
LN's

Endless

Maths Notes



Art:

Kira, by @Ekkoberry

M1. ALGEBRA	6
1.1. Mathematical Syntax and Techniques	6
1.2. Algebraic Identities	10
1.3. Trigonometric Identities	18
1.4. Sequences and Series	21
1.5. Complex Numbers	24
1.6. Discrete Mathematics and Abstract Algebra	26
1.7. Special Functions and Identities	30
M2. GEOMETRY	38
2.1. Properties of 2D and 3D Shapes	38
2.2. Angle, Triangle and Circle Theorems	45
2.3. 2D Coordinate Geometry	54
2.4. Vectors and 3D Geometry	61
M3. CALCULUS	67
3.1. Limits and Numerical Methods	67
3.2. Series Expansions	75
3.3. Differentiation and Integration	79
3.4. Ordinary Differential Equations, Laplace and Z-Transforms	87
3.5. Multivariable and Vector Calculus	98
3.6. Fourier Series and Fourier Transforms	110
3.7. Partial Differential Equations and Variational Calculus	117
M4. LINEAR ALGEBRA	122
4.1. Vector and Matrix Algebra	122
4.2. Transformation Matrices	130
4.3. Fundamental Subspaces and Matrix Decompositions	132
4.4. Matrix and Tensor Calculus	139
M5. STATISTICS	142
5.1. Axioms, Combinatorial Probability and Basic Statistics	142
5.2. Probability Distributions and Random Variables	147
5.3. Hypothesis Testing	162
5.4. Stochastic Processes, Signals and Information Theory	171
5.5. Machine Learning and Computational Statistics	186
5.6. Computer Vision and Computer Graphics	210
5.7. Accounting, Finance and Business	218

M1. ALGEBRA

1.1. Mathematical Syntax and Techniques

1.1.1. Symbols for Relationships and Operators

$x : y$	ratio of x to y , representing quantities $\frac{x}{x+y}$ and $\frac{y}{x+y}$
$\lfloor x \rfloor$	floor of x , $\max\{a \in \mathbb{Z} : a \leq x\}$, round down to integer towards $-\infty$
$\lceil x \rceil$	ceiling of x , $\min\{a \in \mathbb{Z} : a \geq x\}$, round up to integer towards $+\infty$
$\{x\}$	fractional part, $x - \lfloor x \rfloor$
$1.\overline{238307}$	recurring decimal, 1.238307307307307..., or in dot notation 1.238307
\equiv	is identical to; is congruent to
$:=$	is defined as
\therefore	therefore
\because	because
$P \Rightarrow Q$	P implies Q ; if P then Q ; Q is necessary for P , P is sufficient for Q
$P \Leftarrow Q$	P is implied by Q ; if Q then P ; P is necessary for Q , Q is sufficient for P
$P \Leftrightarrow Q$	P and Q are equivalent; P if and only if (iff) Q ; P is necessary and sufficient for Q
$f: A \mapsto B$	function mapping domain A to codomain B
\bar{y}	sample mean of y ; Laplace transform of $y(t)$ into s domain
\hat{y}	estimate of y ; least-squares solution y ; unit vector y
$f_X(x)$	probability density function (pdf) for random variable X , taking value x
$F_X(x)$	cumulative density function (cdf) for random variable X , taking value x
\cong	is isomorphic to; is geometrically congruent to
$\prod_{r=a}^b f(r)$	product over integers; $\prod_{r=a}^b f(r) = f(a) \times f(a+1) \times \dots \times f(b)$
$(f \circ g)(x)$	composition; $fg(x)$; $f(g(x))$
$(f * g)(t)$	convolution of $f(t)$ and $g(t)$
$(f \star g)(t)$	correlation of $f(t)$ and $g(t)$
$[\mathbf{a} \ \mathbf{b} \ \mathbf{c}]$	scalar triple product; $\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c})$
$[a, b, c]$	row vector; $[a, b, c]^T$; alternative way of denoting vectors
$[a; b; c]$	column vector; $[a, b, c]^T$; most common conventional way of denoting vectors
\dot{x}, \ddot{x}	time derivatives of x ; $\frac{dx}{dt}, \frac{d^2x}{dt^2}$
$y^{(n)}(x)$	n th derivative of y with respect to x ; $\frac{d^n y}{dx^n}$
$f_{xy}(x, y)$	partial derivative; $\frac{\partial f}{\partial y \partial x} = \frac{\partial}{\partial y} \left(\frac{\partial f}{\partial x} \right)$
Δf	change in f ; Laplacian of multivariable function f ; $\nabla^2 f$
δf	small change in f

1.1.2. Symbols in Set Theory and Logic

\in	is an element of	$n(A)$	number of elements in A
\notin	is not an element of	\emptyset	the empty set
\subseteq	is a subset of	ξ	the universal set
\subset	is a proper subset of	\forall	for all
$\{x: \dots\}$	the set of all x such that ...	\exists	there exists
A'	complement of A	$A \setminus B$	the set A minus B
$A \cup B$	union of A and B	$A \cap B$	intersection of A and B
\mathbb{P}	prime numbers; $\{2, 3, 5, \dots\}$	\mathbb{Q}	rationals; $\{1, 2, 1/2, \dots\}$
\mathbb{N}	natural numbers; $n; \{1, 2, 3, \dots\}$	\mathbb{R}	reals; x
\mathbb{Z}	integers; $\{\dots, -2, -1, 0, 1, 2, \dots\}$	\mathbb{C}	complex numbers; $z = x + yi$
\mathbb{Z}^*	nonnegative integers; $\{0, 1, 2, \dots\}$	\mathbb{H}	quaternions; $q = w + xi + yj + zk$
\mathbb{R}^n	n -dimensional vectors; \mathbf{x}	$\mathbb{R}^{m \times n}$	$m \times n$ matrices; \mathbf{A}

Algebraic numbers are roots of real-valued polynomials with rational (or integer) coefficients. Transcendental numbers (e, π) are irrational numbers which are **not** algebraic.

Note that $\mathbb{P} \subset \mathbb{N} \subset \mathbb{Z} \subset \mathbb{R} \subset \mathbb{C} \subset \mathbb{H}$. Irrational numbers can be designated as $\mathbb{R} \setminus \mathbb{Q}$.

1.1.3. Greek Alphabet

A α	alpha <i>άλφα</i>	H η	eta <i>ήτα</i>	N ν	nu <i>νι</i>	T τ	tau <i>ταυ</i>
B β	beta <i>βήτα</i>	Θ θ, ϑ	theta <i>θήτα</i>	Ξ ξ	xi <i>ξι</i>	Υ υ	upsilon <i>ύψιλον</i>
Γ γ	gamma <i>γάμμα</i>	I ι	iota <i>ιώτα</i>	O o	omicron <i>όμικρον</i>	Φ φ, ϕ	phi <i>φεϊ</i>
Δ δ	delta <i>δέλτα</i>	K κ	kappa <i>κάπα</i>	Π π, ϖ	pi <i>πι</i>	X χ	chi <i>χϊ</i>
E ϵ	epsilon <i>έψιλο</i>	Λ λ	lambda <i>λάμδα</i>	P ρ	rho <i>ρο</i>	Ψ ψ	psi <i>ψι</i>
Z ζ	zeta <i>ζήτα</i>	M μ	mu <i>μι</i>	Σ σ, ς	sigma <i>σίγμα</i>	Ω ω	omega <i>ωμέγα</i>

- The lowercase letters $\{\vartheta, \iota, o, \varsigma, \upsilon\}$ are typically not used as symbols in mathematics.
- The uppercase letters are typed upright. The lowercase letters are typed in *italic*.
- When using Latin letters ($A a$) as symbols, both uppercase and lowercase are typed *italic*.
- In LaTeX, the Greek letters can be written using e.g. `\gamma` (lower) or `\Gamma` (upper).

1.1.4. Functions

Domain:	the set of all x for which $f(x)$ is defined	
Codomain:	any set containing all the values of $f(x)$.	
Range:	the set containing only the values of $f(x)$, so $\text{range} \subseteq \text{codomain}$.	
Injective:	one-to-one; there is exactly one $f(x)$ for every x ;	$f(x) = f(y) \Leftrightarrow x = y$.
Surjective:	many-to-one; there is at least one x for every $f(x)$;	$f(x) = f(y) \Leftarrow x = y$.
Bijjective:	two-way one-to-one; both injective and surjective;	$y = f(x) \Leftrightarrow x = f^{-1}(y)$.
Involution:	a function whose inverse is identical to itself;	$f(f(x)) = x$
Idempotent:	a function whose nesting is an involution;	$f(f(x)) = f(x)$
Endofunction:	a function whose codomain is identical to its domain	

1.1.5. Rounding of Numbers

A number can be rounded to a given number of decimal places (d.p.) or significant figures (s.f.).

Examples:

- 309.51547 rounded to 2 d.p. is 309.52 (5 rounds up)
- 0.00194105 rounded to 2 s.f. is 0.0019 (leading zeroes are not significant)

When working with physical quantities with finite precision, the least number of available significant figures should be used to most fairly represent the quoted precision.

1.1.6. Standard Form of Numbers

Very large or very small numbers can be written in the form $x = a \times 10^n$

($1 \leq x < 10$: mantissa, integer n : exponent). This number is said to have order of magnitude 10^n .

Numbers in standard form can be added by converting both to the same exponent.

1.1.7. Factorisation of Numbers

A divisor q is a factor of a dividend p if the quotient $\frac{p}{q}$ is an integer (i.e. p is a multiple of q).

Prime factorisation: expressing a number uniquely as a product of powers of prime numbers. The different prime factorisation of two numbers can be represented as a Venn diagram.

Greatest Common Factor (GCF) of a and b : product of intersection of prime factors of a and a .

Lowest Common Multiple (LCM) of a and b : product of union of prime factors of a and b .

The GCF and LCM are related by $lcm(a, b) \times gcd(a, b) = |ab|$.

1.1.8. Power-Law and Exponential Relationships

Variables x and y may be related by nonlinear relationships such as:

Relationship	Equation	Inverse	Linearised form ($Y = MX + C$)	Linear plot
Power Law	$y = k x^n$	$x = \left(\frac{y}{k}\right)^{1/n}$	$\underbrace{\ln y}_Y = \underbrace{n}_M \underbrace{\ln x}_X + \underbrace{\ln k}_C$	$\ln y$ vs $\ln x$
Exponential	$y = k a^x$	$x = \log_a \left(\frac{y}{k}\right)$	$\underbrace{\ln y}_Y = \underbrace{(\ln a)}_M \underbrace{x}_X + \underbrace{\ln k}_C$	$\ln y$ vs x

1.1.9. Dimensional Analysis and Scaling

Dimensions of base SI physical quantities:: mass (**M**), length (**L**), time (**T**), temperature (**Θ**), moles (**N**), electric current (**I**) and luminous intensity (**J**).

The number of parameters in a problem is reduced by expressing the relationship in non-dimensional form. Quantities are $\{q\}_i$. Dimensionless groups are $\{\pi\}_i$.

Buckingham’s π Theorem: For N variables containing M dimensions, the number of dimensionless groups is at least $N - M$.

For a physical equation $f(q_1, q_2, \dots, q_N) = 0$ or $q_1 = f(q_2, \dots, q_N)$, where f may be unknown, there exists an equivalent dimensionless formulation $F(\pi_1, \pi_2, \dots, \pi_M) = 0$ or $\pi_1 = F(\pi_2, \dots, \pi_M)$, where each dimensionless group can be expressed as a power law function of a subset of the physical quantities:

$$\pi_i = (q_1)^{a_{i1}} (q_2)^{a_{i2}} \dots (q_N)^{a_{iN}}$$

For problems with 1 dimensionless group, the simple indicial method can be used, in which $F(\pi) = 0 \rightarrow \pi = C$, some dimensionless constant. Find the powers $\{a\}_i$ such that the expression for π has no dimensions.

Geometric similarity: where all length-ratio dimensionless groups are identical.

Dynamic similarity: where all independent dimensionless groups are identical.

Example with two dimensionless groups: power required to stir a fluid

($[P] = \mathbf{ML}^2\mathbf{T}^{-3}$: power, $[\rho] = \mathbf{ML}^{-3}$: density, $[\mu] = \mathbf{ML}^{-1}\mathbf{T}^{-1}$: viscosity, $[d] = \mathbf{L}$: diameter, $[\omega] = \mathbf{T}^{-1}$: angular speed)

$P = f(\rho, \mu, d, \omega)$: 5 quantities in 3 dimensions \rightarrow at least 2 dimensionless groups (1 dependent, at least 1 independent).

Typical dimensionless groups: $\frac{P}{\rho \omega^3 d^5} = F\left(\frac{\rho \omega d^2}{\mu}\right) \rightarrow$ graphing $\frac{P}{\rho \omega^3 d^5}$ vs $\frac{\rho \omega d^2}{\mu}$ specifies the relationship fully.

1.1.10. Methods of Proof

Proof by Deduction (Direct Proof): use of algebra to show a result.

Example: for all positive integers n , prove that $n^3 - n$ is always divisible by 6.

1. By factoring, $n^3 - n = n(n^2 - 1) = n(n + 1)(n - 1)$.
2. This is a product of three consecutive integer factors.
3. At least one of these factors must be a multiple of 3.
4. At least one of these factors must be a multiple of 2.
5. Since 2 and 3 are coprime, the product of the factors must be a multiple of 6.

Proof by Exhaustion: showing that a result is true for all individual cases.

Example: prove that the square of any positive integer cannot end in the digit 3.

1. The ending digit of a square number is determined only by the last digit of the number:
 $1^2 = 1$, $2^2 = 4$, $3^2 = 9$, $4^2 = 16$, $5^2 = 25$, $6^2 = 36$, $7^2 = 49$, $8^2 = 64$, $9^2 = 81$, $10^2 = 100$
2. Since none of the units digits' squares end in 3, no squared integer ends in 3.

Proof by Contradiction: assume the contrary, find it implies something that contradicts the original assumption, so the assumption must be false (the statement must be true).

Example: prove that $\sqrt{2}$ is irrational.

1. Assume that $\sqrt{2}$ is rational. Then it can be written as $\sqrt{2} = \frac{a}{b}$, where a and b are coprime integers.
2. By algebra, $a = b\sqrt{2} \Rightarrow a^2 = 2b^2 \Rightarrow a^2$ is an even number.
3. If a^2 is an even number, a must also be an even number (also proven by contradiction):
4. Therefore, we can let $a = 2n$ for some integer n . Then $a^2 = 4n^2$.
5. So $4n^2 = 2b^2 \Rightarrow b^2 = 2n^2 \Rightarrow b^2$ is even $\Rightarrow b$ is even.
6. Since a and b are both even, they share a common factor of 2.
7. However, it was assumed that a and b are coprime. This is a contradiction.
8. Therefore, the assumption must be false, so $\sqrt{2}$ is irrational.

Proof by Induction: verify a base case n_0 , state the inductive hypothesis and prove that its validity for case n implies validity for case $n + 1$, and conclude validity for all integers $n \geq n_0$.

Example: prove that $\frac{d^n}{dx^n} \left(e^x \sin \sqrt{3}x \right) = 2^n e^x \sin \left(\sqrt{3}x + \frac{n\pi}{3} \right)$, $n \geq 1$, $n \in \mathbb{N}$.

1. Base case: try $n = 1$.

$$\text{LHS} = \frac{d}{dx} (e^x \sin \sqrt{3}x) = e^x \sin \sqrt{3}x + \sqrt{3} e^x \cos \sqrt{3}x$$

$$\text{RHS} = 2 e^x \sin \left(\sqrt{3}x + \frac{\pi}{3} \right) = 2 e^x \left(\sin \sqrt{3}x \cos \frac{\pi}{3} + \cos \sqrt{3}x \sin \frac{\pi}{3} \right) = e^x \sin \sqrt{3}x + \sqrt{3} e^x \cos \sqrt{3}x$$

Since LHS = RHS, the base case is verified.

2. Assume that the statement is true for some integer $n \geq 1$. Then:

$$\frac{d^{n+1}}{dx^{n+1}} (e^x \sin \sqrt{3}x) = \frac{d}{dx} \left[2^n e^x \sin \left(\sqrt{3}x + \frac{n\pi}{3} \right) \right] = 2^n \left(e^x \sin \left(\sqrt{3}x + \frac{n\pi}{3} \right) + \sqrt{3} e^x \cos \left(\sqrt{3}x + \frac{n\pi}{3} \right) \right) = 2^{n+1} e^x \sin \left(\sqrt{3}x + \frac{(n+1)\pi}{3} \right).$$

3. Since true for $n = 1$, and truth for integer n implies truth for $n + 1$, it is true for all integers $n \geq 1$.

1.2. Algebraic Identities

1.2.1. Factorisation and Common Algebraic Manipulations

Factorisation

Difference of two squares:	$a^2 - b^2 = (a + b)(a - b)$
Sum of two squares (complex):	$a^2 + b^2 = (a + ib)(a - ib)$
Differences of two cubes:	$a^3 - b^3 = (a - b)(a^2 + ab + b^2)$
Sum of two cubes:	$a^3 + b^3 = (a + b)(a^2 - ab + b^2)$
Sum of two fourth-powers:	$a^4 + b^4 = (a^2 + \sqrt{2}ab + b^2)(a^2 - \sqrt{2}ab + b^2)$
Difference of two fourth-powers:	$a^4 - b^4 = (a^2 + b^2)(a + b)(a - b)$
Sophie-Germain identity:	$a^4 + 4b^4 = (a^2 + 2b^2 + 2ab)(a^2 + 2b^2 - 2ab)$

Expansion: derived by 'FOIL (first-outer-inner-last)'

$$(a + b)^2 = a^2 + 2ab + b^2 = (a - b)^2 + 4ab$$

$$(a + b)^3 = a^3 + 3a^2b + 3ab^2 + b^3$$

$$(a + b + c)^2 = a^2 + b^2 + c^2 + 2(ab + bc + ca)$$

$$(a + b + c)^3 = a^3 + b^3 + c^3 + 3(a^2b + a^2c + b^2a + b^2c + c^2a + c^2b) + 6abc$$

Completing the Square:
$$x^2 + bx + c = \left(x + \frac{b}{2}\right)^2 + \left(c - \frac{b^2}{4}\right)$$

Girard-Newton Identities: if $S = \{\alpha, \beta, \gamma, \dots\}$ are a set of K variables then

- $\sum \alpha^2 = (\sum \alpha)^2 - \sum \alpha\beta$
- $\sum \alpha^3 = (\sum \alpha)^3 - 3(\sum \alpha)(\sum \alpha\beta) + \sum \alpha\beta\gamma$
- $\sum \alpha^n = (\sum \alpha) \sum \alpha^{n-1} - (\sum \alpha\beta) \sum \alpha^{n-2} + (\sum \alpha\beta\gamma) \sum \alpha^{n-3} - \dots$

where sums over non-repeated combinations are $\sum \alpha = \alpha + \beta + \gamma + \dots$, $\sum \alpha\beta = \alpha\beta + \beta\gamma + \gamma\alpha + \dots$

These are commonly applied to roots of polynomials (see Vieta's formulas, Section 1.2.5).

If $K > n$ then the $\sum \alpha^{n-i}$ term will include terms such as $\sum \frac{1}{\alpha}$, $\sum \frac{1}{\alpha^2}$, etc. These can be

solved for by substituting previously found expressions recursively. The sum terminates with the $(-1)^{K-1} (\prod \alpha) (\sum \alpha^{n-K})$ term.

Componendo-Dividendo identities (corollaries of cross-multiplication):

If $\frac{a}{b} = \frac{c}{d}$, then $ad = bc$, $\frac{a+b}{b} = \frac{c+d}{d}$, $\frac{a-b}{b} = \frac{c-d}{d}$ and $\frac{a+b}{a-b} = \frac{c+d}{c-d}$.

1.2.2. Binomial Theorem and Trinomial Theorem

For any positive integer n , the binomial and trinomial theorems are, respectively,

$$(a + b)^n = \sum_{r=0}^n {}^n C_r a^r b^{n-r} \quad \text{where } {}^n C_r = \frac{n!}{r!(n-r)!}$$

$$(a + b + c)^n = \sum_{m=0}^n \sum_{k=0}^m {}^n C_m {}^m C_k a^{n-m} b^{m-k} c^k \quad \text{where } {}^n C_m {}^m C_k = \frac{n!}{(n-m)!(m-k)!k!}$$

1.2.3. Properties of Quadratic Polynomials

If $f(x) = ax^2 + bx + c$ then:

(the quadratic formula)

The turning point of $f(x)$ is at $x = -\frac{b}{2a}$. The roots of $f(x)$ are at $x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$.


For real coefficients, the discriminant, $\Delta = b^2 - 4ac$, determines the nature of the roots:

- If $\Delta > 0 \rightarrow$ two distinct real roots.
- If $\Delta < 0 \rightarrow$ two complex roots (complex conjugates).
- If $\Delta = 0 \rightarrow$ a single repeated real root.

1.2.4. Properties of Cubic Polynomials

For a cubic $f(x) = ax^3 + bx^2 + cx + d$, it can be transformed to a ‘depressed cubic’ using the substitution $x = t - \frac{b}{3a}$ into $f(t) = t^3 + pt + q$ where $p = \frac{3ac - b^2}{3a^2}$ and $q = \frac{2b^3 - 9abc + 27a^2d}{27a^3}$.

The inflection point of $f(x)$ is at $x = \frac{-b}{3a}$ i.e. $t = 0$.



The graph of $f(x)$ will have real turning points if $b^2 - 3ac > 0$.

In this case, there is rotational symmetry order 2 about the inflection point, with the lengths shown being

$$\epsilon = \frac{\sqrt{b^2 - 3ac}}{3a} \quad \delta = 2a \epsilon^3$$

For real coefficients, the discriminant is, $\Delta = \frac{4(b^2 - 3ac)^3 - (2b^3 - 9abc + 27a^2d)^2}{27a^2} = \frac{p^3}{27} + \frac{q^2}{4}$.

- If $\Delta > 0 \rightarrow$ three distinct real roots;
- If $\Delta < 0 \rightarrow$ one real root and two complex conjugate roots;
- If $\Delta = 0 \rightarrow$ a repeated root of multiplicity 2 or 3.

The transformed roots of $f(t) = 0$ are $t_1 = \sqrt[3]{-\frac{q}{2} + \Delta} + \sqrt[3]{-\frac{q}{2} - \Delta}$,

$$t_2 = \omega \sqrt[3]{-\frac{q}{2} + \Delta} + \omega^2 \sqrt[3]{-\frac{q}{2} - \Delta} \text{ and } t_3 = \omega^2 \sqrt[3]{-\frac{q}{2} + \Delta} + \omega \sqrt[3]{-\frac{q}{2} - \Delta}$$

where $x_i = t_i - \frac{b}{3a}$ and ω is a primitive cube root of unity:

$$\omega = \frac{-1 + i\sqrt{3}}{2}, \quad \omega^2 = \frac{-1 - i\sqrt{3}}{2}, \quad \omega^3 = 1$$

1.2.5. Relations Between Roots of Polynomials

For a polynomial $P(x)$ of degree n given by

$$P(x) = \sum_{r=0}^n a_r x^r = a_n x^n + a_{n-1} x^{n-1} + a_{n-2} x^{n-2} + \dots + a_1 x + a_0$$

the n roots of the equation $P(x) = 0$, denoted $\alpha_1, \alpha_2, \dots, \alpha_n$, are related by:

- Fundamental theorem of algebra: all α are complex, and number of roots = degree.
- Factor theorem: $P(x) = a_n(x - \alpha_1)(x - \alpha_2)\dots(x - \alpha_n)$
- Rational root theorem: $\alpha = \frac{p}{q} \in \mathbb{Q} \Rightarrow |p|$ is a factor of a_0 and $|q|$ is a factor of a_n
- Vieta formulas:

$$\begin{array}{cccc} \Sigma \alpha = -\frac{a_{n-1}}{a_n} & \Sigma \alpha\beta = \frac{a_{n-2}}{a_n} & \Sigma \alpha\beta\gamma = -\frac{a_{n-3}}{a_n} & \Pi \alpha = (-1)^n \frac{a_0}{a_n} \\ \text{(sum of roots)} & \text{(sum of product pairs)} & \text{(sum of product triplets)} & \text{(product of roots)} \end{array}$$

- Conjugacy:
 - (1) if all a_r are real and $\alpha_1 = u + iv$ is a root then $\alpha_2 = \alpha_1^* = u - iv$ is also a root.
 - (2) if all a_r are rational and $\alpha_1 = u \pm \sqrt{v}$ is an irrational root then $\alpha_2 = u \mp \sqrt{v}$ is also a root.

For a monic polynomial (monocubic, monoquartic, etc), $a_n = 1$, so these formulas simplify.

1.2.6. Division of Polynomials

Polynomials can be divided into the form $\frac{A(x)}{B(x)} = Q(x) + \frac{R(x)}{B(x)}$.

(A : dividend, B : divisor (same or higher degree than A), Q : quotient, R : remainder, deg: degree of)

- $\deg A \geq \deg B$
- $\deg Q = \deg A - \deg B$
- $\deg R < \deg B$

Factorisation: $R = 0$ iff B divides A i.e. B is a factor of A .

Techniques: examples to evaluate $\frac{x^3 - 2x^2 - 4}{x - 3} = x^2 + x + 3 + \frac{5}{x - 3}$:

<p>Polynomial Long Division</p> $ \begin{array}{r} \overline{) x^3 - 2x^2 + 0x - 4} \\ \underline{- x^3 + 3x^2} \\ x^2 + 0x - 4 \\ \underline{- x^2 + 3x} \\ 3x - 4 \\ \underline{- 3x + 9} \\ 5 \end{array} $	<p>Synthetic Division (requires $B = x - a$)</p> $ \begin{array}{r} \overline{) 1 \quad -2 \quad 0 \quad -4} \\ \underline{3 \quad 3 \quad 9} \\ 1 \quad 1 \quad 3 \quad 5 \end{array} $
--	---

If it is known that $R = 0$ beforehand, then an alternative method is to equate coefficients with a general polynomial e.g. $Q(x) = ax^2 + bx + c$, solving for each a , b and c .

1.2.7. Divisibility Rules for Integers

A positive integer n is divisible by...

- | | |
|---|--|
| <ul style="list-style-type: none"> • 2, if the last digit is even • 3, if the sum of the digits is divisible by 3 • 4, if the last two digits form a number divisible by 4 • 5, if the last digit is 5 or 0 | <ul style="list-style-type: none"> • 6, if the number is divisible by both 2 and 3 • 7, if subtracting twice the last digit from the rest of the number gives a multiple of 7 • 8, if the last three digits form a number divisible by 8 • 9, if the sum of the digits is divisible by 9 |
|---|--|

These results can be derived from modular arithmetic.

1.2.8. The Triangle Inequality

For any real or complex numbers a, b : $|a + b| \leq |a| + |b|$

This is also valid for any vectors \mathbf{a} and \mathbf{b} : $|\mathbf{a} + \mathbf{b}| \leq |\mathbf{a}| + |\mathbf{b}|$

The Inverse Triangle Inequality: $|a - b| \geq ||a| - |b||$

$$|\mathbf{a} - \mathbf{b}| \geq ||\mathbf{a}| - |\mathbf{b}||$$

1.2.9. The Cauchy-Schwarz Inequality

For any n -dimensional vectors \mathbf{u} and \mathbf{v} : $\left(\sum_{i=1}^n u_i v_i\right)^2 \leq \left(\sum_{i=1}^n u_i^2\right)\left(\sum_{i=1}^n v_i^2\right)$

For $n = 2$, this can be stated as $(ac + bd)^2 \leq (a^2 + b^2)(c^2 + d^2)$

1.2.10. The Harmonic-Geometric-Arithmetic-Quadratic Mean Inequalities

For any real-valued set of n positive variables x with i th variable x_i :

$$0 < \frac{n}{\sum_{i=1}^n \frac{1}{x_i}} \leq \sqrt[n]{\prod_{i=1}^n x_i} \leq \frac{1}{n} \sum_{i=1}^n x_i \leq \sqrt{\frac{1}{n} \sum_{i=1}^n x_i^2}$$

$$\text{HM} \leq \text{GM} \leq \text{AM} \leq \text{QM}$$

For $n = 2$, this can be stated as: given $a, b > 0$, we have

$$\frac{2ab}{a+b} \leq \sqrt{ab} \leq \frac{a+b}{2} \leq \sqrt{\frac{a^2+b^2}{2}} \quad \text{or} \quad \frac{8a^2b^2}{(a+b)^2} \leq 2ab \leq \frac{1}{2}(a+b)^2 \leq a^2 + b^2$$

Weighted AM-GM inequality: $\frac{x_1 y_1 + x_2 y_2 + \dots + x_n y_n}{x_1 + x_2 + \dots + x_n} \geq \left(y_1^{x_1} y_2^{x_2} \dots y_n^{x_n}\right)^{\frac{1}{x_1 + x_2 + \dots + x_n}}$

1.2.11. Muirhead's Inequality

For two N -length sequences $\{a_i\}$ and $\{b_i\}$, the notation $\{a_i\} > \{b_i\}$ means that $\{a_i\}$ 'majorises' $\{b_i\}$, which is defined by

$$\sum_{i=1}^n a_i \geq \sum_{i=1}^n b_i \quad \forall n \in [1, \dots, N] \quad \text{and} \quad \sum_{i=1}^N a_i = \sum_{i=1}^N b_i \quad \Leftrightarrow \quad \{a_i\} > \{b_i\}.$$

Muirhead's inequality: if $\{a_i\} > \{b_i\}$ then $\sum_{\text{sym } x} x_1^{a_1} x_2^{a_2} \dots x_N^{a_N} \geq \sum_{\text{sym } x} x_1^{b_1} x_2^{b_2} \dots x_N^{b_N}$

where the sum is over all permutations of any chosen set of variables.

Example: the sequence (5, 1) majorises (4, 2). Therefore $x^5 y + x y^5 \geq x^4 y^2 + x^2 y^4$.

1.2.12. Schur's Inequality

For all real a, b, c and all $r > 0$,

$$a^r(a-b)(a-c) + b^r(b-a)(b-c) + c^r(c-a)(c-b) \geq 0.$$

Case $r = 1$: $a^3 + b^3 + c^3 + 3abc \geq a^2b + a^2c + b^2a + b^2c + c^2a + c^2b$

Case $r = 2$: $a^4 + b^4 + c^4 + abc(a+b+c) \geq a^3b + a^3c + b^3a + b^3c + c^3a + c^3b$

Generalisation: for all real a, b, c, x, y, z , where $a \geq b \geq c$ and $(x \geq y \geq z \text{ or } z \geq y \geq x)$, and some convex or monotonic function $f: \mathbb{R} \rightarrow \mathbb{R}^+$, and some constant $k \in \mathbb{Z}^+$,

$$f(x)(a-b)^k(a-c)^k + f(y)(b-a)^k(b-c)^k + f(z)(c-a)^k(c-b)^k \geq 0.$$

1.2.13. Jordan's Inequality

For all $0 \leq x \leq \frac{\pi}{2}$, we have $\frac{2}{\pi}x \leq \sin x \leq x$ i.e. $\frac{2}{\pi} \leq \text{sinc } x \leq 1$

1.2.14. Fermat's Last Theorem

If $n \geq 3$, then there are no positive integer solutions (a, b, c) to the equation $a^n + b^n = c^n$.

(The case of $n = 2$ has infinitely many solutions - the Pythagorean triples.)

1.2.15. Diophantine Equations

The linear Diophantine equation $ax + by = c$ has integer solutions (x, y) iff (a, b, c) are all integers and $c \mid \gcd(a, b)$. If (x, y) is a solution then the other solutions are given by $(x + kv, y - ku)$ where k is an arbitrary integer, and $u = a / \gcd(a, b)$ and $v = b / \gcd(a, b)$.

Proofs on the solutions to Diophantine equations typically include:

- Chinese Remainder Theorem (CRT): see Section 1.6.5.
- Infinite descent: assume a supposedly minimal solution exists, show this implies the existence of a smaller solution, which is a contradiction (similar: Vieta jumping)

1.3. Trigonometric Identities

1.3.1. Trigonometric Functions and Identities

$$\sin(x + y) = \sin x \cos y + \cos x \sin y$$

$$\sin(x - y) = \sin x \cos y - \cos x \sin y$$

$$\cos(x + y) = \cos x \cos y - \sin x \sin y$$

$$\cos(x - y) = \cos x \cos y + \sin x \sin y$$

$$\tan(x + y) = \frac{\tan x + \tan y}{1 - \tan x \tan y}$$

$$\tan(x - y) = \frac{\tan x - \tan y}{1 + \tan x \tan y}$$

$$\sin 2x = 2 \sin x \cos x$$

$$\sin 3x = 3 \sin x - 4 \sin^3 x$$

$$\sin \frac{x}{2} = \pm \sqrt{\frac{1 - \cos x}{2}}$$

$$\cos 2x = 2 \cos^2 x - 1$$

$$\cos 3x = 4 \cos^3 x - 3 \cos x$$

$$\cos \frac{x}{2} = \pm \sqrt{\frac{1 + \cos x}{2}}$$

$$= 1 - 2 \sin^2 x$$

$$= \cos^2 x - \sin^2 x$$

$$\tan 2x = \frac{2 \tan x}{1 - \tan^2 x}$$

$$\tan 3x = \frac{3 \tan x - \tan^3 x}{1 - 3 \tan^2 x}$$

$$\tan \frac{x}{2} = \pm \sqrt{\frac{1 - \cos x}{1 + \cos x}}$$

$$= \frac{\sin x}{1 + \cos x} = \frac{1 - \cos x}{\sin x}$$

$$\sec x = \frac{1}{\cos x}$$

$$\csc x = \frac{1}{\sin x}$$

$$\cot x = \frac{1}{\tan x}$$

$$\sin^2 x + \cos^2 x = 1$$

$$\sec^2 x = 1 + \tan^2 x$$

$$\csc^2 x = 1 + \cot^2 x$$

$$\sin(-x) = -\sin x$$

$$\cos(-x) = \cos x$$

$$\tan(-x) = -\tan x$$

$$\sin\left(\frac{\pi}{2} - x\right) = \cos x$$

$$\cos\left(\frac{\pi}{2} - x\right) = \sin x$$

$$\tan\left(\frac{\pi}{2} - x\right) = \cot x$$

$$\sin(\pi - x) = \sin x$$

$$\cos(\pi - x) = -\cos x$$

$$\tan(\pi - x) = -\tan x$$

$$\sin\left(x + \frac{\pi}{2}\right) = \cos x$$

$$\cos\left(x + \frac{\pi}{2}\right) = -\sin x$$

$$\tan\left(x + \frac{\pi}{2}\right) = \tan x$$

$$\sin(x + \pi) = -\sin x$$

$$\cos(x + \pi) = -\cos x$$

$$\tan(x + \pi) = \tan x$$

$$\sin x + \sin y = 2 \sin \frac{x+y}{2} \cos \frac{x-y}{2}$$

$$\sin x - \sin y = 2 \cos \frac{x+y}{2} \sin \frac{x-y}{2}$$

$$\cos x + \cos y = 2 \cos \frac{x+y}{2} \cos \frac{x-y}{2}$$

$$\cos x - \cos y = -2 \sin \frac{x+y}{2} \sin \frac{x-y}{2}$$

$$2 \sin x \sin y = \cos(x - y) - \cos(x + y)$$

$$2 \cos x \cos y = \cos(x + y) + \cos(x - y)$$

$$2 \sin x \cos y = \sin(x + y) + \sin(x - y)$$

1.3.2. Hyperbolic Functions and Identities

$$\sinh x = \frac{1}{2}(e^x - e^{-x})$$

$$\cosh x = \frac{1}{2}(e^x + e^{-x})$$

$$\tanh x = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

$$\operatorname{csch} x = \frac{1}{\sinh x}$$

$$\operatorname{sech} x = \frac{1}{\cosh x}$$

$$\operatorname{coth} x = \frac{1}{\tanh x}$$

$$\sinh(x + y) = \sinh x \cosh y + \cosh x \sinh y$$

$$\sinh(x - y) = \sinh x \cosh y - \cosh x \sinh y$$

$$\cosh(x + y) = \cosh x \cosh y + \sinh x \sinh y$$

$$\cosh(x - y) = \cosh x \cosh y - \sinh x \sinh y$$

$$\tanh(x + y) = \frac{\tanh x + \tanh y}{1 + \tanh x \tanh y}$$

$$\tanh(x - y) = \frac{\tanh x - \tanh y}{1 - \tanh x \tanh y}$$

$$\sinh 2x = 2 \sinh x \cosh x$$

$$\sinh 3x = 3 \sinh x + 4 \sinh^3 x$$

$$\sinh \frac{x}{2} = \pm \sqrt{\frac{\cosh x - 1}{2}}$$

$$\cosh 2x = 2 \cosh^2 x - 1$$

$$\cosh 3x = 4 \cosh^3 x - 3 \cosh x$$

$$\cosh \frac{x}{2} = \sqrt{\frac{\cosh x + 1}{2}}$$

$$= 1 + 2 \sinh^2 x$$

$$= \cosh^2 x + \sinh^2 x$$

$$\tanh 2x = \frac{2 \tanh x}{1 + \tanh^2 x}$$

$$\tanh 3x = \frac{3 \tanh x - \tanh^3 x}{1 - 3 \tanh^2 x}$$

$$\tanh \frac{x}{2} = \pm \sqrt{\frac{\cosh x - 1}{\cosh x + 1}}$$

$$= \frac{\sinh x}{1 + \cosh x}$$

$$= \frac{\cosh x - 1}{\sinh x}$$

$$\operatorname{sech} x = \frac{1}{\cosh x}$$

$$\operatorname{csch} x = \frac{1}{\sinh x}$$

$$\operatorname{coth} x = \frac{1}{\tanh x}$$

$$\cosh^2 x - \sinh^2 x = 1$$

$$\operatorname{sech}^2 x = 1 - \tanh^2 x$$

$$\operatorname{coth}^2 x - \operatorname{csch}^2 x = 1$$

$$\sinh(-x) = -\sinh x$$

$$\cosh(-x) = \cosh x$$

$$\tanh(-x) = -\tanh x$$

$$\sinh x + \sinh y = 2 \sinh \frac{x+y}{2} \cosh \frac{x-y}{2}$$

$$\sinh x - \sinh y = 2 \cosh \frac{x+y}{2} \sinh \frac{x-y}{2}$$

$$\cosh x + \cosh y = 2 \cosh \frac{x+y}{2} \cosh \frac{x-y}{2}$$

$$\cosh x - \cosh y = -2 \sinh \frac{x+y}{2} \sinh \frac{x-y}{2}$$

$$2 \sinh x \sinh y = \cosh(x+y) - \cosh(x-y)$$

$$2 \cosh x \cosh y = \cosh(x+y) + \cosh(x-y)$$

$$2 \sinh x \cosh y = \sinh(x+y) + \sinh(x-y)$$

Osborn's Rule: for a trig identity, for a product of sines, change the sign, to get the hyperbolic identity.

1.3.3. Inverse Trigonometric Functions

$$\sec^{-1} x = \cos^{-1} \frac{1}{x}$$

$$\csc^{-1} x = \sin^{-1} \frac{1}{x}$$

$$\cot^{-1} x = \tan^{-1} \frac{1}{x}$$

$$\sin^{-1} x + \cos^{-1} x = \frac{\pi}{2}$$

$$\tan^{-1} x + \cot^{-1} x = \frac{\pi}{2}$$

$$\sec^{-1} x + \csc^{-1} x = \frac{\pi}{2}$$

$$\sin^{-1} x \pm \sin^{-1} y = \sin^{-1} \left(x\sqrt{1-y^2} \pm y\sqrt{1-x^2} \right)$$

$$\cos^{-1} x \pm \cos^{-1} y = \cos^{-1} \left(xy \mp \sqrt{1-x^2}\sqrt{1-y^2} \right)$$

$$\sin^{-1} x \pm \cos^{-1} y = \sin^{-1} \left(xy \pm \sqrt{1-x^2}\sqrt{1-y^2} \right) = \cos^{-1} \left(y\sqrt{1-x^2} \mp x\sqrt{1-y^2} \right)$$

$$\tan^{-1} x \pm \tan^{-1} y = \tan^{-1} \left(\frac{x \pm y}{1 \mp xy} \right)$$

1.3.4. Inverse Hyperbolic Functions

$$\sinh^{-1} x = \ln(x + \sqrt{x^2 + 1})$$

$$\cosh^{-1} x = \ln(x + \sqrt{x^2 - 1})$$

$$\tanh^{-1} x = \frac{1}{2} \ln \frac{1+x}{1-x}$$

$$\operatorname{csch}^{-1} x = \sinh^{-1} \frac{1}{x}$$

$$\operatorname{sech}^{-1} x = \cosh^{-1} \frac{1}{x}$$

$$\operatorname{coth}^{-1} x = \tanh^{-1} \frac{1}{x}$$

$$\sinh^{-1} x \pm \sinh^{-1} y = \sinh^{-1} \left(x\sqrt{1+y^2} \pm y\sqrt{1+x^2} \right)$$

$$\cosh^{-1} x \pm \cosh^{-1} y = \cosh^{-1} \left(xy \pm \sqrt{x^2-1}\sqrt{y^2-1} \right)$$

$$\sinh^{-1} x + \cosh^{-1} y = \sinh^{-1} \left(xy + \sqrt{1-x^2}\sqrt{y^2-1} \right) = \cosh^{-1} \left(y\sqrt{1+x^2} + x\sqrt{y^2-1} \right)$$

$$\tanh^{-1} x \pm \tanh^{-1} y = \tanh^{-1} \left(\frac{x \pm y}{1 \pm xy} \right)$$

1.4. Sequences and Series

1.4.1. Arithmetic Sequence

For an arithmetic sequence with first term a and common difference d , terms are

$$a, a + d, a + 2d, a + 3d, \dots$$

n th term: $u_n = a + (n - 1)d$

n th partial sum: $S_n = \sum_{r=1}^n u_n = \frac{n}{2}(2a + (n - 1)d)$

1.4.2. Geometric Sequence

For a geometric sequence with first term a and common ratio r , terms are

$$a, ar, ar^2, ar^3, \dots$$

n th term: $u_n = ar^{n-1}$

n th partial sum: $S_n = \sum_{r=1}^n u_n = \frac{a(1 - r^n)}{1 - r}$

Infinite sum: $S_\infty = \sum_{r=1}^{\infty} u_n = \frac{a}{1 - r}$ convergent if $0 < |r| < 1$

1.4.3. Harmonic Sequence

In a harmonic sequence, terms are the reciprocal of an arithmetic sequence.

$$\frac{1}{a}, \frac{1}{a + d}, \frac{1}{a + 2d}, \frac{1}{a + 3d}, \dots$$

n th term: $u_n = \frac{1}{a + (n - 1)d}$

n th partial sum: $S_n = \sum_{r=1}^n u_n \leq \ln(a + (n - 1)d) + \gamma + \frac{1}{2(a + (n - 1)d)}$

Infinite sum: diverges, always

1.4.4. Arithmetico-Geometric Sequence

In an arithmetico-geometric sequence, terms are a product of an arithmetic and geometric sequence.

$$ab, (a + d)br, (a + 2d)br^2, (a + 3d)br^3, \dots$$

n th term: $u_n = (a + (n - 1)d)br^{n-1}$

n th partial sum: $S_n = \sum_{r=1}^n u_r = \frac{ab - (a + nd)br^n}{1 - r} + \frac{dbr(1 - r^n)}{(1 - r)^2}$

Infinite sum: $S_\infty = \sum_{r=1}^{\infty} u_r = \frac{ab}{1 - r} + \frac{dbr}{(1 - r)^2}$ convergent if $|r| < 1$

1.4.5. Partial Sums of Series

$$\sum_{r=1}^n r = \frac{1}{2}n(n + 1) \qquad \sum_{r=1}^n r^2 = \frac{1}{6}n(n + 1)(2n + 1) \qquad \sum_{r=1}^n r^3 = \frac{1}{4}n^2(n + 1)^2$$

$$\sum_{r=1}^n \frac{1}{r} \geq \ln n + \gamma + O(n^{-1}) \quad \text{where } \gamma = 0.5772156649\dots \text{ is the Euler-Mascheroni constant.}$$

$$\sum_{r=0}^n {}^nC_r = 2^n \qquad \sum_{r=0}^n {}^pC_r \times {}^qC_{n-r} = {}^{p+q}C_n \qquad \text{(Chu-Vandermonde identity)}$$

$$\sum_{r=1}^n \cos r\theta = \frac{\sin \frac{n\theta}{2} \cos \frac{(n+1)\theta}{2}}{\sin \frac{\theta}{2}} \qquad \sum_{r=1}^n \sin r\theta = \frac{\sin \frac{n\theta}{2} \sin \frac{(n+1)\theta}{2}}{\sin \frac{\theta}{2}} \qquad \text{Often useful: } u_n = S_n - S_{n-1}.$$

1.4.6. Infinite Sums of Series

The numerical results below require advanced techniques to prove (e.g. special functions).

$$\sum_{r=1}^{\infty} \frac{1}{r} \text{ diverges} \qquad \sum_{r=1}^{\infty} \frac{1}{r^2} = \frac{\pi^2}{6} \qquad \sum_{r=1}^{\infty} \frac{1}{r^3} = \zeta(3) \approx 1.2021 \qquad \sum_{r=1}^{\infty} \frac{1}{r^4} = \frac{\pi^4}{90}$$

$$\sum_{r=1}^{\infty} \frac{(-1)^{r+1}}{r} = \ln 2 \qquad \sum_{r=1}^{\infty} \frac{(-1)^{r+1}}{r^2} = \frac{\pi^2}{12} \qquad \sum_{r=1}^{\infty} \frac{(-1)^{r+1}}{r^3} = \frac{3}{4}\zeta(3) \approx 0.90154 \qquad \sum_{r=1}^{\infty} \frac{(-1)^{r+1}}{r^4} = \frac{7\pi^4}{720}$$


$$\sum_{r=-\infty}^{\infty} \frac{1}{a + br} = \frac{\pi}{b} \cot \frac{\pi a}{b} \qquad \sum_{r=-\infty}^{\infty} \frac{1}{a + br^2} = \frac{\pi}{\sqrt{ab}} \coth \pi \sqrt{\frac{a}{b}} \qquad \sum_{r=0}^{\infty} \frac{1}{1 + r^2} = \frac{1 + \pi \coth \pi}{2}$$

1.4.7. Common Techniques (Tests) for Proving Convergence or Divergence of Series

- **Divergence Test:** for $\sum_{n=1}^{\infty} a_n$, evaluate $\lim_{n \rightarrow \infty} a_n$.
 If $\lim_{n \rightarrow \infty} a_n \neq 0$ then the series **diverges**. Otherwise, inconclusive.
- **Geometric Series:** for $\sum_{n=1}^{\infty} a r^{n-1}$, evaluate r .
 If $|r| < 1$ then the series **converges** to $\frac{a}{1-r}$. Otherwise the series **diverges**.
- **Power Series:** for $\sum_{n=1}^{\infty} \frac{1}{n^p}$, evaluate p .
 If $p > 1$ then the series **converges** to $\zeta(p)$. Otherwise the series **diverges** (in the usual sense).
- **Comparison Test:** for $\sum_{n=1}^{\infty} a_n$ with all $a_n \geq 0$, compare to a known series $\sum_{n=1}^{\infty} b_n$.
 If there exists some N such that $a_n \leq b_n$ i.e. $\frac{a_n}{b_n} \leq 1$ for all $n > N$, and $\sum_{n=1}^{\infty} b_n$ converges then the series **converges**. Otherwise the series **diverges**.
- **Limit Comparison Test:** for $\sum_{n=1}^{\infty} a_n$ (all $a_n > 0$), compare to known $\sum_{n=1}^{\infty} b_n$ and evaluate $L = \lim_{n \rightarrow \infty} \frac{a_n}{b_n}$.
 If L is finite and nonzero then $\sum_{n=1}^{\infty} a_n$ and $\sum_{n=1}^{\infty} b_n$ either both converge or both diverge.
 If $L = 0$ and $\sum_{n=1}^{\infty} b_n$ converges then the series **converges**.
 If $L \rightarrow \infty$ and $\sum_{n=1}^{\infty} b_n$ diverges then the series **diverges**.
- **Integral Test:** for $\sum_{n=1}^{\infty} a_n$, if $f(n) = a_n$ is positive and decreasing for all $n > N$, evaluate $I = \int_N^{\infty} f(x) dx$.
 If I is finite then the series **converges**. If I diverges then the series **diverges**.
- **Alternating Series Test:** for $\sum_{n=1}^{\infty} (-1)^n a_n$, if $a_{n+1} \leq a_n$ for all $n \geq 1$ and $\lim_{n \rightarrow \infty} a_n = 0$ then the series **converges**.
- **Ratio Test and Root Test:** for any $\sum_{n=1}^{\infty} a_n$, evaluate either $\rho = \lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right|$ or $\rho = \lim_{n \rightarrow \infty} |a_n|^{1/n}$.
 If $0 \leq \rho < 1$ then the series **converges**. If $\rho > 1$ then the series **diverges**. If $\rho = 1$ then inconclusive.

1.5. Complex Numbers

1.5.1. Complex Number Definition



Any complex number $z \in \mathbb{C}$ can be represented as $z = x + yi = r(\cos \theta + i \sin \theta) = r e^{i\theta} = r \text{cis } \theta$ (Euler's formula)

The components are related and defined by $x = \text{Re}(z); y = \text{Im}(z); r = |z|; \theta = \arg z; -\pi < \theta \leq \pi$
 (r : magnitude, θ : argument/phase, x : real part, y : imaginary part)

1.5.2. Complex Conjugate

$$z^* = x - iy = r e^{-i\theta} = r \text{cis}(-\theta)$$

$$zz^* = |z|^2 \text{ is purely real}$$

$$(z_1 + z_2)^* = z_1^* + z_2^*$$

$$(z_1 z_2)^* = z_1^* z_2^*$$

$$\left(\frac{z_1}{z_2}\right)^* = \frac{z_1^*}{z_2^*}$$

1.5.3. Cyclic Nature of Exponentials and Logarithms

- Periodicity: $e^{2\pi ni} = 1, z = r e^{i(\theta + 2\pi n)}$, for any integer n
- De Moivre's theorem: $z = e^{i\theta} \rightarrow z^a = (\cos \theta + i \sin \theta)^a = \exp(ia(\theta + 2\pi n)) = \cos a\theta + i \sin a\theta$
- Natural logarithm: $z = r e^{i\theta} \rightarrow \ln z = \ln r + i(\theta + 2\pi n)$
- n th roots of unity: $\omega_k = e^{i \frac{2\pi k}{n}}, \omega_k^n = 1, \omega_{k+n} = \omega_k, \omega_1^k = \omega_k, \sum_{r=0}^{n-1} \omega^r = 0$

The roots of unity form a regular n -polygon around the origin.

1.5.4. Trigonometric and Hyperbolic Functions

$$\sin x = \frac{e^{ix} - e^{-ix}}{2i}$$

$$\cos x = \frac{e^{ix} + e^{-ix}}{2}$$

$$\tan x = \frac{e^{ix} - e^{-ix}}{e^{ix} + e^{-ix}}$$

$$z = e^{i\theta} \rightarrow z + \frac{1}{z} = 2 \cos \theta$$

$$z - \frac{1}{z} = 2i \sin \theta$$

$$\frac{z^2 - 1}{z^2 + 1} = i \tan \theta$$

$$\sin ix = i \sinh x$$

$$\cos ix = \cosh x$$

$$\sinh ix = i \sin x$$

$$\cosh ix = \cos x$$

$$\sin(x \pm iy) = \sin x \cosh y \pm i \cos x \sinh y$$

$$\cos(x \pm iy) = \cos x \cosh y \mp i \sin x \sinh y$$

$$\sinh(x \pm iy) = \sinh x \cos y \pm i \cosh x \sin y$$

$$\cosh(x \pm iy) = \cosh x \cos y \pm i \sinh x \sin y$$

$$\sin^{-1} z = \ln \left(iz + \sqrt{1 - z^2} \right) = i \sinh^{-1}(-iz)$$

$$\tan^{-1} z = \frac{i}{2} \ln \frac{1 - iz}{1 + iz} = i \tanh^{-1}(-iz)$$

1.5.5. Exponentiation of Complex Numbers

Let $z = a + bi = r e^{i\theta}$. Then, for integers n , the value of $f(z)$ is given by $R e^{i\phi}$, where:

$f(z)$	$R = f(z) $	$\phi = \arg f(z)$
z^i	$e^{-(\theta + 2\pi n)}$	$\ln r$
i^z	$e^{-\frac{4n+1}{2}\pi r \sin \theta}$	$\frac{4n+1}{2}\pi r \cos \theta$
z^{c+di}	$r^c e^{-d(\theta + 2\pi n)}$	$d \ln r + c(\theta + 2\pi n)$

The principal value is for when $n = 0$.

