Perfusion and Wound Diagnostics – Methods and Algorithms for the Determination of Tissue Parameters." *Biomedical Engineering / Biomedizinische Technik* 63 (5): 547–56. https://doi.org/10.1515/bmt-2017-0155.

Holtzman, Ari, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. *The Curious Case of Neural Text Degeneration*. https://doi.org/10.48550/arXiv.1904.09751.

Horst, Fabian, Sebastian Lapuschkin, Wojciech Samek, Klaus-Robert Müller, and Wolfgang I. Schöllhorn. 2019. "Explaining the Unique Nature of Individual Gait Patterns with Deep Learning." *Scientific*

Reports 9 (1): 2391. https://doi.org/10.1038/s41598-019-38748-8.

Hubert, Lawrence, and Phipps Arabie. 1985. "Comparing Partitions." *Journal of Classification* 2 (1): 193–218. https://doi.org/10.1007/BF01908075.

Hundepool, Anco, Josep Domingo-Ferrer, Luisa Franconi, et al. 2012. *Statistical Disclosure Control*. Wiley. https://doi.org/10.1002/9781118348239.

Hurlbert, Stuart H. 1984. "Pseudoreplication and the Design of Ecological Field Experiments." *Ecological Monographs* 54 (2): 187–211. https://doi.org/10.2307/1942661.

Hyndman, Rob J., and George Athanasopoulos. 2021. *Forecasting: Principles and Practice*. 3rd ed. OTexts. https://otexts.com/fpp3/.

Hyndman, Rob J., and Yeasmin Khandakar. 2008. "Automatic Time Series Forecasting: The forecast Package for R." *Journal of Statistical Software* 27 (3): 1–22. https://doi.org/10.18637/jss.v027.i03.

Hyndman, Rob J., and Anne B. Koehler. 2006. "Another Look at Measures of Forecast Accuracy." *International Journal of Forecasting* 22 (4): 679–88. https://doi.org/10.1016/j.ijforecast.2006.03.001.

Inouye, Sharon K., Sidney T. Bogardus, Peter A. Charpentier, et al. 1999. "A Multicomponent Intervention to Prevent Delirium in Hospitalized Older Patients." *New England Journal of Medicine* 340 (9): 669–76. https://doi.org/10.1056/NEJM199903043400901.

International Committee of Medical Journal Editors. 2026. *Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals*. https://www.icmje.org/recommendations/.

International Council for Harmonisation. 2019. *ICH E9(R1): Addendum on Estimands and Sensitivity Analysis in Clinical Trials to the Guideline on Statistical Principles for Clinical Trials*. https://database.ich.org/sites/default/files/E9-R1_Step4_Guideline_2019_1203.pdf.

Ji, Ziwei, Nayeon Lee, Rita Frieske, et al. 2023. "Survey of Hallucination in Natural Language Generation." *ACM Computing Surveys* 55 (12): 1–38. https://doi.org/10.1145/3571730.

Jones, Spencer S., Alun Thomas, R. Scott Evans, et al. 2008. "Forecasting Daily Patient Volumes in the Emergency Department." *Academic Emergency Medicine* 15 (2): 159–70. https://doi.org/10.1111/j.1553-2712.2007.00032.x.

Jurafsky, Daniel, and James H. Martin. 2026. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, with Language Models*. 3rd (online draft). https://web.stanford.edu/~jurafsky/slp3/.

Justice, Amy C., Kenneth E. Covinsky, and Jesse A. Berlin. 1999. "Assessing the Generalizability of Prognostic Information." *Annals of Internal Medicine* 130 (6): 515–24. https://doi.org/10.7326/0003-4819-130-6-199903160-00016.

Kapoor, Sayash, and Arvind Narayanan. 2023. "Leakage and the Reproducibility Crisis in Machine-Learning-Based Science." *Patterns* 4 (9): 100804. https://doi.org/10.1016/j.patter.2023.100804.

Kilkenny, Carol, Nick Parsons, Ed Kadyszewski, et al. 2009. "Survey of the Quality of Experimental Design, Statistical Analysis and Reporting of Research Using Animals." *PLoS ONE* 4 (11): e7824. https://doi.org/10.1371/journal.pone.0007824.

Kuo, Rachel Y. L., Conrad Harrison, Terry-Ann Curran, Benjamin Jones,

et al. 2022. “Artificial Intelligence in Fracture Detection: A Systematic Review and Meta-Analysis.” *Radiology* 304 (1): 50–62. https://doi.org/10.1148/radiol.211785.

Kurland, Brenda F., Laura L. Johnson, Brian L. Egleston, and Paula H. Diehr. 2009. “Longitudinal Data with Follow-up Truncated by Death: Match the Analysis Method to Research Aims.” *Statistical Science* 24 (2): 211–22. https://doi.org/10.1214/09-STS293.

Laird, Nan M., and James H. Ware. 1982. “Random-Effects Models for Longitudinal Data.” *Biometrics* 38 (4): 963–74. https://doi.org/10.2307/2529876.

Lazic, Stanley E. 2010. “The Problem of Pseudoreplication in Neuroscientific Studies: Is It Affecting Your Analysis?” *BMC Neuroscience* 11: 5. https://doi.org/10.1186/1471-2202-11-5.

Lazic, Stanley E., Charlie J. Clarke-Williams, and Marcus R. Munafò. 2018. “What Exactly Is ‘N’ in Cell Culture and Animal Experiments?” *PLOS Biology* 16 (4): e2005282. https://doi.org/10.1371/journal.pbio.2005282.

