BP703T · B.Pharm Semester VII · Theory Notes

AI in Clinical Applications

Unit II — AI in Drug Safety and Pharmacovigilance (6 hours)

Unit II (6 hours)

Pills & Python Notes

From-zero teaching standard: every technical term is defined before use.

Syllabus coverage: Unit II (6 hours).

PILLS & PYTHON

Pharmacy • Data • AI • Healthcare

Instagram @pillsandpython · pillsandpython.com · YouTube @PillsAndPython

CHAPTER 4: Pharmacovigilance Systems & Real Datasets — FAERS/EudraVigilance Structure & Formats, Frequency Analysis & Signal Detection Concepts

Learning Objectives

After studying this chapter, the student will be able to:

1 Define pharmacovigilance and explain why it exists.

2 Describe the flow of an adverse drug reaction report from patient to database.

3 Describe the structure and data formats of FAERS and EudraVigilance in general terms.

4 Explain frequency analysis and signal detection as concepts, without invented statistics.

5 Explain why AI and machine learning are applied to pharmacovigilance data.

Definitions

- **Pharmacovigilance (PV):** The science of detecting, assessing, understanding and preventing adverse effects of medicines. Short memory: "drug safety watch."
- **Adverse Drug Reaction (ADR):** A harmful, unintended response to a medicine given at a normal dose.
- **Adverse Event (AE):** Any harmful medical happening during drug treatment — it may or may not be caused by the drug. (All ADRs are adverse events, but not all adverse events are ADRs.)
- **Spontaneous reporting:** Voluntary reports of suspected ADRs sent by doctors, pharmacists, nurses or patients to the safety authority — the backbone of pharmacovigilance.
- **FAERS:** The United States Food and Drug Administration's spontaneous ADR reporting database — a real, public system where suspected reactions are collected.
- **EudraVigilance:** The European Medicines Agency's system for collecting and managing reports of suspected ADRs for medicines in Europe — a real system with its own structure and formats.
- **Signal (in pharmacovigilance):** A hint from the data suggesting a *new, previously unknown* possible link between a drug and an adverse event — it needs further investigation, not immediate conclusion.
- **Signal detection:** The process of systematically searching report databases for such hints.
- **Frequency analysis:** Counting how often a drug–event pair appears in the database compared with what would be expected — the basic idea behind signal detection.
- **Disproportionality:** The core idea of frequency analysis: a drug–event pair is "disproportionate" if it is reported much more often with that drug than with other drugs.

Terminology Table

Term	Meaning in very simple words
Pharmacovigilance	The organised watch for harmful effects of medicines after they are marketed.
ADR	Harmful, unintended effect of a drug at a normal dose.
Adverse event	Any bad medical event during treatment (cause unknown).
Spontaneous report	A voluntary "I suspect this drug caused this problem" message sent to the authority.
Reporter	The person sending the report — doctor, pharmacist, nurse or patient.
FAERS	The US FDA's public database of suspected ADR reports.
EudraVigilance	The European Medicines Agency's ADR report database system.
Signal	A data hint of a possible new drug–harm link, needing investigation.
Signal detection	Systematic searching of the database for such hints.
Frequency analysis	Counting-based comparison of how often drug–event pairs are reported.
Disproportionality	"Reported more often than expected" — the trigger for a closer look.
Confounding	A hidden factor that distorts the picture (e.g. the disease itself causes the symptom, not the drug).
Under-reporting	The known problem that only a fraction of real ADRs ever gets reported.
Structured data	Information stored in fixed fields/columns (drug name, age, reaction) that computers can count.
Narrative	The free-text story of what happened, written by the reporter.

4.1 Why pharmacovigilance exists — the story

- Clinical trials test a drug in a few thousand selected patients for a short time. Rare harms, harms in children or the elderly, and harms from long use often appear only *after* the drug reaches millions of people.
- Pharmacovigilance is the safety net that catches these late-appearing harms: anyone who suspects a drug caused harm can report it, and experts watch the growing pile of reports for patterns.

Key takeaway: Trials prove a drug works; pharmacovigilance watches whether it keeps being safe in the real world.

4.2 The journey of one ADR report — flowchart

```
Patient takes medicine and feels a suspected harmful effect
        |
        v
Doctor / pharmacist / nurse / patient notices and suspects the drug
        |
        v
A report is written: WHO (age, sex), WHICH DRUG (name, dose),
WHAT HAPPENED (reaction, dates), OUTCOME (recovered? hospitalised?)
        |
        v
Report sent to the national authority
(e.g. to the PvPI — Pharmacovigilance Programme of India — in India)
        |
        v
Report stored in a large structured database (FAERS in the USA,
EudraVigilance in Europe) alongside millions of other reports
        |
        v
Computers + experts run frequency analysis and signal detection:
"Is this drug–reaction pair reported more often than expected?"
        |
        v
Possible signal -> detailed medical review -> regulatory action
(warning label, restricted use, or in serious cases withdrawal)
```

Key takeaway: One report is a whisper; thousands of reports in a database become a voice that computers can analyse.

4.3 Structure and data formats of real systems — FAERS and EudraVigilance (generic description)

Both systems collect the *same kind of information*, organised in structured tables. The description below is generic — it describes the structure and formats, not any specific numbers of records or dates.

What one report (one "case") contains — the standard fields:

Field group	What is recorded	Example of the format
Patient	Age, sex (no names — privacy protected)	Structured fields: age in years, sex as M/F
Drug(s)	Suspected drug name(s), dose, dates of use	Structured fields; drug names standardised
Reaction(s)	The suspected adverse reaction(s)	Coded with a standard medical dictionary
Dates	When the drug was started, when the reaction began	Date fields
Outcome	Recovered, recovering, hospitalised, fatal, etc.	Fixed categories
Reporter	Doctor, pharmacist, patient, etc.	Fixed categories
Narrative	Free-text story of the event	Text field

FAERS (USA):

- The US FDA's spontaneous reporting database.

- Data are released publicly in a structured, table-like format that researchers can download — typically as separate linked tables: one table for patient/reaction information, one for the drugs involved, one for outcomes, and so on, joined by a case number.
- Because it is public and structured, it is one of the most-used real datasets for AI and machine-learning projects in drug safety.

EudraVigilance (Europe):

- The European Medicines Agency's system for managing suspected ADR reports for medicines authorised in Europe.
- Reports follow an internationally agreed electronic format for individual case safety reports, so that reports from many countries fit the same structure.
- Access is tiered: regulators, companies and researchers get different levels of access, with patient privacy protected.

Common features of both systems:

- Reports are *suspected* associations, not proven causes — the database records suspicion, not proof.
- Drug names and reactions are coded with standard dictionaries so that computers can count and compare them (frequency analysis needs standardised words, not free spelling).
- Both suffer from under-reporting and from reports of varying quality — AI models must be built knowing the data are incomplete and noisy.