$i^{\pm i} = e^{\mp\pi/2}$ is purely real.

1.5.6. Root of Unity Filter

For a series $f(x) = \sum_{k=0}^{\infty} a_k x^k$, the value of $\sum_{k=0}^{N-1} f(\omega^k) = N \sum_{k=0}^{\infty} a_{kN}$ where $\omega = e^{\frac{2\pi}{N}i}$ for $N \geq 1$.

This is often useful when working with generating functions (series multisection) whose coefficients are periodic modulo N .

1.6. Discrete Mathematics and Abstract Algebra

1.6.1. Binary Operators

Let $*$ be a binary operator. We say that $*$ is

- Commutative: $a * b = b * a$
- Associative: $(a * b) * c = a * (b * c)$
- Distributive (over +): $a * (b + c) = a * b + a * c$

1.6.2. Axioms of Group Theory

A set S and an operation $*$ form a group $(S, *)$ if and only if all of the following are true:

- Closure: $*$ is a binary operation on S i.e. $a * b$ is in S for every $a, b \in S$.
- Identity: there exists exactly one element a in S such that $z * a = z$ for all z in S .
- Associative: $(a * b) * c = a * (b * c)$ for all $a, b, c \in S$.
- Inverse: every element a in S has exactly one corresponding element b in S such that $a * b$ equals the identity element for the group.

An Abelian group is a group in which $*$ is commutative in S .

Equivalently, the group is Abelian if its Cayley table is symmetric about the leading diagonal.

Examples of groups and their identity elements:

- If S is the integers and $*$ is multiplication, then the identity element is 1.
- If S is the real numbers and $*$ is addition, then the identity element is 0.
- If S is a set of geometric transformations and $*$ is composition, then the identity element is the transformation which does nothing (i.e. the identity matrix, if represented by affine transformation matrices).

For finite field groups, see Section 8.8.10.

For point and space (symmetry) groups and their character tables, see Section 13.2.8.

1.6.3. Axioms of Ring Theory

A set S and two operations $+$ and $*$ form a ring $(S, +, *)$ if and only if all of the following are true:

- $(S, +)$ is an Abelian group.
- S is closed under $*$.
- $*$ is associative.
- $*$ is distributive over $+$.

A commutative ring is a ring in which $*$ is commutative on S .

A field is a ring in which every nonzero element in S has an inverse element under $*$. Equivalently, a field is a group under both $+$ and $*$.

1.6.5. Modular Arithmetic

Definition of modulo operator: $a \equiv b \pmod{n} \Leftrightarrow \frac{a}{n}$ and $\frac{b}{n}$ have the same remainder, where a, b and n are all integers.

If $a \equiv b \pmod{n}$ then $\frac{a-b}{n}$ is an integer.

Euler's totient function: $\phi(n)$ is the number of integers between 1 and n which are coprime with n (no common divisors except 1):

$$\phi(n) = \prod_{a \in \mathbb{N}: a < n, \gcd(a, n) = 1} a = n \prod_{p \in P: p | n} \left(1 - \frac{1}{p}\right)$$

Fermat's little theorem: $a^p \equiv a \pmod{p}$ for prime p which does not divide a

Euler's theorem: $a^{\phi(n)} \equiv 1 \pmod{n}$ for coprime a and n

Wilson's theorem: $(p-1)! \equiv -1 \pmod{p} \Leftrightarrow p$ is prime

Chinese remainder theorem:

A system of N congruences $\bigcap_{i=1}^N \{x \equiv a_i \pmod{n_i}\}$ where all n_i are pairwise coprime has solutions x , any two of which x_i and x_j satisfy $x_i \equiv x_j \pmod{N}$.

The residue class of a modulo n is the set $\overline{a}_n = \{\dots, a-2n, a-n, a, a+n, a+2n, \dots\}$.

The ring of integers modulo n is the set of all residue classes modulo n , represented by

$$\mathbb{Z}/n\mathbb{Z} = \{\overline{a}_n \mid a \in \mathbb{Z}\} = \{\overline{0}_n, \overline{1}_n, \overline{2}_n, \dots, \overline{n-1}_n\}$$

and when $n = 0$, this ring is isomorphic to \mathbb{Z} since $\overline{a}_0 = \{a\}$.

This ring is commutative as $\overline{a}_n \pm \overline{b}_n = \overline{(a \pm b)}_n$ and $\overline{a}_n \overline{b}_n = \overline{(ab)}_n$.

$\mathbb{Z}/n\mathbb{Z}$ is a finite field if and only if n is prime.

1.6.6. Kuratowski's Theorem for Planarity of Graphs

1.6.7. Minimax Cut-Flow Theorem for Networks

1.6.8. Simplex Algorithm for Linear Programming

1.6.9. Lagrange Multipliers for Nonlinear Programming

A typical problem is stated as “Minimise / Maximise $f(\mathbf{x})$ subject to $g(\mathbf{x}) = 0$ ”, where \mathbf{x} is a vector of n scalar independent unknown variables. The Lagrangian is defined as $L(\mathbf{x}, \lambda) = f(\mathbf{x}) - \lambda g(\mathbf{x})$. The solution is given by $\nabla L(\mathbf{x}, \lambda) = 0$ (a system of $n + 1$ equations). Note that these are critical points, not necessarily extrema (may be saddle points).

For multiple constraints, formulated as “Minimise / Maximise $f(\mathbf{x})$ subject to $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ ”, where \mathbf{g} is a vector-valued function for each of the m constraints, the Lagrangian is

$L(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) - \boldsymbol{\lambda} \cdot \mathbf{g}(\mathbf{x}) = f(\mathbf{x}) - \sum_{i=1}^m \lambda_i g_i(\mathbf{x})$ where $\boldsymbol{\lambda}$ is a vector of m unknown multipliers. The solution is again given by $\nabla L(\mathbf{x}, \boldsymbol{\lambda}) = 0$ (a system of $n + m$ equations).

In the Hamiltonian formulation, $H(\mathbf{x}) = f(\mathbf{x}) + \boldsymbol{\lambda} \cdot \mathbf{g}(\mathbf{x})$, which ensures **minima**.

1.6.10. Game Theory for Zero-Order and Higher-Order Games


1.7. Special Functions and Identities

1.7.1. Gamma Function, $\Gamma(x)$ and Digamma Function, $\psi(x)$

- Gamma function as a generalised factorial: $\Gamma(x) = (x - 1)!$ i.e. $x! = x \Gamma(x)$
- Gamma function as an integral:
$$\Gamma(z) = \int_0^{\infty} t^{z-1} e^{-t} dt \quad \text{for } \operatorname{Re}(z) > 0$$
- Reflection identity:
$$\Gamma(z) \Gamma(1 - z) = \frac{\pi}{\sin \pi z} \quad \text{for non-integer } z$$
- Half-integer identity:
$$\Gamma(z) \Gamma\left(z + \frac{1}{2}\right) = 2^{1-2z} \sqrt{\pi} \Gamma(2z) \quad \text{for non-integer } z$$
- Useful exact values:
$$\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi} \quad \Gamma\left(-\frac{1}{2}\right) = -2\sqrt{\pi}$$
- Digamma function:
$$\psi(x) = \frac{d}{dx} \ln \Gamma(x) = \frac{\Gamma'(x)}{\Gamma(x)}$$
- Reflection identity:
$$\psi(1 - x) - \psi(x) = \frac{\pi}{\tan \pi x} \quad \text{for non-integer } x$$
- Integer identity:
$$\psi(x + 1) = \psi(x) + \frac{1}{x}$$

1.7.3. Error Function, $\operatorname{erf} x$

- Error function:
$$\operatorname{erf} x = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$
- Relation to Normal distribution:
$$\operatorname{erf} x = 2 \Phi(\sqrt{2}x) - 1 \quad (\Phi: \text{standard normal c.d.f.})$$
- Complementary error function: $\operatorname{erfc} z = 1 - \operatorname{erf} z$
- Imaginary error function: $\operatorname{erfi} z = -i \operatorname{erf} iz$



Graph of $y = \operatorname{erf} x$

1.7.4. Beta Function, $B(x, y)$

- Beta function:
$$B(x, y) = \int_0^1 t^{x-1} (1 - t)^{y-1} dt$$
- Relation to gamma function:
$$B(x, y) = \frac{\Gamma(x) \Gamma(y)}{\Gamma(x + y)}$$
- Pascal's identity:
$$B(x, y) = B(x, y + 1) + B(x + 1, y)$$
- For integers m, n :
$$B(m, n) = \frac{m + n}{mn} \times \frac{1}{C_m^{m+n}}$$

1.7.5. Hypergeometric Functions, including ${}_2F_1(a, b; c; z)$

- Gaussian hypergeometric function: ${}_2F_1(a, b; c; z) = \sum_{n=0}^{\infty} \frac{(a)_n (b)_n}{(c)_n} \frac{z^n}{n!}$, for $|z| < 1$

This series terminates if b or c is an integer, forming a binomial series.

- Euler's integral formula: $B(b, c - b) {}_2F_1(a, b; c; z) = \int_0^1 x^{b-1} (1-x)^{c-b-1} (1-zx)^{-a} dx$
- Gauss summation theorem: ${}_2F_1(a, b; c; 1) = \frac{\Gamma(c) \Gamma(c-a-b)}{\Gamma(c-a) \Gamma(c-b)}$, for $\text{Re}(c) > \text{Re}(a+b)$
- Barnes' contour integral: ${}_2F_1(a, b; c; z) = \frac{\Gamma(c)}{\Gamma(a) \Gamma(b)} \times \frac{1}{2\pi i} \int_{-i\infty}^{i\infty} \frac{\Gamma(a+s) \Gamma(b+s) \Gamma(-s)}{\Gamma(c+s)} (-z)^s ds$
- Generalised HGF: ${}_pF_q(a_1, \dots, a_p; b_1, \dots, b_q; z) = \sum_{n=0}^{\infty} \frac{(a_1)_n \dots (a_p)_n}{(b_1)_n \dots (b_q)_n} \frac{z^n}{n!}$, for $|z| < 1$
- Regularised HGF: $\widehat{{}_pF_q}(a_1, \dots, a_p; b_1, \dots, b_q; z) = \frac{{}_pF_q(a_1, \dots, a_p; b_1, \dots, b_q; z)}{\Gamma(b_1) \dots \Gamma(b_q)}$
- Kummer's confluent HGF, first kind: $M(a, b; z) = {}_1F_1(a, b; z) = \lim_{c \rightarrow \infty} {}_2F_1(a, c; b; \frac{z}{c})$
- Kummer's confluent HGF, second kind: $U(a, b; z) = z^{-a} {}_2F_0(a, 1+a-b; ; \frac{-1}{z})$

1.7.7. Elliptic Integrals, $K(k)$, $E(k)$ and $\Pi(n, k)$

- Complete elliptic integral, first kind: $K(k) = \int_0^{\pi/2} \frac{d\theta}{\sqrt{1-k^2 \sin^2 \theta}} = \int_0^1 \frac{dt}{\sqrt{(1-t^2)(1-k^2 t^2)}}$
- Complete elliptic integral, second kind: $E(k) = \int_0^{\pi/2} \sqrt{1-k^2 \sin^2 \theta} d\theta = \int_0^1 \frac{\sqrt{1-k^2 t^2}}{\sqrt{1-t^2}} dt$
- Complete elliptic integral, third kind: $\Pi(n, k) = \int_0^{\pi/2} \frac{d\theta}{(1-n \sin^2 \theta) \sqrt{1-k^2 \sin^2 \theta}}$
- Incomplete elliptic integral, first kind: $F(\varphi, k) = \int_0^{\varphi} \frac{d\theta}{\sqrt{1-k^2 \sin^2 \theta}}, F(x; k) = \int_0^x \frac{dt}{\sqrt{(1-t^2)(1-k^2 t^2)}}$
- Incomplete elliptic integral, second kind: $E(\varphi, k) = \int_0^{\varphi} \sqrt{1-k^2 \sin^2 \theta} d\theta, E(x; k) = \int_0^x \frac{\sqrt{1-k^2 t^2}}{\sqrt{1-t^2}} dt$
- Incomplete elliptic integral, third kind: $\Pi(n; \varphi \setminus \alpha) = \int_0^{\varphi} \frac{d\theta}{(1-n \sin^2 \theta) \sqrt{1-(\sin \alpha \sin \theta)^2}}$

Substitutions used above are $t = \sin \theta$ and $x = \sin \varphi$.

Legendre's relation: $K(k) E(\sqrt{1-k^2}) + E(k) K(\sqrt{1-k^2}) - K(k) K(\sqrt{1-k^2}) = \frac{\pi}{2}$

Arithmetic-Geometric mean identity: $K(k) = \frac{\pi}{2 \operatorname{agm}(1, \sqrt{1-k^2})}$

Inverse elliptic integrals (Jacobi functions): with $u = F(\varphi, k) = \int_0^{\varphi} \frac{d\theta}{\sqrt{1-k^2 \sin^2 \theta}}$ then

- Jacobi amplitude function: $\operatorname{am}(u, k) = \varphi$
- Elliptic sine: $\operatorname{sn}(u, k) = \sin(\operatorname{am}(u, k)) = \sin \varphi$
- Elliptic cosine: $\operatorname{cn}(u, k) = \cos(\operatorname{am}(u, k)) = \cos \varphi$
- Delta amplitude: $\operatorname{dn}(u, k) = \frac{d\varphi}{du} \operatorname{am}(u, k) = \frac{d\varphi}{du} = \sqrt{1-k^2 \sin^2 \varphi}$

1.7.8. Zeta Function, $\zeta(z)$

- Zeta function as a series: $\zeta(z) = \sum_{r=1}^{\infty} r^{-z} = \frac{1}{\Gamma(z)} \int_0^{\infty} \frac{x^{z-1}}{e^x - 1} dx, \quad \text{for } \operatorname{Re}(z) > 1$
- Euler's product formula: $\zeta(z) = \prod_{p \in P} \frac{1}{1-p^{-z}},$ a product over all primes p
- Riemann's functional equation: $\zeta(z) = 2^z \pi^{z-1} \sin \frac{\pi z}{2} \Gamma(1-z) \zeta(1-z)$
- Riemann hypothesis: does $\zeta(z) = 0 \Leftrightarrow \operatorname{Re}(z) = \frac{1}{2}$ or $z \in \{-2, -4, -6, \dots\}$?

(critical line)
(trivial zeroes)

1.7.9. Bessel Functions, $J_n(x)$ and $Y_n(x)$, and Hankel Functions, $H_\alpha^{(1)}(x)$ and $H_\alpha^{(2)}(x)$

Bessel functions, $J_n(x)$ and $Y_n(x)$:

1st kind: $J_n(x) = \frac{1}{\pi} \int_0^\pi \cos(n\tau - x \sin \tau) d\tau$

2nd kind: $Y_n(x) = \frac{1}{\pi} \int_0^\pi \sin(x \sin \tau - n\tau) d\tau$

Hankel function, $H_\alpha^{(1)}(x)$ and $H_\alpha^{(2)}(x)$:

1st kind: $H_\alpha^{(1)}(x) = J_\alpha(x) + iY_\alpha(x) = \frac{J_{-\alpha}(x) - e^{-\alpha\pi i} J_\alpha(x)}{i \sin \alpha\pi}$

2nd kind: $H_\alpha^{(2)}(x) = J_\alpha(x) - iY_\alpha(x) = \frac{J_{-\alpha}(x) - e^{\alpha\pi i} J_\alpha(x)}{-i \sin \alpha\pi}$

Modified Bessel Functions, $I_\alpha(x)$ and $K_\alpha(x)$:

1st kind: $I_\alpha(x) = i^{-\alpha} J_\alpha(ix) = \sum_{m=0}^\infty \frac{(x/2)^{2m+\alpha}}{m! \Gamma(m+\alpha+1)}$

2nd kind: $K_\alpha(x) = \frac{\pi}{2} \frac{I_{-\alpha}(x) - I_\alpha(x)}{\sin \alpha\pi}$

Spherical Bessel functions, $j_n(x)$ and $y_n(x)$:

1st kind: $j_n(x) = \sqrt{\frac{\pi}{2x}} J_{n+\frac{1}{2}}(x)$

2nd kind: $y_n(x) = \sqrt{\frac{\pi}{2x}} Y_{n+\frac{1}{2}}(x) = (-1)^{n+1} j_{-n-1}(x)$

1.7.10. Associated Legendre Polynomials, $P_l^m(x)$

Associated Legendre Polynomials: $P_l^m(x) = \frac{(-1)^m}{2^l l!} (1-x^2)^{m/2} \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l, \quad -l \leq m \leq l$

$P_l^m(x)$	$m = 0$	$m = 1$	$m = 2$	$m = 3$
$l = 0$	1			
$l = 1$	x	$-(1-x^2)^{1/2}$		
$l = 2$	$\frac{1}{2}(3x^2-1)$	$-3x(1-x^2)^{1/2}$	$3(1-x^2)$	
$l = 3$	$\frac{1}{2}(5x^3-3x)$	$\frac{3}{2}(1-5x^2)(1-x^2)^{1/2}$	$15x(1-x^2)$	$-15(1-x^2)^{3/2}$

For negative index m : $P_l^{-m}(x) = (-1)^m \frac{(l-m)!}{(l+m)!} P_l^m(x)$

1.7.11. Hermite Polynomials, $H_n(x)$

Hermite polynomials as a derivative: $H_n(x) = (-1)^n e^{x^2} \frac{d^n}{dx^n} e^{-x^2}$

Recurrence relation: $H_{n+1}(x) = 2x H_n(x) - H_n'(x)$

The first few Hermite polynomials are

$$H_0(x) = 1; H_1(x) = 2x; H_2(x) = 4x^2 - 2; H_3(x) = 8x^3 - 12x;$$

$$H_4(x) = 16x^4 - 48x^2 + 12; H_5(x) = 32x^5 - 160x^3 + 120x$$

1.7.12. Generalised Laguerre Polynomials, $L_n^{(\alpha)}(x)$

Recurrence relation: $(k + 1) L_{k+1}^{(\alpha)}(x) = (2k + 1 + \alpha - x) L_k^{(\alpha)}(x) - (k + \alpha) L_{k-1}^{(\alpha)}(x)$

The first few generalised Laguerre polynomials are:

$L_0^{(\alpha)}(x) = 1; L_1^{(\alpha)}(x) = -x + (\alpha + 1); L_2^{(\alpha)}(x) = \frac{1}{2}x^2 + (\alpha + 2)x + \frac{1}{2}(\alpha + 1)(\alpha + 2)$


1.7.13. Airy Functions, $Ai(x)$ and $Bi(x)$

Airy equation: $\frac{d^2 y}{dx^2} - xy = 0$

Linearly independent solutions are:

$Ai(x) = \frac{1}{\pi} \int_0^\infty \cos\left(xt + \frac{t^3}{3}\right) dt$

$Bi(x) = \frac{1}{\pi} \int_0^\infty \exp\left(xt - \frac{t^3}{3}\right) + \sin\left(xt + \frac{t^3}{3}\right) dt$




1.7.14. Fresnel Integrals, $S(x)$ and $C(x)$

Fresnel sine and cosine:

$S(x) = \int_0^x \sin(t^2) dt = \sum_{n=0}^\infty (-1)^n \frac{x^{4n+3}}{(2n+1)!(4n+3)},$

$C(x) = \int_0^x \cos(t^2) dt = \sum_{n=0}^\infty (-1)^n \frac{x^{4n+1}}{(2n)!(4n+1)}.$



Clothoid curve: $\{x(t) = C(t), y(t) = S(t)\}.$

which has a constant rate of change of curvature: $\frac{d\kappa}{ds} = 2$ and $\frac{ds}{dt} = 1.$

Limiting value: $\lim_{x \rightarrow \infty} S(x) = \lim_{x \rightarrow \infty} C(x) = \sqrt{\frac{\pi}{8}}.$

1.7.15. Lambert W Function, $W_k(z)$


The Lambert W function is the inverse function of $f(z) = ze^z$ i.e. $W_k(z) e^{W_k(z)} = z$ ($k \in \mathbb{Z}$).

For real x , the two branches $y = W_0(x)$ and $y = W_{-1}(x)$ are the solutions to $y e^y = x$ for $x \geq 0$ and $-1/e \leq x < 0$ respectively.

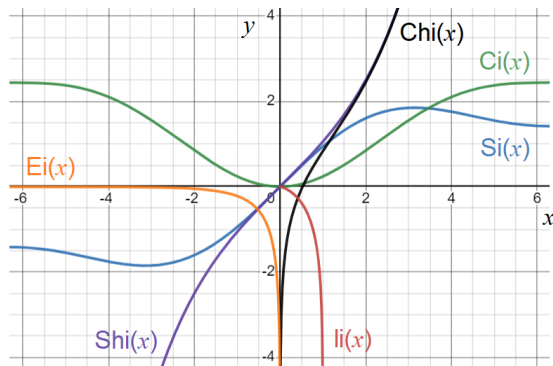
Taylor series: $W_0(x) = \sum_{n=1}^{\infty} \frac{(-n)^{n-1}}{n!} x^n, \quad |x| < 1/e$

Lambert differential equation: $x(1 + y) \frac{dy}{dx} = y$

All branches are solutions to this DE ($y = W_n(x)$).



1.7.16. Trigonometric Integral Functions, Si(x), Ci(x), Shi(x), Chi(x), li(x), Ei(x)



$$\lim_{x \rightarrow \pm\infty} Si(x) = \int_0^{\pm\infty} \frac{\sin t}{t} dt = \pm \frac{\pi}{2}$$

$$\lim_{x \rightarrow \pm\infty} Ci(x) = \int_0^{\pm\infty} \frac{1 - \cos t}{t} dt \text{ diverges to } \infty$$

$$\lim_{x \rightarrow 1} li(x) = \int_0^1 \frac{dt}{\ln t} \text{ diverges to } -\infty$$

- Sine integral: $Si(x) = \int_0^x \text{sinc } t \, dt = \int_0^x \frac{\sin t}{t} dt$
- Cosine integral: $Ci(x) = \int_0^x \frac{1 - \cos t}{t} dt = \gamma + \ln x + \int_x^{\infty} \frac{\cos t}{t} dt$
- Hyperbolic sine integral: $Shi(x) = \int_0^x \frac{\sinh t}{t} dt$ ($\gamma = 0.577216\dots$: Euler-Mascheroni constant)
- Hyperbolic cosine integral: $Chi(x) = \gamma + \ln x - \int_0^x \frac{1 - \cosh t}{t} dt, \quad x > 0$
- Logarithmic integral: $li(x) = \int_0^x \frac{dt}{\ln t}, \quad 0 \leq x < 1$
- Exponential integral: $Ei(x) = - \int_{-x}^{\infty} \frac{e^{-t}}{t} dt = \int_{-\infty}^x \frac{e^t}{t} dt, \quad x < 0$

1.7.17. Spherical Harmonics, $Y_l^m(\theta, \phi)$

Spherical harmonic: $Y_l^m(\theta, \phi) = \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}} P_l^m(\cos \theta) e^{im\phi} = \sqrt{\frac{2l+1}{4\pi}} C_{lm}(\theta, \phi)$ for $|m| \leq l$

The 'Condon-Shortley phase' is due to the term $(-1)^m$, which is included in the definition of P_l^m .

Expressions for C_{lm} in angular spherical coordinates (θ, ϕ) and Cartesian coordinates (x, y, z) are

$$C_{00} = 1; \quad C_{10} = \cos \theta = \frac{z}{r}; \quad C_{1,\pm 1} = \mp \sqrt{\frac{1}{2}} \sin \theta e^{\pm i\phi} = \mp \sqrt{\frac{1}{2}} \frac{x \pm iy}{r}$$

$$C_{20} = \frac{1}{2} (3 \cos^2 \theta - 1) = \frac{3z^2 - r^2}{2r^2} \quad (r = \sqrt{x^2 + y^2 + z^2})$$

$$C_{2,\pm 1} = \mp \sqrt{\frac{3}{2}} \cos \theta \sin \theta e^{\pm i\phi} = \mp \sqrt{\frac{3}{2}} \frac{zx \pm izy}{r^2}; \quad C_{2,\pm 2} = \sqrt{\frac{3}{8}} \sin^2 \theta e^{\pm 2i\phi} = \sqrt{\frac{3}{8}} \frac{x^2 - y^2 \pm 2ixy}{r^2}$$

Real Spherical Harmonics: real-valued alternative definition

$$Y_{lm} = \begin{cases} \frac{i}{\sqrt{2}} (Y_\ell^m - (-1)^m Y_\ell^{-m}) & \text{if } m < 0 \\ Y_\ell^0 & \text{if } m = 0 \\ \frac{1}{\sqrt{2}} (Y_\ell^{-m} + (-1)^m Y_\ell^m) & \text{if } m > 0. \end{cases} = \begin{cases} \sqrt{2} (-1)^m \Im[Y_\ell^{|m|}] & \text{if } m < 0 \\ Y_\ell^0 & \text{if } m = 0 \\ \sqrt{2} (-1)^m \Re[Y_\ell^{|m|}] & \text{if } m > 0. \end{cases}$$

Vector Spherical Harmonics: a complex vector-valued function.

$$\mathbf{Y}_{lm}(\theta, \phi) = Y_l^m(\theta, \phi) \hat{\mathbf{r}}, \quad \mathbf{\Psi}_{lm}(r, \theta, \phi) = r \nabla Y_l^m, \quad \mathbf{\Phi}_{lm}(r, \theta, \phi) = r \hat{\mathbf{r}} \times \nabla Y_l^m$$

Spherical Harmonic Transform: for a function in angular spherical coordinates $f(\theta, \phi)$,

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \hat{a}_{lm} Y_l^m(\theta, \phi) \quad \text{where} \quad \hat{a}_{lm} = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) Y_l^m(\theta, \phi) \sin \theta d\theta d\phi$$

This is analogous to a Fourier series (Section 3.6), with \hat{a}_{lm} as complex coefficients of the basis functions Y_l^m . The Jacobian term is sometimes written as the solid angle $d\Omega = \sin \theta d\theta d\phi$.

$$\text{Normalisation: } \int_0^{2\pi} \int_0^\pi |Y_l^m|^2 d\Omega = 1. \quad \text{Orthogonality: } \int_0^{2\pi} \int_0^\pi (Y_l^m)(Y_{l'}^{m'})^* d\Omega = \delta_{ll'} \delta_{mm'}.$$

1.7.18. Miscellaneous Special Functions


Generalised Marcum Q -function:
$$Q_\nu(a, b) = \frac{1}{a^{\nu-1}} \int_b^\infty x^\nu e^{-\frac{1}{2}(x^2 + a^2)} I_{\nu-1}(ax) dx$$

M2. GEOMETRY

2.1. Properties of 2D and 3D Shapes

2.1.1. Properties of Simple 2D Shapes

Circles and Parts of Circles




Circle

Area: $A = \pi r^2$

Diameter: $d = 2r$


Circumference: $C = 2\pi r = \pi d$



Sector

Area: $A = \frac{1}{2}r^2\theta$


Arc length: $s = r\theta$



Segment

Area: $A = \frac{1}{2}r^2(\theta - \sin \theta)$

Chord length: $c = 2r \sin \frac{\theta}{2}$




Equilateral Triangle

Area: $A = \frac{\sqrt{3}}{4}a^2$

$r = \frac{\sqrt{3}}{6}a$

$R = \frac{\sqrt{3}}{3}a$



Regular n-gon


Area: $A = \frac{na^2}{4 \tan \frac{\pi}{n}}$

$r = \frac{a}{2 \tan \frac{\pi}{n}}, R = \frac{a}{2 \sin \frac{\pi}{n}}$

External angle: $\theta = \frac{2\pi}{n}$

Special Quadrilaterals


(For properties of scalene/right triangles, see Section 2.2.7.)



Kite (Deltoid)

Area: $A = \frac{1}{2}d_1d_2$


diagonals perpendicular



Rhombus

Area: $A = \frac{1}{2}d_1d_2$


diagonals perp. bisectors



Parallelogram

Area: $A = bh = \frac{1}{2}d_1d_2 \sin \theta$

diagonals bisect



Trapezium (Trapezoid)

$A = \frac{1}{2}h(a + b)$

2.1.2. Symmetry

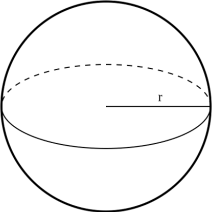
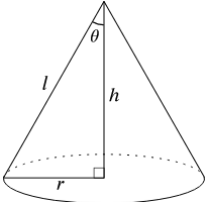
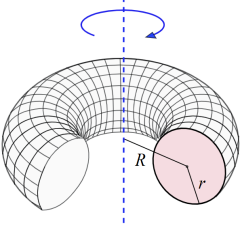
Rotational symmetry of order n : identical after turning through $360^\circ / n$

Reflective/mirror symmetry of order n : identical after reflecting in n different axes

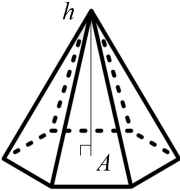
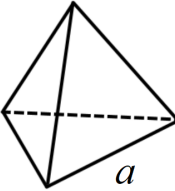
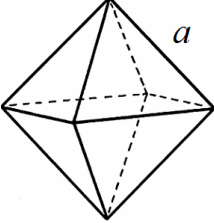
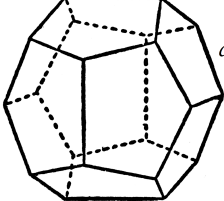
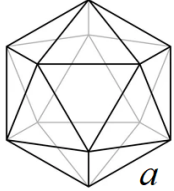
A regular polygon is both rotational and mirror symmetry order n

2.1.3. Volumes and Surface Areas of 3D Solid Figures


Curved bodies:

			
	Sphere	Cone	Torus
Volume	$\frac{4}{3}\pi r^3$	$\frac{1}{3}\pi r^2 h$	$2\pi^2 Rr^2$
Surface Area	$4\pi r^2$	$\pi r l + \pi r^2$	$4\pi^2 Rr$

Pyramidal and Platonic solids (f: faces, v: vertices, e: edges):

					
	Pyramid	Tetrahedron (4 f, 4 v, 6 e)	Octahedron (8 f, 6 v, 12 e)	Dodecahedron (12 f, 20 v, 30 e)	Icosahedron (20 f, 12 v, 30 e)
Volume	$\frac{1}{3}Ah$	$\frac{\sqrt{2}}{12} a^3$	$\frac{\sqrt{2}}{3} a^3$	$\frac{15 + 7\sqrt{5}}{4} a^3$	$\frac{5(3 + \sqrt{5})}{12} a^3$
Surface Area	$\frac{1}{2}pL + A$ (p: base perimeter L: slant length)	$\sqrt{3} a^2$	$2\sqrt{3} a^2$	$3\sqrt{25 + 10\sqrt{5}} a^2$	$5\sqrt{3} a^2$

2.1.4. Sections of Spheres (Sections of Revolution)




Spherical Cap (blue): radius r , flat radius a , height h , half-angle θ :

Volume:

$$V = \frac{\pi h^2}{3} (3r - h) = \frac{\pi h}{6} (3a^2 + h^2) = \frac{\pi r^3}{h} (2 + \cos \theta)(1 - \cos \theta)^2$$

Curved surface area: (circular plane face area: πa^2)

$$A = 2\pi r h = \pi(a^2 + h^2) = 2\pi r^2(1 - \cos \theta)$$



Spherical Sector (blue): radius r , flat radius a , height h , half-angle θ :


Volume:

$$V = \frac{2\pi r^2 h}{3} = \frac{\pi}{6h} (a^2 + h^2)^2 = \frac{2\pi r^3}{3} (1 - \cos \theta)$$

Curved surface area: (cone area: πar)

$$A = 2\pi r h = \Omega r^2 = 2\pi r^2(1 - \cos \theta)$$

(Ω : solid angle, in steradians (by definition))



Spherical Wedge: radius r , dihedral angle θ :


Volume:

$$V = \frac{2\pi r^3 \theta}{3}$$

Curved surface area (lune):

$$A = 2r^2 \theta$$


A plane intersects a sphere in a circle. The maximum area of this circle occurs when the plane cuts the sphere in two equal parts (hemisphere caps separated by a 'great circle').

2.1.5. Some Geometric Results in 3D**3D Quadrilateral Coffin Problem**

AB, BC, CD, AD tangent to a sphere at W, X, Y, Z

W, X, Y, Z are coplanar


2.1.6. Circle-Circle and Sphere-Sphere Intersections



Circle-Circle Intersection

Area of overlapping lens region:

$$A = r^2 \cos^{-1} \frac{d^2 + r^2 - R^2}{2dr} + R^2 \cos^{-1} \frac{d^2 + R^2 - r^2}{2dR} - \frac{1}{2} \sqrt{4d^2 R^2 - (d^2 - r^2 + R^2)^2}$$



Sphere-Sphere Intersection

Volume of overlapping lens region:

$$V = \frac{\pi}{12d} (R + r - d)^2 \times (d^2 + 2dr - 3r^2 + 2dR + 6rR - 3R^2)$$


In both cases,

Distance to intersection chord or intersection plane: $x = \frac{d^2 - r^2 + R^2}{2d}$

Length of intersection chord or diameter of intersection circle: $a = \frac{1}{d} \sqrt{4d^2 R^2 - (d^2 - r^2 + R^2)^2}$

2.1.7. Ellipses and Ellipsoids

Ellipse (2D): an elongated circle.




Ellipse

Area: $A = \pi ab$

Eccentricity: $e^2 = 1 - \frac{b^2}{a^2}$

Perimeter: $P = 4a E(e)$




Elliptical Sector

$\cos^2 \theta_1 + \sin^2 \alpha = 1$ and $\cos^2 \theta_2 + \sin^2 \beta = 1$

Area: $\frac{ab}{2}(\beta - \alpha - \sin(\beta - \alpha))$

Elliptical arc length: $a(E(\beta, e) - E(\alpha, e))$

Ellipsoid (3D): an elongated sphere along two axes (radii a, b, c)



Volume: $\frac{4}{3} \pi abc$ ($a \geq b \geq c$)

Surface area: $2\pi \left[c^2 + \frac{ab}{\sin \varphi} (K(\varphi, k) \cos^2 \varphi + E(\varphi, k) \sin^2 \varphi) \right]$, $\cos \varphi = \frac{c}{a}$ and $k^2 = \frac{a^2(b^2 - c^2)}{b^2(a^2 - c^2)}$

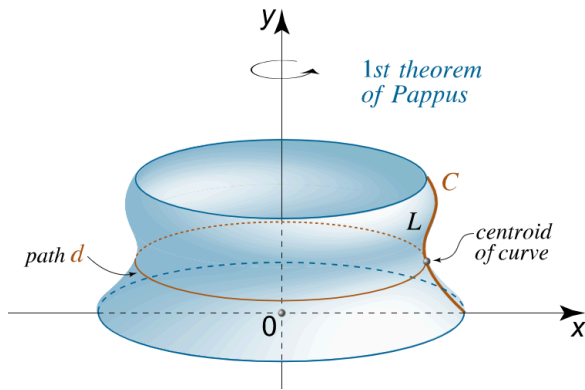
Spheroid: a sphere compressed (oblate) or elongated (prolate) along **one** axis (radii a, a, c)

Surface area, oblate ($c < a$): $2\pi a^2 \left(1 + \frac{1 - e^2}{e} \tanh^{-1} e \right)$, $e^2 = 1 - \frac{c^2}{a^2}$

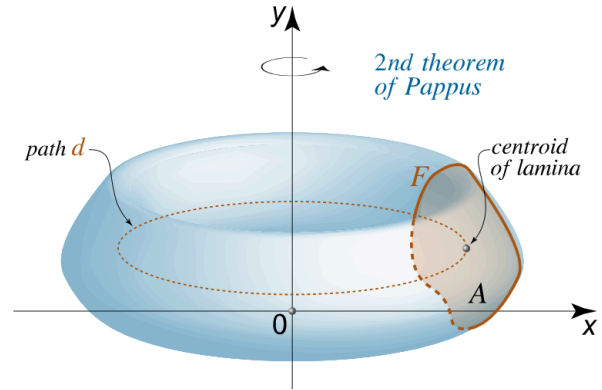
Surface area, prolate ($c > a$): $2\pi a^2 \left(1 + \frac{c}{ae} \sin^{-1} e \right)$, $e^2 = 1 - \frac{a^2}{c^2}$

A plane whose normal is parallel to the axis of stretching intersects the spheroid in a circle.

2.1.8. Pappus' Theorems



Surface Area = $L \times d$




Volume = $A \times d$

For partial revolutions, use the arc length instead of the circumference as d .

Generalisation to Pappus' theorems:

The path traced out by the centroid does not need to be circular: it can be any simple curved path (e.g. linear, parabolic, helical). This will result in a 'swept' solid or surface. The appropriate length d is then the arc length along this path.

Additionally, the curve/lamina being swept may rotate in its plane (torsion: remaining perpendicular to the path) along the path, as long as the angle of twist is continuous. E.g:



Volume = area of lamina \times path length

Surface area = perimeter of lamina \times path length


The resulting solid or surface must not be self-intersecting to produce valid results.

2.2. Angle, Triangle and Circle Theorems


2.2.1. Angle Theorems

Types of angles: acute ($0^\circ < \theta < 90^\circ$), right angle ($\theta = 90^\circ$), obtuse ($90^\circ < \theta < 180^\circ$), straight line ($\theta = 180^\circ$), reflex ($180^\circ < \theta < 360^\circ$)


For angles at a given point,



Angles on a line
 $\alpha + \beta = 180^\circ$




Angles around a point
 $\alpha + \beta + \gamma = 360^\circ$




Opposite angles
are equal


For parallel lines intersected by a single transversal line,



Corresponding Angles ('F')
are equal



Co-Interior Angles ('C')
 $\alpha + \beta = 180^\circ$



Alternate Angles ('Z')
are equal

2.2.2. Measures of Angles

Common measures of angles are: ($1 \text{ rad} = \frac{180}{\pi} \approx 57.29^\circ$)


- Degrees: a full turn is 360° .
- Radians: a full turn is 2π rad. Assumed in all calculations (natural units).
- Gradians: a full turn is 400° (archaic).

Units for small angles include the DMS (degrees-minutes-seconds, $D^\circ M' S''$) system:

- 1 degree = 60 arcminutes ($1^\circ = 60'$)
- 1 arcminute = 60 arcseconds ($1' = 60''$)


Three-figure bearings, used in navigation, are given in degrees clockwise from North, using three digits by convention (e.g. "050" for 50° clockwise from North).

2.2.3. Triangle and Quadrilateral Theorems




Midpoint Theorem

$$|DE| = \frac{1}{2} |AB|$$




Angle Bisector Theorem

$$\frac{b}{a} = \frac{m}{n}$$




Apollonius' Theorem

$$a^2 + b^2 = 2(d^2 + m^2)$$



Stewart's Theorem


$$b^2 n + a^2 m = c(d^2 + mn)$$



Ceva's Theorem

$$\frac{|AF|}{|FB|} \times \frac{|BD|}{|DC|} \times \frac{|CE|}{|EA|} = 1$$


$$\frac{\sin \angle ABE}{\sin \angle CBE} \times \frac{\sin \angle BCF}{\sin \angle ACF} \times \frac{\sin \angle CAD}{\sin \angle BAD} = 1$$



Menelaus' Theorem

$$\frac{|AF|}{|FB|} \times \frac{|BD|}{|DC|} \times \frac{|CE|}{|EA|} = 1$$

in both cases




Ptolemy's Theorem

$$|AB||CD| + |BC||DA| = |AC||BD|$$

If ABCD is **not** cyclic, then:


$$|AB||CD| + |BC||DA| \geq |AC||BD|$$



Routh's Theorem

$$\frac{\Delta PQR}{\Delta ABC} = \frac{(xyz-1)^2}{(xy+y+1)(yz+z+1)(xz+x+1)}$$

where $x = \frac{|BD|}{|DC|}$, $y = \frac{|CE|}{|EA|}$, $z = \frac{|AF|}{|FB|}$




Pythagoras' Theorem


$$a^2 + b^2 = c^2$$

$$\frac{1}{a^2} + \frac{1}{b^2} = \frac{1}{d^2}$$


2.2.4. Circle Theorems




Thales' Theorem
 (Angle in a semicircle)
 AC is a diameter
 $\angle ABC = 90^\circ$




Inscribed Angle Theorems
 Angle at the centre
 $\angle AOC = 2 \angle ABC$




Angles subtended by an arc
 $\angle DAC = \angle DBC$




Alternate Segment Theorem
 ST tangent at A
 $\angle TAC = \angle ABC$




Cyclic Quadrilaterals
 $\angle ABC + \angle ADC = 180^\circ$
 opposite angles add to 180°




Tangent Theorem
 $|AP| = |BP|$, $OA \perp AP$
 OAPB is a cyclic kite




Chord Bisector Theorem
 PQ is a chord
 $OA \perp PQ$, $|AP| = |AQ|$




Butterfly Theorem
 M midpoint of chord PQ
 $|MX| = |MY|$




Intersecting Chords
 $|OA||OC| = |OB||OD|$



Tangent-Secant
 $|OA|^2 = |OB||OC|$



Intersecting Secants
 $|OA||OB| = |OC||OD|$




Carnot-Euler Theorems
 $|OX| + |OY| + |OZ| = R + r$
 $|OI|^2 + r^2 = (R - r)^2$
 distances **negative**
 if entirely outside triangle

Three Cases of the Power of a Point Theorem

2.2.5. Some Special Geometric Constructions


These setups may require unique methods of solving, and are extremely difficult without knowing the technique.



Ford Circles

Three circles tangent to each other, as well as one common tangent line

$$c^{-1/2} = a^{-1/2} + b^{-1/2}$$




Langley's Adventitious Angles

Given a, b, c, d , find e .

In general, it is extremely difficult without trigonometry.

Techniques include: trigonometric Ceva's theorem, identifying congruent/equilateral triangles, three circumcentres method



Triangle Construction Coffin Problem

The line segments between a point in an equilateral triangle and its vertices are used to form a new triangle.

Angles: $\alpha' = \alpha - 60^\circ, \beta' = \beta - 60^\circ, \gamma' = \gamma - 60^\circ$


Can be solved by rotating the diagram 60° about a vertex.

2.2.6. Similarity and Congruence of Triangles

Two triangles are similar ($\triangle ABC \sim \triangle PQR$) if one is an enlargement of the other (AAA).

Two triangles are congruent ($\triangle ABC \cong \triangle PQR$) if they are identical (SSS / SAS / ASA / AAS / RHS).


(S / A / R / H: a side / angle / right-angle / hypotenuse known to be equal in both triangles.)



Fundamental theorem of similarity: $\triangle ABC \sim \triangle AB'C'$

$$\frac{|AB|}{|BB'|} = \frac{|AC|}{|CC'|} \Leftrightarrow \frac{|AB|}{|AB'|} = \frac{|AC|}{|AC'|} \Leftrightarrow BC \parallel B'C'$$

2.2.7. Trigonometry of Right-Angled Triangles



The sides are said to be a (adjacent), o (opposite), h (hypotenuse) relative to acute angle θ .


Definitions: $\sin \theta = \frac{o}{h}$ $\cos \theta = \frac{a}{h}$ $\tan \theta = \frac{o}{a}$ (SohCahToa)

Pythagoras' Theorem: $a^2 + o^2 = h^2$ (more often written $a^2 + b^2 = c^2$)

2.2.8. Trigonometry of Triangles

The results here are valid for any cyclic permutation of $\{a, b, c\}$ and $\{A, B, C\}$.

Area of a triangle, $\Delta ABC = \frac{1}{2}bh = \frac{1}{2}ab \sin C = \frac{c^2 \sin A \sin B}{2 \sin(A+B)}$ (h : height perpendicular to b)



Law of Sines (Sine Rule):


$$\frac{a}{\sin A} = \frac{b}{\sin B} = \frac{c}{\sin C}$$

Law of Cosines (Cosine Rule):

$$c^2 = a^2 + b^2 - 2ab \cos C$$

Law of Tangents (Tangent Rule):

$$\frac{a-b}{a+b} = \frac{\tan \frac{1}{2}(A-B)}{\tan \frac{1}{2}(A+B)}$$



Law of Cotangents (Cotangent Rule): $\frac{\cot \frac{1}{2}A}{s-a} = \frac{\cot \frac{1}{2}B}{s-b} = \frac{\cot \frac{1}{2}C}{s-c} = \frac{1}{r}$

Inscribed Circle Radius:

$$r = \frac{\Delta ABC}{s} = \frac{ab \sin C}{a+b+c}$$

Circumscribed Circle Radius:

$$R = \frac{abc}{4 \Delta ABC} = \frac{a}{2 \sin A}$$


Heron's Theorem: $\Delta ABC = \sqrt{s(s-a)(s-b)(s-c)}$

Mollweide's formulas: $\frac{a+b}{c} = \frac{\cos \frac{1}{2}(A-B)}{\sin \frac{1}{2}C}$ and $\frac{a-b}{c} = \frac{\sin \frac{1}{2}(A-B)}{\cos \frac{1}{2}C}$

(s : semiperimeter, $s = \frac{a+b+c}{2}$, ΔABC : area of triangle ABC , r : inradius, R : circumradius)

2.2.9. Rational Trigonometry

Rational trigonometry is an alternative formulation of trigonometry in Euclidean geometry that uses 'spreads' and 'quadrances' instead of angles and lengths, which avoids the use of transcendental functions and irrational numbers.




A line with equation $ax + by + c = 0$ is represented as $\langle a : b : c \rangle$.

$$\text{Spread: } s(l_1, l_2) = \frac{(a_1 a_2 - b_1 b_2)^2}{(a_1^2 + b_1^2)(a_2^2 + b_2^2)} \quad (\text{equal to } \sin^2 \theta)$$

Quadrance: Q is equal to distance squared. (equal to a^2)

A triangle is considered a set of three lines. Identities are:

- Pythagorean Theorem: $Q_1 + Q_2 = Q_3 \Leftrightarrow Q_1 \perp Q_2$.
- Triple Spread Formula: $(s_1 + s_2 + s_3)^2 = 2(s_1^2 + s_2^2 + s_3^2) + 4s_1 s_2 s_3$ (angle sum)
- Spread Law: $\frac{s_1}{Q_1} = \frac{s_2}{Q_2} = \frac{s_3}{Q_3}$ (sine rule)
- Cross Law: $(Q_1 + Q_2 - Q_3)^2 = 4Q_1 Q_2 (1 - s_3)$ (cosine rule)


2.2.10. Projective Geometry

The cross ratio (anharmonic ratio) of any four collinear points is invariant under perspective projection:

$$\text{Cross ratio} = \frac{|AC||BD|}{|BC||AD|} = \frac{|ac||bd|}{|bc||ad|}$$

2.2.11. Mass Point Geometry (Barycentric Coordinates)

2.2.12. Spherical Geometry and Trigonometry (Non-Euclidean Geometry)



For a triangle made from three **great circle arcs of unit radius**:

a, b, c represent both (arc) lengths **and** angles subtended from the centre of the sphere O (in radians) i.e. $\angle AOB = c, \angle BOC = a, \angle COA = b$.

Spherical Cosine Rule: $\cos a = \cos b \cos c + \sin b \sin c \cos A$

Spherical Sine Rule: $\frac{\sin A}{\sin a} = \frac{\sin B}{\sin b} = \frac{\sin C}{\sin c}$

Inverse Cosine Rule: $\cos A = \sin B \sin C \cos a - \cos B \cos C$

Area of triangle $\triangle ABC$ (on sphere) = $A + B + C - \pi$ (Girard's Theorem)

Solid Angles: trihedral angles measured from O (units: steradians [sr]; full sphere = 4π sr.)

- Spherical triangle, ABC from O : $\Omega_o = A + B + C - \pi$
- Cone, vertex O , apex angle 2θ : $\Omega_o = 4\pi \sin^2 \frac{\theta}{2}$
- Irregular tetrahedron $OABC$: $\tan \frac{\Omega_o}{2} = \frac{\overline{OA} \cdot \overline{OB} \times \overline{OC}}{|\overline{OA}||\overline{OB}||\overline{OC}| + (\overline{OA} \cdot \overline{OB})|\overline{OC}| + (\overline{OB} \cdot \overline{OC})|\overline{OA}| + (\overline{OC} \cdot \overline{OA})|\overline{OB}|}$
 $\cos \Omega_o = \frac{1}{3}(\cos \angle AOB + \cos \angle BOC + \cos \angle COA)$

The solid angle of a polyhedron is the sum of solid angles of the non-overlapping tetrahedra sharing the vertex (e.g. compute from Delauney tetrahedral mesh of 3D point cloud).

2.2.13. Hyperbolic Geometry (Non-Euclidean, Lobachevsky Geometry)

Klein disk model (projective model): points represented as being inside a unit disk.

2.3. 2D Coordinate Geometry

2.3.1. Coordinates

A point P can lie in the xy Cartesian coordinate plane with origin O (frame xOy).

The coordinates of P can be written as $P(x, y)$. (x : abscissa of P , y : ordinate of P)

2.3.2. Equations of Lines

Gradient-Intercept:

$$y = mx + c$$


Point-Slope:

$$y - y_1 = m(x - x_1)$$

Two-point Interpolation:

$$\frac{y - y_1}{x - x_1} = \frac{y_2 - y_1}{x_2 - x_1}$$


2.3.3. Ratio Division of a Line Segment



Internal Division:

If $|AP| : |PB| = m : n$, then

$$P = \left(\frac{nx_A + mx_B}{m + n}, \frac{ny_A + my_B}{m + n} \right)$$



External Division:

If $|AQ| : |BQ| = m : n$, then

$$Q = \left(\frac{mx_B - nx_A}{m + n}, \frac{my_B - ny_A}{m + n} \right)$$

2.3.4. Tangential Angle and Angle Between Lines

For a tangent line of gradient m , the angle with the x -axis is ψ , where

$$m = \frac{dy}{dx} = \tan \psi$$

$$\Delta\psi = \tan^{-1} \left| \frac{m_1 - m_2}{1 + m_1 m_2} \right|$$

2.3.5. Parallel and Perpendicular Lines

If lines L_1 and L_2 have gradients m_1 and m_2 then

$$m_1 = m_2 \Leftrightarrow L_1 \parallel L_2 \quad \text{and} \quad m_1 m_2 = -1 \Leftrightarrow L_1 \perp L_2.$$

2.3.6. Area of an Irregular Plane Polygon From Coordinates

Shoelace formula: if an n -sided irregular polygon has vertices (ordered cyclically anticlockwise) at coordinates $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ then the area enclosed is

$$A = \frac{1}{2} \sum_{i=1}^n \begin{vmatrix} x_i & x_{i+1} \\ y_i & y_{i+1} \end{vmatrix} = \frac{1}{2} \left(\begin{vmatrix} x_1 & x_2 \\ y_1 & y_2 \end{vmatrix} + \begin{vmatrix} x_2 & x_3 \\ y_2 & y_3 \end{vmatrix} + \dots + \begin{vmatrix} x_n & x_1 \\ y_n & y_1 \end{vmatrix} \right)$$

In the case of a triangle, $n = 3$, this is equivalent to

$$A = \frac{1}{2} \begin{vmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{vmatrix} = \frac{1}{2} (x_1y_2 - x_2y_1 + x_2y_3 - x_3y_2 + x_3y_1 - x_1y_3)$$

2.3.7. Collinearity of Points and Concurrency of Lines

Collinearity: three points $(x_1, y_1), (x_2, y_2), (x_3, y_3)$ lie on the same line if

$$\begin{vmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{vmatrix} = 0 \quad \Leftrightarrow \quad x_1y_2 - x_2y_1 + x_2y_3 - x_3y_2 + x_3y_1 - x_1y_3 = 0.$$

Concurrency: three lines $a_1x + b_1y + c_1 = 0$, $a_2x + b_2y + c_2 = 0$ and $a_3x + b_3y + c_3 = 0$ intersect at a single point if

$$\begin{vmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ a_3 & b_3 & c_3 \end{vmatrix} = 0 \quad \Leftrightarrow \quad a_1(b_2c_3 - b_3c_2) - b_1(a_2c_3 - a_3c_2) + c_1(a_2b_3 - a_3b_2) = 0$$

In homogeneous coordinates (Section 4.2.3), there is a duality between collinearity of points (X, Y, Z) and lines (A, B, C) .