Leeflang, Mariska M. G., Anne W. S. Rutjes, Johannes B. Reitsma, Lotty Hooft, and Patrick M. M. Bossuyt. 2013. “Variation of a Test’s Sensitivity and Specificity with Disease Prevalence.” *Canadian Medical Association Journal* 185 (11): E537–44. https://doi.org/10.1503/cmaj.121286.

Leek, Jeffrey T., Robert B. Scharpf, Héctor Corrada Bravo, David Simcha, et al. 2010. “Tackling the Widespread and Critical Impact of Batch Effects in High-Throughput Data.” *Nature Reviews Genetics* 11 (10): 733–39. https://doi.org/10.1038/nrg2825.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, et al. 2020. *Retrieval-*

Augmented Generation for Knowledge-Intensive NLP Tasks. https://doi.org/10.48550/arXiv.2005.11401.

Little, John D. C. 1961. "A Proof for the Queuing Formula: $L = \lambda W$." *Operations Research* 9 (3): 383–87. https://doi.org/10.1287/opre.9.3.383.

Little, Roderick J., Ralph D'Agostino, Michael L. Cohen, Kay Dickersin, Scott S. Emerson, et al. 2012. "The Prevention and Treatment of Missing Data in Clinical Trials." *New England Journal of Medicine* 367 (14): 1355–60. https://doi.org/10.1056/NEJMsr1203730.

Liu, Xiaoxuan, Samantha Cruz Rivera, David Moher, et al. 2020. "Reporting Guidelines for Clinical Trial Reports for Interventions Involving Artificial Intelligence: The CONSORT-AI Extension." *Nature Medicine* 26 (9): 1364–74. https://doi.org/10.1038/s41591-020-1034-x.

Lu, Guolan, and Baowei Fei. 2014. "Medical Hyperspectral Imaging: A Review." *Journal of Biomedical Optics* 19 (1): 010901. https://doi.org/10.1117/1.JBO.19.1.010901.

Marcantonio, Edward R., Jonathan M. Flacker, R. John Wright, and Neil M. Resnick. 2001. "Reducing Delirium After Hip Fracture: A Randomized Trial." *Journal of the American Geriatrics Society* 49 (5): 516–22. https://doi.org/10.1046/j.1532-5415.2001.49108.x.

Mathis, Alexander, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, et al. 2018. "DeepLabCut: Markerless Pose Estimation of User-Defined Body Parts with Deep Learning." *Nature Neuroscience* 21 (9): 1281–89. https://doi.org/10.1038/s41593-018-0209-y.

Matthews, J. N. S., Douglas G. Altman, M. J. Campbell, and Patrick Royston. 1990. "Analysis of Serial Measurements in Medical Research." *BMJ* 300 (6719): 230–35. https://doi.org/10.1136/bmj.300.6719.230.

McInnes, Leland. 2026. *Using UMAP for Clustering*. https://umap-learn.readthedocs.io/en/latest/clustering.html.

McInnes, Leland, John Healy, Nathaniel Saul, and Lukas Großberger. 2018. "UMAP: Uniform Manifold Approximation and Projection." *Journal of Open Source Software* 3 (29): 861. https://doi.org/10.21105/joss.00861.

Medical Device Coordination Group. 2025. *MDCG 2019-11 Rev. 1: Qualification and Classification of Software – Regulation (EU) 2017/745 and Regulation (EU) 2017/746*. https://health.ec.europa.eu/system/files/2025-06/mdcg_2019_11_en.pdf.

Miller, Mark Roman, Fernanda Herrera, Hanseul Jun, James A. Landay, and Jeremy N. Bailenson. 2020. "Personal Identifiability of User Tracking Data During Observation of 360-Degree VR Video." *Scientific Reports* 10: 17404. https://doi.org/10.1038/s41598-020-74486-y.

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, et al. 2019. "Model Cards for Model Reporting." *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–29. https://doi.org/10.1145/3287560.3287596.

Monks, Thomas, Christine S. M. Currie, Bhakti Stephan Onggo, Stewart Robinson, Martin Kunc, and Simon J. E. Taylor. 2019. "Strengthening the Reporting of Empirical Simulation Studies: Introducing the STRESS Guidelines." *Journal of Simulation* 13 (1): 55–67. https://doi.org/10.1080/17477778.2018.1442155.

Moons, Karel G. M., Johanna A. A. Damen, Tabea Kaul, et al. 2025. "PROBAST+AI: An Updated Quality, Risk of Bias, and Applicability Assessment Tool for Prediction Models Using Regression or Artificial Intelligence Methods." *BMJ* 388: e082505. https://doi.org/10.1136/bmj-2024-082505.

Mould, D. R., and R. N. Upton. 2012. "Basic Concepts in Population Modeling, Simulation, and Model-Based Drug Development." *CPT: Pharmacometrics & Systems Pharmacology* 1 (9): e6. https://doi.org/10.1038/psp.2012.4.

Mould, D. R., and R. N. Upton. 2013. "Basic Concepts in Population Modeling, Simulation, and Model-Based Drug Development – Part 2: Introduction to Pharmacokinetic Modeling Methods." *CPT: Pharmacometrics & Systems Pharmacology* 2 (4): e38. https://doi.org/10.1038/psp.2013.14.