Key takeaway: FAERS and EudraVigilance are giant structured tables of *suspected* drug–harm pairs — the raw material on which AI-based signal detection works.

4.4 Frequency analysis — the counting idea (conceptual only)

- Imagine a huge table with drugs as rows and reactions as columns. Each cell counts how many reports link that drug with that reaction.
- **Frequency analysis** asks, for each drug–reaction pair: "Is this pair reported *more often than we would expect* given how often the drug is reported overall and how often the reaction is reported overall?"

- **Disproportionality** is the name of this "more often than expected" idea. A strongly disproportionate pair becomes a *candidate signal* — a hint, not a verdict.

Worked toy example — the idea with small numbers (illustrative toy dataset; invented only to show the concept):

- In a tiny teaching database of 1,000 reports: Drug A appears in 100 reports; rash appears in 50 reports overall.
- If drug and rash were unrelated, we would roughly expect about (100 × 50) ÷ 1000 = 5 reports linking Drug A with rash by chance.
- Suppose we actually see 25 reports of Drug A + rash — five times the expected number. This *disproportionality* (25 observed vs ~5 expected) is the kind of pattern frequency analysis flags for expert review.
- The expert then checks: is it real, or is it confounding (e.g. the disease treated by Drug A itself causes rashes)? Only after review does a hint become a confirmed signal.

Key takeaway: Frequency analysis = systematic counting that asks "is this pair surprisingly common?" — the computer counts, the human expert judges.

4.5 Signal detection — from hint to action (conceptual overview)

```
Step 1: COUNT        Frequency analysis flags disproportionate
                     drug–reaction pairs (candidate signals)
        |
Step 2: FILTER       Remove pairs already known, duplicates,
                     and obvious confounded cases
        |
Step 3: REVIEW       Medical experts examine the cases, the
                     narratives, and the scientific literature
        |
Step 4: DECIDE       Signal confirmed? -> regulatory action
                     (label warning, restriction, further study)
                     Not confirmed? -> keep watching
```

- **Important:** A signal is never proof. It is the *start* of an investigation, not the end. This is why pharmacovigilance pairs computer counting with human medical judgement.
- **Where AI fits:** Machine-learning models can scan millions of reports faster than humans, combine structured fields with the free-text narrative, adjust for confounders, and rank candidate signals so experts review the most suspicious ones first. (Later chapters cover logistic regression for ADR risk prediction and performance metrics for these models.)

4.6 Classification of pharmacovigilance data sources

Source	Nature	Strength
Spontaneous reports (FAERS, EudraVigilance)	Voluntary suspicions	Huge volume; catches rare events
Clinical trial data	Controlled experiments	Clean, but small and short

Source	Nature	Strength
Electronic health records	Routine care data	Real-world, but messy
Literature case reports	Published stories	Detailed, but scattered

INFOGRAPHIC CONCEPT: A pipeline infographic titled "From a patient's rash to a drug label warning." Five connected panels with arrows: (1) a patient and pharmacist with a speech bubble "I think this tablet caused my rash"; (2) a report form with labelled fields — patient, drug, reaction, dates, outcome; (3) a giant database cylinder labelled "FAERS / EudraVigilance — millions of structured reports"; (4) a computer screen showing a highlighted drug–reaction pair with the caption "reported 5× more often than expected — candidate signal"; (5) a medical expert with a magnifying glass and a drug label gaining a new warning box. Bottom tagline: "Count systematically. Judge medically. Act responsibly." Side note: "Reports record suspicion, not proof."

4.7 Analogy

Pharmacovigilance is like a city's complaint helpline for a new water supply. One complaint about funny-tasting water could be anything — but if complaints about the same taste cluster in neighbourhoods served by one pipeline, engineers investigate that pipeline. Frequency analysis is the clustering; the engineers' inspection is the expert review; fixing the pipe is the regulatory action.

4.8 Pharmacy relevance

- Pharmacists are among the most frequent ADR reporters — reporting a suspected reaction is a core professional duty.
- A good report needs the four essentials: patient details, drug details (name, dose, dates), reaction details (what, when), and outcome — exactly the structured fields databases expect.
- Understanding signal detection helps pharmacists interpret safety alerts: a new warning usually began as a statistical hint in these databases.

4.9 India context

- India's Pharmacovigilance Programme (PvPI) collects ADR reports from across the country and contributes to global drug safety monitoring.
- Indian pharmacists and doctors can report suspected ADRs through the national system; every report strengthens the database that signal-detection methods analyse.
- Because many medicines in India are used in large, diverse populations, Indian reports add valuable real-world evidence to global safety knowledge.

Key points

1 Pharmacovigilance = detecting, assessing, understanding and preventing ADRs; spontaneous reporting is its backbone.

2 One report flows: patient → reporter → national authority → structured database (FAERS/EudraVigilance) → analysis → possible regulatory action.

3 FAERS (US FDA) is a public structured database of suspected ADR reports in linked tables; EudraVigilance (EMA) uses an internationally agreed electronic case format with tiered access.

4 Standard fields: patient, drug(s), coded reaction(s), dates, outcome, reporter type, narrative. Reports record suspicion, not proof.

5 Frequency analysis counts drug–reaction pairs and flags disproportionality ("reported more often than expected") — toy illustration: 25 observed vs ~5 expected.

6 A signal is a hint needing expert review, never proof by itself; AI helps by scanning millions of reports and ranking candidates.

Common confusions

- **Confusion:** "A drug in FAERS with many rash reports definitely causes rashes." **Truth:** The database records *suspicions*; many reports can reflect reporting bias or the underlying disease, not proof of causation.
- **Confusion:** "Signal detection proves a new side effect." **Truth:** It generates a hypothesis; medical review and further study decide.
- **Confusion:** "FAERS and EudraVigilance contain complete ADR data." **Truth:** Under-reporting means they capture only a fraction of real events, and report quality varies.

Common mistakes

1 Treating a disproportionality flag as a confirmed ADR — it is only a candidate signal.

2 Writing ADR reports without the four essentials (patient, drug + dates, reaction + dates, outcome), making the report unusable for analysis.

3 Ignoring confounding — the disease or another drug may explain the event.

4 Inventing or quoting record counts/dates for these databases from memory — describe structure and format, not statistics you cannot verify.

Mnemonic

"REPORT" for a complete ADR report: **R**eaction (what happened), **E**xposure (which drug, dose, dates), **P**atient (age, sex), **O**utcome (recovered/hospitalised), **R**eporter (who you are), **T**ime (dates of drug start and reaction start). Or the pipeline chant: "**C**ount, **F**ilter, **R**eview, **D**ecide."

Exam summary

- Pharmacovigilance: science of ADR detection, assessment, understanding, prevention.
- Report journey: suspicion → structured report (patient/drug/reaction/dates/outcome) → authority → database → frequency analysis → expert review → action.
- FAERS: US FDA public spontaneous-report database, structured linked tables. EudraVigilance: EMA system, standardised electronic case format, tiered access.