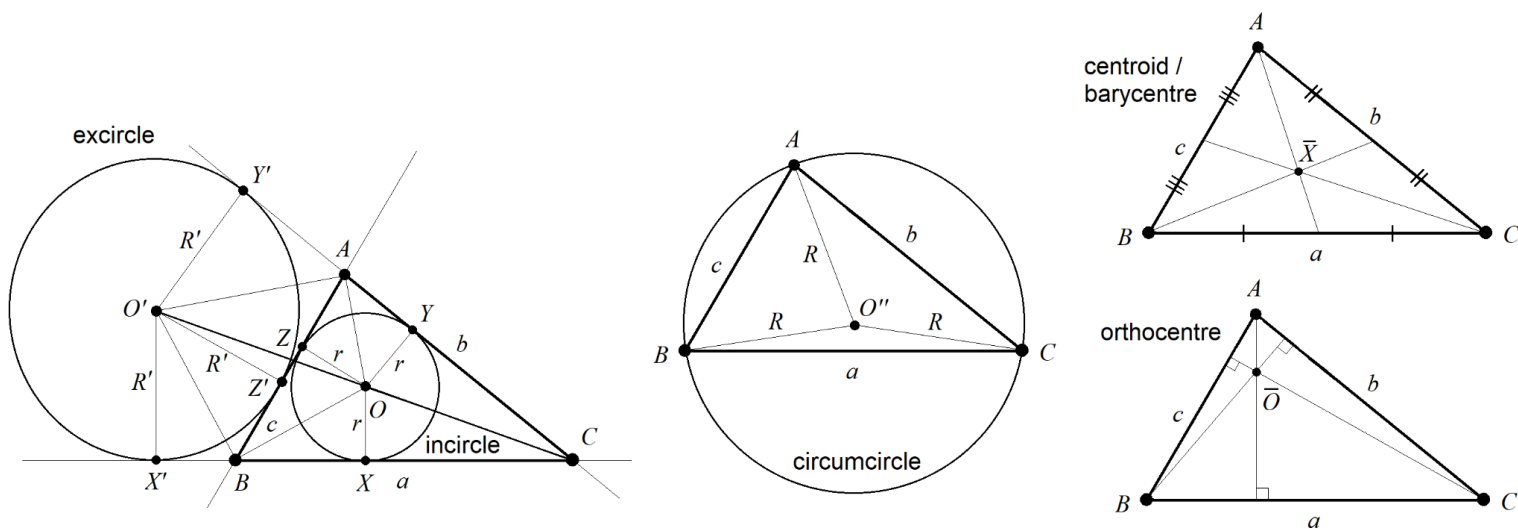
2.3.8. Equation of a Circle

For a circle with centre (x_0, y_0) and radius r , the equation is $(x - x_0)^2 + (y - y_0)^2 = r^2$

Standard form: $\frac{(x - x_0)^2}{r^2} + \frac{(y - y_0)^2}{r^2} = 1$

2.3.9. Centres of a Triangle: Incentre, Excentres, Circumcentre, Barycentre, Orthocentre

For a triangle with vertices $A = (x_A, y_A)$, $B = (x_B, y_B)$, $C = (x_C, y_C)$ opposite sides of length a, b, c :



Incentre: $O = \left(\frac{ax_A + bx_B + cx_C}{a + b + c}, \frac{ay_A + by_B + cy_C}{a + b + c} \right)$, Inradius: $r = \frac{2 \Delta_{ABC}}{a + b + c}$

Excentre (of C): $O' = \left(\frac{ax_A + bx_B - cx_C}{a + b - c}, \frac{ay_A + by_B - cy_C}{a + b - c} \right)$, Exradius: $R' = \frac{2 \Delta_{ABC}}{a + b - c}$ (on AB)

Circumcentre: $O'' = \left(\frac{(x_A^2 + y_A^2)(y_B - y_C) + (x_B^2 + y_B^2)(y_C - y_A) + (x_C^2 + y_C^2)(y_A - y_B)}{2(x_A(y_B - y_C) + x_B(y_C - y_A) + x_C(y_A - y_B))}, \frac{(x_A^2 + y_A^2)(x_C - x_B) + (x_B^2 + y_B^2)(x_A - x_C) + (x_C^2 + y_C^2)(x_B - x_A)}{2(x_A(y_B - y_C) + x_B(y_C - y_A) + x_C(y_A - y_B))} \right)$

Circumradius: $R = \frac{abc}{4 \Delta_{ABC}}$

Barycentre: $\bar{X} = \left(\frac{x_A + x_B + x_C}{3}, \frac{y_A + y_B + y_C}{3} \right)$ $\frac{2}{3}$ along any median line from vertex

Orthocentre: $\bar{O} = \left(\frac{x_A \tan A + x_B \tan B + x_C \tan C}{\tan A + \tan B + \tan C}, \frac{y_A \tan A + y_B \tan B + y_C \tan C}{\tan A + \tan B + \tan C} \right)$ (A, B, C : angles)

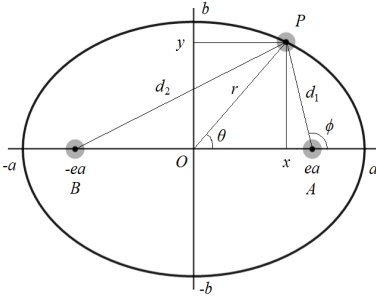
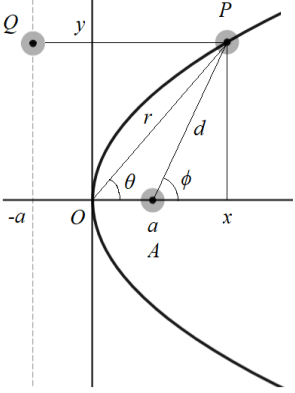
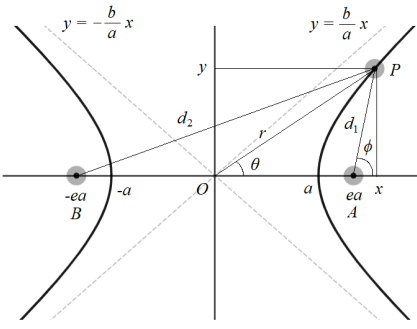
- Every triangle has a unique incentre, circumcentre, barycentre (centroid) and orthocentre.
- Every triangle has **three** distinct excentres.
- The size order of the excircles follows the size order of the lengths of the tangent edges, or equivalently, the size order of the opposite internal angles.
- The incentre, excentre and external point (O, O', C) are collinear.
- Area relations for incircle and excircle:

$$\Delta_{ABC} = \Delta_{AOB} + \Delta_{BOC} + \Delta_{COA} = \frac{1}{2}ar + \frac{1}{2}br + \frac{1}{2}cr = sr$$

$$\Delta_{ABC} = \Delta_{O'BC} + \Delta_{O'AC} - \Delta_{O'AB} = \frac{1}{2}aR' + \frac{1}{2}bR' - \frac{1}{2}cR' = (s - c)R'$$

- The Apollonius circle to the three excircles is internally tangent to all three and has radius $\frac{r^2 + s^2}{4r}$.

2.3.10. Properties of Conic Sections

	Ellipse	Parabola	Hyperbola
Diagram			
Cartesian equation	$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$	$y^2 = 4ax$	$\frac{x^2}{a^2} - \frac{y^2}{b^2} = 1$
Polar equations	$r = \frac{b}{\sqrt{1 - e^2 \cos^2 \theta}}$ (pole at O) $d_1 = \frac{b^2}{a(1 + e \cos \phi)}$ (pole at A)	$r = \frac{4a \cos \theta}{\sin^2 \theta}$ (pole at O) $d = \frac{2a}{1 - \cos \phi}$ (pole at A)	$r = \frac{b}{\sqrt{e^2 \cos^2 \theta - 1}}$ (pole at O) $d_1 = \frac{b^2}{a(1 - e \cos \phi)}$ (pole at A)
Parametric equations	$x = a \cos t, y = b \sin t$ $x = \pm a \operatorname{sech} t, y = b \tanh t$	$x = \frac{1}{4a} t^2, y = t$	$x = a \sec t, y = b \tan t$ $x = \pm a \cosh t, y = b \sinh t$
Definitions	O: centre A, B: focal points (foci) a: semi-major axis b: semi-minor axis e: eccentricity	O: vertex A: focal point (focus) a: focal length e: eccentricity $x = -a$: directrix	O: centre A, B: focal points (foci) a: semi-major axis e: eccentricity $y = \pm \frac{b}{a}x$: asymptotes
Eccentricity	$e^2 = 1 - \left(\frac{b}{a}\right)^2; 0 \leq e < 1$	$e = 1$	$e^2 = 1 + \left(\frac{b}{a}\right)^2; e > 1$
Plane-cone intersection	Plane gradient shallower than the cone.	Plane gradient equals that of the cone.	Plane gradient steeper than the cone.
Reflective property	Internal rays on AP are reflected into B.	Incident rays on PQ are reflected into A.	External rays parallel to AP are reflected towards B.
Distance property	$ AP + BP = d_1 + d_2 = 2a$	$ AP = PQ = d = x + a$	$ AP - BP = d_1 - d_2 = 2a$

2.3.11. Conic-Line Intersections

For a line $L: y = mx + c$, and conics translated to the origin given by

$$\text{ellipse } E: \frac{x^2}{a^2} + \frac{y^2}{b^2} = 1, \quad \text{hyperbola } H: \frac{x^2}{a^2} - \frac{y^2}{b^2} = 1, \quad \text{parabola } P: y^2 = 4ax$$

Then L makes either zero, one or two intersections with any conic, with x -coordinates of all intersections at the roots of

- Ellipse-line: $(a^2 m^2 + b^2)x^2 + 2a^2 mc x + a^2(c^2 - b^2) = 0$
- Hyperbola-line: $(a^2 m^2 - b^2)x^2 + 2a^2 mc x + a^2(c^2 + b^2) = 0$
- Parabola-line: $m^2 x^2 + (2mc - 4a)x + c^2 = 0$

Condition for tangency: discriminant of the quadratic is zero.

2.3.12. Canonical Matrix Equation of Conic Sections (Homogeneous Coordinates)

Equation of a Conic Section: any conic plane curve can be represented by the equation

$$A x^2 + B y^2 + 2C xy + 2D x + 2E y + F = 0$$

which can be represented in homogeneous coordinates (see Section 4.2.3) as

$$\mathbf{x}^T \mathbf{Q} \mathbf{x} = 0$$

where $\mathbf{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$ and $\mathbf{Q} = \begin{bmatrix} A & B & D \\ B & C & E \\ D & E & F \end{bmatrix}$ is a symmetric 3×3 matrix.

Conic Equation from Points: five points on a conic define the conic: (x_i, y_i) ($i = 1, 2, 3, 4, 5$).

$$\det \begin{bmatrix} x^2 & y^2 & xy & x & y & 1 \\ x_1^2 & y_1^2 & x_1 y_1 & x_1 & y_1 & 1 \\ x_2^2 & y_2^2 & x_2 y_2 & x_2 & y_2 & 1 \\ x_3^2 & y_3^2 & x_3 y_3 & x_3 & y_3 & 1 \\ x_4^2 & y_4^2 & x_4 y_4 & x_4 & y_4 & 1 \\ x_5^2 & y_5^2 & x_5 y_5 & x_5 & y_5 & 1 \end{bmatrix} = 0$$

The coefficients $A \dots F$ can be found by evaluating the six 5×5 sub-determinants across the top row.

Note that C, D, E have an extra factor of 2.

Affine Transformations on Conic Sections: transformations can be applied to a conic using

$$\mathbf{Q}' = (\mathbf{R}^{-1})^T \mathbf{Q} \mathbf{R}^{-1} \quad \text{where } \mathbf{R} \text{ is a } 3 \times 3 \text{ affine transformation matrix mapping a conic represented by } \mathbf{Q} \text{ into a conic represented by } \mathbf{Q}'$$

For the general form of the affine transformation matrix \mathbf{R} , see Section 4.2.3.

Geometric Parameters from a Conic Canonical Matrix: using singular value decomposition.

For any conic section $\mathbf{x}^T \mathbf{Q}' \mathbf{x} = 0$, define the matrix \mathbf{M} as the 2×2 matrix formed from the first two rows and columns of \mathbf{Q}' (i.e. $\mathbf{M} : \mathbf{Q} \rightarrow \mathbf{Q}'$ is a linear transformation matrix from a conic \mathbf{Q} aligned with the $\{\mathbf{i}, \mathbf{j}\}$ axes into the conic represented by \mathbf{Q}' translated to the origin).

If the singular value decomposition (SVD, see Section 4.3.7) is written as $\mathbf{M} = \mathbf{U} \Sigma \mathbf{V}^T$, then:

- The columns of \mathbf{U} represent the normalised **principal axes** of \mathbf{Q}' .
- The columns of \mathbf{V} (rows of \mathbf{V}^T) represent the vectors on the **original** quadric \mathbf{Q} which are mapped to the principal axes of \mathbf{Q}' .
- The singular values are the linear **scale factors** in the corresponding axes.

2.3.13. Radius of Curvature

A curve in 2D space has a curvature κ , with associated circular radius of curvature R , where

Cartesian, $y(x)$:
$$R = \frac{(1 + (y')^2)^{3/2}}{y''} \quad (x: \text{independent var, } y: \text{dependent var})$$

Parametric, $\{x(t), y(t)\}$:
$$R = \frac{((x')^2 + (y')^2)^{3/2}}{x'y'' - x''y'} \quad (t: \text{parameter})$$

Polar, $r(\theta)$:
$$R = \frac{(r^2 + (r')^2)^{3/2}}{r^2 + 2(r')^2 - rr''} \quad (r: \text{radial distance, } \theta: \text{polar angle})$$

Intrinsic, $s(\psi)$:
$$R = s' \quad (s: \text{arc length, } \psi: \text{tangential angle})$$

where f' is the derivative with respect to its argument.

2.3.14. Areas, Arc Lengths and Centres of Mass for Plane Curves by Integration


	Area A	Arc Length s	Centre of Mass (of plane region under curve)
Cartesian	$\int_a^b y \, dx$	$\int_a^b \sqrt{1 + \left(\frac{dy}{dx}\right)^2} \, dx$	$\bar{x} = \frac{1}{A} \int_a^b xy \, dx, \quad \bar{y} = \frac{1}{2A} \int_a^b y^2 \, dx$
Parametric	$\int_{t_1}^{t_2} y \frac{dx}{dt} \, dt$	$\int_{t_1}^{t_2} \sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2} \, dt$	$\bar{x} = \frac{1}{A} \int_{t_1}^{t_2} xy \frac{dx}{dt} \, dt, \quad \bar{y} = \frac{1}{2A} \int_{t_1}^{t_2} y^2 \frac{dx}{dt} \, dt$
Polar	$\frac{1}{2} \int_{\theta_1}^{\theta_2} r^2 \, d\theta$	$\int_{\theta_1}^{\theta_2} \sqrt{r^2 + \left(\frac{dr}{d\theta}\right)^2} \, d\theta$	(use Cartesian substitutions)

2.3.15. Tessellation (Tiling) and Partitioning of the Plane

Tessellation uses a fixed set of tile shapes to tile the plane completely, using only rotation and translation of the tiles to form unit cells which span the plane.


Delaunay triangulation (DT) is a mapping from a set of vertex points $\{P\}_i$ to a set of triangles $\{T\}_j$ connecting the points $\{P\}_i$ such that no point in $\{P\}_i$ is inside any triangle in the set. DT results in the maximisation of the smallest angle in every triangle. The convex hull of the vertices is the smallest convex polygon containing all vertices. If there are n vertices, of which h are on the convex hull, then there are $2n - 2 - h$ triangles and $3n - 3 - h$ edges on the Delaunay triangulation.

The **Voronoi diagram** of a set of vertex points $\{P\}_i$ is a partitioning of the plane into convex irregular polygonal cells. The vertices of these polygons are the circumcentres of the corresponding Delaunay triangles (duality). Gradually expanding circles at equal rates from each vertex produces the Voronoi diagram when any two circles 'collide'. The closest vertex to any given point is the vertex contained within the cell.



Delaunay Triangulation

vertices in black, circumcentres in red



Voronoi Diagram

vertices in black, cells (partitions) in red

The Delaunay and Voronoi partitions are useful in modelling a wide variety of phenomena. They can also be extended to higher dimensions.

The Bowyer-Watson algorithm computes the Delaunay triangulation of a vertex set.

2.4. Vectors and 3D Geometry

2.4.1. Direction Cosines

If a vector $\mathbf{a} = (a_x, a_y, a_z)$ makes angles α, β, γ with an orthogonal set of x -, y - and z -axes, then:

- the quantities $\cos \alpha, \cos \beta$ and $\cos \gamma$ are the direction cosines of \mathbf{a} .
- $\cos \alpha = a_x / |\mathbf{a}|$, etc.
- $\mathbf{a} = |\mathbf{a}| (\cos \alpha \mathbf{i} + \cos \beta \mathbf{j} + \cos \gamma \mathbf{k})$.

The direction cosines represent the component of a unit vector along \mathbf{a} parallel to each axis.

2.4.2. Scalar Product (Dot Product) and Vector Product (Cross Product) Algebra

The scalar and vector products are defined by components as

$$\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^\top \mathbf{b} = \sum_i a_i b_i = a_x b_x + a_y b_y + a_z b_z$$

$$\mathbf{a} \times \mathbf{b} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ a_x & a_y & a_z \\ b_x & b_y & b_z \end{vmatrix} = (a_y b_z - a_z b_y) \mathbf{i} + (a_z b_x - a_x b_z) \mathbf{j} + (a_x b_y - a_y b_x) \mathbf{k}$$

In terms of magnitudes and angles,

$$\mathbf{a} \cdot \mathbf{b} = |\mathbf{a}| |\mathbf{b}| \cos \theta$$

$$\mathbf{a} \times \mathbf{b} = |\mathbf{a}| |\mathbf{b}| \sin \theta \hat{\mathbf{n}}$$

Commutative / anticommutative properties:

$$\mathbf{a} \cdot \mathbf{b} = \mathbf{b} \cdot \mathbf{a}$$

$$\mathbf{a} \times \mathbf{b} = -\mathbf{b} \times \mathbf{a}$$

Useful identities: (for the triple product identities see Section 2.4.4)


$$|\mathbf{a} \times \mathbf{b}|^2 + |\mathbf{a} \cdot \mathbf{b}|^2 = (|\mathbf{a}| |\mathbf{b}|)^2 \quad \text{(Lagrange's identity)}$$

$$|\mathbf{a} \pm \mathbf{b}|^2 = \mathbf{a} \cdot \mathbf{a} + \mathbf{b} \cdot \mathbf{b} \pm 2(\mathbf{a} \cdot \mathbf{b}) \quad \text{(from cosine rule)}$$


$$\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) + \mathbf{c} \times (\mathbf{a} \times \mathbf{b}) + \mathbf{b} \times (\mathbf{c} \times \mathbf{a}) = \mathbf{0} \quad \text{(Jacobi identity)}$$

$$(\mathbf{a} \times \mathbf{b}) \cdot (\mathbf{c} \times \mathbf{d}) = (\mathbf{a} \cdot \mathbf{c})(\mathbf{b} \cdot \mathbf{d}) - (\mathbf{b} \cdot \mathbf{c})(\mathbf{a} \cdot \mathbf{d}) \quad \text{(Binet-Cauchy identity)}$$

2.4.3. Right-Hand Rules for Vector Product Orientation



Right hand rule:



Right hand grip rule:
for rotations

2.4.4. Triple Products

Scalar triple product:

$$\begin{aligned}
 \mathbf{a} \cdot (\mathbf{b} \times \mathbf{c}) &= \begin{vmatrix} a_x & a_y & a_z \\ b_x & b_y & b_z \\ c_x & c_y & c_z \end{vmatrix} = \begin{vmatrix} a_x & b_x & c_x \\ a_y & b_y & c_y \\ a_z & b_z & c_z \end{vmatrix} \\
 &= \mathbf{b} \cdot (\mathbf{c} \times \mathbf{a}) = \mathbf{c} \cdot (\mathbf{a} \times \mathbf{b}) \\
 &= -\mathbf{a} \cdot (\mathbf{c} \times \mathbf{b}) = -\mathbf{c} \cdot (\mathbf{b} \times \mathbf{a}) = -\mathbf{b} \cdot (\mathbf{a} \times \mathbf{c})
 \end{aligned}$$

Vector triple product:

$$\begin{aligned}
 \mathbf{a} \times (\mathbf{b} \times \mathbf{c}) &= (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c} \\
 (\mathbf{a} \times \mathbf{b}) \times \mathbf{c} &= (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{b} \cdot \mathbf{c})\mathbf{a}
 \end{aligned}$$

2.4.5. Vector Products for Areas and Volumes

Area of triangle spanned by \mathbf{a} and \mathbf{b} : $A = \frac{1}{2} |\mathbf{a} \times \mathbf{b}|$

Volume of parallelepiped spanned by \mathbf{a} , \mathbf{b} and \mathbf{c} : $V = |\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c})|$

Volume of tetrahedron spanned by \mathbf{a} , \mathbf{b} and \mathbf{c} : $V = \frac{1}{6} |\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c})|$

Volume of a polyhedron of N faces, triangulated as $\bigcap_{i=1}^N \{\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i\}$: $V = \frac{1}{6} \left| \sum_{i=1}^N \mathbf{a}_i \cdot (\mathbf{b}_i \times \mathbf{c}_i) \right|$

2.4.6. Equations of Lines and Planes

Line ($\mathbf{p} = [p_x \ p_y \ p_z]^T$: point on line, $\mathbf{d} = [d_x \ d_y \ d_z]^T$: direction vector):

$$\frac{x-p_x}{d_x} = \frac{y-p_y}{d_y} = \frac{z-p_z}{d_z} (= t) \quad \mathbf{r} = \mathbf{p} + t \mathbf{d} \quad (\mathbf{r} - \mathbf{p}) \times \mathbf{d} = \mathbf{0}$$

scalar (if not parallel to any axis) parametric non-parametric

Plane (\mathbf{p} : point on plane, $\{\mathbf{a}, \mathbf{b}\}$: vectors in plane, $\mathbf{n} = [n_x \ n_y \ n_z]^T$: normal vector to plane):

$$n_x x + n_y y + n_z z = d \quad \mathbf{r} = \mathbf{p} + s \mathbf{a} + t \mathbf{b} \quad (\mathbf{r} - \mathbf{p}) \cdot \mathbf{n} = 0$$

scalar parametric non-parametric

Plane-Plane Intersection: $\Pi_1: (\mathbf{r} - \mathbf{p}_1) \cdot \mathbf{n}_1 = 0$ and $\Pi_2: (\mathbf{r} - \mathbf{p}_2) \cdot \mathbf{n}_2 = 0$ intersect in a line with direction vector given by $(\mathbf{n}_1 \times \mathbf{n}_2)$. In a view parallel to this vector, the two planes are projected as straight lines parallel to all viewing planes, which can simplify problems.

2.4.7. Shortest Distances Between Points, Planes and Lines

Point \mathbf{c} to line $\mathbf{r} = \mathbf{p} + t \mathbf{d}$:

$$d_{\min} = |(\mathbf{c} - \mathbf{p}) \times \mathbf{d}|$$

Point \mathbf{c} to plane $(\mathbf{r} - \mathbf{p}) \cdot \mathbf{n} = 0$:

$$d_{\min} = \frac{|(\mathbf{c} - \mathbf{p}) \cdot \mathbf{n}|}{|\mathbf{n}|}$$

Skew lines $\mathbf{r}_1 = \mathbf{p}_1 + t \mathbf{d}_1$ and $\mathbf{r}_2 = \mathbf{p}_2 + t \mathbf{d}_2$:

$$d_{\min} = \frac{|(\mathbf{p}_1 - \mathbf{p}_2) \cdot (\mathbf{d}_1 \times \mathbf{d}_2)|}{|\mathbf{d}_1 \times \mathbf{d}_2|}$$

To find the point(s) of closest approach, define generalised point(s) (\mathbf{r}_1 and \mathbf{r}_2) on the object(s) in parametric form and assert perpendicularity: solve $(\mathbf{r}_1 - \mathbf{r}_2) \cdot \mathbf{d}_1 = 0$ and $(\mathbf{r}_1 - \mathbf{r}_2) \cdot \mathbf{d}_2 = 0$ for parameters t_1 and t_2 .

Shortest distance to an (external) sphere is shortest distance to its centre, minus the radius of the sphere.

Shortest distance to an (external) cylinder is the shortest distance to its axis, minus the radius of the cylinder.

2.4.8. Vector Equations of Curved Surfaces

Sphere, centre \mathbf{c} , radius R :

$$(\mathbf{r} - \mathbf{c}) \cdot (\mathbf{r} - \mathbf{c}) = R^2 \quad \text{or} \quad |\mathbf{r} - \mathbf{c}| = R$$

Double cone, axis \mathbf{n} , opening half-angle θ :

$$(\mathbf{r} \cdot \mathbf{n})^2 = \mathbf{r} \cdot \mathbf{r} \cos^2 \theta \quad \text{or} \quad \mathbf{r} \cdot \mathbf{n} = |\mathbf{r}| \cos \theta$$

Cylinder, axis \mathbf{n} , radius R :

$$(\mathbf{r} \times \mathbf{n}) \cdot (\mathbf{r} \times \mathbf{n}) = R^2 \quad \text{or} \quad |\mathbf{r} \times \mathbf{n}| = R$$

Ellipsoid of revolution,

$$|\mathbf{r} - \mathbf{f}_1| + |\mathbf{r} - \mathbf{f}_2| = 2a$$

foci \mathbf{f}_1 and \mathbf{f}_2 , semi-major axis a :

Hyperboloid of revolution (two sheets),

$$||\mathbf{r} - \mathbf{f}_1| - |\mathbf{r} - \mathbf{f}_2|| = 2a$$

foci \mathbf{f}_1 and \mathbf{f}_2 , semi-major axis a :

2.4.9. Surfaces and Volumes of Revolution and their Centres of Mass

For a plane curve rotated 360° about a horizontal axis to produce an axisymmetric solid or surface (shell), properties are:

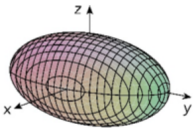
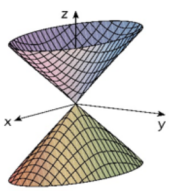
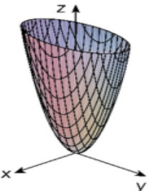
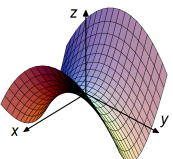
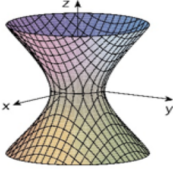
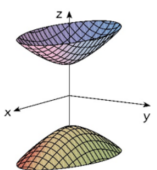
	Volume	COM (solid)	Surface area	COM (shell)
Cartesian, $y(x)$ revolving around the x -axis	$\pi \int_a^b y^2 dx$	$\frac{\int_a^b xy^2 dx}{\int_a^b y^2 dx}$	$2\pi \int_a^b y \sqrt{1 + (y')^2} dx$	$\frac{\int_a^b y^2 \sqrt{1 + (y')^2} dx}{\int_a^b y \sqrt{1 + (y')^2} dx}$
Parametric, $\{x(t), y(t)\}$ revolving around the x -axis	$\pi \int_a^b y^2 x' dt$	$\frac{\int_a^b xy^2 x' dt}{\int_a^b y^2 x' dt}$	$2\pi \int_a^b y \sqrt{(x')^2 + (y')^2} dt$	$\frac{\int_a^b x'y^2 \sqrt{(x')^2 + (y')^2} dt}{\int_a^b x'y \sqrt{(x')^2 + (y')^2} dt}$

The COM (centre of mass) is given as its x -coordinate, \bar{x} , and assumes uniform density. The other ordinates are $\bar{y} = \bar{z} = 0$.

For COMs of common geometric figures, see Section 6.3.

For Pappus' theorems for solids of revolution and solids swept along a curve, see Section 2.1.5.

2.4.10. Quadric Surfaces (3D Extensions of Conic Sections)

Surface	Cartesian	Parametric	Canonical Matrix Q
 <p>Ellipsoid</p>	$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1$ <p>a, b, c: semi-axes</p>	$\begin{aligned} x &= a \cos u \sin v \\ y &= b \sin u \sin v \\ z &= c \cos v \end{aligned}$ <p>$0 \leq u < 2\pi, 0 \leq v < \pi$</p>	$\begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & \frac{1}{b^2} & 0 & 0 \\ 0 & 0 & \frac{1}{c^2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$
 <p>Cone</p>	$\frac{z^2}{c^2} = \frac{x^2}{a^2} + \frac{y^2}{b^2}$ <p>$\frac{c}{a}, \frac{c}{b}$: slopes in xz and yz planes</p>	$\begin{aligned} x &= av \cos u \\ y &= bv \sin u \\ z &= cv \end{aligned}$ <p>$0 \leq u < 2\pi$</p>	$\begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & \frac{1}{b^2} & 0 & 0 \\ 0 & 0 & -\frac{1}{c^2} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$
 <p>Elliptic Paraboloid</p>	$\frac{z}{c} = \frac{x^2}{a^2} + \frac{y^2}{b^2}$ <p>$\sqrt{\frac{z}{c}} a, \sqrt{\frac{z}{c}} b$: semi-axes of ellipse cross-section at z</p>	$\begin{aligned} x &= av \cos u \\ y &= bv \sin u \\ z &= cv^2 \end{aligned}$ <p>$0 \leq u < 2\pi, v \geq 0$</p>	$\begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & \frac{1}{b^2} & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{2c} \\ 0 & 0 & -\frac{1}{2c} & 0 \end{bmatrix}$
 <p>Hyperbolic Paraboloid</p>	$\frac{z}{c} = \frac{x^2}{a^2} - \frac{y^2}{b^2}$ <p>$\sqrt{\frac{z}{c}} a$: focal length of hyperbola cross-section at z</p>	$\begin{aligned} x &= av \cosh u \\ y &= bv \sinh u \\ z &= cv^2, \end{aligned}$ <p>(only for $\frac{y}{x} \leq \frac{b}{a}$)</p>	$\begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & -\frac{1}{b^2} & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{2c} \\ 0 & 0 & -\frac{1}{2c} & 0 \end{bmatrix}$
 <p>Hyperboloid, one sheet</p>	$\frac{x^2}{a^2} + \frac{y^2}{b^2} - \frac{z^2}{c^2} = 1$ <p>$\sqrt{1 + \frac{z^2}{c^2}} a, \sqrt{1 + \frac{z^2}{c^2}} b$: semi-axes of ellipse cross-section at z</p>	$\begin{aligned} x &= a \cos u \cosh v \\ y &= b \sin u \cosh v \\ z &= c \sinh v \end{aligned}$ <p>$0 \leq u < 2\pi$</p>	$\begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & \frac{1}{b^2} & 0 & 0 \\ 0 & 0 & -\frac{1}{c^2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$
 <p>Hyperboloid, two sheets</p>	$-\frac{x^2}{a^2} - \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1$ <p>$\sqrt{\frac{z^2}{c^2} - 1} a, \sqrt{\frac{z^2}{c^2} - 1} b$: semi-axes of ellipse cross-section at z</p>	$\begin{aligned} x &= a \cos u \sinh v \\ y &= b \sin u \sinh v \\ z &= \pm c \cosh v \end{aligned}$ <p>$0 \leq u < \pi$</p>	$\begin{bmatrix} -\frac{1}{a^2} & 0 & 0 & 0 \\ 0 & -\frac{1}{b^2} & 0 & 0 \\ 0 & 0 & \frac{1}{c^2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$

2.4.11. Quadrics in Homogeneous Coordinates

Any quadric surface can be represented by the equation

$$a_{xx}x^2 + a_{yy}y^2 + a_{zz}z^2 + 2a_{xy}xy + 2a_{yz}yz + 2a_{xz}xz + 2a_x x + 2a_y y + 2a_z z + a_1 = 0$$

which can be represented in matrix form as $\mathbf{x}^T \mathbf{Q} \mathbf{x} = 0$ (canonical form) where

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad \text{and} \quad \mathbf{Q} = \begin{bmatrix} a_{xx} & a_{xy} & a_{xz} & a_x \\ a_{xy} & a_{yy} & a_{yz} & a_y \\ a_{xz} & a_{yz} & a_{zz} & a_z \\ a_x & a_y & a_z & a_1 \end{bmatrix} \quad \text{is a symmetric } 4 \times 4 \text{ matrix.}$$

Quadric Equation from Points: nine points on a quadric define the quadric: (x_i, y_i) ($i = 1, 2, \dots, 9$).

The coefficients $a_{xx}, a_{xy}, \dots, a_1$ can be found by evaluating the ten 9×9 sub-determinants across the top row of the 10×10 matrix \mathbf{A} , where row 1 is $[x^2 \ y^2 \ z^2 \ xy \ yz \ xz \ x \ y \ z \ 1]$ and rows 2-10 are:

$$[x_i^2 \ y_i^2 \ z_i^2 \ x_i y_i \ y_i z_i \ x_i z_i \ x_i \ y_i \ z_i \ 1]$$

The equation satisfies $|\mathbf{A}| = 0$. Note that the off-diagonal coefficients in \mathbf{Q} have a factor of $\frac{1}{2}$.

Quadric-Line Intersection (Raytracing of Quadric Surfaces)

A line $\mathbf{x} = \mathbf{p} + \lambda \mathbf{n}$ intersects a quadric $\mathbf{x}^T \mathbf{Q} \mathbf{x} = 0$ at values of λ satisfying the quadratic

$$(\mathbf{n}^T \mathbf{Q} \mathbf{n}) \lambda^2 + (2 \mathbf{n}^T \mathbf{Q} \mathbf{p}) \lambda + (\mathbf{p}^T \mathbf{Q} \mathbf{p}) = 0.$$

A quadric generally intersects a **plane** in a conic section curve in the plane.

Affine Transformations on Quadric Surfaces

Transformations can be applied to a quadric using

$$\mathbf{Q}' = (\mathbf{R}^{-1})^T \mathbf{Q} \mathbf{R}^{-1} \quad \text{where } \mathbf{R} \text{ is a } 4 \times 4 \text{ affine transformation matrix mapping a quadric represented by } \mathbf{Q} \text{ into a quadric represented by } \mathbf{Q}'$$

For the general form of the affine transformation matrix \mathbf{R} , see Section 4.2.3.

Singular Value Decomposition of a Quadric Canonical Matrix

For any quadric $\mathbf{x}^T \mathbf{Q}' \mathbf{x} = 0$, define the matrix \mathbf{M} as the 3×3 matrix formed from the first three rows and columns of \mathbf{Q}' (i.e. $\mathbf{M} : \mathbf{Q} \rightarrow \mathbf{Q}'$ is a linear transformation matrix from a quadric \mathbf{Q} aligned with the $\{\mathbf{i}, \mathbf{j}, \mathbf{k}\}$ axes into the quadric represented by \mathbf{Q}' translated to the origin).

If the singular value decomposition (SVD, see Section 4.3.7) is written as $\mathbf{M} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$, then:

- The columns of \mathbf{U} represent the normalised **principal axes** of \mathbf{Q}' .
- The columns of \mathbf{V} (rows of \mathbf{V}^T) represent the vectors on the **original** quadric \mathbf{Q} which are mapped to the principal axes of \mathbf{Q}' .
- The singular values are the linear **scale factors** in the corresponding axes.

M3. CALCULUS

3.1. Limits and Numerical Methods

3.1.1. Formal Definition of Limits

Let $f(x)$ be defined for all real $x \neq a$ over an open interval containing a . We say that

$$\lim_{x \rightarrow a} f(x) = L \quad (\text{the (two-sided) limit of } f(x) \text{ as } x \text{ approaches } a \text{ is } L)$$

if, for every $\varepsilon > 0$, there exists some $\delta > 0$, such that if $0 < |x - a| < \delta$, then $|f(x) - L| < \varepsilon$:


$$\lim_{x \rightarrow a} f(x) = L \Leftrightarrow (\forall \varepsilon > 0)(\exists \delta > 0)(\forall x \in D): (0 < |x - a| < \delta \Rightarrow |f(x) - L| < \varepsilon)$$

The one-sided limits are $\lim_{x \rightarrow a^+} f(x)$ (right-sided) and $\lim_{x \rightarrow a^-} f(x)$ (left-sided).

- For limits to infinity, the condition is $x > \delta$ (if a is $+\infty$) or $x < -\delta$ (if a is $-\infty$).
- For one-sided limits, use $0 < x - a < \delta$ (right-sided) or $0 < a - x < \delta$ (left-sided).

3.1.2. Limits at Discontinuities and Circle Notation

A graph of a function $y = f(x)$ should include open circles \circ for limiting values and closed circles \bullet for defined values. For example:



$$\lim_{x \rightarrow -2} f(x) = 0 \quad (\text{removable discontinuity})$$

$$\lim_{x \rightarrow -1} f(x) = -1 \quad \text{but} \quad f(-1) = 1 \quad (\text{jump discontinuity})$$

$$\lim_{x \rightarrow 0^-} f(x) = 0 \quad \text{and} \quad \lim_{x \rightarrow 0^+} f(x) = f(0) = 1$$

$$\lim_{x \rightarrow 1^-} f(x) = 0 \quad \text{and} \quad \lim_{x \rightarrow 1^+} f(x) = 3 \quad \text{and} \quad f(1) = 2$$

$$\lim_{x \rightarrow 0} f(x) \quad \text{and} \quad \lim_{x \rightarrow 1} f(x) \quad \text{do not exist.}$$

A discontinuity at $x = a$ is said to be removable if $\lim_{x \rightarrow a} f(x)$ exists. Continuity can be

established by including $x = a$ in the domain of $f(x)$, at which $f(a) := \lim_{x \rightarrow a} f(x) = \lim_{x \rightarrow a^\pm} f(x)$.

3.1.2. Standard Limits

Asymptotic ($x \rightarrow \infty$) growth order: $x^x \gg x! \gg a^x \gg x^a, x^2 \gg x \log x \gg x \gg \log x \gg 1$.

Limits to values:

- | | |
|---|---|
| <ul style="list-style-type: none"> • $\lim_{x \rightarrow \infty} (x^a p^x) = 0 \quad p < 1, \forall a \in \mathbb{R}.$ • $\lim_{x \rightarrow 0} (x^a \ln x) = 0, \quad \forall a > 0.$ • $\lim_{x \rightarrow 0} \frac{\sin ax}{x} = \lim_{x \rightarrow 0} \frac{\tan ax}{x} = a$ • $\lim_{x \rightarrow \infty} \left(\frac{a^x}{x!} \right) = 0.$ | <ul style="list-style-type: none"> • $\lim_{x \rightarrow a} \frac{x^n - a^n}{x - a} = na^{n-1}, \quad n \in \mathbb{Q}.$ • $\lim_{x \rightarrow 0} \frac{e^x - 1}{x} = 1.$ • $\lim_{x \rightarrow 0} \frac{a^x - 1}{x} = \ln a, \quad a > 0.$ |
|---|---|

Limits to functions:

- $\lim_{n \rightarrow \infty} \left(1 + \frac{x}{n} \right)^n = \lim_{n \rightarrow 0} (1 + nx)^{1/n} = e^x$
- $\lim_{n \rightarrow \infty} \cos^n \left(\frac{x}{\sqrt{n}} \right) = e^{-\frac{1}{2}x^2}$
- $\lim_{n \rightarrow \infty} \cosh^n \left(\frac{x}{\sqrt{n}} \right) = e^{\frac{1}{2}x^2}$

Derivative and integral as limit definitions:

- $\lim_{n \rightarrow \infty} n \left(f \left(x + \frac{a}{n} \right) - f(x) \right) = af'(x), \quad \forall a \in \mathbb{R} \quad \text{(forward difference)}$
- $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n f \left(\frac{k}{n} \right) = \int_0^1 f(x) \, dx \quad \text{(Riemann summation: rectangular rule)}$

For Stirling's formula involving asymptotic expressions for $n!$ and $\ln n!$, see Section 1.7.1.

3.1.3. Numerical Differentiation

Difference Quotients (finite difference approximations to derivatives):

- First derivative, forward difference: $f'(x) \approx \frac{f(x+h) - f(x)}{h} + O(h)$
- First derivative, backward difference: $f'(x) \approx \frac{f(x) - f(x-h)}{h} + O(h)$
- First derivative, central difference: $f'(x) \approx \frac{f(x+h) - f(x-h)}{2h} + O(h^2)$
- Second derivative, central difference: $f''(x) \approx \frac{f(x+h) - 2f(x) + f(x-h)}{h^2} + O(h^2)$

For partial derivatives (central differences):

- First partials: $\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x+h, y) - f(x-h, y)}{2h}$ and $\frac{\partial f(x, y)}{\partial y} \approx \frac{f(x, y+h) - f(x, y-h)}{2h}$
- Second partials: $\frac{\partial^2 f(x, y)}{\partial x^2} \approx \frac{f(x+h, y) - 2f(x, y) + f(x-h, y)}{h^2}$
- Mixed partials: $\frac{\partial^2 f(x, y)}{\partial x \partial y} \approx \frac{f(x+h, y+h) - f(x-h, y+h) - f(x+h, y-h) + f(x-h, y-h)}{4h^2}$

For discrete sequences:


- One-sided: $\frac{u_{n+1} - u_n}{\Delta t} = \frac{du}{dt} + \frac{d^2u}{dt^2} \frac{\Delta t}{2!} + \dots$
- Two-sided: $\frac{u_{n+1} - u_{n-1}}{2\Delta t} = \frac{du}{dt} + \frac{d^3u}{dt^3} \frac{\Delta t^2}{3!} + \dots$

Python (SciPy): `f` is a callable function

```
from scipy.misc import derivative
gradient = derivative(f, x, dx=h)
```

3.1.4. Numerical Integration by Riemann Summation

A definite integral $I = \int_a^b f(x) dx$ can be approximated by splitting the interval $[a, b]$ into N equally-spaced intervals containing $N + 1$ ordinates from $x_0 = a$ to $x_N = b$ inclusive.



- Midpoint Rule:
$$I \approx \sum_{i=1}^N f(x_i^*) \cdot \Delta x$$
- Trapezium Rule:
$$I \approx \frac{\Delta x}{2} \left(f(x_0) + f(x_N) + 2 \sum_{i=1}^{N-1} f(x_i) \right)$$
- Simpson's Rule:
$$I \approx \frac{\Delta x}{2} \left(f(x_0) + f(x_N) + 4 \sum_{i=1}^{\frac{N}{2}} f(x_{2i-1}) + 2 \sum_{i=1}^{\frac{N}{2}-1} f(x_{2i}) \right)$$

 (Simpsons' pattern of factors: 1, 4, 2, 4, 2, ..., 4, 1)

Maximum error bounds: if E is the absolute error in the approximation to I , then

Midpoint Error	Trapezium Error	Simpson's Error
$ E_M \leq \frac{K(b-a)^3}{24N^2}$,	$ E_T \leq \frac{K(b-a)^3}{12N^2}$,	$ E_S \leq \frac{M(b-a)^5}{180N^4}$.

where $|f''(x)| \leq K$ and $|f^{(4)}(x)| \leq M$ for all $a \leq x \leq b$.

Python (SciPy): `f` is a callable function representing the integrand


```
from scipy import integrate
val, abserr = integrate.quad(f, a, b)
```

3.1.5. Fixed Point Iteration for Solving Algebraic Equations


A root $x = \alpha$ to a single-variable equation $f(x) = 0$ can sometimes be found by writing the equation in the form $g(x) = x$ (so that $g(x) = f(x) + x$) and iterating (for suitable x_1):

$$x_{n+1} = g(x_n) \quad \Leftrightarrow \quad \lim_{n \rightarrow \infty} x_n = \alpha = g(\alpha) \text{ if } |g'(\alpha)| \leq 1$$

The iteration process, if it converges, can be represented as:



Convergent Staircase



Convergent Cobweb

Convergence Behaviour: the type of iteration depends on the gradient of $g(x)$ in the region between the initial point and the true root. Convergence requires $|g'(x)| < 1$. If the true root is $x = \alpha$ then the behaviour around the root is:

- Convergent Staircase: when $0 \leq g'(\alpha) \leq 1$.
- Convergent Cobweb: when $-1 \leq g'(\alpha) < 0$.
- Divergent Staircase: when $g'(\alpha) > 1$.
- Divergent Cobweb: when $g'(\alpha) < -1$.

If $0 < |g'(x)| < 1$ then convergence is linear:

$$\lim_{n \rightarrow \infty} \frac{x_{n+1} - \alpha}{x_n - \alpha} = g'(\alpha).$$

If $g'(\alpha) = 0$ then convergence is quadratic:


$$\lim_{n \rightarrow \infty} \frac{|x_{n+1} - \alpha|}{|x_n - \alpha|^2} = \lambda \text{ (some constant).}$$

3.1.6. Newton-Raphson Method (Gradient Descent) For Solving Algebraic Equations

A root $x = \alpha$ to a single-variable equation $f(x) = 0$ can be found by iterating (for suitable x_1):

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \quad \Leftrightarrow \quad \lim_{n \rightarrow \infty} x_n = \alpha \quad (\text{if convergent})$$

The iteration process forms a pattern as follows:



Convergence Behaviour: convergence of Newton's method is generally difficult to predict.

If the root is a single root then convergence is quadratic.

If the root is repeated (algebraic multiplicity m) then convergence is linear, but can be accelerated to quadratic by using the iteration $x_{n+1} = x_n - m \times \frac{f(x_n)}{f'(x_n)}$ instead.

Generalisation to Multivariable Equations: for systems of algebraic equations of the form $\mathbf{f}(\mathbf{x}) = \mathbf{0}$, where \mathbf{f} is a vector-valued function containing each equation and \mathbf{x} is a vector of variables, Newton's method is $\mathbf{x}_{n+1} = \mathbf{x}_n - \mathbf{J}^{-1}(\mathbf{x}_n) \mathbf{f}(\mathbf{x}_n)$ where \mathbf{J} is the Jacobian matrix (Section 3.5.2) of \mathbf{f} . For increased efficiency and numerical stability, instead of computing \mathbf{J}^{-1} , the system of linear equations $\mathbf{J}(\mathbf{x}_n) (\mathbf{x}_{n+1} - \mathbf{x}_n) = -\mathbf{f}(\mathbf{x}_n)$ can be solved for $\mathbf{x}_{n+1} - \mathbf{x}_n$ at each iteration.

Python: f is a callable function representing the system (input: array \mathbf{x} ; output: array $\mathbf{f}(\mathbf{x})$)

```
from scipy.optimize import fsolve
root = fsolve(f, x0)
```


3.1.7. Numerical Methods for ODEs

An ODE integrator operates on either an individual ODE ($y' = f(t, y)$) or a system of coupled ODEs ($\mathbf{y}' = \mathbf{f}(t, \mathbf{y})$). The iteration is of fixed step size h , i.e. $t_{n+1} = t_n + h$.

Euler's Method:

$$y_{n+1} = y_n + hf(t_n, y_n)$$

Euler's Improved Method:

$$y_{n+1} = y_{n-1} + 2hf(t_n, y_n)$$

Predictor-Corrector (Heun's method):

$$\hat{y}_{n+1} = y_n + hf(t_n, y_n)$$

$$y_{n+1} = y_n + \frac{h}{2}[f(t_n, y_n) + f(t_{n+1}, \hat{y}_{n+1})]$$

Störmer-Verlet Integrator (symplectic: ideal for position-velocity-acceleration equations)

For a **2nd-order** ODE given by $\mathbf{x}'' = \mathbf{A}(t, \mathbf{x}, \mathbf{v})$, (with $\mathbf{v} = \mathbf{x}'$ and $t_{n+1} = t_n + \Delta t$)

1. set $\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{v}_0 \Delta t + \frac{1}{2} \mathbf{A}(\mathbf{x}_0) \Delta t^2$,
2. for $n = 1, 2, \dots$ iterate

$$\mathbf{x}_{n+1} = 2\mathbf{x}_n - \mathbf{x}_{n-1} + \mathbf{A}(\mathbf{x}_n) \Delta t^2.$$

For the more sophisticated Gauss-Jackson integration algorithm, see Section 9.1.6.

Runge-Kutta Method (RK4; implicit 4th order):

$$y_{n+1} = y_n + \frac{h}{6} (k_{n1} + 2k_{n2} + 2k_{n3} + k_{n4})$$

$$k_{n1} = f(t_n, y_n)$$

$$k_{n2} = f(t_n + 0.5h, y_n + 0.5hk_{n1})$$

$$k_{n3} = f(t_n + 0.5h, y_n + 0.5hk_{n2})$$

$$k_{n4} = f(t_n + h, y_n + hk_{n3})$$

Butcher Tableau for RK4

(see Section 3.1.7.)

0				
1/2	1/2			
1/2	0	1/2		
1	0	0	1	
	1/6	1/3	1/3	1/6

Programming:

MATLAB: implements RK4. f is a callable function for the RHS.

```
[t, y] = ode45(f, [t_start, t_end], y0)
```

Python (SciPy): f is a callable function representing the RHS (inputs: t and \mathbf{y} ; output: \mathbf{y}')

```
from scipy.integrate import odeint
y = odeint(f, y0, t_array)
```

3.1.8. Butcher Tableau for Generalised Runge-Kutta ODE Integrators

A general RK method has an iteration step of the form $y_{n+1} = y_n + h \sum_{i=1}^s b_i k_{ni}$.

(s : order of the method, b_i : coefficients, k_{ni} evaluations of f near t_n and y_n (below)).

Explicit Scheme

$$k_{ni} = f\left(t_n + c_i h, y_n + h \sum_{j=1}^{i-1} a_{ij} k_{nj}\right)$$

0					
c_2	a_{21}				
c_3	a_{31}	a_{32}			
\vdots	\vdots		\ddots		
c_s	a_{s1}	a_{s2}	\cdots	$a_{s,s-1}$	
	b_1	b_2	\cdots	b_{s-1}	b_s

Implicit Scheme

$$k_{ni} = f\left(t_n + c_i h, y_n + h \sum_{j=1}^s a_{ij} k_{nj}\right)$$

c_1	a_{11}	a_{12}	\cdots	a_{1s}
c_2	a_{21}	a_{22}	\cdots	a_{2s}
\vdots	\vdots		\ddots	\vdots
c_s	a_{s1}	a_{s2}	\cdots	$a_{s,s}$
	b_1	b_2	\cdots	b_s

This allows any RK method to be represented as a matrix \mathbf{a} , a vector \mathbf{b} and a vector \mathbf{c} , allowing for efficient computation.

3.2. Series Expansions

3.2.1. Single Variable Definitions of Maclaurin Series and Taylor Series

The **Maclaurin series** about $x = 0$:

$$f(x) = \sum_{r=0}^{\infty} \frac{f^{(r)}(0)}{r!} \cdot x^r = f(0) + f'(0)x + \frac{1}{2}f''(0)x^2 + \frac{1}{6}f'''(0)x^3 + \dots$$

The sequence of coefficients of x^n is given by $\mathcal{Z}^{-1}[f(x^{-1})](n)$ where \mathcal{Z}^{-1} is the inverse z -transform (see Section 3.4.12).