Nagendran, Myura, Yang Chen, Christopher A. Lovejoy, et al. 2020. "Artificial Intelligence Versus Clinicians: Systematic Review of Design, Reporting Standards, and Claims of Deep Learning Studies." *BMJ* 368: m689. https://doi.org/10.1136/bmj.m689.

NC3Rs. 2026. *The 3Rs*. https://nc3rs.org.uk/who-we-are/3rs.

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." *Science* 366 (6464): 447–53. https://doi.org/10.1126/science.aax2342.

OWASP Foundation. 2024. *OWASP Top 10 for LLM Applications 2025*. https://genai.owasp.org/llm-top-10/.

Paszke, Adam, Sam Gross, Francisco Massa, et al. 2019. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." *Advances in Neural Information Processing Systems 32*, 8024–35. https://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine*

Learning Research 12: 2825–30. https://jmlr.org/papers/v12/pedregosa11a.html.

Peng, Roger D. 2011. "Reproducible Research in Computational Science." *Science* 334 (6060): 1226–27. https://doi.org/10.1126/science.1213847.

Percie du Sert, Nathalie, Viki Hurst, Amrita Ahluwalia, et al. 2020. "The ARRIVE Guidelines 2.0: Updated Guidelines for Reporting Animal Research." *PLOS Biology* 18 (7): e3000410. https://doi.org/10.1371/journal.pbio.3000410.

Petzold, Linda. 1983. "Automatic Selection of Methods for Solving Stiff and Nonstiff Systems of Ordinary Differential Equations." *SIAM Journal on Scientific and Statistical Computing* 4 (1): 136–48. https://doi.org/10.1137/0904010.

Quarto. n.d. *Creating a Book*. https://quarto.org/docs/books/.

R Core Team. 2026. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. https://www.R-project.org/.

Rackauckas, Christopher, and Qing Nie. 2017. "DifferentialEquations.jl – a Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia." *Journal of Open Research Software* 5 (1): 15. https://doi.org/10.5334/jors.151.

Raue, A., C. Kreutz, T. Maiwald, et al. 2009. "Structural and Practical Identifiability Analysis of Partially Observed Dynamical Models by Exploiting the Profile Likelihood." *Bioinformatics* 25 (15): 1923–29. https://doi.org/10.1093/bioinformatics/btp358.

Riley, Richard D., Joie Ensor, Kym I. E. Snell, et al. 2020. "Calculating

the Sample Size Required for Developing a Clinical Prediction Model." *BMJ* 368: m441. https://doi.org/10.1136/bmj.m441.

Roberts, David R., Volker Bahn, Simone Ciuti, et al. 2017. "Cross-Validation Strategies for Data with Temporal, Spatial, Hierarchical, or Phylogenetic Structure." *Ecography* 40 (8): 913–29. https://doi.org/10.1111/ecog.02881.

Roberts, Michael, Derek Driggs, Matthew Thorpe, et al. 2021. "Common Pitfalls and Recommendations for Using Machine Learning to Detect and Prognosticate for COVID-19 Using Chest Radiographs and CT Scans." *Nature Machine Intelligence* 3 (3): 199–217. https://doi.org/10.1038/s42256-021-00307-0.

Robin, Xavier, Natacha Turck, Alexandre Hainard, et al. 2011. "pROC: An Open-Source Package for R and S+ to Analyze and Compare ROC Curves." *BMC Bioinformatics* 12: 77. https://doi.org/10.1186/1471-2105-12-77.

Rocher, Luc, Julien M. Hendrickx, and Yves-Alexandre de Montjoye. 2019. "Estimating the Success of Re-Identifications in Incomplete Datasets Using Generative Models." *Nature Communications* 10: 3069. https://doi.org/10.1038/s41467-019-10933-3.

Rousseeuw, Peter J. 1987. "Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis." *Journal of Computational and Applied Mathematics* 20: 53–65. https://doi.org/10.1016/0377-0427(87)90125-7.

Rubin, Donald B. 1976. "Inference and Missing Data." *Biometrika* 63 (3): 581–92. https://doi.org/10.1093/biomet/63.3.581.

Sale, Mark, and Eric A. Sherer. 2015. "A Genetic Algorithm Based Global Search Strategy for Population Pharmacokinetic/Pharmacodynamic

Model Selection." *British Journal of Clinical Pharmacology* 79 (1): 28–39. https://doi.org/10.1111/bcp.12179.

Sandve, Geir Kjetil, Anton Nekrutenko, James Taylor, and Eivind Hovig. 2013. "Ten Simple Rules for Reproducible Computational Research." *PLoS Computational Biology* 9 (10): e1003285. https://doi.org/10.1371/journal.pcbi.1003285.

Sargent, Robert G. 2013. "Verification and Validation of Simulation Models." *Journal of Simulation* 7 (1): 12–24. https://doi.org/10.1057/jos.2012.20.

Schapiro, Denis, Hartland W. Jackson, Swetha Raghuraman, Jana R. Fischer, et al. 2017. "histoCAT: Analysis of Cell Phenotypes and Interactions in Multiplex Image Cytometry Data." *Nature Methods* 14 (9): 873–76. https://doi.org/10.1038/nmeth.4391.

Schillinger, Dean, John Piette, Kevin Grumbach, et al. 2003. "Closing the Loop: Physician Communication with Diabetic Patients Who Have Low Health Literacy." *Archives of Internal Medicine* 163 (1): 83–90. https://doi.org/10.1001/archinte.163.1.83.

