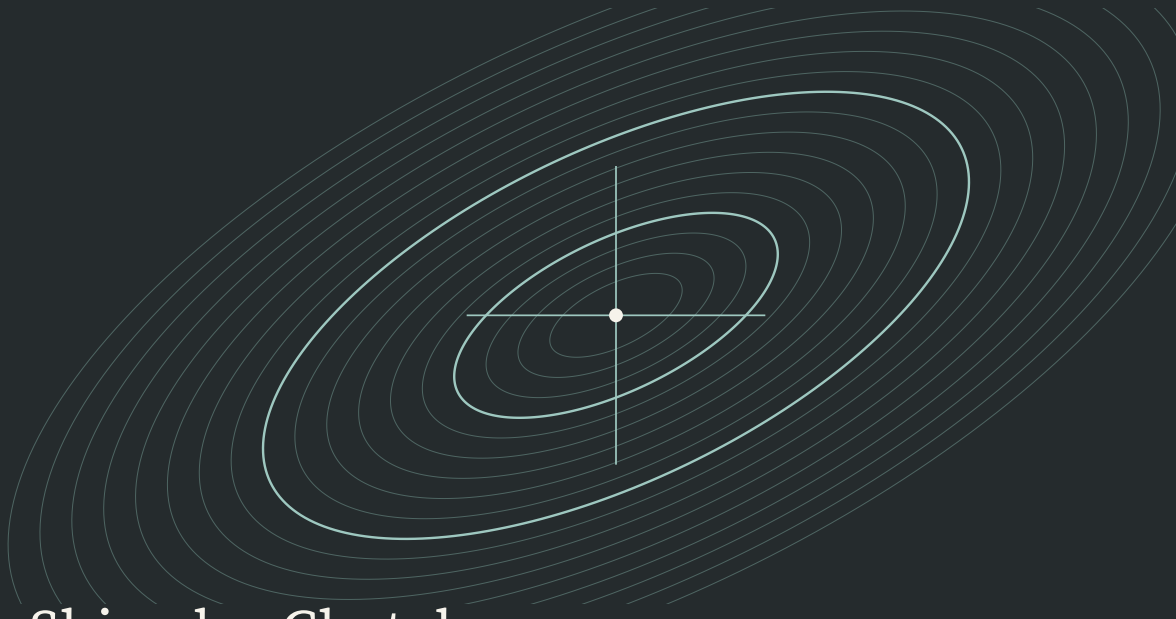
INNOWORKS PRESS

CONCEPTS / METHODS / WORKED PROBLEMS

Before the Algorithm

Mathematics and Engineering for Intelligent Systems



Akash Shivadas Chatake

19 CHAPTERS

128 WORKED PROBLEMS

WORKING STUDY EDITION / 2026

INNOWORKS PRESS

Before the Algorithm

Mathematics and Engineering for Intelligent Systems

A concise companion to concepts and worked problems

Akash Shivadas Chatake

Working study edition • November 2026

Before the Algorithm
Mathematics and Engineering for Intelligent Systems

Akash Shivadas Chatake

Published by InnoWorks Press, an imprint of Chatake Innoworks Private Limited.

Copyright © 2026 Akash Shivadas Chatake. All rights reserved.

Working study edition, November 2026. No ISBN has been assigned.

This independent educational text is not an official university publication, an endorsed entrance guide or a collection of previous-year examination questions. Its original study scope was informed by the supplied September 2026 IIT Bombay C-MInDS entrance syllabus; the scope map is retained at the end. Engineering coverage is introductory and selective.

This working edition is open for reading and personal study through the publisher's website. That access does not grant a general republication or redistribution licence. Bibliographic references identify further reading; those books have not been reproduced here.

Editorial status: author and technical review remain in progress. AI-assisted drafting, explanation and production tools were used in preparing the study material. Twenty-six representative symbolic and numerical checks have been run; they are not exhaustive verification or external peer review.

Publisher's edition and updates:
https://innoworkspress.org/books/before-the-algorithm/

Before the algorithm

A shadow moves across a wall. Water gathers in a vessel. A warm cup slowly gives its heat to a room. Long before we write an equation, the world offers us relationships: between position and time, flow and accumulation, difference and change. Mathematics begins when we ask which of these relationships we can describe clearly enough to examine.

There is something quietly remarkable in this. A few symbols can connect events that seem to belong to different worlds. The derivative that describes a moving body also describes how a prediction responds to a parameter. The projection of one vector onto a space becomes a way of fitting imperfect observations. A balance of what enters and leaves gives order to a circuit, a reaction or a thermal process. The connections do not make these things identical. They let an idea learned in one place become useful in another.

An algorithm arrives later. Before its steps can mean anything, we must decide what the quantities represent, what counts as an answer, and which assumptions permit the calculation. A machine may carry out a million operations without hesitation. Choosing the right question still deserves a moment of human hesitation.

This small book is meant for that moment. It brings together foundations that are easily scattered across courses and easily forgotten between examinations. You can enter at a particular problem, follow its method and check its result. The short openings offer a little perspective before the arithmetic begins. They are invitations to understand, not obstacles to a quick revision.

Read with a pencil. Ask what is changing, what is held fixed and what would happen at an extreme. Notice the units. Distrust an answer that becomes more impressive as its meaning becomes less clear. When an idea seems difficult, look for a simpler case that still contains its structure.

We do not need to know everything before we begin. We need a few dependable ideas, a willingness to return to them, and the patience to discover how much they connect.

How to use this companion

For a quick glance, read a chapter's opening and definitions, then choose a worked problem. Each problem keeps its explanation, calculation and check together. Cover the calculation when you want to practise; uncover it when you want to diagnose a missed step.

For a longer revision, read Chapters 1–3 for calculus, 4–6 for linear algebra and 7–9 for probability and statistics. Chapters 10–11 cover algorithms; 12–13 connect the foundations to machine learning. Chapters 14–16 introduce selected engineering principles. Chapters 17–19 provide mixed problems, figure reading and algebraic tools.

The aim is usable understanding within a compact format. This companion does not replace a full course in any of its six fields. Its 128 worked problems are independently prepared study examples. No official examination weighting, question prediction, university endorsement or admission outcome is implied.

Notation and algebra before you begin

A variable is a symbol whose value may vary. A parameter is held fixed while a calculation is performed, although it may be estimated from data later. In $y = mx + b$, x is the input, y the output, and m, b specify the line. Parentheses control order: $(a+b)^2$ means $(a+b)(a+b)$, so

$$(a+b)^2 = a^2 + 2ab + b^2, \qquad (a-b)^2 = a^2 - 2ab + b^2.$$

The symbol $\Rightarrow$ means "therefore" or "implies". An equality sign asserts that the expressions on both sides have the same value. Never use an equality sign merely to mean "and then".

A fraction a/b requires $b \neq 0$. Multiplying both sides of an equation by the same nonzero number preserves its solutions. Factoring can be checked by multiplication. For example,

$$x^2 - 5x + 6 = (x-2)(x-3).$$

Thus $x^2 - 5x + 6 = 0$ implies $x = 2$ or $x = 3$. The zero-product rule works because a product of real numbers is zero only when at least one factor is zero.

An exponent denotes repeated multiplication when it is a positive integer. Thus $x^3 = xxx$, while $x^0 = 1$ for $x \neq 0$ and $x^{-1} = 1/x$. A square root is nonnegative: $\sqrt{x^2} = |x|$. The natural logarithm is the inverse of the exponential: $\ln(e^x) = x$ and $e^{\ln x} = x$ for $x > 0$. Useful identities are

$$\ln(ab) = \ln a + \ln b, \qquad \ln(a/b) = \ln a - \ln b, \qquad \ln(a^r) = r \ln a$$

when the arguments are positive.

Vectors will usually be columns; x^T is the transpose. The symbol $\sum$ adds a collection of terms: $\sum_{i=1}^{3} x_i = x_1 + x_2 + x_3$. A derivative measures local change; an integral accumulates contributions. Both are developed rather than assumed in the first chapters.

Contents

Before the algorithm ii

How to use this companion iii

Notation and algebra before you begin iv

1 Functions, limits and differentiation 1

1.1 What a function tells you . . . 1
1.2 Limits and the derivative . . . 2
1.3 Rules and how to recognise them . . . 3
1.4 Stationary points, curvature and approximation . . . 4
1.5 When a method fails . . . 5
1.6 A boundary and a nondifferentiable point . . . 5

2 Integration and accumulated quantities 7

2.1 From rectangles to an integral . . . 7
2.2 Substitution as the reverse chain rule . . . 8
2.3 Integration by parts . . . 9
2.4 Double integrals and the region . . . 9
2.5 Polar coordinates and the Jacobian . . . 10
2.6 An improper integral and its convergence . . . 11

3 Multivariable calculus and optimisation 12

3.1 Holding one input fixed . . . 12
3.2 When every variable changes . . . 13
3.3 The Hessian and local shape . . . 14
3.4 Constraints and feasible directions . . . 15
3.5 Connection to learning . . . 16

4 Vectors, linear systems and spaces **17**

4.1 Vectors as objects and coordinates . . . 17
4.2 Solving simultaneous equations . . . 18
4.3 Span, independence and basis . . . 19
4.4 Rank, nullity and determinant . . . 20
4.5 Why dependence removes uniqueness . . . 21

5 Orthogonality, projections and least squares **22**

5.1 Angles and distances . . . 22
5.2 Projection derived, not memorised . . . 23
5.3 Least squares as a projection . . . 24
5.4 Orthonormal coordinates and numerical conditioning . . . 25

6 Eigenvalues, quadratic forms and curvature **27**

6.1 Directions preserved by a transformation . . . 27
6.2 Diagonalisation and powers . . . 28
6.3 Quadratic forms and positive definiteness . . . 29
6.4 Condition numbers and sensitivity . . . 30
6.5 Why this belongs in machine learning . . . 30

7 Probability, counting and conditioning **31**

7.1 Sample spaces and events . . . 31
7.2 Counting arrangements . . . 32
7.3 Conditional probability . . . 32
7.4 Total probability and Bayes' rule . . . 33
7.5 Complements and repeated trials . . . 34

8 Random variables and distributions **36**

8.1 A numerical description of uncertainty . . . 36
8.2 Expectation, variance and transformation . . . 37
8.3 Bernoulli and binomial models . . . 37
8.4 Continuous uniform and exponential distributions . . . 38
8.5 Normal standardisation . . . 39
8.6 Covariance and dependence . . . 40
8.7 Uncorrelated does not mean independent . . . 40

9 Sampling, estimation and regression **42**

9.1 A parameter is not an estimate . . . 42
9.2 Bias, variance and mean squared error . . . 43
9.3 Sample variance and degrees of freedom . . . 44
9.4 Maximum likelihood from first principles . . . 44
9.5 Simple linear regression . . . 45
9.6 Intervals and interpretation . . . 46

10 Algorithmic growth and recurrences **47**

10.1 What is being counted? . . . 47
10.2 Logarithms and repeated halving . . . 48
10.3 Nested loops and sums . . . 49
10.4 Divide and conquer . . . 50
10.5 Space and recursion depth . . . 51

11 Searching, sorting, heaps and trees **52**

11.1 Searching and its invariant . . . 52
11.2 Insertion sort . . . 53
11.3 Merge sort and merging . . . 54
11.4 Quicksort and bad partitions . . . 54
11.5 Binary search trees and heaps . . . 55

12 Learning tasks, losses and fitted models **56**

12.1 What is learned? . . . 56
12.2 Squared and absolute losses . . . 57
12.3 Empirical risk and fitting . . . 57
12.4 Binary classification and the sigmoid . . . 58
12.5 Overfitting and underfitting . . . 59
12.6 Data preparation is part of the model . . . 60
12.7 Derive the sigmoid derivative . . . 60

13 Optimisation, regularisation and evaluation **62**

13.1 Why gradient descent moves downhill . . . 62
13.2 A stability condition for a quadratic . . . 63
13.3 Regularisation changes the objective . . . 64
13.4 Confusion counts and metrics . . . 65

13.5 Time-series splitting and uncertainty 65
13.6 Batch and stochastic gradients 66

14 Mechanics, structures and materials **68**
14.1 Units are part of the calculation 68
14.2 Force and acceleration 68
14.3 Moments and support reactions 69
14.4 Stress, strain and elastic response 70
14.5 Work, kinetic energy and power 71
14.6 Hydrostatic pressure 72
14.7 Static friction as a bound 72

15 Electrical circuits and computing foundations **73**
15.1 Voltage, current, resistance and power 73
15.2 Series and parallel connections 74
15.3 Kirchhoff's laws 75
15.4 Energy units and efficiency 76
15.5 Bits, binary numbers and Boolean reasoning 76

16 Thermodynamics and process balances **78**
16.1 Systems, boundaries and accounting 78
16.2 Heat, work and internal energy 79
16.3 Ideal gas calculations 80
16.4 Mass balances and mixtures 81
16.5 Residence time and rate 81
16.6 Model assumptions and physical interpretation 82

17 Mixed problems with complete solutions **83**
17.1 Calculus and linear algebra together 83
17.2 Probability, estimation and prediction 85
17.3 Algorithm and learning calculations 87
17.4 Engineering and modelling together 89
17.5 How to review a completed solution 90

18 Reading mathematical figures **92**
18.1 Slope is different from height 93

18.2 Area and changing bounds 94
18.3 Contours and gradient direction 95
18.4 Projection and residual geometry 96
18.5 Probability is mass or area 97
18.6 Growth and optimisation plots 98
18.7 Scores, flow and physical dimensions 100

19 Algebra and numerical tools for self-study **102**
19.1 Fractions and equations 102
19.2 Quadratics and factorisation 103
19.3 Exponentials and logarithms 104
19.4 Scientific notation and units 104
19.5 Ratios, percentages and changes 105
19.6 Rounding and exact answers 105
19.7 Normalising a vector 106

Syllabus map and reading directions **107**

Chapter 1

Functions, limits and differentiation

The idea before the calculation

A falling stone does not carry a formula. We introduce one to describe a relation between quantities: time and distance, temperature and output, an input and a prediction. A function is this act of description made precise. Its domain says where the description is allowed to speak.

A derivative asks a narrower question: what does change look like very close to here? A curved path becomes approximately straight when examined locally. The approximation is powerful precisely because its reach is limited. A slope measured at one point need not describe the whole journey. In the limit defining a derivative, the increment approaches zero; we never divide by an increment already set equal to zero.

1.1 What a function tells you

A function assigns one output to each allowed input. If $f(x) = x^2$, the input 3 produces 9, and the input -3 also produces 9. Different inputs may share an output; a single allowed input cannot have two different outputs under the same function. The domain is the collection of allowed inputs. For $1/x$, zero is excluded. For $\ln x$, real-valued inputs must be positive.

The graph consists of points $(x, f(x))$. Its height gives a value; its slope tells us how fast that value is changing. These are different quantities. At the bottom of a parabola, its height need not be zero, but its slope is zero. In a prediction model, a function might map temperature to estimated gas output. Its derivative would express the local output sensitivity to temperature, not the output itself.

1.2 Limits and the derivative

A limit describes nearby behaviour. The expression $(x^2-1)/(x-1)$ is undefined at $x=1$, but factoring its numerator reveals what happens near that point:

$$\frac{x^2-1}{x-1} = \frac{(x-1)(x+1)}{x-1} = x+1 \quad (x \neq 1).$$

Consequently its limit as $x \to 1$ is 2. Cancelling a factor does not retroactively define the original expression at the excluded point.

The derivative is a limit of average slopes:

$$f'(x) = \lim_{h \to 0} \frac{f(x+h)-f(x)}{h}. \tag{1.1}$$

The numerator is the change in output; the denominator is the change in input. For $f(x)=x^2$, expanding before cancelling gives

$$\begin{aligned} \frac{f(x+h)-f(x)}{h} &= \frac{x^2+2xh+h^2-x^2}{h} \\ &= 2x+h \longrightarrow 2x. \end{aligned}$$

Thus $f'(x)=2x$. At $x=3$ the slope is 6; at $x=-3$ it is -6. Negative slope means the output decreases as the input increases locally.

Worked problem 1.1 A removable discontinuity

Problem. Find $\lim_{x \to 2}(x^2-4)/(x-2)$.

Understanding the method. Direct substitution gives $0/0$, which is an indeterminate form, not the answer. Factor the difference of squares so that a common factor can be cancelled away from $x=2$.

Calculation.

$$\begin{aligned} x^2-4 &= (x-2)(x+2), \\ \frac{x^2-4}{x-2} &= x+2 \qquad (x \neq 2), \\ \lim_{x \to 2} \frac{x^2-4}{x-2} &= \lim_{x \to 2}(x+2) = \boxed{4}. \end{aligned}$$

Check and meaning. At $x=1.99$ the simplified value is 3.99 and at 2.01 it is 4.01. Both approach 4. The original quotient remains undefined exactly at 2.

1.3 Rules and how to recognise them

The power rule is $\frac{d}{dx}x^n = nx^{n-1}$ on domains where the expression and derivative exist. Constants differentiate to zero. A constant multiplying a function stays outside the derivative. Sums differentiate term by term. The exponential and logarithm satisfy

$$\frac{d}{dx}e^x = e^x, \qquad \frac{d}{dx}\ln x = \frac{1}{x} \quad (x > 0).$$

A product needs two terms, because either factor can change:

$$(uv)' = u'v + uv'.$$

A quotient uses $(u/v)' = (u'v - uv')/v^2$ where $v \neq 0$. A composition is a function inside another function. Its derivative is the outer derivative evaluated at the inner function, multiplied by the inner derivative:

$$\frac{d}{dx}f(g(x)) = f'(g(x))g'(x).$$

Recognise the structure before calculating. In $\ln(1 + x^2)$ the outside is logarithm and the inside is $1 + x^2$. In x^2e^x there are two multiplied factors, so the product rule is needed.

Worked problem 1.2 Product and chain rules together

Problem. Differentiate $f(x) = x^2e^{3x}$ and calculate $f'(1)$.

Understanding the method. The outer structure is a product. The exponential factor itself contains the inner function $3x$, so its derivative needs the chain rule.

Calculation.

$$\begin{aligned} u &= x^2, & u' &= 2x, \\ v &= e^{3x}, & v' &= e^{3x}\frac{d(3x)}{dx} = 3e^{3x}, \\ f'(x) &= (2x)e^{3x} + x^2(3e^{3x}) \\ &= e^{3x}(2x + 3x^2), \\ f'(1) &= e^3(2 + 3) = \boxed{5e^3}. \end{aligned}$$

Check and meaning. Both contributions contain e^{3x}. Factoring it is optional, but losing the multiplier 3 from the chain rule changes the answer.

Worked problem 1.3 A logarithmic composition

Problem. Find the derivative of $g(x) = \ln(1 + e^x)$ and evaluate it at zero.

Understanding the method. The derivative of $\ln u$ is u'/u. The inner expression $u = 1+e^x$ is positive for every real x, so the logarithm is defined throughout the real line.

Calculation.

$$\begin{aligned} u &= 1 + e^x, & u' &= e^x, \\ g'(x) &= \frac{e^x}{1+e^x} = \frac{1}{1+e^{-x}}, \\ g'(0) &= \frac{1}{1+1} = \boxed{\tfrac{1}{2}}. \end{aligned}$$

Check and meaning. The derivative lies between zero and one. The same function appears later as the sigmoid in binary classification.

1.4 Stationary points, curvature and approximation

A stationary point satisfies $f'(x) = 0$. It is a candidate for an interior extremum, but zero slope alone does not prove a minimum: x^3 has zero derivative at zero and no local maximum or minimum there. If $f''(a) > 0$ at a stationary point, nearby slopes increase through zero and the point is a strict local minimum. If $f''(a) < 0$, it is a strict local maximum. If the second derivative is zero, use further reasoning.

For a differentiable function, a small input change h has approximate effect $f(x+h) \approx f(x) + hf'(x)$. Taylor expansion adds curvature:

$$f(x+h) = f(x) + hf'(x) + \tfrac{1}{2}h^2 f''(x) + \cdots .$$

The approximation improves as h becomes small under suitable smoothness conditions. It is not an exact identity unless higher terms vanish.

Worked problem 1.4 A quadratic minimum from two viewpoints

Problem. Find the minimum of $f(x) = x^2 - 6x + 5$ over all real x.

Understanding the method. Differentiate to find an interior stationary point, then classify it. Completing the square gives an independent global check.

Calculation.

$$\begin{aligned} f'(x) &= 2x - 6 = 0 \quad \Rightarrow \quad x = 3, \\ f''(x) &= 2 > 0, \\ f(3) &= 9 - 18 + 5 = \boxed{-4}. \end{aligned}$$

Moreover,

$$f(x) = x^2 - 6x + 9 - 4 = (x-3)^2 - 4 \geq -4.$$

Equality holds exactly at $x = 3$, so the minimum is global.

Check and meaning. The squared term cannot be negative. This establishes more than the local second-derivative test alone.

Worked problem 1.5 A tangent-line estimate

Problem. Estimate $\sqrt{4.1}$ by linear approximation at $x = 4$.

Understanding the method. Use a point where the function is easy to evaluate. With $f(x) = \sqrt{x}$, the derivative is $f'(x) = 1/(2\sqrt{x})$.

Calculation.

$$f(4) = 2, \qquad f'(4) = \frac{1}{4}, \qquad h = 0.1,$$

$$\begin{aligned} f(4.1) &\approx f(4) + hf'(4) \\ &= 2 + (0.1)(0.25) = \boxed{2.025}. \end{aligned}$$

Check and meaning. Squaring 2.025 gives 4.100625, close to 4.1. Because $\sqrt{x}$ is concave, its tangent lies above the graph, so a slight overestimate is expected.

1.5 When a method fails

For optimisation on a closed interval, check endpoints as well as stationary points. A derivative can fail to exist at a corner; $|x|$ has a minimum at zero despite no derivative there. Logarithmic expressions require positive arguments. Quotient rules require nonzero denominators. Recognising these restrictions is part of solving the problem, not a separate formality.

For $f(x) = x^2$ on $[1, 2]$, the unconstrained stationary point 0 is outside the interval. The values $f(1) = 1$ and $f(2) = 4$ show that the interval minimum is at 1. A memorised rule "set derivative to zero" without a domain check would miss the answer.

1.6 A boundary and a nondifferentiable point

Differentiation identifies local smooth behaviour. It does not remove the need to read the domain. At a boundary, movement may be allowed only in one direction; at a corner, a tangent slope may not exist. Both situations can still produce a minimum.

Worked problem 1.6 An absolute-value minimum

Problem. Find the minimum of $f(x) = |x - 2| + 1$. Does setting $f'(x) = 0$ find it?

Understanding the method. Absolute value is distance. Its piecewise definition exposes the slopes on either side of the corner.

Calculation.

$$f(x) = \begin{cases} 3 - x, & x < 2, \\ x - 1, & x \geq 2. \end{cases}$$

For $x < 2$, the slope is -1; for $x > 2$, the slope is 1. Neither side has zero derivative. At 2, the left and right derivatives differ, so $f'(2)$ does not exist. Nevertheless $|x - 2| \geq 0$ gives

$$\boxed{f_{\min} = 1 \text{ at } x = 2}.$$

Check and meaning. The function falls toward the corner and rises after it. The answer follows from distance or a sign change in one-sided slopes, not from a zero derivative.

Chapter 2

Integration and accumulated quantities

The idea before the calculation

A river's flow rate tells us what passes now; an integral tells us what has passed over an interval. Accumulation turns a succession of local contributions into a total. The same idea joins velocity to displacement, power to energy, and a probability density to a probability.

Units help reveal the meaning: litres per second, multiplied by seconds and accumulated, become litres. Signed contributions can cancel, so an integral of velocity gives displacement rather than necessarily the distance travelled. The fundamental theorem of calculus connects this accounting of accumulated change with differentiation. The constant in an indefinite integral remembers something a rate cannot tell us: where the accumulation began.

2.1 From rectangles to an integral

An integral sums many small contributions. If a flow rate $q(t)$ is measured in litres per minute, then $q(t)\,\Delta t$ approximates volume transferred over a short time interval. Adding such products and taking finer intervals leads to $\int_a^b q(t)\,\mathrm{d}t$. The units are rate times time, hence litres. This interpretation is more useful than treating integration as a collection of symbol rules.

The definite integral is signed: portions below the horizontal axis subtract from portions above it. Geometric area requires integrating absolute height or splitting at sign changes. The indefinite integral denotes a family of antiderivatives; its arbitrary constant disappears when definite endpoint differences are taken.

The fundamental theorem gives

$$\int_a^b f(x)\,\mathrm{d}x = F(b) - F(a), \qquad F' = f.$$

This turns accumulation into an endpoint calculation. Always evaluate the upper endpoint minus the lower endpoint with parentheses intact.

Worked problem 2.1 Polynomial integration with endpoints

Problem. Evaluate $\int_0^2 (3x^2 + 2x)\,\mathrm{d}x$.