- Frequency analysis: counting drug–reaction pairs vs expected; disproportionality flags candidate signals (toy: 25 observed vs ~5 expected).
- Signal = hint needing investigation, not proof. AI/ML scales the counting and ranking for experts.

One-minute revision

Drug harms after marketing are caught by pharmacovigilance. Anyone suspecting harm reports it; reports carry patient, drug, reaction, dates, outcome in structured fields. Millions of reports sit in FAERS (USA, public tables) and EudraVigilance (Europe, standard electronic format). Frequency analysis counts pairs and flags "more often than expected" (disproportionality) — a candidate signal. Experts review; a confirmed signal can change the drug label. Suspicion, not proof.

Short-answer questions

1 Define pharmacovigilance and ADR. How does an adverse event differ from an ADR?

2 List the standard fields recorded in one ADR case report.

3 What is disproportionality in frequency analysis?

4 Why is a signal not considered proof of causation?

Long-answer questions

1 Trace the journey of a suspected ADR from the patient to a regulatory action, with a flowchart. Describe the structure and data formats of FAERS and EudraVigilance in general terms.

2 Explain frequency analysis and signal detection conceptually, using a small numerical illustration of observed-vs-expected reporting. Discuss the roles of the computer and the human expert, and the limitations of spontaneous-report data.

MCQs (with answer key)

1 Pharmacovigilance is mainly concerned with:

a) Drug manufacturing b) Detecting and preventing adverse effects of medicines c) Drug pricing d) Hospital management

1 FAERS is:

a) A European clinical trial registry b) The US FDA's spontaneous ADR reporting database c) An Indian drug price list d) A pharmacy textbook

1 EudraVigilance is managed by:

a) The US FDA b) The European Medicines Agency c) The World Bank d) A private company

1 A "signal" in pharmacovigilance means:

a) Proven proof that a drug causes harm b) A data hint of a possible new drug–harm link needing investigation c) A radio message from a hospital d) A marketing signal for drug sales

1 Frequency analysis in signal detection is based on:

a) Random guessing b) Counting how often a drug–reaction pair is reported versus expected c) Animal experiments only d) Doctor opinions without data

Answer key: 1-b, 2-b, 3-b, 4-b, 5-b

Viva questions

1 Who can report a suspected ADR?

2 What are the four essentials of a good ADR report?

3 Why do FAERS and EudraVigilance use coded dictionaries for drug names and reactions?

4 What is under-reporting, and why does it matter for AI models?

Application-based questions

1 A patient develops a severe skin rash nine days after starting a new medicine. As the reporting pharmacist, write out the structured information you would include in the ADR report, grouped under the standard database fields.

2 A hospital committee sees an AI dashboard flag: "Drug X + liver injury reported 6× more often than expected." Write a short note explaining what this flag means, what it does NOT mean, and the next three steps the committee should take.

References

- Strom, Kimmel & Hennessy — Textbook of Pharmacoepidemiology
- Aggarwal & Reddy — Healthcare Data Analytics
- Steyerberg — Clinical Prediction Models

BP703T — AI in Clinical Applications: Chapters 5–8

CHAPTER 5: Logistic Regression for ADR Risk Prediction

Learning Objectives

After studying this chapter, you will be able to:

1 Define logistic regression, odds, odds ratio, sigmoid function, threshold, and binary outcome in simple words.

2 Explain why ordinary (linear) regression cannot be used when the answer is only "yes" or "no".

3 Describe how logistic regression converts patient information into a probability of an adverse drug reaction (ADR).

4 Interpret a logistic regression coefficient as an odds ratio with a worked example.

5 State the pharmacy relevance of ADR risk prediction models.

Definitions

- **Adverse Drug Reaction (ADR):** Any harmful, unwanted effect of a medicine taken at a normal dose. Example: nausea after taking an antibiotic, or a skin rash after taking a painkiller.
- **Risk prediction:** Guessing (with mathematics, not with luck) how likely something is to happen before it happens. "ADR risk prediction" means estimating the chance that a particular patient will get an ADR.
- **Logistic regression:** A mathematical method that takes patient information (like age, dose, kidney function) and gives out a probability between 0 and 1 (for example, 0.25 = 25% chance) that a "yes/no" event will happen (for example, "will this patient get an ADR — yes or no?").
- **Binary outcome:** An outcome with only two possible answers, like yes/no, pass/fail, ADR/no ADR. The word "binary" means "made of two".
- **Probability:** A number between 0 and 1 that tells how likely an event is. 0 = it will never happen. 1 = it will always happen. 0.5 = fifty-fifty chance.
- **Odds:** Another way of writing chance. Odds = probability ÷ (1 − probability). If the probability of rain is 0.75 (75%), the odds of rain are 0.75 ÷ 0.25 = 3, spoken as "3 to 1".
- **Odds ratio:** How many times bigger the odds become when one factor changes. An odds ratio of 2 means "the odds double". An odds ratio of 1 means "no change".
- **Sigmoid function (logistic function):** The S-shaped curve used in logistic regression. It squeezes any number into a value between 0 and 1, so the output is always a valid probability.
- **Threshold:** A cut-off value chosen by us. For example, "if the predicted probability is 0.5 or more, we call the patient high risk". Everything above the line is "yes", everything below is "no".
- **Predictor (feature, independent variable):** A piece of information used to make the prediction, such as dose in mg, age in years, or kidney function. "Feature" is the computer-science word for the same thing.
- **Coefficient:** A number in the model's equation that tells how strongly one predictor pushes the answer up or down.

Terminology Table (read this first)

Term	Simple meaning	Pharmacy example
Logistic regression	A formula that predicts a yes/no answer as a probability (0 to 1)	Predicting whether a patient will develop nausea from a drug

Term	Simple meaning	Pharmacy example
Binary outcome	Only two possible answers	ADR happens (1) or does not happen (0)
Probability	Chance of an event, from 0 (never) to 1 (always)	0.25 means 25% chance of ADR
Odds	Probability ÷ (1 − probability)	If p = 0.25, odds = 0.25 ÷ 0.75 = 1/3 ≈ 0.33
Odds ratio	Factor by which odds change when a predictor increases by one unit	Odds ratio 1.5 for dose = odds of ADR are 1.5× for each extra unit of dose
Sigmoid function	S-shaped curve that converts any number into 0–1	Converts the model's score into a probability
Log-odds (logit)	The logarithm of the odds; the straight-line part of the model	log-odds = −4.0 + 0.02 × dose
Threshold	Cut-off probability for a yes/no decision	0.5: above = "high risk", below = "low risk"
Predictor / feature	Input information used for prediction	Age, dose, weight, kidney function
Coefficient	Weight of one predictor in the equation	0.02 in "0.02 × dose"
Training data	Past patient records used to build the model	Old hospital records of dose and ADR
Fitted model	The equation after the computer has found the best numbers	logit(p) = −4.0 + 0.02 × dose