Scrucca, Luca, Michael Fop, T. Brendan Murphy, and Adrian E. Raftery. 2016. "Mclust 5: Clustering, Classification and Density Estimation Using Gaussian Finite Mixture Models." *The R Journal* 8 (1): 289–317. https://doi.org/10.32614/RJ-2016-021.

Seelig, Julian, Raban A. Heller, Julian Hackler, et al. 2020. "Selenium and Copper Status – Potential Signposts for Neurological Remission After Traumatic Spinal Cord Injury." *Journal of Trace Elements in Medicine and Biology* 57: 126415. https://doi.org/10.1016/j.jtemb.2019.126415.

Seok, Junhee, H. Shaw Warren, Alex G. Cuenca, et al. 2013. "Genomic Responses in Mouse Models Poorly Mimic Human Inflammatory Dis-

eases." *Proceedings of the National Academy of Sciences* 110 (9): 3507–12. https://doi.org/10.1073/pnas.1222878110.

Seyyed-Kalantari, Laleh, Haoran Zhang, Matthew B. A. McDermott, Irene Y. Chen, and Marzyeh Ghassemi. 2021. "Underdiagnosis Bias of Artificial Intelligence Algorithms Applied to Chest Radiographs in Under-Served Patient Populations." *Nature Medicine* 27 (12): 2176–82. https://doi.org/10.1038/s41591-021-01595-0.

Shmueli, Galit. 2010. "To Explain or to Predict?" *Statistical Science* 25 (3): 289–310. https://doi.org/10.1214/10-STS330.

Shokri, Reza, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. "Membership Inference Attacks Against Machine Learning Models." *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18. https://doi.org/10.1109/SP.2017.41.

Singhal, Karan, Shekoofeh Azizi, Tao Tu, et al. 2023. "Large Language Models Encode Clinical Knowledge." *Nature* 620 (7972): 172–80. https://doi.org/10.1038/s41586-023-06291-2.

Sjoding, Michael W., Robert P. Dickson, Theodore J. Iwashyna, et al. 2020. "Racial Bias in Pulse Oximetry Measurement." *New England Journal of Medicine* 383 (25): 2477–78. https://doi.org/10.1056/NEJMc2029240.

Soetaert, Karline, Thomas Petzoldt, and R. Woodrow Setzer. 2010. "Solving Differential Equations in R: Package deSolve." *Journal of Statistical Software* 33 (9): 1–25. https://doi.org/10.18637/jss.v033.i09.

Sounderajah, Viknesh, Ahmad Guni, Xiaoxuan Liu, et al. 2025. "The STARD-AI Reporting Guideline for Diagnostic Accuracy Studies Using Artificial Intelligence." *Nature Medicine* 31: 3283–89. https://doi.org/10.1038/s41591-025-03953-8.

Souppaya, Murugiah, Karen Scarfone, and Donna Dodson. 2022. *Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities*. NIST SP 800-218. National Institute of Standards; Technology. https://doi.org/10.6028/NIST.SP.800-218.

Spiegelhalter, David J., Keith R. Abrams, and Jonathan P. Myles. 2004. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. Wiley. https://doi.org/10.1002/0470092602.

Squair, Jordan W., Matthieu Gautier, Claudia Kathe, Mark A. Anderson, et al. 2021. "Confronting False Discoveries in Single-Cell Differential Expression." *Nature Communications* 12 (1): 5692. https://doi.org/10.1038/s41467-021-25960-2.

Stenum, Jan, Cristina Rossi, and Ryan T. Roemmich. 2021. "Two-Dimensional Video-Based Analysis of Human Gait Using Pose Estimation." *PLOS Computational Biology* 17 (4): e1008935. https://doi.org/10.1371/journal.pcbi.1008935.

Steyerberg, Ewout W., Andrew J. Vickers, Nancy R. Cook, et al. 2010. "Assessing the Performance of Prediction Models: A Framework for Traditional and Novel Measures." *Epidemiology* 21 (1): 128–38. https://doi.org/10.1097/EDE.0b013e3181c30fb2.

Stringer, Carsen, Tim Wang, Michalis Michaelos, and Marius Pachitariu. 2021. "Cellpose: A Generalist Algorithm for Cellular Segmentation." *Nature Methods* 18 (1): 100–106. https://doi.org/10.1038/s41592-020-01018-x.

Swanson, Marianne, Pauline Bowen, Amy Wohl Phillips, Dean Gallup, and David Lynes. 2010. *Contingency Planning Guide for Federal Information Systems*. NIST SP 800-34 Rev. 1. National Institute of Standards; Technology. https://doi.org/10.6028/NIST.SP.800-34r1.

Sweeney, Latanya. 2002. “K-Anonymity: A Model for Protecting Privacy.” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10 (5): 557–70. https://doi.org/10.1142/S0218488502001648.

Takao, Keizo, and Tsuyoshi Miyakawa. 2015. “Genomic Responses in Mouse Models Greatly Mimic Human Inflammatory Diseases.” *Proceedings of the National Academy of Sciences* 112 (4): 1167–72. https://doi.org/10.1073/pnas.1401965111.

Tashman, Leonard J. 2000. “Out-of-Sample Tests of Forecasting Accuracy: An Analysis and Review.” *International Journal of Forecasting* 16 (4): 437–50. https://doi.org/10.1016/S0169-2070(00)00065-0.

Team SimPy. n.d. *SimPy: Discrete Event Simulation for Python*. https://simpy.readthedocs.io/en/latest/.