Understanding the method. Integrate a sum term by term. The antiderivative of x^n is $x^{n+1}/(n+1)$ for $n \neq -1$.

Calculation.

$$\begin{aligned}\int (3x^2 + 2x)\,\mathrm{d}x &= 3\frac{x^3}{3} + 2\frac{x^2}{2} + C = x^3 + x^2 + C,\\ \int_0^2 (3x^2 + 2x)\,\mathrm{d}x &= [x^3 + x^2]_0^2\\ &= (8 + 4) - (0 + 0) = \boxed{12}.\end{aligned}$$

Check and meaning. The integrand is nonnegative throughout the interval, so a positive integral is expected. Differentiating $x^3 + x^2$ recovers $3x^2 + 2x$.

2.2 Substitution as the reverse chain rule

When an integrand contains a function and a multiple of its derivative, try substituting the function as a new variable. The differential must be transformed too. For a definite integral the bounds belong to the integration variable: if the variable changes to u, change the bounds to values of u.

This is why $\int 2x/(1 + x^2)\,\mathrm{d}x$ becomes $\int (1/u)\,\mathrm{d}u$ when $u = 1 + x^2$. It does not become $\int 2x/u\,\mathrm{d}u$, because $2x\,\mathrm{d}x$ has already been replaced by $\mathrm{d}u$.

Worked problem 2.2 An exponential substitution

Problem. Evaluate $I = \int_0^1 x e^{x^2}\,\mathrm{d}x$.

Understanding the method. The inner exponent x^2 has derivative $2x$. The integrand has x instead of $2x$, so a factor $1/2$ must be retained.

Calculation.

$$\begin{aligned}u = x^2, \qquad & \mathrm{d}u = 2x\,\mathrm{d}x, \qquad && x\,\mathrm{d}x = \tfrac{1}{2}\,\mathrm{d}u,\\ x = 0 \Rightarrow u = 0, \qquad & x = 1 \Rightarrow u = 1,\end{aligned}$$

$$\begin{aligned}I &= \frac{1}{2}\int_0^1 e^u\,\mathrm{d}u = \frac{1}{2}[e^u]_0^1\\ &= \boxed{\frac{e - 1}{2}} \approx 0.8591.\end{aligned}$$

Check and meaning. The antiderivative $\frac{1}{2}e^{x^2}$ differentiates to xe^{x^2}, confirming the factor.

2.3 Integration by parts

The product rule rearranges into the integration-by-parts identity

$$\int u\,\mathrm{d}v = uv - \int v\,\mathrm{d}u.$$

It is useful when differentiating one factor makes it simpler. For a polynomial times an exponential, differentiating the polynomial lowers its degree while integrating the exponential is straightforward. A good choice reduces the complexity of the remaining integral.

Worked problem 2.3 A polynomial multiplied by an exponential

Problem. Evaluate $\int_0^1 xe^x\,\mathrm{d}x$.

Understanding the method. Choose $u = x$ and $\mathrm{d}v = e^x\,\mathrm{d}x$, giving $\mathrm{d}u = \mathrm{d}x$ and $v = e^x$.

Calculation.

$$\begin{aligned}\int_0^1 xe^x\,\mathrm{d}x &= [xe^x]_0^1 - \int_0^1 e^x\,\mathrm{d}x \\ &= (e-0) - [e^x]_0^1 \\ &= e - (e-1) = \boxed{1}.\end{aligned}$$

Check and meaning. The primitive is $xe^x - e^x = (x-1)e^x$. Differentiating it gives xe^x. The cancellation of the two e terms is genuine.

2.4 Double integrals and the region

A double integral sums over an area. In an iterated integral, the inner variable changes while the outer one is held constant. For a rectangular region, both bounds can be constants. For a triangular or curved region, an inner bound may depend on the outer variable. Sketching the region avoids integrating over the wrong area.

Worked problem 2.4 A rectangular region

Problem. Compute $\int_0^1 \int_0^2 (x+y)\,\mathrm{d}y\,\mathrm{d}x$.

Understanding the method. The order $\mathrm{d}y\,\mathrm{d}x$ means integrate in y first. During that inner integration, x acts like a constant.

Calculation.

$$\int_0^2 (x+y)\,\mathrm{d}y = [xy + y^2/2]_0^2 = 2x + 2,$$
$$\int_0^1 (2x+2)\,\mathrm{d}x = [x^2 + 2x]_0^1 = \boxed{3}.$$

Check and meaning. The area is 2. The average value of $x+y$ on this rectangle is $0.5 + 1 = 1.5$, so area times average is $2(1.5) = 3$.

Worked problem 2.5 A triangular region

Problem. Integrate $x + y$ over $0 \le x \le 1$, $0 \le y \le 1 - x$.

Understanding the method. For each fixed x, the vertical interval stops at the line $y = 1 - x$. Replacing that bound by 1 would incorrectly use a square.

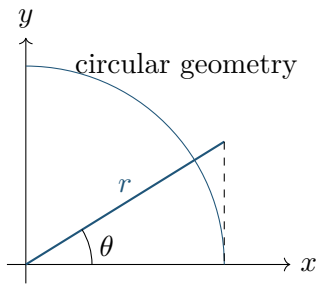
Calculation.

$$\begin{aligned} I &= \int_0^1 \left[xy + \frac{y^2}{2} \right]_{y=0}^{1-x} \mathrm{d}x \\ &= \int_0^1 \left[x(1-x) + \frac{(1-x)^2}{2} \right] \mathrm{d}x \\ &= \int_0^1 \left[\frac{1}{2} - \frac{x^2}{2} \right] \mathrm{d}x \\ &= \left[\frac{x}{2} - \frac{x^3}{6} \right]_0^1 = \boxed{\frac{1}{3}}. \end{aligned}$$

Check and meaning. The triangle has area $1/2$ and centroid $(1/3, 1/3)$. A linear function averages to its value at the centroid, so $(1/2)(2/3) = 1/3$.

2.5 Polar coordinates and the Jacobian

For circles, use $x = r\cos\theta$, $y = r\sin\theta$. A tiny polar region has radial thickness $\mathrm{d}r$ and arc length approximately $r\,\mathrm{d}\theta$, so its area is $r\,\mathrm{d}r\,\mathrm{d}\theta$. This explains the Jacobian factor geometrically. It cannot be omitted.



Worked problem 2.6 A radial integral over a disk

Problem. Evaluate $\iint_D (x^2 + y^2)\,\mathrm{d}A$ over the disk $x^2 + y^2 \le 4$.

Understanding the method. The boundary is a circle of radius 2, and the integrand becomes r^2. Multiplying by the area element supplies another factor r.

Calculation.

$$\begin{aligned} I &= \int_0^{2\pi} \int_0^2 r^2 (r\,\mathrm{d}r\,\mathrm{d}\theta) \\ &= \int_0^{2\pi} \left[\frac{r^4}{4}\right]_0^2 \mathrm{d}\theta \\ &= \int_0^{2\pi} 4\,\mathrm{d}\theta = \boxed{8\pi}. \end{aligned}$$

Check and meaning. The disk area is 4π, and the mean squared radius is 2, giving 8π. Using $\mathrm{d}r\,\mathrm{d}\theta$ without r would describe the wrong measure.

2.6 An improper integral and its convergence

An infinite integration bound is handled through a limit. A formal antiderivative is not enough: the endpoint limit must exist as a finite number for convergence.

Worked problem 2.7 A decaying exponential over an infinite interval

Problem. For $\lambda > 0$, evaluate $\int_0^\infty \lambda e^{-\lambda x}\,\mathrm{d}x$.

Understanding the method. First integrate to a finite upper bound b, then take $b \to \infty$. Positive λ makes the exponential decay.

Calculation.

$$\begin{aligned} \int_0^b \lambda e^{-\lambda x}\,\mathrm{d}x &= [-e^{-\lambda x}]_0^b = 1 - e^{-\lambda b}, \\ \int_0^\infty \lambda e^{-\lambda x}\,\mathrm{d}x &= \lim_{b\to\infty} (1 - e^{-\lambda b}) = \boxed{1}. \end{aligned}$$

Check and meaning. This calculation verifies that the exponential density in Chapter 8 integrates to one. The assumption $\lambda > 0$ is essential; it ensures decay and a nonnegative density.

Chapter 3

Multivariable calculus and optimisation

The idea before the calculation

Most interesting systems have several ways to change. A reaction depends on temperature and concentration; a fitted model depends on many parameters. A partial derivative asks what happens along one coordinate while the others stay fixed. The gradient collects these local sensitivities into a vector.

Picture a landscape, but treat the picture carefully. A stationary point may be a valley, a hilltop or a saddle. Zero gradient alone cannot distinguish them. Moreover, "steepest" depends on how distance is measured: rescaling a coordinate changes the geometry. Before optimising, identify the variables, their units and any constraints. The mathematics searches within the problem we formulate; it cannot choose the right formulation for us.

3.1 Holding one input fixed

A function $f(x, y)$ may depend on two inputs, such as temperature and pressure. A partial derivative changes only one input while holding the other fixed. For

$$f(x, y) = x^2 y + 3y^2,$$

we get $f_x = 2xy$ because y is constant during differentiation in x. Differentiating in y instead gives $f_y = x^2 + 6y$. These expressions measure different sensitivities.

The gradient collects partial derivatives into a vector:

$$\nabla f = \begin{pmatrix} f_x \\ f_y \end{pmatrix}.$$

For a unit direction u, the directional derivative is $\nabla f \cdot u$. The requirement that u have length one matters: otherwise its magnitude changes the answer. The gradient points toward the direction of greatest first-order increase when it is nonzero.

Worked problem 3.1 A gradient and a directional derivative

Problem. For $f = x^2y + 3y^2$, find the gradient at $(1, 2)$ and the derivative in direction $(3, 4)$.

Understanding the method. Calculate the gradient first. A direction vector must be normalised before taking its inner product with the gradient.

Calculation.

$$\begin{aligned} f_x &= 2xy, & f_y &= x^2 + 6y, \\ \nabla f(1,2) &= (4,13)^T, & & \\ \|(3,4)\| &= \sqrt{9+16} = 5, & u &= (3/5, 4/5)^T, \\ D_u f(1,2) &= (4,13)\cdot(3/5,4/5) & & \\ &= 12/5 + 52/5 = \boxed{64/5}. & & \end{aligned}$$

Check and meaning. Using $(3, 4)$ directly would produce 64, which is five times the rate per unit distance requested.

3.2 When every variable changes

If x and y themselves depend on t, changes in both affect f. The total derivative is

$$\frac{df}{dt} = f_x \frac{dx}{dt} + f_y \frac{dy}{dt}.$$

This is the multivariable chain rule. Its terms represent the separate pathways through which t influences the output. Substituting the parameterisation into f first is often an independent way to check it.

Worked problem 3.2 Two pathways through a function

Problem. Let $z = x^2 + y^2$, $x = t$, $y = 2t$. Find dz/dt at $t = 1$.

Understanding the method. Both x and y change, so both chain-rule terms are needed.

Calculation.

$$\begin{aligned} z_x &= 2x, & z_y &= 2y, & x' &= 1, & y' &= 2, \\ \frac{dz}{dt} &= 2x(1) + 2y(2). & & & & & \end{aligned}$$

At $t = 1$, $(x, y) = (1, 2)$, hence

$$\frac{dz}{dt} = 2(1) + 2(2)(2) = \boxed{10}.$$

Direct substitution gives $z = t^2 + (2t)^2 = 5t^2$, whose derivative $10t$ also equals 10 at $t = 1$.

3.3 The Hessian and local shape

The Hessian is the matrix of second partial derivatives:

$$H = \begin{pmatrix} f_{xx} & f_{xy} \\ f_{yx} & f_{yy} \end{pmatrix}.$$

For sufficiently smooth functions the mixed partials agree. Around a stationary point a, the leading change is a quadratic form, approximately $\frac{1}{2}h^T H(a)h$. If it is positive in every nonzero direction, the stationary point is a strict local minimum. If negative in every direction, it is a local maximum. If it takes both signs, the point is a saddle.

For a symmetric two-by-two Hessian, $f_{xx} > 0$ and $\det H > 0$ imply positive definiteness. A negative determinant implies an indefinite form and hence a saddle at a stationary point. A zero determinant requires more analysis; the test is inconclusive.

Worked problem 3.3 A quadratic minimum with a mixed term

Problem. Find and classify the stationary point of $f = x^2 + 2xy + 3y^2 - 4x - 8y$.

Understanding the method. Set both first partial derivatives to zero, solve the resulting linear system, then examine the Hessian.

Calculation.

$$\begin{aligned} f_x &= 2x + 2y - 4 = 0 && \Rightarrow x + y = 2, \\ f_y &= 2x + 6y - 8 = 0 && \Rightarrow x + 3y = 4. \end{aligned}$$

Subtracting gives $2y = 2$, so $y = 1$ and $x = 1$. The Hessian is

$$H = \begin{pmatrix} 2 & 2 \\ 2 & 6 \end{pmatrix}, \qquad \det H = 12 - 4 = 8 > 0, \quad H_{11} = 2 > 0.$$

It is positive definite, so $(1, 1)$ is a strict local minimum. The value is

$$f(1, 1) = 1 + 2 + 3 - 4 - 8 = \boxed{-6}.$$

Check and meaning. Because the Hessian is constant and positive definite, this quadratic is strictly convex and this minimum is global.

Worked problem 3.4 Why zero gradient is not enough

Problem. Classify the stationary point of $f(x, y) = x^2 - y^2$.

Understanding the method. Examine changes along different directions. If some increase the function and others decrease it, the point cannot be a maximum or minimum.

Calculation.

$$\nabla f = (2x, -2y)^T = 0 \quad \Rightarrow \quad (x, y) = (0, 0).$$

Along $y = 0$, $f(x, 0) = x^2 > 0$ for $x \neq 0$. Along $x = 0$, $f(0, y) = -y^2 < 0$ for $y \neq 0$. Since $f(0, 0) = 0$, arbitrarily close points lie on both sides of this value. Thus the origin is a $\boxed{\text{saddle point}}$.

Equivalently $H = \operatorname{diag}(2, -2)$ has determinant $-4 < 0$.

3.4 Constraints and feasible directions

A constraint restricts which points may be considered. On $x + y = 2$, one may substitute $y = 2 - x$. Alternatively a Lagrange multiplier solves $\nabla f = \lambda \nabla g$ for an equality constraint $g = 0$, subject to regularity such as nonzero constraint gradient. The multiplier approach says that at a constrained optimum, first-order change along every allowed tangent direction vanishes.

Worked problem 3.5 The same constrained problem by two methods

Problem. Minimise $x^2 + y^2$ subject to $x + y = 2$.

Understanding the method. Substitution reduces the problem to one variable. Lagrange multipliers provide a check and illustrate a method that extends to several constraints.

Calculation.

With $y = 2 - x$,

$$\begin{aligned} f(x, 2 - x) &= x^2 + (2 - x)^2 \\ &= 2x^2 - 4x + 4 = 2(x - 1)^2 + 2. \end{aligned}$$

The squared term is smallest at $x = 1$, giving $y = 1$ and minimum $\boxed{2}$.

For multipliers, use $g = x + y - 2 = 0$. Then

$$(2x, 2y) = \lambda(1, 1) \quad \Rightarrow \quad x = y.$$

Combined with $x + y = 2$, this again yields $(1, 1)$.

Check and meaning. The square-completion argument proves global minimality over the whole constraint line; a multiplier equation alone only supplies candidates.

3.5 Connection to learning

A loss function depends on many parameters. Its gradient gives the local direction in which loss increases fastest; moving against that gradient is the principle of gradient descent. A Hessian describes curvature and helps explain why a learning rate that works in one direction may overshoot in another. These ideas are revisited after linear algebra supplies the language of positive definite matrices.

Chapter 4

Vectors, linear systems and spaces

The idea before the calculation

A vector lets us hold several quantities in one mathematical object. A matrix describes how those quantities combine. In a chemical balance, its columns can represent reactions and its rows the substances whose amounts change. In a linear system, each row records a condition the unknowns must satisfy.

The question is not always "what is the answer?" It may first be "is an answer possible, and is it unique?" Rank counts independent information. A null-space direction changes the unknowns without changing the observations. This is why more equations do not automatically mean more knowledge: some merely repeat what was already known. Solving a system is an investigation of structure before it is an exercise in arithmetic.

4.1 Vectors as objects and coordinates

A vector can represent a displacement, a collection of measurements, or model coefficients. A vector in $\mathbb{R}^n$ has n real components. Addition and scalar multiplication act componentwise. The zero vector has every component zero. The transpose changes a column into a row, enabling matrix products with compatible dimensions.

For $u = (u_1, \dots, u_n)^T$ and $v = (v_1, \dots, v_n)^T$, the dot product is $u^T v = \sum_i u_i v_i$, a scalar. The Euclidean norm is $\|u\| = \sqrt{u^T u}$. A matrix A with m rows and n columns maps a vector in $\mathbb{R}^n$ to one in $\mathbb{R}^m$. Thus Ax is a weighted combination of the columns of A; the weights are the entries of x.

Worked problem 4.1 Matrix multiplication as a combination of columns

Problem. For $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and $x = (2, -1)^T$, calculate Ax.

Understanding the method. Each output component is a row of A dotted with x. Equivalently, combine the columns using coefficients 2 and -1.

Calculation.

$$Ax = \begin{pmatrix} 1(2) + 2(-1) \\ 3(2) + 4(-1) \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$$

By columns,

$$2\begin{pmatrix} 1 \\ 3 \end{pmatrix} - \begin{pmatrix} 2 \\ 4 \end{pmatrix} = \boxed{\begin{pmatrix} 0 \\ 2 \end{pmatrix}}.$$

Check and meaning. The dimensions are $(2 \times 2)(2 \times 1) = 2 \times 1$. Multiplying corresponding entries without summing is a different operation.

4.2 Solving simultaneous equations

The system $Ax = b$ asks for input weights that produce b. Elimination subtracts multiples of one equation from another to remove variables. Swapping equations, multiplying an equation by a nonzero constant, and adding a multiple of another equation preserve the solution set.

A pivot is the leading nonzero entry used to eliminate entries below or above it. The number of pivots is the rank. A row $0 =$ nonzero means inconsistency; a free variable in a consistent system leads to infinitely many solutions. Square systems are not automatically uniquely solvable.

Worked problem 4.2 A unique solution by elimination

Problem. Solve $x + 2y = 5$ and $3x + 4y = 11$.

Understanding the method. Multiply the first equation by 3 so that subtracting it from the second eliminates x.

Calculation.

$$\begin{aligned} 3x + 6y &= 15, \\ 3x + 4y &= 11, \\ (3x + 6y) - (3x + 4y) &= 15 - 11, \\ 2y = 4 \quad &\Rightarrow \quad y = 2, \\ x + 2(2) = 5 \quad &\Rightarrow \quad x = 1. \end{aligned}$$

Thus $\boxed{(x, y) = (1, 2)}$. Substitution checks $3(1) + 4(2) = 11$.

Worked problem 4.3 Infinite solutions and inconsistency

Problem. Compare $x + 2y = 3$, $2x + 4y = 6$ with the system whose second right-hand side is 7.

Understanding the method. The second left-hand side is twice the first. Consistency requires the right-hand side to have the same relation.

Calculation.

For the first system, subtract twice the first equation from the second:

$$(2x + 4y) - 2(x + 2y) = 6 - 2(3) \quad \Rightarrow \quad 0 = 0.$$

There is only one independent constraint. Let $y = t$. Then

$$\boxed{(x, y) = (3 - 2t, t), \quad t \in \mathbb{R}}.$$

For the modified system,

$$(2x + 4y) - 2(x + 2y) = 7 - 6 \quad \Rightarrow \quad 0 = 1,$$

so there is $\boxed{\text{no solution}}$.

Check and meaning. The coefficient matrix has the same rank in both systems. The right-hand side determines whether the dependent equations are compatible.

4.3 Span, independence and basis

The span of vectors is the set of all their linear combinations. Vectors are linearly independent if the only combination yielding zero uses all zero coefficients. A basis is an independent collection spanning a space. Its number of vectors is the dimension.

A subspace must contain zero and be closed under addition and scalar multiplication. A line through the origin in $\mathbb{R}^2$ is a subspace; a shifted line such as $x + y = 2$ is not, because it excludes zero. The solution set of a homogeneous system $Ax = 0$ is always a subspace, called the null space.

The column space contains the possible outputs Ax. The equation $Ax = b$ is consistent precisely when b lies in that column space. This geometric viewpoint unifies the elimination cases: a target may be reachable uniquely, reachable by multiple input combinations, or unreachable.

Worked problem 4.4 Testing a proposed subspace

Problem. Is $S = \{(x, y) : y = 2x\}$ a subspace? Is $T = \{(x, y) : y = 2x + 1\}$?

Understanding the method. Test zero, addition and scalar multiplication. One failed requirement is sufficient to reject a subspace.

Calculation.

Every element of S has form $t(1,2)$. Zero occurs at $t=0$. If $u=a(1,2)$ and $v=b(1,2)$, then

$$u+v=(a+b)(1,2)\in S, \qquad cu=(ca)(1,2)\in S.$$

Thus $\boxed{S \text{ is a subspace}}$ with basis $(1,2)$ and dimension one.

For T, the origin would require $0=2(0)+1$, which is false. Therefore $\boxed{T \text{ is not a subspace}}$.

4.4 Rank, nullity and determinant

For an $m\times n$ matrix, the rank-nullity theorem states

$$\operatorname{rank}(A)+\dim \mathcal{N}(A)=n.$$

The right side is the number of input coordinates, not the number of equations. A matrix with more columns than pivots necessarily has nonzero null vectors.

For a two-by-two matrix,

$$\det\begin{pmatrix} a & b \\ c & d \end{pmatrix}=ad-bc.$$

A square matrix is invertible exactly when its determinant is nonzero. Invertibility means $Ax=b$ has a unique solution for every b. Determinants are not defined for arbitrary rectangular matrices; rank and nullity still are.

Worked problem 4.5 Find a null-space basis

Problem. Find the rank and null space of $A=\begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{pmatrix}$.

Understanding the method. The second row is twice the first. Solve the remaining homogeneous equation by choosing the two free variables as parameters.

Calculation.

$$x_1+2x_2+3x_3=0.$$

Let $x_2=s$ and $x_3=t$. Then $x_1=-2s-3t$, so

$$x=s\begin{pmatrix} -2 \\ 1 \\ 0 \end{pmatrix}+t\begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix}.$$

The two displayed vectors are independent: the second and third coordinates force any zero combination to have both coefficients zero. Therefore

$$\boxed{\operatorname{rank} A=1, \quad \dim \mathcal{N}(A)=2}.$$

Check and meaning. Rank plus nullity is $1+2=3$, the number of columns. Each basis vector can be checked by direct multiplication with A.

Worked problem 4.6 An inverse and a solution

Problem. Find the inverse of $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and use it to solve $Ax = (5, 11)^T$.

Understanding the method. For a nonsingular two-by-two matrix, exchange the diagonal entries, negate the off-diagonal entries, then divide by the determinant.

Calculation.

$$\begin{aligned} \det A &= 4 - 6 = -2, \\ A^{-1} &= \frac{1}{-2} \begin{pmatrix} 4 & -2 \\ -3 & 1 \end{pmatrix} = \begin{pmatrix} -2 & 1 \\ 3/2 & -1/2 \end{pmatrix}, \\ x &= A^{-1}b = \begin{pmatrix} -10 + 11 \\ 15/2 - 11/2 \end{pmatrix} = \boxed{\begin{pmatrix} 1 \\ 2 \end{pmatrix}}. \end{aligned}$$

Check and meaning. The result agrees with elimination. For large numerical systems, solving by factorisation is generally preferable to explicitly calculating an inverse.

4.5 Why dependence removes uniqueness

If $Az = 0$ for a nonzero vector z and $Ax_0 = b$, then every $x_0 + tz$ also solves the same equation: $A(x_0 + tz) = Ax_0 + tAz = b$. This is the algebraic explanation of a whole line of solutions, not just a parameter trick.

Worked problem 4.7 Generate all solutions from one

Problem. For $A = \begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}$ and $b = (3, 6)^T$, use a particular solution and a null vector to describe all solutions.

Understanding the method. Find one convenient solution and add every vector that A maps to zero.

Calculation.

A particular solution is $x_0 = (3, 0)^T$. The null vector $z = (-2, 1)^T$ satisfies $Az = 0$. Hence

$$x = x_0 + tz = \begin{pmatrix} 3 - 2t \\ t \end{pmatrix}, \qquad t \in \mathbb{R}.$$

Substituting gives $x_1 + 2x_2 = 3 - 2t + 2t = 3$, while the second equation is twice the first. Therefore this parameterisation supplies every solution.