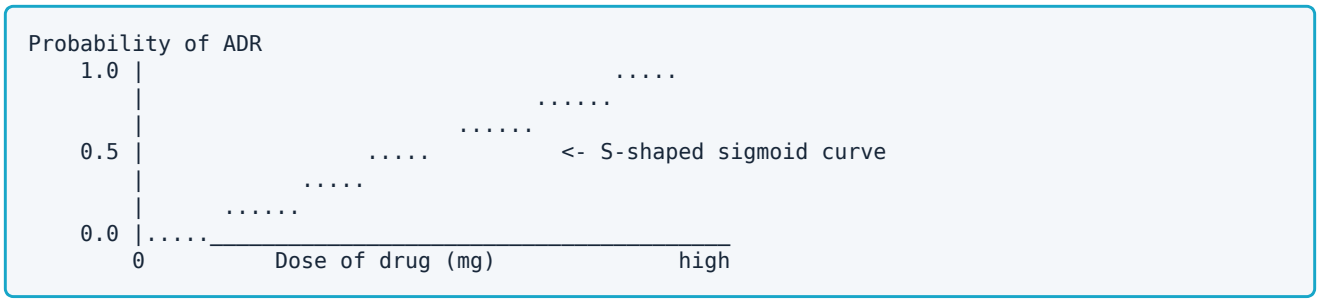
5.1 Why ordinary regression fails for yes/no answers

Technical concept: Simple linear regression draws a straight line through data. It works when the answer is a number (like blood concentration in mg/L). It fails when the answer is only "yes" or "no" (like ADR happened or not).

Simple explanation: Imagine drawing a straight line to predict "will it rain — yes or no?" The line can give silly answers like "rain = 1.4" or "rain = −0.2". Those numbers have no meaning for a yes/no question. We need a method whose answer always stays between 0 and 1, like a proper probability.

Example: We want to predict whether patients get nausea (yes = 1, no = 0) from different doses of a drug. A straight line through the data would predict "nausea = 1.3" for a very high dose — meaningless, because the answer can only be 0 or 1.

Diagram:



A straight line would shoot past 1.0 and below 0.0 (wrong). The S-curve (sigmoid) gently bends and stays between 0 and 1 (correct).

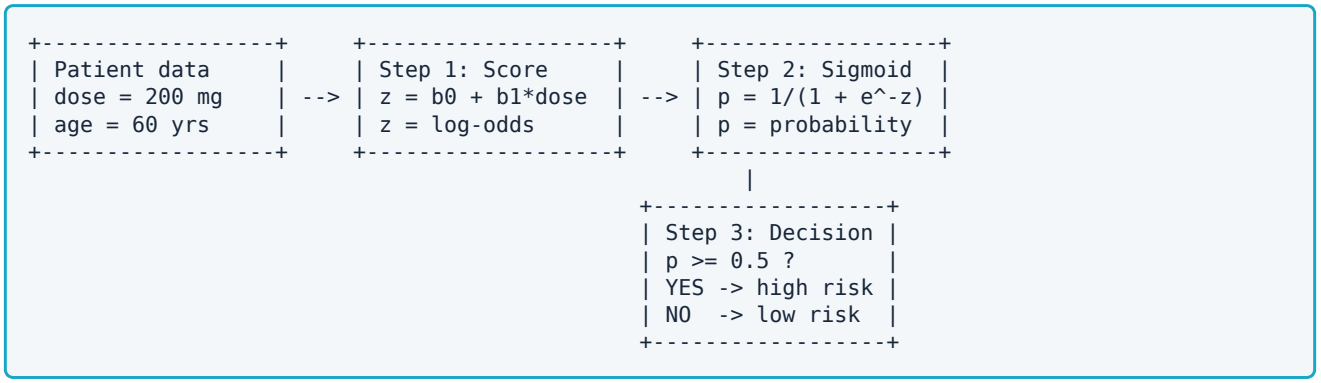
Key takeaway: For yes/no outcomes, replace the straight line with an S-shaped curve. That S-curve method is logistic regression.

5.2 From patient information to probability: the mechanism

Technical concept: Logistic regression works in three steps: (1) combine the predictors into one score using a straight-line equation (the log-odds), (2) pass the score through the sigmoid function to get a probability between 0 and 1, (3) compare the probability with a threshold to make a yes/no decision.

Simple explanation: Think of it like a weighing machine. Each patient factor (dose, age, kidney function) adds some weight to a score. Then the machine squeezes the score into a percentage. Finally, we compare the percentage with our cut-off line.

Flowchart:



The two equations you must know:

1 Score (log-odds): $z = b_0 + b_1 \times (\text{predictor})$

2 Sigmoid: $p = 1 \div (1 + e^{-z})$, where $e \approx 2.718$

Example: A fitted model says: log-odds = −4.0 + 0.02 × dose. For a patient taking 200 mg: z = −4.0 + 0.02 × 200 = −4.0 + 4.0 = 0. Then p = 1 ÷ (1 + e^0) = 1 ÷ (1 + 1) = 0.5. So the predicted ADR probability is 50%.

Key takeaway: Logistic regression = a straight-line score squeezed through an S-curve into a probability.

5.3 WORKED EXAMPLE — predicting ADR probability (illustrative toy dataset)

The dataset below is an **illustrative toy dataset** (invented for teaching; it is not real patient data and not from any real study).

Toy data: 6 patients, one predictor (dose in mg), one binary outcome (ADR: 1 = yes, 0 = no).

Patient	Dose (mg)	ADR
1	50	0
2	100	0
3	150	1
4	200	0
5	250	1

Patient	Dose (mg)	ADR
6	300	1

Suppose a computer has already fitted this **illustrative** model to the toy data:

log-odds (z) = −4.0 + 0.02 × dose

Question A: What is the predicted ADR probability for a patient taking 150 mg?

Step 1 — compute the score z: z = −4.0 + 0.02 × 150 z = −4.0 + 3.0 z = −1.0

Step 2 — apply the sigmoid. We need e^(−z) = e^(−(−1.0)) = e^1.0 ≈ 2.718. p = 1 ÷ (1 + 2.718) p = 1 ÷ 3.718 p ≈ 0.269

Answer: about 0.27, i.e., a 27% predicted chance of ADR.

Question B: What is the predicted ADR probability for a patient taking 300 mg?

Step 1: z = −4.0 + 0.02 × 300 = −4.0 + 6.0 = 2.0 Step 2: e^(−z) = e^(−2.0) ≈ 0.135 p = 1 ÷ (1 + 0.135) = 1 ÷ 1.135 ≈ 0.881

Answer: about 0.88, i.e., an 88% predicted chance of ADR.

Question C: At what dose does the model say "high risk" (probability ≥ 0.5)?