Textor, Johannes, Benito van der Zander, Mark S. Gilthorpe, Maciej Liśkiewicz, and George T. H. Ellison. 2016. “Robust Causal Inference Using Directed Acyclic Graphs: The R Package ‘Dagitty’.” *International Journal of Epidemiology* 45 (6): 1887–94. https://doi.org/10.1093/ije/dyw341.

Tibshirani, Robert, and Guenther Walther. 2005. “Cluster Validation by Prediction Strength.” *Journal of Computational and Graphical Statistics* 14 (3): 511–28. https://doi.org/10.1198/106186005X59243.

Ucar, Iñaki, Bart Smeets, and Arturo Azcorra. 2019. “simmer: Discrete-Event Simulation for R.” *Journal of Statistical Software* 90 (2): 1–30. https://doi.org/10.18637/jss.v090.i02.

Van Calster, Ben, David J. McLernon, Maarten van Smeden, Laure Wynants, and Ewout W. Steyerberg. 2019. “Calibration: The Achilles Heel of Predictive Analytics.” *BMC Medicine* 17: 230. https://doi.org/10.1186/s12916-019-1466-7.

Varma, Sudhir, and Richard Simon. 2006. "Bias in Error Estimation When Using Cross-Validation for Model Selection." *BMC Bioinformatics* 7: 91. https://doi.org/10.1186/1471-2105-7-91.

Vasey, Baptiste, Myura Nagendran, Bruce Campbell, et al. 2022. "Reporting Guideline for the Early Stage Clinical Evaluation of Decision Support Systems Driven by Artificial Intelligence: DECIDE-AI." *BMJ* 377: e070904. https://doi.org/10.1136/bmj-2022-070904.

Vassilev, Apostol, Alina Oprea, Alie Fordyce, Hyrum Anderson, Xander Davies, and Maia Hamin. 2025. *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations*. NIST AI 100-2e2025. National Institute of Standards; Technology. https://doi.org/10.6028/NIST.AI.100-2e2025.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, et al. 2017. *Attention Is All You Need*. https://doi.org/10.48550/arXiv.1706.03762.

Venables, W. N., and B. D. Ripley. 2002. *Modern Applied Statistics with S*. 4th ed. Springer. https://doi.org/10.1007/978-0-387-21706-2.

Vickers, Andrew J., and Elena B. Elkin. 2006. "Decision Curve Analysis: A Novel Method for Evaluating Prediction Models." *Medical Decision Making* 26 (6): 565–74. https://doi.org/10.1177/0272989X06295361.

Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, et al. 2020. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." *Nature Methods* 17: 261–72. https://doi.org/10.1038/s41592-019-0686-2.

Walt, Stéfan van der, Johannes L. Schönberger, Juan Nunez-Iglesias, et al. 2014. "Scikit-Image: Image Processing in Python." *PeerJ* 2: e453. https://doi.org/10.7717/peerj.453.

Walters, William H., and Esther Isabelle Wilder. 2023. "Fabrication and

Errors in the Bibliographic Citations Generated by ChatGPT." *Scientific Reports* 13: 14045. https://doi.org/10.1038/s41598-023-41032-5.

Wickham, Hadley. 2014. "Tidy Data." *Journal of Statistical Software* 59 (10): 1–23. https://doi.org/10.18637/jss.v059.i10.

Wilkinson, Mark D., Michel Dumontier, IJsbrand Jan Aalbersberg, et al. 2016. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data* 3: 160018. https://doi.org/10.1038/sdata.2016.18.

Wilson, Edwin B. 1927. "Probable Inference, the Law of Succession, and Statistical Inference." *Journal of the American Statistical Association* 22 (158): 209–12. https://doi.org/10.1080/01621459.1927.10502953.

Wilson, Greg, Jennifer Bryan, Karen Cranston, et al. 2017. "Good Enough Practices in Scientific Computing." *PLoS Computational Biology* 13 (6): e1005510. https://doi.org/10.1371/journal.pcbi.1005510.

Wong, Andrew, Erkin Otles, John P. Donnelly, et al. 2021. "External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients." *JAMA Internal Medicine* 181 (8): 1065–70. https://doi.org/10.1001/jamainternmed.2021.2626.

World Health Organization. 2024. *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*. World Health Organization. https://www.who.int/publications/i/item/9789240084759.

Worp, H. Bart van der, David W. Howells, Emily S. Sena, et al. 2010. "Can Animal Models of Disease Reliably Inform Human Studies?" *PLoS Medicine* 7 (3): e1000245. https://doi.org/10.1371/journal.pmed.1000245.

Zech, John R., Marcus A. Badgeley, Manway Liu, et al. 2018. "Variable Generalization Performance of a Deep Learning Model to Detect Pneumonia in Chest Radiographs: A Cross-Sectional Study." *PLOS Medicine* 15 (11): e1002683. https://doi.org/10.1371/journal.pmed.1002683.

Ziemann, Mark, Yotam Eren, and Assam El-Osta. 2016. "Gene Name Errors Are Widespread in the Scientific Literature." *Genome Biology* 17: 177. https://doi.org/10.1186/s13059-016-1044-7.

Zimmerman, Kip D., Mark A. Espeland, and Carl D. Langefeld. 2021. "A Practical Solution to Pseudoreplication Bias in Single-Cell Studies." *Nature Communications* 12 (1): 738. https://doi.org/10.1038/s41467-021-21038-1.