Check and meaning. The free parameter corresponds to nullity one. The fitted output stays fixed while coefficients move along the null direction.

Chapter 5

Orthogonality, projections and least squares

The idea before the calculation

Measurements rarely agree perfectly. Several observations may ask incompatible things of the same model. Least squares chooses a fitted answer by making the sum of squared residuals as small as possible. Geometrically, it projects the observations onto the space of outputs the model can represent.

At the optimum, the residual is orthogonal to that space under the chosen inner product. This is a statement about the fitting geometry, not proof that the model explains the cause of the observations. Correlated predictors can also make individual coefficients unstable even when predictions look reasonable. Ask both questions: how close is the fit, and how much can we trust the quantities inferred from it?

5.1 Angles and distances

For nonzero vectors, $u^T v = \|u\|\|v\| \cos\theta$. Thus a zero dot product means a right angle. Orthogonality is central because it separates independent geometric directions. A vector of length one is a unit vector; normalising u means dividing it by $\|u\|$.

A collection is orthonormal when its members are mutually orthogonal unit vectors. If the columns of Q are orthonormal, $Q^T Q = I$. The reverse product QQ^T is the identity only if those columns span the whole output space; otherwise it is a projection.

Worked problem 5.1 Normalise and test orthogonality

Problem. For $u = (1, 2)^T$ and $v = (2, -1)^T$, find their lengths and unit directions.

Understanding the method. Use the square root of the sum of squared components. A zero dot product verifies that the directions are perpendicular.

Calculation.

$$\|u\| = \sqrt{1+4} = \sqrt{5}, \qquad \|v\| = \sqrt{4+1} = \sqrt{5},$$
$$u^T v = 1(2) + 2(-1) = 0,$$
$$\hat{u} = \frac{1}{\sqrt{5}}(1,2)^T, \qquad \hat{v} = \frac{1}{\sqrt{5}}(2,-1)^T.$$

Each normalised vector has squared length $(1+4)/5 = 1$, and their dot product remains zero. They form an orthonormal basis for $\mathbb{R}^2$.

5.2 Projection derived, not memorised

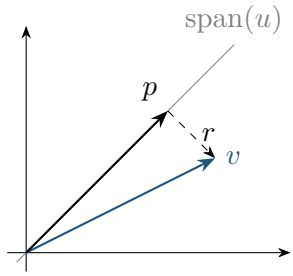
Suppose we approximate v by a point on the line through a nonzero vector u. Every such point has form αu. At the nearest point the residual $r = v - \alpha u$ is perpendicular to the line. Therefore

$$u^T(v - \alpha u) = 0 \quad \Rightarrow \quad \alpha = \frac{u^T v}{u^T u}.$$

The projection vector is $p = \alpha u$, and its matrix form is

$$p = Pv, \qquad P = \frac{uu^T}{u^T u}.$$

The coefficient α is a scalar; p is a vector. Mixing them is a frequent error.



Worked problem 5.2 Projection onto a line

Problem. Project $v = (2,1)^T$ onto $u = (1,1)^T$ and find the squared residual length.

Understanding the method. Calculate the scalar coefficient, multiply it by the direction vector, then subtract from v.

Calculation.

$$u^T v = 3, \qquad u^T u = 2, \qquad \alpha = 3/2,$$
$$p = (3/2, 3/2)^T,$$
$$r = (2,1)^T - (3/2, 3/2)^T = (1/2, -1/2)^T,$$
$$\|r\|^2 = (1/2)^2 + (-1/2)^2 = \boxed{1/2}.$$

Check and meaning. $u^T r = 1/2 - 1/2 = 0$. Also $\|v\|^2 = 5 = \|p\|^2 + \|r\|^2 = 9/2 + 1/2$, consistent with a right-angle decomposition.

5.3 Least squares as a projection

When $Ax = b$ has no exact solution, least squares chooses $\hat{x}$ minimising $\|b - Ax\|^2$. The fitted output $A\hat{x}$ lies in the column space of A. The residual must be perpendicular to every column, giving

$$A^T(b - A\hat{x}) = 0 \quad \Rightarrow \quad A^T A\hat{x} = A^T b.$$

These are the normal equations. If the columns of A are independent, $A^T A$ is invertible and the coefficient solution is unique. If the columns are dependent, fitted outputs can still be unique while coefficient representations are not.

The normal equations also follow by differentiating the squared-error objective. Expanding $\|b - Ax\|^2$ gives $b^T b - 2x^T A^T b + x^T A^T Ax$, whose gradient is $-2A^T b + 2A^T Ax$. Setting it to zero gives the same equations.

Worked problem 5.3 A constant fitted to three observations

Problem. Find the least-squares constant for $b = (1, 3, 8)^T$.

Understanding the method. A constant model predicts c for every observation. Its design matrix is the single column $A = (1, 1, 1)^T$.

Calculation.

$$A^T A = 1 + 1 + 1 = 3,$$
$$A^T b = 1 + 3 + 8 = 12,$$
$$3\hat{c} = 12 \quad \Rightarrow \quad \boxed{\hat{c} = 4},$$
$$r = b - A\hat{c} = (-3, -1, 4)^T,$$
$$\|r\|^2 = 9 + 1 + 16 = 26.$$

Check and meaning. The residual sum $-3 - 1 + 4 = 0$ is the orthogonality condition against the constant column. This explains why the arithmetic mean minimises squared error for a constant model.

Worked problem 5.4 Fit a straight line with an intercept

Problem. Fit $y \approx a + bx$ to $(0,1),(1,2),(2,2)$ by least squares.

Understanding the method. The two columns are a constant column and the observed x values. Each row corresponds to one data point.

Calculation.

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix}, \quad y = \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix}, \quad A^T A = \begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix}, \quad A^T y = \begin{pmatrix} 5 \\ 6 \end{pmatrix}.$$

The normal equations are

$$\begin{aligned} 3a + 3b &= 5, \\ 3a + 5b &= 6. \end{aligned}$$

Subtracting gives $2b = 1$, hence $b = 1/2$ and $a = (5 - 3/2)/3 = 7/6$. The fitted line is

$$\boxed{\hat{y} = \frac{7}{6} + \frac{1}{2}x}.$$

Predictions are $(7/6, 10/6, 13/6)$, so residuals are $(-1/6, 2/6, -1/6)$. Their sum is zero, and their dot product with the x column is $0 + 2/6 - 2/6 = 0$.

Check and meaning. Both design columns are orthogonal to the residual, confirming the normal equations numerically.

5.4 Orthonormal coordinates and numerical conditioning

For orthonormal columns Q, the projection of b onto their span is $QQ^T b$. Its coefficients are simply $Q^T b$, because $Q^T Q = I$. This is why QR factorisation is useful in least squares: if $A = QR$, the coefficient system reduces to $R\hat{x} = Q^T b$ when A has independent columns.

Conditioning concerns sensitivity of the mathematical problem. Nearly parallel columns can make coefficients highly sensitive to small perturbations in observations. An algorithm's numerical stability is a separate property. Forming $A^T A$ can worsen conditioning, so the normal equations are valuable for hand derivations but QR or SVD are often preferred in numerical work.

Worked problem 5.5 Projection using an orthonormal pair

Problem. Let $q_1 = (1,0,0)^T$, $q_2 = (0,1,0)^T$ and $b = (3,4,5)^T$. Find the projection onto their span.

Understanding the method. The pair spans the coordinate plane $z = 0$. In an orthonormal basis the coefficients are dot products.

Calculation.

$$q_1^T b = 3, \qquad q_2^T b = 4,$$
$$p = 3q_1 + 4q_2 = (3, 4, 0)^T,$$
$$r = b - p = (0, 0, 5)^T.$$

Thus $\boxed{\|r\| = 5}$. Both $q_1^T r$ and $q_2^T r$ are zero. Here $QQ^T = \mathrm{diag}(1, 1, 0)$, not the three-by-three identity.

Chapter 6

Eigenvalues, quadratic forms and curvature

The idea before the calculation

Some directions retain their identity when a linear transformation acts: they are stretched, shrunk or reversed without being turned into a different line. These are eigenvector directions. Their eigenvalues tell us the corresponding factors of change. In a vibrating system, related ideas describe natural modes; in optimisation, they describe local curvature.

For real symmetric matrices, an orthonormal eigenvector basis makes this geometry especially clear. General matrices need not offer that same convenient structure. In a positive-definite quadratic problem, very unequal eigenvalues produce a narrow valley: progress can be easy along one direction and painfully slow along another. Eigenvalues therefore help explain behaviour that a list of matrix entries leaves hidden.

6.1 Directions preserved by a transformation

An eigenvector is a nonzero vector whose direction is preserved by a square matrix: $Av = \lambda v$. The number λ is the eigenvalue. A negative eigenvalue reverses the direction, a magnitude above one stretches it, and zero sends that eigenvector to zero. The zero vector is excluded because it would satisfy the equation for every λ without identifying a direction.

Rearranging gives $(A-\lambda I)v = 0$. A nonzero solution exists only if $A-\lambda I$ is singular, hence $\det(A - \lambda I) = 0$. This polynomial equation supplies eigenvalues; solving a homogeneous system for each eigenvalue supplies eigenvectors.

Worked problem 6.1 Find and verify two eigenpairs

Problem. Find the eigenvalues and eigenvectors of $A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$.

Understanding the method. First solve the characteristic equation. Then substitute each root into $A - \lambda I$ and find a nonzero null vector.

Calculation.

$$\begin{aligned}\det(A - \lambda I) &= (2 - \lambda)^2 - 1 \\ &= \lambda^2 - 4\lambda + 3 = (\lambda - 1)(\lambda - 3).\end{aligned}$$

Thus $\lambda = 1$ or 3. For $\lambda = 1$, the first row gives $v_1 + v_2 = 0$; choose $v = (1, -1)^T$. For $\lambda = 3$, $-v_1 + v_2 = 0$; choose $v = (1, 1)^T$. Direct verification is

$$A\begin{pmatrix}1\\-1\end{pmatrix} = \begin{pmatrix}1\\-1\end{pmatrix}, \qquad A\begin{pmatrix}1\\1\end{pmatrix} = 3\begin{pmatrix}1\\1\end{pmatrix}.$$

Check and meaning. The eigenvalue sum is 4, the trace; the product is 3, the determinant. The two eigenvectors are perpendicular because A is symmetric and the eigenvalues are distinct.

6.2 Diagonalisation and powers

If a matrix has enough independent eigenvectors to form an invertible matrix P, then $A = PDP^{-1}$ where D is diagonal with eigenvalues on its diagonal. This turns repeated multiplication into scalar powers: $A^k = PD^kP^{-1}$. Not every square matrix is diagonalizable, but a real symmetric matrix has an orthonormal eigenbasis, so $A = QDQ^T$.

The trace always equals the sum of eigenvalues and the determinant their product, counting algebraic multiplicities. These identities do not require the matrix to be symmetric. Symmetry matters when using eigenvalues to classify a quadratic form.

Worked problem 6.2 Apply a matrix repeatedly

Problem. For the previous matrix, calculate $A^3(1, 1)^T$ without multiplying three full matrices.

Understanding the method. The vector is an eigenvector with eigenvalue 3, so each multiplication scales it by another factor 3.

Calculation.

$$\begin{aligned}Av &= 3v, \\ A^2v &= A(3v) = 3Av = 3^2v, \\ A^3v &= 3^3v = 27(1, 1)^T = \boxed{(27, 27)^T}.\end{aligned}$$

Check and meaning. This shortcut applies because the input is a known eigenvector. An arbitrary vector must first be decomposed into eigenvector components.

6.3 Quadratic forms and positive definiteness

A quadratic form is $x^T Ax$. If A is symmetric, it is positive definite when this quantity is positive for every nonzero x. Positive semidefinite permits zero. An indefinite matrix permits both positive and negative values. These properties classify local curvature in optimisation.

For a symmetric matrix, the eigenvalues reveal the classification. Writing $x = Qz$ gives $x^T Ax = z^T Dz = \sum_i \lambda_i z_i^2$. Therefore all eigenvalues positive implies positive definiteness; all nonnegative implies positive semidefiniteness; mixed positive and negative eigenvalues imply indefiniteness.

For $\begin{pmatrix} a & b \\ b & c \end{pmatrix}$, positive definiteness is equivalent to $a > 0$ and $ac - b^2 > 0$. A positive determinant alone is insufficient: $-I$ has positive determinant in dimension two but is negative definite.

Worked problem 6.3 Classify a symmetric matrix

Problem. Classify $B = \begin{pmatrix} 2 & 2 \\ 2 & 6 \end{pmatrix}$.

Understanding the method. Use the two-by-two positive-definiteness test, then verify by completing a square in the quadratic form.

Calculation.

$$B_{11} = 2 > 0, \qquad \det B = 12 - 4 = 8 > 0.$$

Thus B is positive definite. Directly,

$$\begin{aligned} \begin{pmatrix} x & y \end{pmatrix} B \begin{pmatrix} x \\ y \end{pmatrix} &= 2x^2 + 4xy + 6y^2 \\ &= 2(x+y)^2 + 4y^2. \end{aligned}$$

Both terms are nonnegative and can vanish simultaneously only at $(0, 0)$. Hence the quadratic form is positive for every nonzero vector.

Worked problem 6.4 A zero eigenvalue changes the conclusion

Problem. A real symmetric matrix has eigenvalues 4 and 0. Is it positive definite?

Understanding the method. A positive-definite form must be strictly positive in every nonzero direction. A zero eigenvalue supplies a nonzero eigenvector along which the form is zero.

Calculation.

If $Av = 0$ for a nonzero eigenvector v, then

$$v^T Av = v^T 0 = 0.$$

The strict inequality fails. Since the eigenvalues are both nonnegative, the matrix is $\boxed{\text{positive semidefinite, not positive definite}}$.

Check and meaning. The distinction matters: a quadratic loss may have a flat direction and multiple minimisers when curvature is only semidefinite.

6.4 Condition numbers and sensitivity

For an invertible symmetric positive-definite matrix, the spectral condition number is $\kappa_2(A) = \lambda_{\max}/\lambda_{\min}$. More generally it is the ratio of largest to smallest singular value. A large value indicates a direction that is weak relative to another, allowing some input perturbations to produce large relative solution changes.

Worked problem 6.5 A poorly scaled system

Problem. Solve $Ax = b$ for $A = \operatorname{diag}(1, 0.001)$ and $b = (1, 0.001)^T$. Then change the second component of b to 0.002.

Understanding the method. A diagonal system divides each right-hand-side component by its corresponding diagonal entry. The small diagonal entry amplifies an absolute change.

Calculation.

Initially,

$$x_1 = 1/1 = 1, \qquad x_2 = 0.001/0.001 = 1.$$

After the perturbation,

$$\tilde{x}_1 = 1, \qquad \tilde{x}_2 = 0.002/0.001 = 2.$$

The absolute right-hand-side change is 0.001 in the second component, while the solution change there is 1. The condition number is

$$\boxed{\kappa_2(A) = 1/0.001 = 1000}.$$

Check and meaning. The example illustrates sensitivity of the system. It does not mean every perturbation attains the worst-case relative amplification.

6.5 Why this belongs in machine learning

For squared-error regression, the Hessian contains $X^T X$. If the columns of X are dependent, this Hessian is singular. Adding a positive ridge penalty adds positive diagonal curvature and can make the system uniquely solvable. Eigenvalues also explain gradient-descent stability: strongly curved directions require smaller steps than weakly curved ones. These are mathematical links, not separate facts to memorise.

Chapter 7

Probability, counting and conditioning

The idea before the calculation

Uncertainty begins with a question and a description of possible outcomes. Probability gives that description rules. Conditioning then changes the information available: the chance of rain before seeing clouds need not equal the chance after seeing them. The event has not travelled backwards in time; our information has changed.

A rare condition can remain unlikely after a positive test when false positives are common enough. Bayes' rule keeps this background frequency in the calculation. Independence has an equally precise meaning: learning one event does not change the probability of the other. It is different from saying two events cannot happen together. Clear language about events prevents errors before any fraction is evaluated.

7.1 Sample spaces and events

A sample space lists the possible outcomes of an experiment. An event is a subset of outcomes. For a fair die the outcomes are $1, \dots, 6$, and the event "even" is $\{2, 4, 6\}$. When all outcomes are equally likely, probability equals favourable outcomes divided by total outcomes. That counting shortcut is invalid when the outcomes have unequal probabilities.

Probabilities lie between zero and one, the whole sample space has probability one, and disjoint events have additive probabilities. The complement A^c is the event that A does not occur, so $P(A^c) = 1 - P(A)$. For overlapping events,

$$P(A \cup B) = P(A) + P(B) - P(A \cap B).$$

The subtraction removes double counting of the overlap.

Worked problem 7.1 A union of events

Problem. A fair die is rolled. Find the probability that the result is even or greater than four.

Understanding the method. "Or" corresponds to a union. The events overlap at outcome 6, which must be counted only once.

Calculation.

$$A = \{2, 4, 6\}, \qquad B = \{5, 6\}, \qquad A \cap B = \{6\},$$
$$P(A \cup B) = 3/6 + 2/6 - 1/6 = \boxed{2/3}.$$

Directly the union is $\{2, 4, 5, 6\}$, four outcomes out of six.

7.2 Counting arrangements

The factorial $n!$ counts arrangements of n distinct objects: $n(n-1)\cdots 1$. Choosing k objects without regard to order uses the binomial coefficient

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}.$$

The denominator removes the internal rearrangements of the chosen and unchosen groups. If order matters, the count is instead $n!/(n-k)!$. Before using a formula, decide whether replacement is allowed and whether order distinguishes outcomes.

Worked problem 7.2 Unordered selection

Problem. How many different three-person committees can be formed from six people?

Understanding the method. A committee has no ordered roles: choosing A, B, C gives the same committee as C, A, B. Divide the ordered count by 3!.

Calculation.

$$\binom{6}{3} = \frac{6!}{3!3!} = \frac{6 \cdot 5 \cdot 4}{3 \cdot 2 \cdot 1} = \boxed{20}.$$

If the three roles were president, secretary and treasurer, the count would be $6 \cdot 5 \cdot 4 = 120$, because role assignment matters.

7.3 Conditional probability

Conditioning restricts attention to outcomes where some event has occurred. For $P(B) > 0$,

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}.$$

The denominator is the probability of the restricted population. Rearranging gives the multiplication rule $P(A \cap B) = P(A \mid B)P(B)$.

Events A and B are independent when $P(A \cap B) = P(A)P(B)$. For a conditioning event of positive probability, this is equivalent to $P(A \mid B) = P(A)$. Mutually exclusive positive-probability events are not independent: occurrence of one rules out the other.

Worked problem 7.3 Drawing without replacement

Problem. A box has three red and two blue balls. Two are drawn without replacement. Find the probability that both are red.

Understanding the method. The second draw uses a changed box. After a red ball is removed, only two red balls remain among four total.

Calculation.

$$\begin{aligned} P(R_1) &= 3/5, \\ P(R_2 \mid R_1) &= 2/4, \\ P(R_1 \cap R_2) &= (3/5)(2/4) = \boxed{3/10}. \end{aligned}$$

As an unordered count, $\binom{3}{2}/\binom{5}{2} = 3/10$, agreeing with the sequential calculation.

Common mistake. Using $(3/5)^2$ assumes replacement or independent draws. Neither holds here.

7.4 Total probability and Bayes' rule

If mutually exclusive cases B_i cover the whole sample space, sum over them:

$$P(A) = \sum_i P(A \mid B_i)P(B_i).$$

Bayes' rule reverses the direction of conditioning:

$$P(B_j \mid A) = \frac{P(A \mid B_j)P(B_j)}{\sum_i P(A \mid B_i)P(B_i)}.$$

In a diagnostic problem, sensitivity $P(+ \mid D)$ is not the same as the probability of disease among positive tests $P(D \mid +)$. The latter depends on prevalence $P(D)$ as well as test behaviour.

Worked problem 7.4 Base rates and a positive test

Problem. Disease prevalence is 1%, sensitivity is 90%, and the false-positive rate is 10%. Find $P(D \mid +)$.

Understanding the method. Split positives into true positives and false positives. The denominator must include both sources of a positive result.

Calculation.

$$\begin{aligned} P(D) &= 0.01, & P(D^c) &= 0.99, \\ P(+ \cap D) &= 0.90(0.01) = 0.009, \\ P(+ \cap D^c) &= 0.10(0.99) = 0.099, \\ P(+) &= 0.009 + 0.099 = 0.108, \\ P(D \mid +) &= \frac{0.009}{0.108} = \boxed{\frac{1}{12} \approx 8.33\%}. \end{aligned}$$

Check and meaning. Among 10,000 people, expect 90 true positives and 990 false positives. Then $90/(90 + 990) = 1/12$. The small prevalence explains why most positives can be false positives despite high sensitivity.

Worked problem 7.5 A mixture of production lines

Problem. Line A makes 60% of items and has defect probability 0.02. Line B makes 40% and has defect probability 0.05. Given a defective item, find the probability it came from B.

Understanding the method. The production shares are prior probabilities. Multiply each share by its conditional defect rate to obtain the contributions to all defects.

Calculation.

$$\begin{aligned} P(D) &= P(D \mid A)P(A) + P(D \mid B)P(B) \\ &= 0.02(0.60) + 0.05(0.40) \\ &= 0.012 + 0.020 = 0.032, \\ P(B \mid D) &= \frac{0.05(0.40)}{0.032} = \boxed{0.625}. \end{aligned}$$

Check and meaning. B makes fewer items but has a higher defect rate, so it contributes 62.5% of all defects.

7.5 Complements and repeated trials

For independent trials, the probability that all specified outcomes occur is the product of their probabilities. “At least one” is often easiest through its complement, “none”. This avoids listing cases of one, two, three or more occurrences.

Worked problem 7.6 At least one occurrence

Problem. An event has probability 0.2 on each of five independent trials. Find the probability it occurs at least once.

Understanding the method. No occurrences has probability $(1-0.2)^5$. The desired event is its complement.

Calculation.

$$P(\text{none}) = 0.8^5 = 0.32768,$$
$$P(\text{at least one}) = 1 - 0.32768 = \boxed{0.67232}.$$

Check and meaning. The result exceeds the single-trial probability but remains below one. Simply adding $5(0.2)$ would count overlapping occurrence patterns more than once.

Chapter 8

Random variables and distributions

The idea before the calculation

A random variable turns an outcome into a number: the count of arrivals, a measured lifetime, or the result of a noisy experiment. Its distribution describes how probability is spread across possible values. An expectation is a probability-weighted average, not a promise about the next observation.

For a continuous variable, probability belongs to intervals. A density is a rate of probability per unit of the variable; its height can exceed one while its total area remains one. Variance measures spread in squared units, and standard deviation returns to the original units. Choosing a distribution means choosing assumptions about the mechanism. A familiar curve is useful only when those assumptions are reasonable.

8.1 A numerical description of uncertainty

A random variable assigns a number to each outcome. A discrete variable has a probability mass function $p(x)$ with $\sum_x p(x) = 1$. A continuous variable uses a density $f(x)$ satisfying $f(x) \geq 0$ and $\int f(x)\,\mathrm{d}x = 1$. For a continuous variable, probability is an integral over an interval; the probability of one exact value is zero even when the density there is positive.

The cumulative distribution function $F(t) = P(X \leq t)$ works for both types. For a continuous density, $F(t) = \int_{-\infty}^{t} f(x)\,\mathrm{d}x$. A density's units are reciprocal to the units of x, allowing its integral to be dimensionless.

8.2 Expectation, variance and transformation

Expectation is a probability-weighted mean:

$$\mathbb{E}[X] = \sum_x xp(x) \quad \text{or} \quad \mathbb{E}[X] = \int x f(x)\,\mathrm{d}x.$$

It need not equal a possible realised value: the expected die result is 3.5. Linearity gives $\mathbb{E}[aX + bY + c] = a\mathbb{E}[X] + b\mathbb{E}[Y] + c$, without requiring independence.

Variance measures squared deviation about the mean:

$$\operatorname{Var}(X) = \mathbb{E}[(X - \mathbb{E}X)^2] = \mathbb{E}[X^2] - (\mathbb{E}X)^2.$$

It is nonnegative and has squared units. Standard deviation is its square root. Constants do not add uncertainty, so $\operatorname{Var}(aX + b) = a^2 \operatorname{Var}(X)$.

Worked problem 8.1 Expectation and variance from a table

Problem. Let $P(X = 0) = 0.2$, $P(X = 1) = 0.5$, and $P(X = 2) = 0.3$. Find $\mathbb{E}[X]$ and $\operatorname{Var}(X)$.

Understanding the method. Calculate the first and second moments separately, then subtract the square of the mean.