When p = 0.5, the score z = 0 (because 1 ÷ (1 + e^0) = 1/2). Set −4.0 + 0.02 × dose = 0 0.02 × dose = 4.0 dose = 4.0 ÷ 0.02 = 200 mg

Answer: at 200 mg and above, the model classifies the patient as high risk (with threshold 0.5).

Key takeaway: Given the equation, prediction is just two steps — multiply and add to get z, then squeeze with the sigmoid.

5.4 Interpreting coefficients as odds ratios

Technical concept: In logistic regression, e^(coefficient) is the odds ratio for a one-unit increase in that predictor. It tells how the odds of the event change.

Simple explanation: The coefficient 0.02 looks tiny, but its real meaning is: for every extra 1 mg of dose, the odds of ADR are multiplied by e^0.02 ≈ 1.0202 — about a 2% increase in odds per mg.

Worked mini-example (illustrative numbers):

- For a 100 mg dose increase: odds multiply by e^(0.02 × 100) = e^2.0 ≈ 7.39.
- Meaning: a patient on 100 mg more of the drug has about 7.4 times the odds of ADR compared with the lower dose, according to this illustrative model.

Interpretation table:

Odds ratio	Meaning
= 1	No effect; the predictor does not change the odds
> 1	Risk factor; higher predictor value → higher odds of ADR
< 1	Protective factor; higher predictor value → lower odds of ADR

Example: In a model predicting liver injury, suppose the coefficient for "age above 60" gives an odds ratio of 2.5. That means elderly patients have 2.5 times the odds of liver injury compared with younger patients, holding other factors fixed.

Analogy: An odds ratio is like a volume knob for risk. Odds ratio 1 = knob does nothing. Above 1 = knob turns risk up. Below 1 = knob turns risk down.

Key takeaway: Always convert the coefficient with e^(coefficient) before explaining it to a doctor — odds ratios are the language clinicians understand.

5.5 Classification of logistic regression use in drug safety

Type	What it predicts	Pharmacy example
Single-predictor model	ADR risk from one factor	Nausea risk from dose alone
Multiple-predictor model	ADR risk from many factors together	Bleeding risk from dose + age + kidney function
Signal detection support	Which drug–ADR pairs deserve investigation	Flagging a drug–rash pair from spontaneous reports for expert review

5.6 Model building steps (mechanism, in order)

1. **Collect data:** Past records with predictors (dose, age, kidney function) and the outcome (ADR yes/no).
2. **Clean data:** Remove impossible values and handle missing entries.
3. **Choose predictors:** Pick factors that make medical sense (Chapter 6 explains bias if we choose badly).
4. **Fit the model:** The computer finds the coefficients that make the S-curve fit the data best (a method called maximum likelihood — meaning "the curve that makes the observed data most likely").
5. **Interpret:** Convert coefficients to odds ratios.
6. **Validate:** Test the model on new patients it has never seen (Chapter 6 shows how to measure performance).

INFOGRAPHIC CONCEPT: A thermometer-style "ADR Risk Meter". Left side: patient factors flow in as colored pills (dose = red pill, age = blue pill, kidney function = green pill). Middle: the pills drop into a funnel labeled "log-odds score", then through an S-shaped pipe labeled "sigmoid". Right side: a vertical meter from 0% (green) to 100% (red) with a needle pointing at the predicted probability, and a dashed line at the 50% threshold marked "decision line". Caption: "Patient factors in → score → probability out."

5.7 Pharmacy relevance

- **Safer prescribing:** A pharmacist can flag a patient whose predicted ADR probability is high and suggest a lower dose, a different drug, or extra monitoring.

- **Prioritising monitoring:** In a busy hospital, predicted-risk scores help decide which patients need daily blood tests and which need only weekly checks.
- **Supporting pharmacovigilance:** Risk models help ADR monitoring centres decide which reported cases need urgent expert review first.

5.8 India context

- India runs the Pharmacovigilance Programme of India (PvPI), which collects ADR reports from monitoring centres across the country. Risk prediction models can be built on such collected data to find high-risk patient groups — for example, elderly patients taking multiple medicines.
- Indian patients differ from Western trial populations in body weight, diet, and genetics, so a model trained abroad should be re-checked on Indian data before use here (more on this bias problem in Chapter 6).

5.9 Key points

- Logistic regression predicts yes/no outcomes as probabilities between 0 and 1.
- It uses an S-shaped sigmoid curve, not a straight line.
- The equation has two parts: log-odds score, then sigmoid squeeze.
- Coefficients are interpreted as odds ratios using e^(coefficient).
- A threshold (often 0.5) turns the probability into a yes/no decision.

5.10 Common confusions

- **Odds vs probability:** Probability 0.75 is "75%". Odds of 3 means "3 to 1". They are two languages for the same chance. Odds = p ÷ (1 − p).
- **Logistic regression vs linear regression:** Linear regression predicts numbers (straight line); logistic regression predicts yes/no as probabilities (S-curve). Despite its name, logistic regression is a classification method.
- **Correlation vs causation:** If dose and ADR move together, the model only says they are associated. It does not prove the dose caused the ADR. (Chapter 6 explains confounding.)
- **Probability 0.7 does not mean "70% of this patient"** — it means that among 100 similar patients, about 70 would be expected to get the ADR.

5.11 Common mistakes

- Forgetting that the output is a probability, and treating 0.8 as "the patient will definitely get the ADR".
- Interpreting the raw coefficient (0.02) directly instead of converting to an odds ratio.
- Using a threshold of 0.5 blindly when missing an ADR is dangerous — sometimes a lower threshold is safer (Chapter 6).

- Building a model on too few patients, so the S-curve wobbles and cannot be trusted.

5.12 Mnemonics

- **"S for Sigmoid, S for S-shaped, S for Stays between 0 and 1."**
- **"LOGIT = Log Odds Get IT"** — logistic regression first computes the log-odds, then gets the probability.
- **"Odds ratio > 1: Risk is UP; odds ratio < 1: risk goes DOWN; equals 1: nothing's NEW."**

5.13 Exam summary

Logistic regression is the standard method for predicting binary (yes/no) outcomes such as ADR occurrence. A linear score (log-odds) is computed from patient predictors, then transformed by the sigmoid function into a probability between 0 and 1. Coefficients are interpreted as odds ratios (e^coefficient): above 1 = risk factor, below 1 = protective, equal to 1 = no effect. A threshold converts the probability into a high/low risk decision. In pharmacy it supports safer prescribing, targeted monitoring, and pharmacovigilance prioritisation.

5.14 One-minute revision

Yes/no question → use S-curve, not straight line. Score $z = b_0 + b_1x$. Probability $p = 1/(1 + e^{(-z)})$. Coefficient → odds ratio = e^coefficient. Threshold 0.5 → high/low risk decision.