The **Taylor series** about some value a :

$$f(x) = \sum_{r=0}^{\infty} \frac{f^{(r)}(a)}{r!} \cdot (x - a)^r = f(a) + f'(a)(x - a) + \frac{1}{2}f''(a)(x - a)^2 + \dots$$

This can also be written in the form (expanding about a fixed x)

$$f(x + \delta x) = \sum_{r=0}^{\infty} \frac{f^{(r)}(x)}{r!} \cdot h^r = f(x) + f'(x)\delta x + \frac{1}{2}f''(x)(\delta x)^2 + \dots$$

3.2.3. Maclaurin Series Expansions of Common Functions

Exponentials and Logarithms:

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots = \sum_{n=0}^{\infty} \frac{x^n}{n!} \quad \text{for all complex } x$$

$$e^{-x} = 1 - x + \frac{x^2}{2!} - \frac{x^3}{3!} + \cdots = \sum_{n=0}^{\infty} (-1)^n \frac{x^n}{n!} \quad \text{for all complex } x$$

$$\ln(1+x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \cdots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{n+1}}{n+1} \quad \text{for all } |x| \leq 1, x \neq -1, \text{ principal value}$$

$$\ln\left(\frac{1}{1-x}\right) = x + \frac{x^2}{2} + \frac{x^3}{3} + \cdots = \sum_{n=0}^{\infty} \frac{x^{n+1}}{n+1} \quad \text{for all } |x| \leq 1, x \neq -1, \text{ principal value}$$

Generalised binomial series expansion:

$$(1+x)^n = 1 + nx + \frac{n(n-1)}{2!}x^2 + \frac{n(n-1)(n-2)}{3!}x^3 + \cdots = \sum_{r=0}^{\infty} {}^n C_r x^r \quad \text{for all } |x| < 1$$

Trigonometric and hyperbolic functions:

$$\sin x = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \cdots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)!} \quad \text{for all complex } x$$

$$\cos x = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \cdots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n}}{(2n)!} \quad \text{for all complex } x$$

$$\tan x = x + \frac{1}{3}x^3 + \frac{2}{15}x^5 + \frac{17}{315}x^7 + \cdots \quad \text{for all } |x| < \frac{\pi}{2}$$

$$\sec x = 1 + \frac{1}{2}x^2 + \frac{5}{24}x^4 + \frac{61}{720}x^6 + \cdots \quad \text{for all } |x| < \frac{\pi}{2}$$

$$\sinh x = x + \frac{x^3}{3!} + \frac{x^5}{5!} + \cdots = \sum_{n=0}^{\infty} \frac{x^{2n+1}}{(2n+1)!} \quad \text{for all complex } x$$

$$\cosh x = 1 + \frac{x^2}{2!} + \frac{x^4}{4!} + \cdots = \sum_{n=0}^{\infty} \frac{x^{2n}}{(2n)!} \quad \text{for all complex } x$$

$$\tanh x = x - \frac{1}{3}x^3 + \frac{2}{15}x^5 - \frac{17}{315}x^7 + \cdots \quad \text{for all } |x| < \frac{\pi}{2}$$

$$\operatorname{sech} x = 1 - \frac{1}{2}x^2 + \frac{5}{24}x^4 - \frac{61}{720}x^6 + \cdots \quad \text{for all } |x| < \frac{\pi}{2}$$

Inverse trigonometric and inverse hyperbolic functions:

$$\sin^{-1} x = x + \frac{1}{6}x^3 + \frac{3}{40}x^5 + \frac{5}{112}x^7 + \dots \quad \text{for all } |x| \leq 1$$

$$\cos^{-1} x = \frac{\pi}{2} - x - \frac{1}{6}x^3 - \frac{3}{40}x^5 - \frac{5}{112}x^7 - \dots \quad \text{for all } |x| \leq 1$$

$$\tan^{-1} x = x - \frac{x^3}{3} + \frac{x^5}{5} - \frac{x^7}{7} + \dots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{2n+1} \quad \text{for all } |x| \leq 1, x \neq \pm i$$

$$\cot^{-1} x = \frac{\pi}{2} - x + \frac{x^3}{3} - \frac{x^5}{5} + \frac{x^7}{7} = \frac{\pi}{2} - \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{2n+1} \quad \text{for all } |x| \leq 1, x \neq \pm i$$

$$\sinh^{-1} x = x - \frac{1}{6}x^3 + \frac{3}{40}x^5 - \frac{5}{112}x^7 + \dots = \sum_{n=0}^{\infty} \frac{(-1)^n (2n)!}{2^{2n} (n!)^2 (2n+1)} x^{2n+1} \quad \text{for all } |x| < 1$$

$$\tanh^{-1} x = x + \frac{x^3}{3} + \frac{x^5}{5} + \frac{x^7}{7} + \dots = \sum_{n=0}^{\infty} \frac{x^{2n+1}}{2n+1} \quad \text{for all } |x| < 1$$

Special Functions:

$$\operatorname{erf} x = \frac{2}{\sqrt{\pi}}x - \frac{2}{3\sqrt{\pi}}x^3 + \frac{1}{5\sqrt{\pi}}x^5 - \frac{1}{21\sqrt{\pi}}x^7 + \dots = \frac{2}{\sqrt{\pi}} \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)n!} \quad \text{for all } x$$

3.2.4. Lagrange's Inversion Theorem

If $y = f(x)$ and therefore $x = f^{-1}(y)$ then a Taylor series for x about a is given by

$$f^{-1}(y) = x = a + \sum_{n=1}^{\infty} \frac{g_n}{n!} (y - f(a))^n$$

The coefficients are $\frac{g_n}{n!}$ where $g_n = \lim_{x \rightarrow a} \frac{d^{n-1}}{dx^{n-1}} \left(\frac{x-a}{f(x)-f(a)} \right)^n$.

For shifting of Maclaurin series, the inverse function of $f(x+a)$ is $f^{-1}(x) - a$.

The radius of convergence is not easily determined from the function alone.

3.2.5. Laurent Series

Laurent series allow for inclusion of poles by summing over all integer powers of x , often used for complex functions $f(z)$. The Laurent series for $f(z)$ about c is given by

$$f(z) = \sum_{n=-\infty}^{\infty} g_n (z - c)^n$$

The coefficients are $g_n = \frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)}{(z-c)^{n+1}} dz$.

(Cauchy contour integral around γ : counterclockwise Jordan curve where $f(z)$ holomorphic)

3.3. Differentiation and Integration

3.3.1. Continuity, Differentiability and Smoothness

A function $f(x)$ is said to be

- Continuous at $x = a$: if $\lim_{x \rightarrow a} f(x) = f(a)$
- Differentiable at $x = a$: if $\lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}$ exists
- Smooth (infinitely differentiable) at $x = a$: if $f^{(n)}(x)$ is continuous at $x = a$ for all nonnegative integer n .

If these terms are used without specifying a point $x = a$, then the condition must hold for all values of a in the domain of f .

3.3.2. Limit Definition of a Derivative

The first and second derivatives as limits are (as forward differences, assuming they exist):

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} \qquad f''(x) = \frac{f(x+h) - 2f(x) + f(x-h)}{h^2}$$

3.3.3. Limit Definition of a Definite Integral (Riemann Sum)

In general,
$$\int_a^b f(x) \, dx = \lim_{n \rightarrow \infty} \sum_{r=1}^n f(x_r) \Delta x$$

where $\Delta x = \frac{b-a}{n}$ and $x_r = a + (r-1)\Delta x$.

3.3.4. Mean Value Theorem and Intermediate Value Theorem

Mean Value Theorem: for a monotonically increasing function $f(x)$, we have $f(a) \leq M \leq f(b)$

where $M = \frac{1}{b-a} \int_a^b f(x) \, dx$ is the mean value of $f(x)$ on $a < x < b$.

Intermediate Value Theorem: for a continuous function $f(x)$ on the domain $[a, b]$, for all y such that $f(a) \leq y \leq f(b)$, there exists some $a \leq x \leq b$ such that $y = f(x)$.

Bolzano's Theorem: if a continuous function has values of opposite sign inside an interval, then it has a root in that interval.

3.3.5. Derivatives and Integrals of Functions

Algebraic, Exponential, Logarithmic, Trigonometric and Hyperbolic functions:

Function, $f(x)$	Derivative, $f'(x)$	Integral, $F(x) (+ C)$
1	0	x
x^n	nx^{n-1}	$\frac{1}{n+1}x^{n+1}$
\sqrt{x}	$\frac{1}{2\sqrt{x}}$	$\frac{2}{3}x^{3/2}$
e^x	e^x	e^x
a^x	$(\ln a)a^x$	$\frac{x^a}{\ln a}$
$\ln x $	$\frac{1}{x}$	$x(\ln x - 1)$
$\sin x$	$\cos x$	$-\cos x$
$\cos x$	$-\sin x$	$\sin x$
$\tan x$	$\sec^2 x$	$\ln \sec x $
$\sec x$	$\sec x \tan x$	$\ln \sec x + \tan x = \ln \left \tan \frac{x}{2} + \frac{\pi}{4} \right $
$\csc x$	$-\csc x \cot x$	$-\ln \csc x + \cot x = \ln \left \tan \frac{x}{2} \right $
$\cot x$	$-\csc^2 x$	$\ln \sin x $
$\sinh x$	$\cosh x$	$\cosh x$
$\cosh x$	$\sinh x$	$\sinh x$
$\tanh x$	$\operatorname{sech}^2 x$	$\ln \cosh x$
$\operatorname{sech} x$	$-\operatorname{sech} x \tanh x$	$2 \tan^{-1} \tanh \frac{x}{2} = \tan^{-1} \sinh x$
$\operatorname{csch} x$	$-\operatorname{csch} x \coth x$	$-\ln \operatorname{csch} x + \coth x = \ln \left \tanh \frac{x}{2} \right $
$\coth x$	$-\operatorname{csch}^2 x$	$\ln \sinh x $

Rationals with Radicals and Inverse Trigonometric / Inverse Hyperbolic Functions:

Function, $f(x)$	Derivative, $f'(x)$
$\sin^{-1} x$	$\frac{1}{\sqrt{1-x^2}}$
$\cos^{-1} x$	$\frac{-1}{\sqrt{1-x^2}}$
$\tan^{-1} x$	$\frac{1}{1+x^2}$
$\sec^{-1} x$	$\frac{1}{ x \sqrt{x^2-1}}$
$\sinh^{-1} x$	$\frac{1}{\sqrt{1+x^2}}$
$\cosh^{-1} x$	$\frac{1}{\sqrt{x^2-1}}$
$\tanh^{-1} x,$ $\coth^{-1} x$	$\frac{1}{1-x^2}$

Function, $f(x)$	Integral, $F(x) (+ C)$
$\frac{1}{\sqrt{a^2-x^2}}$	$\sin^{-1} \frac{x}{a}$
$\frac{a}{\sqrt{x^4-a^2x^2}}$	$\sec^{-1} \frac{x}{a}$
$\frac{1}{\sqrt{a^2+x^2}}$	$\sinh^{-1} \frac{x}{a}$
$\frac{1}{\sqrt{x^2-a^2}}$	$\cosh^{-1} \frac{x}{a}$
$\frac{1}{a^2+x^2}$	$\frac{1}{a} \tan^{-1} \frac{x}{a}$
$\frac{1}{a^2-x^2} (x < a)$	$\frac{1}{a} \tanh^{-1} \frac{x}{a}$
$\frac{1}{a^2-x^2} (x > a)$	$\frac{1}{a} \coth^{-1} \frac{x}{a}$

Integrals of radical functions:

$$\int \sqrt{a^2+x^2} dx = \frac{1}{2} \left(x\sqrt{a^2+x^2} + a^2 \sinh^{-1} \frac{x}{a} \right) + C$$

$$\int \sqrt{a^2-x^2} dx = \frac{1}{2} \left(x\sqrt{a^2-x^2} + a^2 \sin^{-1} \frac{x}{a} \right) + C$$

$$\int \sqrt{x^2-a^2} dx = \frac{1}{2} \left(x\sqrt{x^2-a^2} - a^2 \cosh^{-1} \frac{x}{a} \right) + C$$

Other common useful integrands:

$$\int \sec^3 x dx = \frac{1}{2} (\sec x \tan x + \ln |\sec x + \tan x|) + C$$

$$\int e^{ax} \sin bx dx = \frac{e^{ax} (a \sin bx - b \cos bx)}{a^2 + b^2} + C = \frac{e^{ax} \sin(bx - \tan^{-1} \frac{b}{a})}{\sqrt{a^2 + b^2}} + C$$

$$\int e^{ax} \cos bx dx = \frac{e^{ax} (a \cos bx + b \sin bx)}{a^2 + b^2} + C = \frac{e^{ax} \cos(bx - \tan^{-1} \frac{b}{a})}{\sqrt{a^2 + b^2}} + C$$

3.3.6. Indefinite Integral Reduction Formulas

Integral, I_n	Recurrence Relation
$I_n = \int \frac{x^n}{\sqrt{ax+b}} dx$	$I_n = \frac{2x^n \sqrt{ax+b}}{a(2n+1)} - \frac{2nb}{a(2n+1)} I_{n-1}$
$I_n = \int \frac{dx}{x^n \sqrt{ax+b}}$	$I_n = -\frac{\sqrt{ax+b}}{(n-1)bx^{n-1}} - \frac{a(2n-3)}{2b(n-1)} I_{n-1}$
$I_n = \int x^n \sqrt{ax+b} dx$	$I_n = \frac{2x^n \sqrt{(ax+b)^3}}{a(2n+3)} - \frac{2nb}{a(2n+3)} I_{n-1}$
$I_n = \int \frac{dx}{(px+q)^n \sqrt{ax+b}}$	$I_n = -\frac{\sqrt{ax+b}}{(n-1)(aq-bp)(px+q)^{n-1}} + \frac{a(2n-3)}{2(n-1)(aq-bp)} I_{n-1}$
$I_{m,n} = \int \frac{x^m dx}{(ax^2+bx+c)^n}$	$I_{m,n} = -\frac{x^{m-1}}{a(2n-m-1)(ax^2+bx+c)^{n-1}} - \frac{b(n-m)}{a(2n-m-1)} I_{m-1,n} + \frac{c(m-1)}{a(2n-m-1)} I_{m-2,n}$
$I_{m,n} = \int \frac{dx}{x^m(ax^2+bx+c)^n}$	$-c(m-1)I_{m,n} = \frac{1}{x^{m-1}(ax^2+bx+c)^{n-1}} + a(m+2n-3)I_{m-2,n} + b(m+n-2)I_{m-1,n}$
$I_n = \int \sin^n ax dx$	$I_n = -\frac{1}{an} \sin^{n-1} ax \cos ax + \frac{n-1}{n} I_{n-2}$
$I_n = \int \cos^n ax dx$	$I_n = \frac{1}{an} \sin ax \cos^{n-1} ax + \frac{n-1}{n} I_{n-2}$
$I_n = \int e^{ax} \sin^n bx dx$	$I_n = \frac{e^{ax} \sin^{n-1} bx}{a^2 + (bn)^2} (a \sin bx - bn \cos bx) + \frac{n(n-1)b^2}{a^2 + (bn)^2} I_{n-2}$
$I_n = \int e^{ax} \cos^n bx dx$	$I_n = \frac{e^{ax} \cos^{n-1} bx}{a^2 + (bn)^2} (a \cos bx + bn \sin bx) + \frac{n(n-1)b^2}{a^2 + (bn)^2} I_{n-2}$
$I_{m,n} = \int \sin^m ax \cos^n bx dx$	$I_{m,n} = \begin{cases} \frac{1}{a(n-1) \sin^{m-1} ax \cos^{n-1} ax} + \frac{m+n-2}{n-1} I_{m,n-2} \\ -\frac{1}{a(m-1) \sin^{m-1} ax \cos^{n-1} ax} + \frac{m+n-2}{m-1} I_{m-2,n} \end{cases}$

3.3.7. Special Definite and Improper Integrals

$$\int_{-\infty}^{\infty} \exp(-ax^2) \, dx = \sqrt{\frac{\pi}{a}}$$

$$\int_{-\infty}^{\infty} x^{2n} \exp(-ax^2) \, dx = \frac{(2n-1)!!}{(2a)^n} \sqrt{\frac{\pi}{a}}$$

$$\int_0^{\infty} x^n \exp(-ax) \, dx = \frac{n!}{a^{n+1}}$$

$$\int_{-\infty}^{\infty} \frac{\sin ax}{x} \, dx = \pi$$

$$\int_0^{\pi/2} \sin^m x \cos^n x \, dx = \frac{(m-1)!!(n-1)!!}{(m+n)!!} \times C, \quad C = \begin{cases} \pi/2 & \text{if } m, n \text{ both even} \\ 1 & \text{otherwise} \end{cases}$$

$$\int_{-\infty}^{\infty} \frac{dx}{a^n + x^n} = \frac{2\pi}{na^{n-1}} \csc \frac{\pi}{n} \quad (\text{for even } n)$$

$k!!$ is the double factorial, see Section 1.7.7.

3.3.8. Differentiation Rules: Product Rule, Quotient Rule, Chain Rule

If u, v, w, \dots are functions then:

- Product rule: $(uv)' = uv' + u'v$ $(uvw)' = uvw' + uv'w + u'vw$

- Leibniz rule for repeated differentiation of a product:

$$(uv)^{(n)} = u^{(n)}v + nu^{(n-1)}v' + \dots + {}^nC_p u^{(n-p)}v^{(p)} + \dots + uv^{(n)} = \sum_{p=0}^n {}^nC_p u^{(n-p)}v^{(p)}$$

- Quotient rule: $\left(\frac{u}{v}\right)' = \frac{u'v - v'u}{v^2}$

- Chain rule and Implicit differentiation: if $z = u(v(x))$ then

$$z' = (u \circ v)' = u(v)' = u'(v) v' \quad \frac{dz}{dx} = \frac{dz}{dv} \times \frac{dv}{dx}$$

3.3.9. Integration Rules: Integration by Parts, Integration by Substitution

Integration by parts: $\int_a^b u \frac{dv}{dx} dx = [uv]_a^b - \int_a^b v \frac{du}{dx} dx$

Integration by substitution: $\int_a^b u(v) \frac{dv}{dx} dx = \int_{v(a)}^{v(b)} u dv$

Leibniz Integral Rule for differentiation under the integral sign (Feynman's Technique):

$$\frac{d}{dx} \int_{a(x)}^{b(x)} f(x, y) dy = f(x, b) \frac{db}{dx} - f(x, a) \frac{da}{dx} + \int_{a(x)}^{b(x)} \frac{\partial f(x, y)}{\partial x} dy$$

3.3.10. Dirac Delta Functions (Impulse Function) and the Sifting Theorem

The delta function $\delta(x)$ is zero for all $x \neq 0$ and 'spikes' to $+\infty$ at $x = 0$, such that $\int_{-\infty}^{\infty} \delta(x) dx = 1$.

Integral of delta function: $\int_{-\infty}^x \delta(x - a) dx = H(x - a)$ (H : Heaviside unit step function)

Sifting property: $\int_{-\infty}^{\infty} f(x) \delta(x - a) dx = f(a)$ (Convolution: $f(x) * \delta(x - a) = f(x - a)$.)

For unilateral convolutions (integrating from 0 to x), the RHS is multiplied by $H(x)$.

3.3.11. Standard Substitutions for Integration

Integrals of radicals should use trigonometric (trig sub) or hyperbolic substitutions:

- $(a^2 - x^2)$ or $\sqrt{a^2 - x^2}$ → let $x = a \sin \theta$ or $x = a \cos \theta$
- $(a^2 + x^2)$ or $\sqrt{a^2 + x^2}$ → let $x = a \tan \theta$ or $x = a \sinh \theta$
- $(x^2 - a^2)$ or $\sqrt{x^2 - a^2}$ → let $x = a \sec \theta$ or $x = a \cosh \theta$

Integrals of rational functions of x and radicals should use:

- $\frac{1}{(ax + b)\sqrt{px + q}}$ → let $u^2 = px + q$
- $\frac{1}{(ax + b)\sqrt{px^2 + qx + r}}$ → let $\frac{1}{u} = ax + b$

More complicated rational functions of x and $\sqrt{px^2 + qx + r}$ should use Euler substitutions:

- if $p > 0$: → let $\sqrt{px^2 + qx + r} = u \pm x\sqrt{p}$ → $x = \frac{r - u^2}{\pm 2u\sqrt{p} - q}$
- if $r > 0$: → let $\sqrt{px^2 + qx + r} = xu \pm \sqrt{r}$ → $x = \frac{\pm 2u\sqrt{r} - q}{p - u^2}$
- if $q^2 - 4pr > 0$: → let $\sqrt{px^2 + qx + r} = \sqrt{p(x - \alpha)(x - \beta)} = (x - \alpha)u$ → $x = \frac{p\beta - \alpha u^2}{p - u^2}$

Integrals of rational functions of $(\sin x$ and/or $\cos x)$ or $(\sinh x$ and/or $\cosh x)$ should use the Weierstrass substitution (tangent half-angle substitution):

- $\frac{P(\sin x, \cos x)}{Q(\sin x, \cos x)}$ → let $t = \tan \frac{x}{2}$ → $\sin x = \frac{2t}{1 + t^2}$, $\cos x = \frac{1 - t^2}{1 + t^2}$, $dx = \frac{2 dt}{1 + t^2}$
- $\frac{P(\sinh x, \cosh x)}{Q(\sinh x, \cosh x)}$ → let $t = \tanh \frac{x}{2}$ → $\sinh x = \frac{2t}{1 - t^2}$, $\cosh x = \frac{1 + t^2}{1 - t^2}$, $dx = \frac{2 dt}{1 - t^2}$

3.3.12. Identities for Definite Integrals

Simple identities:
$$\int_a^b f(x) dx = - \int_b^a f(x) dx, \quad \int_a^b f(x) dx = \int_a^b f(y) dy$$

Reflections (King's rules):
$$\int_a^b f(x) dx = \int_a^b f(a+b-x) dx, \quad \int_a^b \frac{f(x)}{f(x)+f(a+b-x)} dx = \frac{b-a}{2}$$

Periodic function, T :
$$\int_a^{a+nT} f(x) dx = n \int_b^{b+T} f(x) dx \text{ for any } a, b \text{ and integers } n$$

Parity:
$$\text{odd: } \int_{-a}^a f(x) dx = 0, \quad \text{even: } \int_{-a}^a f(x) dx = 2 \int_0^a f(x) dx$$

Absolute values:
$$\int_a^b f(x) dx \leq \left| \int_a^b f(x) dx \right| \leq \int_a^b |f(x)| dx$$

Cauchy-Schwarz inequality:
$$\left| \int_a^b f(x) g(x) dx \right|^2 \leq \left(\int_a^b f(x)^2 dx \right) \left(\int_a^b g(x)^2 dx \right)$$

Monotonically increasing:
$$\int_a^b f(x) dx + \int_{f(a)}^{f(b)} f^{-1}(x) dx = b f(b) - a f(a)$$

3.4. Ordinary Differential Equations, Laplace and Z-Transforms

3.4.1. Classification of Ordinary Differential Equations (ODEs)

Ordinary differential equation (ODE): an equation relating a dependent variable y and its derivatives (y' , y'' , etc) with respect to a single independent variable x .

(The dependent function is sometimes also written as $y(x)$ or $x(t)$, or any other variables.)

Linear ODE: $\sum_{r=0}^n a_r(x) y^{(r)}(x) = f(x)$ (n : order, a_r : coefficient functions)

Homogeneous ODE: linear ODE with $f(x) = 0$. (**nonhomogeneous:** $f(x) \neq 0$)

Autonomous ODEs have no explicit dependence on x e.g. $y \, dy/dx = 1 - y$.

Nonlinear ODEs may have functions of the derivatives or products of variables e.g. xy , y^2 , $\exp y'$.

The **degree** of a nonlinear ODE is the exponent on the highest-order derivative

e.g. $x(y'')^3 - (y')^4 = 1$ is a nonlinear second-order ordinary differential equation with degree 3.

3.4.2. Separable DEs (First Order, Nonlinear)

For a separated ODE of the form $f(y) \, dy = g(x) \, dx$, the solution can be found by integrating

both sides. Initial conditions $y(x_0) = y_0$ can be applied with $\int_{y_0}^y f(y) \, dy = \int_{x_0}^x g(x) \, dx$.

3.4.3. Linear DEs (First Order, Linear)

To solve an ODE of the form $\frac{dy}{dx} + P(x)y = Q(x)$, multiply both sides by the integrating factor $I(x) = \exp(\int P(x) \, dx)$ and use the product rule so that the solution is given by $I(x)y(x) = \int I(x)Q(x) \, dx$. The computation of $I(x)$ does not require an arbitrary constant $+C$, even if initial conditions are not given.

3.4.4. Homogeneous DEs (First Order, Nonlinear) and Common Substitutions

To solve an ODE of the form $\frac{dy}{dx} = f\left(\frac{y}{x}\right)$, substitute $u = \frac{y}{x}$ so that $\frac{dy}{dx} = u + x \frac{du}{dx}$, which yields $\frac{du}{f(u) - u} = \frac{dx}{x}$, which is a separable DE in $u(x)$.

To solve an ODE of the form $\frac{dy}{dx} = f(ax + by + c)$, substitute $u = ax + by + c$ so that $\frac{dy}{dx} = \frac{1}{b} \frac{du}{dx} - \frac{a}{b}$, which yields $\frac{du}{dx} = bf(u) + a$, which is a separable DE in $u(x)$.

To solve an ODE of the form $\frac{dy}{dx} = \frac{a_1x + b_1y + c_1}{a_2x + b_2y + c_2}$, substitute $x = u + h$ and $y = v + k$, where h and k are the constant solutions to the system $\{a_1h + b_1k + c_1 = 0, a_2h + b_2k + c_2 = 0\}$.

Then, $\frac{dy}{dx} = \frac{dv}{du}$ and the DE becomes $\frac{dv}{du} = -\frac{a_1 + b_1 \frac{v}{u}}{a_2 + b_2 \frac{v}{u}}$, which is a homogeneous DE in $\frac{v}{u}$.

3.4.5. Linear DEs with Constant Coefficients (Second Order, Linear)

To solve an ODE of the form $a \frac{d^2 y}{dx^2} + b \frac{dy}{dx} + c y = f(x)$: (or with higher order)

- Solve the characteristic equation, $a\lambda^2 + b\lambda + c = 0$. (or with higher order)
- Depending on the nature of the roots λ , find the complementary function, y_{CF} :

- If λ_1 and λ_2 are **real and distinct**,

$$y_{CF}(x) = A e^{\lambda_1 x} + B e^{\lambda_2 x} \quad \text{or} \quad y_{CF}(x) = C \cosh \lambda_1 x + D \sinh \lambda_2 x$$

- If $\lambda_1 = \lambda_2 = \lambda$ is the **real repeated root**,

$$y_{CF}(x) = (A + Bt)e^{\lambda t}$$

- If $\lambda_1 = \alpha + i\beta$ and $\lambda_2 = \alpha - i\beta$ are the **distinct complex conjugate roots**,

$$y_{CF}(x) = e^{\alpha x} (A \cos \beta x + B \sin \beta x) \quad \text{or} \quad y_{CF} = C e^{\alpha x} \sin(\beta x - D)$$

- Find the particular integral y_{PI} using one of the following methods:
(note that if $f(x) = 0$ (homogeneous) then $y_{PI}(x) = 0$.)

- **Method of Undetermined Coefficients:** choose a suitable trial function based on the form of $f(x)$ from the table in Section 3.4.5, substituting it into the differential equation and equating linearly independent terms to solve for the coefficients.

- **Variation of Parameters:** evaluate the Wronskian, where y_1 and y_2 are the basis functions of y_{CF} .

$$W(x) = \begin{vmatrix} y_1(x) & y_2(x) \\ y_1'(x) & y_2'(x) \end{vmatrix}$$

The particular integral is then

$$y_{PI}(x) = -y_1(x) \int \frac{y_2(x) f(x)}{W(x)} dx + y_2(x) \int \frac{y_1(x) f(x)}{W(x)} dx$$

- By superposition, the solution is $y(x) = y_{CF}(x) + y_{PI}(x)$,
- The remaining constants in the $y_{CF}(x)$ term can be found using initial/boundary conditions.

Alternative methods without solving the characteristic equation are:

- **Laplace transform** (Section 3.4.15.): take LT of both sides, rearrange for $Y(s)$, take ILT

- **Convolution:** if the impulse response $g(t)$ is known, then $y(t) = (f * g)(t)$.

Note that for an LTI system with $y(0) = 0$ the impulse response is the derivative of the step response i.e. let $f(x) = 1$ and differentiate the solution. This is the 1D Green's function approach.

3.4.6. Trial Functions for Nonhomogeneous Differential Equations

For linear differential equations with constant coefficients, where $f(x)$ is linearly independent of the complementary function:

$f(x)$	Trial function
1	C
x^n , for integer n	$Cx^n + Dx^{n-1} + \dots + C_0$
k^x	Ck^x
e^{kx}	Ce^{kx}
$x e^{kx}$	$(Cx + D)e^{kx}$
$x^n e^{kx}$	$(Cx^n + Dx^{n-1} + \dots + C_0)e^{kx}$
$\sin px$ or $\cos px$	$C \sin px + D \cos px$
$e^{kx} \sin px$ or $e^{kx} \cos px$	$(C \sin px + D \cos px)e^{kx}$
$x^n e^{kx} \sin px$ or $x^n e^{kx} \cos px$	$(Cx^n + Dx^{n-1} + \dots + C_0)(C_S \sin px + C_C \cos px)e^{kx}$

where C, D, \dots, C_C, C_S are undetermined coefficients.

If $f(x)$ has a component which is not linearly independent of $y_{CF}(x)$, then the corresponding component in the trial function must be multiplied by x , or by x^2 in the case where this is still not linearly independent (i.e. repeated real roots solution with the same form as $f(x)$.)

3.4.7. Cauchy-Euler DEs (Second Order, Linear)

To solve an ODE of the form $ax^2 \frac{d^2y}{dx^2} + bx \frac{dy}{dx} + cy = f(x)$, substitute $u = \ln x$ so that

$$\frac{dy}{dx} = \frac{1}{x} \frac{dy}{du} \quad \text{and} \quad \frac{d^2y}{dx^2} = \frac{1}{x^2} \left(\frac{d^2y}{du^2} - \frac{dy}{du} \right),$$

which yields a second-order differential equation with constant coefficients. The resulting RHS will be $f(e^u)$, for which a particular integral may often be found.

3.4.8. Bernoulli DEs (First Order, Nonlinear)

To solve an ODE of the form $\frac{dy}{dx} + P(x)y(x) = Q(x)[y(x)]^n$, substitute $u(x) = y^{1-n}(x)$ so that

$$\frac{dy}{dx} = \frac{1}{1-n} u^{-\frac{n}{1-n}} \frac{du}{dx}, \quad \text{which yields a linear ODE in } u(x): \frac{du}{dx} + (1-n)P(x)u = (1-n)Q(x).$$

3.4.9. Exact DEs (First Order, Nonlinear)

An ODE $M(x, y) dx + N(x, y) dy = 0 \leftrightarrow N(x, y) \frac{dy}{dx} + M(x, y) = 0$ is exact if there exists a 'potential function' $F(x, y)$ such that $\frac{\partial F}{\partial x} = M(x, y)$ and $\frac{\partial F}{\partial y} = N(x, y)$.

The condition for exactness is met if $\frac{\partial M}{\partial y} = \frac{\partial N}{\partial x}$ (by Clairhaut's theorem).

To solve, find $F(x, y)$ by integrating M and N , separating the components for each variable. Use unknown functions $f(x)$ and $g(y)$ for the arbitrary constants of integration, and solve to make the antiderivatives equal to each other. The solutions are the contour lines of $F(x, y)$, implicitly satisfying $F(x, y) = C$ for some arbitrary constant C .

Almost Exact DEs:

An ODE $M(x, y) dx + N(x, y) dy = 0$ that is **not** exact can sometimes be multiplied by an integrating factor $\mu(x, y)$ on both sides to make an ODE $\widehat{M}(x, y) dx + \widehat{N}(x, y) dy = 0$ that is exact i.e. $\frac{\partial \widehat{M}}{\partial y} = \frac{\partial \widehat{N}}{\partial x}$ (where $\widehat{M}(x, y) = \mu(x, y) M(x, y)$ and $\widehat{N}(x, y) = \mu(x, y) N(x, y)$).

Techniques for finding such an integrating factor, if it exists:

- If $\frac{1}{N(x, y)} \left(\frac{\partial M}{\partial y} - \frac{\partial N}{\partial x} \right)$ is a function of x only, then $\mu(x) = \exp \int \frac{1}{N(x, y)} \left(\frac{\partial M}{\partial y} - \frac{\partial N}{\partial x} \right) dx$.
- If $\frac{1}{M(x, y)} \left(\frac{\partial N}{\partial x} - \frac{\partial M}{\partial y} \right)$ is a function of y only, then $\mu(y) = \exp \int \frac{1}{M(x, y)} \left(\frac{\partial N}{\partial x} - \frac{\partial M}{\partial y} \right) dy$.

3.4.10. Power Series Solution of DEs (Taylor Series Expansions) (Any Order, Linear)

To find the power series expansion $y(x) = \sum_{n=0}^{\infty} a_n x^n$ of the solution to a DE of the form

$$\sum_{k=0}^N p_k(x) y^{(k)}(x) = f(x), \text{ valid in some neighbourhood around } x = 0:$$

Power Series Method:

- Let $y(x) = \sum_{n=0}^{\infty} a_n x^n \Rightarrow y'(x) = \sum_{n=0}^{\infty} n a_n x^{n-1} \Rightarrow y''(x) = \sum_{n=0}^{\infty} n(n-1) a_n x^{n-2} \dots$ in the DE.
- Write the power series for each $p_k(x)$ and $f(x)$ and absorb these powers into the $y^{(n)}$ series.
- Re-index the summations to make them all have the same exponents of x .
- Pull out the first few terms of summations to make them all start at the same index n .
- Combine the summations and factor out the x^n term.
- Set everything inside the summation to zero to yield a recurrence relation in a_n , and set the pulled out terms to zero.
- Use initial conditions e.g. $a_n = \frac{y^{(n)}(0)}{n!}$, or let $(a_0, a_1, \dots) \in \{(1, 0), (0, 1), \dots\}$ for a linearly independent set of basis solutions.

Leibniz-Maclaurin Method:

- Differentiate both sides of the differential equation with respect to x , n times, using the general Leibniz rule for differentiating products.
- Let $x = 0$, convert the derivatives to series coefficients i.e. $a_n = \frac{y^{(n)}(0)}{n!}$ to yield a recurrence relation in a_n .
- If there are undetermined coefficients, evaluate the original DE at $x = 0$ to find them.

Frobenius Method: used when any $p_k(x)$ and $f(x)$ is not infinitely differentiable at $x = 0$.

For the DE $y'' + p(x)y' + q(x)y = 0$:

- Solve the indicial equation, $r(r-1) + u_0r + v_0 = 0$, where u_0 and v_0 are the constant terms in the Taylor series expansion of $u(x) = xp(x)$ and $v(x) = x^2q(x)$ respectively, for r .
- Case 1: Distinct real roots where r_1 and r_2 do not differ by an integer:
 - Use power series method with $y = \sum_{n=0}^{\infty} a_n x^{n+r}$, find recurrence relation in terms of r
 - Sub in each root: $y = A \sum_{n=0}^{\infty} a_n x^{n+r_1} + B \sum_{n=0}^{\infty} b_n x^{n+r_2}$
- Case 2: repeated roots. $y = (A + B \ln x) \sum_{n=0}^{\infty} a_n x^{n+r} + B \sum_{n=1}^{\infty} b_n x^{n+r}$, $b_n = \frac{da_n}{dr}(0)$
- Case 3: roots that differ by an integer. $y = (A + B \ln x) \sum_{n=0}^{\infty} a_n x^{n+r_1} + B \sum_{n=0}^{\infty} b_n x^{n+r_1}$, $b_n = \frac{da_n}{dr}(0)$

Fuch's theorem: radius of convergence of Frobenius series, $R \geq \min\{R_{p(x)}, R_{q(x)}, R_{r(x)}\}$.

If $x = 0$ is the only 'regular singular point' ($u(x)$ and $v(x)$ infinitely differentiable at $x = 0$) then the Frobenius series converges everywhere. Otherwise, R is the distance to the nearest singular point.

3.4.11. Higher Order DEs as Systems of First Order DEs

An N th order ODE $\sum_{n=0}^N a_n(x) y^{(n)}(x) = f(x)$ can be written as a set of N first-order ODEs.

Label the N new dependent variables $\left\{ y \rightarrow y_0, \frac{dy}{dx} \rightarrow y_1, \dots, \frac{d^{N-1}y}{dx^{N-1}} \rightarrow y_{N-1} \right\}$.

The system is then $\left\{ \frac{dy_0}{dx} = y_1, \frac{dy_1}{dx} = y_2, \dots, \frac{dy_{N-2}}{dx} = y_{N-1}, \frac{dy_{N-1}}{dx} = f(x) - \sum_{n=0}^{N-1} a_n(x) y_n \right\}$.

Initial conditions correspond directly to the initial condition for each equation in the system.

The method still works for nonlinear higher order ODEs but the resulting system will be nonlinear.

ODEs in this form are readily solved by computer/numerical methods.

3.4.12. Systems of Differential Equations

A system of n first-order ODEs can be written in the form $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t)$ where \mathbf{x} is a vector of unknown functions: $\mathbf{x} = [x_1(t), x_2(t), \dots, x_n(t)]^T$.

An autonomous system is one in which $\mathbf{f}(\mathbf{x}, t) = \mathbf{f}(\mathbf{x})$ i.e. explicitly independent of t .

Linear homogeneous systems with constant coefficients:

- Standard form: $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ (\mathbf{A} : square $n \times n$ matrix of constant coefficients)
- Ansatz: $\mathbf{x} = \exp(\mathbf{A}t) \mathbf{x}_0$ (\mathbf{x}_0 : initial conditions at $t = 0$)
- General solution (for a 2×2 system): (\mathbf{u} : eigenvectors of \mathbf{A} , λ : eigenvalues of \mathbf{A})

$$\mathbf{x}(t) = \begin{cases} c_1 e^{\lambda_1 t} \mathbf{u}_1 + c_2 e^{\lambda_2 t} \mathbf{u}_2 & \text{if } \lambda_{1,2} \text{ are real} \\ c_1 e^{\alpha t} (\mathbf{u}_1 \cos \beta t + \mathbf{u}_2 \sin \beta t) + c_2 e^{\alpha t} (\mathbf{u}_1 \sin \beta t - \mathbf{u}_2 \cos \beta t) & \text{if } \lambda_{1,2} = \alpha \pm \beta i \text{ are complex} \\ c_1 e^{\lambda t} \mathbf{u} + c_2 e^{\lambda t} (\mathbf{u}t + \mathbf{v}), \text{ for any } \mathbf{v} : (\mathbf{A} - \lambda \mathbf{I})\mathbf{v} = \mathbf{u} & \text{if } \lambda \text{ is a repeated defective eigenvalue} \end{cases}$$

(\mathbf{u} : eigenvectors of \mathbf{A} , λ : eigenvalues of \mathbf{A})


Linear nonhomogeneous systems with constant coefficients:

- Standard form: $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{f}(t)$

Solution methods include the method of undetermined coefficients (using the complementary solution from the homogeneous case) or variation of parameters. The formula for variation of parameters is

$$\mathbf{x}_{PI}(t) = \mathbf{X} \int \mathbf{X}^{-1} \mathbf{f}(t) dt \quad (\mathbf{X}: \text{matrix where each column is a linearly independent part of the complementary solution})$$

Phase plane and equilibrium point stability: equilibrium point(s) occur when $d\mathbf{x}/dt = \mathbf{0}$.



Type	Eigenvalues	Stability
Node	Real λ , same signs	$\lambda < 0$, stable $\lambda > 0$, unstable
Saddle	Real λ opposite signs	Mostly Unstable
Center	λ pure imaginary	—
Focus/Spiral	Complex λ , $\text{Re}(\lambda) \neq 0$	$\text{Re}(\lambda) < 0$, stable $\text{Re}(\lambda) > 0$, unstable
Degenerate Node	Repeated roots,	$\lambda > 0$, stable
Lines of Equilibria	One zero eigenvalue	$\lambda < 0$, stable

- For linear homogeneous systems, the origin is the only equilibrium point.
- The eigenvectors of \mathbf{A} are directed along the asymptotic trajectories of the system (**nullclines**).
- An equilibrium point is the intersection of the x and y nullclines, for which $\frac{dx}{dt} = 0$ and $\frac{dy}{dt} = 0$.
- For nonlinear systems, the nullclines may be curved.


3.4.13. Graphical Representations of Differential Equations

Slope Field (Direction Field) of a Differential Equation

For a differential equation $\frac{dy}{dx} = f(x, y)$, the slope field is a unit vector field in the x - y plane given by $\mathbf{u}(x, y) = \frac{1}{\sqrt{1 + f(x, y)^2}} \mathbf{i} + \frac{f(x, y)}{\sqrt{1 + f(x, y)^2}} \mathbf{j}$.

Every solution to the differential equation is a **field line** of the vector field.

Curves with $\frac{dy}{dx} = f(x, y) = k$ for constant k are 'isoclines'.




Example: slope field for $\frac{dy}{dx} = x^2 - x - 2$.

The curves represent solutions for initial conditions $y(0) = 4$, $y(0) = 0$ and $y(0) = -4$.

Phase Plane of a System of Differential Equations

For a system of two ODEs $\left\{ \frac{dx}{dt} = f(x, y, t), \frac{dy}{dt} = g(x, y, t) \right\}$, the phase space plot is a vector field in the x - y plane given by $\mathbf{u}(x, y, t) = \frac{f(x, y, t)}{\sqrt{f(x, y, t)^2 + g(x, y, t)^2}} \mathbf{i} + \frac{g(x, y, t)}{\sqrt{f(x, y, t)^2 + g(x, y, t)^2}} \mathbf{j}$.

If the system is autonomous (no t dependence), this is a static vector field.



Example: phase portrait for $\left\{ \frac{dx}{dt} = x(7 - x - 2y), \frac{dy}{dt} = y(5 - y - x) \right\}$

Lines for $\frac{dx}{dt} = 0$ are shown in blue (x -nullclines: $x = 0, x + 2y = 7$)

Lines for $\frac{dy}{dt} = 0$ are shown in red (y -nullclines: $y = 0, x + y = 5$)

A fixed point occurs at the intersection of nullclines.

3.4.14. Special Forms of DEs (Second Order, Nonlinear)**Bessel's differential equation** (generalised):

$$x^2 \frac{d^2 y}{dx^2} + (2p+1)x \frac{dy}{dx} + (\lambda^2 x^{2q} + \alpha^2) y = 0 \Rightarrow y = x^{-p} \left(A \cdot J_{\frac{\sqrt{p^2 - \alpha^2}}{q}} \left(\frac{\lambda}{q} x^q \right) + B \cdot Y_{\frac{\sqrt{p^2 - \alpha^2}}{q}} \left(\frac{\lambda}{q} x^q \right) \right)$$

Spherical Bessel's differential equation (a particular case of the above):

$$x^2 \frac{d^2 y}{dx^2} + 2x \frac{dy}{dx} + (\lambda^2 x^2 - n(n+1)) y = 0$$

$$\Rightarrow y = A \cdot \frac{J_{n+\frac{1}{2}}(\lambda x)}{\sqrt{\lambda x}} + B \cdot \frac{Y_{n+\frac{1}{2}}(\lambda x)}{\sqrt{\lambda x}} = A' \cdot j_n(\lambda x) + B' \cdot y_n(\lambda x)$$

In the case $n = 0$, the solution is $y = A' \frac{\sin \lambda x}{\lambda x} - B' \frac{\cos \lambda x}{\lambda x}$, for which $B' = 0$ if $y(0)$ is finite.**Generalised Laguerre differential equation:**

$$x \frac{d^2 y}{dx^2} + (\alpha + 1 - x) \frac{dy}{dx} + ny = 0 \Rightarrow y = A \cdot L_n^{(\alpha)}(x) + B \cdot U(-n, \alpha + 1, x)$$

Hypergeometric differential equation:

$$x(1-x) \frac{d^2 y}{dx^2} + (c - (a+b+1)x) \frac{dy}{dx} - aby = 0$$

$$\Rightarrow y = A \cdot {}_2F_1(a, b; c; x) + B \cdot (-x)^{1-c} {}_2F_1(a-c+1, b-c+1; 2-c; x)$$

Confluent Hypergeometric differential equation:

$$x \frac{d^2 y}{dx^2} + (c-x) \frac{dy}{dx} - ay = 0 \Rightarrow y = A \cdot {}_1F_1(a; c; x) + B \cdot U(a, c, x)$$

Hermite's differential equation:

$$\frac{d}{dx} \left(e^{-\frac{1}{2}x^2} \frac{dy}{dx} \right) + \lambda e^{-\frac{1}{2}x^2} y = 0 \Rightarrow y = A \cdot H_\lambda \left(\frac{x}{\sqrt{2}} \right) + B \cdot {}_1F_1 \left(-\frac{\lambda}{2}; \frac{1}{2}; \frac{x^2}{2} \right)$$

(For the special function definitions, see Section 1.7.)

3.4.15. Laplace Transforms

$$F(s) = \mathcal{L}\{f(t)\} = \int_0^{\infty} f(t) e^{-st} dt$$

Derivatives, Integrals, Deltas and Algebraic Functions:

$f(t)$	$F(s)$
$e^{-at}x(t)$	$\bar{x}(s+a)$
$x(t-\tau)H(t-\tau)$	$e^{-s\tau}\bar{x}(s)$
$\frac{dx(t)}{dt} = x'(t)$	$s\bar{x}(s) - x(0)$
$\frac{d^2x(t)}{dt^2} = x''(t)$	$s^2\bar{x}(s) - sx(0) - x'(0)$
$\frac{d^nx(t)}{dt^n} = x^{(n)}(t)$	$s^n\bar{x}(s) - s^{n-1}x(0) - s^{n-2}x'(0) - \dots - sx^{(n-2)}(0) - x^{(n-1)}(0)$
$\int_0^t x(\tau) d\tau$	$s^{-1}\bar{x}(s)$
$\int_0^t x_1(\tau)x_2(t-\tau) d\tau$	$\bar{x}_1(s)\bar{x}_2(s)$
$tx(t)$	$-\frac{d}{ds}\bar{x}(s)$
$1 = H(t)$	s^{-1}
$\delta(t)$	1
$H(t-\tau)$	$s^{-1}e^{-s\tau}$
$\delta(t-\tau)$	$e^{-s\tau}$

Powers, Exponential, Trigonometric and Hyperbolic:

$f(t)$	$F(s)$	$f(t)$	$F(s)$
t	s^{-2}	t^n	$n!s^{-n-1}$
e^{-at}	$(s+a)^{-1}$	$t^n e^{-at}$	$\frac{n!}{(s+a)^{n+1}}$
$\sin \omega t$	$\frac{\omega}{s^2 + \omega^2}$	$\cos \omega t$	$\frac{s}{s^2 + \omega^2}$
$e^{-at} \sin \omega t$	$\frac{\omega}{(s+a)^2 + \omega^2}$	$e^{-at} \cos \omega t$	$\frac{(s+a)}{(s+a)^2 + \omega^2}$
$t \sin \omega t$	$\frac{2s\omega}{(s^2 + \omega^2)^2}$	$t \cos \omega t$	$\frac{s^2 - \omega^2}{(s^2 + \omega^2)^2}$
$\sinh \omega t$	$\frac{\omega}{s^2 - \omega^2}$	$\cosh \omega t$	$\frac{s}{s^2 - \omega^2}$

Initial Value / Final Value Theorem: $f(0^+) = \lim_{s \rightarrow \infty} s F(s)$ and $\lim_{t \rightarrow \infty} f(t) = \lim_{s \rightarrow 0} s F(s)$.

3.4.16. Convolution Theorem

$$(f * g)(t) = \int_0^t f(t - \tau)g(\tau) d\tau$$

The Laplace transform of a convolution is the product of their transforms:

$$\mathcal{L}\{(f * g)(t)\} = F(s) \cdot G(s) \quad \text{equivalently} \quad \mathcal{L}^{-1}\{F(s)G(s)\} = (f * g)(t)$$

The convolution theorem also applies to Fourier transforms (Section 3.6.5).

3.4.17. Inverse Laplace Transform by the Cauchy Residue Theorem

The inverse Laplace transform is defined as (Fourier-Mellin formula)

$$f(t) = \mathcal{L}^{-1}\{F(s)\} = \frac{1}{2\pi i} \int_{\gamma-i\infty}^{\gamma+i\infty} F(s) e^{st} ds$$

where γ is a constant larger than the real part of any pole of $F(s)$.

If $\gamma = 0$ (i.e. no unstable poles: $\text{Re}(s_k) < 0$) then this is similar to the inverse Fourier transform.

Using the Residue Theorem and Jordan's Lemma (using a semicircular contour), this is equivalent to (by complex analysis):

$$f(t) = 2\pi i \sum_k \text{Res}[F(s), s_k]$$

with the sum over all residues at the poles of $F(s)$. The residue is defined as

$$\text{Res}[F(s), s_k] = \frac{1}{(n-1)!} \lim_{s \rightarrow s_k} \frac{d^{n-1}}{ds^{n-1}} (s - s_k)^n F(s)$$

where n is the multiplicity of pole s_k .

If $n = 1$ then the residue simplifies to $\text{Res}(F, s_k) = \lim_{s \rightarrow s_k} (s - s_k)F(s)$.

which is the formalised 'cover-up method' of partial fractions if F is rational.

3.4.18. Linear Difference Equations

Difference Equations are discretised differential equations, expressed as a recurrence relation between terms of a sequence $\{y\}_n$ for $n = 0, 1, 2, \dots$

To solve a second-order difference equation of the form (or higher order)

$$a y_n + b y_{n-1} + c y_{n-2} = f(n),$$

- Solve the characteristic equation, $a\lambda^2 + b\lambda + c = 0$. (or higher order)
- Depending on the nature of the roots λ , find the complementary function, $y_n^{(CF)}$:

- If λ_1 and λ_2 are **real and distinct**,

$$y_n^{(CF)} = A \lambda_1^n + B \lambda_2^n$$

- If $\lambda_1 = \lambda_2 = \lambda$ is the **real repeated root**,

$$y_n^{(CF)} = A \lambda_1^n + B n \lambda_2^n$$

- If $\lambda_1 = R \exp i\theta$ and $\lambda_2 = R \exp -i\theta$ are the **distinct complex conjugate roots**,

$$y_n^{(CF)} = R^n (A \cos n\theta + B \sin n\theta)$$

- Use the Method of Undetermined Coefficients to determine the particular 'integral', $y_n^{(PI)}$ (note that if $f(n) = 0$ then $y_n^{(PI)} = 0$). The trial functions are identical to the case of a nonhomogeneous differential equation (Section 3.4.5.), with x replaced by n .
- By superposition, the solution is $y_n = y_n^{(CF)} + y_n^{(PI)}$ where the remaining constants can be found using given conditions.

Alternative methods without solving the characteristic equation are:

- **Z-transform / generating function** (Section 3.4.19): $y_n = Z^{-1}(Y(z))$, where $Y(z)$ is the generating function with $x = z^{-1}$ given by $Y(z) = \sum_n y_n z^n$ (the Z transform). (The generating function is $Y(x) = \sum_n y_n x^n$.)
- **Convolution**: if the impulse response g_n is known, then $y_n = (f * g)[n]$. For the definition of the discrete convolution, see Section 5.4.7.

3.4.19. Z-Transforms, Inverse Z Transforms and Generating Functions

Sequence:

$$y_n = \mathcal{Z}^{-1}[Y(z)][n], \quad n = 0, 1, 2, \dots$$

Z-transform:

$$Y(z) = \mathcal{Z}[y_n](z) = \sum_{n=0}^{\infty} y_n z^{-n}$$

Generating function:

$$Y(x) = \mathcal{Z}[y_n](x^{-1}) = \sum_{n=0}^{\infty} y_n x^n$$

1 (unit step)	$\frac{1}{1 - z^{-1}}$	$\frac{1}{1 - x}$
nT	$\frac{Tz^{-1}}{(1 - z^{-1})^2}$	$\frac{Tx}{(1 - x)^2}$
$\frac{(n + m - 1)!}{n!(m - 1)!} = {}^{n+m-1}C_n$	$\frac{1}{(1 - z^{-1})^m}$	$\frac{1}{(1 - x)^m}$
e^{-anT}	$\frac{1}{1 - e^{-aT}z^{-1}}$	$\frac{1}{1 - e^{-aT}x}$
$\sin(\omega_0 nT)$	$\frac{\sin(\omega_0 T)z^{-1}}{1 - 2 \cos(\omega_0 T)z^{-1} + z^{-2}}$	$\frac{\sin(\omega_0 T)x}{1 - 2 \cos(\omega_0 T)x + x^2}$
$\cos(\omega_0 nT)$	$\frac{1 - \cos(\omega_0 T)z^{-1}}{1 - 2 \cos(\omega_0 T)z^{-1} + z^{-2}}$	$\frac{1 - \cos(\omega_0 T)x}{1 - 2 \cos(\omega_0 T)x + x^2}$
$\frac{r^{n-1}}{\sin(\omega_0 T)} [r \sin(\omega_0(n + 1)T) - a \sin(\omega_0 nT)]$	$\frac{1 - az^{-1}}{1 - 2r \cos(\omega_0 T)z^{-1} + r^2 z^{-2}}$	$\frac{1 - ax}{1 - 2r \cos(\omega_0 T)x + r^2 x^2}$
$r^n [A \cos(\omega_0 nT) + B \sin(\omega_0 nT)]$	$\frac{A + rz^{-1}(B \sin(\omega_0 T) - A \cos(\omega_0 T))}{1 - 2r \cos(\omega_0 T)z^{-1} + r^2 z^{-2}}$	$\frac{A + rx(B \sin(\omega_0 T) - A \cos(\omega_0 T))}{1 - 2r \cos(\omega_0 T)x + r^2 x^2}$
$r^n y_n$	$Y(r^{-1}z)$	$Y(r^{-1}x^{-1})$
y_{n+1}	$zY(z) - zy_0$	$x^{-1}Y(x^{-1}) - x^{-1}y_0$
y_{n-1}	$z^{-1}Y(z) + y_{-1}$	$xY(x^{-1}) + y_{-1}$
y_{n+m}	$z^m G(z) - (z^m y_0 + \dots + zy_{m-1})$	$x^{-m}G(x^{-1}) - (x^{-m}y_0 + \dots + x^{-1}y_{m-1})$
y_{n-m}	$z^{-m}G(z) + (z^{-(m-1)}y_{-1} + \dots + y_{-m})$	$x^m G(x^{-1}) + (x^{m-1}y_{-1} + \dots + y_{-m})$

Initial Value / Final Value Theorem: $y_0 = \lim_{z \rightarrow \infty} Y(z)$ and $\lim_{n \rightarrow \infty} y_n = \lim_{z \rightarrow 1} (z - 1) Y(z)$

Note: the final value theorem requires the poles of $(z - 1) Y(z)$ to have $|z| < 1$.