Calculation.

$$\begin{aligned}
\mathbb{E}[X] &= 0(0.2) + 1(0.5) + 2(0.3) = 1.1,\\
\mathbb{E}[X^2] &= 0^2(0.2) + 1^2(0.5) + 2^2(0.3) = 1.7,\\
\operatorname{Var}(X) &= 1.7 - (1.1)^2 = 1.7 - 1.21 = \boxed{0.49}.
\end{aligned}$$

Thus the standard deviation is $\sqrt{0.49} = 0.7$.

Check and meaning. The probabilities add to one, and the mean lies between the smallest and largest values. Negative variance would indicate an arithmetic or model error.

8.3 Bernoulli and binomial models

A Bernoulli variable equals one for success and zero for failure. If success probability is p, then $\mathbb{E}[X] = p$ and $\operatorname{Var}(X) = p(1 - p)$. Summing n independent Bernoulli trials with the same p gives a binomial variable:

$$P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}, \quad \mathbb{E}[X] = np, \quad \operatorname{Var}(X) = np(1 - p).$$

The coefficient counts the placements of the k successes. Independence and constant probability are assumptions, not automatic properties of any repeated observations.

Worked problem 8.2 An exact binomial probability

Problem. A fair coin is tossed four times independently. Find the probability of exactly two heads.

Understanding the method. Each sequence with two heads has probability $(1/2)^4$. Count how many such sequences are possible.

Calculation.

$$\begin{aligned}\binom{4}{2} &= \frac{4!}{2!2!} = 6,\\ P(X=2) &= 6\left(\frac{1}{2}\right)^2\left(\frac{1}{2}\right)^2\\ &= \frac{6}{16} = \boxed{\frac{3}{8} = 0.375}.\end{aligned}$$

Check and meaning. The six arrangements are HHTT, HTHT, HTTH, THHT, THTH and TTHH. Omitting the coefficient counts only one arrangement.

Worked problem 8.3 Binomial moments

Problem. If $X \sim \text{Binomial}(10, 0.3)$, find its mean, variance and standard deviation.

Understanding the method. Treat X as the sum of ten independent success indicators. Means add, and independence makes variances add.

Calculation.

$$\begin{aligned}\mathbb{E}[X] &= 10(0.3) = \boxed{3},\\ \text{Var}(X) &= 10(0.3)(0.7) = \boxed{2.1},\\ \text{SD}(X) &= \sqrt{2.1} \approx \boxed{1.449}.\end{aligned}$$

Check and meaning. The variance is not np^2. The factor $1-p$ reflects the Bernoulli uncertainty, which vanishes when success is certain or impossible.

8.4 Continuous uniform and exponential distributions

A uniform variable on $[a, b]$ has density $1/(b-a)$ there, mean $(a+b)/2$ and variance $(b-a)^2/12$. Interval probability is interval length divided by total length.

An exponential variable with rate $\lambda > 0$ has density $\lambda e^{-\lambda x}$ for $x \geq 0$. The survival probability is

$$P(X > t) = \int_t^\infty \lambda e^{-\lambda x}\,\mathrm{d}x = e^{-\lambda t} \quad (t \geq 0).$$

Its mean is $1/\lambda$. A rate has reciprocal-time units when X is a waiting time.

Worked problem 8.4 Waiting-time probabilities

Problem. A waiting time is exponential with rate 0.5 per minute. Find $P(X > 4)$ and $P(2 < X \leq 4)$.

Understanding the method. Use survival probabilities. To be between two times, the waiting time must survive the earlier boundary but not the later one.

Calculation.

$$\begin{aligned} P(X > 4) &= e^{-0.5(4)} = \boxed{e^{-2} \approx 0.1353}, \\ P(2 < X \leq 4) &= P(X > 2) - P(X > 4) \\ &= e^{-1} - e^{-2} \approx \boxed{0.2325}. \end{aligned}$$

Check and meaning. The rate times time is dimensionless, as an exponent must be. Reversing the difference would produce an impossible negative probability.

8.5 Normal standardisation

For $X \sim N(\mu, \sigma^2)$ with $\sigma > 0$, the standardised variable $Z = (X - \mu)/\sigma$ has mean zero and variance one. Standardisation converts a threshold into a number of standard deviations. Computing a numerical tail probability then requires a normal table, a calculator function, or a supplied approximation; the z-score itself is not a probability.

Worked problem 8.5 A normal threshold

Problem. Measurements are normal with mean 10 and standard deviation 2. Express $P(X > 14)$ using the standard normal distribution.

Understanding the method. Subtract the mean and divide by the standard deviation on both sides of the inequality. The divisor is positive, so the inequality direction stays unchanged.

Calculation.

$$P(X > 14) = P\left(\frac{X - 10}{2} > \frac{14 - 10}{2}\right) = P(Z > 2) = \boxed{1 - \Phi(2)}.$$

If $\Phi(2) \approx 0.97725$ is supplied, the tail is approximately 0.02275.

Check and meaning. A threshold two standard deviations above the mean should have a small upper-tail probability, not a probability near one.

8.6 Covariance and dependence

Covariance is $\operatorname{Cov}(X,Y)=\mathbb{E}[(X-\mathbb{E}X)(Y-\mathbb{E}Y)]$. For finite second moments,

$$\operatorname{Var}(aX+bY)=a^2\operatorname{Var}(X)+b^2\operatorname{Var}(Y)+2ab\operatorname{Cov}(X,Y).$$

Independence makes covariance zero, but the converse is not generally true. Covariance carries units of X times units of Y; correlation divides by standard deviations to make a dimensionless measure between -1 and 1 when both variances are positive.

Worked problem 8.6 Variance of a difference

Problem. Given $\operatorname{Var}(X)=4$, $\operatorname{Var}(Y)=9$ and $\operatorname{Cov}(X,Y)=1$, compute $\operatorname{Var}(2X-Y)$.

Understanding the method. The coefficient of Y is -1, so the covariance contribution is negative. Variances themselves do not acquire a negative sign.

Calculation.

$$\begin{aligned}\operatorname{Var}(2X-Y)&=2^2(4)+(-1)^2(9)+2(2)(-1)(1)\\&=16+9-4=\boxed{21}.\end{aligned}$$

Check and meaning. Under independence the covariance would be zero and the result would be 25. The stated positive covariance reduces uncertainty in this particular difference.

8.7 Uncorrelated does not mean independent

Zero covariance removes linear association, but nonlinear dependence can remain. Independence is a stronger statement about the joint distribution.

Worked problem 8.7 A dependent pair with zero covariance

Problem. Let X be equally likely to be $-1,0,1$ and let $Y=X^2$. Show that covariance is zero but the variables are dependent.

Understanding the method. Compute $\mathbb{E}[X]$ and $\mathbb{E}[XY]$ by symmetry. Then exhibit a conditional probability changed by knowing X.

Calculation.

$$\begin{aligned}\mathbb{E}[X]&=(-1+0+1)/3=0,\\\mathbb{E}[Y]&=(1+0+1)/3=2/3,\\\mathbb{E}[XY]&=\mathbb{E}[X^3]=(-1+0+1)/3=0,\\\operatorname{Cov}(X,Y)&=\mathbb{E}[XY]-\mathbb{E}[X]\mathbb{E}[Y]=0.\end{aligned}$$

But $P(Y = 0) = 1/3$, whereas $P(Y = 0 \mid X = 0) = 1$. Knowing X determines Y, so they are not independent.

Check and meaning. This example explains why a correlation close to zero does not establish the absence of every relationship.

Chapter 9

Sampling, estimation and regression

The idea before the calculation

A sample is a small window onto a larger population. Enlarging the window can reduce random fluctuation, but it cannot repair a window pointed in the wrong direction. Selection bias remains bias even in a very large dataset.

Estimation asks what the observations suggest and how uncertain that suggestion is. A frequentist confidence procedure with 95% coverage captures the fixed population parameter in 95% of repeated samples under its assumptions; it does not assign a 95% probability to that fixed parameter after one interval is computed. Regression also requires care: a relationship in observed data need not be causal. Good inference keeps the sampling process beside the formula, rather than letting the formula hide it.

9.1 A parameter is not an estimate

A parameter describes a population or model, such as its unknown mean μ. An estimator is a rule computed from a random sample, such as $\bar{X} = n^{-1}\sum_i X_i$. The realised numerical value is an estimate. Before data are observed, the estimator varies from sample to sample. That variation is its sampling uncertainty.

For independent observations with common mean μ and variance σ^2,

$$\mathbb{E}[\bar{X}] = \mu, \qquad \operatorname{Var}(\bar{X}) = \frac{\sigma^2}{n}, \qquad \operatorname{SE}(\bar{X}) = \frac{\sigma}{\sqrt{n}}.$$

The variance calculation follows because the factor $1/n$ is squared and n variances add. If observations are correlated, covariance terms must be included, so the familiar $1/n$ reduction may fail.

Worked problem 9.1 Sample size and uncertainty

Problem. Independent measurements have standard deviation 6. What is the standard error of their average for $n = 9$ and for $n = 36$?

Understanding the method. The standard deviation of one measurement differs from the standard deviation of the average. Divide the former by $\sqrt{n}$.

Calculation.

$$\mathrm{SE}_9 = 6/\sqrt{9} = 6/3 = \boxed{2},$$
$$\mathrm{SE}_{36} = 6/\sqrt{36} = 6/6 = \boxed{1}.$$

Check and meaning. Quadrupling sample size halves the standard error. Doubling sample size would reduce it by $\sqrt{2}$, not by two.

9.2 Bias, variance and mean squared error

An estimator $\hat{\theta}$ is unbiased if $\mathbb{E}[\hat{\theta}] = \theta$. Its bias is $\mathbb{E}[\hat{\theta}] - \theta$. Accuracy is not determined by bias alone because a highly variable unbiased estimator can be less useful than a mildly biased stable one. The estimation mean squared error decomposes as

$$\mathbb{E}[(\hat{\theta} - \theta)^2] = \mathrm{Var}(\hat{\theta}) + \left(\mathbb{E}[\hat{\theta}] - \theta\right)^2.$$

This identity follows by adding and subtracting $\mathbb{E}[\hat{\theta}]$ inside the square; the cross term has expectation zero. It is an estimation identity, related to but distinct from the predictive bias–variance discussion later.

Worked problem 9.2 Compare two estimators

Problem. Estimator A is unbiased with variance 4. Estimator B has bias 1 and variance 2. Which has smaller MSE?

Understanding the method. Add variance to squared bias, not to the unsquared bias.

Calculation.

$$\mathrm{MSE}(A) = 4 + 0^2 = 4,$$
$$\mathrm{MSE}(B) = 2 + 1^2 = 3.$$

Thus $\boxed{B \text{ has the smaller MSE}}$ despite being biased.

Check and meaning. If B's bias were 2, its MSE would be $2 + 4 = 6$, reversing the comparison. The actual magnitudes matter.

9.3 Sample variance and degrees of freedom

For observations $x_1, \dots, x_n$, the unbiased sample variance uses

$$s^2 = \frac{1}{n-1}\sum_i (x_i - \bar{x})^2 \quad (n > 1).$$

One degree of freedom is used by estimating the mean: the residuals sum to zero, so only $n-1$ can vary independently. In contrast, the average squared deviation about the fitted mean divides by n. Both calculations exist but answer different questions; do not silently switch denominators.

Worked problem 9.3 Variance of three measurements

Problem. For data $2, 4, 6$, calculate the mean, unbiased sample variance and average squared deviation about the mean.

Understanding the method. Compute deviations first, then apply the requested denominator.

Calculation.

$$\begin{aligned} \bar{x} &= (2+4+6)/3 = 4, \\ (x_i - \bar{x}) &= (-2, 0, 2), \\ \sum_i (x_i - \bar{x})^2 &= 4 + 0 + 4 = 8, \\ s^2 &= 8/(3-1) = \boxed{4}, \\ \frac{1}{n}\sum_i (x_i - \bar{x})^2 &= \boxed{8/3}. \end{aligned}$$

Check and meaning. The sample standard deviation is $\sqrt{4} = 2$. Calling $8/3$ the unbiased sample variance would be incorrect.

9.4 Maximum likelihood from first principles

Likelihood views the observed data as fixed and asks which parameter makes them most probable under the model. For independent Bernoulli observations with k successes in n trials, likelihood is proportional to $p^k(1-p)^{n-k}$. The logarithm converts products to sums without changing the maximiser because it is strictly increasing.

Worked problem 9.4 Estimate a success probability

Problem. Eight independent Bernoulli trials contain three successes. Find the maximum-likelihood estimate of p.

Understanding the method. Maximise the logarithm of the likelihood. Here $0 < k < n$, so an interior solution exists. Boundary cases $k = 0$ or $k = n$ must be treated at the endpoints.

Calculation.

$$\begin{aligned} L(p) &= p^3(1-p)^5, \\ \ell(p) &= 3\ln p + 5\ln(1-p), \\ \ell'(p) &= \frac{3}{p} - \frac{5}{1-p} = 0, \\ 3(1-p) = 5p \quad &\Rightarrow \quad p = \boxed{3/8}. \end{aligned}$$

Also $\ell''(p) = -3/p^2 - 5/(1-p)^2 < 0$ for $0 < p < 1$, so the stationary point maximises log likelihood.

Check and meaning. The estimate equals the observed success fraction. This conclusion relies on the Bernoulli model with a shared probability across trials.

9.5 Simple linear regression

The line $\hat{y} = a + bx$ estimates a response using an input. Ordinary least squares minimises $\sum_i (y_i - a - bx_i)^2$. Differentiation gives

$$b = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sum_i (x_i - \bar{x})^2}, \qquad a = \bar{y} - b\bar{x},$$

provided the x_i are not all identical. The numerator measures joint variation; the denominator measures input variation. The formula is equivalent to the normal equations developed in Chapter 5.

Worked problem 9.5 Compute a regression line from centred sums

Problem. Fit a line to $(1, 2), (2, 3), (3, 5)$.

Understanding the method. Centre the inputs and responses, compute the two sums, then recover the intercept.

Calculation.

$$\bar{x} = 2, \qquad \bar{y} = 10/3.$$

The deviations are $x - \bar{x} = (-1, 0, 1)$ and $y - \bar{y} = (-4/3, -1/3, 5/3)$. Hence

$$\begin{aligned} S_{xx} &= (-1)^2 + 0^2 + 1^2 = 2, \\ S_{xy} &= (-1)(-4/3) + 0(-1/3) + 1(5/3) = 3, \\ b &= 3/2, \\ a &= 10/3 - (3/2)(2) = 1/3. \end{aligned}$$

Thus $\boxed{\hat{y} = 1/3 + (3/2)x}$. Predictions are $(11/6, 10/3, 29/6)$ and residuals $(1/6, -1/3, 1/6)$, whose sum is zero.

Check and meaning. The fitted line passes through $(\bar{x}, \bar{y})$, a property of ordinary least squares with an unconstrained intercept.

9.6 Intervals and interpretation

A confidence interval is produced by a procedure designed to cover a fixed parameter at a specified long-run rate under its assumptions. A prediction interval targets a new observation and usually includes more variability. Neither is obtained by attaching a percentage to an arbitrary range.

For normally distributed independent observations with known population standard deviation, a mean interval has form $\bar{x} \pm z_{1-\alpha/2}\sigma/\sqrt{n}$. When σ is estimated, a suitable t-based procedure can be used under its assumptions. The distinction between known and estimated spread matters.

Worked problem 9.6 A mean interval with a supplied critical value

Problem. Suppose $\bar{x} = 20$, known $\sigma = 4$, $n = 16$, and a normal-model 95% interval uses $z = 1.96$. Compute it.

Understanding the method. The margin is the critical value times the standard error, not times the observation standard deviation.

Calculation.

$$\begin{aligned} \text{SE} &= 4/\sqrt{16} = 1, \\ \text{margin} &= 1.96(1) = 1.96, \\ \text{interval} &= [20 - 1.96, 20 + 1.96] = \boxed{[18.04, 21.96]}. \end{aligned}$$

Check and meaning. This is an interval for the population mean under the stated model, not a range intended to contain 95% of individual observations.

Chapter 10

Algorithmic growth and recurrences

The idea before the calculation

An algorithm can be mathematically correct and still be unusable at the scale we need. Complexity studies how its resource requirements grow as the input grows. Doubling an input can mean twice the work, four times the work, or something far more severe.

Big-O notation expresses an asymptotic upper bound within a chosen cost model. It is not a stopwatch reading, and an upper bound need not be tight. Constants, memory access and implementation still matter for actual inputs. Recurrences make repeated structure visible: the cost of a problem can be written in terms of the smaller problems it creates. Begin by identifying that structure, then solve the recurrence; guessing from the final formula reverses the reasoning.

10.1 What is being counted?

An algorithm is a finite method for transforming inputs into outputs. Complexity describes how required resources grow with input size. Before using a growth class, identify the input-size variable and the elementary operation being counted. Reading one element, comparing two numbers, and adding fixed-size numbers are often treated as constant-time actions in introductory analysis.

An exact operation count may be $3n + 2$. Its asymptotic class is linear because the n term dominates for large n. Constants still matter in implementation, but growth classes let us compare scaling without depending on a particular processor. A faster asymptotic method may be slower on tiny inputs due to overhead.

Big-O is an upper bound: $f(n) = O(g(n))$ if $f(n) \leq cg(n)$ for some positive c and all sufficiently large n. Big-Omega is a lower bound. Theta gives both. Saying an algorithm is $O(n^2)$ does not necessarily say its tight growth is quadratic; a linear function also satisfies that loose upper bound.

Worked problem 10.1 Find a tight class

Problem. Classify $f(n) = 3n^2 + 7n + 5$ for $n \geq 1$.

Understanding the method. Bound lower-order terms by multiples of n^2 to establish an upper bound, and retain the leading positive term for a lower bound.

Calculation.

Since $n \leq n^2$ and $1 \leq n^2$,

$$3n^2 \leq f(n) \leq 3n^2 + 7n^2 + 5n^2 = 15n^2.$$

Therefore $\boxed{f(n) = \Theta(n^2)}$.

Check and meaning. The constants 3 and 15 need not be the best possible; they establish the definition. Merely deleting terms without understanding the bound hides the reasoning.

10.2 Logarithms and repeated halving

The logarithm $\log_2 n$ counts how many times a positive power of two can be halved before reaching one. If $n = 2^k$, then after k halvings the remaining size is $n/2^k = 1$. This explains binary search's logarithmic growth.

Changing the base of a logarithm multiplies it by a constant: $\log_a n = \log_b n / \log_b a$. Thus base changes do not change a Theta class, although they do change exact numerical counts. Typical increasing growth rates are 1, $\log n$, n, $n \log n$, n^2, 2^n and $n!$ for sufficiently large n.

Worked problem 10.2 Halving a search interval

Problem. A search begins with 1024 candidates. How many ideal halvings leave one candidate?

Understanding the method. Solve $1024/2^k = 1$. This asks for a logarithm, not for $1024/2$.

Calculation.

$$2^k = 1024 = 2^{10} \quad \Rightarrow \quad \boxed{k = 10}.$$

Check and meaning. This is the number of ideal size halvings. Exact element-comparison counts depend on midpoint and termination conventions and can include a final candidate comparison.

10.3 Nested loops and sums

For independent full loops of n iterations each, multiply their counts to get n^2. But if the inner bound depends on the outer index, sum that bound. A triangular loop has count

$$\sum_{i=1}^{n} i = \frac{n(n+1)}{2}.$$

The formula follows by adding the series to its reversed order: every pair sums to $n+1$, and there are n pairs in the doubled sum.

Worked problem 10.3 A triangular loop

Problem. For each $i = 1, \dots, n$, an inner instruction runs for $j = 1, \dots, i$. Give its exact count and tight order.

Understanding the method. The inner loop executes i times, so the total is a sum rather than n repetitions of a fixed inner length.

Calculation.

$$\begin{aligned} S &= 1 + 2 + \cdots + n, \\ S &= n + (n-1) + \cdots + 1, \\ 2S &= n(n+1), \\ S &= \boxed{\frac{n(n+1)}{2} = \Theta(n^2)}. \end{aligned}$$

For $n = 4$, the count is $1 + 2 + 3 + 4 = 10$, confirming the formula.

Worked problem 10.4 An inner loop that doubles

Problem. An outer loop runs n times. On each pass an inner variable takes values $1, 2, 4, \dots, 2^k \le n$. Find the order of total work.

Understanding the method. The inner variable doubles instead of increasing by one. Its number of values is $\lfloor \log_2 n \rfloor + 1$.

Calculation.

$$T(n) = n\big(\lfloor \log_2 n \rfloor + 1\big) = \boxed{\Theta(n \log n)}.$$

Check and meaning. Counting n inner iterations would be wrong because most integers are skipped. At $n = 8$ the inner values are $1, 2, 4, 8$, only four iterations.

10.4 Divide and conquer

A recurrence expresses the work on size n in terms of smaller subproblems. In $T(n) = aT(n/b)+f(n)$, there are a subproblems of size n/b, while $f(n)$ counts splitting and combining outside the recursive calls. Include a base case to determine exact values.

For standard equal-sized recurrences, compare $f(n)$ with $n^{\log_b a}$. If they match up to constant factors in the basic Master theorem case, the result gains a logarithmic factor. If the combining work is polynomially smaller, leaves dominate; if polynomially larger and a regularity condition holds, combining work dominates. Do not apply the theorem to recurrences outside its assumptions without checking.

Worked problem 10.5 Derive the merge-sort recurrence

Problem. Solve $T(n) = 2T(n/2) + n$ for powers of two with $T(1) = 1$.

Understanding the method. At each internal level, the number of subproblems doubles and their sizes halve. Their total combining work remains n.

Calculation.

At level j, there are 2^j nodes of size $n/2^j$, so the level cost is

$$2^j \frac{n}{2^j} = n.$$

There are $\log_2 n$ internal levels and n leaves costing one each. Therefore

$$\boxed{T(n) = n \log_2 n + n = \Theta(n \log n)}.$$

For $n = 8$, $T(2) = 4$, $T(4) = 12$ and $T(8) = 32 = 8(3) + 8$.

Worked problem 10.6 One recursive subproblem

Problem. Solve $T(n) = T(n/2) + 1$, $T(1) = 1$, for powers of two.

Understanding the method. Repeatedly expand the recurrence. Only one path is followed, unlike the two-branch recurrence above.

Calculation.

$$\begin{aligned} T(n) &= T(n/2) + 1 \\ &= T(n/4) + 2 \\ &= T(n/2^k) + k. \end{aligned}$$

Choose $k = \log_2 n$ so $n/2^k = 1$. Thus

$$\boxed{T(n) = 1 + \log_2 n = \Theta(\log n)}.$$

Check and meaning. For $n = 16$, values $T(1), T(2), T(4), T(8), T(16)$ are $1, 2, 3, 4, 5$.

10.5 Space and recursion depth

Time counts operations; auxiliary space counts additional storage. A recursive algorithm may keep one stack frame per active call. Binary search can be iterative with $O(1)$ auxiliary space; a straightforward recursive version uses $O(\log n)$ stack space. Merge sort's recursion depth is logarithmic, but conventional array merging also uses linear auxiliary storage. Resource measures must be named explicitly.

Chapter 11

Searching, sorting, heaps and trees

The idea before the calculation

A useful data structure stores more than values: it maintains a relationship that makes certain questions easier to answer. Sorted order permits binary search. A heap keeps a priority extreme accessible without sorting everything. A search tree offers paths through ordered choices, though a badly unbalanced tree can become a long chain.

The central question is what remains true after every operation. This invariant is the quiet reason an algorithm works. Before tracing code or counting comparisons, name it. Also ask which operations matter most: frequent insertion, fast lookup, ordered traversal or retrieving the smallest item. There is no universally best structure; there is a structure whose promises fit the task.

11.1 Searching and its invariant

Linear search scans entries until it finds a target or reaches the end. Its worst-case comparison count is n. Binary search exploits sorted order: if the middle value is smaller than the target, every value to its left can be excluded. An invariant is a fact maintained throughout execution. Here it is: if the target is present, it lies inside the current interval.

A useful inclusive-index convention starts with $\ell = 0$, $r = n - 1$ and continues while $\ell \leq r$. Choose $m = \ell + \lfloor (r - \ell)/2 \rfloor$. If $a_m < t$, set $\ell = m + 1$; if $a_m > t$, set $r = m - 1$. The updates must exclude the already-tested middle position or termination can fail.

Worked problem 11.1 A successful binary search

Problem. Find 17 in $[2, 5, 8, 13, 17, 21, 30]$ using zero-based indices.

Understanding the method. Record interval endpoints and midpoint on each comparison; this exposes off-by-one mistakes.

Calculation.

ℓ	r	m	a_m	Decision
0	6	3	13	$17 > 13$, set $\ell = 4$
4	6	5	21	$17 < 21$, set $r = 4$
4	4	4	17	found

The answer is $\boxed{\text{index } 4}$ after three element comparisons.

Check and meaning. Index four is the fifth entry. Confusing a position counted from one with a zero-based index can make a correct trace appear wrong.