Questions

Short-answer questions:

1 Define logistic regression and state one situation where it is preferred over linear regression.

2 What is an odds ratio? Interpret an odds ratio of 2.5 in one sentence.

3 Write the sigmoid function and explain what it does.

4 Why is a threshold needed after logistic regression gives a probability?

Long-answer questions:

1 Explain, with a labelled diagram, why linear regression is unsuitable for binary outcomes and how logistic regression solves the problem.

2 A model predicts ADR risk with log-odds = −3.0 + 0.03 × dose. Calculate the predicted probability for a 100 mg dose and a 200 mg dose, showing all steps. Interpret the odds ratio for a 50 mg dose increase.

MCQs (answer key at the end of the chapter):

1 Logistic regression is used when the outcome variable is:

a) A continuous number like blood pressure b) A binary yes/no outcome like ADR present or absent c) Always normally distributed d) Measured only in percentages

1 The output of the sigmoid function always lies between:

a) −1 and +1 b) 0 and 100 c) 0 and 1 d) −infinity and +infinity

1 If a coefficient is 0.693, the odds ratio for a one-unit increase is approximately:

a) 0.5 b) 1.0 c) 2.0 d) 6.9

1 An odds ratio of 0.6 for a predictor means the predictor is:

a) A risk factor b) A protective factor c) Has no effect d) Cannot be interpreted

1 With log-odds = −4.0 + 0.02 × dose and threshold 0.5, a patient on 250 mg is classified as:

a) Low risk, because probability is below 0.5 b) High risk, because probability is above 0.5 c) Cannot be classified without more data d) Low risk, because dose is below 200 mg

Answer key: 1-b, 2-c, 3-c, 4-b, 5-b

Viva questions:

1 What does "binary outcome" mean? Give a pharmacy example.

2 What is the difference between odds and probability?

3 Why is the sigmoid curve S-shaped?

4 What happens if we apply a straight line to yes/no data?

Application-based questions:

1 You are a hospital pharmacist. A logistic model gives a patient an 82% predicted probability of bleeding with an anticoagulant. What three actions would you take, and why?

2 A study reports "odds ratio 3.2 for kidney impairment" for a drug's ADR. Explain this number to a junior nurse in simple words.

References

- Aggarwal, C. C. and Reddy, C. K. Healthcare Data Analytics.
- Campbell, M. J., Machin, D. and Walters, S. J. Medical Statistics: A Textbook for the Health Sciences.
- Steyerberg, E. W. Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating.
- Strom, B. L., Kimmel, S. E. and Hennessy, S. Pharmacoepidemiology.

CHAPTER 6: Confusion Matrix and Clinical Performance Metrics — Sensitivity, Specificity, Precision, Recall; Bias and Confounding in Observational Data

Learning Objectives

After studying this chapter, you will be able to:

1 Define confusion matrix, true positive, false positive, true negative, false negative, sensitivity, specificity, precision, recall, accuracy, bias, confounding, and confounder.

2 Build a confusion matrix from model predictions and calculate all four clinical metrics step by step.

3 Explain the trade-off between sensitivity and specificity when the threshold changes.

4 Distinguish bias from confounding, with pharmacy examples of each.

5 Explain why observational pharmacovigilance data needs special caution before trusting a model.

Definitions

- **Confusion matrix:** A 2 × 2 table that compares the model's predictions against the truth. It counts four things: true positives, false positives, true negatives, false negatives. It is the starting point for all performance metrics.
- **True positive (TP):** The model said "yes" and the truth was "yes". Example: the model predicted ADR, and the patient really got the ADR. A correct alarm.
- **False positive (FP):** The model said "yes" but the truth was "no". A false alarm. Example: the model predicted ADR, but the patient was fine.
- **True negative (TN):** The model said "no" and the truth was "no". Correctly saying "all clear".
- **False negative (FN):** The model said "no" but the truth was "yes". A missed case — the most dangerous error in drug safety.
- **Sensitivity (recall):** Of all the patients who TRULY had the event, how many did the model catch? Formula: TP ÷ (TP + FN).
- **Specificity:** Of all the patients who TRULY did NOT have the event, how many did the model correctly clear? Formula: TN ÷ (TN + FP).
- **Precision:** Of all the patients the model FLAGGED as positive, how many truly had the event? Formula: TP ÷ (TP + FP).
- **Recall:** Another name for sensitivity (used more in computer science; "sensitivity" is the medical word).
- **Accuracy:** Overall fraction correct: (TP + TN) ÷ total.
- **Bias:** A systematic (repeatable, one-directional) error that pushes results away from the truth. Not random noise — a tilted weighing scale, not a shaky hand.
- **Confounding:** A mixing of effects. A third factor (the confounder) is linked to both the predictor and the outcome, so the predictor falsely looks responsible.
- **Confounder:** That third factor itself. Classic example: age can confound the link between "number of medicines taken" and "ADR risk".

- **Observational data:** Data collected by watching what happens in real life (hospital records, ADR reports), without a planned experiment. Spontaneous ADR reports are observational data.

Terminology Table (read this first)

Term	Simple meaning	Pharmacy example
Confusion matrix	2×2 table of predicted vs actual	Checking an ADR model's predictions against real outcomes
True positive (TP)	Predicted yes, actually yes	Model flagged ADR; patient got ADR
False positive (FP)	Predicted yes, actually no	Model flagged ADR; patient was fine (false alarm)
True negative (TN)	Predicted no, actually no	Model said safe; patient was safe
False negative (FN)	Predicted no, actually yes	Model said safe; patient got ADR (missed case)
Sensitivity / recall	Caught ÷ all real cases = TP/(TP+FN)	Caught 6 of 8 real ADR cases = 75%
Specificity	Cleared ÷ all real non-cases = TN/(TN+FP)	Cleared 9 of 12 safe patients = 75%
Precision	True alarms ÷ all alarms = TP/(TP+FP)	6 of 9 alarms were real = 66.7%
Accuracy	All correct ÷ total = (TP+TN)/total	15 of 20 correct = 75%
Threshold	Cut-off probability for yes/no	Lowering it catches more cases but raises false alarms
Bias	Systematic tilt away from truth	Only serious ADRs get reported, so mild ADRs look rare
Confounding	Third factor mixes up the cause	Elderly take more drugs AND get more ADRs — age confounds the drug–ADR link
Observational data	Real-life records, no experiment	Spontaneous ADR reports, hospital records

6.1 The confusion matrix: the 2 × 2 table

Technical concept: Every yes/no prediction falls into one of four boxes. Rows = the truth (actual), columns = the model's prediction.

Simple explanation: Think of a security guard at a gate. Four things can happen: guard stops a real thief (TP), guard stops an innocent person (FP), guard lets an innocent person pass (TN), guard lets a real thief pass (FN).