The Laplace-analogous residue formula for the Inverse Z-Transform is

$$g_n = \mathcal{Z}^{-1}\{G(z)\} = \frac{1}{2\pi i} \oint_C G(z) z^{n-1} dz = \sum_k \text{Res} [G(z)z^{n-1}, z_k]$$

where the residue Res is defined in Section 3.4.10. Note the extra factor of z^{n-1} .

3.5. Multivariable and Vector Calculus

3.5.1. Differentiation of Vector Products

For vector-valued functions $\mathbf{a}(t)$, $\mathbf{b}(t)$, and scalar-valued functions $u(t)$ of a single variable,

$$(\mathbf{a} \cdot \mathbf{b})' = \mathbf{a}' \cdot \mathbf{b} + \mathbf{a} \cdot \mathbf{b}' \qquad (\mathbf{a} \times \mathbf{b})' = \mathbf{a}' \times \mathbf{b} + \mathbf{a} \times \mathbf{b}' \qquad (u\mathbf{a})' = u'\mathbf{a} + u\mathbf{a}'$$

3.5.2. Jacobian Matrix

For a **vector valued** function \mathbf{f} of n variables x_1, \dots, x_n , the Jacobian \mathbf{J} is $J_{ij} = \frac{\partial f_i}{\partial x_j}$:

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \mathbf{f}}{\partial x_1} & \cdots & \frac{\partial \mathbf{f}}{\partial x_n} \end{bmatrix} = \begin{bmatrix} \nabla^T f_1 \\ \vdots \\ \nabla^T f_m \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}$$

3.5.3. Hessian Matrix

For a **scalar valued** function f of n variables x_1, \dots, x_n , the Hessian \mathbf{H} is $H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$:

$$\mathbf{H}_f = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix},$$

The Hessian for a vector-valued function $\mathbf{f} = [f_1, f_2, \dots, f_n]$ is the third-rank Hessian tensor whose elements are $[\mathbf{H}_{f_1}, \mathbf{H}_{f_2}, \dots, \mathbf{H}_{f_n}]$, where \mathbf{H}_{f_i} is the Hessian matrix of f_i .

3.2.3. Multivariable Taylor Series

For a **scalar-valued** function $f(\mathbf{x})$ about \mathbf{x}_0 ,

$$f(\mathbf{x}_0 + \mathbf{h}) = f(\mathbf{x}_0) + (\mathbf{h} \cdot \nabla) f(\mathbf{x}_0) + \frac{1}{2!} (\mathbf{h} \cdot \nabla)(\mathbf{h} \cdot \nabla) f(\mathbf{x}_0) + \frac{1}{3!} (\mathbf{h} \cdot \nabla)(\mathbf{h} \cdot \nabla)(\mathbf{h} \cdot \nabla) f(\mathbf{x}_0) + \dots$$

In the case of a two-variable scalar function f about (x_0, y_0) ,

$$f(x_0 + h, y_0 + k) = \sum_{n=0}^{\infty} \left[\frac{1}{n!} \sum_{i=0}^n \binom{n}{i} h^i k^{n-i} \frac{\partial^n f}{\partial x^i \partial y^{n-i}} \right]$$

$$f(x_0 + h, y_0 + k) = f(x_0, y_0) + \left[h \frac{\partial f}{\partial x} + k \frac{\partial f}{\partial y} \right] + \frac{1}{2!} \left[h^2 \frac{\partial^2 f}{\partial x^2} + 2hk \frac{\partial^2 f}{\partial x \partial y} + k^2 \frac{\partial^2 f}{\partial y^2} \right] + \dots$$

The quadratic approximation is

$$f(\mathbf{x}_0 + \mathbf{h}) \approx f(\mathbf{x}_0) + \mathbf{h}^T \nabla f + \frac{1}{2!} \mathbf{h}^T \mathbf{H}(\mathbf{x}_0) \mathbf{h} + \dots$$

where $\mathbf{H}(\mathbf{x}_0)$ is the Hessian matrix of f (Section 3.5.3) at \mathbf{x}_0 . The linear term $\mathbf{h}^T \nabla f$ is the directional derivative of f in the direction of \mathbf{h} , also written as $D_{\mathbf{h}} f(\mathbf{x}_0) = \nabla f \cdot \mathbf{h}$.

For a **vector-valued** function $\mathbf{f}(\mathbf{x})$ about \mathbf{x}_0 ,

The quadratic approximation is

$$\mathbf{f}(\mathbf{x}_0 + \mathbf{h}) \approx \mathbf{f}(\mathbf{x}_0) + \mathbf{J}(\mathbf{x}_0) \mathbf{h} + \frac{1}{2!} \mathbf{h}^T \mathbf{H}(\mathbf{x}_0) \mathbf{h} + \dots$$

where $\mathbf{H}(\mathbf{x}_0)$ is the Hessian tensor of \mathbf{f} (Section 3.5.3) at \mathbf{x}_0 , and $\mathbf{J}(\mathbf{x}_0)$ is the Jacobian matrix of \mathbf{f} at \mathbf{x}_0 . In Einstein summation notation (Section 4.4.1), this quadratic approximation is

$$f_i(\mathbf{x}_0 + \mathbf{h}) \approx f_i(\mathbf{x}_0) + J_{ij}(\mathbf{x}_0) h_j + \frac{1}{2!} H_{ijk}(\mathbf{x}_0) h_j h_k + \dots$$

3.5.4. Stationary Points of a Scalar-Valued Multivariable Function

A function $\phi(x_1, \dots, x_n)$ has a **stationary point** when $\nabla\phi = \mathbf{0}$ i.e. $\frac{\partial\phi}{\partial x_1} = \frac{\partial\phi}{\partial x_2} = \dots = \frac{\partial\phi}{\partial x_n} = 0$.

If the determinant of the Hessian matrix $\Delta = |\mathbf{H}| \neq 0$ at a stationary point, then

- **Minimum point:** $\Delta > 0$ and all $\frac{\partial^2\phi}{\partial x_i^2} > 0$. (ϕ is locally **convex**)
- **Maximum point:** $\Delta < 0$ and all $\frac{\partial^2\phi}{\partial x_i^2} < 0$. (ϕ is locally **concave**)
- **Saddle point:** all other cases for which $\Delta \neq 0$.

The case $\Delta = 0$ can be a maximum, a minimum, a saddle point, or none of these.

For two variables, $\phi(x, y)$, $\Delta = \frac{\partial^2\phi}{\partial x^2} \frac{\partial^2\phi}{\partial y^2} - \left(\frac{\partial^2\phi}{\partial x \partial y}\right)^2$.

Second partial derivatives are symmetric (Clairhaut's theorem): $\frac{\partial^2\phi}{\partial x \partial y} = \frac{\partial^2\phi}{\partial y \partial x}$.

3.5.5. Total Differentials

For a function $\phi(x, y, z, \dots)$, $d\phi = \frac{d\phi}{dx} dx + \frac{d\phi}{dy} dy + \frac{d\phi}{dz} dz + \dots$

If $f(x, y) dx + g(x, y) dy = d\phi$, then $\frac{\partial f}{\partial y} = \frac{\partial g}{\partial x}$ (an exact differential).

3.5.6. Multivariable Chain Rule


If x, y, z are functions of u, v, w, \dots

$$\left(\frac{\partial\phi}{\partial u}\right)_{v,w,\dots} = \frac{\partial\phi}{\partial x} \left(\frac{\partial x}{\partial u}\right)_{v,w,\dots} + \frac{\partial\phi}{\partial y} \left(\frac{\partial y}{\partial u}\right)_{v,w,\dots} + \frac{\partial\phi}{\partial z} \left(\frac{\partial z}{\partial u}\right)_{v,w,\dots} + \dots$$

3.5.7. Derivatives on Curved Lines and Curved Surfaces

For a **curve** defined parametrically as $\mathbf{r}(t) = [x(t) \ y(t) \ z(t)]^T$, the unit tangent vector \mathbf{T} , unit normal vector \mathbf{N} and unit binormal vector \mathbf{B} are given by

$$\hat{\mathbf{T}}(t) = \frac{\dot{\mathbf{r}}}{|\dot{\mathbf{r}}|} = \frac{\frac{d\mathbf{r}}{dt}}{\left|\frac{d\mathbf{r}}{dt}\right|} \quad \hat{\mathbf{N}}(t) = \frac{\frac{d\hat{\mathbf{T}}}{dt}}{\left|\frac{d\hat{\mathbf{T}}}{dt}\right|} \quad \hat{\mathbf{B}}(t) = \hat{\mathbf{T}}(t) \times \hat{\mathbf{N}}(t)$$



so that $\{\mathbf{T}, \mathbf{N}, \mathbf{B}\}$ forms a right-handed orthonormal set.

The equation of the tangent line at $\mathbf{r} = \mathbf{r}_0$ is then $(\mathbf{r} - \mathbf{r}_0) \times \mathbf{r}' = \mathbf{0}$.

The vectors \mathbf{T} , \mathbf{N} and \mathbf{B} vary with arc length s along the curve by (Frenet-Serret formulas):

$$\frac{d}{ds} \begin{bmatrix} \hat{\mathbf{T}} \\ \hat{\mathbf{N}} \\ \hat{\mathbf{B}} \end{bmatrix} = \begin{bmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{T}} \\ \hat{\mathbf{N}} \\ \hat{\mathbf{B}} \end{bmatrix} \quad (\kappa: \text{curvature}, \quad \tau: \text{torsion})$$

$$\kappa = \frac{|\dot{\mathbf{r}} \times \ddot{\mathbf{r}}|}{|\dot{\mathbf{r}}|^3} \quad \text{and} \quad \tau = \frac{(\dot{\mathbf{r}} \times \ddot{\mathbf{r}}) \cdot \ddot{\mathbf{r}}}{|\dot{\mathbf{r}} \times \ddot{\mathbf{r}}|^2}$$

The associated radius of curvature is $R = \kappa^{-1}$ and the ‘osculating circle’ lies in the plane spanned by \mathbf{T} and \mathbf{N} , with \mathbf{B} as its normal.

For a **surface** defined implicitly as $\phi(x, y, z) = 0$, the unit tangent vector \mathbf{T} (defined as being the projection of some vector \mathbf{u} in (x, y) -space onto the surface) and unit normal vector \mathbf{N} are

$$\hat{\mathbf{T}}(x, y, z) = \frac{(D_{\hat{\mathbf{k}}}\phi)\hat{\mathbf{u}} - (D_{\hat{\mathbf{u}}}\phi)\hat{\mathbf{k}}}{|\nabla\phi|} \quad \hat{\mathbf{N}}(x, y, z) = \frac{\nabla\phi}{|\nabla\phi|}$$

where \mathbf{k} is the unit vector in the z -direction and $D_{\mathbf{a}}f$ is the **directional derivative**, defined as $D_{\mathbf{a}}f = \nabla f \cdot \mathbf{a}$, representing the component of the gradient parallel to \mathbf{a} .

The equation of the tangent plane at $\mathbf{r} = \mathbf{r}_0$ is then $(\mathbf{r} - \mathbf{r}_0) \cdot \nabla\phi(\mathbf{r}_0) = 0$.

If \mathbf{N} is evaluated at a vector \mathbf{r}_0 which does not lie on the surface, then \mathbf{N} can instead be interpreted as the direction of steepest ascent for ϕ at $\mathbf{r} = \mathbf{r}_0$ (since $\phi \neq 0$ off the surface).

For a scalar-valued function $\phi(\mathbf{r})$, the regions of constant ϕ are called **isosurfaces** (contour surfaces; level surfaces) in 3D or **isolines** (contour lines) in 2D.

These results are easily generalisable to other dimensional functions, except the binormal vector which is only uniquely defined in \mathbb{R}^3 .

3.5.8. Reduction of a Multiple Integral with Common Bounds to a Single Integral

Double integral to single integral:
$$\int_a^x \int_a^u f(t) dt du = \int_a^x (x - t) f(t) dt$$

Triple integral to single integral:
$$\int_a^x \int_a^u \int_a^v f(t) dt dv du = \frac{1}{2} \int_a^x (x - t)^2 f(t) dt$$

These can be useful for simplifying numerical integration of multiple integrals.

3.5.9. Change of Variables for Multiple Integration

Surface Integrals:

For a change of variables in a surface integral from $(x, y) \rightarrow (u(x, y), v(x, y))$,

$$\iint_S f(x, y) dx dy = \iint_{S'} F(u, v) |J| du dv \quad J = \frac{\partial(x, y)}{\partial(u, v)} = \begin{vmatrix} \frac{\partial x}{\partial u} & \frac{\partial x}{\partial v} \\ \frac{\partial y}{\partial u} & \frac{\partial y}{\partial v} \end{vmatrix}$$

For surface integrals involving vector normals, where the sign is chosen to preserve the sense.

$$\mathbf{n} dA = \mathbf{n} dx dy = \pm \frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v} du dv$$

Volume integrals:


For a change of variables in a volume integral from $(x, y, z) \rightarrow (u(x, y, z), v(x, y, z), w(x, y, z))$,

$$\iiint_V f(x, y, z) dx dy dz = \iiint_{V'} F(u, v, w) |J| du dv dw \quad J = \frac{\partial(x, y, z)}{\partial(u, v, w)} = \begin{vmatrix} \frac{\partial x}{\partial u} & \frac{\partial x}{\partial v} & \frac{\partial x}{\partial w} \\ \frac{\partial y}{\partial u} & \frac{\partial y}{\partial v} & \frac{\partial y}{\partial w} \\ \frac{\partial z}{\partial u} & \frac{\partial z}{\partial v} & \frac{\partial z}{\partial w} \end{vmatrix}$$

The inverse Jacobian determinant is the same as that of the inverse substitution:

$$\frac{1}{J} = \frac{\partial(u, v, \dots)}{\partial(x, y, \dots)}$$

3.5.10. Vector Calculus in Cartesian Coordinates



Parameters: (x, y, z) , all real ordinates

Scalar field: $r = x\hat{i} + y\hat{j} + z\hat{k}$

Vector field: $u = u_x\hat{i} + u_y\hat{j} + u_z\hat{k}$

Kinematics:

Position: $r = x\hat{i} + y\hat{j} + z\hat{k}$

Velocity: $\dot{r} = \dot{x}\hat{i} + \dot{y}\hat{j} + \dot{z}\hat{k}$

Acceleration: $\ddot{r} = \ddot{x}\hat{i} + \ddot{y}\hat{j} + \ddot{z}\hat{k}$

Angular velocity: $\omega = 0$ (fixed)

Unit vectors: $\dot{\hat{i}} = 0, \dot{\hat{j}} = 0, \dot{\hat{k}} = 0$ (fixed)

Differential Elements

Line element: $dr = dx\hat{i} + dy\hat{j} + dz\hat{k}$

Volume element: $dV = dx dy dz$

Surface elements: $dS_x = dy dz, dS_y = dx dz, dS_z = dx dy$

Distance:

$$d = \sqrt{x^2 + y^2 + z^2}$$

Vector Operators:

Gradient: $\nabla f = \frac{\partial f}{\partial x}\hat{i} + \frac{\partial f}{\partial y}\hat{j} + \frac{\partial f}{\partial z}\hat{k}$

Divergence: $\nabla \cdot u = \frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} + \frac{\partial u_z}{\partial z}$

Curl: $\nabla \times u = \begin{vmatrix} \hat{i} & \hat{j} & \hat{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ u_x & u_y & u_z \end{vmatrix}$

Laplacian: $\nabla^2 f = \Delta f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} + \frac{\partial^2 f}{\partial z^2}$

Biharmonic: $\nabla^4 f = \Delta^2 f = \frac{\partial^4 f}{\partial x^4} + \frac{\partial^4 f}{\partial y^4} + \frac{\partial^4 f}{\partial z^4} + 2\frac{\partial^4 f}{\partial x^2 \partial y^2} + 2\frac{\partial^4 f}{\partial y^2 \partial z^2} + 2\frac{\partial^4 f}{\partial x^2 \partial z^2}$

3.5.11. Vector Calculus in Spherically Symmetric (Radial) Coordinates

Parameters: $r \geq 0$ (radial coordinate): uniform in every direction


Gradient: $\nabla f = \frac{df}{dr}\hat{r}$ Divergence: $\nabla \cdot u = \frac{1}{r^2} \frac{d}{dr} (r^2 u_r)$ Curl: $\nabla \times u = 0$

Laplacian: $\nabla^2 f = \Delta f = \frac{d^2 f}{dr^2} + \frac{2}{r} \frac{df}{dr} = \frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{df}{dr} \right)$ (auxiliary function: $u(r) = r f(r)$)

Biharmonic: $\nabla^4 f = \Delta^2 f = \frac{d^4 f}{dr^4} + \frac{2}{r} \frac{d^3 f}{dr^3} - \frac{1}{r^2} \frac{d^2 f}{dr^2} + \frac{1}{r^3} \frac{df}{dr}$ $\rightarrow \nabla^2 f = \frac{1}{r} \frac{d^2 u}{dr^2}$

Volume element: $dV = 4\pi r^2 dr$ (shell element)

3.5.12. Vector Calculus in Cylindrical Coordinates



Parameters: r : radius, $0 \leq \theta \leq 2\pi$: polar angle, z : elevation

Scalar field: $\mathbf{r} = r \hat{\mathbf{r}} + z \hat{\mathbf{z}}$

Vector field: $\mathbf{u} = u_r \hat{\mathbf{r}} + u_\theta \hat{\boldsymbol{\theta}} + u_z \hat{\mathbf{z}}$

Coordinate Conversions to and from Cartesian (x, y, z) :

$$x = r \cos \theta, \quad y = r \sin \theta, \quad z = z$$

$$r = \sqrt{x^2 + y^2}, \quad \theta = \text{atan2}(y, x), \quad z = z$$

Unit Vector Conversions to and from Cartesian $(\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}})$:

$$\hat{\mathbf{r}} = \cos \theta \hat{\mathbf{i}} + \sin \theta \hat{\mathbf{j}}, \quad \hat{\boldsymbol{\theta}} = -\sin \theta \hat{\mathbf{i}} + \cos \theta \hat{\mathbf{j}}, \quad \hat{\mathbf{z}} = \hat{\mathbf{k}}$$

$$\hat{\mathbf{i}} = \cos \theta \hat{\mathbf{r}} - \sin \theta \hat{\boldsymbol{\theta}}, \quad \hat{\mathbf{j}} = \sin \theta \hat{\mathbf{r}} + \cos \theta \hat{\boldsymbol{\theta}}, \quad \hat{\mathbf{k}} = \hat{\mathbf{z}}$$

Kinematics: time derivatives of displacement \mathbf{r}

Position: $\mathbf{r} = r \hat{\mathbf{r}} + z \hat{\mathbf{z}}$

Velocity: $\dot{\mathbf{r}} = \dot{r} \hat{\mathbf{r}} + r \dot{\theta} \hat{\boldsymbol{\theta}} + \dot{z} \hat{\mathbf{z}}$

Acceleration: $\ddot{\mathbf{r}} = (\ddot{r} - r \dot{\theta}^2) \hat{\mathbf{r}} + (2\dot{r} \dot{\theta} + r \ddot{\theta}) \hat{\boldsymbol{\theta}} + \ddot{z} \hat{\mathbf{z}}$

Angular Velocity: $\boldsymbol{\omega} = \dot{\theta} \hat{\mathbf{z}}$

Unit vectors: $\dot{\hat{\mathbf{r}}} = \dot{\theta} \hat{\boldsymbol{\theta}}, \quad \dot{\hat{\boldsymbol{\theta}}} = -\dot{\theta} \hat{\mathbf{r}}, \quad \dot{\hat{\mathbf{z}}} = \mathbf{0}$

Differential Elements

Line element: $d\mathbf{r} = dr \hat{\mathbf{r}} + r d\theta \hat{\boldsymbol{\theta}} + dz \hat{\mathbf{z}}$

Volume element: $dV = r dr d\theta dz$

(Jacobian: $\frac{\partial(x, y, z)}{\partial(r, \theta, z)} = r$)

Surface elements: $dS_r = r d\theta dz, \quad dS_\theta = r dr dz, \quad dS_z = r dr d\theta$

Distance: $d = \sqrt{r_1^2 + r_2^2 - 2r_1 r_2 \cos(\theta_1 - \theta_2) + (z_1 - z_2)^2}$

Vector Operators

Gradient: $\nabla f = \frac{\partial f}{\partial r} \hat{\mathbf{r}} + \frac{1}{r} \frac{\partial f}{\partial \theta} \hat{\boldsymbol{\theta}} + \frac{\partial f}{\partial z} \hat{\mathbf{z}}$


Divergence: $\nabla \cdot \mathbf{u} = \frac{1}{r} \frac{\partial(r u_r)}{\partial r} + \frac{1}{r} \frac{\partial u_\theta}{\partial \theta} + \frac{\partial u_z}{\partial z}$

Curl: $\nabla \times \mathbf{u} = \frac{1}{r} \begin{vmatrix} \hat{\mathbf{r}} & r \hat{\boldsymbol{\theta}} & \hat{\mathbf{z}} \\ \frac{\partial}{\partial r} & \frac{\partial}{\partial \theta} & \frac{\partial}{\partial z} \\ u_r & r u_\theta & u_z \end{vmatrix}$

Laplacian: $\nabla^2 f = \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial f}{\partial r} \right) + \frac{1}{r^2} \frac{\partial^2 f}{\partial \theta^2} + \frac{\partial^2 f}{\partial z^2}$

Biharmonic: $\Delta^2 f = \nabla^4 f = \frac{\partial^4 f}{\partial r^4} + \frac{2}{r^2} \frac{\partial^4 f}{\partial r^2 \partial \theta^2} + \frac{1}{r^4} \frac{\partial^4 f}{\partial \theta^4} + \frac{2}{r} \frac{\partial^3 f}{\partial r^3} - \frac{2}{r^3} \frac{\partial^3 f}{\partial r \partial \theta^2} - \frac{1}{r^2} \frac{\partial^2 f}{\partial r^2} + \frac{4}{r^4} \frac{\partial^2 f}{\partial \theta^2} + \frac{1}{r^3} \frac{\partial f}{\partial r} + \frac{\partial^4 f}{\partial z^4}$

3.5.13. Vector Calculus in Spherical Coordinates



Parameters: r : radius, $0 \leq \theta \leq 2\pi$: azimuth/longitude angle, $0 \leq \phi \leq \pi$: zenith/colatitude angle

Vector field: $\mathbf{u} = u_r \hat{\mathbf{r}} + u_\theta \hat{\boldsymbol{\theta}} + u_\phi \hat{\boldsymbol{\phi}}$

Coordinate Conversions to and from Cartesian (x, y, z) :

$$x = r \sin \phi \cos \theta, \quad y = r \sin \phi \sin \theta, \quad z = r \cos \phi$$

$$r = \sqrt{x^2 + y^2 + z^2}, \quad \theta = \text{atan2}(y, x), \quad \phi = \cos^{-1} \frac{z}{r}$$

Unit Vector Conversions to and from Cartesian $(\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}})$:

$$\begin{aligned} \hat{\mathbf{r}} &= \sin \phi \cos \theta \hat{\mathbf{i}} + \sin \phi \sin \theta \hat{\mathbf{j}} + \cos \phi \hat{\mathbf{k}}, & \hat{\mathbf{i}} &= \cos \theta \sin \phi \hat{\mathbf{r}} - \sin \theta \hat{\boldsymbol{\theta}} + \cos \theta \cos \phi \hat{\boldsymbol{\phi}} \\ \hat{\boldsymbol{\theta}} &= -\sin \theta \hat{\mathbf{i}} + \cos \theta \hat{\mathbf{j}}, & \hat{\mathbf{j}} &= \sin \theta \sin \phi \hat{\mathbf{r}} + \cos \theta \hat{\boldsymbol{\theta}} + \sin \theta \cos \phi \hat{\boldsymbol{\phi}} \\ \hat{\boldsymbol{\phi}} &= \cos \theta \cos \phi \hat{\mathbf{i}} + \sin \theta \cos \phi \hat{\mathbf{j}} - \sin \phi \hat{\mathbf{k}}, & \hat{\mathbf{k}} &= \cos \phi \hat{\mathbf{r}} - \sin \phi \hat{\boldsymbol{\phi}} \end{aligned}$$

Kinematics: time derivatives of displacement \mathbf{r}

Position: $\mathbf{r} = r \hat{\mathbf{r}}$

Velocity: $\dot{\mathbf{r}} = \dot{r} \hat{\mathbf{r}} + r \dot{\theta} \sin \phi \hat{\boldsymbol{\theta}} + r \dot{\phi} \hat{\boldsymbol{\phi}}$

Acceleration: $\ddot{\mathbf{r}} = (\ddot{r} - r \dot{\theta}^2 \sin^2 \phi - r \dot{\phi}^2) \hat{\mathbf{r}} + ((r \ddot{\theta} + 2 \dot{r} \dot{\theta}) \sin \phi + 2 r \dot{\theta} \dot{\phi} \cos \phi) \hat{\boldsymbol{\theta}} + (r \ddot{\phi} + 2 \dot{r} \dot{\phi} - r \dot{\theta}^2 \sin \phi \cos \phi) \hat{\boldsymbol{\phi}}$

Angular Velocity: $\boldsymbol{\omega} = \dot{\theta} \cos \phi \hat{\mathbf{r}} + \dot{\phi} \hat{\boldsymbol{\theta}} - \dot{\theta} \sin \phi \hat{\boldsymbol{\phi}} = \dot{\phi} \hat{\boldsymbol{\theta}} + \dot{\theta} \hat{\mathbf{k}}$

Unit vectors: $\dot{\hat{\mathbf{r}}} = -\dot{\theta} \sin \phi \hat{\boldsymbol{\theta}} + \dot{\phi} \hat{\boldsymbol{\phi}}, \quad \dot{\hat{\boldsymbol{\theta}}} = -\dot{\theta} \sin \phi \hat{\mathbf{r}} + \dot{\theta} \cos \phi \hat{\boldsymbol{\phi}}, \quad \dot{\hat{\boldsymbol{\phi}}} = \dot{\phi} \hat{\mathbf{r}} - \dot{\theta} \cos \phi \hat{\boldsymbol{\theta}}$

Differential Elements

Line element: $d\mathbf{r} = dr \hat{\mathbf{r}} + r \sin \phi d\theta \hat{\boldsymbol{\theta}} + r d\phi \hat{\boldsymbol{\phi}}$

Volume element: $dV = r^2 \sin \phi dr d\theta d\phi$ (Jacobian: $\frac{\partial(x, y, z)}{\partial(r, \phi, \theta)} = r^2 \sin \phi$)

Surface elements: $dS_r = r^2 \sin \phi d\theta d\phi, \quad dS_\theta = r dr d\phi, \quad dS_\phi = r \sin \phi dr d\theta$

Distance: $d = \sqrt{r_1^2 + r_2^2 - 2r_1 r_2 (\sin \phi_1 \sin \phi_2 \cos(\theta_1 - \theta_2) + \cos \phi_1 \cos \phi_2)}$

Vector Operators

Gradient: $\nabla f = \frac{\partial f}{\partial r} \hat{\mathbf{r}} + \frac{1}{r} \frac{\partial f}{\partial \phi} \hat{\boldsymbol{\phi}} + \frac{1}{r \sin \phi} \frac{\partial f}{\partial \theta} \hat{\boldsymbol{\theta}}$ Divergence: $\nabla \cdot \mathbf{u} = \frac{1}{r^2} \frac{\partial(r^2 u_r)}{\partial r} + \frac{1}{r \sin \phi} \frac{\partial(\sin \phi u_\phi)}{\partial \phi} + \frac{1}{r \sin \phi} \frac{\partial u_\theta}{\partial \theta}$

Curl: $\nabla \times \mathbf{u} = \frac{1}{r^2 \sin \phi} \begin{vmatrix} \hat{\mathbf{r}} & r \hat{\boldsymbol{\phi}} & r \sin \phi \hat{\boldsymbol{\theta}} \\ \partial/\partial r & \partial/\partial \phi & \partial/\partial \theta \\ u_r & r u_\phi & r \sin \phi u_\theta \end{vmatrix}$

Laplacian: $\nabla^2 f = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial f}{\partial r} \right) + \frac{1}{r^2 \sin \phi} \frac{\partial}{\partial \phi} \left(\sin \phi \frac{\partial f}{\partial \phi} \right) + \frac{1}{r^2 \sin^2 \phi} \frac{\partial^2 f}{\partial \theta^2}$

3.5.14. Vector Fields

Field Lines of a Vector Field

The plot of a vector-valued function $\mathbf{u} = \mathbf{f}(x, y, z) = \mathbf{f}(\mathbf{r}) = u_x \mathbf{i} + u_y \mathbf{j} + u_z \mathbf{k}$ is a vector field.

The equations of the field lines (curves with tangent vectors \mathbf{u}) are given by $\frac{dx}{u_x} = \frac{dy}{u_y} = \frac{dz}{u_z}$,

so that e.g. the field lines in the plane $z = 0$ satisfy $\frac{dy}{dx} = \frac{u_y(x, y, 0)}{u_x(x, y, 0)}$.

Potentials of Vector Fields

- **Irrotational (conservative) field:** if $\nabla \times \mathbf{u} = 0$ (curl free). In this case, then there exists a scalar potential f such that $\mathbf{u} = \nabla f$. For some applications it is more natural to use $\mathbf{u} = -\nabla f$. The isosurfaces of f are the equipotentials of \mathbf{u} , and \mathbf{u} is perpendicular to these isosurfaces (the normal vectors of \mathbf{u}) everywhere.
- **Solenoidal (incompressible) field:** if $\nabla \cdot \mathbf{u} = 0$ (divergence free). In this case, then there exists a vector potential \mathbf{A} such that $\mathbf{u} = \nabla \times \mathbf{A}$.
 \mathbf{A} is usually chosen so that $\nabla \cdot \mathbf{A} = 0$. Then, \mathbf{u} is the vorticity field of \mathbf{A} .

Decompositions of Vector Fields

- **Helmholtz Decomposition:** a field \mathbf{u} can be written as $\mathbf{u} = -\nabla f + \nabla \times \mathbf{A}$ (irrotational part plus solenoidal part). The expressions for f and \mathbf{A} are

$$4\pi f(\mathbf{r}) = \iiint_V \frac{\nabla_{\mathbf{r}'} \cdot \mathbf{u}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dV' - \iint_S \hat{\mathbf{n}}' \cdot \frac{\mathbf{u}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS', \quad 4\pi \mathbf{A}(\mathbf{r}) = \iiint_V \frac{\nabla_{\mathbf{r}'} \times \mathbf{u}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dV' - \iint_S \hat{\mathbf{n}}' \times \frac{\mathbf{u}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS'$$

- **Poloidal-Toroidal (Chandrasekhar-Kendall) Decomposition:** if $\nabla \cdot \mathbf{u} = 0$ (solenoidal) and \mathbf{u} is defined in spherical coordinates $\{\mathbf{e}_r, \mathbf{e}_\theta, \mathbf{e}_\phi\}$, then the vector potential $\mathbf{u} = \nabla \times \mathbf{A}$ can be written as $\mathbf{A} = \mathbf{T} + \mathbf{P}$, where $\mathbf{T} = T \mathbf{e}_r$ (toroidal part) and $\mathbf{P} = \nabla \times (P \mathbf{e}_r) = \nabla P \times \mathbf{e}_r$ (poloidal part).

The toroidal part satisfies $\mathbf{e}_r \cdot \mathbf{T} = 0$. The poloidal part satisfies $\mathbf{e}_r \cdot (\nabla \times \mathbf{P}) = 0$.

$T(\mathbf{r})$ and $P(\mathbf{r})$ are scalar fields that satisfy the Poisson equations $\mathbf{e}_r \cdot (\nabla \times \mathbf{u}) = -\Delta_H T$ and $\mathbf{e}_r \cdot \mathbf{u} = -\Delta_H P$, where Δ_H is the scalar Laplacian containing only the $\{\theta, \phi\}$ ('horizon') terms.

3.5.15. Vector Calculus Identities

Properties of Vector Calculus Operators with Vector Operators: for scalar-valued functions f and vector-valued functions \mathbf{u} ,

$$\begin{aligned} \nabla(f_1 + f_2) &= \nabla f_1 + \nabla f_2 & \nabla \cdot (\mathbf{u}_1 + \mathbf{u}_2) &= \nabla \cdot \mathbf{u}_1 + \nabla \cdot \mathbf{u}_2 \\ \nabla \times (\mathbf{u}_1 + \mathbf{u}_2) &= \nabla \times \mathbf{u}_1 + \nabla \times \mathbf{u}_2 & \nabla \cdot (f\mathbf{u}) &= f \nabla \cdot \mathbf{u} + (\nabla f) \cdot \mathbf{u} \\ \nabla \times (f\mathbf{u}) &= f \nabla \times \mathbf{u} + (\nabla f) \times \mathbf{u} & \nabla \cdot (\mathbf{u}_1 \times \mathbf{u}_2) &= \mathbf{u}_2 \cdot \nabla \times \mathbf{u}_1 - \mathbf{u}_1 \cdot \nabla \times \mathbf{u}_2 \\ \nabla \cdot \nabla \times \mathbf{u} &= 0 & \nabla \times \nabla f &= \mathbf{0} \\ \nabla \times (\nabla \times \mathbf{u}) &= \nabla(\nabla \cdot \mathbf{u}) - \nabla^2 \mathbf{u} & \nabla \times (\mathbf{u}_1 \times \mathbf{u}_2) &= \mathbf{u}_1 \nabla \cdot \mathbf{u}_2 - \mathbf{u}_2 \nabla \cdot \mathbf{u}_1 + (\mathbf{u}_2 \cdot \nabla) \mathbf{u}_1 - (\mathbf{u}_1 \cdot \nabla) \mathbf{u}_2 \\ \mathbf{u} \times (\nabla \times \mathbf{u}) + (\mathbf{u} \cdot \nabla) \mathbf{u} &= \frac{1}{2} \nabla(u^2) & (\nabla^2 \mathbf{u} &= [\nabla^2 u_x, \nabla^2 u_y, \nabla^2 u_z]^T: \text{vector Laplacian.}) \end{aligned}$$

Gauss Theorem (divergence theorem): for a closed surface S enclosing a volume V , with the outward normal taken for $d\mathbf{A}$, the total emitted flux Φ is equal to the net internal divergence

$$\Phi = \iiint_V \nabla \cdot \mathbf{u} \, dV = \oiint_S \mathbf{u} \cdot d\mathbf{A}$$

Stokes Theorem (curl theorem): for an open surface S bounded by a closed curve C circulating S anti-clockwise, the circulation Γ is equal to the net enclosed rotation (flux of vorticity):

$$\Gamma = \oint_C \mathbf{u} \cdot d\mathbf{l} = \iint_S \nabla \times \mathbf{u} \cdot d\mathbf{A}$$

Green's Theorem: in planar 2D space, Stokes' theorem reduces to:

$$\Gamma = \oint_C \mathbf{u} \cdot d\mathbf{l} = \oint_C u_x \, dx + u_y \, dy = \iint_S \frac{\partial u_y}{\partial x} - \frac{\partial u_x}{\partial y} \, dx \, dy$$

Green's First Identity: divergence theorem with $\mathbf{u} = f \nabla g$ and $\nabla \cdot (f \nabla g) = \nabla f \cdot \nabla g + f \nabla \cdot \nabla g$.

$$\oiint_S f \nabla g \cdot d\mathbf{A} = \iiint_V (\nabla f \cdot \nabla g) \, dV + \iiint_V f \nabla^2 g \, dV$$

Green's Second Identity: difference of symmetric forms of first identity.

$$\oiint_S (f \nabla g - g \nabla f) \cdot d\mathbf{A} = \iiint_V (f \nabla^2 g - g \nabla^2 f) \, dV$$

Green's Third Identity: in the second identity, let $f = G$ (a Green's function, Section 3.7.4), chosen suitably for the PDE to be solved such that $G = 0$ on the boundary. Substitute solutions for ∇G and $\nabla^2 G$.

3.5.16. Differential Operators

Expressions of derivatives of functions can be written as operators acting on functions.

$$n\text{th partial differential operator: } D_x^n = \frac{\partial^n}{\partial x^n} \Rightarrow D_x^n y = \frac{\partial^n y}{\partial x^n}$$

$$\text{Example: } Dx D = \frac{d}{dx} \left(x \frac{d}{dx} \right) = \frac{d}{dx} + x \frac{d^2}{dx^2} = D + x D^2$$

$$\text{General operators can be constructed e.g. } L = 2D_x^2 - x D_x D_y \Rightarrow L\phi = 2 \frac{\partial^2 \phi}{\partial x^2} - x \frac{\partial \phi}{\partial x} \frac{\partial \phi}{\partial y}$$

Differential operators are generally not commutative.

Linear differential operators with constant coefficients are commutative.

Separation of a coupled system of linear partial differential equations: using operators

If L_i are **commutative** linear differential operators with $L_1 L_3 = L_3 L_1$ and $L_2 L_4 = L_4 L_2$ then the coupled system of PDEs $\{L_1 u + L_2 v = f; L_3 u + L_4 v = g\}$ can be uncoupled to yield the two PDEs

$$\{(L_4 L_1 - L_2 L_3)u = L_4 f - L_2 g; (L_3 L_2 - L_1 L_4)v = L_3 f - L_1 g\}$$

($u = u(x, y, \dots)$, $v = v(x, y, \dots)$): dependent variables, L_i : linear differential operators of $\{x, y, \dots\}$ with constant coefficients, $\{f = f(x, y, \dots), g = g(x, y, \dots)\}$: differentiable functions)

3.5.17. Multivariable (Spatial) Continuous Fourier Transform

For the ordinary Fourier transform, see Section 3.6.4. When the input domain is multi-dimensional (e.g. position space $\mathbf{x} = [x, y, z]$ rather than time t), the output frequency domain is also multi-dimensional (e.g. spatial frequency $\mathbf{k} = [k_x, k_y, k_z]$ rather than temporal frequency ω).

$$F(\mathbf{k}) = \int_{C^n} f(\mathbf{x}) \exp(-j \mathbf{k} \cdot \mathbf{x}) d\mathbf{x}$$

Forward Fourier Transform

$$f(\mathbf{x}) = \frac{1}{(2\pi)^n} \int_{C^n} F(\mathbf{k}) \exp(j \mathbf{k} \cdot \mathbf{x}) d\mathbf{k}$$

Inverse Fourier Transform

The integration is generally performed over all $\mathbf{x} \in C^n$ and $\mathbf{k} \in C^n$. In many practical applications, the function $f(\mathbf{x})$ is real and even so $F(\mathbf{k})$ is real and even so $\mathbf{x} \in R^n$ and $\mathbf{k} \in R^n$.

In quantum mechanics $f = \psi \in C$ and \mathbf{k} -space is momentum space (since $\mathbf{p} = \hbar \mathbf{k}$).

Multivariable Parseval's Theorem:

$$\int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} |f(\mathbf{x})|^2 d\mathbf{x} = \frac{1}{(2\pi)^n} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} |F(\mathbf{k})|^2 d\mathbf{k}$$

3.5.18. Multivariable (Spatial) Discrete Fourier Transform

For the ordinary discrete Fourier transform (DFT), see Section 3.6.5.

When the input domain is discrete multi-dimensional (e.g. pixel space $\mathbf{w} = [u, v]$), the output frequency domain is also discrete multi-dimensional (e.g. discrete spatial frequency $\mathbf{k} = [k_u, k_v]$).

$$F(\mathbf{k}) = \sum_{n_1=0}^{N_1-1} \dots \sum_{n_m=0}^{N_m-1} f(\mathbf{w}) \exp\left(-j \sum_{a=1}^m \frac{2\pi i k_a n_a}{N_a}\right)$$

Forward Discrete Fourier Transform

$$f(\mathbf{w}) = \frac{1}{N_1 \dots N_m} \sum_{k_1=0}^{N_1-1} \dots \sum_{k_m=0}^{N_m-1} F(\mathbf{k}) \exp\left(j \sum_{a=1}^m \frac{2\pi i k_a n_a}{N_a}\right)$$

Inverse Discrete Fourier Transform

Multivariable Convolution Theorem: $F(\mathbf{k}) G(\mathbf{k}) = \text{DFT}\{(f * g)(\mathbf{w})\}$ (DFT: discrete FT)

where the convolution is circular (periodic). Common application: image filtering (Section 5.6.1).

3.5.18. Multivariable Z-Transform

For the ordinary Z-transform, see Section 3.4.19.

Forward Z-transform:

$$F(z_1, \dots, z_m) = \sum_{n_1=-\infty}^{\infty} \dots \sum_{n_m=-\infty}^{\infty} f(n_1, \dots, n_m) \times (z_1^{n_1} \dots z_m^{n_m})$$

3.6. Fourier Series and Fourier Transforms

3.6.1. General Fourier Series Definition

Real-Valued Fourier Series

The real-valued Fourier series is defined for a function $f(t)$ on $0 \leq t < T$ as

$$f(t) = d + \sum_{n=1}^{\infty} \left(a_n \cos \frac{2\pi nt}{T} + b_n \sin \frac{2\pi nt}{T} \right)$$

$$\text{where } d = \frac{1}{T} \int_0^T f(t) dt, \quad a_n = \frac{2}{T} \int_0^T f(t) \cos \frac{2\pi nt}{T} dt, \quad b_n = \frac{2}{T} \int_0^T f(t) \sin \frac{2\pi nt}{T} dt$$

If the function $f(t)$ is periodic, of period T , then these relationships are valid for all t . The integrals may then be taken over any range of T .

- If $f(t)$ is even then $b_n = 0$.
- If $f(t)$ is odd then $a_n = 0$.
- If $f(t)$ has zero mean value then $d = 0$.

The rate of convergence, $O(n^{-(k+1)})$, is such that the k -th derivative of $f(t)$ is discontinuous.

Complex-Valued Fourier Series

Equivalently, the complex-valued Fourier series is

$$f(t) = \sum_{n=-\infty}^{\infty} c_n e^{i2\pi nt/T} \quad \text{where} \quad c_n = \frac{1}{T} \int_0^T f(t) e^{-i2\pi nt/T} dt$$

$$\text{i.e. } f(t) = \sum_{n=-\infty}^{\infty} c_n e^{in\omega_0 t} \quad \text{where} \quad c_n = \frac{1}{T} \int_0^T f(t) e^{-in\omega_0 t} dt$$

The relationship between the complex and real forms of the coefficients is

$$c_n = \begin{cases} \frac{1}{2}(a_n - ib_n) & \text{for } n > 0 \\ d & \text{for } n = 0 \end{cases}$$

and, for real functions $f(t)$, we have $c_{-n} = c_n^*$.

The (scientific) fundamental frequency is $\omega_0 = \frac{2\pi}{T}$ and the (scientific) n th harmonic is $n\omega_0$.

3.6.2. Half-Range Fourier Series Definition

If a Fourier series representation of $f(x)$ is required to be valid only in $0 \leq x \leq L$, then it only needs to contain either the sine terms alone or the cosine terms alone. For example

$$f(x) = \sum_{n=1}^{\infty} b_n \sin \frac{n\pi x}{L}$$

where $b_n = \frac{2}{L} \int_0^L f(x) \sin \frac{n\pi x}{L} dx$

Note that the wavelength of the first term in the series ($n = 1$) is $2L$ rather than L (as would be the case for the full-range series).


3.6.3. Fourier Series of Common Waveforms

Half-wave Rectified Cosine

$$f(t) = \frac{1}{\pi} + \frac{1}{2} \cos \omega_0 t + \frac{2}{\pi} \sum_{m=1}^{\infty} (-1)^{m+1} \frac{\cos 2m\omega_0 t}{4m^2 - 1}$$

$$f(t) = \frac{1}{\pi} + \frac{1}{4} e^{i\omega_0 t} + \frac{1}{4} e^{-i\omega_0 t} + \frac{1}{\pi} \sum_{\substack{n=-\infty \\ n \text{ even} \\ n \neq 0}}^{\infty} (\pm 1) \frac{e^{in\omega_0 t}}{n^2 - 1}$$

signs alternate, + for $n = 2$




Full-wave Rectified Cosine

$$f(t) = \frac{2}{\pi} \left[1 + 2 \sum_{m=1}^{\infty} (-1)^{m+1} \frac{\cos 2m\omega_0 t}{4m^2 - 1} \right]$$

$$f(t) = \frac{2}{\pi} \left[1 + \sum_{\substack{n=-\infty \\ n \text{ even} \\ n \neq 0}}^{\infty} (\pm 1) \frac{e^{in\omega_0 t}}{n^2 - 1} \right]$$


signs alternate, + for $n = 2$



Square Wave

$$f(t) = \frac{4}{\pi} \sum_{m=1}^{\infty} \frac{\sin(2m-1)\omega_0 t}{2m-1}$$

$$f(t) = \sum_{\substack{n=-\infty \\ n \text{ odd}}}^{\infty} \frac{2}{i\pi n} e^{in\omega_0 t}$$




Triangular Wave

$$f(t) = \frac{8}{\pi^2} \sum_{m=1}^{\infty} (-1)^{m+1} \frac{\sin(2m-1)\omega_0 t}{(2m-1)^2}$$

$$f(t) = \frac{4}{i\pi^2} \sum_{\substack{n=-\infty \\ n \text{ odd}}}^{\infty} (\pm 1) \frac{e^{in\omega_0 t}}{n^2}$$


signs alternate, + for $n = 1$

**Sawtooth Wave**

$$f(t) = \frac{2}{\pi} \sum_{n=1}^{\infty} (-1)^{n+1} \frac{\sin n\omega_0 t}{n}$$


$$f(t) = \frac{1}{i\pi} \sum_{\substack{n=-\infty \\ n \neq 0}}^{\infty} (\pm 1) \frac{e^{in\omega_0 t}}{n}$$

signs alternate, + for $n = 1$

**Pulse Wave**

$$f(t) = \frac{a}{T} \left[1 + 2 \sum_{n=1}^{\infty} \frac{\sin \frac{n\pi a}{T}}{\frac{n\pi a}{T}} \cos n\omega_0 t \right]$$

$$f(t) = \frac{a}{T} \left[1 + \sum_{\substack{n=-\infty \\ n \neq 0}}^{\infty} \frac{\sin \frac{n\pi a}{T}}{\frac{n\pi a}{T}} e^{in\omega_0 t} \right]$$



3.6.4. Fourier Transforms

The Fourier transform maps a continuous time domain t to a continuous frequency domain ω :

$$\hat{y}(\omega) = \int_{-\infty}^{\infty} y(t) e^{-i\omega t} dt$$

Forward Fourier Transform

$$y(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{y}(\omega) e^{i\omega t} d\omega$$

Inverse Fourier Transform

- Some sources handle the 2π factor differently and define transforms with differences in signs of the exponent. All transform theorems are valid as long as it is done consistently.
- Fourier transforms are sometimes written in terms of frequency $f = \omega / 2\pi$ (as done below in the analogous discrete Fourier transform).
- The Fourier transform is a slice along the imaginary axis of the Laplace transform: $s = i\omega$; $\omega = \text{Im}\{s\}$.

3.6.5. Discrete Fourier Transforms

The discrete Fourier transform (DFT) maps a finite sequence ($x_n, n = 0, 1, \dots, N-1$) to a finite sequence of discrete frequencies ($X_k, k = 0, 1, \dots, N-1$)

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i k n / N}$$

Forward Discrete Fourier Transform

$$x_n = \frac{1}{N} \sum_{k=0}^{N-1} X_k e^{2\pi i k n / N}$$

Inverse Discrete Fourier Transform

Notes:

- The DFT gives a discrete approximation to the frequency spectrum of a continuous function passing through all points x_n containing frequency components no higher than $\frac{1}{2T}$, where T is the sampling period. The Nyquist frequency is equal to twice the highest frequency contained in the original signal before sampling. Sampling at a rate below the Nyquist frequency leads to high-frequency information loss and aliasing effects (distorsion).
- Time and frequency parameter relation: $f_k = \frac{k}{NT}$ [] and $t_n = nT$ [s] for integers $0 \leq n, k < N$.
- Total sampling time: NT . Fundamental frequency: $f_1 = \frac{1}{NT}$. Sampling rate: $f_N = \frac{1}{T}$.
- For real sequences, $X_k^* = X_{N-k}$, since the cosine waves at these frequencies pass the same points.

Common DFTs (for $N = 4$):

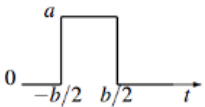
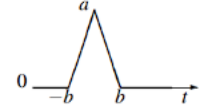
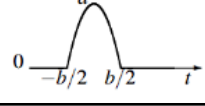
- Cosine: $x_n = \cos \frac{n\pi}{2}$ i.e. $x_n = [1, 0, -1, 0] \rightarrow X_k = [0, 2, 0, 2]$ ($f_1 = \frac{1}{4T}, f_3 = \frac{3}{4T}$)
- Sine: $x_n = \sin \frac{n\pi}{2}$ i.e. $x_n = [0, 1, 0, -1] \rightarrow X_k = [0, -2i, 0, 2i]$

For infinite discrete sequences ($N \rightarrow \infty$ and summing in both directions), the DFT is called the

'Discrete-Time Fourier Transform' (DTFT), $X(\omega) = \sum_{n=-\infty}^{\infty} x_n e^{-i\omega n}$, which is equivalent to the the

bilateral Z-transform $\hat{X}(z)$ of x_n with $z = e^{i\omega}$.

3.6.6. Fourier Transforms of Common Signals

Time-domain waveform $g(t)$	Frequency spectrum $G(\omega)$
1 (DC level)	$2\pi \delta(\omega) = \delta(f)$
$H(t)$ (Heaviside step)	$\pi \delta(\omega) + \frac{1}{j\omega}$
$e^{j\omega_0 t}$ (complex sinusoid)	$2\pi \delta(\omega - \omega_0)$
$\cos \omega_0 t$	$\pi [\delta(\omega - \omega_0) + \delta(\omega + \omega_0)]$
$\sin \omega_0 t$	$\frac{\pi}{j} [\delta(\omega - \omega_0) - \delta(\omega + \omega_0)]$
e^{-at^2} (Gaussian)	$\sqrt{\frac{\pi}{a}} e^{-\frac{1}{4a}\omega^2}$
$\sum_{n=-\infty}^{\infty} \delta(t - nT)$ (impulse train)	$\frac{2\pi}{T} \sum_{m=-\infty}^{\infty} \delta(\omega - \frac{2\pi m}{T})$
 $a \text{ rect } \frac{t}{b}$ (rectangle function)	$ab \text{ sinc } \frac{\omega b}{2}$ where $\text{sinc } x := (\sin x) / x$
 $a \text{ tri } \frac{t}{b}$ (triangle function)	$ab \text{ sinc}^2 \frac{\omega b}{2}$ main lobe bandwidth: $1 / b$
 $a \cos \frac{\pi t}{b} \cdot H(\frac{t}{b} - x)$	$\frac{ab}{2} [\text{sinc } \frac{\omega b - \pi}{2} + \text{sinc } \frac{\omega b + \pi}{2}]$
$a f(t) + b g(t)$	$a F(\omega) + b G(\omega)$ (linearity)
$g(t - t_0)$ (time shift)	$e^{-j\omega t_0} G(\omega)$
$e^{j\omega_0 t} g(t)$	$G(\omega - \omega_0)$ (frequency shift)
$\frac{d^n g(t)}{dt^n}$ (differentiation)	$(j\omega)^n G(\omega)$
$(f * g)(t) = \int_{-\infty}^{\infty} f(\tau) g(t - \tau) d\tau$ (convolution)	$F(\omega) G(\omega)$ (multiplication)
$f(t) g(t)$ (multiplication)	$\frac{1}{2\pi} (F * G)(\omega)$ (convolution)
$G(t)$ (duality)	$2\pi g(-\omega)$
if $g(t) \rightarrow G(\omega)$ then $G(t) \rightarrow 2\pi g(-\omega)$	
$g(t)^*$ (complex conjugate)	$G(-\omega)^*$
real-valued even function $g(t)$	real-valued even function $G(\omega)$
real-valued odd function $g(t)$	imaginary-valued odd function $G(\omega)$

3.6.7. Signal Energy, Power and Parseval's Theorem of Energy Conservation

The energy of a signal $g(t)$ is defined as $E_g = \int_{-\infty}^{\infty} |g(t)|^2 dt$.

The power of a signal $g(t)$ is defined as $P_g = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} |g(t)|^2 dt$.