Worked problem 11.2 An unsuccessful binary search

Problem. Search for 6 in $[2, 5, 8, 13, 17, 21, 30]$.

Understanding the method. A search can terminate with an empty interval. The invariant then establishes absence.

Calculation.

Initially $\ell = 0, r = 6, m = 3$, and $13 > 6$, so $r = 2$. Next $m = 1$ and $5 < 6$, so $\ell = 2$. Now $m = 2$ and $8 > 6$, so $r = 1$. The endpoints satisfy $\ell = 2 > r = 1$, meaning the interval is empty. Thus $\boxed{6 \text{ is absent}}$.

Check and meaning. The search did not need to compare the target to every entry. Sorted order justified each elimination.

11.2 Insertion sort

Insertion sort maintains a sorted prefix. It takes the next item and shifts larger prefix items right until the insertion position is available. On already sorted data it uses linear comparisons; on reverse-sorted data, the total shifts are $1 + 2 + \cdots + (n - 1) = n(n - 1)/2$, giving quadratic worst-case time. A standard implementation needs constant auxiliary array space.

Worked problem 11.3 Trace insertion sort completely

Problem. Sort $[4, 2, 3, 1]$ by insertion sort and count shifts.

Understanding the method. A shift moves an existing prefix item one place right. Treat the saved key separately so it is not overwritten.

Calculation.

Insert key 2: shift 4, obtaining $[2, 4, 3, 1]$; shifts $= 1$.

Insert key 3: shift 4, then insert after 2, obtaining $[2, 3, 4, 1]$; shifts $= 1$.

Insert key 1: shift 4, then 3, then 2, obtaining $[1, 2, 3, 4]$; shifts $= 3$.

$$\boxed{\text{Total shifts} = 1 + 1 + 3 = 5}.$$

Check and meaning. Shift count is not identical to comparison count; the latter depends on whether unsuccessful terminating comparisons are counted.

11.3 Merge sort and merging

Merge sort divides the input, sorts both halves, and merges the sorted halves. Merging compares the next unconsumed entry from each half and outputs the smaller. Once one half is empty, the remaining entries are copied. For two lists of lengths a and b, the worst-case number of comparisons is $a + b - 1$ when both are nonempty.

The recurrence $T(n) = 2T(n/2) + \Theta(n)$ yields $\Theta(n \log n)$ worst-case time. Conventional array merging uses $\Theta(n)$ auxiliary storage. A stable merge chooses the left item first when keys are equal, preserving relative order of equal-valued records.

Worked problem 11.4 A merge and exact comparisons

Problem. Merge $[1, 4, 7]$ and $[2, 3, 8]$, counting comparisons.

Understanding the method. Compare only the two current heads. Advance the pointer from which the smaller value was taken.

Calculation.

The head comparisons are

$$1 < 2, \quad 4 > 2, \quad 4 > 3, \quad 4 < 8, \quad 7 < 8.$$

They output $1, 2, 3, 4, 7$. The left list is now empty; append 8 without another head comparison. Therefore

$$\boxed{[1, 2, 3, 4, 7, 8], \quad 5 \text{ comparisons}}.$$

Check and meaning. The maximum for lengths three and three is $3 + 3 - 1 = 5$, attained here.

11.4 Quicksort and bad partitions

Quicksort chooses a pivot and partitions elements around it. Balanced partitions yield $\Theta(n \log n)$ growth. Repeatedly placing only one element on one side gives $T(n) = T(n - 1) + \Theta(n)$, and the sum $n + (n - 1) + \cdots + 1$ is quadratic. Randomisation gives good expected performance under the algorithm's model, but ordinary quicksort's worst case remains quadratic.

Worked problem 11.5 A worst-case comparison sum

Problem. A quicksort variant compares a pivot to every other item and repeatedly produces subproblem sizes $n - 1$ and 0. Find comparisons for $n = 5$.

Understanding the method. The nontrivial partitions compare the pivot with $4, 3, 2, 1$ other entries.

Calculation.

$$C(5) = 4 + 3 + 2 + 1 = \boxed{10}.$$

In general,

$$C(n) = \sum_{k=1}^{n-1} k = \frac{n(n-1)}{2} = \Theta(n^2).$$

Check and meaning. The result describes the specified partition-comparison model. Other implementation details can alter exact counts without changing the worst-case growth class.

11.5 Binary search trees and heaps

A binary search tree orders keys so all keys in a node's left subtree are smaller and all in the right subtree are larger, under a chosen duplicate convention. Searching follows one root-to-leaf path. A balanced tree has logarithmic height; a chain-shaped tree can have linear height.

A binary max-heap instead guarantees each parent is at least as large as its children. It does not globally sort all nodes. In a zero-based array, children of index i are $2i + 1$ and $2i + 2$ when present. The maximum is at the root. Extracting it moves the final item to the root and sifts down, taking $O(\log n)$ time.

Bottom-up heap construction is $O(n)$, not the loose $O(n \log n)$ bound obtained by assigning maximum depth to every node. Most nodes are leaves or near leaves. The work can be bounded by a sum proportional to $n \sum_{h \geq 0} h/2^h$, which is linear because that series converges.

Worked problem 11.6 Verify a max-heap

Problem. Does array $[9, 7, 8, 3, 2, 5]$ satisfy the max-heap property?

Understanding the method. Check parents against their existing children using array indices. Leaves require no child checks.

Calculation.

Index 0 has value 9, with children 7 and 8; both are smaller. Index 1 has value 7, with children 3 and 2. Index 2 has value 8, with the sole existing child 5. All inequalities hold, so $\boxed{\text{yes}}$.

Check and meaning. The array is not sorted, since $7 < 8$ occurs. Heap order and sorted order are different properties.

Chapter 12

Learning tasks, losses and fitted models

The idea before the calculation

Learning from data begins with a modest ambition: use observed examples to construct a rule that may work on examples not yet seen. Inputs, targets, model family and loss turn that ambition into a mathematical problem. Each choice decides which similarities the learner notices and which errors it treats as costly.

A smaller training loss establishes a better fit to those training observations under that loss. It does not by itself establish better future performance. That further claim needs appropriate evaluation and assumptions about how future data relate to the past. A model is an organised simplification. Its usefulness depends on what it preserves, what it leaves out, and the situation in which we ask it to predict.

12.1 What is learned?

Supervised learning uses input–target pairs (x_i, y_i) to fit a rule for predicting targets from inputs. Regression predicts numerical values; classification predicts categories or their probabilities. Unsupervised learning seeks structure without supplied target labels, as in clustering. The task must be defined before choosing an algorithm.

A model family gives possible prediction rules. A parameter is fitted from data, such as a line's slope. A hyperparameter controls the fitting procedure or model family, such as a regularisation strength. Training data fit parameters; validation data select modelling choices; a final test set assesses the selected procedure. Reusing the test set repeatedly for selection compromises its role as an independent assessment.

In a forecasting example, "predict gas production" is incomplete. Specify the prediction horizon, output units, observations available at that time, and whether the test represents a new time period, a new plant, or both. A physically impossible information pathway can make excellent-looking scores meaningless.

12.2 Squared and absolute losses

For residual $e_i = \hat{y}_i - y_i$, mean squared error is $n^{-1}\sum e_i^2$, and mean absolute error is $n^{-1}\sum |e_i|$. Squaring penalises large residuals more strongly; absolute loss grows linearly. RMSE is the square root of MSE and returns to target units. The denominator and sign convention must be clear, although squared and absolute losses do not depend on the residual sign convention.

Worked problem 12.1 Calculate three regression metrics

Problem. Observed outputs are $(2, 4, 6)$ and predictions are $(3, 3, 4)$. Find MSE, RMSE and MAE.

Understanding the method. Make a residual table before calculating aggregate metrics. This prevents omission of an observation.

Calculation.

y	$\hat{y}$	e	e^2	$\lvert e\rvert$
2	3	1	1	1
4	3	−1	1	1
6	4	−2	4	2

$$\begin{aligned} \text{MSE} &= (1+1+4)/3 = \boxed{2}, \\ \text{RMSE} &= \boxed{\sqrt{2} \approx 1.414}, \\ \text{MAE} &= (1+1+2)/3 = \boxed{4/3}. \end{aligned}$$

Check and meaning. MSE has squared output units; RMSE and MAE have output units. Neither MSE nor MAE can be negative.

12.3 Empirical risk and fitting

Training minimises an average loss on observed data, called empirical risk. A small training loss does not guarantee small expected loss on future data. A sufficiently flexible model can memorise a finite dataset, including noise. Generalisation requires evaluation on data not used to fit or select the final model.

For a constant model under squared loss, differentiation shows the fitted constant is the mean. Under absolute loss, a median is a minimiser. These facts explain how loss functions encode the meaning of a "best" prediction rather than merely changing numerical reporting.

Worked problem 12.2 Squared versus absolute loss for a constant

Problem. For responses $1, 2, 9$, find a squared-loss-minimising constant and an absolute-loss-minimising constant.

Understanding the method. The squared loss is differentiable. Absolute loss is piecewise linear, and its minimiser is a median.

Calculation.

$$\begin{aligned} S(c) &= (c-1)^2 + (c-2)^2 + (c-9)^2, \\ S'(c) &= 2(c-1) + 2(c-2) + 2(c-9) = 6c - 24, \\ S'(c) = 0 &\Rightarrow \boxed{c = 4}. \end{aligned}$$

The sorted data have median 2, so the absolute-loss minimiser is $\boxed{c = 2}$. Its absolute loss is $1 + 0 + 7 = 8$, while at 4 it is $3 + 2 + 5 = 10$.

Check and meaning. The large response 9 pulls the mean farther than the median. Both answers are correct for their respective loss functions.

12.4 Binary classification and the sigmoid

A linear score $z = w^T x + b$ can be transformed into a probability-like output by the sigmoid

$$\sigma(z) = \frac{1}{1 + e^{-z}}.$$

It lies between zero and one, equals $1/2$ at zero, and has derivative $\sigma(z)(1-\sigma(z))$. A threshold turns a score into a category; the threshold is a decision rule and can be varied according to costs. A score should not be treated as a calibrated probability without assessment.

For a Bernoulli target $y \in \{0, 1\}$ and predicted probability p, negative log likelihood is

$$\ell(p, y) = -\{y \ln p + (1 - y) \ln(1 - p)\}.$$

It penalises confidently wrong predictions strongly. A reported zero loss for finite non-extreme probabilities would be suspect.

Worked problem 12.3 Sigmoid and a classification threshold

Problem. A model score is $z = \ln 3$. Find its sigmoid output and the prediction under threshold 0.5.

Understanding the method. Use the inverse relation between logarithm and exponential: $e^{-\ln 3} = 1/3$.

Calculation.

$$\sigma(\ln 3) = \frac{1}{1 + 1/3} = \frac{3}{4} = \boxed{0.75}.$$

Since $0.75 \geq 0.5$, the threshold rule predicts the positive class.

Check and meaning. The positive score gives an output above one-half. The output is not $\ln 3$ itself; the score and probability are different quantities.

Worked problem 12.4 Log loss for one observation

Problem. A positive observation receives predicted probability $p = 0.8$. Calculate its log loss. Compare with $p = 0.2$.

Understanding the method. For $y = 1$, the $(1 - y)$ term vanishes, leaving $-\ln p$.

Calculation.

$$\ell(0.8, 1) = -\ln 0.8 \approx \boxed{0.2231},$$
$$\ell(0.2, 1) = -\ln 0.2 \approx \boxed{1.6094}.$$

Check and meaning. Assigning less probability to the event that actually occurred gives a larger loss. Natural logarithms are used; changing the log base rescales the loss.

12.5 Overfitting and underfitting

Underfitting means the model or procedure cannot capture relevant structure even in training. Overfitting means strong performance on training data fails to carry over. A large training–validation gap suggests overfitting, but a changed validation distribution or flawed split may also cause the gap. Diagnose before prescribing a remedy.

A simple baseline is essential. A constant mean predictor, majority classifier or last-value forecast provides a reference. A complex model's value comes from improvement on a relevant held-out task, not from complexity alone. Repeatedly selecting the best model on the same small validation set can itself overfit the validation process.

Worked problem 12.5 A majority-class baseline

Problem. A test set has 95 negatives and 5 positives. A classifier always predicts negative. Find accuracy and positive recall.

Understanding the method. Fill the confusion counts before using formulas. A missed positive is a false negative.

Calculation.

$$TN = 95, \quad FN = 5, \quad TP = 0, \quad FP = 0.$$

Then

$$\text{accuracy} = \frac{TP + TN}{100} = \boxed{95\%}, \qquad \text{recall} = \frac{TP}{TP + FN} = \boxed{0}.$$

Check and meaning. The high accuracy hides complete failure to detect positives. Precision is undefined here because no positive predictions were made; a software convention may replace that undefined value with zero.

12.6 Data preparation is part of the model

Scaling, imputation, feature selection and dimensionality reduction learn from data. Fit them inside the training portion of each split. Applying a scaler estimated from the entire dataset lets test-distribution information influence training. Temporal features need an additional check: were all their ingredients available when the prediction would have been issued? A rolling mean that includes future values is leakage even if the model itself was fitted only on earlier dates.

12.7 Derive the sigmoid derivative

The sigmoid derivative is often quoted as a formula. It follows directly from the chain rule and makes sense geometrically: the curve is steepest near its centre and flattens near its limiting values.

Worked problem 12.6 Differentiate a probability transformation

Problem. Derive $\sigma'(z)$ for $\sigma(z) = (1 + e^{-z})^{-1}$ and evaluate it at zero.

Understanding the method. Differentiate the outer reciprocal power and then the inner exponential. There are two minus signs, which cancel.

Calculation.

$$\begin{aligned}
\sigma'(z) &= -(1 + e^{-z})^{-2}(-e^{-z}) \\
&= \frac{e^{-z}}{(1 + e^{-z})^2} \\
&= \frac{1}{1 + e^{-z}} \left(\frac{e^{-z}}{1 + e^{-z}} \right) \\
&= \sigma(z)(1 - \sigma(z)).
\end{aligned}$$

Since $\sigma(0) = 1/2$, $\sigma'(0) = (1/2)(1/2) = \boxed{1/4}$.

Check and meaning. The derivative is positive, consistent with an increasing curve. A negative result would reveal a lost minus sign.

Chapter 13

Optimisation, regularisation and evaluation

The idea before the calculation

Optimisation asks how to improve an objective; evaluation asks whether the improvement matters beyond the data used to obtain it. These are connected questions with different answers. A model may fit the training set more closely while becoming less dependable elsewhere.

Regularisation deliberately restrains flexibility, often trading some training fit for greater stability. Its strength depends on feature scale and on the precise objective, so a penalty parameter has no universal meaning in isolation. Validation helps choose such settings; the test set should remain separate from those choices. Repeatedly consulting it turns it into another development set. Even a careful test measures performance under its sampling conditions, not under every possible future change.

13.1 Why gradient descent moves downhill

A differentiable loss $L(w)$ has local approximation $L(w+h) \approx L(w) + \nabla L(w)^T h$. Choosing $h = -\eta \nabla L(w)$ makes the linear term $-\eta \|\nabla L(w)\|^2$, nonpositive when $\eta > 0$. This is the local reason for moving against the gradient. It does not guarantee improvement for an arbitrarily large step, because curvature and higher-order terms then matter.

The update is $w_{k+1} = w_k - \eta \nabla L(w_k)$. The gradient must be evaluated at the current iterate. A learning rate is a step-size control, not a probability. Small rates may converge slowly; large ones can oscillate or diverge.

Worked problem 13.1 Two full gradient updates

Problem. For $L(w) = (w-3)^2$, start at $w_0 = 0$ and use $\eta = 0.1$. Compute two updates and all three losses.

Understanding the method. Differentiate once symbolically, then substitute the current parameter value at every update.

Calculation.

$$\begin{aligned} L'(w) &= 2(w-3), \\ L'(w_0) &= 2(0-3) = -6, \\ w_1 &= 0 - 0.1(-6) = 0.6, \\ L'(w_1) &= 2(0.6-3) = -4.8, \\ w_2 &= 0.6 - 0.1(-4.8) = 1.08. \end{aligned}$$

The losses are

$$L(0) = 9, \qquad L(0.6) = 2.4^2 = 5.76, \qquad L(1.08) = 1.92^2 = \boxed{3.6864}.$$

Check and meaning. Both updates move closer to the minimiser $w = 3$, and the loss decreases. Keeping the original gradient -6 for the second step would be a different, incorrect update.

13.2 A stability condition for a quadratic

Consider $L(w) = \frac{1}{2}aw^2$ with $a > 0$. Its gradient is aw, and the update becomes $w_{k+1} = (1 - \eta a)w_k$. Repeated updates shrink toward zero exactly when $|1 - \eta a| < 1$, giving $0 < \eta < 2/a$. If the multiplier is negative but its magnitude is below one, iterates alternate sign while shrinking. If its magnitude exceeds one, they diverge unless started exactly at zero.

For a positive-definite quadratic in several variables, each eigenvector direction has its own curvature eigenvalue. Stability therefore requires a rate small enough for the largest curvature, in this simple fixed-step quadratic setting. This connects optimisation to Chapter 6.

Worked problem 13.2 A learning rate that diverges

Problem. For $L(w) = \frac{1}{2}(4)w^2$ and $w_0 = 1$, compare $\eta = 0.2$ and $\eta = 0.6$.

Understanding the method. The multiplier is $1 - 4\eta$. Its magnitude determines convergence of this exact scalar recurrence.

Calculation.

For $\eta = 0.2$,

$$1 - 4\eta = 0.2, \quad w_1 = 0.2, \quad w_2 = 0.04,$$

which approaches zero. For $\eta = 0.6$,

$$1 - 4\eta = -1.4, \quad w_1 = -1.4, \quad w_2 = 1.96,$$

whose magnitude grows. The stability interval is $\boxed{0 < \eta < 0.5}$.

Check and meaning. A negative multiplier alone does not imply divergence. It is the magnitude above one that makes the second sequence grow.

13.3 Regularisation changes the objective

Regularisation adds a preference to the fitting criterion. Ridge uses a nonnegative coefficient multiplying a squared norm. For a scalar,

$$J(w) = (w - y)^2 + \lambda w^2.$$

Differentiating gives $2(w - y) + 2\lambda w = 0$, hence $\hat{w} = y/(1 + \lambda)$. As λ increases, this estimate shrinks toward zero. In multiple dimensions, scaling affects the meaning of equal penalties on coefficients, and an intercept is commonly left unpenalised.

A ridge penalty can reduce variance and improve generalisation, but it introduces bias. It is not a universal guarantee of improvement. Choose its strength using validation data appropriate to the task.

Worked problem 13.3 Fit with and without a penalty

Problem. Minimise $J(w) = (w - 2)^2 + 0.5w^2$. Compare with the unpenalised minimiser.

Understanding the method. Differentiate both the data-fit term and the penalty. A penalty influences the optimum, not merely the reported score.

Calculation.

$$\begin{aligned} J'(w) &= 2(w - 2) + 2(0.5)w = 3w - 4, \\ J'(w) = 0 &\Rightarrow \boxed{\hat{w} = 4/3}, \\ J''(w) &= 3 > 0. \end{aligned}$$

Without the penalty, $2(w - 2) = 0$ gives $w = 2$. At the penalised optimum,

$$J(4/3) = (-2/3)^2 + \tfrac{1}{2}(4/3)^2 = 4/9 + 8/9 = \boxed{4/3}.$$

Check and meaning. The fitted value is smaller in magnitude than 2. Comparing only the data-fit term would miss the purpose of the changed objective.

13.4 Confusion counts and metrics

For a binary classifier, true positives are actual positives predicted positive, false positives are actual negatives predicted positive, and false negatives are actual positives predicted negative. Precision asks what fraction of positive predictions are correct; recall asks what fraction of actual positives are detected:

$$P = \frac{TP}{TP + FP}, \quad R = \frac{TP}{TP + FN}, \quad F_1 = \frac{2PR}{P + R} = \frac{2TP}{2TP + FP + FN}.$$

The formulas require meaningful denominators. Accuracy uses both classes and is $(TP + TN)/N$. Changing a threshold changes these counts; a lower threshold usually finds more positives while creating more false alarms, though exact changes depend on the scores.

Worked problem 13.4 Metrics from a detector

Problem. A detector has $TP = 30$, $FP = 10$, $FN = 20$, and $TN = 40$. Find precision, recall, F1 and accuracy.

Understanding the method. Construct each denominator from the population that the metric describes.

Calculation.

$$\begin{aligned} P &= 30/(30 + 10) = 0.75, \\ R &= 30/(30 + 20) = 0.60, \\ F_1 &= \frac{2(0.75)(0.60)}{0.75 + 0.60} = \frac{0.9}{1.35} = \boxed{2/3}, \\ N &= 30 + 10 + 20 + 40 = 100, \\ \text{accuracy} &= (30 + 40)/100 = \boxed{0.70}. \end{aligned}$$

Check and meaning. The alternative F1 formula gives $60/(60 + 10 + 20) = 2/3$. Accuracy is not the arithmetic mean of precision and recall.

13.5 Time-series splitting and uncertainty

Forecast evaluation must respect time. Training on later observations and evaluating on earlier ones does not reproduce prospective prediction. Even chronological splitting can leak if a feature uses future measurements or a scaler is fitted on the entire timeline. A gap may be necessary when windows or targets overlap.

A prediction interval should be assessed through coverage and width. Empirical coverage is the fraction of held-out observations lying within their corresponding intervals. A very wide interval can cover almost everything while being uninformative, so coverage alone is

insufficient. A confidence interval for a mean and a prediction interval for a new observation answer different questions.

Worked problem 13.5 Evaluate interval coverage

Problem. Five observations are $[2, 4, 6, 8, 10]$. Their intervals are $[1, 3]$, $[3, 5]$, $[4, 5]$, $[7, 9]$ and $[9, 11]$. Find coverage and mean width.

Understanding the method. Check membership separately for each observation and interval, then average the binary coverage indicators and widths.

Calculation.

The membership indicators are $(1, 1, 0, 1, 1)$ because 6 lies outside $[4, 5]$. Thus

$$\text{coverage} = 4/5 = \boxed{80\%}.$$

Widths are $(2, 2, 1, 2, 2)$, so mean width is

$$(2 + 2 + 1 + 2 + 2)/5 = \boxed{1.8}.$$

Check and meaning. This is empirical coverage on five cases, not proof of an 80% coverage guarantee on future data. The small sample itself has substantial uncertainty.

13.6 Batch and stochastic gradients

For an average loss $L(w) = n^{-1} \sum_i \ell_i(w)$, the full gradient averages the individual gradients. A stochastic update uses one sampled example; a mini-batch averages a subset. Random sampling can make the stochastic gradient an unbiased estimate of the full gradient under the sampling scheme, but a single stochastic step need not decrease the full loss.

Worked problem 13.6 A full gradient versus one example

Problem. Two targets are 0 and 4, and a constant model predicts w. Let $L(w) = \frac{1}{2}[w^2 + (w - 4)^2]$. At $w = 1$, compare the full gradient with each single-example gradient.

Understanding the method. Differentiate each per-example squared loss, then average them. Keep the factor $1/2$ in the objective.

Calculation.

$$\begin{aligned} \ell_1'(1) &= 2(1 - 0) = 2, \\ \ell_2'(1) &= 2(1 - 4) = -6, \\ L'(1) &= \tfrac{1}{2}(2 - 6) = \boxed{-2}. \end{aligned}$$

With learning rate 0.1, the full update gives $1 - 0.1(-2) = 1.2$. A sampled first example gives 0.8; a sampled second gives 1.6. Their equally weighted average is 1.2.

Check and meaning. The full optimum is the target mean 2. One individual example can push the parameter away from that optimum even though the average gradient points toward it.

Chapter 14

Mechanics, structures and materials

The idea before the calculation

A bridge, a spring and a falling object each invite a different level of description. Mechanics begins by choosing the body or system, drawing the forces and deciding which effects can be neglected. A free-body diagram is therefore an argument about the model, not decoration beside it.

Force, stress and strain answer different questions: applied interaction, force per area, and relative deformation. Linear elasticity is an approximation within a material's appropriate operating range. Units, directions and boundary conditions are part of each equation's meaning. A numerically tidy result can still describe the wrong physical situation. The calculations here develop foundational reasoning; actual structural design requires the relevant standards, material data and professional assessment.

14.1 Units are part of the calculation

Force, energy, pressure and power are different quantities. In SI, force is measured in newtons, $1\,\mathrm{N} = 1\,\mathrm{kg\,m\,s^{-2}}$. Work and energy use joules, $1\,\mathrm{J} = 1\,\mathrm{N\,m}$. Power is energy per time, $1\,\mathrm{W} = 1\,\mathrm{J\,s^{-1}}$. Stress and pressure use pascals, $1\,\mathrm{Pa} = 1\,\mathrm{N\,m^{-2}}$.