Diagram (the matrix):

```
                    Model predicted
                  YES (positive)   NO (negative)
               +---------------+---------------+
Actual  YES    | TRUE POSITIVE | FALSE NEGATIVE|
(truth) (pos)  |   (TP)        |   (FN)        |
               | correct alarm | missed case   |
               +---------------+---------------+
Actual  NO     | FALSE POSITIVE| TRUE NEGATIVE |
(truth) (neg)  |   (FP)        |   (TN)        |
               | false alarm   | correctly     |
               |               | cleared       |
               +---------------+---------------+
```

Memory rule: Read it as "Predicted X, Actually Y". True = prediction matched truth. False = prediction missed truth. Positive/Negative = what the model said.

Key takeaway: All four clinical metrics are just different fractions built from these four boxes.

6.2 The four metrics: what each one answers

Technical concept: Each metric answers a different practical question. No single metric tells the whole story.

Simple explanation with the questions they answer:

Metric	Formula	Question it answers
Sensitivity (recall)	TP ÷ (TP + FN)	"Of all real cases, how many did we catch?"
Specificity	TN ÷ (TN + FP)	"Of all safe patients, how many did we correctly clear?"
Precision	TP ÷ (TP + FP)	"When the alarm rings, how often is it real?"
Accuracy	(TP + TN) ÷ total	"Overall, how often were we right?"

Analogy: A smoke detector. Sensitivity = does it ring for every real fire? Specificity = does it stay silent when there is no fire? Precision = when it rings, is there really a fire (or just burnt toast)?

Key takeaway: Sensitivity protects patients (catch every case); specificity protects resources (avoid false alarms); precision tells you how much to trust each alarm.

6.3 WORKED EXAMPLE — building the matrix and computing metrics (illustrative toy dataset)

The dataset below is an **illustrative toy dataset** (invented for teaching; not real patient data).

A logistic regression ADR model was tested on 20 patients it had never seen. Threshold = 0.5.

Group	Count	Detail
Actually had ADR (actual yes)	8	Model predicted yes for 6, predicted no for 2

Group	Count	Detail
Actually no ADR (actual no)	12	Model predicted yes for 3, predicted no for 9

Step 1 — fill the four boxes:

- TP (predicted yes, actually yes) = 6
- FN (predicted no, actually yes) = 2
- FP (predicted yes, actually no) = 3
- TN (predicted no, actually no) = 9
- Check total: 6 + 2 + 3 + 9 = 20 ✓

Step 2 — accuracy: Accuracy = (TP + TN) ÷ total = (6 + 9) ÷ 20 = 15 ÷ 20 = 0.75 = 75%

Step 3 — sensitivity (recall): Sensitivity = TP ÷ (TP + FN) = 6 ÷ (6 + 2) = 6 ÷ 8 = 0.75 = 75% Meaning: the model caught 75% of the real ADR cases; it missed 2 out of 8.

Step 4 — specificity: Specificity = TN ÷ (TN + FP) = 9 ÷ (9 + 3) = 9 ÷ 12 = 0.75 = 75% Meaning: the model correctly cleared 75% of the safe patients; it falsely alarmed 3 out of 12.

Step 5 — precision: Precision = TP ÷ (TP + FP) = 6 ÷ (6 + 3) = 6 ÷ 9 ≈ 0.667 = 66.7% Meaning: when the model raised an alarm, it was right about two-thirds of the time.

Worked comparison — what if we lower the threshold to 0.3? (illustrative) Lowering the threshold makes the model say "yes" more easily. Suppose the new boxes become TP = 7, FN = 1, FP = 5, TN = 7 (total still 20).

- New sensitivity = 7 ÷ 8 = 0.875 = 87.5% (better — fewer missed cases)
- New specificity = 7 ÷ 12 ≈ 0.583 = 58.3% (worse — more false alarms)
- New precision = 7 ÷ 12 ≈ 0.583 = 58.3% (worse — alarms less trustworthy)

Key takeaway: There is no free lunch. Lowering the threshold catches more real cases (sensitivity up) but creates more false alarms (specificity and precision down). In drug safety, missing an ADR (FN) is usually more dangerous than a false alarm (FP), so we often prefer higher sensitivity.

6.4 Which metric matters most? (classification by situation)

Situation	Most important metric	Why
Screening for a dangerous ADR	Sensitivity	A missed case can harm or kill; false alarms are acceptable
Confirming before stopping a life-saving drug	Specificity / precision	A false alarm could wrongly stop needed treatment
Comparing two models overall	Accuracy + all four	One number can hide a dangerous weakness
Rare ADR (very few real cases)	Sensitivity + precision	Accuracy looks high even for a useless model (see common mistakes)

6.5 Bias in observational datasets

Technical concept: Bias is a systematic error in how data are collected or selected, which tilts every conclusion in one direction.

Simple explanation: If you only ask happy customers about a restaurant, your review is biased — not because you lied, but because unhappy customers were never asked.

Classification of common biases in drug-safety data:

Bias	Simple meaning	Pharmacy example
Reporting bias	Only some events get reported	Doctors report serious ADRs but skip mild nausea, so mild ADRs look rarer than they are
Selection bias	The studied group is not like the real group	Model trained on city hospital data applied to rural patients
Recall bias	People misremember the past	Patients forget mild side effects when asked months later
Measurement bias	The measuring tool is tilted	A faulty BP machine always reads 10 points high

Diagram — reporting bias:

```
All ADRs that really happen
[ mild ][ mild ][ mild ][ serious ][ mild ][ serious ][ mild ]
   |       |                   |              |
   X       X  (not reported)    v  (reported)  X
                      What the database sees:
                      [ serious ][ serious ]  -> "ADRs are mostly serious!"
```

Key takeaway: A model trained on biased data learns the bias. Garbage in, gospel out — the computer cannot know what was never recorded.

6.6 Confounding: the hidden third factor

Technical concept: A confounder is associated with both the predictor and the outcome, creating a false impression that the predictor causes the outcome.

Simple explanation: Ice-cream sales and drowning deaths rise together in summer. Ice cream does not cause drowning — hot weather (the confounder) causes both.

Diagram — the confounder triangle:

```
            Age (confounder)
            /              \
           /                \
          v                  v
Number of medicines      ADR risk
taken (predictor)        (outcome)
```

Elderly patients take more medicines AND have higher ADR risk because of weaker kidneys and liver. If we ignore age, "more medicines" falsely looks like the full cause of the extra ADRs.

Pharmacy example (illustrative toy scenario): In a toy hospital record, patients on Drug X had double the ADR rate of others. But Drug X was mostly given to elderly patients, and elderly patients get more ADRs from every drug. After adjusting for age, Drug X's extra risk disappeared — age was the confounder, not the drug.

How to handle confounding:

1 **Measure it:** Record the suspected confounder (age, kidney function, other diseases).

2 **Adjust for it:** Include it as a predictor in the logistic regression model (multiple-predictor model from Chapter 5).

3 **Stratify:** Analyse young and old patients separately and check the result holds in both groups.