The power spectral density (power spectrum) is $S_{gg}(\omega) = \frac{1}{2\pi} |G(\omega)|^2 = |G(f)|^2$.

Parseval's Theorem: The energy in the time and frequency domain must be the same:

$$E_g = \int_{-\infty}^{\infty} |g(t)|^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |G(\omega)|^2 d\omega$$

For more on signal analysis, see Section 5.4.

3.7. Partial Differential Equations and Variational Calculus

3.7.1. Classification of Linear Second-Order Partial Differential Equations (PDEs)

A PDE of the form

$$A \frac{\partial^2 u}{\partial x^2} + 2B \frac{\partial^2 u}{\partial x \partial y} + C \frac{\partial^2 u}{\partial y^2} + D \frac{\partial u}{\partial x} + E \frac{\partial u}{\partial y} + Fu + G = 0$$

is said to be

- **Elliptic:** if $B^2 < AC$ (e.g. Laplace equation, Poisson equation)
- **Parabolic:** if $B^2 = AC$ (e.g. Heat equation, Diffusion equation)
- **Hyperbolic:** if $B^2 > AC$ (e.g. Wave equation)

3.7.2. Classification of Boundary Conditions

A PDE requires initial conditions (ICs: $u = f(x, 0)$) and boundary conditions (BCs) to be fully specified. The BCs constrain the value of the dependent variable on the boundary of the region satisfying the PDE. The types of BCs are:

- **Dirichlet:** dependent variable specified on boundary
e.g. $u = 0$ when $x = 0$; $u = 1 - e^{-3t}$ when $x = 1$
- **Neumann:** gradient of dependent variable specified on boundary
e.g. $\frac{du}{dx} = 0$ when $x = 0$ and $x = 1$
- **Robin:** an ODE for the dependent variable is specified on the boundary
e.g. $\frac{du}{dx} + 2u = x$ when $|x| = 1$, defined for $|x| \leq 1$.
- **Mixed:** the boundary is split into several parts, each with different conditions
e.g. $u = 4t$ when $x = 0$; $\frac{du}{dx} = 1 - u$ when $x = 1$, defined on $0 \leq x \leq 1$.

3.7.3. Common PDEs and their 2D General Solutions

Partial Differential Equations

$\frac{\partial T}{\partial t} = \alpha \nabla^2 T + \frac{\alpha}{k} \dot{q}_{\text{gen}}$	$\frac{\partial c}{\partial t} = D \nabla^2 c + \dot{c}_{\text{gen}}$	$\nabla^2 \phi = \rho_{\text{gen}}$	$\frac{\partial^2 \psi}{\partial t^2} = c^2 \nabla^2 \psi$
Heat Equation	Diffusion Equation	Poisson's Equation	Wave Equation
(T : temperature, α : thermal diffusivity, k : thermal conductivity, q : heat source per volume)	(c : concentration, D : diffusion coefficient)	(Laplace's Equation if $\rho_{\text{gen}} = 0$) (ϕ : scalar potential, ρ : source density)	(ψ : scalar displacement, c : phase velocity)

Heat Equation / Diffusion Equation: 2D Solutions

Separation of variables using $T(x, t) = X(x) \bar{T}(t)$ with a negative separation constant $-\lambda^2$ gives $X'' + \lambda^2 X = 0$ and $\bar{T}' + \alpha \lambda^2 \bar{T} = 0$ with solutions $X(x) = A \sin \lambda x + B \cos \lambda x$ and $\bar{T}(t) = C \exp(-\alpha \lambda^2 t)$.

Self-similar solution: let $\eta = \frac{x}{\sqrt{\alpha t}}$, separate as $T(x, t) = H(\eta) \bar{T}(t)$. Find an ODE in $H(\eta)$.

Heat flux: $\mathbf{q} = -k \nabla T$; Diffusion flux: $\mathbf{J} = -D \nabla c$; Convection boundary condition: $k \frac{\partial T}{\partial n_{\text{out}}} = h(T_{\infty} - T)$

Normalisation condition: all solutions satisfy $\frac{d}{dt} \int_{-\infty}^{\infty} T(x, t) dx = 0$ (conservation law).

Laplace's Equation: 2D Solutions

General solutions may be

$$\begin{aligned} \phi(x, y) = X(x) Y(y) &= (A \sinh \lambda x + B \cosh \lambda x)(C \sin \lambda y + D \cos \lambda y) \\ &= (A \sin \lambda x + B \cos \lambda x)(C \sinh \lambda y + D \cosh \lambda y) \\ &= (Ax + B)(Cy + D) \end{aligned}$$

Separation constant

using $+\lambda^2$
using $-\lambda^2$
using 0

Wave Equation: 1D Solutions

Separation of variables using $Y(x, t) = X(x) T(t)$ with a negative separation constant $-\lambda^2$ gives $X'' + \lambda^2 X = 0$ and $\bar{T}' + \alpha \lambda^2 \bar{T} = 0$ with solutions $X(x) = A \sin \lambda x + B \cos \lambda x$ and $T(t) = C \cos \lambda ct + D \sin \lambda ct$.

D'Alembert's solution: consider the PDEs $\frac{\partial y}{\partial t} \pm c \frac{\partial y}{\partial x} = 0$ and let $\eta = ct \mp x$. It is clear that any $y = f(\eta)$ is a solution, so the general solution is $y(x, t) = f(ct - x) + g(ct + x)$. The lines $ct - x = \text{constant}$ and $ct + x = \text{constant}$, along which the right and left running waves move in the (x, t) plane, are the characteristics of the PDE. By differentiation, these PDEs are jointly equivalent to the wave equation.

3.7.4. Techniques for Solving Partial Differential Equations

For a typical PDE for a multivariable function u in terms of space x and time t ,

Separation of Variables: assume $u(x, t) = X(x) T(t)$.

1. Substitute $u = XT$ where X is a function of x only and T is a function of t only. Compute the derivatives as e.g. $u_{xx} = X''(x) T(x)$, $u_t = X(x) T'(t)$ etc.
2. Rearrange so that all X terms and all T terms are on opposite sides. Set both sides equal to a separation constant, initially $\pm\lambda^2$.
3. Solve each ODE for $X(x)$ and $T(t)$ in terms of λ , selecting the appropriate sign (or zero) depending on whether the behaviour can meet the boundary conditions: it may be oscillatory, exponential, or linear.
4. Use initial conditions and boundary conditions to constrain or discretise λ and the undetermined coefficients.
5. Use superposition to sum over remaining undetermined indices, if required.

Laplace Transforms: defined such that $L\{u(x, t)\} = U(x, s) = \int_0^{\infty} u(x, t) e^{-st} dt$

1. Transform the PDE using e.g. $L\{u_x\} = U_x(x, s)$, $L\{u_t\} = s U(x, s) - u(x, 0)$, etc. Also transform the boundary conditions e.g. $\{u(0, t) = f(t)\} \rightarrow \{U(0, s) = F(s)\}$, and substitute the initial condition for $u(x, 0)$.
2. Solve the resulting ODE for $U(x, s)$ and apply boundary conditions.
3. The inverse Laplace transform of $U(x, s)$ (considering x as a constant) is $u(x, t)$.

Fourier and other integral transforms may also be used.

Green's Functions: multivariable convolution of input function with the impulse response to solve the equation $L[u] = f$.

1. Identify independent linear differential operators L and the forcing functions (nonhomogeneous components) f .
2. Find or look up (Section 3.7.5) the Green's function $G(\mathbf{x})$ for the operator L .
3. Apply initial conditions and boundary conditions to constrain G .
4. Apply the convolution theorem as $u(\mathbf{x}) = \int f(\boldsymbol{\eta}) G(\mathbf{x} - \boldsymbol{\eta}) d\boldsymbol{\eta}$.

3.7.5. Green's Functions

A Green's function G is the impulse response of a linear differential operator: $L [G] = \delta(\mathbf{x}, t)$

In the table, $r^2 = x^2 + y^2 (+ z^2)$ in 2D (or 3D). (H : step function, I : Bessel functions)

Differential operator L	Green's function $G(x, y, [z, t])$
∇_{2D}^2 2D Poisson equation	$\frac{1}{2\pi} \ln r$
∇_{3D}^2 3D Poisson equation	$-\frac{1}{4\pi r}$
$\nabla_{3D}^2 + k^2$ Schrodinger equation	$-\frac{e^{-ikr}}{4\pi r}$
$\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial x^2}$ 1D wave equation	$\frac{1}{2c} H\left(t - \left \frac{x}{c}\right \right)$
$\frac{\partial^2}{\partial t^2} - c^2 \nabla_{2D}^2$ 2D wave equation	$\frac{1}{2\pi c \sqrt{c^2 t^2 - r^2}} H\left(t - \frac{r}{c}\right)$
$\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla_{3D}^2$ 3D wave equation	$\frac{1}{4\pi r} \delta\left(t - \frac{r}{c}\right)$
$\frac{\partial}{\partial t} - k \frac{\partial^2}{\partial x^2}$ 1D diffusion equation	$\left(\frac{1}{4\pi kt}\right)^{1/2} \exp\left(-\frac{x^2}{4kt}\right) H(t)$
$\frac{\partial}{\partial t} - k \nabla_{2D}^2$ 2D diffusion equation	$\frac{1}{4\pi kt} \exp\left(-\frac{r^2}{4kt}\right) H(t)$
$\frac{\partial}{\partial t} - k \nabla_{3D}^2$ 3D diffusion equation	$\left(\frac{1}{4\pi kt}\right)^{3/2} \exp\left(-\frac{r^2}{4kt}\right) H(t)$
$\frac{\partial^2}{\partial t^2} + 2\gamma \frac{\partial}{\partial t} - c^2 \frac{\partial^2}{\partial x^2}$ Telegrapher's equation	$\frac{1}{2} e^{-\gamma t} \left[\delta(ct - x) + \delta(ct + x) + H(ct - x) \left(\frac{\gamma}{c} I_0\left(\frac{\gamma u}{c}\right) + \frac{\gamma t}{u} I_1\left(\frac{\gamma u}{c}\right) \right) \right]$ where $u = \sqrt{c^2 t^2 - r^2}$

3.7.6. Fundamental Lemma of Calculus of Variations

If a continuous function f on an open interval (a, b) satisfies $\int_a^b f(x)h(x) dx = 0$ for all compactly supported smooth functions h on (a, b) , then $f(x) \equiv 0$.

Corollary: if $\int_a^b f(x)h(x) + g(x)h'(x) dx = 0$ then $g'(x) = f(x)$ (same conditions as above).

3.7.7. Euler-Lagrange Equation

A 'variation' (or perturbation) $f + \delta f$ of a function $f(x)$ constrained to satisfy $f(a) = A$ and $f(b) = B$ can be written as $g_\varepsilon(x) = f(x) + \varepsilon \eta(x)$ where ε is small and $\eta(x)$ is a differentiable function satisfying $\eta(a) = \eta(b) = 0$.

Let S be a functional defined as $S[f] = \int_a^b f(t, x, \dot{x}) dt$

The function f for which S is stationary i.e. local gradient to small variations is zero, satisfies the differential equation

$$\frac{\partial f}{\partial x} = \frac{d}{dt} \left(\frac{\partial f}{\partial \dot{x}} \right).$$

If f is explicitly a function of higher derivatives of x , then the equation is

$$\frac{\partial f}{\partial x} - \frac{d}{dt} \left(\frac{\partial f}{\partial \dot{x}} \right) + \frac{d^2}{dt^2} \left(\frac{\partial f}{\partial \ddot{x}} \right) - \dots + (-1)^k \frac{d^k}{dt^k} \left(\frac{\partial f}{\partial x^{(k)}} \right) = 0.$$

Note that x, x' etc. are considered **independent** variables.

Common applications:

- Minimum action (Lagrangian / Hamiltonian for mechanical systems, see Section 6.2.12).
- Minimum distance on a curved surface (geodesics)
- Minimum energy configuration of a system

M4. LINEAR ALGEBRA

4.1. Vector and Matrix Algebra

4.1.1. Properties of Real Matrices

Zero, Ones and Identity Matrices:

- Zero matrix: $\mathbf{0}$; Ones matrix $\mathbf{1}$; Identity matrix \mathbf{I}_n ($n \times n$ matrix with 1's on leading diagonal)

Singular Matrices:

- If $|\mathbf{A}| = 0$ then \mathbf{A} is singular (non-invertible) and \mathbf{A}^{-1} does not exist.

Symmetric and Antisymmetric Matrices

- If $\mathbf{A} = \mathbf{A}^T$ then \mathbf{A} is symmetric ($A_{ij} = A_{ji}$) and \mathbf{A}^{-1} exists.
- If $\mathbf{A} = -\mathbf{A}^T$ then \mathbf{A} is antisymmetric (skew-symmetric) ($A_{ij} = -A_{ji}$), \mathbf{A}^{-1} does **not** exist and \mathbf{A} has zeros on the leading diagonal.

Diagonal, Triangular and Sparse Matrices: special cases of square matrices

- If $A_{ij} = 0$ for all $i \neq j$ then \mathbf{A} is diagonal.
- If $A_{ij} = 0$ for all $i > j$ then \mathbf{A} is upper-triangular.
- If $A_{ij} = 0$ for all $i < j$ then \mathbf{A} is lower-triangular.
- If $A_{ij} = 0$ for 'most' $i \neq j$ then \mathbf{A} is sparse (informal definition but computationally useful)

Orthogonal Matrices:

- If $\mathbf{A}\mathbf{A}^T = \mathbf{A}^T\mathbf{A} = \mathbf{I}$ then \mathbf{A} is orthogonal (orthonormal), $\mathbf{A} = \mathbf{A}^T = \mathbf{A}^{-1}$, the rows and columns of \mathbf{A} are orthonormal vectors.

Full-Rank Matrices:

- If $\text{rank}(\mathbf{A}) = \min(\text{dim}(\mathbf{A}))$ then \mathbf{A} is full-rank.

Defective Matrices:

- If \mathbf{A} does not have a full set of eigenvectors then \mathbf{A} is defective (non-diagonalisable). A defective $n \times n$ matrix \mathbf{A} has at least one eigenvalue with algebraic multiplicity $m > 1$, with less than m associated eigenvectors.

Idempotent Matrices: for square matrices \mathbf{A} ,

- If $\mathbf{A}^2 = \mathbf{A}$ then \mathbf{A} is idempotent.
- All $\mathbf{A} = \mathbf{I}_n$ are idempotent. If $\mathbf{A} \neq \mathbf{I}_n$ is idempotent, then \mathbf{A} is singular.

Positive Definite and Positive Semi-Definite Matrices: for symmetric matrices $\mathbf{A} = \mathbf{A}^T$,

- If $\mathbf{x}^T \mathbf{A} \mathbf{x} > 0$ for all $\mathbf{x} \neq \mathbf{0}$ then \mathbf{A} is positive definite ($\mathbf{A} > 0$), and \mathbf{A} has all positive eigenvalues.
- If $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$ for all \mathbf{x} then \mathbf{A} is positive semi-definite ($\mathbf{A} \geq 0$), and \mathbf{A} has all nonnegative eigenvalues.

Row Echelon Form and Reduced Row Echelon Form:

- \mathbf{A} is in row echelon form if 1) all rows having only zero entries are at the bottom and 2) the leading entry (the left-most nonzero entry) of every nonzero row (the pivot) is on the right of the leading entry of every row above.
- \mathbf{A} is in reduced row echelon form (RREF) if 1) \mathbf{A} is in row echelon form, 2) the leading entry in each nonzero row is 1 (a 'leading one') and 3) each column containing a leading 1 has zeros in all its other entries.

4.1.2. Properties of Complex Matrices

Normal and Unitary Matrices:

- If $\mathbf{A} \mathbf{A}^* = \mathbf{A}^* \mathbf{A}$ then \mathbf{A} is normal.
- If $\mathbf{A} \mathbf{A}^* = \mathbf{A}^* \mathbf{A} = \mathbf{I}$ then \mathbf{A} is unitary and $\mathbf{A} = \mathbf{A}^* = \mathbf{A}^{-1}$. The matrix can be written as $\mathbf{A} = \exp(i\mathbf{H})$ where \mathbf{H} is a Hermitian matrix, or diagonalised to $\mathbf{A} = \mathbf{U} \mathbf{D} \mathbf{U}^*$ where \mathbf{U} is unitary and \mathbf{D} is diagonal and unitary.

Hermitian and Anti-Hermitian Matrices:

- If $\mathbf{A} = \mathbf{A}^*$ ($\mathbf{A} = \mathbf{A}^H$) ($A_{ij} = \overline{A_{ji}}$) then \mathbf{A} is Hermitian (self-adjoint). $\mathbf{x}^* \mathbf{A} \mathbf{x}$ is real for all complex vectors \mathbf{x} . \mathbf{A} has spectral decomposition $\mathbf{A} = \mathbf{U} \mathbf{D} \mathbf{U}^*$ where \mathbf{U} is unitary and \mathbf{D} is diagonal. $|\mathbf{A}|$ is real.
- If $\mathbf{A} = -\mathbf{A}^*$ ($\mathbf{A} = -\mathbf{A}^H$) then \mathbf{A} is anti-Hermitian (skew-Hermitian) ($A_{ij} = -\overline{A_{ji}}$). The entries of \mathbf{A} on the leading diagonal have no real part.

For matrix decompositions, see Section 4.3.

4.1.2. Simple Properties of Matrix Operations

Matrix Multiplication:

- $\mathbf{ABC} = (\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$ (grouping / associative)
- $\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{AB} + \mathbf{AC}$ (factorisation / distributive over addition)
- $\mathbf{AB} \neq \mathbf{BA}$ (**not** commutative in general)

Matrix Algebra: identical to scalar algebra, **except:**

- multiplication is not commutative e.g. $(\mathbf{A} + \mathbf{B})^2 = \mathbf{A}^2 + \mathbf{AB} + \mathbf{BA} + \mathbf{B}^2 \neq \mathbf{A}^2 + 2\mathbf{AB} + \mathbf{B}^2$
- vectors squared become $\mathbf{x}^T \mathbf{x}$
- divisions become matrix inverses \mathbf{A}^{-1}

Determinants, Transpose, Conjugate Transpose, Inverse (if it exists) and Trace:

- $\text{tr}(\mathbf{A})$ is the sum of all leading diagonal elements of \mathbf{A} .
- \mathbf{A}^* is the conjugate transpose: element-wise complex conjugate **and** transpose

$$|\mathbf{A}^T| = |\mathbf{A}| \quad |\mathbf{A}^*| = |\mathbf{A}|^* \quad |\mathbf{A}^{-1}| = |\mathbf{A}|^{-1}$$

$$|\mathbf{AB}| = |\mathbf{A}||\mathbf{B}| \quad (\mathbf{AB})^T = \mathbf{B}^T \mathbf{A}^T \quad (\mathbf{AB})^* = \mathbf{B}^* \mathbf{A}^* \quad (\mathbf{AB})^{-1} = \mathbf{B}^{-1} \mathbf{A}^{-1}$$

$$(\mathbf{A}^T)^T = (\mathbf{A}^*)^* = (\mathbf{A}^{-1})^{-1} = \mathbf{A}$$

$$\text{tr}(a\mathbf{A} + b\mathbf{B}) = a \text{tr}(\mathbf{A}) + b \text{tr}(\mathbf{B}) \quad \text{tr}(\mathbf{AB}) = \text{tr}(\mathbf{BA}) \quad \text{tr}(\mathbf{ab}^T) = \mathbf{b}^T \mathbf{a}$$

(\mathbf{a} and \mathbf{b} are column vectors of equal length)

4.1.3. Determinant of a Matrix

The determinant of a matrix \mathbf{A} is written $|\mathbf{A}|$ or $\det \mathbf{A}$.

For a 2×2 matrix:
$$\begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc$$

For a 3×3 matrix:
$$\begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = a(ei - fh) - b(di - fg) + c(dh - eg)$$

In general, the determinant of a larger square matrix is the alternating sum of products of entries along any row or column with the determinant of the matrix formed by the remaining entries (not in that row or column). **Any** row or column can be chosen: in the above 3×3 expression, the top row was used.

Properties of Determinants:

1. Invariance under transpose and conjugation: $|\mathbf{A}^T| = |\mathbf{A}|$ and $|\mathbf{A}^*| = |\mathbf{A}|^*$
2. Determinant of an inverse: $|\mathbf{A}^{-1}| = |\mathbf{A}|^{-1}$
3. If \mathbf{A} is either diagonal or upper/lower triangular, then $|\mathbf{A}|$ is the product of the diagonals.
4. When a single row or column is multiplied by a , the determinant is multiplied by a .
5. When a matrix is multiplied by a , the determinant is multiplied by a^n .
6. Replacing a row/column with itself plus another column does not change the determinant.
7. Swapping two rows/columns multiplies the determinant by -1 .
8. If an entire row or column is zero, the determinant is zero.
9. The determinant is distributive over multiplication: $|\mathbf{A}\mathbf{B}| = |\mathbf{A}||\mathbf{B}|$.
10. If one row or column is written as a sum, then the determinant can be written as the sum of two determinants:

$$\begin{vmatrix} a_{11} + \alpha & a_{12} & a_{13} \\ a_{21} + \beta & a_{22} & a_{23} \\ a_{31} + \gamma & a_{32} & a_{33} \end{vmatrix} = \begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} + \begin{vmatrix} \alpha & a_{12} & a_{13} \\ \beta & a_{22} & a_{23} \\ \gamma & a_{32} & a_{33} \end{vmatrix}$$

Properties 3-7 are often used to help factorise algebraic determinants, exploiting the circular permutation symmetry of the symbols, for example:

$$\begin{aligned} \begin{vmatrix} 1 & a^2 & bc \\ 1 & b^2 & ca \\ 1 & c^2 & ab \end{vmatrix} &= \begin{vmatrix} 1 & a^2 & bc \\ 0 & b^2 - a^2 & ca - bc \\ 0 & c^2 - a^2 & ab - bc \end{vmatrix} = \begin{vmatrix} b^2 - a^2 & c(a - b) \\ c^2 - a^2 & b(a - c) \end{vmatrix} = (b-a)(c-a) \begin{vmatrix} b+a & -c \\ c+a & -b \end{vmatrix} \\ &= (b-a)(c-a) \begin{vmatrix} b+a & -c \\ c-b & c-b \end{vmatrix} = (b-a)(c-a)(c-b) \begin{vmatrix} b+a & -c \\ 1 & 1 \end{vmatrix} = (b-a)(c-a)(c-b)(a+b+c) \end{aligned}$$

4.1.4. Inverse Matrix, Transpose and Conjugate Transpose

Inverse of a 2×2 matrix: $\mathbf{A} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \Rightarrow \mathbf{A}^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$

Inverse of a 3×3 matrix: $\mathbf{A}^{-1} = \frac{1}{|\mathbf{A}|} \mathbf{C}^T$

where \mathbf{C} is the matrix of cofactors, which is the determinant of the 2×2 matrix formed by the elements not in the corresponding row or column, multiplied by an alternating sign of +1 or -1 (starting with +1 in the top left).

Transpose: $\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}^T = \begin{bmatrix} a & d & g \\ b & e & h \\ c & f & i \end{bmatrix} \Leftrightarrow (\mathbf{A}^T)_{ij} = \mathbf{A}_{ji}$

Conjugate Transpose (Hermitian transpose):

$$\begin{bmatrix} \alpha + \beta i & \gamma + \delta i \\ \varepsilon + \zeta i & \eta + \kappa i \end{bmatrix}^* = \begin{bmatrix} \alpha - \beta i & \varepsilon - \zeta i \\ \gamma - \delta i & \eta - \kappa i \end{bmatrix} \Leftrightarrow (\mathbf{A}^*)_{ij} = \mathbf{A}_{ji}^*$$

Results for complex matrices typically use \mathbf{A}^* instead of \mathbf{A}^T . For real matrices, $\mathbf{A}^* = \mathbf{A}^T$.

4.1.5. Outer Product

Outer product: $\mathbf{a} \otimes \mathbf{b} = \mathbf{a}\mathbf{b}^T$ so that the elements are $(\mathbf{a} \otimes \mathbf{b})_{ij} = a_i b_j$.

Self outer product: If \mathbf{n} is a column vector, then $\mathbf{n} \otimes \mathbf{n} = \mathbf{n}\mathbf{n}^T$ is a square, symmetric matrix.

$$\mathbf{n}\mathbf{n}^T = \begin{bmatrix} n_1 \\ n_2 \\ n_3 \end{bmatrix} \begin{bmatrix} n_1 & n_2 & n_3 \end{bmatrix} = \begin{bmatrix} n_1^2 & n_1 n_2 & n_1 n_3 \\ n_1 n_2 & n_2^2 & n_2 n_3 \\ n_1 n_3 & n_2 n_3 & n_3^2 \end{bmatrix}$$

4.1.6. Cross Product Matrix

The cross product matrix of a vector \mathbf{n} is denoted $[\mathbf{n}]_{\times}$ and has the property $[\mathbf{n}]_{\times} \mathbf{a} = \mathbf{n} \times \mathbf{a}$ for any vector $\mathbf{a} \in \mathbb{R}^3$. i.e. it represents a cross product as a linear transformation. It is defined as

$$[\mathbf{n}]_{\times} = \begin{bmatrix} 0 & -n_3 & n_2 \\ n_3 & 0 & -n_1 \\ -n_2 & n_1 & 0 \end{bmatrix}$$

The cross product matrix $[\mathbf{n}]_{\times}$ is always skew-symmetric. The columns of $[\mathbf{n}]_{\times}$ are the cross products of \mathbf{n} with each unit basis vector $\{\mathbf{i}, \mathbf{j}, \mathbf{k}\}$.

4.1.7. Eigenvalues and Eigenvectors

Definitions: for a square $n \times n$ matrix \mathbf{A} , the eigenvalues are denoted λ and eigenvectors are denoted \mathbf{u} .

- Definition: $\mathbf{A}\mathbf{u} = \lambda\mathbf{u} \leftrightarrow (\mathbf{A} - \lambda\mathbf{I})\mathbf{u} = \mathbf{0} \leftrightarrow |\mathbf{A} - \lambda\mathbf{I}| = 0$ (characteristic equation, degree n)
- Non-defective matrix: every λ (including repeated) has a unique corresponding \mathbf{u} .
- Algebraic multiplicity a : the exponent of the factor $(\lambda - e)^a$ in the characteristic polynomial.
- Geometric multiplicity g : the number of linearly independent eigenvectors associated i.e. the dimension of the null space of $\mathbf{A} - e\mathbf{I}$.
- Eigenbasis / eigenspace: the space spanned by the eigenvectors. If \mathbf{A} is symmetric and the eigenvectors are normalised, then the eigenbasis is an orthonormal basis.

Geometric interpretation: $\mathbf{A} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is viewed as a linear transformation (Sections 4.2.2. and 4.3).

- The eigenvectors point along the invariant lines under \mathbf{A} . These vectors are not rotated (only scaled).
- If $\lambda = 1$ then the corresponding line is a line of invariant points under \mathbf{A} .
- A repeated eigenvalue represents an invariant plane under \mathbf{A} (spanned by the eigenvectors).

Relationships Between Eigenvalues and Eigenvectors:

- **Sums and products of eigenvalues:** $\lambda_1 + \lambda_2 + \dots + \lambda_n = \text{tr}(\mathbf{A})$ and $\lambda_1\lambda_2\dots\lambda_n = \det(\mathbf{A})$.
Sum of $(n-1)$ -permutations: $\lambda_2\lambda_3\dots\lambda_n + \lambda_1\lambda_3\dots\lambda_n + \lambda_1\lambda_2\dots\lambda_{n-1} = \lambda_1\lambda_2\lambda_3\dots\lambda_n (\lambda_1^{-1} + \lambda_2^{-1} + \dots + \lambda_n^{-1}) = \text{tr}(\mathbf{C})$
where \mathbf{C} is the matrix of cofactors of \mathbf{A} . The Faddeev-LeVerrier algorithm computes the coefficients of the characteristic polynomial via Vieta's formulas by applying Newton's identities to these expressions.
- **Linearity:** eigenvalues of $a\mathbf{A}$ are $a\lambda$; eigenvalues of $\mathbf{A} + a\mathbf{I}$ are $\lambda + a$; eigenvalues of \mathbf{A}^n are λ^n (including $n = -1$ as the inverse).
- **Commutative matrices:** if \mathbf{A} and \mathbf{B} have the same eigenvectors, then $\mathbf{AB} = \mathbf{BA}$.
- **Polynomial of a matrix:** for any polynomial $f(x)$, the eigenvalues of $f(\mathbf{A})$ are $f(\lambda)$ (with the constant term taken to be a multiple of \mathbf{I}).
- **Cayley-Hamilton theorem:** a matrix \mathbf{A} satisfies its own characteristic equation.
- **Rayleigh's quotient:** if \mathbf{A} is Hermitian, then the quantity $C = \frac{\mathbf{x}^\top \mathbf{A} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} = \hat{\mathbf{x}} \cdot \mathbf{A} \hat{\mathbf{x}} = \hat{\mathbf{x}}^\top \mathbf{A} \hat{\mathbf{x}}$ is bounded by $\lambda_1 \leq C \leq \lambda_n$, where λ_1 and λ_n are the smallest and largest eigenvalues of \mathbf{A} respectively, and \mathbf{x} is **any** vector. Also, if $\mathbf{x} \approx \mathbf{u}$ then $C \approx \lambda$, the (approximate) corresponding eigenvalue to eigenvector \mathbf{u} of \mathbf{A} .
- **Spectral radius:** the smallest circle in the complex plane containing all eigenvalues, $\rho(\mathbf{A}) = \max |\lambda_i|$.
For any integer $k \geq 1$ and norm (Section 4.1.9), $\rho(\mathbf{A}) \leq \|\mathbf{A}^k\|^{1/k}$, with equality in the limit of $k \rightarrow \infty$.
- **Gershgorin circle theorem:** every complex eigenvalue of \mathbf{A} lies within the union of circles in the complex plane, centred at the diagonal entries of \mathbf{A} , with radii given by the sum of the magnitudes of all off-diagonal entries in that row i.e. $\lambda \in \{z : \bigcup_{i=1}^n |z - A_{ii}| \leq \sum_{j \neq i} |A_{ij}|\}$.
- **Singular values:** the singular values σ of \mathbf{A} are the square roots of the eigenvalues of $\mathbf{A}\mathbf{A}^\top$ or $\mathbf{A}^\top\mathbf{A}$.

4.1.8. Approximations for Eigenvalues and Eigenvectors

Shifted Inverse Power Method

To approximate an eigenvalue of \mathbf{A} :

- Choose an initial approximation α to the target eigenvalue of \mathbf{A} .
- Calculate the matrix $\mathbf{B} = \mathbf{A} - \alpha\mathbf{I}$.
- Choose any initial normalised vector \mathbf{r}_0 .
- Iterate: starting with $n = 0$ and incrementing,
 - Solve the system $\mathbf{B} \mathbf{r}_{n+1} = \mathbf{r}_n$, typically using efficient LU decomposition.
 - Calculate $\mu_{n+1} = \mathbf{r}_{n+1} \cdot \mathbf{r}_n$ as an estimate for the shifted eigenvalue.
 - Normalise \mathbf{r}_{n+1} (in-place) and continue.
- Once sufficient convergence is achieved, calculate $\lambda = \frac{1}{\mu_\infty} + \alpha$ as the target eigenvalue of \mathbf{A} , where μ_∞ is the limiting (converged) value of sequence μ_n . The unit vector \mathbf{r}_n has simultaneously converged to the corresponding eigenvector.

Note that if \mathbf{r}_0 happens to be chosen as a different eigenvector of \mathbf{A} , the method will not converge.

An initial estimate for α can often be found using the Gershgorin Circle Theorem (see Section 4.1.7.), which works best when the matrix is sparse or near-diagonal (most off-diagonal entries are zero or much smaller than the diagonal entries).

Rayleigh's quotient (Section 4.1.7) can also be used to approximate the eigenvalue from an eigenvector approximation, and bounds the eigenvalues.

4.1.9. Normed Vector Spaces and Matrix Norms

Norms are scalar-valued measures of the ‘typical size’ of an object (vector/function/matrix).

Vector Norms: if \mathbf{x} is a vector with elements x_i , four common norms are

- 1-norm: $\|\mathbf{x}\|_1 = \sum_i |x_i|$ (L_1 norm / l_1 norm / Manhattan norm / Taxicab norm)
- 2-norm: $\|\mathbf{x}\|_2 = \sqrt{\sum_i x_i^2}$ (L_2 norm / l_2 norm / quadratic norm / Euclidean norm)
Inner product form: $\|\mathbf{x}\|_2^2 = \mathbf{x} \cdot \mathbf{x}$ (in a Hilbert space)
- p -norm: $\|\mathbf{x}\|_p = \left(\sum_i |x_i|^p \right)^{1/p}$ (for $p \geq 1$; l^p norm / Lebesgue norm)
- ∞ -norm: $\|\mathbf{x}\|_\infty = \max \{x_i\}$ (infinity norm / supremum norm)

Generalisation to Infinite Dimensions (Banach Spaces)

Discrete or continuous functions $f: \mathbb{R} \rightarrow \mathbb{R}$ can be considered infinite-dimensional vectors.

- l_p norm: $\|x\|_p = \left(\sum_i |x_i|^p \right)^{1/p}$
- L_p norm: $\|f\|_p = \left(\int_\Omega |f(x)|^p dx \right)^{1/p}$ (f defined in a Lebesgue space $L_p(\Omega)$)

Matrix Norms: if \mathbf{A} is a matrix with elements a_{ij} , four common norms are

- 1-norm: the maximum absolute column sum, $\|\mathbf{A}\|_1 = \max_j \sum_i |a_{ij}|$
- Infinity-norm: the maximum absolute row sum, $\|\mathbf{A}\|_\infty = \max_i \sum_j |a_{ij}|$
- Euclidean norm: $\|\mathbf{A}\|_E = \sqrt{\sum_i \sum_j |a_{ij}|^2}$ (Frobenius norm)
- 2-norm: the largest singular value of \mathbf{A} . (spectral norm)

The condition number of an invertible square matrix \mathbf{A} is defined by $\kappa = \|\mathbf{A}\| \|\mathbf{A}^{-1}\|$ evaluated using one of these norms (the same one in both places. If the 2-norm is used, the condition number is the ratio of the largest to the smallest singular value of \mathbf{A}).

4.2. Transformation Matrices

4.2.1. Rotations, Reflections, Shears and Projections as Transformation Matrices

Rotation Matrices: for a counterclockwise rotation about the origin, angle θ ,

$$\begin{array}{cccc}
 \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} & \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix} & \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix} & \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \\
 \text{2D} & \text{3D, about } x\text{-axis} & \text{3D, about } y\text{-axis} & \text{3D, about } z\text{-axis}
 \end{array}$$

The general rotation about line $\mathbf{r} \times \mathbf{n} = \mathbf{0}$ is given by $(\cos \theta) \mathbf{I} + (\sin \theta) [\mathbf{n}]_{\times} + (1 - \cos \theta) (\mathbf{n} \otimes \mathbf{n})$.

In 3D, the rotation is considered counterclockwise when viewed inwards from the positive axis, towards the origin. The eigenvalues are $\{1, e^{i\theta}, e^{-i\theta}\}$. The real eigenvector corresponding to $\lambda = 1$ is the axis of rotation. The trace is $1 + 2 \cos \theta$. The determinant is 1.

Reflection Matrices:

In 2D, the reflection in the line $y = (\tan \theta)x$ is given by $\begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix}$. Determinant is -1.

In 3D, the reflection in the plane $\mathbf{r} \cdot \mathbf{n} = 0$ is given by $\mathbf{I} - 2\mathbf{nn}^T = \begin{bmatrix} 1 - 2n_1^2 & -2n_1n_2 & -2n_1n_3 \\ -2n_1n_2 & 1 - 2n_2^2 & -2n_2n_3 \\ -2n_1n_3 & -2n_2n_3 & 1 - 2n_3^2 \end{bmatrix}$
 where \mathbf{n} is the **unit** normal vector of the plane.

Shear Matrices: shear factor λ is given by $\mathbf{I} + \lambda \mathbf{n} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$: angle of shear: $\tan \gamma = \lambda$

$$\begin{array}{ccc}
 \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 0 \\ \lambda & 1 \end{bmatrix} & \begin{bmatrix} 1 + \lambda n_1 & \lambda n_1 & \lambda n_1 \\ \lambda n_2 & 1 + \lambda n_2 & \lambda n_2 \\ \lambda n_3 & \lambda n_3 & 1 + \lambda n_3 \end{bmatrix} \\
 \text{in } x\text{-axis} & \text{in } y\text{-axis} & \text{general 3D shear in line } \mathbf{r} \times \mathbf{n} = \mathbf{0}
 \end{array}$$

Orthogonal Projection: the projection matrix of \mathbb{R}^3 onto the plane $\mathbf{r} \cdot \mathbf{n} = 0$, where \mathbf{n} is the **unit** normal vector of the plane, is

$$\mathbf{I} - \mathbf{nn}^T = \begin{bmatrix} 1 - n_1^2 & -n_1n_2 & -n_1n_3 \\ -n_1n_2 & 1 - n_2^2 & -n_2n_3 \\ -n_1n_3 & -n_2n_3 & 1 - n_3^2 \end{bmatrix}.$$

This is a singular matrix since it involves a reduction in dimensionality ($3 \rightarrow 2$).

4.2.2. Geometric Interpretation of Invariance, Eigenvalues and Eigenvectors

For all linear transformations, the origin is invariant, which is not considered to 'count' here.

In 2D (or higher):

Invariant point: if $\mathbf{A}\mathbf{p} = \mathbf{p}$ then \mathbf{p} is invariant under \mathbf{A} (\mathbf{p} does not move).

\mathbf{p} is an eigenvector of \mathbf{A} , with eigenvalue 1.

Line of invariant points: a line $\mathbf{r} = \lambda\mathbf{n}$ for which all points on the line are invariant.

\mathbf{n} is an eigenvector of \mathbf{A} with eigenvalue 1.

Invariant line: a line $\mathbf{r} = \lambda\mathbf{n}$ for which points on the line are mapped to another point on the line, so that the line as a whole does not move.

\mathbf{n} is an eigenvector of \mathbf{A} , with points being scaled by the eigenvalue.

In 3D (or higher):

Plane of invariant points: a plane $\mathbf{r} = \lambda\mathbf{u} + \mu\mathbf{v}$ for which all points on the plane are invariant.

\mathbf{u} and \mathbf{v} are eigenvectors with repeated eigenvalue 1.

Invariant plane: a line $\mathbf{r} = \lambda\mathbf{u} + \mu\mathbf{v}$ for which points on the plane are mapped to another point on the plane, so that the plane as a whole does not move.

\mathbf{u} and \mathbf{v} are eigenvectors with the same repeated eigenvalue.

4.2.3. Affine Transformations

An affine transformation represents the combination of a linear transformation followed by a translation in space. In 3D, they can be represented as 4×4 matrices, with space vectors taking the form $[x \ y \ z \ 1]^T$, known as homogeneous coordinates (with $w = 1$, WLOG).

A 3D affine transformation matrix \mathbf{R} has the form

$$\mathbf{R} = \begin{bmatrix} M_{ii} & M_{ji} & M_{ki} & \Delta x \\ M_{ij} & M_{jj} & M_{kj} & \Delta y \\ M_{ik} & M_{jk} & M_{kk} & \Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & \Delta x \\ 0 & 1 & 0 & \Delta y \\ 0 & 0 & 1 & \Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{translations}} \underbrace{\begin{bmatrix} M_{ii} & M_{ji} & M_{ki} & 0 \\ M_{ij} & M_{jj} & M_{kj} & 0 \\ M_{ik} & M_{jk} & M_{kk} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{rotations, shears, projections...}}$$

where M_{xx} are the entries of a 3×3 linear transformation matrix \mathbf{M} and $(\Delta x, \Delta y, \Delta z)$ are translations parallel to the coordinate axes.

4.3. Fundamental Subspaces and Matrix Decompositions

4.3.1. Fundamental Subspaces

For any $m \times n$ matrix A of rank r , with $y = Ax$, (dims = number of dimensions in the space)

	Subspace	Form	Dims	Basis	Projection matrix onto subspace
Input space (dim = n)	Row space (domain)	$C(A^T)$	r_{col}	nonzero rows of $rref(A)$	$A^T(AA^T)^{-1}A$
	Null space (kernel)	$N(A)$	$n - r_{col}$	$Ax = 0$ i.e. $x^T A = 0^T$	$I - A^T(AA^T)^{-1}A$
Output space (dim = m)	Column space (image; range)	$C(A)$	r_{row}	columns of A corresponding to columns of $rref(A)$ with leading 1's	$A(A^T A)^{-1}A^T$
	Left null space (cokernel)	$N(A^T)$	$m - r_{col}$	$y^T A = 0^T$ i.e. $A^T y = 0$	$I - A(A^T A)^{-1}A^T$

For any matrix A acting as a transformation of space R^n into space R^m ,

- Vectors in the row space are mapped to the column space.
- Vectors in the null space are mapped to the origin.
- No vector is mapped to the left null space.
- The eigenvectors span the column space.

Orthogonal complements:

- The row space and the null space are orthogonal.
- The column space and the left null space are orthogonal.

Ranks:

- The column rank is the dimension of the row space. Full column rank: $r = n$.
- The row rank is the dimension of the column space. Full row rank: $r = m$.
- If A has full row rank then a solution x to the system $Ax = b$ exists.
- If A has full rank (i.e. square, invertible) then the solution x to $Ax = b$ is unique.

4.3.2. The Gram-Schmidt Orthonormalisation Process

Given a set of vectors a_i , the Gram-Schmidt process gives a set of orthonormal vectors q_i which span the same space as a_i , by subtracting off components parallel to each q_i .

Taking normalised a_1 as q_1 , the vector q_i is the normalised vector of $q_i' = a_i - \sum_{k=1}^{i-1} (a_i \cdot q_k) q_k$.

The complete set q_i is obtained when the resulting vector is 0 .

4.3.3. LU Decomposition

For any $m \times n$ matrix \mathbf{A} ,

$$\mathbf{PA} = \mathbf{LU}$$

where \mathbf{P} is a permutation matrix (sometimes omitted), \mathbf{L} a square $m \times m$ lower triangular matrix and \mathbf{U} an $m \times n$ echelon matrix.

Manual computation:

- To compute \mathbf{U} , use Gaussian elimination to convert \mathbf{A} into row echelon form, using row operations of the form $\mathbf{r}_j' = \mathbf{r}_j - a \mathbf{r}_i$ where $j > i$. If using partial pivoting, swap rows **before** each elimination such that the pivot (diagonal entry of \mathbf{U}) is maximised (by magnitude), equivalently ensuring that the multiplier (entry in \mathbf{L}) has $|a| \leq 1$.
- To compute \mathbf{L} , let each lower-triangular entry be equal to the coefficient a used in the Gaussian elimination step to form a zero in the corresponding position in \mathbf{U} . The diagonal entries of \mathbf{L} are all 1 and the upper-triangular entries are all 0. If using partial pivoting, swap entries in \mathbf{L} **only** when they are below the diagonal of the column, and only in the column in which the swap was performed.
- To compute \mathbf{P} , start with \mathbf{I} and swap the same rows as done during the process of computing \mathbf{U} . If partial pivoting was **not** used, then $\mathbf{P} = \mathbf{I}$ and can be omitted.

Programming functions:

- MATLAB: `[L, U, P] = lu(A)`
- Python: `P, L, U = scipy.linalg.lu(A)` where \mathbf{A} is a NumPy array

Basis of subspaces:

- Column space: the columns of \mathbf{L} corresponding to nonzero rows of \mathbf{U}
- Left nullspace: the nullspace of \mathbf{A}^T
- Row space: the nonzero rows of \mathbf{U}
- Nullspace: the nullspace of \mathbf{U}

To solve a system of equations of the form $\mathbf{Ax} = \mathbf{b}$:

- Transform $\mathbf{Ax} = \mathbf{b}$ into $\mathbf{Ux} = \mathbf{c}$ where $\mathbf{Lc} = \mathbf{b}$ can be solved by forward-substitution.
- Set all free variables to zero and find a particular solution \mathbf{x}_0 .
- Set the RHS to zero, give each free variable in turn the value 1 while the others are zero, and solve to find a set of vectors which span the null space of \mathbf{A} : $\mathbf{n}_1, \mathbf{n}_2$, etc.
- The general solution is $\mathbf{x} = \mathbf{x}_0 + \lambda \mathbf{n}_1 + \mu \mathbf{n}_2 + \dots$, where λ, μ , etc. are arbitrary.

4.3.4. Cholesky Decomposition

For a Hermitian (if real, then symmetric), positive-definite matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{L}\mathbf{L}^* \quad \text{in general, or} \quad \mathbf{A} = \mathbf{L}\mathbf{L}^\top \quad \text{for real matrices}$$

where \mathbf{L} is a lower triangular matrix. \mathbf{L}^* is the conjugate transpose.

Manual computation:

- For each row of \mathbf{L} , compute the diagonal entry using this first formula below, starting with the top left.
- Then compute the remaining entries on that row using the second formula below, starting from the left and moving right until the diagonal. All other entries are 0.
- Move down to the next row, repeating until all rows of \mathbf{L} are filled.

$$L_{j,j} = (\pm) \sqrt{A_{j,j} - \sum_{k=1}^{j-1} L_{j,k}^2},$$

$$L_{i,j} = \frac{1}{L_{j,j}} \left(A_{i,j} - \sum_{k=1}^{j-1} L_{i,k} L_{j,k} \right) \quad \text{for } i > j.$$

for real matrices

$$L_{j,j} = \sqrt{A_{j,j} - \sum_{k=1}^{j-1} L_{j,k} L_{j,k}^*},$$

$$L_{i,j} = \frac{1}{L_{j,j}} \left(A_{i,j} - \sum_{k=1}^{j-1} L_{i,k} L_{j,k}^* \right) \quad \text{for } i > j.$$

for complex matrices

Programming functions:

- MATLAB: $\mathbf{L} = \text{chol}(\mathbf{A})$
- Python: $\mathbf{L} = \text{scipy.linalg.cholesky}(\mathbf{A})$ where \mathbf{A} is a NumPy array

To solve a system of equations of the form $\mathbf{Ax} = \mathbf{b}$:

- Transform $\mathbf{Ax} = \mathbf{b}$ into $\mathbf{L}^*\mathbf{x} = \mathbf{c}$ where $\mathbf{Lc} = \mathbf{b}$ can be solved by forward-substitution.
- Solve $\mathbf{L}^*\mathbf{x} = \mathbf{c}$ by back-substitution.

This method is approximately twice as fast as a solution using LU decomposition.

4.3.5. QR Decomposition and Least Squares Fitting

For any $m \times n$ matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{QR}$$

where \mathbf{Q} is an orthonormal $m \times r$ matrix and \mathbf{R} is an invertible upper-triangular $r \times n$ matrix, where r is the rank of \mathbf{A} .

Manual computation:

- To compute \mathbf{Q} , use the Gram-Schmidt process (Section 4.3.2.) to find an orthonormal set from the columns of \mathbf{A} . These vectors form the columns of \mathbf{Q} .
- To compute \mathbf{R} , use

$$\mathbf{R}_{ij} = \mathbf{q}_i \cdot \mathbf{a}_j \quad \text{if } i \leq j; \quad \text{otherwise } \mathbf{R}_{ij} = 0$$

where \mathbf{q}_i is the i th column vector of \mathbf{Q} and \mathbf{a}_j is the j th column vector of \mathbf{A} .

Programming functions:

- MATLAB: `[Q, R] = qr(A)`
- Python: `Q, R = scipy.linalg.qr(A)` where \mathbf{A} is a NumPy array

To solve a least-squares system $\mathbf{Ax}^* = \mathbf{b}$: in general, $\mathbf{x}^* = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$, or more efficiently,

- Find $\mathbf{A} = \mathbf{QR}$.
- Solve $\mathbf{Rx}^* = \mathbf{Q}^T \mathbf{b}$ by back-substitution.
- The solution $\mathbf{x} = \mathbf{x}^*$ satisfies $\min \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{Ax} - \mathbf{b}\|_2^2\} = \min \{\mathbf{x} \in \mathbb{R}^n : \sum_i |\mathbf{A}_i \cdot \mathbf{x} - b_i|^2\}$.

Least-squares system in tabulated form: mapping x_i to y_i ($1 \leq i \leq n$), the coefficients of a best-fit line $y = a + bx$ or parabola $y = a + bx + cx^2$ can be found by solving the systems

$$\begin{array}{l} \text{Straight line: } y = a + bx \\ \text{Quadratic: } y = a + bx + cx^2 \end{array} \quad \left\{ \begin{array}{l} an + b \sum_i x_i = \sum_i y_i \\ a \sum_i x_i + b \sum_i x_i^2 = \sum_i x_i y_i \\ \\ an + b \sum_i x_i + c \sum_i x_i^2 = \sum_i y_i \\ a \sum_i x_i + b \sum_i x_i^2 + c \sum_i x_i^3 = \sum_i x_i y_i \\ a \sum_i x_i^2 + b \sum_i x_i^3 + c \sum_i x_i^4 = \sum_i x_i^2 y_i \end{array} \right.$$

4.3.6. Eigendecomposition (Diagonalisation)

If \mathbf{A} is a square $n \times n$ matrix with n linearly independent eigenvectors (diagonalisable, nondefective), then

$$\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^{-1} \quad \text{and} \quad \mathbf{A}^n = \mathbf{S}\mathbf{\Lambda}^n\mathbf{S}^{-1}$$

where \mathbf{S} has the eigenvectors of \mathbf{A} as its columns, and $\mathbf{\Lambda}$ is a diagonal matrix with the corresponding eigenvalues along the diagonal. Matrices \mathbf{A} and $\mathbf{\Lambda}$ are said to be similar, so that they represent the same geometric transformation with a change of basis given by \mathbf{S} .

Spectral theorem: If \mathbf{A} is Hermitian then $\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^*$

Spectral decomposition: If \mathbf{A} is real and symmetric then $\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^T$.

where \mathbf{S} the orthonormal matrix of **normalised** eigenvectors as its columns.

Programming functions (where $\mathbf{D} \leftrightarrow \mathbf{\Lambda}$)

- MATLAB: `[S, D] = eig(A)`.
- Python: `D, S = np.linalg.eig(A)` # can use `eigh(A)` if Hermitian

4.3.7. Singular Value Decomposition (SVD)

For any $m \times n$ matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{Q}_1\mathbf{\Sigma}\mathbf{Q}_2^T$$

- \mathbf{Q}_1 is an $m \times m$ orthonormal matrix with the normalised eigenvectors of $\mathbf{A}\mathbf{A}^T$ as its columns.
- \mathbf{Q}_2 is an $n \times n$ orthonormal matrix with the normalised eigenvectors of $\mathbf{A}^T\mathbf{A}$ as its columns.
- $\mathbf{\Sigma}$ is an $m \times n$ diagonal matrix containing the r singular values, arranged in descending order on the leading diagonal, which are the square roots of the non-zero eigenvalues of both $\mathbf{A}\mathbf{A}^T$ and $\mathbf{A}^T\mathbf{A}$, where r is the rank of \mathbf{A} .

Programming functions:

- MATLAB: `[Q1, S, Q2] = svd(A)`
- Python: `Q1, S, Q2_T = np.linalg.svd(A)` where \mathbf{A} is a NumPy array

Basis of subspaces of \mathbf{A} :

- Column space: the first r columns of \mathbf{Q}_1
- Left nullspace: the last $m - r$ columns of \mathbf{Q}_1
- Row space: the first r columns of \mathbf{Q}_2
- Nullspace: the last $n - r$ columns of \mathbf{Q}_2

4.3.8. Other Matrix Decompositions

Schur Decomposition: for any square matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{U}\mathbf{T}\mathbf{U}^*$$

where \mathbf{U} is unitary, and \mathbf{T} is upper-triangular with the eigenvalues of \mathbf{A} along its diagonal. If \mathbf{A} is normal then \mathbf{T} is diagonal and the Schur form matches the spectral decomposition.

Polar Decomposition: for any square matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{U}\mathbf{P}$$

where \mathbf{U} is unitary and \mathbf{P} is positive semi-definite and Hermitian. \mathbf{P} is unique with $\mathbf{P}^2 = \mathbf{A}\mathbf{A}^*$. If \mathbf{P} has spectral decomposition $\mathbf{P} = \mathbf{V}\mathbf{D}\mathbf{V}^*$ and $\mathbf{W} = \mathbf{U}\mathbf{V}$ then the SVD is $\mathbf{A} = \mathbf{W}\mathbf{D}\mathbf{V}^*$.

Toeplitz decomposition: for any square matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{H} + \mathbf{G}$$

where \mathbf{H} is Hermitian and \mathbf{G} is skew-Hermitian, with $\mathbf{H} = (\mathbf{A} + \mathbf{A}^*)/2$ and $\mathbf{G} = (\mathbf{A} - \mathbf{A}^*)/2$.