Convert units before combining numbers. Because $1\,\mathrm{mm} = 10^{-3}\,\mathrm{m}$, an area of one square millimetre is $10^{-6}\,\mathrm{m^2}$. Therefore $1\,\mathrm{N/mm^2} = 10^{6}\,\mathrm{Pa} = 1\,\mathrm{MPa}$. Squared and cubed conversion factors are a common source of large errors.

14.2 Force and acceleration

Newton's second law is $\sum F = ma$: it uses the net force, not a selected individual force. A free-body diagram isolates the object and shows external forces. Weight is mg downward;

a support can exert a normal reaction; friction acts tangentially and opposes the relevant relative sliding tendency under its model.

Static equilibrium requires zero net force and zero net moment. A body at constant velocity also has zero net force, even though it is moving. Zero velocity at one instant does not by itself imply zero acceleration.

Worked problem 14.1 Net force with opposing forces

Problem. A 5 kg body has a 30 N horizontal pull to the right and 10 N resistance to the left. Find its acceleration.

Understanding the method. Choose right as positive and add signed forces before dividing by mass.

Calculation.

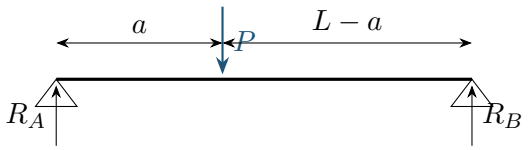
$$F_{\text{net}} = 30 - 10 = 20\,\text{N},$$
$$a = F_{\text{net}}/m = 20/5 = \boxed{4\,\text{m/s}^2 \text{ right}}.$$

Check and meaning. Dividing the pull alone by mass would give $6\,\text{m/s}^2$ and ignore the stated resistance.

14.3 Moments and support reactions

The magnitude of a moment is force times perpendicular distance from the reference point to the force's line of action. The sign depends on the chosen rotation convention. Taking moments about a support eliminates its own reaction from that equation, simplifying beam problems.

For a simply supported beam with a downward point load, vertical force balance says the reactions sum to the load. Moment balance determines how that total is shared. The reactions are not necessarily equal; symmetry is required for an equal split.



Worked problem 14.2 An off-centre beam load

Problem. A simply supported beam is 4 m long. A downward 12 kN load acts 1 m from the left support. Neglect beam weight. Find the vertical reactions.

Understanding the method. Take counterclockwise moment about the left support as positive. The right reaction has a 4 m arm; the load has a 1 m arm.

Calculation.

$$\begin{aligned}\sum M_A = 0: \quad & 4R_B - 12(1) = 0,\\ R_B = 3\,\text{kN},&\\ \sum F_y = 0: \quad & R_A + R_B - 12 = 0,\\ R_A = 12 - 3 = &\boxed{9\,\text{kN}}.\end{aligned}$$

Thus $\boxed{R_A = 9\,\text{kN},\ R_B = 3\,\text{kN}}$ upward.

Check and meaning. The closer support carries more of the load. The reactions sum to 12 kN and satisfy the moment equation.

14.4 Stress, strain and elastic response

Average axial normal stress is $\sigma = F/A$ for a member under an appropriately distributed axial force. Engineering strain is $\varepsilon = \Delta L/L$, a dimensionless ratio. Within a linear elastic regime, Hooke's law is $\sigma = E\varepsilon$, where E is Young's modulus. Combining the relations gives extension $\Delta L = FL/(AE)$ for a uniform bar under the stated assumptions.

Stiffness and strength are different. A material's modulus describes elastic resistance to strain; yield strength marks the onset of significant plastic behaviour under a specified test. Toughness concerns energy absorption before fracture. A high modulus alone does not establish high failure load. Geometry matters as well as material properties.

Worked problem 14.3 Stress, strain and inferred modulus

Problem. A bar of length 2 m extends by 1 mm under 20 kN. Its area is 500 mm^2. Find stress and strain, and infer E if response is linear elastic.

Understanding the method. Use consistent units. Stress can conveniently be calculated in N/mm^2, while the length ratio must use matching length units.

Calculation.

$$\begin{aligned}F &= 20{,}000\,\text{N},\\ \sigma &= 20{,}000/500 = 40\,\text{N/mm}^2 = 40\,\text{MPa},\\ L &= 2000\,\text{mm},\\ \varepsilon &= 1/2000 = 0.0005,\\ E &= 40/0.0005 = 80{,}000\,\text{MPa} = \boxed{80\,\text{GPa}}.\end{aligned}$$

Check and meaning. The modulus inference needs the elastic assumption. The observed load and extension alone do not establish that the material is within its elastic range.

Worked problem 14.4 Predict an axial extension

Problem. A uniform bar has $F = 10$ kN, $L = 1$ m, $A = 100\,\text{mm}^2$, and $E = 200$ GPa. Find its elastic extension.

Understanding the method. Use $\Delta L = FL/(AE)$ after converting all quantities to one consistent unit system.

Calculation.

Using N and mm,

$$F = 10{,}000\,\text{N}, \qquad L = 1000\,\text{mm}, \qquad E = 200{,}000\,\text{N/mm}^2,$$

$$\Delta L = \frac{10{,}000(1000)}{100(200{,}000)} = \boxed{0.5\,\text{mm}}.$$

Check and meaning. Stress is 100 MPa, strain is $100/200{,}000 = 0.0005$, and multiplying strain by 1000 mm gives the same 0.5 mm.

14.5 Work, kinetic energy and power

For a constant force parallel to displacement, work is $W = Fs$. More generally it is the integral of force dotted with displacement. Kinetic energy is $K = mv^2/2$, and net work changes kinetic energy. Power is the rate of doing work. Torque and work share dimensions N m, but torque represents a rotational effect and is not itself energy transferred without angular displacement.

Worked problem 14.5 Work and final speed

Problem. A net constant force of 10 N acts over 5 m on a 2 kg body initially at rest. Find its final speed, neglecting other energy changes.

Understanding the method. The net work changes kinetic energy. Because the initial speed is zero, initial kinetic energy vanishes.

Calculation.

$$W_{\text{net}} = Fs = 10(5) = 50\,\text{J},$$

$$\frac{1}{2}mv^2 - 0 = 50,$$

$$\frac{1}{2}(2)v^2 = 50 \quad \Rightarrow \quad v = \boxed{\sqrt{50} \approx 7.071\,\text{m/s}}.$$

Check and meaning. From dynamics, $a = 10/2 = 5$ and $v^2 = 2as = 50$, agreeing with the energy method.

14.6 Hydrostatic pressure

For a fluid of approximately constant density at rest, gauge pressure at depth h is $p = \rho g h$. Absolute pressure includes the pressure at the free surface. Gauge and absolute readings must not be confused, especially in gas-law calculations that require absolute pressure.

Worked problem 14.6 Pressure at depth

Problem. Find water gauge pressure at depth 2 m using $\rho = 1000\,\mathrm{kg/m^3}$ and $g = 9.81\,\mathrm{m/s^2}$.

Understanding the method. Multiply density, gravitational acceleration and vertical depth. The pressure depends on depth, not on container shape under the simple static model.

Calculation.

$$p = 1000(9.81)(2) = 19{,}620\,\mathrm{Pa} = \boxed{19.62\,\mathrm{kPa}}.$$

Check and meaning. Adding atmospheric pressure would produce absolute pressure; it is not required for the gauge pressure asked here.

14.7 Static friction as a bound

A simple dry-friction model distinguishes the maximum available static friction from the actual static friction. While an object remains at rest, friction adjusts to balance the required tangential force up to its maximum $\mu_s N$. It is not always equal to that maximum.

Worked problem 14.7 Will a block move?

Problem. A 10 kg block on a horizontal surface has static-friction coefficient 0.3. A horizontal force of 20 N is applied. Use $g = 9.81\,\mathrm{m/s^2}$. Does it start sliding under this idealised model?

Understanding the method. Calculate the normal reaction and maximum static friction. Compare the required balancing force with this limit.

Calculation.

$$N = mg = 10(9.81) = 98.1\,\mathrm{N},$$
$$f_{s,\max} = \mu_s N = 0.3(98.1) = 29.43\,\mathrm{N}.$$

Since $20 < 29.43$, static friction can balance the applied force. The block remains at rest, with actual friction magnitude $\boxed{20\,\mathrm{N}}$, opposite the pull.

Check and meaning. Using 29.43 N as the actual friction would incorrectly create a net force on a body that can remain stationary.

Chapter 15

Electrical circuits and computing foundations

The idea before the calculation

A circuit diagram compresses a physical arrangement into relationships we can calculate. Kirchhoff's current law expresses charge balance at a node. Kirchhoff's voltage law describes loop voltages within the applicable lumped-circuit model. Neither law removes the need to choose signs consistently and recognise where an ideal approximation becomes inadequate.

Voltage is not power: power depends on both voltage and current, with the sign convention deciding whether a component absorbs or supplies it. Computing introduces another useful abstraction. A logical state is represented by a physical device, but logical operations can be studied without describing every moving charge. Understanding the boundary between these levels helps connect circuit reasoning to the machines that execute algorithms.

15.1 Voltage, current, resistance and power

Current is rate of charge flow, measured in amperes. Voltage is energy transferred per unit charge. For an ideal resistor obeying Ohm's law, $V = IR$. Power is $P = VI$; combining the relations gives $P = I^2R = V^2/R$. These expressions describe dissipated power for a positive resistor under consistent passive sign convention.

A resistor value is not automatically constant in all physical conditions; temperature and nonlinear components can alter behaviour. Introductory calculations assume ideal constant resistances unless stated otherwise. Units provide immediate checks: volts divided by ohms give amperes, and volts times amperes give watts.

Worked problem 15.1 One resistor and energy

Problem. A $6\,\Omega$ resistor is connected across 12 V for 10 s. Find current, power and energy.

Understanding the method. Find current from Ohm's law, power from voltage times current, and energy from power times time.

Calculation.

$$I = V/R = 12/6 = 2\,\text{A},$$
$$P = VI = 12(2) = 24\,\text{W},$$
$$E = Pt = 24(10) = \boxed{240\,\text{J}}.$$

Check and meaning. $V^2/R = 144/6 = 24$ W independently verifies power. Watts and joules cannot be interchanged without a time factor.

15.2 Series and parallel connections

Series resistors carry the same current, and their voltage drops add. Hence $R_{\text{eq}} = \sum R_i$. Parallel branches share voltage, and their currents add. Hence $1/R_{\text{eq}} = \sum 1/R_i$. For two resistors this becomes $R_1R_2/(R_1 + R_2)$.

An equivalent series resistance exceeds each positive member; a parallel equivalent is smaller than the smallest branch resistance. These inequalities are useful sanity checks.

Worked problem 15.2 Series voltage drops and power balance

Problem. A $2\,\Omega$ and a $4\,\Omega$ resistor are in series across 12 V. Find each voltage drop and dissipated power.

Understanding the method. Find the common current from total resistance. Then apply Ohm's law separately to each member.

Calculation.

$$R_{\text{eq}} = 2 + 4 = 6\,\Omega, \qquad I = 12/6 = 2\,\text{A},$$
$$V_1 = IR_1 = 2(2) = 4\,\text{V}, \qquad V_2 = 2(4) = 8\,\text{V},$$
$$P_1 = I^2R_1 = 4(2) = 8\,\text{W}, \qquad P_2 = 4(4) = 16\,\text{W}.$$

The drops sum to 12 V and the powers sum to 24 W, matching source power $12(2) = 24$ W.

Worked problem 15.3 Parallel branches

Problem. A $6\,\Omega$ and a $3\,\Omega$ resistor are in parallel across 12 V. Find branch currents and equivalent resistance.

Understanding the method. Each branch receives the full source voltage. Calculate currents separately before adding.

Calculation.

$$I_1 = 12/6 = 2\,\text{A}, \qquad I_2 = 12/3 = 4\,\text{A},$$
$$I_{\text{total}} = 2 + 4 = 6\,\text{A},$$
$$R_{\text{eq}} = 12/6 = \boxed{2\,\Omega}.$$

Also $1/R_{\text{eq}} = 1/6 + 1/3 = 1/2$. Total power is $12(6) = 72$ W.

Check and meaning. The parallel equivalent $2\,\Omega$ is below both $3\,\Omega$ and $6\,\Omega$, as expected.

15.3 Kirchhoff's laws

Current conservation at a node says incoming current equals outgoing current. Voltage balance around a closed loop says signed voltage changes sum to zero in the basic lumped-circuit model. These laws lead to simultaneous equations, so the linear algebra already developed becomes a circuit-solving tool.

For nodal analysis, choose a reference potential and express branch current as a voltage difference divided by resistance. Consistent sign conventions matter more than guessing each current direction. A negative solved current simply means actual flow is opposite to the assumed direction.

Worked problem 15.4 A node connected to two sources

Problem. A node at unknown voltage v connects to 10 V through $2\,\Omega$, to 4 V through $3\,\Omega$, and to ground through $6\,\Omega$. Find v.

Understanding the method. Assume all three currents leave the node. Their algebraic sum is zero because no charge accumulates there.

Calculation.

$$\frac{v-10}{2} + \frac{v-4}{3} + \frac{v}{6} = 0,$$
$$3(v-10) + 2(v-4) + v = 0,$$
$$6v - 38 = 0 \quad \Rightarrow \quad \boxed{v = 19/3\,\text{V}}.$$

At this voltage the three assumed outward currents are $-11/6$, $7/9$ and $19/18$ A. Their sum is $(-33 + 14 + 19)/18 = 0$.

Check and meaning. The negative first current means current enters the node from the 10 V source. It does not invalidate the chosen sign convention.

15.4 Energy units and efficiency

A kilowatt-hour is power multiplied by time. Since $1\,\text{kW} = 1000\,\text{J/s}$ and one hour is 3600 s, $1\,\text{kWh} = 3.6 \times 10^6$ J. Efficiency is useful output energy divided by input energy, measured on a clearly stated basis. Efficiencies from different stages multiply when the output of one stage is the input of the next and their definitions are compatible.

Worked problem 15.5 Electrical use and useful energy

Problem. A 2 kW device operates for 3 h at 80% efficiency. Find electrical input energy and useful output energy.

Understanding the method. Compute input from rated power and time, then apply the efficiency ratio.

Calculation.

$$E_{\text{in}} = 2(3) = 6\,\text{kWh} = 21.6\,\text{MJ},$$
$$E_{\text{useful}} = 0.8(6) = \boxed{4.8\,\text{kWh}} = 17.28\,\text{MJ}.$$

The difference is 1.2 kWh under this accounting boundary.

Check and meaning. An efficiency below one cannot produce useful output greater than the stated input under the same energy definition.

15.5 Bits, binary numbers and Boolean reasoning

Computing stores information using bits with values zero or one. A binary numeral represents a weighted sum of powers of two. For example, $1011_2 = 1 \cdot 8 + 0 \cdot 4 + 1 \cdot 2 + 1 = 11$. One byte contains eight bits. Memory units may use decimal or binary prefixes; state the convention if an exact size is requested.

Boolean AND is true only if both inputs are true; OR is true if at least one is true; NOT reverses truth. Exclusive OR is true when the inputs differ. De Morgan's laws are $\neg(A \wedge B) = \neg A \vee \neg B$ and $\neg(A \vee B) = \neg A \wedge \neg B$. These are useful in conditions and complement events alike.

Worked problem 15.6 Binary conversion and a condition

Problem. Convert 11010_2 to decimal. Evaluate $(A \wedge \neg B) \vee C$ for $A = 1, B = 1, C = 0$.

Understanding the method. The leftmost bit has the largest power-of-two weight. For the Boolean expression, apply NOT before AND, then OR according to the parentheses.

Calculation.

$$11010_2 = 16 + 8 + 0 + 2 + 0 = \boxed{26}.$$

With the given truth values,

$$(A \wedge \neg B) \vee C = (1 \wedge 0) \vee 0 = 0 \vee 0 = \boxed{0}.$$

Check and meaning. The Boolean symbols are logical operations, not ordinary addition and multiplication in every context. The truth-table definitions determine the result.

Chapter 16

Thermodynamics and process balances

The idea before the calculation

Before writing a balance, draw a boundary. What can cross it: matter, heat, work? The answer decides which terms belong in the equation. A closed system and a flowing process may obey the same conservation principles while requiring different accounts of energy transfer.

Energy conservation does not say every transformation is equally possible. Entropy helps describe the direction and limitations of processes; it is not simply energy that has disappeared. Temperature ratios in thermodynamic relations require an absolute scale. The habit running through these calculations is careful accounting: define the system, state the assumptions, choose consistent units and check what enters and leaves. A predictive model must respect these constraints if it is to describe a physically possible process.

16.1 Systems, boundaries and accounting

A thermodynamic system is the material or region selected for study. A closed system exchanges energy but not mass with its surroundings. A control volume may have material crossing its boundary. Define that boundary before writing a balance: a statement about one vessel is not automatically a statement about an entire plant.

A general accounting structure is accumulation = input − output + generation − consumption. Total mass is conserved in ordinary chemical-process balances, so there is no net generation of total mass. A chemical species can be generated or consumed by reaction. Energy can cross as heat, work or energy carried by flowing material. A steady-state assumption sets accumulation to zero; it does not mean all flows vanish.

16.2 Heat, work and internal energy

For a closed system with negligible bulk kinetic and potential energy changes, use

$$\Delta U = Q - W,$$

where Q is heat into the system and W is work done by the system. Another sign convention may put work into the system as positive; then the equation changes accordingly. The physical accounting must remain consistent.

Heat and work describe transfers, not stored substances. Internal energy is a state property. Under a constant-specific-heat approximation without phase change, sensible heating is $Q = mc\Delta T$. Temperature differences in kelvin and Celsius have the same numerical size; absolute-temperature relations such as the ideal gas law require kelvin.

Worked problem 16.1 A closed-system energy balance

Problem. A closed system receives 150 J of heat while doing 45 J of work. Find its internal-energy change.

Understanding the method. Apply the stated sign convention. Positive heat raises internal energy; positive work done by the system removes energy.

Calculation.

$$\begin{aligned} \Delta U &= Q - W \\ &= 150 - 45 = \boxed{105\,\mathrm{J}}. \end{aligned}$$

Check and meaning. The accounting is $150 = 105 + 45$: received heat is divided between increased storage and work delivered.

Worked problem 16.2 Heating liquid water

Problem. Estimate the heat needed to warm 2 kg of water from 20°C to 25°C, using $c = 4.2\,\mathrm{kJ/(kg\,K)}$. Neglect losses and phase change.

Understanding the method. The temperature rise is 5 K. Multiply mass, specific heat and rise, retaining units.

Calculation.

$$\begin{aligned} \Delta T &= 25 - 20 = 5\,\mathrm{K}, \\ Q &= mc\Delta T \\ &= (2\,\mathrm{kg})(4.2\,\mathrm{kJ/(kg\,K)})(5\,\mathrm{K}) \\ &= \boxed{42\,\mathrm{kJ}}. \end{aligned}$$

Check and meaning. The kg and K units cancel. A real heater would generally require more supplied energy if losses or container heating were included.

16.3 Ideal gas calculations

For an ideal gas, $pV = nRT$. Use absolute pressure and absolute temperature. With $R = 8.314\,\mathrm{J/(mol\,K)}$, pressure must be in pascals and volume in cubic metres because one Pa $\mathrm{m^3}$ equals one joule. The model is an approximation; its suitability depends on pressure, temperature and gas behaviour.

For fixed amount of gas, pV/T is constant. If temperature remains constant, pressure varies inversely with volume. If volume remains constant, pressure is proportional to absolute temperature. Ratios of Celsius temperatures have no such meaning.

Worked problem 16.3 Pressure of an ideal gas

Problem. One mole of ideal gas occupies 0.024 $\mathrm{m^3}$ at 300 K. Find pressure using $R = 8.314$.

Understanding the method. Rearrange the ideal gas equation for pressure and use matching SI units.

Calculation.

$$\begin{aligned} p &= \frac{nRT}{V} = \frac{1(8.314)(300)}{0.024} \\ &= 103{,}925\,\mathrm{Pa} \approx \boxed{103.9\,\mathrm{kPa}}. \end{aligned}$$

Check and meaning. The result is an absolute pressure under the ideal-gas model. It is not gauge pressure unless a reference pressure is subtracted.

Worked problem 16.4 Constant-volume warming

Problem. An ideal gas at 100 kPa absolute and 300 K is heated at constant volume to 360 K. Find its new pressure.

Understanding the method. For fixed amount and volume, p/T is constant.

Calculation.

$$\frac{p_2}{p_1} = \frac{T_2}{T_1} \quad \Rightarrow \quad p_2 = 100\frac{360}{300} = \boxed{120\,\mathrm{kPa}}.$$

Check and meaning. Absolute temperature rises by 20%, so pressure does too. Using Celsius in the ratio would give a fundamentally different and incorrect result.

16.4 Mass balances and mixtures

For a nonreacting component in a steady mixer, component mass input equals component mass output. Concentrations must be stated on a consistent basis: mass fractions, mole fractions and volume fractions cannot generally be interchanged. Fractions sum to one on their stated basis.

Worked problem 16.5 Mix two liquid streams

Problem. A 2 kg/min stream contains 10% solute by mass. Another 3 kg/min stream contains 20%. Find the outlet mass flow and solute fraction for a steady nonreacting mixer.

Understanding the method. Write a total mass balance and a separate solute balance. Weighted averaging is a consequence of these balances.

Calculation.

$$\begin{aligned} \dot{m}_{\text{out}} &= 2 + 3 = 5\,\text{kg/min}, \\ \dot{m}_{\text{solute}} &= 2(0.10) + 3(0.20) = 0.2 + 0.6 = 0.8\,\text{kg/min}, \\ w_{\text{out}} &= 0.8/5 = \boxed{0.16 = 16\%}. \end{aligned}$$

Check and meaning. The outlet fraction lies between 10% and 20%, closer to the stronger stream because it has the larger flow. The arithmetic mean 15% ignores unequal flows.

16.5 Residence time and rate

For an idealised vessel of working volume V and volumetric throughput Q, nominal hydraulic residence time is $\tau = V/Q$. Units must match. It is a characteristic volume-to-flow ratio, not a guarantee that every fluid parcel spends exactly that time in a mixed vessel.

A first-order decay model $dC/dt = -kC$ assumes rate proportional to remaining concentration. Separating variables gives $C(t) = C_0 e^{-kt}$. The rate constant has reciprocal-time units. Its half-life solves $C(t) = C_0/2$, yielding $t_{1/2} = \ln 2/k$. These are idealised model calculations, not validated predictions for any particular biogas system.

Worked problem 16.6 Residence time

Problem. A vessel has working volume 120 m^3 and throughput 8 m^3/day. Find nominal residence time.

Understanding the method. Divide stored working volume by volume passing through each day.

Calculation.

$$\tau = \frac{120\,\text{m}^3}{8\,\text{m}^3/\text{day}} = \boxed{15\ \text{days}}.$$

Check and meaning. Using total shell volume instead of working liquid volume can change the result. The specified working volume is the relevant quantity here.

Worked problem 16.7 First-order decay and half-life

Problem. For $C(t) = 100e^{-0.2t}$ with t in hours, find $C(5)$ and the half-life.

Understanding the method. The exponent must be dimensionless, so $k = 0.2$ per hour. The half-life is the time at which the concentration halves.

Calculation.

$$\begin{aligned} C(5) &= 100e^{-0.2(5)} = 100e^{-1} \approx \boxed{36.79}, \\ 50 &= 100e^{-0.2t_{1/2}}, \\ \ln(1/2) &= -0.2t_{1/2}, \\ t_{1/2} &= \frac{\ln 2}{0.2} \approx \boxed{3.466\ \text{h}}. \end{aligned}$$

Check and meaning. Five hours is longer than one half-life, so a value below 50 is expected. Concentration units are inherited from the stated initial value.

16.6 Model assumptions and physical interpretation

An equation can be mathematically solved while its physical assumptions are false. Check whether flow is steady, mixing is adequate, density is approximately constant, losses matter, and pressure is absolute or gauge. Dimensional agreement is necessary but not sufficient for validity. A model giving plausible-looking numbers is not automatically calibrated to a plant.

These checks also matter for research proposals. A software prediction should distinguish measured inputs, derived features, fitted parameters and assumptions. In an entrance calculation, clearly stating the assumptions given in the question often prevents selecting an attractive but inappropriate formula.

Chapter 17

Mixed problems with complete solutions

The idea before the calculation

In a chapter, the heading lends a hint. Outside it, a problem arrives without announcing which tool will help. Mixed work develops the judgement that formulas alone cannot supply: deciding what kind of question is being asked.