Zou, Hui, and Trevor Hastie. 2005. "Regularization and Variable Selection via the Elastic Net." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67 (2): 301–20. https://doi.org/10.1111/j.1467-9868.2005.00503.x.

A. Edition notes, software and imprint

R. Heller · Date: 2026-11-08 · Last changed: 2026-11-09

A.1. Edition

Edition	First public edition, November 2026
History	Cases, worked examples and chapter structure developed from February 2025, revised regularly until July 2026
Legal and literature status	November 2026 (individual reference dates are given in the chapters)
Online	https://cttir.github.io/clinical-data-science-in-r/
Source	https://github.com/CTTIR/clinical-data-science-in-r
Companion course	https://cttir.github.io/clinical-data-science-in-r-course/

Changes between editions are recorded in the repository history. Corrections are welcome as issues.

A.2. Software

All R results in this edition were computed with the versions below. The build stores the results in `_freeze/`, so the published version can be reproduced without re-running every chapter; a fresh run with newer package versions may differ in the last digits of some estimates.

Component	Version
R	4.6.1
ggplot2	4.0.3
dplyr	1.2.1
tidyr	1.3.2
lme4	2.1.0
nlme	3.1.169
deSolve	1.42
glmnet	5.2
pROC	1.19.1
dagitty	0.3.4
class	7.3.23
cluster	2.1.8.2
mclust	6.1.3
simmer	4.4.7
forecast	9.0.2
digest	0.6.39
knitr	1.52
rmarkdown	2.32

Quarto renders the book; the version used for the published build is set in the repository's GitHub Actions workflow.

A.3. Legal status

The chapters on regulation, data protection and security (Chapter 22 to Chapter 26) give the reference date and source for each legal statement. Statements about dates of application refer to the version of the legal text in force in November 2026. Local responsibilities and national implementation must be checked for each institution.

A.4. Use of AI tools

Used for support, including Chat AI (KISSKI/GWDG) and locally hosted models, some self-trained or weight-edited (Linux, Intel/NVIDIA; no data shared with third parties). All content and references were revised and verified by the author.

A.5. Imprint

Clinical Data Science in R: Methods in Clinical AI and Translational Trauma Research

© 2025–2026 Raban Heller. ORCID: 0000-0001-8006-9742.

Text and figures: CC BY 4.0. Code: MIT licence.

This book is provided “as is”, without warranty of any kind. The author assumes no responsibility for errors or omissions. The content does not replace individual medical, legal or statistical advice, and no example describes a real patient, product or institution.

Contact for corrections and suggestions: https://github.com/CTTIR/clinical-data-science-in-r/issues

B. Chapter register

R. Heller · Date: 2026-11-08 · Last changed: 2026-11-09

The register lists every chapter with its guiding question, intended readers, prerequisites, learning activity, course workshop and primary source. It is also available as a machine-readable file: `editorial/chapter-register.csv`.

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
I	Clinical decisions and the translational chain	Who will do something differently because of the output?	Medical students, clinical researchers and doctoral candidates new to data analysis	Required (W1 core reading)	Yes: first R steps, pathway data frame, first ggplot2 graphic with expected output, decision-map completeness check, rows vs independent units of both teaching datasets, timing counter-check (4 stopifnot)	fig-b01-pathway (pathway timeline), tbl-b01-map (decision map)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
I	Describe, predict, explain	Which target quantity answers the actual question?	Clinicians and researchers planning analyses of routine or registry data	Required (W5 core reading)	Yes: dagitty adjustment sets, simulation of 4,000 patients, crude vs adjusted vs ICU-only risk differences, explanation of the sign flip, 100 repetitions, prognostic logistic model, counter-check without the delay-to-ICU arrow (6 stopifnot)	fig-b02-dag (DAG), tbl-b02-questions (three questions)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
I	How data arise: missingness and denominators	Which patient, animal or sample is missing from the table?	Students and clinical researchers who collect or receive clinical or experimental data	Required (W1 core reading)	Yds: patient and visit tables, dictionary, counts and means per window, two-path count of patients, measurements and linked records, spaghetti plot, denominator figure, CSV export, HSI pseudoreplication, k-anonymity of a registry extract, counter-check with hidden values (7 stopifnot)	fig-b03-trajectories, fig-b03-denominators, tbl-b03-status, tbl-b03-dictionary

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
I	Reconstructing a research project in R	Can a colleague regenerate the number in the figure?	Anyone handing over, reviewing or reusing an analysis, R beginners	Supplementary (recommended reading)	Yes: R objects, relative path, base graphic, long and wide, five scripts executed in memory, two-path counts, target value V4 = 0.56 m/s (n = 7 of 12), SHA-256 comparison with the B03 file, sessionInfo, Quarto report and Git listings, artefact register, firmware counter-check (9 stopifnot)	fig-b04-windows (base R), tbl-b04-artefacts

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
II	From baseline to learning model	Which additional effort delivers an improvement that matters for a decision?	Clinicians, doctoral researchers and reviewers building or appraising a first prediction model	Supplement (recommended reading for W1 and W3)	Yes: synthetic delirium cohort (1,200 patients, constant IDs, 800/400 split), exported to data/b05-delirium-cohort.csv, 8 stopifnot checks	Fig. 5.1 test AUC with bootstrap intervals and cases found at 30% capacity (n = 400 test patients, 103 events), Table 5.1 where each step happens
II	Accuracy, calibration and clinical utility	When does a good ranking lead to the wrong action?	Clinicians, nurses, doctoral researchers and reviewers who set thresholds	Required (W3 core reading)	Yes: reads the B05 test patients (checksum of test IDs), 9 stopifnot checks	Fig. 6.1 calibration plot (10 groups of 40, n = 400), Fig. 6.2 decision curve, Tables 6.1 and 6.2