4.3.9. Matrix Exponentials

For a square $m \times m$ matrix \mathbf{A} , the matrix exponential $\exp(\mathbf{A})$ is defined as

$$\exp(\mathbf{A}) = e^{\mathbf{A}} = \sum_{n=0}^{\infty} \frac{\mathbf{A}^n}{n!} = \mathbf{I} + \mathbf{A} + \frac{1}{2}\mathbf{A}^2 + \dots$$

To compute $\exp(\mathbf{A})$ from its eigenvalues and eigenvectors, let $\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^{-1}$ (eigendecomposition: Section 4.3.6). Then $\exp(\mathbf{A}) = \mathbf{S} \exp(\mathbf{\Lambda}) \mathbf{S}^{-1}$.

Properties:

- $\exp(\mathbf{A}^T) = \exp(\mathbf{A})^T$ and $\exp(\mathbf{A}^*) = \exp(\mathbf{A})^*$
- If $\mathbf{A}\mathbf{B} = \mathbf{B}\mathbf{A}$ then $e^{\mathbf{A}}e^{\mathbf{B}} = e^{\mathbf{B}}e^{\mathbf{A}} = e^{\mathbf{A}+\mathbf{B}}$

4.3.10. Statement of a Convex Optimisation Problem

Objective: to find \mathbf{x} that minimises $f_0(\mathbf{x})$ such that $f_i(\mathbf{x}) \leq b_i$ and $h_i(\mathbf{x}) = 0$, where

- $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$ is an input state vector
- $f_0(\mathbf{x})$ is the objective function (scalar)
- $f_i(\mathbf{x}) \leq b_i$ are bounds for valid regions (feasible region) of the input \mathbf{x}
- $h_i(\mathbf{x}) = 0$ are constraints for valid contour lines (level sets) of the input \mathbf{x}

Convex function: $f_i(\alpha\mathbf{x} + \beta\mathbf{y}) \leq \alpha f_i(\mathbf{x}) + \beta f_i(\mathbf{y})$ for all $\alpha + \beta = 1$, $\alpha \geq 0$, $\beta \geq 0$

In a **convex** optimisation problem, the objective and inequality constraint functions f_0 and f_i are convex functions. Convex optimisation problems are guaranteed to have no more than one local minima: if a minimum exists, then it is the global minimum.

The global minimum solution is denoted $\mathbf{x} = \mathbf{x}^*$, so that $f_0(\mathbf{x}^*) = \min \{\mathbf{x} \in \mathbb{R}^n : f_0(\mathbf{x})\}$ within the specified constraints.

Least-Squares Optimisation

Objective: minimise $f_0(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$ (l_2 -norm of residual vector)

Constraints: (arbitrary)

The solution is given by $\mathbf{x}^* = \mathbf{A}^+\mathbf{b}$ ($\mathbf{A}^+ = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$: Moore-Penrose pseudoinverse).

Linear Programming (LP Problems)

Objective: minimise $f_0(\mathbf{x}) = \mathbf{c}^T\mathbf{x}$ (linear combination of variables)

Constraints: $\mathbf{a}_i^T\mathbf{x} \leq b_i$ and $\mathbf{x} \geq 0$

The solution \mathbf{x}^* is guaranteed to lie on the vertex of the boundary of the constraint region.

Danzig's simplex algorithm: visualise feasible region as a polytope (n -D polyhedron with hyperplane faces) in a uniform gradient field $\nabla f_0 = \mathbf{c}$. Traverse edges of the polytope until optimum vertex found.

Quadratic Programming

Objective: minimise $f_0(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{P}\mathbf{x} + \mathbf{q}^T\mathbf{x} + \mathbf{r}$ (\mathbf{P} : positive semi-definite matrix)

Constraints: $\mathbf{G}\mathbf{x} \leq \mathbf{h}$ and $\mathbf{A}\mathbf{x} = \mathbf{b}$

4.4. Matrix and Tensor Calculus

4.4.1. Indicial Tensor Notation

Einstein summation notation: repeated indices imply summation over indices.

- Component form: $\mathbf{a} = a_i \mathbf{e}_i = \sum_i a_i \mathbf{e}_i$
- Inner product: $\mathbf{a} \cdot \mathbf{b} = a_i b_i$
- Outer product: if $\mathbf{A} = \mathbf{a} \otimes \mathbf{b}$ then $A_{ij} = a_i b_j$
- Kronecker delta: $\delta_{ij} = 1$ if $i = j$ else $\delta_{ij} = 0$; for orthonormal \mathbf{e} , then $\delta_{ij} = \mathbf{e}_i \cdot \mathbf{e}_j$.
- Permutation symbol:
 - $e_{ijk} = 1$, when i, j, k are in cyclic order
 - $e_{ijk} = -1$, when i, j, k are in anti-cyclic order
 - $e_{ijk} = 0$, when any indices repeat
- Triple product contraction: $e_{ijk} e_{ipq} = \delta_{jp} \delta_{kq} - \delta_{jq} \delta_{kp}$ (epsilon-delta identity)
- Gradient field: $\nabla \phi = (\partial \phi / \partial x_i) \mathbf{e}_i = \phi_{,i} \mathbf{e}_i$
- Divergence field: $\nabla \cdot \mathbf{v} = v_{i,i}$
- Curl field: $\nabla \times \mathbf{v} = e_{ijk} v_{k,j} \mathbf{e}_i$
- Gauss' divergence theorem: $\int_V \frac{\partial A_{ij}}{\partial x_j} dV = \oint_S A_{ij} n_j dS$
- Stokes' curl theorem: $\int_S e_{ijk} \frac{\partial A_{pk}}{\partial x_j} n_i dS = \oint_C A_{pk} dx_k$

4.4.2. Differentiation with Respect to Vectors and Matrices

$$\left[\frac{\partial \mathbf{x}}{\partial y} \right]_i = \frac{\partial x_i}{\partial y}$$

Vector-valued function $\mathbf{x}(y)$

$$\left[\frac{\partial x}{\partial \mathbf{y}} \right]_i = \frac{\partial x}{\partial y_i}$$

Multivariable function $x(\mathbf{y})$

$$\left[\frac{\partial \mathbf{x}}{\partial \mathbf{y}} \right]_{ij} = \frac{\partial x_i}{\partial y_j}$$

Vector field $\mathbf{x}(\mathbf{y})$

Differential of matrix expressions (\mathbf{X}, \mathbf{Y} : matrix-valued functions)

$$\begin{aligned} \partial \mathbf{A} &= 0 && (\mathbf{A} \text{ is a constant}) \\ \partial(\alpha \mathbf{X}) &= \alpha \partial \mathbf{X} \\ \partial(\mathbf{X} + \mathbf{Y}) &= \partial \mathbf{X} + \partial \mathbf{Y} \\ \partial(\text{Tr}(\mathbf{X})) &= \text{Tr}(\partial \mathbf{X}) \\ \partial(\mathbf{X}\mathbf{Y}) &= (\partial \mathbf{X})\mathbf{Y} + \mathbf{X}(\partial \mathbf{Y}) \\ \partial(\mathbf{X} \circ \mathbf{Y}) &= (\partial \mathbf{X}) \circ \mathbf{Y} + \mathbf{X} \circ (\partial \mathbf{Y}) \\ \partial(\mathbf{X} \otimes \mathbf{Y}) &= (\partial \mathbf{X}) \otimes \mathbf{Y} + \mathbf{X} \otimes (\partial \mathbf{Y}) \\ \partial(\mathbf{X}^{-1}) &= -\mathbf{X}^{-1}(\partial \mathbf{X})\mathbf{X}^{-1} \\ \partial(\det(\mathbf{X})) &= \text{Tr}(\text{adj}(\mathbf{X})\partial \mathbf{X}) \\ \partial(\det(\mathbf{X})) &= \det(\mathbf{X})\text{Tr}(\mathbf{X}^{-1}\partial \mathbf{X}) \\ \partial(\ln(\det(\mathbf{X}))) &= \text{Tr}(\mathbf{X}^{-1}\partial \mathbf{X}) \\ \partial \mathbf{X}^T &= (\partial \mathbf{X})^T \\ \partial \mathbf{X}^H &= (\partial \mathbf{X})^H \end{aligned}$$

Derivative of determinant: $\frac{\partial \det(\mathbf{Y})}{\partial x} = \det(\mathbf{Y})\text{Tr} \left[\mathbf{Y}^{-1} \frac{\partial \mathbf{Y}}{\partial x} \right]$

Derivative of inverse: $\frac{\partial \mathbf{Y}^{-1}}{\partial x} = -\mathbf{Y}^{-1} \frac{\partial \mathbf{Y}}{\partial x} \mathbf{Y}^{-1}$

Derivative of eigenvalues: $\frac{\partial}{\partial \mathbf{X}} \sum \text{eig}(\mathbf{X}) = \frac{\partial}{\partial \mathbf{X}} \text{Tr}(\mathbf{X}) = \mathbf{I}$
 $\frac{\partial}{\partial \mathbf{X}} \prod \text{eig}(\mathbf{X}) = \frac{\partial}{\partial \mathbf{X}} \det(\mathbf{X}) = \det(\mathbf{X})\mathbf{X}^{-T}$

Derivatives of linear forms: $\frac{\partial \mathbf{x}^T \mathbf{a}}{\partial \mathbf{x}} = \frac{\partial \mathbf{a}^T \mathbf{x}}{\partial \mathbf{x}} = \mathbf{a}$ $\frac{\partial \mathbf{a}^T \mathbf{X} \mathbf{b}}{\partial \mathbf{X}} = \mathbf{a} \mathbf{b}^T$ $\frac{\partial \mathbf{a}^T \mathbf{X}^T \mathbf{b}}{\partial \mathbf{X}} = \mathbf{b} \mathbf{a}^T$

Derivatives of quadratic forms: $\frac{\partial \mathbf{x}^T \mathbf{B} \mathbf{x}}{\partial \mathbf{x}} = (\mathbf{B} + \mathbf{B}^T)\mathbf{x}$

Gradient and Hessian: if $f = \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x}$ then $\nabla_{\mathbf{x}} f = \frac{\partial f}{\partial \mathbf{x}} = (\mathbf{A} + \mathbf{A}^T)\mathbf{x} + \mathbf{b}$; $\frac{\partial^2 f}{\partial \mathbf{x} \partial \mathbf{x}^T} = \mathbf{A} + \mathbf{A}^T$

Derivative of trace: if $f(x) = \frac{dF(\mathbf{X})}{dx}$ then $\frac{\partial \text{Tr}(F(\mathbf{X}))}{\partial \mathbf{X}} = f(\mathbf{X})^T$

4.4.3. Quadratic Forms

A quadratic form is a scalar-valued function of the form

$$f(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x}, \quad f: \mathbb{R}^n \rightarrow \mathbb{R}, \quad \mathbf{A} = \mathbf{A}^T \in \mathbb{R}^{n \times n}, \quad \mathbf{b} \in \mathbb{R}^n.$$

If $\mathbf{x} = [x_1 \dots x_n]$ then the $\mathbf{x}^T \mathbf{A} \mathbf{x}$ term expands to a weighted sum over all combinations of $x_i x_j$ and the $\mathbf{b}^T \mathbf{x}$ term expands to a weighted sum over all x_i .

Iff $\mathbf{A} > 0$ then f is a strictly convex function.

4.4.4. Block Matrices

A block matrix is a matrix in which the entries can be matrices, vectors or scalars.

$$\mathbf{P} = \begin{bmatrix} 1 & 2 & 2 & 7 \\ 1 & 5 & 6 & 2 \\ 3 & 3 & 4 & 5 \\ 3 & 3 & 6 & 7 \end{bmatrix} \quad \mathbf{P}_{11} = \begin{bmatrix} 1 & 2 \\ 1 & 5 \end{bmatrix}, \quad \mathbf{P}_{12} = \begin{bmatrix} 2 & 7 \\ 6 & 2 \end{bmatrix}, \quad \mathbf{P}_{21} = \begin{bmatrix} 3 & 3 \\ 3 & 3 \end{bmatrix}, \quad \mathbf{P}_{22} = \begin{bmatrix} 4 & 5 \\ 6 & 7 \end{bmatrix}.$$

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \\ \mathbf{P}_{21} & \mathbf{P}_{22} \end{bmatrix}.$$

Block matrix \mathbf{P} is said to be $(m \times n)$ if it has m row partitions and n column partitions. Its full dimensions are at least as large as $m \times n$.

- **Multiplication:** if \mathbf{A} is $(m \times n)$ and \mathbf{B} is $(p \times n)$ then $\mathbf{C} = \mathbf{A}\mathbf{B}$ is $(m \times n)$: $\mathbf{C}_{qr} = \mathbf{A}_{qi} \mathbf{B}_{ir}$.
All partitions must be conformable such that # columns in $\mathbf{A}_{qi} = \#$ rows in \mathbf{B}_{ir} for all i .

M5. STATISTICS

5.1. Axioms, Combinatorial Probability and Basic Statistics

5.1.1. Axioms of Probability

Axioms: The probability P is a measure that verifies the following:

- Probability of an event: $P(A) \in R, P(A) \geq 0, \forall A \subseteq \Omega$
- Sample space is a certain event: $P(\Omega) = 1$
- Additivity for disjoint sets: $P(A \cup B) = P(A) + P(B), \text{ if } A \cap B = \emptyset$

Immediate Consequences: these can be demonstrated easily from a Venn diagram.

- Monotonicity: if $A \subseteq B$ then $P(A) \leq P(B)$.
- Empty set: $P(\emptyset) = 0$.
- Complement: $P(\bar{A}) = 1 - P(A)$
- Bound: $0 \leq P(A) \leq 1, \forall A \subseteq \Omega$.

5.1.2. Rules of Probability

- Addition: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$.
- Sum rule: $P(A \cap B) + P(A \cap \bar{B}) = P(A)$.
- Conditional probability: $P(A | B) = \frac{P(A \cap B)}{P(B)}$ if $P(B) \neq 0$.
- Law of total probability: $P(A) = P(A | B) P(B) + P(A | \bar{B}) P(\bar{B})$.
- Bayes' rule: $P(B | A) = \frac{P(B)}{P(A)} P(A | B) = \frac{P(A | B) P(B)}{P(A | B) P(B) + P(A | \bar{B}) P(\bar{B})}$.

5.1.3. Combinatorial (Frequentist) Definition of Probability

Assuming that the event occurs with equal likelihood in all outcomes:

$$\text{Probability of an event } A = \frac{\text{number of outcomes in which } A \text{ occurs}}{\text{total number of outcomes}}.$$

5.1.4. Definitions of Mean and Variance

For a single variable dataset $\{x_1, x_2, \dots, x_n\}$, the sample mean and variance are

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 = \frac{1}{n-1} \sum_{i=1}^n x_i^2 - \frac{1}{n(n-1)} \left(\sum_{i=1}^n x_i \right)^2$$

If $n \rightarrow N$ is large and representative of the population, then the population mean and variance are

$$\mu = \frac{1}{N} \sum_{i=1}^N x_i \quad \sigma^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2 = \frac{1}{N} \sum_{i=1}^N x_i^2 - \frac{1}{N^2} \left(\sum_{i=1}^N x_i \right)^2$$

The quantities $\sum_{i=1}^n x_i$ and $\sum_{i=1}^n x_i^2$ are known as summary statistics.

Standard deviation: σ

Coefficient of Variation: $\text{CoV} = \frac{\sigma}{\mu}$

5.1.5. Data Presentation

For categorical or discrete data:

- Pie chart: angle of the pie represents fraction of total frequency.
- Bar chart: bar height represents frequency. May have error bars, be grouped and/or stacked.
- Frequency table: lists the frequencies explicitly.
- Pictogram / tally chart: shows icons representing a given unit frequency.
- Choropleth map: colour-coded values or buckets. Often used for geographical data.

For numerical and continuous data:

- Histogram: shows frequency density = frequency (area) / bin size, of the intervals
- Line chart: shows values as X on the graph, connected. Can also be used for discrete.
- Stem and leaf plot: lists of numerical data grouped by the most significant digit.
- Box and whisker plot: shows the min, max, quartiles and mean.
Conventionally, outliers are identified as $x > \text{UQ} + 1.5 \times \text{IQR}$ or $x < \text{LQ} - 1.5 \times \text{IQR}$.
(UQ: upper quartile, LQ: lower quartile, IQR = UQ - LQ: interquartile range)

For bivariate data:

- Scatter graph: shows values as X on the graph, not connected.

5.1.6. Common Numbers and Operators in Combinatorics

- **Factorial:** $n! = \prod_{k=1}^n k = n(n - 1)(n - 2)\dots 2 \times 1$ for positive integers n , and $0! = 1$.

- **Double factorial:**

$$n!! = n(n - 2)(n - 4)\dots 2 \text{ if } n \text{ even; } \quad n!! = n(n - 2)(n - 4)\dots 3 \text{ if } n \text{ odd.}$$

$$n!! = 2^{n/2} \left(\frac{n}{2}\right)! \text{ if } n \text{ even; } \quad n!! = 2^{\left(\frac{1-n}{2}\right)} \times \frac{n!}{\left(\frac{n-1}{2}\right)!} \text{ if } n \text{ odd.}$$

- **Asymptotic growth:** $n^k \ll n!! \ll n! \ll n^n$ as $n \rightarrow \infty$

- **Stirling's approximation:** $\ln n! \sim n \ln n - n$ and $n! \sim \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$ as $n \rightarrow \infty$

- **Double factorial ratio:** $\frac{(2n)!!}{(2n-1)!!} \sim \sqrt{\pi n}$ as $n \rightarrow \infty$

- **Rising factorial:** $x^{(n)} = x^{\overline{n}} = \prod_{k=0}^{n-1} (x + k) = x(x + 1)(x + 2)\dots (x + n - 1) = \frac{(x+n-1)!}{(x-1)!}$

- **Falling factorial:** $(x)_n = \prod_{k=0}^{n-1} (x - k) = x(x - 1)(x - 2)\dots (x - n + 1) = \frac{x!}{(x-n)!}$

- **Derangement (subfactorial):** $!n = n! \times \sum_{k=0}^n \frac{(-1)^k}{k!} = \text{round}\left(\frac{n!}{e}\right)$, with $!0 = 1$

- **Combinations (binomial coefficient):** $\binom{n}{r} = {}^n C_r = \frac{n!}{r!(n-r)!}$

- **Permutations:** ${}^n P_r = \frac{n!}{(n-r)!} = r! \times {}^n C_r$

- **Bell numbers:** $B_{n+1} = \sum_{k=0}^n {}^n C_k \times B_k$

- Explicit formula (Dobiński's formula): $B_n = \frac{1}{e} \sum_{k=0}^{\infty} \frac{k^n}{k!}$

- First few Bell numbers: $B_0 = 1, B_1 = 1, B_2 = 2, B_3 = 5, B_4 = 15, B_5 = 52, B_6 = 203\dots$

- Exponential generating function: $EG(B_n; z) = \sum_{n=0}^{\infty} \frac{B_n}{n!} z^n = e^{e^z - 1}$

- **Harmonic numbers:** $H_n = \sum_{r=1}^n \frac{1}{r}$

- **Stirling numbers of the 1st kind:** coefficients s such that $(x)_n = \sum_{k=0}^n s(n, k) x^k$

- **Stirling numbers of the 2nd kind:** $S(n, k) = \sum_{i=0}^k \frac{(-1)^{k-i} i^n}{(k-i)! i!}$

- **Lah numbers:** $L(n, k) = {}^{n-1} C_{k-1} \times \frac{n!}{k!}$

5.1.7. Counting Problems of Unordered Sets (Combinations and Partitions)

Consider a set S of n objects, all of which are **unique**, e.g. $S = \{\text{red circle, yellow heart, green heart, blue circle, purple heart}\}$

A **combination** is an **unordered** non-empty subset of S of any length r , e.g. $\{\text{red circle, green heart, blue circle}\}$.

A **partition** is a superset of S containing any number k of non-overlapping (mutually disjoint) combinations of S , e.g. $\{\{\text{red circle, purple heart}\}, \{\text{yellow heart, green heart}\}, \{\text{blue circle}\}\}$.

A combination may (if specified) include the same element(s) of S multiple times. If this is the case, we say it has “repeats” or “with replacement”.

- Number of r -length **combinations without replacement** $= {}^n C_r = \frac{n!}{r!(n-r)!}$
- Number of r -length **combinations with replacement** $= {}^{n+r-1} C_r = \frac{(n+r-1)!}{r!(n-1)!}$
- Number of **partitions** $= B_n$ (Bell numbers)
- Number of **partitions into k combinations** $= S(n, k)$ (Stirling numbers of the second kind)

5.1.8. Counting Problems of Ordered Sets (Permutations and Derangements)

Consider a set S of n objects, all of which are **unique**, e.g. $S = \{\text{red circle, yellow heart, green heart, blue circle, purple heart}\}$.

A **permutation** is a linearly **ordered** non-empty subset of S of any length r , e.g. $(\text{red circle, blue circle, yellow heart})$.

A **cyclic permutation** is a permutation where left/right shifting is considered an identical permutation.

A **derangement** is an n -length permutation without replacement of an **ordered** set S such that no object remains in its original position, e.g. $(\text{red circle, yellow heart, green heart, blue circle, purple heart}) \rightarrow (\text{yellow heart, blue circle, red circle, purple heart, green heart})$.

A permutation may (if specified) include the same element(s) of S multiple times. If this is the case, we say it has “repeats” or “with replacement”.

- Number of r -length **permutations without replacement** $= {}^n P_r = \frac{n!}{(n-r)!}$
- Number of r -length **permutations with replacement** $= n^r$
- Number of **derangements** $= !n$ (derangement)
- Number of **partitions into k cyclic permutations** $= s(n, k)$ (Stirling numbers of the 1st kind)
- Number of **partitions into k permutations** $= L(n, k)$ (Lah numbers / Stirling numbers of the 3rd kind)

If the set contains items considered to be **identical** e.g. $S = \{\text{red circle, red circle, yellow heart, green heart, green heart, green heart, blue circle, blue circle, purple heart}\}$, where each item $\text{red circle, yellow heart, green heart} \dots$ appears n_1, n_2, n_3, \dots times respectively ($n_1 + n_2 + n_3 + \dots = n$):

- Number of n -length **multiset permutations without replacement** $= \frac{n!}{n_1! n_2! n_3! \dots}$

5.1.9. Common Counting Problems

Coupon Collector Problem: each box contains one coupon. Each coupon can come in n different kinds, uniformly distributed. Let T be the number of boxes opened once at least one of each coupon kind has been found, where $T \geq n$.

- Probability distribution: $P(T = t) = \frac{n!}{n^t} S(t - 1, n - 1) = n^{1-t} \sum_{r=0}^{n-1} {}^{n-1}C_r (-1)^{n-r-1} r^{t-1}$
- Cumulative distribution: $P(T \leq t) = \frac{n!}{n^t} S(t, n) = \sum_{i=0}^n {}^n C_i (-1)^{n-i} \left(\frac{i}{n}\right)^t$
(S : Stirling numbers of the 2nd kind)
- Expected number of boxes required: $E[T] = n H_n$ (H_n : harmonic numbers)
- Variance in number of boxes required: $Var[T] = n^2 \left(\sum_{r=1}^n \frac{1}{r^2} \right) - n H_n$
- Expected remaining number of boxes required given $m < n$ already found: $E[T_m] = n H_{n-m}$
- Asymptotic limit for large n : $E[T] \sim n \ln n + \gamma n + \frac{1}{2} + O(n^{-1})$ ($\gamma \approx 0.5772\dots$: Euler-Mascheroni constant)

Hat Check Problem: a group of n men put their n unique hats into a box. Afterwards, the men then randomly take back one hat each from the box without replacement. Let M be the number of men who correctly retrieved their own hat, where $0 \leq M \leq n$.

- Probability distribution: $P(M = m) = \frac{{}^n C_m \times !(n-m)}{n!}$ ($!$: derangement)
- Probability all incorrect: $P(M = 0) = \frac{!n}{n!}$, with $\lim_{n \rightarrow \infty} P(M = 0) = \frac{1}{e}$.
- Probability all correct: $P(M = n) = \frac{1}{n!}$.
- Expected number of correct hats: $E[M] = 1$
- Variance in number of correct hats: $Var[M] = 1$

Birthday Problem: in a room of n people, find the probability that any two share a birthday.

- $P(\text{at least two shared birthdays}) = 1 - \frac{{}^{365}P_n}{365^n}$.
- For two groups (a men and b women, with $a + b = n$):

$$P(\text{a man shares a birthday with a woman}) = \frac{1}{365^n} \sum_{i=1}^a \sum_{j=1}^b S(a, i) S(b, j) \prod_{k=0}^{i+j-1} (365 - k)$$

Intersection of Random Sets: let $S_X = \{1, 2, \dots, M\}$ and $S_Y = \{1, 2, \dots, N\}$. Without replacement, uniform randomly and independently, choose a elements from S_X and b elements from S_Y , and put them in sets X and Y respectively.

- The count of common elements has distribution $P(|X \cap Y| = z) = \frac{{}^M C_z \times {}^{M-z} C_{b-z} \times {}^{N-b} C_{a-z}}{{}^M C_b \times {}^N C_a}$.

5.2. Probability Distributions and Random Variables

5.2.1. Discrete Probability Distributions

Distribution	$P(X=r)$ PMF	$E(X)$ expectation	$\text{Var}(X)$ variance	$G_X(z)$ PGF	$H(X)$ differential entropy
Bernoulli $r \in \{0, 1\}$ $X \sim \text{Ber}(p)$	$p^r q^{1-r}$ $q = 1 - p$	p	pq	$q + pz$	$-q \log q + p \log p$
Discrete Uniform $r \in \{a, \dots, b\}$ $X \sim U(a, b)$	$\frac{1}{n}$ $n = b - a + 1$	$\frac{a+b}{2}$	$\frac{n^2 - 1}{12}$	$\frac{z^a - z^{b+1}}{n(1-z)}$	$\log n$
Binomial $r \in \{0, 1, 2, \dots, n\}$ $X \sim B(n, p)$	${}^n C_r p^r q^{n-r}$	np	npq	$(q + pz)^n$	$\log \sqrt{2\pi npq} + O(n^{-1})$
Negative Binomial $r \in \{0, 1, 2, \dots\}$ $X \sim \text{NB}(n, p)$	${}^{r+n-1} C_r p^n q^r$	$\frac{nq}{p}$	$\frac{nq}{p^2}$	$\left(\frac{p}{1-qz}\right)^n$	-
Beta-Binomial $r \in \{0, 1, 2, \dots, n\}$ $X \sim \text{BetaBin}(n, \alpha, \beta)$	$\frac{{}^n C_r B(r+\alpha, n-r+\beta)}{B(\alpha, \beta)}$	$\frac{n\alpha}{\alpha + \beta}$	$\frac{n\alpha\beta(\alpha+\beta+n)}{(\alpha+\beta)^2(\alpha+\beta+1)}$	-	-
Geometric $r \in \{1, 2, 3, \dots\}$ $X \sim \text{Geo}(p)$	$q^{r-1} p$	p^{-1}	qp^{-2}	$\frac{pz}{1-qz}$	$\frac{-q}{p} \log q - \log p$
Hypergeometric $r \in \{\max(0, n+K-N), \dots, \min(n, K)\}$ $X \sim \text{HG}(N, K, n)$	$\frac{{}^K C_r {}^{N-K} C_{n-r}}{{}^N C_n}$	$\frac{nK}{N}$	$\frac{nK(N-K)(N-n)}{N^2(N-1)}$	$\frac{{}^{N-K} C_n}{{}^N C_n} \times {}_2F_1(-n, -K; N-K-n+1; z)$	-
Negative Hypergeometric $r \in \{0, \dots, K\}$ $X \sim \text{NHG}(N, K, n)$	$\frac{{}^{r+n-1} C_r {}^{N-n-r} C_{K-r}}{{}^N C_K}$	$\frac{nK}{N-K+1}$	$\frac{nK(N+1)(N-K+1-n)}{N-K+2}$	-	-
Poisson $r \in \{0, 1, 2, \dots\}$ $X \sim \text{Po}(\lambda)$	$\frac{e^{-\lambda} \lambda^r}{r!}$	λ	λ	$e^{\lambda(z-1)}$	$\log \sqrt{2\pi e \lambda} + O(\lambda^{-1})$

5.2.2. Continuous Probability Distributions

Distribution	$f_X(x)$ PDF	$E(X)$ expectation	$Var(X)$ variance	$M_X(s)$ MGF	$H(X)$ differential entropy
Rectangular $x \in [a, b]$ $X \sim \text{Rect}(a, b)$	$\frac{1}{b-a}$	$\frac{a+b}{2}$	$\frac{(b-a)^2}{12}$	$\frac{e^{bs} - e^{as}}{s(b-a)}$	$\log(b-a)$
Normal $x \in \mathbb{R}$ $X \sim N(\mu, \sigma^2)$	$\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$	μ	σ^2	$e^{\mu s + \frac{1}{2}\sigma^2 s^2}$	$\frac{1}{2}(1 + \log(2\pi\sigma^2))$
Student's t $x \in \mathbb{R}$ $X \sim t_\nu$	$\frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi} \Gamma(\frac{\nu}{2})} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{1}{2}(\nu+1)}$	0	$\frac{\nu}{\nu-2}$	(undefined for tail ranges of t)	$\frac{\nu+1}{2}(\psi(\frac{\nu+1}{2}) - \psi(\frac{\nu}{2})) + \log(\nu^{1/2} B(\frac{\nu}{2}, \frac{1}{2}))$
Exponential $x \geq 0$ $X \sim \text{Exp}(\lambda)$	$\lambda e^{-\lambda x}$	$\frac{1}{\lambda}$	$\frac{1}{\lambda^2}$	$\frac{\lambda}{\lambda-s}$	$1 - \log \lambda$
Rayleigh $x \geq 0$ $X \sim \text{Rayleigh}(\sigma)$	$\frac{x}{\sigma^2} e^{-\frac{x^2}{2\sigma^2}}$	$\sqrt{\frac{\pi\sigma^2}{2}}$	$\frac{4-\pi}{2} \sigma^2$	-	$1 + \frac{\gamma}{2} + \log \frac{\sigma}{\sqrt{2}}$
Beta $x \in [0, 1]$ $X \sim \text{Beta}(\alpha, \beta)$	$\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1} (1-x)^{\beta-1}$	$\frac{\alpha}{\alpha+\beta}$	$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$	$1 + \sum_{k=1}^{\infty} \left(\prod_{r=0}^{k-1} \frac{\alpha+r}{\alpha+\beta+r} \right) \frac{s^k}{k!}$	$\log B(\alpha, \beta) - (\alpha-1)\psi(\alpha) - (\beta-1)\psi(\beta) + (\alpha+\beta-2)\psi(\alpha+\beta)$

For the multivariate Normal distribution, see Section 5.4.1.

5.2.3. Typical Modelling Cases for Probability Distributions

Urn Model (With Replacements / Repeats): Consider an urn (container) of balls, of which a proportion p are **green** (success) and the rest $(1 - p)$ are **red** (failure). Suppose that balls are sampled **with replacement** (each ball is returned before taking the next ball).

Binomial	A total of n balls are sampled. The number of green balls sampled is distributed as $X \sim B(n, p)$.
Geometric	Samples are drawn until the first green ball is sampled. The total number of balls sampled (including the green) is distributed as $X \sim \text{Geo}(p)$. The number of red balls sampled is then $X - 1$.
Negative Binomial	Samples are drawn until the first n green balls are sampled. The number of red balls sampled is distributed as $X \sim \text{NB}(n, p)$.
Beta	A fixed number of balls are drawn. It is observed that $\alpha - 1$ are green and that $\beta - 1$ are red . The prior for the number of green balls is uniform. The (posterior) probability of drawing green is distributed as $p \sim \text{Beta}(\alpha, \beta)$.
Beta Binomial	A total of n balls are sampled. The probability p of drawing green is unknown and is distributed as $p \sim \text{Beta}(\alpha, \beta)$. The number of green balls sampled is distributed as $X \sim \text{BetaBin}(n, \alpha, \beta)$.

Urn Model (Without Replacement / Repeats): Consider an urn of N balls, of which K are **green** and the rest $(N - K)$ are **red**. Suppose that balls are sampled **without replacement** (the combination of balls is taken at once, so the sample must contain unique balls).

Hypergeometric	A total of n balls are sampled. The number of green balls sampled is distributed as $X \sim \text{HG}(N, K, n)$. If $n \ll N \rightarrow \infty$ then $X \sim B(n, K/N)$ ($K = pN$).
Negative Hypergeometric	Samples are drawn until the first n red balls are sampled. The number of green balls sampled is distributed as $X \sim \text{NHG}(N, K, n)$. If $n \ll N \rightarrow \infty$ then $X \sim \text{BetaBin}(K, n, N - K - n + 1)$ ($K = \frac{\alpha}{p}$ and $N - K = \frac{\beta}{1 - p}$).

Event Model: Consider a period over which point events can occur at an average rate λ events per unit interval (often time, sometimes distance or number of transitions).

Poisson	The number of events per unit interval is distributed as $X \sim \text{Po}(\lambda)$.
Exponential	The interval between two consecutive events is distributed as $X \sim \text{Exp}(\lambda)$.

5.2.4. Sampling From Normal Distributions

If X is Normally distributed as $X \sim N(\mu, \sigma^2)$, then:

- The standard score $Z = \frac{X - \mu}{\sigma}$ is distributed as $Z \sim N(0, 1)$.

For a random sample of n observations X_n from X with sample standard deviation s :

- The sample mean \bar{X} is distributed as $\bar{X} \sim N(\mu, \frac{\sigma^2}{n})$. ($Var(\bar{X}) = \frac{\sigma_x^2}{n}$ for any iid sample)
- The sample mean standard score $\bar{Z} = \frac{\bar{X} - \mu}{s^2/n}$ is distributed as $\bar{Z} \sim t_{n-1}$.
- If n is sufficiently large ($n > 30$) then $s \approx \sigma$ and $\bar{Z} \sim N(0, 1)$.

5.2.5. Central Limit Theorem

For a set of n independent random variables X_1, X_2, \dots, X_n , each having means and variances $(\mu_1, \sigma_1^2), (\mu_2, \sigma_2^2), \dots, (\mu_n, \sigma_n^2)$, the CLT states that

- The random variable $S_n = \sum_{i=1}^n X_i$ is approximately Normally distributed (weak CLT),
- $\lim_{n \rightarrow \infty} S_n \sim N\left(\sum_{i=1}^n \mu_i, \sum_{i=1}^n \sigma_i^2\right)$ i.e. the approximation is exact as $n \rightarrow \infty$.

If all X_i are i.i.d. with mean μ and variance σ^2 then $\frac{S_n - n\mu}{\sigma\sqrt{n}} \sim N(0, 1)$ as $n \rightarrow \infty$.

These results hold **regardless** of the distribution of X .

The Berry-Esseen Theorem (improved by Shevtsova, 2010) bounds the error in the weak CLT Normal approximation by the value of its CDF (z is the value of standardised S_n):

$$\left| \underbrace{P\left(\frac{\sum_{i=1}^n (X_i - \mu_i)}{\sqrt{\sum_{i=1}^n \sigma_i^2}} \leq z\right)}_{\text{exact CDF at } z} - \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-x^2/2} dx}_{\text{approximate CDF at } z} \right| \leq \underbrace{0.56 \times \frac{\max_i \{E(|X_i|^3)\}}{\sqrt{n} \cdot \max_i \{\sigma_i^3\}}}_{\text{upper bound for error}}$$

Therefore error $\sim n^{-1/2}$. Practically, the CLT is 'good' when the sample size is $n \geq 30$.

5.2.6. Expectation from Probability Density Function

Expectation (mean) of a random variable given PMF or PDF:

$$\text{Discrete: } E[X] = \sum_{x=-\infty}^{\infty} x P(X = x) = \mu \quad \text{Continuous: } E[X] = \int_{-\infty}^{\infty} x f_X(x) dx = \mu$$

Tail formula for nonnegative random variables ($X \geq 0$):

$$\text{Discrete: } E[X] = \sum_{x=0}^{\infty} P(X \geq x) \quad \text{Continuous: } E[X] = \int_0^{\infty} (1 - F_X(x)) dx$$

Law of the Unconscious Statistician (LOTUS theorem): if $Y = g(X)$ then

$$\text{Discrete: } E[Y] = \sum_{x=-\infty}^{\infty} g(x) P(X = x) \quad \text{Continuous: } E[Y] = \int_{-\infty}^{\infty} g(x) f_X(x) dx$$

Variance of a random variable given PMF or PDF and mean (2nd central moment):

$$\text{Discrete: } Var[X] = \sum_{x=-\infty}^{\infty} (x - \mu)^2 P(X = x) \quad \text{Continuous: } Var[X] = \int_{-\infty}^{\infty} (x - \mu)^2 f_X(x) dx$$

5.2.7. Moments of Distributions

Moments of a random variable X are defined as:

- **n th raw moment:** $\mu_n' = E[X^n]$
- **n th central moment:** $\mu_n = E[(X - \mu)^n] = \sum_{r=0}^n (-1)^r \binom{n}{r} \mu^r \mu_{n-r}'$
- **n th central moment, standardised:** $\frac{\mu_n}{\sigma^n} = \frac{E[(X - \mu)^n]}{E[(X - \mu)^2]^{n/2}}$

Some important moment-related measures are:

- **Mean:** $\mu = E[X] = \mu_1'$ (central tendency)
- **Variance:** $\sigma^2 = Var[X] = E[X^2] - E[X]^2 = \mu_2$ (spread about the mean)
- **Skewness:** $\gamma = Skew[X] = \mu_3 / \sigma^3$ (asymmetry)
 $\gamma > 0$: smaller values more likely
 $\gamma < 0$: larger values more likely
- **Kurtosis:** $\kappa = Kurt[X] = \mu_4 / \sigma^4$ (tailedness / outlier frequency)
- **Excess Kurtosis:** $\bar{\kappa} = \kappa - 3$ (Fisher kurtosis)
 mesokurtic: $\bar{\kappa} = 0$ (Gaussian)
 leptokurtic: $\bar{\kappa} > 0$ (tail-heavy)
 platykurtic: $\bar{\kappa} < 0$ (tail-light)

5.2.8. Generating Functions

For a **discrete** random variable X :

- **Probability generating function (PGF):**
$$G_X(z) = E[z^X] = \sum_{k=-\infty}^{\infty} P(X = k) z^k$$

For any variable X , the PGF is a polynomial in z where the coefficients give the PMF. The PGF is the Z -transform (discretised Mellin transform) of the PMF $P(X = k)$.

For a **continuous** random variable X :

- **Moment generating function (MGF):**
$$M_X(s) = E[e^{sX}] = \int_{-\infty}^{\infty} f_X(x) e^{sx} dx$$

The MGF is the bilateral Laplace transform of $f_X(x)$ (or ordinary LT if X has support $X \geq 0$). The PGF and MGF definitions are related by $z = e^s$.

- **Characteristic function (CF):**
$$\varphi_X(t) = E[e^{itX}] = \int_{-\infty}^{\infty} f_X(x) e^{ixt} dx$$

The CF is the Fourier transform of $f_X(x)$ (signs reversed; using t as the Fourier 'frequency'). The MGF and CF definitions are related by $s = it$ so that $M_X(t) = \varphi_X(-it)$.

Identities with Generating Functions: manipulating random variables

- Shift and scale: $Y = aX + b \Rightarrow M_Y(s) = e^{sb} M_X(as)$
- Sum of RVs: $Y = X_1 + X_2 \Rightarrow G_Y(z) = G_{X_1}(z) G_{X_2}(z); M_Y(s) = M_{X_1}(s) M_{X_2}(s)$
- Difference of RVs: $Y = X_1 - X_2 \Rightarrow G_Y(z) = G_{X_1}(z) G_{X_2}\left(\frac{1}{z}\right); M_Y(s) = M_{X_1}(s) M_{X_2}(-s)$
- Joint GF of X, Y : $G_{X,Y}(z, w) = E[z^X w^Y]; M_{X,Y}(s, t) = E[e^{sX+tY}]$
- Marginal GF: $G_X(z) = G_{X,Y}(z, 0); M_X(s) = M_{X,Y}(s, 0)$

Important Quantities (Moments) from the PGF and MGF: see Section 5.2.7 for definitions

The coefficient of z^k in the expansion of the PGF $G_X(z)$ is the PMF of X i.e. $P(X = k)$.

The coefficient of s^n in the expansion of the MGF $M_X(s)$ is proportional to the n th raw moment of X :

$$M_X(s) = \sum_{n=0}^{\infty} \frac{s^n \mu_n'}{n!} \text{ since } \mu_n' = M_X^{(n)}(s)$$

Expectation: $E[X] = M_X'(0) = G_X'(1) \quad \text{since } E[X^n] = M_X^{(n)}(0) = G_X^{(n)}(1)$

Variance:
$$\begin{aligned} \text{Var}[X] &= E[X^2] - E[X]^2 \\ &= M_X''(0) - M_X'(0)^2 \\ &= G_X''(1) + G_X'(1) - (G_X'(1))^2 \end{aligned}$$

5.2.10. Inverse Transform Sampling

Let U be a uniformly distributed random variable on $0 \leq U \leq 1$. If X is a random variable with cdf or cmf $F_X(x)$, then $X = F_X^{-1}(U)$. This can be used to generate samples with a given pdf or cdf, given a system to generate uniformly distributed random numbers.

To generate RVs with a truncated distribution $a \leq X \leq b$, let $F(a) \leq U \leq F(b)$ and $X = F_X^{-1}(U)$.

5.2.11. Inequalities for the Expectation of Random Variables

- **Markov's Inequality:** $P(X \geq a) \leq \frac{E[X]}{a}$ for nonnegative X i.e. $X \geq 0$
- **Jensen's Inequality:** $E[g(X)] \geq g(E[X])$ for convex functions g i.e. $g''(x) \geq 0$
- **Chebyshev's Inequality:** $P(|X - E[X]| > a) \leq \frac{\text{Var}[X]}{a^2}$ for $a > 0$
- **Minkowski's Inequality:** $(E[|X + Y|^p])^{1/p} \leq (E[|X|^p])^{1/p} + (E[|Y|^p])^{1/p}$ for $p \geq 1$
- **Hölder's Inequality:** $E[|XY|] \leq (E[|X|^p])^{1/p} (E[|Y|^q])^{1/q}$ for $p, q > 1$ with $\frac{1}{p} + \frac{1}{q} = 1$

5.2.12. Distributions of Functions of Random Variables

To find the PDF of Y , $f_Y(y)$, write down $F_Y(y) = P(Y \leq y)$ and let $Y = g(X)$. Solve the resulting inequality for X to write $F_Y(y)$ in terms of $F_X(y)$. Finally, use $f_Y(y) = \frac{dF_Y}{dy}$ and $F_X(x) = \int_{-\infty}^x f_X(x) dx$.

Common functions:

- If $Y = aX + b$ then $f_Y(y) = \frac{1}{a} f_X\left(\frac{y-b}{a}\right)$ where $a > 0$
- If $Y = X^2$ then $f_Y(y) = \frac{f_X(\sqrt{y}) + f_X(-\sqrt{y})}{2\sqrt{y}}$
- If $Y = X^3$ then $f_Y(y) = \frac{1}{3} y^{-2/3} f_X(y^{1/3})$
- If $Y = \sqrt{X}$ then $f_Y(y) = 2y f_X(y^2)$
- If $Y = \min(X, a)$ then $F_Y(y) = F_X(y)$ if $y \leq a$ else 1 ($f_Y(y)$ contains a delta function)
- If $Y = \max(X, a)$ then $F_Y(y) = F_X(y)$ if $y \geq a$ else 0 ($f_Y(y)$ contains a delta function)

For a multivariable function e.g. two RVs, $Y = g(X_1, X_2)$, use (for the CDF):

$$F_Y(y) = \iint_S f_{X_1 X_2}(x_1, x_2) dx_1 dx_2, \quad \text{where } S = \{(x_1, x_2) \in \mathbb{R}^2 : g(x_1, x_2) \leq y\}$$

Common functions:

- If $Y = \max\{X_1, X_2, \dots, X_n\}$ then $F_Y(y) = F_{X_1}(y) F_{X_2}(y) \dots F_{X_n}(y)$ (**parallel system**)
- If $Y = \min\{X_1, X_2, \dots, X_n\}$ then $F_Y(y) = \left(1 - F_{X_1}(y)\right) \left(1 - F_{X_2}(y)\right) \dots \left(1 - F_{X_n}(y)\right)$
(**series system**: models 'weakest link' failure)

5.2.13. Manipulating Independent Random Variables and their Distributions

Let X_1 and X_2 be **independent** random variables. Then:

Linear Transformations (green: also valid when X_1 and X_2 are not independent.)

Transform	$E[Y]$	$Var[Y]$	$f_Y(y)$	$M_Y(s)$
Scale $Y = aX$	$a E[X]$	$a^2 Var[X]$	$\frac{1}{a} f_X\left(\frac{y}{a}\right)$	$M_X(as)$
Scale and Shift $Y = aX + b$	$a E[X] + b$	$a^2 Var[X]$	$\frac{1}{a} f_X\left(\frac{y-b}{a}\right)$	$e^{sb} M_X(as)$
Sum $Y = X_1 + X_2$	$E[X_1] + E[X_2]$	$Var[X_1] + Var[X_2]$	$(f_{X_1} * f_{X_2})(y)$	$M_{X_1}(s) M_{X_2}(s)$

Nonlinear Transformations (all the below require **independence**)

- **Product:** if $Y = X_1 X_2$ then

$$E[Y] = E[X_1] E[X_2]$$

$$Var[Y] = Var[X_1] Var[X_2] + Var[X_1] E[X_2]^2 + Var[X_2] E[X_1]^2$$

The new distribution is $f_Y(y) = \int_{-\infty}^{\infty} f_{X_1}(x) f_{X_2}\left(\frac{y}{x}\right) \frac{1}{|x|} dx$ (product distribution)

- **Ratio:** if $Y = X_1 / X_2$ then

$E[Y]$ can be found by LOTUS, and $Var[Y] = E[X_1^2] E\left[\frac{1}{X_2^2}\right] - E[X_1]^2 E\left[\frac{1}{X_2}\right]^2$.

The new distribution is $f_Y(y) = \int_{-\infty}^{\infty} |x| f_{X_1}(xy) f_{X_2}(x) dx$ (ratio distribution)

Combinations of Common Distributions: for independent random variables,

- If $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$ then $X_1 + X_2 \sim N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$.
- If $X_1 \sim Po(\lambda_1)$ and $X_2 \sim Po(\lambda_2)$ then $X_1 + X_2 \sim Po(\lambda_1 + \lambda_2)$.
- If $X_1 \sim B(n_1, p)$ and $X_2 \sim B(n_2, p)$ then $X_1 + X_2 \sim B(n_1 + n_2, p)$.
- If $X \sim B(N, p)$ and $N \sim Po(\lambda)$ then $X \sim Po(p\lambda)$.
- If $X_1 \sim N(0, \sigma^2)$ and $X_2 \sim N(0, \sigma^2)$ then $X_1^2 + X_2^2 \sim Exp\left(\frac{1}{2\sigma^2}\right)$.
- If $X_i \sim N(0, 1)$ then $\sum_{i=1}^N X_i^2 \sim \chi^2_N$ (Chi-Square distribution with N degrees of freedom).
- If $X_1 \sim N(0, \sigma^2)$ and $X_2 \sim N(0, \sigma^2)$ then $\sqrt{X_1^2 + X_2^2} \sim Rayleigh(\sigma)$.

Reciprocal Normal Distribution: if $X \sim N(\mu, \sigma^2)$ and $Y = 1/X$ then $f_Y(y) = \frac{1}{\sqrt{2\pi\sigma y^2}} \exp\left(-\frac{(1-\mu y)^2}{2\sigma^2 y^2}\right)$.

5.2.14. Multivariable Probability: Joint and Marginal Distributions

For jointly distributed random variables (X, Y) , let $p(x, y) = f_{XY}(x, y)$ denote the joint pmf if (X, Y) are discrete, or the joint pdf if (X, Y) are jointly continuous. $p(x) = f_X(x)$ is the marginal pmf/pdf of X . $p(x | y) = f_{X|Y}(x | y)$ is the pmf/pdf of X , conditioned on some value of $Y = y$.