Translate the words into quantities and relationships before reaching for a method. A rough estimate, a limiting case or a unit check can give the calculation a direction. If the result is surprising, investigate both arithmetic and assumptions. Perhaps the algebra is correct and the model is not. The aim is to recognise a familiar structure inside an unfamiliar setting, then explain why the chosen method belongs there. This chapter deliberately combines topics. Read the method explanation whenever a problem crosses a subject boundary. These are original practice problems; no official question distribution or past-paper status is implied.

17.1 Calculus and linear algebra together

Worked problem 17.1 Differentiate a squared-error objective

Problem. Let $L(a,b) = (a+b-2)^2 + (a+2b-3)^2$. Find its minimiser.

Understanding the method. Each residual depends on both parameters. Differentiate each squared residual by the chain rule. The resulting equations are linear in a, b.

Calculation.

$$L_a = 2(a+b-2) + 2(a+2b-3) = 4a + 6b - 10,$$
$$L_b = 2(a+b-2) + 4(a+2b-3) = 6a + 10b - 16.$$

Setting both to zero gives $2a + 3b = 5$ and $3a + 5b = 8$. Multiply the first by 3 and the second by 2:

$$6a + 9b = 15, \qquad 6a + 10b = 16.$$

Subtracting yields $b = 1$, then $a = 1$. Both residuals become zero, so the nonnegative loss is globally minimised:

$$\boxed{a = 1,\ b = 1,\ L = 0}.$$

Check and meaning. Zero squared error establishes global minimality directly. The Hessian is $\begin{pmatrix} 4 & 6 \\ 6 & 10 \end{pmatrix}$ with positive first principal minor and determinant 4, also confirming strict convexity.

Worked problem 17.2 A quadratic form from a matrix product

Problem. For $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$, compute $A^T A$ and show that its quadratic form is positive for every nonzero vector.

Understanding the method. $x^T A^T A x = (Ax)^T (Ax) = \|Ax\|^2$. It is strictly positive if A has no nonzero null vector.

Calculation.

$$A^T A = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix}.$$

For $x = (s, t)^T$,

$$x^T A^T A x = s^2 + 4st + 5t^2 = (s + 2t)^2 + t^2.$$

Both squares vanish only if $t = 0$ and then $s = 0$. Therefore the form is strictly positive for nonzero x, so $\boxed{A^T A \text{ is positive definite}}$.

Worked problem 17.3 Optimise with a boundary

Problem. Minimise $f(x) = (x - 3)^2$ over $0 \le x \le 2$.

Understanding the method. The unconstrained stationary point lies outside the feasible interval. Check the derivative sign or compare the boundaries.

Calculation.

$$f'(x) = 2(x - 3) < 0 \quad \text{for } 0 \le x \le 2.$$

Thus the function decreases throughout the interval and is smallest at the right endpoint. Since $f(2) = 1$ and $f(0) = 9$,

$$\boxed{x = 2, \quad f_{\min} = 1}.$$

Check and meaning. Reporting $x = 3$ would ignore the domain. A stationary-point equation supplies no feasible candidate here.

Worked problem 17.4 Integrate a probability density

Problem. Let $f(x) = cx$ for $0 \le x \le 2$, zero elsewhere. Find c, $P(X \le 1)$ and $\mathbb{E}[X]$.

Understanding the method. A density must integrate to one. Use that condition before calculating probabilities or moments.

Calculation.

$$\begin{aligned} 1 &= \int_0^2 cx \, \mathrm{d}x = c[x^2/2]_0^2 = 2c \quad \Rightarrow \quad c = 1/2, \\ P(X \le 1) &= \int_0^1 \frac{x}{2} \, \mathrm{d}x = [x^2/4]_0^1 = \boxed{1/4}, \\ \mathbb{E}[X] &= \int_0^2 x\frac{x}{2} \, \mathrm{d}x = \frac{1}{2}[x^3/3]_0^2 = \boxed{4/3}. \end{aligned}$$

Check and meaning. The density increases toward the right, so a mean above the midpoint 1 is reasonable. A density value exceeding one would not alone be an error, although probabilities must remain at most one.

17.2 Probability, estimation and prediction

Worked problem 17.5 Independent sample averages

Problem. Independent X_1, X_2, X_3 have mean 5 and variance 12. Find mean and variance of $Y = (X_1 + X_2 + X_3)/3$.

Understanding the method. Linearity determines the mean. Independence removes covariance terms from the variance.

Calculation.

$$\begin{aligned} \mathbb{E}[Y] &= \tfrac{1}{3}(5 + 5 + 5) = \boxed{5}, \\ \mathrm{Var}(Y) &= \tfrac{1}{9}(12 + 12 + 12) = \boxed{4}, \\ \mathrm{SD}(Y) &= 2. \end{aligned}$$

Check and meaning. The average has the same expected value but less variance than one observation. Without independence, the covariance terms would need to be supplied or modelled.

Worked problem 17.6 A complement with unequal trial probabilities

Problem. Two independent detectors have detection probabilities 0.8 and 0.7. Find the probability at least one detects the event.

Understanding the method. Use the probability both fail, then subtract from one. Independence applies to the two detector outcomes as stated.

Calculation.

$$P(\text{both fail}) = (1-0.8)(1-0.7) = 0.2(0.3) = 0.06,$$
$$P(\text{at least one detects}) = 1 - 0.06 = \boxed{0.94}.$$

Check and meaning. The result exceeds either individual detection probability. Adding $0.8 + 0.7$ would exceed one and double-count joint detections.

Worked problem 17.7 Conditional expectation

Problem. Suppose $\mathbb{E}[Y \mid X] = 2X + 1$ and $\mathbb{E}[X] = 3$. Find $\mathbb{E}[Y]$.

Understanding the method. The law of total expectation says the overall mean equals the mean of the conditional mean. It does not require replacing a random variable inside an arbitrary nonlinear function by its mean. Here the function is linear.

Calculation.

$$\begin{aligned}\mathbb{E}[Y] &= \mathbb{E}[\mathbb{E}[Y \mid X]] \\ &= \mathbb{E}[2X+1] = 2\mathbb{E}[X] + 1 \\ &= 2(3) + 1 = \boxed{7}.\end{aligned}$$

Check and meaning. For a nonlinear conditional mean such as X^2, knowledge of $\mathbb{E}[X]$ alone would not determine $\mathbb{E}[Y]$.

Worked problem 17.8 Estimate an exponential rate

Problem. Independent positive waiting times $1, 2, 3$ follow an exponential model with unknown rate λ. Find its maximum-likelihood estimate.

Understanding the method. Multiply the densities, take logarithms, then differentiate with respect to the rate. The observations are fixed during this optimisation.

Calculation.

$$
\begin{aligned}
L(\lambda) &= \lambda^3 e^{-\lambda(1+2+3)} = \lambda^3 e^{-6\lambda}, \\
\ell(\lambda) &= 3\ln\lambda - 6\lambda, \\
\ell'(\lambda) &= 3/\lambda - 6 = 0 \quad \Rightarrow \quad \boxed{\hat{\lambda} = 1/2}, \\
\ell''(\lambda) &= -3/\lambda^2 < 0.
\end{aligned}
$$

Check and meaning. The fitted mean waiting time is $1/\hat{\lambda} = 2$, equal to the sample mean. Rate and mean waiting time have reciprocal units.

17.3 Algorithm and learning calculations

Worked problem 17.9 Compare two operation counts

Problem. At $n = 1024$, compare $n\log_2 n$ operations with n^2 operations, assuming equal operation costs.

Understanding the method. This is a numerical comparison under the stated equal-cost model, not a complete runtime benchmark.

Calculation.

$$
\begin{aligned}
\log_2 1024 &= 10, \\
n\log_2 n &= 1024(10) = 10{,}240, \\
n^2 &= 1{,}048{,}576, \\
\frac{n^2}{n\log_2 n} &= \frac{1024}{10} = \boxed{102.4}.
\end{aligned}
$$

Check and meaning. Under the assumptions, the quadratic count is 102.4 times larger. Real implementation constants and memory behaviour can affect elapsed time.

Worked problem 17.10 A two-parameter gradient update

Problem. For $L(w_1, w_2) = (w_1 - 1)^2 + 2(w_2 + 1)^2$, start at $(0, 0)$ and use $\eta = 0.1$. Find one gradient update and its loss.

Understanding the method. Differentiate with respect to each parameter while holding the other fixed, then update both using the original gradient.

Calculation.

$$\begin{aligned}
\nabla L &= (2(w_1 - 1), 4(w_2 + 1))^T, \\
\nabla L(0,0) &= (-2, 4)^T, \\
w_{\text{new}} &= (0,0)^T - 0.1(-2,4)^T = (0.2, -0.4)^T, \\
L(w_{\text{new}}) &= (-0.8)^2 + 2(0.6)^2 \\
&= 0.64 + 0.72 = \boxed{1.36}.
\end{aligned}$$

The starting loss was $1 + 2 = 3$, so this update reduces it.

Check and meaning. Updating one coordinate and then recalculating a gradient for the next would be a different procedure from the simultaneous gradient update requested.

Worked problem 17.11 Standardise training and test values

Problem. Training data have mean 10 and standard deviation 2. Standardise a test value 14. Should the test value be included when recomputing the scaler?

Understanding the method. The fitted scaler is part of the training procedure and is applied unchanged to held-out data.

Calculation.

$$z = \frac{x - \mu_{\text{train}}}{\sigma_{\text{train}}} = \frac{14 - 10}{2} = \boxed{2}.$$

The test value should not be used to refit the scaler for this evaluation. Its role is to assess the procedure already fitted from training data.

Check and meaning. A z-score of two means two training standard deviations above the training mean. It is not a probability of two.

Worked problem 17.12 A confusion matrix after a threshold change

Problem. A classifier has $TP = 8, FP = 2, FN = 2, TN = 8$. A lower threshold changes one false negative to true positive and two true negatives to false positive. Recompute precision and recall.

Understanding the method. Update the counts before recomputing ratios. The total number of actual positives and negatives stays unchanged.

Calculation.

$$\begin{aligned}
&TP' = 9, \qquad FN' = 1, \qquad FP' = 4, \qquad TN' = 6, \\
&P' = 9/(9+4) = \boxed{9/13 \approx 0.6923}, \\
&R' = 9/(9+1) = \boxed{0.9}.
\end{aligned}$$

Originally precision and recall were both 0.8. The lower threshold increases recall but lowers precision in this case.

Check and meaning. Both before and after, there are ten actual positives and ten actual negatives. The threshold changes predictions, not the labels.

17.4 Engineering and modelling together

Worked problem 17.13 A combined resistor network

Problem. A $2\,\Omega$ resistor is in series with a parallel pair of $6\,\Omega$ and $3\,\Omega$. The source is 12 V. Find source current and the current through each parallel branch.

Understanding the method. Reduce the parallel group first, then solve the series circuit. Use the group's voltage to recover its branch currents.

Calculation.

$$\begin{aligned} R_p &= \frac{6(3)}{6+3} = 2\,\Omega, \\ R_{\text{total}} &= 2 + 2 = 4\,\Omega, \\ I_{\text{source}} &= 12/4 = 3\,\text{A}, \\ V_p &= I_{\text{source}} R_p = 3(2) = 6\,\text{V}, \\ I_6 &= 6/6 = 1\,\text{A}, \qquad I_3 = 6/3 = 2\,\text{A}. \end{aligned}$$

Check and meaning. The branch currents sum to 3 A, matching the series current. Using the full 12 V across the parallel group would ignore the first resistor's voltage drop.

Worked problem 17.14 A heater with losses

Problem. Water requires 42 kJ of useful heating. A heater is 70% efficient and draws 1 kW. Find required electrical energy and heating time under constant-power assumptions.

Understanding the method. Efficiency relates useful output to input. Calculate input energy before dividing by input power.

Calculation.

$$\begin{aligned} E_{\text{in}} &= E_{\text{useful}}/\eta = 42/0.70 = 60\,\text{kJ}, \\ P &= 1\,\text{kW} = 1\,\text{kJ/s}, \\ t &= E_{\text{in}}/P = 60/1 = \boxed{60\ \text{s}}. \end{aligned}$$

Check and meaning. The input exceeds useful heat because efficiency is below one. Dividing 42 kJ by the electrical power directly would silently assume perfect efficiency.

Worked problem 17.15 A vessel with accumulation

Problem. A tank initially contains 10 kg. Inflow is 3 kg/min and outflow is 2 kg/min for four minutes, with no reaction. Find final mass.

Understanding the method. This is unsteady: input exceeds output, so accumulation is not zero.

Calculation.

$$\begin{aligned}\frac{dM}{dt} &= 3 - 2 = 1\,\text{kg/min},\\ \Delta M &= 1(4) = 4\,\text{kg},\\ M(4) &= 10 + 4 = \boxed{14\,\text{kg}}.\end{aligned}$$

Check and meaning. A steady-state assumption would contradict the specified unequal flows unless another flow or effect were present.

Worked problem 17.16 An exponential model and measured units

Problem. A synthetic process model is $q(t) = 20e^{-0.1t}$ m^3/day with t in days. Find total predicted volume from day zero to day ten.

Understanding the method. Flow is a rate, so total volume is its time integral. The model is supplied for calculation; no measured plant validity is assumed.

Calculation.

$$\begin{aligned}V &= \int_0^{10} 20e^{-0.1t}\,\mathrm{d}t\\ &= 20\left[-10e^{-0.1t}\right]_0^{10}\\ &= 200(1 - e^{-1}) \approx \boxed{126.42\,\text{m}^3}.\end{aligned}$$

Check and meaning. The starting rate is 20 m^3/day and then decreases. Therefore the ten-day total must be below 200 m^3, which it is.

17.5 How to review a completed solution

A complete written answer makes the method visible. Check that all given quantities were used appropriately, every denominator is allowed, units match, and a numerical substitution did not change the original problem. Distinguish an exact expression from a decimal approximation. When an assumption is needed but absent, state the conditional answer rather than inventing a fact.

If a result seems impossible, first check signs, bounds and units. A negative probability, negative variance, resistor power with inconsistent units, or an infeasible optimiser is a signal to revisit the calculation. Independent checks such as substitution, conservation, a graph, a limiting case or a second method are often faster than repeating every line blindly.

Chapter 18

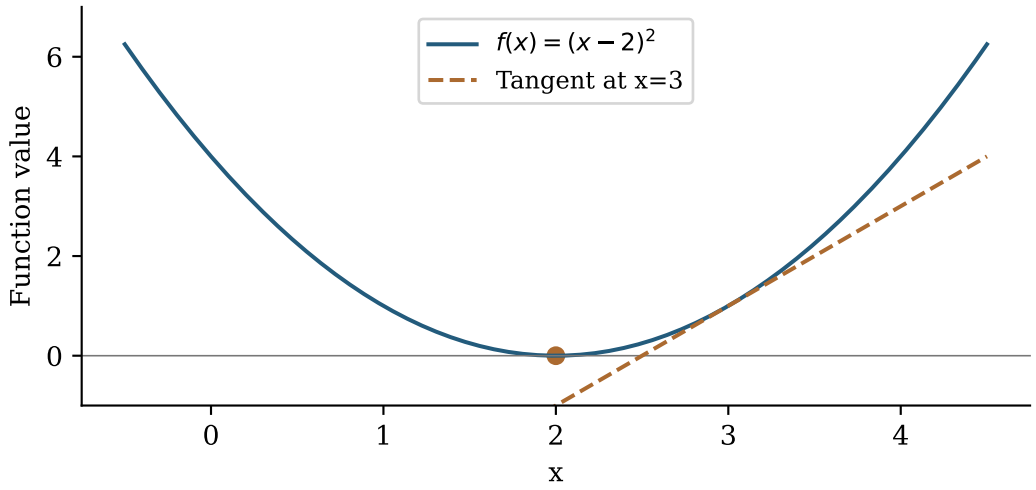
Reading mathematical figures

The idea before the calculation

A figure is a way of thinking with space. It can reveal a trend that a table conceals, but it can also make the same numbers look very different. Begin with the axes, units, ranges and scale. Equal distances on a logarithmic axis represent equal ratios, rather than equal differences.

A slope on the page depends on the axis scales; a physical rate must be read with its units. Distinguish observations, fitted curves and illustrative diagrams, and look for uncertainty and provenance where measurements are claimed. The figures in this book illustrate mathematical relationships and worked examples. Their purpose is to make structure visible, while the equations and assumptions explain what that structure means. A figure is useful only when its axes, quantities and relationship to a calculation are understood. This chapter revisits the mathematical ideas through exact conceptual plots. They are generated from stated equations, not from experimental plant measurements. Reading a figure is an additional way to check a solution, not a replacement for the derivation.

18.1 Slope is different from height



The curve is $f(x) = (x-2)^2$. Its lowest height is zero at $x = 2$. The dashed tangent at $x = 3$ has slope $f'(3) = 2$ and equation $y = f(3) + 2(x-3) = 2x - 5$. The derivative is the slope of this tangent, not the curve's height $f(3) = 1$.

Worked problem 18.1 Recover a tangent equation

Problem. For $f(x) = (x-2)^2$, find the tangent at $x = 1$ and its intercept with the vertical axis.

Understanding the method. A line with slope m through $(a, f(a))$ has equation $y - f(a) = m(x-a)$. Obtain m by differentiation.

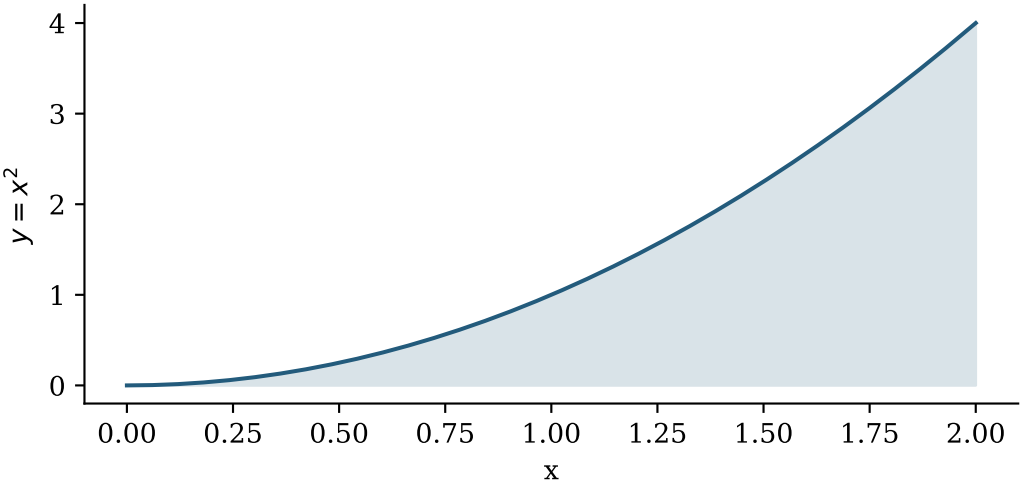
Calculation.

$$f(1) = 1, \qquad f'(x) = 2(x-2), \qquad f'(1) = -2,$$
$$y - 1 = -2(x-1) \quad \Rightarrow \quad y = -2x + 3.$$

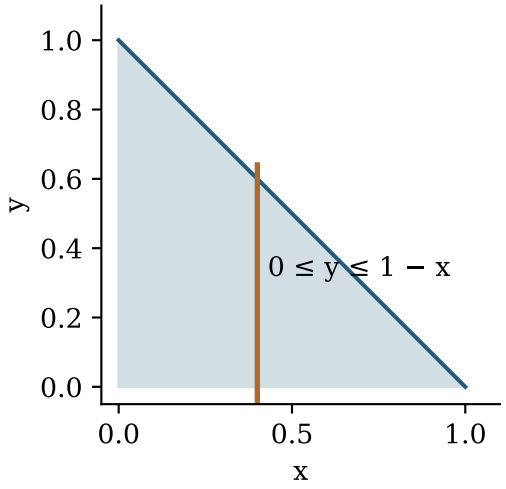
At $x = 0$, $y = 3$. Thus the tangent is $\boxed{y = -2x + 3}$ and its vertical intercept is $\boxed{3}$.

Check and meaning. The negative slope agrees with the falling left side of the parabola. The tangent passes through $(1, 1)$ as required.

18.2 Area and changing bounds



The shaded area under x^2 from zero to two is $[x^3/3]_0^2 = 8/3$. It has area units if both axes carry length units. The shape widens toward large x, so using a single endpoint height times width would overestimate it.



The vertical slice of the triangular region has upper boundary $y = 1 - x$. At $x = 0.4$ its height is 0.6. The slice changes as x changes, which is why the inner integration bound depends on the outer variable.

Worked problem 18.2 Reverse the order of a triangular integral

Problem. Express $\int_0^1 \int_0^{1-x} (x + y)\,\mathrm{d}y\,\mathrm{d}x$ with x as the inner variable.

Understanding the method. Describe the same region in the other direction. A horizontal slice at fixed y runs from $x = 0$ to $x = 1 - y$.

Calculation.

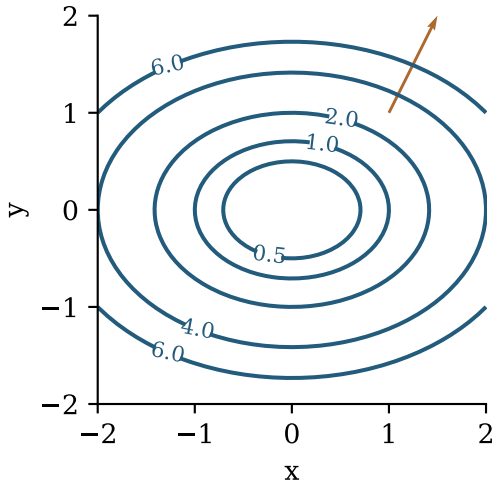
$$I = \int_0^1 \int_0^{1-y} (x + y)\,\mathrm{d}x\,\mathrm{d}y.$$

Evaluating,

$$\begin{aligned} I &= \int_0^1 \left[\frac{x^2}{2} + yx \right]_0^{1-y} \mathrm{d}y \\ &= \int_0^1 \left[\frac{(1-y)^2}{2} + y(1-y) \right] \mathrm{d}y \\ &= \int_0^1 \left(\frac{1}{2} - \frac{y^2}{2} \right) \mathrm{d}y = \boxed{1/3}. \end{aligned}$$

Check and meaning. The value agrees with the earlier order. Changing order changes the bounds, not the region or the final integral.

18.3 Contours and gradient direction



Contours join points with equal $f(x, y) = x^2 + 2y^2$. Moving along a contour does not change the value to first order. The gradient $(2x, 4y)$ is perpendicular to a contour at smooth points where it is nonzero. At $(1, 1)$ it points in direction $(2, 4)$. It is not generally directed along the line from the origin with equal scaling in both axes, because curvature differs by direction.

Worked problem 18.3 A tangent direction to a contour

Problem. At $(1, 1)$ for $f = x^2 + 2y^2$, show that direction $(2, -1)$ is tangent to the contour to first order.

Understanding the method. A tangent direction has zero dot product with the gradient. For a directional rate per unit distance, normalise it; zero remains zero after normalisation.

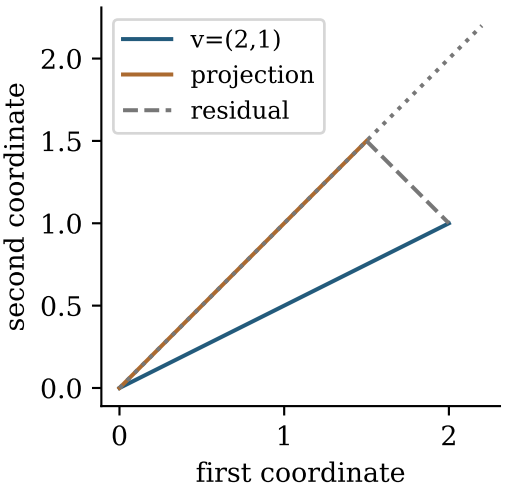
Calculation.

$$\nabla f(1,1) = (2,4), \qquad (2,4)\cdot(2,-1) = 4-4 = \boxed{0}.$$

For $u = (2,-1)/\sqrt{5}$, the directional derivative is also zero.

Check and meaning. This says the first-order change vanishes. A finite straight movement need not remain on the curved contour.

18.4 Projection and residual geometry



The projection $(1.5, 1.5)$ lies on the line spanned by $(1,1)$. The dashed residual points from the projection to the original vector $(2,1)$. It is perpendicular to the line. This right-angle condition is the geometric source of the normal equations, not an additional coincidence.

Worked problem 18.4 Distance to a subspace

Problem. Find the shortest distance from $(2,1)$ to the line $y = x$.