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
II	Reading an AI study critically	Which findings support measurement quality, and which support better care?	Clinicians preparing a journal club, decision-support committees, reviewers	Required (W2 core reading)	Yes: reconstruction from published counts with stopifnot (3 checks), journal-club card as data frame	Table 7.1 tools and checks, Table 7.2 checked statement, Table 7.3 reconstructed 2 x 2 table (n = 38,455 hospitalisations, published)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
II	Language models in clinic and science	Which statement can be checked after generation?	Clinicians, researchers, teachers and students who use or review generated text	Required (W2 core reading)	Yes: stored answers exported to data/b08-stored-answers.csv, 8 stopifnot checks	Fig. 8.1 share of false, true and other sentences at three temperatures (2,000 generated sentences each), Table 8.1 evidence card
III	A fracture alert in the emergency department	What does an alert mean when the prevalence changes?	Clinicians, medical students and reviewers acting on an image-analysis alert	required (W3 core reading)	Base R: fourfold counts, PPV/NPV, alerts per shift, counter-check with specificity 80%. stopifnot: 45/140 and 180/260, 32.1% and 69.2%. Download: data/b09-fracture-alert-cohorts.csv	fig-b09-icons (icon array, n = 1,000 patients per setting), fig-b09-ppv (PPV vs prevalence), tbl-b09-fourfold

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
III	Hyperspectral imaging	What separates a biological signal from illumination and device characteristics?	Doctoral researchers and preclinical teams analysing spectral images, reviewers	required (W4 core reading)	Base R plus class::knn and prcomp. stopifnot: grey-tile tolerance, pixel vs animal p-value pattern, index accuracy >= model accuracy on new animals, 50% under stale calibration, brightness-only counter-check	fig-b10-chain (measurement chain), fig-b10-index (index per animal, n = 8), tbl-b10-units

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
III	Measuring gait from video	Are two measurement methods similar enough for this rehabilitation decision?	Physiotherapists, rehabilitation and trauma clinicians, researchers comparing measurement methods	supplementary (W4 gait station)	Base R plus ggplot2. stopifnot: two hand-checked angles, step-level limits outside and 10-step limits inside tolerance, wider bias CI with patients as unit, bias shift with 30 degree yaw. Download: data/b11-gait-steps.csv	fig-b11-chain, fig-b11-perspective (error vs camera yaw), fig-b11-ba (Bland-Altman, n = 24 patients)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
III	Spatial biology and image analysis	How does a segmentation decision change the biological conclusion?	Researchers in pathology and translational immunology, doctoral researchers, reviewers of spatial analyses	supplementary (W4 tissue station)	Base R plus ggplot2. stopifnot: design (12 patients, 36 fields, minimum distance), monotone shares, cell-level $p < 0.01$ but patient-level $p > 0.05$ at 4 micrometres, ratio test $p > 0.05$, edge correction raises share	fig-b12-chain, fig-b12-field (detail at 1 and 4 micrometres), fig-b12-shift (share and ratio vs expansion, n = 12 patients)
IV	Longitudinal biomarkers	Whose course does the estimated slope describe?	Clinical researchers, doctoral students and reviewers working with repeated laboratory values	required (core reading W5)	Yes: synthetic cohort of 36 persons (186 of 216 planned values), 21 stopifnot conditions, CSV data/b13-selenium-synthetic.csv	fig-b13-spaghetti (spaghetti plot with persons measured per visit), fig-b13-delta (tipping-point curve)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
IV	Cluster gradient or artefact	When does a figure become a testable subtype hypothesis?	Translational researchers, doctoral students and reviewers of patient cluster maps	supplementary (recommended reading W5)	Yes: 300 synthetic persons from one continuous gradient, 12 stopifnot conditions	fig-b14-slices (PCA map and stacked PC1 histogram), fig-b14-countercheck (gradient vs gap)
IV	Mechanistic models and pharmacometrics	Which sampling times make it possible to separate clearance from volume?	Clinical pharmacology trainees, doctoral students planning sampling, reviewers of PK analyses	required (core reading W6, alternative to B18)	Yes: virtual parameters only (dose 100 mg, V 20 L, CL 5 L/h), 16 stopifnot conditions	fig-b15-identifiability (curves coinciding at 4 h, normalised sensitivities)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
IV	Bayes, small samples and the 3Rs	What uncertainty remains even when an event rate looks striking?	Preclinical researchers, doctoral students, animal welfare bodies, reviewers of small experiments	supplement (recommended reading W5 and W6)	Yes: 10 fictitious rats (7/9 bridged), 16 stopifnot conditions	fig-b16-posteriors (priors, posteriors and intervals)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
V	Forecasting patient flow	Is the forecast interval of an emergency department usable for next week?	ED and nursing leads planning rosters, hospital operations analysts, clinicians reading forecast dashboards	Required	Yes: synthetic negative-binomial hourly arrivals (26 weeks), forecast::snaive checked against hand computation, temporal train/test check; writes data/b17-ed-arrivals.csv	fig-b17-coverage (weekly coverage, n = 18 forecast weeks), fig-b17-surge, fig-b17-shiftmap, tbl-b17-summary