Key takeaway: Association is not causation. Always ask: "Is there a third factor linked to both sides?"

INFOGRAPHIC CONCEPT: A balance scale titled "Is your model fair?". Left pan holds a 2×2 confusion matrix with the four boxes colored (TP green, TN green, FP orange, FN red). Right pan holds four small gauges labeled Sensitivity, Specificity, Precision, Accuracy. Below the scale, a cracked magnifying glass labeled "Bias" and a triangle diagram labeled "Confounder" with a warning sign. Caption: "Measure all four metrics — and check your data for bias and confounding before you trust them."

6.7 Pharmacy relevance

- **Model evaluation before use:** A pharmacist on a hospital AI committee must read sensitivity and specificity — not just "accuracy 95%" — before allowing an ADR alert system on the wards.
- **Alert fatigue:** A model with low precision fires too many false alarms; doctors start ignoring all alerts, including real ones. Precision matters for daily practice.
- **Confounding in practice:** Before blaming a new drug for ADRs, check whether the patients receiving it were already sicker or older.

6.8 India context

- Spontaneous ADR reporting in India faces under-reporting: busy clinics may report only serious or unusual reactions. Models built on such data will underestimate common, mild ADRs — a reporting bias to disclose in any project (Unit V).
- India's PvPI network collects reports from many centres; differences in reporting habits between centres can create selection bias if one centre dominates the data.

6.9 Key points

- The confusion matrix has four boxes: TP, FP, TN, FN.
- Sensitivity = TP/(TP+FN): catches real cases. Specificity = TN/(TN+FP): clears safe patients. Precision = TP/(TP+FP): trustworthiness of alarms. Accuracy = (TP+TN)/total.

- Lowering the threshold raises sensitivity but lowers specificity and precision.
- Bias = systematic tilt in data collection (reporting, selection, recall, measurement).
- Confounding = a third factor linked to both predictor and outcome; adjust by including it in the model or stratifying.

6.10 Common confusions

- **Sensitivity vs precision:** Sensitivity asks "of all REAL cases, how many caught?" Precision asks "of all ALARMS, how many real?" A model can have 100% sensitivity (catch everything) with terrible precision (alarm constantly).
- **Recall = sensitivity.** They are two names for the same formula, used by different fields (computer science vs medicine).
- **Bias vs confounding:** Bias is an error in how data were collected; confounding is a real third factor mixing up cause and effect. Both mislead, but the fixes differ.
- **Accuracy can lie:** If only 2 of 100 patients get an ADR, a lazy model predicting "no ADR" for everyone scores 98% accuracy but catches zero cases. Always report sensitivity alongside accuracy.

6.11 Common mistakes

- Reporting only accuracy for a rare ADR.
- Confusing the rows and columns of the matrix (rows = actual truth, columns = prediction).
- Treating a false negative as harmless — in drug safety it is the costliest error.
- Claiming "the drug causes ADRs" from observational data without checking confounders like age.

6.12 Mnemonics

- **"SnNout: a SeNsitive test, when Negative, rules OUT disease."** (High sensitivity → a negative result is trustworthy.)
- **"SpPin: a SPecific test, when Positive, rules IN disease."** (High specificity → a positive result is trustworthy.)
- **"SE = Sick-catching: TP ÷ all truly Sick. SP = Safe-spotting: TN ÷ all truly safe."**
- **"BIAS = Built-In Asymmetric Sampling"** — the tilt is built into how data were gathered.

6.13 Exam summary

The confusion matrix (TP, FP, TN, FN) is the foundation for evaluating classification models. Sensitivity = TP/(TP+FN) measures case-catching ability; specificity = TN/(TN+FP) measures correct clearance; precision = TP/(TP+FP) measures alarm trustworthiness; accuracy = (TP+TN)/total measures overall correctness. Threshold choice trades sensitivity against specificity. In observational pharmacovigilance data, reporting bias, selection bias, recall bias, and confounding (e.g., by age) can mislead models; confounders must be measured and adjusted for.

6.14 One-minute revision

Matrix: rows = actual, columns = predicted. TP/FP/TN/FN → sensitivity = TP/(TP+FN), specificity = TN/(TN+FP), precision = TP/(TP+FP), accuracy = (TP+TN)/total. Lower threshold → sensitivity up, specificity down. Bias = tilted data collection; confounding = hidden third factor (adjust or stratify).

Questions

Short-answer questions:

1 Draw a confusion matrix and label all four boxes with a one-line meaning each.

2 Define sensitivity and specificity with formulas.

3 What is a confounder? Give one pharmacy example.

4 Why is a false negative more dangerous than a false positive in ADR screening?

Long-answer questions:

1 An ADR model tested on 40 patients gives TP = 14, FN = 6, FP = 4, TN = 16. Compute accuracy, sensitivity, specificity, and precision showing all steps, and advise whether this model is fit for screening a dangerous ADR.

2 Explain reporting bias and confounding in observational pharmacovigilance data, with one example of each and one method to handle each.

MCQs (answer key at the end of the chapter):

1 In a confusion matrix, a false negative means:

a) Model predicted yes, actually yes b) Model predicted no, actually yes c) Model predicted yes, actually no d) Model predicted no, actually no

1 Sensitivity is calculated as:

a) TN ÷ (TN + FP) b) TP ÷ (TP + FN) c) TP ÷ (TP + FP) d) (TP + TN) ÷ total

1 Lowering the decision threshold of a model will generally:

a) Increase sensitivity and decrease specificity b) Decrease sensitivity and increase specificity c) Increase both sensitivity and specificity d) Leave both unchanged

1 Age acting as a link between "number of medicines" and "ADR risk" is an example of:

a) Reporting bias b) Confounding c) Measurement bias d) Recall bias

1 Precision answers the question:

a) Of all real cases, how many did we catch? b) When the alarm rings, how often is it real? c) Of all safe patients, how many did we clear? d) Overall, how often were we right?

Answer key: 1-b, 2-b, 3-a, 4-b, 5-b

Viva questions:

1 What are the rows and columns of a confusion matrix?

2 Why do we need four metrics instead of just accuracy?

3 Give one example of reporting bias in ADR data.

4 How is confounding handled in a logistic regression model?

Application-based questions:

1 Your hospital wants an AI alert for a rare but fatal ADR. Would you set the threshold high or low? Justify using sensitivity and precision.

2 A model trained on spontaneous ADR reports shows Drug Y has triple the ADR rate. List two biases or confounders you would check before warning doctors, and how you would check each.

References

- Aggarwal, C. C. and Reddy, C. K. Healthcare Data Analytics.
- Campbell, M. J., Machin, D. and Walters, S. J. Medical Statistics: A Textbook for the Health Sciences.
- Steyerberg, E. W. Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating.
- Strom, B. L., Kimmel, S. E. and Hennessy, S. Pharmacoepidemiology.