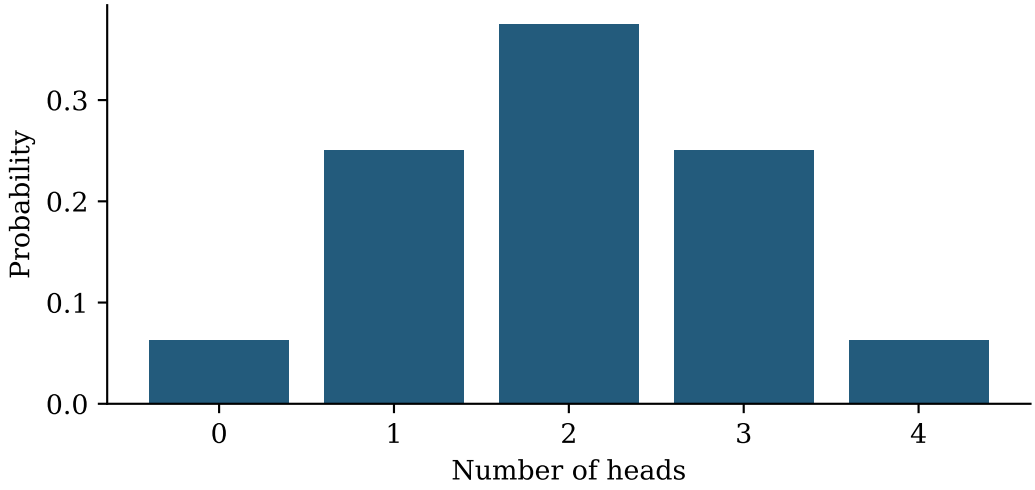
Understanding the method. Use the length of the orthogonal residual, not the length of the projection.

Calculation.

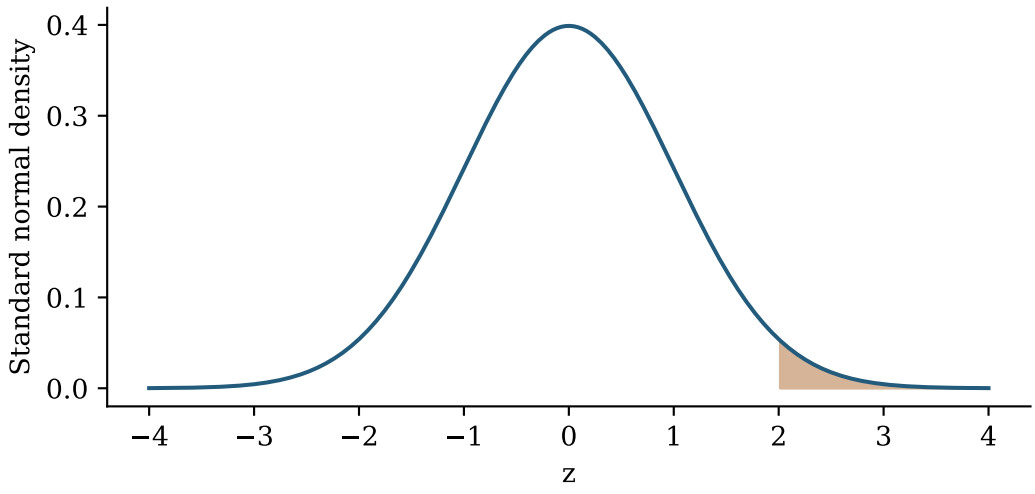
$$\begin{aligned} p &= (3/2, 3/2), \\ r &= (1/2, -1/2), \\ \text{distance} &= \sqrt{(1/2)^2 + (-1/2)^2} = \boxed{1/\sqrt{2}}. \end{aligned}$$

Check and meaning. The point is not on the line, so the distance is positive. The line equation $x - y = 0$ independently gives $|2 - 1|/\sqrt{1^2 + (-1)^2} = 1/\sqrt{2}$.

18.5 Probability is mass or area



The binomial plot shows probabilities for $0, 1, 2, 3, 4$ heads in four fair tosses. Its bar heights are $1/16, 4/16, 6/16, 4/16, 1/16$, summing to one. Discrete probability is obtained by adding relevant masses, not integrating a smooth curve between the bars.



For a continuous normal variable, probability is area under the density. The shaded right tail begins at $z = 2$. The density height at 2 is not the tail probability: the latter integrates all density from 2 to infinity.

Worked problem 18.5 A two-sided normal tail

Problem. Using $\Phi(2) = 0.97725$, find $P(|Z| > 2)$ for a standard normal variable.

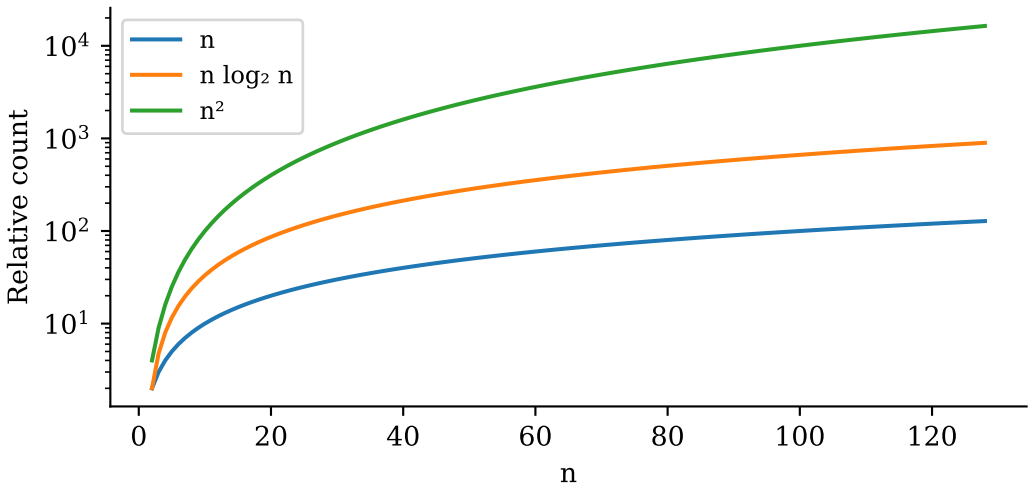
Understanding the method. Symmetry gives equal tails. The event is the disjoint union $Z < -2$ or $Z > 2$.

Calculation.

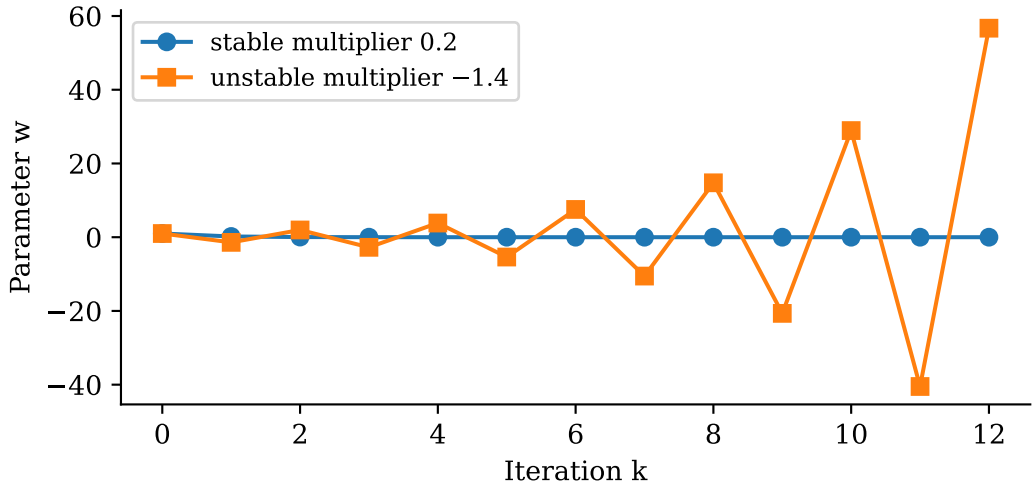
$$\begin{aligned} P(Z > 2) &= 1 - 0.97725 = 0.02275, \\ P(Z < -2) &= 0.02275, \\ P(|Z| > 2) &= 2(0.02275) = \boxed{0.0455}. \end{aligned}$$

Check and meaning. The two-sided tail is twice the one-sided tail because of symmetry; this shortcut would not automatically apply to an asymmetric distribution.

18.6 Growth and optimisation plots



The vertical axis is logarithmic, so equal vertical steps correspond to multiplicative changes. It allows n, $n \log_2 n$ and n^2 to be compared on one plot. A curve's appearance depends on axis scaling; the algebraic complexity classification does not.



The stable sequence has multiplier 0.2 and shrinks. The unstable one has multiplier -1.4, alternates sign and grows in magnitude. This directly visualises the quadratic stability condition $|1 - \eta a| < 1$.

Worked problem 18.6 An alternating but convergent update

Problem. If $w_{k+1} = -0.5 w_k$ and $w_0 = 2$, calculate three updates and identify the limit.

Understanding the method. A negative multiplier reverses sign, but a magnitude below one shrinks each update.

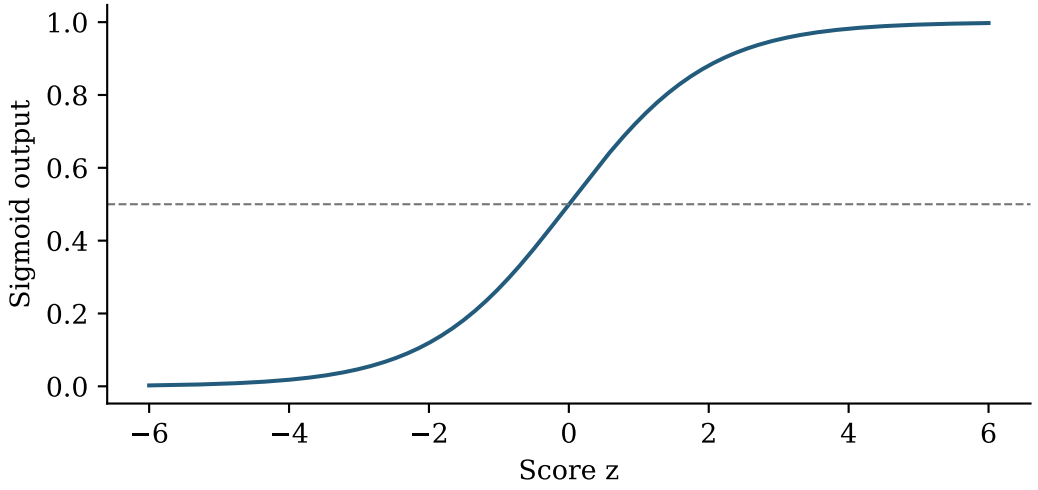
Calculation.

$$w_1 = -1, \qquad w_2 = 0.5, \qquad w_3 = -0.25.$$

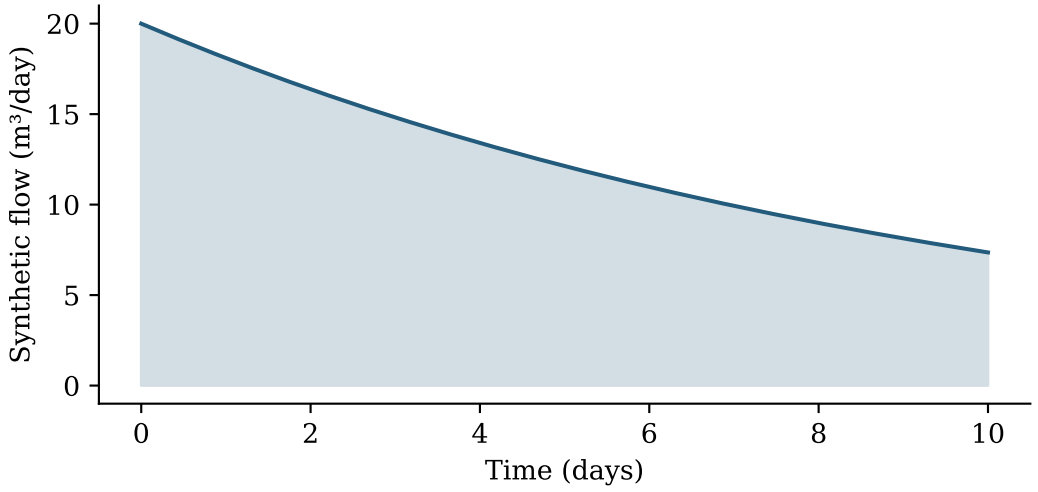
In general $w_k = 2(-0.5)^k$, so $\boxed{w_k \to 0}$.

Check and meaning. Oscillation is compatible with convergence. Divergence requires growing magnitude, not simply changing sign.

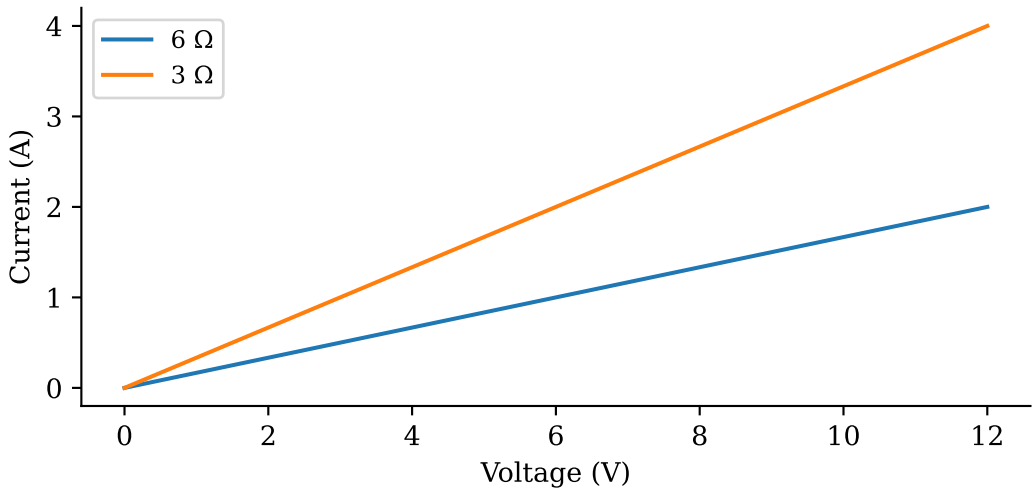
18.7 Scores, flow and physical dimensions



A score of zero maps to probability output one-half. Large positive and negative scores approach one and zero, respectively. The curve is smooth; applying a threshold creates a separate discontinuous class decision.



The area under the synthetic flow curve is volume, because a rate in m^3/day is multiplied by days during integration. The curve itself is neither accumulated volume nor measured plant evidence.



Current versus voltage has slope $1/R$ for an ideal resistor. The $3\,\Omega$ line is steeper than the $6\,\Omega$ line. If voltage were plotted vertically against current horizontally, the slope would instead be R. Always read axis labels before interpreting a slope.

Chapter 19

Algebra and numerical tools for self-study

The idea before the calculation

Algebra is the craft of changing an expression while keeping track of what remains true. Each step carries conditions. Dividing by a quantity assumes it is nonzero; squaring an equation can introduce candidates that were not solutions of the original equation. Checking the final result in the original problem restores those forgotten conditions.

Numerical computation adds a second distinction: the mathematical quantity and its finite-precision representation. Subtracting nearly equal rounded numbers can lose useful digits; an ill-conditioned problem can amplify small input errors. More decimal places in the answer do not create better measurements. Treat algebra, units and numerical checks as companions: together they make a calculation easier to understand and harder to fool.

This reference chapter supports the calculations used throughout the book. Return here when the difficulty is not the new concept but an algebraic manipulation, exponent, fraction or unit conversion. These are prerequisites for using the syllabus mathematics, rather than a separate claim about examination topics.

19.1 Fractions and equations

To add fractions, use a common denominator. To divide by a nonzero fraction, multiply by its reciprocal. In a multistep equation, perform the same reversible operation on both sides. Keep exact fractions until the final decimal approximation when possible.

Worked problem 19.1 Add and divide fractions

Problem. Calculate $(\frac{1}{2} + \frac{1}{3})/(\frac{3}{4})$.

Understanding the method. Add the numerator using denominator six, then divide by multiplying by the reciprocal of 3/4.

Calculation.

$$\frac{1}{2}+\frac{1}{3}=\frac{3}{6}+\frac{2}{6}=\frac{5}{6},$$
$$\frac{5/6}{3/4}=\frac{5}{6}\cdot\frac{4}{3}=\frac{20}{18}=\boxed{\frac{10}{9}}.$$

Check and meaning. The numerator is about 0.8333 and the divisor is 0.75, so the quotient should be slightly above one. Adding denominators directly would not represent fraction addition.

Worked problem 19.2 Solve a linear equation with denominators

Problem. Solve $(x-1)/2+(x+2)/3=4$.

Understanding the method. Multiply every term by the common denominator six. Parentheses must remain until multiplication is distributed.

Calculation.

$$\begin{aligned}3(x-1)+2(x+2)&=24,\\ 3x-3+2x+4&=24,\\ 5x+1&=24,\\ x&=\boxed{23/5}.\end{aligned}$$

Substitution gives $(3.6)/2+(6.6)/3=1.8+2.2=4$.

19.2 Quadratics and factorisation

A quadratic equation $ax^2+bx+c=0$ with $a\neq 0$ has roots

$$x=\frac{-b\pm\sqrt{b^2-4ac}}{2a}.$$

The discriminant b^2-4ac determines whether real roots exist. A positive value gives two distinct real roots; zero gives a repeated root; negative gives no real roots. Factoring can be faster when simple factors are visible, but the quadratic formula applies generally.

Worked problem 19.3 Use the quadratic formula

Problem. Solve $2x^2-5x+2=0$.

Understanding the method. Identify signed coefficients before substituting. Here $a = 2$, $b = -5$, $c = 2$.

Calculation.

$$\begin{aligned} b^2 - 4ac &= (-5)^2 - 4(2)(2) = 25 - 16 = 9, \\ x &= \frac{5 \pm 3}{4}, \\ x_1 &= \boxed{2}, \qquad\qquad x_2 = \boxed{1/2}. \end{aligned}$$

The factorisation $(2x - 1)(x - 2) = 0$ verifies both roots.

Check and meaning. The denominator is $2a = 4$, not merely 2. The minus sign in $-b$ turns -5 into $+5$.

19.3 Exponentials and logarithms

An exponential equation is often solved by taking logarithms. The argument of a real logarithm must be positive. If $e^{kt} = c$ with $c > 0$ and $k \neq 0$, then $kt = \ln c$ and $t = \ln c/k$.

Worked problem 19.4 Find the time to a specified fraction

Problem. For $C(t) = 50e^{-0.25t}$, find the time at which $C = 10$.

Understanding the method. First divide by the initial value. Then take natural logarithms to bring the exponent down.

Calculation.

$$\begin{aligned} 10 &= 50e^{-0.25t}, \\ 0.2 &= e^{-0.25t}, \\ \ln 0.2 &= -0.25t, \\ t &= -\frac{\ln 0.2}{0.25} = \frac{\ln 5}{0.25} \approx \boxed{6.438}. \end{aligned}$$

Check and meaning. The time is positive because $\ln 0.2$ is negative. Its units are the reciprocal of the stated rate's units.

19.4 Scientific notation and units

Scientific notation writes a nonzero quantity as a number between one and ten in magnitude times a power of ten. Multiplication adds exponents; division subtracts them. Unit conversions must apply to every power: converting area from mm^2 to m^2 squares the millimetre conversion.

Worked problem 19.5 Convert an area before calculating stress

Problem. Convert $250\,\mathrm{mm}^2$ to m^2 and find stress under 5000 N.

Understanding the method. Because 1 mm is 10^{-3} m, one square millimetre is $(10^{-3})^2 = 10^{-6}$ m^2.

Calculation.

$$\begin{aligned} A &= 250 \times 10^{-6} = 2.5 \times 10^{-4}\,\mathrm{m}^2, \\ \sigma &= \frac{5000}{2.5 \times 10^{-4}} = 2.0 \times 10^{7}\,\mathrm{Pa} \\ &= \boxed{20\,\mathrm{MPa}}. \end{aligned}$$

Check and meaning. The direct calculation $5000/250 = 20$ $\mathrm{N/mm}^2$ gives the same result. Using only 10^{-3} for the area conversion would create a factor-of-one-thousand error.

19.5 Ratios, percentages and changes

A percentage is a fraction of one hundred. An increase from 20 to 25 is an absolute change of 5 and a relative change of $(25-20)/20 = 25\%$. The reverse change from 25 to 20 is -20%, because its reference value is different. A percentage-point difference between percentages is another quantity: increasing accuracy from 80% to 90% is ten percentage points, or a 12.5% relative increase in the numerical accuracy.

Worked problem 19.6 Relative improvement in error

Problem. A baseline MSE is 10 and a new model has MSE 7. Calculate the relative reduction.

Understanding the method. Use the baseline as the denominator because the question asks improvement relative to the baseline.

Calculation.

$$\text{relative reduction} = \frac{10-7}{10} = 0.3 = \boxed{30\%}.$$

Check and meaning. The absolute reduction is 3 squared target units. Dividing by 7 would measure a different relative comparison.

19.6 Rounding and exact answers

An exact answer such as $1/3$ differs from the decimal approximation 0.3333. Use an approximation sign when rounding. Retain enough internal precision so early rounding does

not distort a later cancellation or ratio. In objective questions, inspect whether an answer is requested as an expression, integer, percentage or decimal.

Worked problem 19.7 Avoid premature rounding

Problem. Compare $3(1/3)$ exactly with a calculation that rounds $1/3$ to 0.33 before multiplication.

Understanding the method. An exact rational representation retains information that a short decimal discards.

Calculation.

$$3\left(\frac{1}{3}\right) = \boxed{1}, \qquad 3(0.33) = 0.99.$$

The absolute discrepancy is 0.01, or 1% of the exact result. Rounding only after the exact multiplication would give 1.00.

Check and meaning. Not every approximation creates a serious error, but delaying rounding is a simple way to avoid unnecessary loss of accuracy.

19.7 Normalising a vector

Normalisation changes length while preserving direction. Divide a nonzero vector by its norm. The zero vector cannot be normalised because its norm is zero.

Worked problem 19.8 A unit direction

Problem. Find a unit vector in the direction $v = (3, -4)^T$ and verify its length.

Understanding the method. The Euclidean norm is the square root of the sum of squared components.

Calculation.

$$\begin{aligned} \|v\| &= \sqrt{3^2 + (-4)^2} = \sqrt{25} = 5, \\ u &= v/5 = \boxed{(3/5, -4/5)^T}, \\ \|u\|^2 &= 9/25 + 16/25 = 1. \end{aligned}$$

Check and meaning. The signs are retained, so the unit vector points in the same direction. Changing a sign would reflect a component, not merely rescale the vector.

Syllabus map and reading directions

The original study scope was informed by the supplied *C-MInDS MS by Research and PhD Admissions 2026–27 Test Syllabus*, September 2026. Its references are Gilbert Strang, *Calculus*, second edition, Chapters 1–15; Strang, *Introduction to Linear Algebra*, third edition, Chapters 1–6; Sheldon M. Ross, *Introduction to Probability and Statistics for Engineers and Scientists*, fifth edition, Chapters 3, 4, 6, 7 and 9; Cormen, Leiserson, Rivest and Stein, *Introduction to Algorithms*, third edition, Chapters 1, 2, 3, 4, 6, 7 and 12; and Goodfellow, Bengio and Courville, *Deep Learning*, Chapter 5.

Chapters 1–3 of this book map to calculus; 4–6 to linear algebra; 7–9 to probability and statistics; 10–11 to algorithms; 12–13 to machine learning; and 14–16 to engineering core. Chapter 17 provides mixed original problems. Supplementary details are included when they help explain a listed topic; their inclusion is not a prediction of an examination question.

Official and author-hosted reading entrances:

- https://www.minds.iitb.ac.in/admissions/test-syllabus.pdf
- https://ocw.mit.edu/courses/18-01sc-single-variable-calculus-fall-2010/
- https://ocw.mit.edu/courses/18-06-linear-algebra-spring-2010/
- https://ocw.mit.edu/courses/6-006-introduction-to-algorithms-fall-2011/
- https://www.deeplearningbook.org/contents/ml.html

The open courses are supplementary, not the official C-MInDS reading list. Their examination papers belong to their own courses. The original problems in this book should be used to learn methods; actual selection, difficulty and topic balance remain with the Institute.

BEFORE THE ALGORITHM

Begin with the idea. Then make it work.

Before a calculation becomes useful, its symbols must mean something. This concise companion brings the foundations of intelligent systems into one place, connecting mathematical ideas with the physical and computational questions they help us answer.

Move from calculus, linear algebra and probability to algorithms, machine learning and selected engineering principles. Each worked problem keeps the method, calculation and check together, so a quick revision can become a moment of understanding.

Inside: 19 chapters, 128 worked problems, short conceptual openings, mathematical figures, mixed practice and algebraic tools.

AKASH SHIVADAS CHATAKE

AIML, Data Science and Mathematics Lecturer. His teaching connects mathematical reasoning with scientific intuition and practical computation.

Read the working edition

innoworkspress.org

Independent study material. Author and technical review in progress.

INNOWORKS PRESS

Chatake Innoworks Private Limited