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
V	Mass casualty incidents, operating capacity and simulation	Which resource limits care in the scenario under study?	Trauma surgeons, anaesthetists and nurses writing mass casualty plans, emergency planners, operations analysts	Required (elective with B15)	Yes: simmer::add_dataframe with stored arrivals/service times, digest fingerprint of event lists; writes data/b18-event-list.csv	fig-b18-waits (wait for OR by triage and scenario, 200 replications), tbl-b18-summary

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
V	From one centre or species to the next	To which new case should the claim refer?	Doctoral researchers in preclinical imaging, clinical teams planning validations, reviewers	Required	Yds: base R hand calculation with stopifnot target values, class::knn on synthetic spectra (seed 4419, separate knn_ objects); writes data/b19-spectra.csv incl. fold columns	fig-b19-splitplan (printable split plan, n = 4 animals), fig-b19-spectra, tbl-b19-levels, tbl-b19-splitplan, tbl-b19-repeat

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
V	Pilot operation, impact and monitoring	Which observation would stop the introduction?	Clinical leads of pilots, physicians and nurses acting as human oversight, quality and patient safety officers, reviewers of early clinical evaluations	Supplementary	Yes, base R simulation (seed 2003), qbeta exact intervals, Poisson stop rule, suppression check	fig-b20-pchart (weekly night-shift alert share, n = 20 weeks), tbl-b20-sheet (monitoring sheet)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VI	Ethics at the bedside and in study design	Who bears the harm of a wrong decision, and who can object?	Clinicians, nurses, researchers, reviewers	Supplement	Yes, synthetic group counts for setting A, Wilson intervals checked against prop.test(), counter-checks for setting B and larger denominators (stopifnot)	fig-b21-subgroups (error rates by group, n = 1,000 patients)
VI	Regulating research software and clinical AI	Which change in the intended purpose changes the regulatory question?	Research teams, clinical IT, regulatory staff, clinicians, reviewers	Required (W7)	No lab (decision exercise), timeline figure with date checks (stopifnot)	fig-b22-timeline (legal dates, status 8 November 2026)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VI	Data protection along the entire data flow	Who receives which data, for what purpose, and for how long?	Clinical researchers, data stewards, teachers, reviewers	Required (W7)	Yes: data-flow register with rule checks, two counter-check versions, 72-hour clock (stopifnot)	fig-b23-flow (data flow with identifiability, boundaries and deletion points, fictional)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VI	Information security and recovery	How does a team protect findings data, the model version and availability?	Clinical IT, research software engineers, data stewards, clinical leads	Supplementary	Yes tally-based access test, SHA-256 release check with digest, hash-chained log, restore test with three synthetic cases, two counter-checks (stopifnot)	fig-b24-threat (threat model with trust boundaries and attack surface, fictional)

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VII	Security risks of learning and generative systems	Which instruction or input may cross a trust boundary?	clinical teams introducing an assistant or model service, research software engineers, information security officers, reviewers	supplement-security	yes, simulated access service (8 synthetic requests, 3 configurations), checksum check of a model artefact with digest, minimal log, stopifnot	fig-b25-control-matrix, tbl-b25-boundaries

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VII	Secondary use and responsible open science	Which form of scientific traceability holds up when the data are sensitive?	clinical and doctoral researchers who publish, data stewards, ethics and data access committees, reviewers	supplement	synthetic registry extract (600 records, rare group of 3), equivalence classes, homogeneity, primary and secondary suppression, data dictionary check, stopifnot	fig-b26-classes, tbl-b26-access-matrix, tbl-b26-ehds

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VII	Model passport, study report and decision	Which evidence does an independent team need for exactly this target statement?	project groups, doctoral researchers before colloquium or submission, reviewers, decision makers	required	yds: B19 teaching case (4/4 spectra known animals, 3/6 spectra and 1/2 animals new animals, ID overlap 0), exact intervals with qbeta, passport built from the analysis, stopifnot	fig-b27-denominators, tbl-b27-reporting-paths, tbl-b27-passport, tbl-b27-decision-sheet, tbl-b27-counter-review, tbl-b27-blank-passport, tbl-b27-blank-decision

Table B.1.: Chapters B01–B28 with guiding question, course workshop and lab.

B-ID	Part	Title	Guiding question	W-ID	Path status	R lab
VII	Case atlas, glossary and transfer	From which decision do I find the right method again?	all readers starting a project, course participants in W1 and W8, supervisors	required	yes: tables stored as data, consistency checks (all 28 chapters reachable, no dangling IDs, at least 60 glossary terms), route search for two reader tasks, path extension when the target changes, stopifnot	fig-b28-atlas-map, tbl-b28-risk, tbl-b28-goal, tbl-b28-data, tbl-b28-glossary, tbl-b28-mistakes, tbl-b28-next

C. Checklist for contributions and revisions

R. Heller · Date: 2026-05-17 · Last changed: 2026-05-17

For every new edition or revised chapter:

- The purpose, readership and prerequisites of the chapter are stated.
- Authorship and licences of reused sources are correct.
- Every claim, citation, figure caption and cross-reference has been checked; DOIs resolve.
- Every number is labelled as synthetic, published, assumption or open.
- The examples run in a fresh R session; the software status is documented in Appendix A.
- Legal statements distinguish publication, entry into force and date of application, and state the reference date.
- Navigation, contrast, alternative text, print view and narrow screens have been checked.
- HTML and EPUB output have been inspected before distribution.
- Open issues and the exact edition are recorded in the edition notes.