- Definition of joint density: $P(a \leq X \leq b \cap c \leq Y \leq d) = \int_c^d \int_a^b p(x, y) dx dy$
- Definition from joint CDF: $p(x, y) = \frac{\partial^2}{\partial x \partial y} F_{XY}(x, y)$ (also written $f_{XY}(x, y)$)
- Conditional PMF / PDF: $p(x | y) = \frac{p(x, y)}{p(y)}$, $p(y) \neq 0$ (also written $f_{X|Y}(x | y)$)
- Bayes' rule: $p(x | y) = \frac{p(y | x) p(x)}{\sum_x p(x) p(x | y)}$ ($p(x)$: prior distribution)

Independence between random variables $\mathbf{X} = \{X_1, X_2, \dots\}$: (\mathbf{X} is a random vector)



- Conditional independence: $p(x_1, x_2 | y) = p(x_1 | y) p(x_2 | y)$, $p(y) \neq 0$
- Pairwise independence: $p(x_i, x_j) = p(x_i) p(x_j)$ for every pair i, j
- (Mutual) independence: $p(\text{union of all } x_i) = \text{product of all } p(x_i)$ (stronger than pairwise)
- Conditional expectation: $E[X | Y] = \sum_x x p(x | y)$ (= $E[X]$ if independent)
- Linearity of conditionals: $E[(aX + bY) | Z] = a E[X | Z] + b E[Y | Z]$

Quantification of dependence and association:

- Covariance: $\text{Cov}[X, Y] = E[(X - E[X])(Y - E[Y])] = E[XY] - E[X] E[Y]$
- Law of total probability: $p(x) = \sum_y p(x, y)$
- Law of total expectation: $E[X] = E[E[X | Y]]$
- Law of total variance: $\text{Var}[X] = E[\text{Var}[Y | X]] + \text{Var}[E[Y | X]]$
- Law of total covariance: $\text{Cov}[X, Y] = E[\text{Cov}[X, Y | Z]] + \text{Cov}[E[X | Z], E[Y | Z]]$
- Linearity of covariance: $\text{Cov}[aX, Y] = \text{Cov}[X, aY] = a \text{Cov}[X, Y]$, and $\text{Cov}[X + Y, Z] = \text{Cov}[X, Z] + \text{Cov}[Y, Z]$
- Pearson Correlation Coefficient: $\rho_{XY} = \text{Corr}[X, Y] = \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}(X) \text{Var}(Y)}} = \frac{\text{Cov}(X, Y)}{\sigma_X \sigma_Y}$ (PMCC)
- General variance of sums: $\text{Var}[aX \pm bY + c] = a^2 \text{Var}[X] + b^2 \text{Var}[Y] \pm 2ab \text{Cov}[X, Y]$

5.2.15. Standard Normal Distribution: Critical Values (z table and its Inverse)

The table gives the values of $P(Z \leq z)$ for a given z , with $Z \sim N(0, 1)$. These are left tail probabilities.



z	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09
-3.4	.0003	.0003	.0003	.0003	.0003	.0003	.0003	.0003	.0003	.0002
-3.3	.0005	.0005	.0005	.0004	.0004	.0004	.0004	.0004	.0004	.0003
-3.2	.0007	.0007	.0006	.0006	.0006	.0006	.0006	.0005	.0005	.0005
-3.1	.0010	.0009	.0009	.0009	.0008	.0008	.0008	.0008	.0007	.0007
-3.0	.0013	.0013	.0013	.0012	.0012	.0011	.0011	.0011	.0010	.0010
-2.9	.0019	.0018	.0018	.0017	.0016	.0016	.0015	.0015	.0014	.0014
-2.8	.0026	.0025	.0024	.0023	.0023	.0022	.0021	.0021	.0020	.0019
-2.7	.0035	.0034	.0033	.0032	.0031	.0030	.0029	.0028	.0027	.0026
-2.6	.0047	.0045	.0044	.0043	.0041	.0040	.0039	.0038	.0037	.0036
-2.5	.0062	.0060	.0059	.0057	.0055	.0054	.0052	.0051	.0049	.0048
-2.4	.0082	.0080	.0078	.0075	.0073	.0071	.0069	.0068	.0066	.0064
-2.3	.0107	.0104	.0102	.0099	.0096	.0094	.0091	.0089	.0087	.0084
-2.2	.0139	.0136	.0132	.0129	.0125	.0122	.0119	.0116	.0113	.0110
-2.1	.0179	.0174	.0170	.0166	.0162	.0158	.0154	.0150	.0146	.0143
-2.0	.0228	.0222	.0217	.0212	.0207	.0202	.0197	.0192	.0188	.0183
-1.9	.0287	.0281	.0274	.0268	.0262	.0256	.0250	.0244	.0239	.0233
-1.8	.0359	.0351	.0344	.0336	.0329	.0322	.0314	.0307	.0301	.0294
-1.7	.0446	.0436	.0427	.0418	.0409	.0401	.0392	.0384	.0375	.0367
-1.6	.0548	.0537	.0526	.0516	.0505	.0495	.0485	.0475	.0465	.0455
-1.5	.0668	.0655	.0643	.0630	.0618	.0606	.0594	.0582	.0571	.0559
-1.4	.0808	.0793	.0778	.0764	.0749	.0735	.0721	.0708	.0694	.0681
-1.3	.0968	.0951	.0934	.0918	.0901	.0885	.0869	.0853	.0838	.0823
-1.2	.1151	.1131	.1112	.1093	.1075	.1056	.1038	.1020	.1003	.0985
-1.1	.1357	.1335	.1314	.1292	.1271	.1251	.1230	.1210	.1190	.1170
-1.0	.1587	.1562	.1539	.1515	.1492	.1469	.1446	.1423	.1401	.1379
-0.9	.1841	.1814	.1788	.1762	.1736	.1711	.1685	.1660	.1635	.1611
-0.8	.2119	.2090	.2061	.2033	.2005	.1977	.1949	.1922	.1894	.1867
-0.7	.2420	.2389	.2358	.2327	.2296	.2266	.2236	.2206	.2177	.2148
-0.6	.2743	.2709	.2676	.2643	.2611	.2578	.2546	.2514	.2483	.2451
-0.5	.3085	.3050	.3015	.2981	.2946	.2912	.2877	.2843	.2810	.2776
-0.4	.3446	.3409	.3372	.3336	.3300	.3264	.3228	.3192	.3156	.3121
-0.3	.3821	.3783	.3745	.3707	.3669	.3632	.3594	.3557	.3520	.3483
-0.2	.4207	.4168	.4129	.4090	.4052	.4013	.3974	.3936	.3897	.3859
-0.1	.4602	.4562	.4522	.4483	.4443	.4404	.4364	.4325	.4286	.4247
-0.0	.5000	.4960	.4920	.4880	.4840	.4801	.4761	.4721	.4681	.4641

z	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09
0.0	.5000	.5040	.5080	.5120	.5160	.5199	.5239	.5279	.5319	.5359
0.1	.5398	.5438	.5478	.5517	.5557	.5596	.5636	.5675	.5714	.5753
0.2	.5793	.5832	.5871	.5910	.5948	.5987	.6026	.6064	.6103	.6141
0.3	.6179	.6217	.6255	.6293	.6331	.6368	.6406	.6443	.6480	.6517
0.4	.6554	.6591	.6628	.6664	.6700	.6736	.6772	.6808	.6844	.6879
0.5	.6915	.6950	.6985	.7019	.7054	.7088	.7123	.7157	.7190	.7224
0.6	.7257	.7291	.7324	.7357	.7389	.7422	.7454	.7486	.7517	.7549
0.7	.7580	.7611	.7642	.7673	.7704	.7734	.7764	.7794	.7823	.7852
0.8	.7881	.7910	.7939	.7967	.7995	.8023	.8051	.8078	.8106	.8133
0.9	.8159	.8186	.8212	.8238	.8264	.8289	.8315	.8340	.8365	.8389
1.0	.8413	.8438	.8461	.8485	.8508	.8531	.8554	.8577	.8599	.8621
1.1	.8643	.8665	.8686	.8708	.8729	.8749	.8770	.8790	.8810	.8830
1.2	.8849	.8869	.8888	.8907	.8925	.8944	.8962	.8980	.8997	.9015
1.3	.9032	.9049	.9066	.9082	.9099	.9115	.9131	.9147	.9162	.9177
1.4	.9192	.9207	.9222	.9236	.9251	.9265	.9279	.9292	.9306	.9319
1.5	.9332	.9345	.9357	.9370	.9382	.9394	.9406	.9418	.9429	.9441
1.6	.9452	.9463	.9474	.9484	.9495	.9505	.9515	.9525	.9535	.9545
1.7	.9554	.9564	.9573	.9582	.9591	.9599	.9608	.9616	.9625	.9633
1.8	.9641	.9649	.9656	.9664	.9671	.9678	.9686	.9693	.9699	.9706
1.9	.9713	.9719	.9726	.9732	.9738	.9744	.9750	.9756	.9761	.9767
2.0	.9772	.9778	.9783	.9788	.9793	.9798	.9803	.9808	.9812	.9817
2.1	.9821	.9826	.9830	.9834	.9838	.9842	.9846	.9850	.9854	.9857
2.2	.9861	.9864	.9868	.9871	.9875	.9878	.9881	.9884	.9887	.9890
2.3	.9893	.9896	.9898	.9901	.9904	.9906	.9909	.9911	.9913	.9916
2.4	.9918	.9920	.9922	.9925	.9927	.9929	.9931	.9932	.9934	.9936
2.5	.9938	.9940	.9941	.9943	.9945	.9946	.9948	.9949	.9951	.9952
2.6	.9953	.9955	.9956	.9957	.9959	.9960	.9961	.9962	.9963	.9964
2.7	.9965	.9966	.9967	.9968	.9969	.9970	.9971	.9972	.9973	.9974
2.8	.9974	.9975	.9976	.9977	.9977	.9978	.9979	.9979	.9980	.9981
2.9	.9981	.9982	.9982	.9983	.9984	.9984	.9985	.9985	.9986	.9986
3.0	.9987	.9987	.9987	.9988	.9988	.9989	.9989	.9989	.9990	.9990
3.1	.9990	.9991	.9991	.9991	.9992	.9992	.9992	.9992	.9993	.9993
3.2	.9993	.9993	.9994	.9994	.9994	.9994	.9994	.9995	.9995	.9995
3.3	.9995	.9995	.9995	.9996	.9996	.9996	.9996	.9996	.9996	.9997
3.4	.9997	.9997	.9997	.9997	.9997	.9997	.9997	.9997	.9997	.9998


Inverse z Table: $z = \Phi^{-1}(P(Z \leq z))$

$P(Z \leq z)$	0.001	0.005	0.01	0.05	0.1	0.5
z	-3.090232	-2.575829	-2.326347	-1.644854	-1.281552	0

$P(Z \leq z)$	0.9	0.9995	0.95	0.99	0.995	0.999
z	1.281552	3.290527	1.644854	2.326347	2.575829	3.090232

5.2.16. Student's *t*-Distribution: Critical Values (Inverse *t* table)

The table gives the values of x satisfying $P(X \leq x) = p$, where X is a random variable having the t distribution with ν degrees of freedom.




<i>p</i>	0.9	0.95	0.975	0.99	0.995
<i>v</i>					
1	3.078	6.314	12.706	31.821	63.657
2	1.886	2.920	4.303	6.965	9.925
3	1.638	2.353	3.182	4.541	5.841
4	1.533	2.132	2.776	3.747	4.604
5	1.476	2.015	2.571	3.365	4.032
6	1.440	1.943	2.447	3.143	3.707
7	1.415	1.895	2.365	2.998	3.499
8	1.397	1.860	2.306	2.896	3.355
9	1.383	1.833	2.262	2.821	3.250
10	1.372	1.812	2.228	2.764	3.169
11	1.363	1.796	2.201	2.718	3.106
12	1.356	1.782	2.179	2.681	3.055
13	1.350	1.771	2.160	2.650	3.012
14	1.345	1.761	2.145	2.624	2.977
15	1.341	1.753	2.131	2.602	2.947
16	1.337	1.746	2.121	2.583	2.921
17	1.333	1.740	2.110	2.567	2.898
18	1.330	1.734	2.101	2.552	2.878
19	1.328	1.729	2.093	2.539	2.861
20	1.325	1.725	2.086	2.528	2.845
21	1.323	1.721	2.080	2.518	2.831
22	1.321	1.717	2.074	2.508	2.819
23	1.319	1.714	2.069	2.500	2.807
24	1.318	1.711	2.064	2.492	2.797
25	1.316	1.708	2.060	2.485	2.787
26	1.315	1.706	2.056	2.479	2.779
27	1.314	1.703	2.052	2.473	2.771
28	1.313	1.701	2.048	2.467	2.763

<i>p</i>	0.9	0.95	0.975	0.99	0.995
<i>v</i>					
29	1.311	1.699	2.045	2.462	2.756
30	1.310	1.697	2.042	2.457	2.750
31	1.309	1.696	2.040	2.453	2.744
32	1.309	1.694	2.037	2.449	2.738
33	1.308	1.692	2.035	2.445	2.733
34	1.307	1.691	2.032	2.441	2.728
35	1.306	1.690	2.030	2.438	2.724
36	1.306	1.688	2.028	2.434	2.719
37	1.305	1.687	2.026	2.431	2.715
38	1.304	1.686	2.024	2.429	2.712
39	1.304	1.685	2.023	2.426	2.708
40	1.303	1.684	2.021	2.423	2.704
45	1.301	1.679	2.014	2.412	2.690
50	1.299	1.676	2.009	2.403	2.678
55	1.297	1.673	2.004	2.396	2.668
60	1.296	1.671	2.000	2.390	2.660
65	1.295	1.669	1.997	2.385	2.654
70	1.294	1.667	1.994	2.381	2.648
75	1.293	1.665	1.992	2.377	2.643
80	1.292	1.664	1.990	2.374	2.639
85	1.292	1.663	1.988	2.371	2.635
90	1.291	1.662	1.987	2.368	2.632
95	1.291	1.661	1.985	2.366	2.629
100	1.290	1.660	1.984	2.364	2.626
125	1.288	1.657	1.979	2.357	2.616
150	1.287	1.655	1.976	2.351	2.609
200	1.286	1.653	1.972	2.345	2.601
∞	1.282	1.645	1.960	2.326	2.576

PDF of the t -distribution:
$$f_t(x; \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

5.2.17. Chi-Squared (χ^2) Distribution: Critical Values (Inverse Chi square table)

The table gives the values of x satisfying $P(X \leq x) = p$, where X is a random variable having the χ^2 distribution with ν degrees of freedom.




p	0.005	0.01	0.025	0.05	0.1	0.9	0.95	0.975	0.99	0.995	p
ν											ν
1	0.00004	0.0002	0.001	0.004	0.016	2.706	3.841	5.024	6.635	7.879	1
2	0.010	0.020	0.051	0.103	0.211	4.605	5.991	7.378	9.210	10.597	2
3	0.072	0.115	0.216	0.352	0.584	6.251	7.815	9.348	11.345	12.838	3
4	0.207	0.297	0.484	0.711	1.064	7.779	9.488	11.143	13.277	14.860	4
5	0.412	0.554	0.831	1.145	1.610	9.236	11.070	12.833	15.086	16.750	5
6	0.676	0.872	1.237	1.635	2.204	10.645	12.592	14.449	16.812	18.548	6
7	0.989	1.239	1.690	2.167	2.833	12.017	14.067	16.013	18.475	20.278	7
8	1.344	1.646	2.180	2.733	3.490	13.362	15.507	17.535	20.090	21.955	8
9	1.735	2.088	2.700	3.325	4.168	14.684	16.919	19.023	21.666	23.589	9
10	2.156	2.558	3.247	3.940	4.865	15.987	18.307	20.483	23.209	25.188	10
11	2.603	3.053	3.816	4.575	5.578	17.275	19.675	21.920	24.725	26.757	11
12	3.074	3.571	4.404	5.226	6.304	18.549	21.026	23.337	26.217	28.300	12
13	3.565	4.107	5.009	5.892	7.042	19.812	22.362	24.736	27.688	29.819	13
14	4.075	4.660	5.629	6.571	7.790	21.064	23.685	26.119	29.141	31.319	14
15	4.601	5.229	6.262	7.261	8.547	22.307	24.996	27.488	30.578	32.801	15
16	5.142	5.812	6.908	7.962	9.312	23.542	26.296	28.845	32.000	34.267	16
17	5.697	6.408	7.564	8.672	10.085	24.769	27.587	30.191	33.409	35.718	17
18	6.265	7.015	8.231	9.390	10.865	25.989	28.869	31.526	34.805	37.156	18
19	6.844	7.633	8.907	10.117	11.651	27.204	30.144	32.852	36.191	38.582	19
20	7.434	8.260	9.591	10.851	12.443	28.412	31.410	34.170	37.566	39.997	20
21	8.034	8.897	10.283	11.591	13.240	29.615	32.671	35.479	38.932	41.401	21
22	8.643	9.542	10.982	12.338	14.041	30.813	33.924	36.781	40.289	42.796	22
23	9.260	10.196	11.689	13.091	14.848	32.007	35.172	38.076	41.638	44.181	23
24	9.886	10.856	12.401	13.848	15.659	33.196	36.415	39.364	42.980	45.559	24
25	10.520	11.524	13.120	14.611	16.473	34.382	37.652	40.646	44.314	46.928	25
26	11.160	12.198	13.844	15.379	17.292	35.563	38.885	41.923	45.642	48.290	26
27	11.808	12.879	14.573	16.151	18.114	36.741	40.113	43.195	46.963	49.645	27
28	12.461	13.565	15.308	16.928	18.939	37.916	41.337	44.461	48.278	50.993	28
29	13.121	14.256	16.047	17.708	19.768	39.087	42.557	45.722	49.588	52.336	29
30	13.787	14.953	16.791	18.493	20.599	40.256	43.773	46.979	50.892	53.672	30
31	14.458	15.655	17.539	19.281	21.434	41.422	44.985	48.232	52.191	55.003	31
32	15.134	16.362	18.291	20.072	22.271	42.585	46.194	49.480	53.486	56.328	32
33	15.815	17.074	19.047	20.867	23.110	43.745	47.400	50.725	54.776	57.648	33
34	16.501	17.789	19.806	21.664	23.952	44.903	48.602	51.996	56.061	58.964	34
35	17.192	18.509	20.569	22.465	24.797	46.059	49.802	53.203	57.342	60.275	35
36	17.887	19.223	21.336	23.269	25.643	47.212	50.998	54.437	58.619	61.581	36
37	18.586	19.960	22.106	24.075	26.492	48.363	52.192	55.668	59.892	62.883	37
38	19.289	20.691	22.878	24.884	27.343	49.513	53.384	56.896	61.162	64.181	38
39	19.996	21.426	23.654	25.695	28.196	50.660	54.572	58.120	62.428	65.476	39
40	20.707	22.164	24.433	26.509	29.051	51.805	55.758	59.342	63.691	66.766	40
45	24.311	25.901	28.366	30.612	33.350	57.505	61.656	65.410	69.957	73.166	45
50	27.991	29.707	32.357	34.764	37.689	63.167	67.505	71.420	76.154	79.490	50
55	31.735	33.570	36.398	38.958	42.060	68.796	73.311	77.380	82.292	85.749	55
60	35.534	37.485	40.482	43.188	46.459	74.397	79.082	83.298	88.379	91.952	60
65	39.383	41.444	44.603	47.450	50.883	79.973	84.821	89.177	94.422	98.105	65
70	43.275	45.442	48.758	51.739	55.329	85.527	90.531	95.023	100.425	104.215	70
75	47.206	49.475	52.942	56.054	59.795	91.061	96.217	100.839	106.393	110.286	75
80	51.172	53.540	57.153	60.391	64.278	96.578	101.879	106.629	112.329	116.321	80
85	55.170	57.634	61.389	64.749	68.777	102.079	107.522	112.393	118.236	122.325	85
90	59.196	61.754	65.647	69.126	73.291	107.565	113.145	118.136	124.116	128.299	90
95	63.250	65.898	69.925	73.520	77.818	113.038	118.752	123.858	129.973	134.247	95
100	67.328	70.065	74.222	77.929	82.358	118.498	124.342	129.561	135.807	140.169	100

PDF of the Chi-Square distribution:

$$f_{\chi^2}(x; \nu) = \frac{x^{\nu/2-1} e^{-x/2}}{2^{\nu/2} \Gamma(\frac{\nu}{2})}$$

5.3.18. F-Distribution: Critical Values (Inverse F table)

The table gives the values of x satisfying $P(X \geq x) = p$, where X is a random variable having the F distribution formed by $X = \frac{\chi_1^2/\nu_1}{\chi_2^2/\nu_2}$, where χ_1 and χ_2 are χ^2 -distributed with ν_1 and ν_2 degrees of freedom respectively.



degrees of freedom in the numerator, ν_1

		degrees of freedom in the denominator, ν_2									
		p	1	2	3	4	5	6	7	8	9
1	.100	39.86	49.50	53.59	55.83	57.24	58.20	58.91	59.44	59.86	60.19
	.050	161.45	199.50	215.71	224.58	230.16	233.99	236.77	238.88	240.54	241.88
	.025	647.79	799.50	864.16	899.58	921.85	937.11	948.22	956.66	963.28	968.63
	.010	4052.2	4999.5	5403.4	5624.6	5763.6	5859.0	5928.4	5981.1	6022.5	6055.8
	.001	405284	500000	540379	562500	576405	585937	592873	598144	602284	605621
2	.100	8.53	9.00	9.16	9.24	9.29	9.33	9.35	9.37	9.38	9.39
	.050	18.51	19.00	19.16	19.25	19.30	19.33	19.35	19.37	19.38	19.40
	.025	38.51	39.00	39.17	39.25	39.30	39.33	39.36	39.37	39.39	39.40
	.010	98.50	99.00	99.17	99.25	99.30	99.33	99.36	99.37	99.39	99.40
	.001	998.50	999.00	999.17	999.25	999.30	999.33	999.36	999.37	999.39	999.40
3	.100	5.54	5.46	5.39	5.34	5.31	5.28	5.27	5.25	5.24	5.23
	.050	10.13	9.55	9.28	9.12	9.01	8.94	8.89	8.85	8.81	8.79
	.025	17.44	16.04	15.44	15.10	14.88	14.73	14.62	14.54	14.47	14.42
	.010	34.12	30.82	29.46	28.71	28.24	27.91	27.67	27.49	27.35	27.23
	.001	167.03	148.50	141.11	137.10	134.58	132.85	131.58	130.62	129.86	129.25
4	.100	4.54	4.32	4.19	4.11	4.05	4.01	3.98	3.95	3.94	3.92
	.050	7.71	6.94	6.59	6.39	6.26	6.16	6.09	6.04	6.00	5.96
	.025	12.22	10.65	9.98	9.60	9.36	9.20	9.07	8.98	8.90	8.84
	.010	21.20	18.00	16.69	15.98	15.52	15.21	14.98	14.80	14.66	14.55
	.001	74.14	61.25	56.18	53.44	51.71	50.53	49.66	49.00	48.47	48.05
5	.100	4.06	3.78	3.62	3.52	3.45	3.40	3.37	3.34	3.32	3.30
	.050	6.61	5.79	5.41	5.19	5.05	4.95	4.88	4.82	4.77	4.74
	.025	10.01	8.43	7.76	7.39	7.15	6.98	6.85	6.76	6.68	6.62
	.010	16.26	13.27	12.06	11.39	10.97	10.67	10.46	10.29	10.16	10.05
	.001	47.18	37.12	33.20	31.09	29.75	28.83	28.16	27.65	27.24	26.92
6	.100	3.78	3.46	3.29	3.18	3.11	3.05	3.01	2.98	2.96	2.94
	.050	5.99	5.14	4.76	4.53	4.39	4.28	4.21	4.15	4.10	4.06
	.025	8.81	7.26	6.60	6.23	5.99	5.82	5.70	5.60	5.52	5.46
	.010	13.75	10.92	9.78	9.15	8.75	8.47	8.26	8.10	7.98	7.87
	.001	35.51	27.00	23.70	21.92	20.80	20.03	19.46	19.03	18.69	18.41
7	.100	3.59	3.26	3.07	2.96	2.88	2.83	2.78	2.75	2.72	2.70
	.050	5.59	4.74	4.35	4.12	3.97	3.87	3.79	3.73	3.68	3.64
	.025	8.07	6.54	5.89	5.52	5.29	5.12	4.99	4.90	4.82	4.76
	.010	12.25	9.55	8.45	7.85	7.46	7.19	6.99	6.84	6.72	6.62
	.001	29.25	21.69	18.77	17.20	16.21	15.52	15.02	14.63	14.33	14.08
8	.100	3.46	3.11	2.92	2.81	2.73	2.67	2.62	2.59	2.56	2.54
	.050	5.32	4.46	4.07	3.84	3.69	3.58	3.50	3.44	3.39	3.35
	.025	7.57	6.06	5.42	5.05	4.82	4.65	4.53	4.43	4.36	4.30
	.010	11.26	8.65	7.59	7.01	6.63	6.37	6.18	6.03	5.91	5.81
	.001	25.41	18.49	15.83	14.39	13.48	12.86	12.40	12.05	11.77	11.54
9	.100	3.36	3.01	2.81	2.69	2.61	2.55	2.51	2.47	2.44	2.42
	.050	5.12	4.26	3.86	3.63	3.48	3.37	3.29	3.23	3.18	3.14
	.025	7.21	5.71	5.08	4.72	4.48	4.32	4.20	4.10	4.03	3.96
	.010	10.56	8.02	6.99	6.42	6.06	5.80	5.61	5.47	5.35	5.26
	.001	22.86	16.39	13.90	12.56	11.71	11.13	10.70	10.37	10.11	9.89
10	.100	3.29	2.92	2.73	2.61	2.52	2.46	2.41	2.38	2.35	2.32
	.050	4.96	4.10	3.71	3.48	3.33	3.22	3.14	3.07	3.02	2.98
	.025	6.94	5.46	4.83	4.47	4.24	4.07	3.95	3.85	3.78	3.72
	.010	10.04	7.56	6.55	5.99	5.64	5.39	5.20	5.06	4.94	4.85
	.001	21.04	14.91	12.55	11.28	10.48	9.93	9.52	9.20	8.96	8.75

PDF of the F -distribution:

$$f_F(x; \nu_1, \nu_2) = \frac{1}{B\left(\frac{\nu_1}{2}, \frac{\nu_2}{2}\right)} \left(\frac{\nu_1}{\nu_2}\right)^{\nu_1/2} x^{\nu_1/2-1} \left(1 + \frac{\nu_1}{\nu_2}x\right)^{-\frac{\nu_1+\nu_2}{2}} \quad 165$$

5.2.19. Product Moment Correlation Coefficient (PMCC): Critical Values (Inverse rho table)

The table gives the critical values, for different significance levels, of the Pearson's product moment correlation coefficient (PMCC), ρ , for varying sample sizes, n .

One tail Two tail	10% 20%	5% 10%	2.5% 5%	1% 2%	0.5% 1%	One tail Two tail
<i>n</i>						<i>n</i>
4	0.8000	0.9000	0.9500	0.9800	0.9900	4
5	0.6870	0.8054	0.8783	0.9343	0.9587	5
6	0.6084	0.7293	0.8114	0.8822	0.9172	6
7	0.5509	0.6694	0.7545	0.8329	0.8745	7
8	0.5067	0.6215	0.7067	0.7887	0.8343	8
9	0.4716	0.5822	0.6664	0.7498	0.7977	9
10	0.4428	0.5494	0.6319	0.7155	0.7646	10
11	0.4187	0.5214	0.6021	0.6851	0.7348	11
12	0.3981	0.4973	0.5760	0.6581	0.7079	12
13	0.3802	0.4762	0.5529	0.6339	0.6835	13
14	0.3646	0.4575	0.5324	0.6120	0.6614	14
15	0.3507	0.4409	0.5140	0.5923	0.6411	15
16	0.3383	0.4259	0.4973	0.5742	0.6226	16
17	0.3271	0.4124	0.4821	0.5577	0.6055	17
18	0.3170	0.4000	0.4683	0.5425	0.5897	18
19	0.3077	0.3887	0.4555	0.5285	0.5751	19
20	0.2992	0.3783	0.4438	0.5155	0.5614	20
21	0.2914	0.3687	0.4329	0.5034	0.5487	21
22	0.2841	0.3598	0.4227	0.4921	0.5368	22
23	0.2774	0.3515	0.4132	0.4815	0.5256	23
24	0.2711	0.3438	0.4044	0.4716	0.5151	24
25	0.2653	0.3365	0.3961	0.4622	0.5052	25
26	0.2598	0.3297	0.3882	0.4534	0.4958	26
27	0.2546	0.3233	0.3809	0.4451	0.4869	27
28	0.2497	0.3172	0.3739	0.4372	0.4785	28
29	0.2451	0.3115	0.3673	0.4297	0.4705	29
30	0.2407	0.3061	0.3610	0.4226	0.4629	30
31	0.2366	0.3009	0.3550	0.4158	0.4556	31
32	0.2327	0.2960	0.3494	0.4093	0.4487	32
33	0.2289	0.2913	0.3440	0.4032	0.4421	33
34	0.2254	0.2869	0.3388	0.3972	0.4357	34
35	0.2220	0.2826	0.3338	0.3916	0.4296	35
36	0.2187	0.2785	0.3291	0.3862	0.4238	36
37	0.2156	0.2746	0.3246	0.3810	0.4182	37
38	0.2126	0.2709	0.3202	0.3760	0.4128	38
39	0.2097	0.2673	0.3160	0.3712	0.4076	39
40	0.2070	0.2638	0.3120	0.3665	0.4026	40
41	0.2043	0.2605	0.3081	0.3621	0.3978	41
42	0.2018	0.2573	0.3044	0.3578	0.3932	42
43	0.1993	0.2542	0.3008	0.3536	0.3887	43
44	0.1970	0.2512	0.2973	0.3496	0.3843	44
45	0.1947	0.2483	0.2940	0.3457	0.3801	45
46	0.1925	0.2455	0.2907	0.3420	0.3761	46
47	0.1903	0.2429	0.2876	0.3384	0.3721	47
48	0.1883	0.2403	0.2845	0.3348	0.3683	48
49	0.1863	0.2377	0.2816	0.3314	0.3646	49
50	0.1843	0.2353	0.2787	0.3281	0.3610	50
60	0.1678	0.2144	0.2542	0.2997	0.3301	60
70	0.1550	0.1982	0.2352	0.2776	0.3060	70
80	0.1448	0.1852	0.2199	0.2597	0.2864	80
90	0.1364	0.1745	0.2072	0.2449	0.2702	90
100	0.1292	0.1654	0.1966	0.2324	0.2565	100

5.3. Hypothesis Testing

5.3.1. Decision and Estimation Theory

Definitions:

- For a set of i iid observations $\mathbf{x} = [x_1, x_2, \dots, x_i]$ of a random variable $\mathbf{X} = [X_1, X_2, \dots, X_i]$, the distribution of each X_i depends on x and some unknown parameter(s) θ .
- The estimate (decision) for θ is a function of the observations: $\hat{\theta}(\mathbf{x})$.
- The estimator (decision rule) for a parameter θ is given by $\hat{\Theta} = \hat{\theta}(\mathbf{X})$.

5.3.2. Bayesian Statistics

In Bayesian statistics, the unknown parameter θ is viewed as the value of a random variable Θ . The distribution of the sample is then interpreted as the conditional distribution $f_{\mathbf{X}|\Theta}(\mathbf{x} | \theta)$.

- Prior function: $f_{\Theta}(\theta)$ (prior to any measurements)
- Likelihood function: $f_{\mathbf{X}|\Theta}(\mathbf{x} | \theta)$.
- Posterior function: $f_{\Theta|\mathbf{X}}(\theta | \mathbf{x})$ (after the measurements)

5.3.3. Estimator Metrics

Maximum likelihood estimator (ML): $\hat{\theta}_{\text{ML}}(\mathbf{x}) = \underset{\theta}{\operatorname{argmax}} f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)$

Maximum a posteriori estimator (MAP): $\hat{\theta}_{\text{MAP}}(\mathbf{x}) = \underset{\theta}{\operatorname{argmax}} f_{\Theta|\mathbf{X}}(\theta|\mathbf{x})$

Minimum mean squared error (MMSE): $\hat{\theta}_{\text{MMSE}}(\mathbf{x}) = \mathbb{E}[\Theta|\mathbf{X} = \mathbf{x}] = \int \theta f_{\Theta|\mathbf{X}}(\theta|\mathbf{x}) d\theta$
which minimises the value of $\mathbb{E}[(\theta - \Theta)^2|\mathbf{X} = \mathbf{x}]$.

The posterior function is obtained from Bayes' rule (see Section 5.2.9) as

$$f_{\Theta|\mathbf{X}}(\theta|\mathbf{x}) = \frac{f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta) \cdot f_{\Theta}(\theta)}{f_{\mathbf{X}}(\mathbf{x})}.$$

Note that if $f_{\Theta}(\theta) = \text{constant}$ (i.e. Θ is uniformly distributed) then $\hat{\theta}_{\text{ML}} = \hat{\theta}_{\text{MAP}}$.

5.3.4. Principles of Hypothesis Testing

Hypothesis testing involves assessing the probability of observing a given dataset, given the assumption of a particular null hypothesis (H_0).


Alternative hypothesis (H_1): a potential alternative statement to explain a dataset, attributed to an external effect influencing the data.

- p-value (p): the probability of obtaining a result at least as extreme as the given data, under the assumption that H_0 is true.
- $p = P(\text{data or further from } H_0 \mid H_0)$. The p-value is **not** the probability of H_0 .
- Significance level (α): if $p \leq \alpha$, the result is unlikely to happen under H_0 so we reject H_0 and the result is said to be statistically significant.

Analysis of a Hypothesis Test

If H_1 is of the form $\theta \neq a$, the p-value must be found on the basis of a two-tailed test in which the critical region is central to the distribution of θ under H_0 i.e. use an effective significance level of $\alpha/2$ in each tail.

The diagram illustrates the conditional distributions $f_{\theta}(\theta \mid H_0)$ and $f_{\theta}(\theta \mid H_1)$.



- **Type I error:** falsely **rejecting** H_0 when in reality, H_0 is true ($\alpha = P(\text{reject } H_0 \mid H_0)$)
- **Type II error:** falsely **accepting** H_0 when in reality, H_1 is true ($\beta = P(\text{accept } H_0 \mid H_1)$)
- **Critical (rejection) region** under H_0 : $\theta \geq \theta^*$ such that $P(\theta \geq \theta^* \mid H_0) \leq \alpha$
- **Power** of a test = $1 - \beta$

5.3.5. Specified Distribution Hypothesis Tests

After sampling a random variable X , a test can be performed to investigate whether the sample is expected under a particular null hypothesis distribution. Commonly used for Binomial, Normal...

Assumptions: all assumptions made by the underlying distribution being proposed.

The p -value is the probability of observing the sample **or more extreme** given the null hypothesis.

Example (Binomial): a coin is flipped 16 times and lands tails 11 times. Investigate at 5% significance whether the coin is fair or biased. Find the power of the test if the true probability of tails is 0.6.

- Assume that the number of tails X is binomially distributed as $X \sim B(16, p)$ (independent trials, binary outcomes, constant probability of success p).
- $H_0: p = 0.5$ (the coin is unbiased / fair).
- $H_1: p \neq 0.5$ (the coin is biased / unfair) (two-tailed test).
- Under H_0 , $X \sim B(16, 0.5)$. The test statistic is $X = 11$.
- $P(X \geq 11 | H_0) = 0.1051$ (p -value).
- Since $0.1051 > 0.025$ (half the significance level), we accept H_0 .
- There is insufficient evidence to suggest that the coin is biased.
- Critical region: $X \geq 13$ and $X \leq 3$, Acceptance region: $4 \leq X \leq 12$. Critical values: 3 and 13.
- $P(\text{Type I error}) = 0.05$ (significance level).
- $P(\text{Type II error} | \text{true } p = 0.6) = P(\text{accept } H_0 | p = 0.6) = P(4 \leq X \leq 12 | p = 0.6) = 0.9339$.
- Power of the test = $1 - 0.9339 = 0.0661$.

Example (Normal): a particular make and model of car is known to have an average fuel mileage of 25.0 miles per gallon (mpg) with variance 6.1 mpg². When a new additive is added to the fuel of 35 cars, their mean mileage rises to 25.9 mpg. Test at 5% significance whether the additive increased the mean mileage, and construct a 99% confidence interval for the population mean mileage with the additive.

- Assume that the mileage X is Normally distributed as $X \sim N(\mu, 6.1)$.
- $H_0: \mu = 25$ (the mileage has not increased).
- $H_1: \mu > 25$ (the mileage has increased).
- Under H_0 , $X \sim N(25, 6.1)$ and therefore the sample mean $\bar{X} \sim N(25, \frac{6.1}{35}) = N(25, 0.1743)$.
- $p = P(\bar{X} > 25.9 | H_0) = 0.0156$ (or using z -statistic: $z = \frac{25.9 - 25}{\sqrt{0.1743}} = 2.1557 \rightarrow p = 1 - \Phi(z) = 0.0156$)
- Since $0.0156 < 0.05$ (significance level), we reject H_0 .
- There is sufficient evidence to suggest the population mean mileage has increased.
- Critical region: $z > \Phi^{-1}(0.95) \rightarrow z > 1.6449 \rightarrow \bar{X} > 25 + 1.6449 \times \sqrt{0.1743} \rightarrow \bar{X} > 25.6867$.
Acceptance region: $\bar{X} < 25.6867$. Critical value: 25.6867.
- $P(\text{Type I error}) = 0.05$ (significance level).
- 99% confidence interval for new mean: $z = \Phi^{-1}(0.995) = 2.5758$
 $\rightarrow \mu \in (25.9 - 2.5758\sqrt{0.1743}, 25.9 + 2.5758\sqrt{0.1743}) \rightarrow \mu \in (24.8246, 26.9754)$.
This interval captures the true mean mileage with fuel additives with 99% confidence.

5.3.6. Chi-Square Tests

Chi-Square Test for Goodness of Fit or Association

A Chi-Square (χ^2) test can be performed to investigate whether a discrete categorical random variable X obeys a particular distribution, or for association between two categorical variables X and Y (a non-parametric test).

H_0 : there is no association between X and Y

H_1 : there is an association between X and Y

Degrees of freedom for an $a \times b$ observed contingency table O : $\nu = (a - 1)(b - 1)$

Expected value under H_0 (no association): $E_{ij} = \frac{(\text{row total})_i \times (\text{column total})_j}{\text{grand total}}$

Test statistic: $\chi^2 = \sum_{i,j} \frac{(O_{ij} - E_{ij})^2}{E_{ij}}$

Critical value, χ^2_c , found from the table in Section 5.2.12.

If $\chi^2 < \chi^2_c$ then we accept H_0 (evidence to suggest independence)

If $\chi^2 > \chi^2_c$ then we reject H_0 (evidence to suggest association)

Modifications to the test methodology include:

- Pooling: if observed frequencies are less than 5, rows / columns should be combined.
- Yates' continuity correction: if the contingency table is 2×2 (1 degree of freedom) and at least one frequency is less than 5 (and so cannot be pooled), the test statistic is

$$\chi^2_{\text{Yates}} = \sum_{i,j} \frac{(|O_{ij} - E_{ij}| - 0.5)^2}{E_{ij}}$$

Yates' correction is not universally accepted. Applying it decreases the test statistic, increases the p-value and increases the probability of a type II error.

Software implementations:

Python: `scipy.stats.chisquare(f_obs, f_exp=None)`

R: `chisq.test(data) # `correct=False` to disable Yates`

Excel: `=CHITEST(obs_range, exp_range)`

Chi-Square Test for Variance

After sampling a random variable X , a Chi-Square (χ^2) test can be performed to investigate whether X has a given population variance.

$H_0: \sigma^2_X = a^2$

$H_1: \sigma^2_X \neq a^2$ (two tailed test) or $\sigma^2_X < a^2$ or $\sigma^2_X > a^2$ (one tailed test).

Test statistic: $\chi^2 = \frac{(N - 1) s^2}{a^2}$ (s^2 : sample variance)

Critical value, χ^2_c , found from the table in Section 5.2.12.

If $\chi^2 < \chi^2_c$ then we accept H_0 ; If $\chi^2 > \chi^2_c$ then we reject H_0

5.3.7. Inferential Parametric Tests: t -Tests and Analysis of Variance (ANOVA)

One-Sample t -Test for the Mean of a Normal Distribution

One sample, Normal distribution, unknown variance. $H_0: \mu = a$; $H_1: \mu \neq a$ (if two-tail).

Test statistic: $t = \frac{\bar{x} - a}{S/\sqrt{n}}$ (\bar{x} : sample mean, S : sample std.dev, n : sample size)

The test statistic has a t -distribution with $\nu = n - 1$ degrees of freedom.

Critical values in table in Section 5.2.16.

Student's Two-Sample t -Test for Independent Means of Homoscedastic Normal Distributions

Two samples A and B , Normal distributions, equal variance. $H_0: \mu_A = \mu_B$; $H_1: \mu_A \neq \mu_B$ (if two-tail)

Test statistic: $t = \frac{\bar{x}_A - \bar{x}_B}{\sqrt{\frac{S^2}{n_A} + \frac{S^2}{n_B}}}$ ($\bar{x}_A = \frac{\sum x_i}{n_A}$: sample mean of A , $S^2 = \frac{\sum(x_i - \bar{x}_A)^2 + \sum(x_j - \bar{x}_B)^2}{n_A + n_B - 2}$: pooled sample variance)

The test statistic has a t -distribution with $\nu = n_A + n_B - 2$ degrees of freedom.

Critical values in table in Section 5.2.16.

Welch's Two-Sample t -Test for Independent Means of Heteroscedastic Normal Distributions

Two samples A and B , Normal distributions, unequal variance. $H_0: \mu_A = \mu_B$; $H_1: \mu_A \neq \mu_B$ (if two-tail)

Test statistic: $t = \frac{\bar{x}_A - \bar{x}_B}{\sqrt{\frac{S_A^2}{n_A} + \frac{S_B^2}{n_B}}}$. The test statistic has a t -distribution with $\nu = \frac{\left(\frac{S_A^2}{n_A} + \frac{S_B^2}{n_B}\right)^2}{\frac{S_A^4}{n_A^2(n_A-1)} + \frac{S_B^4}{n_B^2(n_B-1)}}$ dof.

Critical values in table in Section 5.2.16.

Fisher's ANOVA Test for Independent Means of Homoscedastic Normal Distributions

N variables $\{X_j\}_i$ (i th observation of X_j , $1 \leq j \leq N$), Normal distributions, equal variance. H_0 : all μ_i equal; H_1 : not all μ_i equal.

Test statistic: $F = \frac{S^2_{between}}{S^2_{within}}$ ($S^2_{between} = \frac{\sum_j (\bar{X}_j - \bar{X})^2}{\left(\sum_j n_j\right) - N}$, $S^2_{within} = \frac{\sum_{j,i} (X_{ji} - \bar{X}_j)^2}{N - 1}$)

The test statistic has an F -distribution with $\nu_1 = \left(\sum_j n_j\right) - N$ and $\nu_2 = N - 1$ degrees of freedom.

Critical values in table in Section 5.2.18. **Post-hoc analysis:** Tukey's range test.

Welch's ANOVA Test for Independent Means of Heteroscedastic Normal Distributions

N variables $\{X_j\}_i$ (i th observation of X_j , $1 \leq j \leq N$), Normal distributions, unequal variance. H_0 : all μ_i equal; H_1 : not all μ_i equal.

Test statistic: $F = \frac{S^2_{between}}{S^2_{within}}$ ($S^2_{between} = \frac{\sum_j \frac{n_j}{s_j^2} \left(\bar{X}_j - \frac{\sum_j \frac{n_j \bar{X}_j}{s_j^2}}{\sum_j \frac{n_j}{s_j^2}} \right)^2}{N - 1}$, $S^2_{within} = 1 + \frac{2(N-2)}{N^2-1} T$, $T = \sum_j \frac{1}{n_j-1} \left(1 - \frac{\frac{n_j}{s_j^2}}{\sum_j \frac{n_j}{s_j^2}} \right)^2$)

The test statistic has an F -distribution with $\nu_1 = N - 1$ and $\nu_2 = \frac{N^2-1}{3T}$ degrees of freedom.

Critical values in table in Section 5.2.18. **Post-hoc analysis:** Games-Howell test.

5.3.8. Other Inferential Hypothesis Tests

Mann-Whitney U -Test (Wilcoxon Rank Sum Test)

Two samples A and B , unknown distributions. H_0 : A and B from same distributions; H_1 : A and B from different distributions.

Kruskal-Wallis H -Test (Non-Parametric One-Way ANOVA)

N variables $\{X_j\}_i$ (i th observation of X_j , $1 \leq j \leq N$), unknown distributions. H_0 : all distributions equal; H_1 : not all distributions equal. Post-hoc analysis: Dunn's test or Conover-Iman test.

Shapiro-Wilk Test for Normality

One sample X , unknown distribution. H_0 : X has a Normal distribution, H_1 : X does not have a Normal distribution.

Pearson's Test for Linear Correlation

Paired (bivariate) dataset $\mathbf{X} = \{X, Y\}$, Normal joint distribution. H_0 : X and Y are uncorrelated, H_1 : X and Y are correlated.

Test statistic: Pearson's PMCC:

Critical values in table in Section 5.2.19.

$$r_{xy} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}, \quad |r_{xy}| \leq 1.$$

Spearman's Rank-Order Test for Monotonic Correlation

Paired (bivariate) dataset $\mathbf{X} = \{X, Y\}$, unknown distributions. H_0 : X and Y are uncorrelated, H_1 : X and Y are correlated.

Test statistic: Spearman's rho,

Similar alternative: Kendall's Tau Test.

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)} \quad (d_i: \text{rank difference of observation } i) \\ \text{rank is the index in the ordered list.}$$

Levene's Test for Homoscedasticity (Equal Variances)

N variables $\{X_j\}_i$ (i th observation of X_j , $1 \leq j \leq N$), unknown distributions. H_0 : all distributions equal; H_1 : not all distributions equal. Post-hoc analysis: Dunn's test or Conover-Iman test.

Similar alternative: Brown-Forsythe test.

Cohen's d Statistic for Mean Effect Size

The d metric is related to the t -statistic in testing for independent means by $d = \sqrt{\frac{1}{N_1} + \frac{1}{N_2}} \times t$.

It is typically used in the context of quantifying the effect size of a test group against a control group.

Egger's Regression Test for Intervention Effects

For a collection of univariate datasets, a funnel plot is made of standard error σ_x / \sqrt{N} against a measure of effect size: depending on context, this may be a raw mean value, a correlation coefficient, odds ratio, or Cohen's d metric. One point is made per dataset. H_0 : funnel plot regression line is vertical (implies no bias), H_1 : funnel plot regression line is sloped (implies bias). Egger's test is commonly used in meta-analyses to test for publication bias.

5.3.9. Common Fallacies in Statistical Inference, Interpretation and Discourse

Interpretation of Statistical Tests: issues arising when concluding and communicating.

- **Texas Sharpshooter:** the cherry-picking of a cluster of data to fit a conclusion, or asserting that a pattern has an underlying cause other than randomness. A related concept in data misuse is 'p-hacking', in which the same (or slightly modified) tests are conducted on the same dataset until statistical significance is found (the multiple comparisons problem).
- **False Cause:** asserting that correlation implies causation, rather than randomness or a common cause.
- **Gambler's Fallacy:** assuming that the outcome of an event occurring after a series of the event has already been observed is lower than observing the event in general, when in fact they are independent.
- **Prosecutor's Fallacy:** assuming that the probability of observing an outcome given some evidence is the same as the probability of observing the evidence given the outcome.
- **Composition Fallacy:** assuming that the properties of the parts of a system completely determine those of the whole system, when they may be different (e.g. interaction effects, 'emergent properties').
- **Slippery Slope:** asserting that if one event happens, then a subsequent chain of events will also happen, without clearly establishing the validity of these links.
- **Begging the Question:** presenting a circular argument; presupposing the conclusion within the premise.
- **Ambiguity:** using language with multiple meanings, from which readers from different target audiences may interpret in different ways. Can occur when using terms of art.
- **Confirmation Bias:** favouring arguments and evidence which align with a person's existing beliefs while downplaying opposing evidence, regardless of its merit.
- **Post-hoc Rationalisation:** constructing (often improvised) arguments to justify behaviour or beliefs that are otherwise incompatible with one's beliefs. Often used when one is experiencing 'cognitive dissonance', in which a person simultaneously supports two logically irreconcilable beliefs, sometimes without being consciously aware of it. Also used to justify hypocrisy.

Methodology of Experiments and Studies: issues arising when conducting a study.

- **Loaded Question:** posing a question with a strong built-in bias towards a known outcome, which can hamper the validity of the testing methodology.
- **Survivorship Bias:** the sample under study may be self-selecting, making the results unrepresentative.
- **Observer Effect:** subjects under study may alter their behaviour if they know they are under study, in captivity, or respond differently depending on who is studying them.
- **Placebo Effect:** common in medicine. If subjects are told they can expect to see an effect from taking a treatment, they may genuinely experience the effect, even if the treatment itself does nothing.
- **Replication Crisis:** the observation that many studies, particularly in the social sciences which study complex population-level interactions, can often not be reproduced precisely. However, this does not imply the results are always invalid, rather, if different studies can investigate hypotheses from different perspectives and come to similar conclusions (triangulation), it implies that there is a deeper effect at play. Neglecting this nuance and extrapolating it to where it does not apply can lead to undue public distrust in the scientific method and the body of science as a whole.
- **Confounding variables:** factors which may influence a study but were not controlled for, either because they were not considered when the study was conducted or the control measures taken were inadequate to suppress their influence.

Argumentation and Discourse: logical fallacies and rhetorical techniques.

- **Strawman:** misrepresenting a statement in order to make it easier to argue against.
- **Anecdotes:** personal experiences or isolated incidents are statistically meaningless (typically $n \sim 1$ sample size) but are often far more compelling due to their tangibility or use of emotive language.
- **Appeal to Authority:** claims made by a perceived authority should not be considered valid based solely on their status of authority, but rather the rigorousness of their investigative methodology and access to empirical evidence e.g. peer-reviewed scientific literature.
- **Bandwagon:** the validity of a claim is not inherently dependent on how popular the claim is.
- **False Dichotomy:** asserting that there are only two possible outcomes (binary decision), when there may be more than two. The reverse of this is assuming a middle ground in a *true* dichotomy.
- **Non Sequitur:** any statement that can be demonstrated to be formally illogical or self-contradictory.
- **Slander / Libel:** public defamation by making false statements aimed at damaging one's reputation.
- **Burden of Proof:** the burden of proof lies with the individual making the claim against the current consensus i.e. what is presented without evidence can be dismissed without evidence. This is sometimes referred to as '(dis)proof beyond reasonable doubt'. If there is no current consensus, then all sides have a burden of proof.
- **Argumentum Ad Hominem:** attempting to discredit an opposing view by attacking irrelevant qualities of the person (whether true or false) who holds that view, without attacking the view itself.
- **Gish Gallop:** presenting a large number of claims in a short amount of time, making it seem as if one has an endless list of strong arguments, without allowing time to respond, and without explaining anything that could reveal that the arguments are not independent and/or strong.
- **Tu Quoque (Whataboutism):** claiming that one's opponent is a hypocrite because they committed the same act that one is being accused of, without actually defending oneself against the accusation. Whataboutism is the general propagandistic tactic of diverting attention to another scenario, without elaborating on whether such a comparison is valid to make.
- **Motte and Bailey Fallacy:** presenting a more outlandish (less well-supported) claim before falling back to a more well-established claim once it is criticised, and implying they use the same reasoning.
- **Socratic Method:** the use of open-ended questions and well-defined terminology to promote a non-confrontational discussion where opinionated people can reflect on their own perspectives. It can challenge presuppositions and expose unrealised self-contradictions.
- **Hegel's Dialectics:** presenting an initial idea (thesis), a contradictory idea (antithesis) and a higher-level resolution that integrates ideas discussed in each (synthesis).

5.3.10. The Scientific Method

The scientific method is the empirical process of reliably acquiring new knowledge about the natural world.

1. **Question:** identify something in the natural world, and ask a question about it.
2. **Fact-Finding:** consult existing scientific literature to research the topic at hand. Gather preliminary information that will be useful in studying the topic.
3. **Hypothesis:** formulate a potential explanation or answer to the question based on initial knowledge gained.
4. **Predictions:** before investigations begin, make testable, falsifiable predictions as to what the expected outcome would be if the hypothesis put forward is correct.
5. **Test:** design an experiment to investigate the question. Conduct the experiment in a safe, ethical and reproducible manner to investigate the question and record all observations.
6. **Analysis:** process the results to obtain useful data. If appropriate, perform statistical tests to quantify the likelihood of these results under a null and alternative hypothesis.
7. **Interpretations:** draw conclusions from the data analysis. These conclusions may serve as the starting point for new investigations.

Writing Scientific Literature:

For experimental work, the paper should outline the ‘story’ of how the topic is introduced:

1) abstract (succinct statement of the problem, approach and results), **2) introduction**, **3) materials and methods**, **4) results and discussion**, **5) conclusions**, **6) references** (from existing primary scientific literature, cited in a standard style). Sections are field-dependent.

Submit the work conducted in the form of a paper to a peer-reviewed journal. Designate a ‘corresponding author’ who can be contacted to answer questions about the work. When in proceedings, respond to suggestions and criticism from peer-reviewers and be open to assisting others in replicating your work.

Reading Scientific Literature:

An often useful approach when researching a topic comprehensively is to find a ‘review’ paper of the topic via Google Scholar. For papers, read: 1) abstract, 2) look at the figures, 3) conclusions, 4) the rest of the paper, 5) search for author’s discussions of their work in other sources. If seeking to examine methodology, check for any supplementary materials.

5.4. Stochastic Processes, Signals and Information Theory

5.4.1. Identities for Random Vectors

$$\mathbf{X} = \{X_1 \ X_2 \ \dots\}$$

These identities hold when X_i are scalar RVs. If X_i are n -vectors, then variances should be divided by $\frac{1}{n}$ for population quantities and by $\frac{1}{n-1}$ for sample quantities (as in Section 5.4.1).

- Covariance matrix (autocovariance): $\Sigma_{\mathbf{X}\mathbf{X}} = \text{Var}[\mathbf{X}] = \text{Cov}[\mathbf{X}, \mathbf{X}] = E[\mathbf{X}\mathbf{X}^T] - E[\mathbf{X}] E[\mathbf{X}]^T$; $\Sigma_{ij} = \text{Cov}[X_i, X_j]$
- Cross-Covariance (joint variance): $\Sigma_{\mathbf{X}\mathbf{Y}} = \text{Cov}[\mathbf{X}, \mathbf{Y}] = E[\mathbf{X}\mathbf{Y}^T] - E[\mathbf{X}] E[\mathbf{Y}]^T$; $\Sigma_{ij} = \text{Cov}[X_i, Y_j]$
- Autocorrelation and cross-correlation: $\mathbf{R}_{\mathbf{X}\mathbf{X}} = E[\mathbf{X}\mathbf{X}^T]$; $\mathbf{R}_{\mathbf{X}\mathbf{Y}} = E[\mathbf{X}\mathbf{Y}^T]$; $(\mathbf{R}_{\mathbf{X}\mathbf{Y}})_{ij} = \text{Corr}[X_i, Y_j]$
- Covariance of a sum: if $\mathbf{Z} = \mathbf{X} + \mathbf{Y}$ then $\Sigma_{\mathbf{Z}\mathbf{Z}} = \Sigma_{\mathbf{X}\mathbf{X}} + \Sigma_{\mathbf{X}\mathbf{Y}} + \Sigma_{\mathbf{Y}\mathbf{X}} + \Sigma_{\mathbf{Y}\mathbf{Y}}$ (note that $\Sigma_{\mathbf{X}\mathbf{Y}} = \Sigma_{\mathbf{Y}\mathbf{X}}^T$)
- Covariance of a transformation: if $\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{b}$ (\mathbf{A}, \mathbf{b} : constants) then $\Sigma_{\mathbf{Y}\mathbf{Y}} = \mathbf{A}\Sigma_{\mathbf{X}\mathbf{X}}\mathbf{A}^T$

5.4.2. Convolution, Cross-Correlation and Autocorrelation on LTI Systems

For discrete signals f_n and g_n , and continuous signals $f(t)$ and $g(t)$ (where $t = nT$ and T is the sampling period), and random variables X and Y :

	Convolution $f * g$	Cross-Correlation $f \star g$	Autocorrelation $f \star f$
Discrete	$\sum_{m=-\infty}^{\infty} f_m g_{n-m}$	$\sum_{m=-\infty}^{\infty} f_m g_{n+m}$	$\sum_{m=-\infty}^{\infty} f_m f_{n+m}$
Continuous	$\int_{-\infty}^{\infty} f(\tau) g(t - \tau) d\tau$	$\int_{-\infty}^{\infty} f(\tau) g(t + \tau) d\tau$	$\int_{-\infty}^{\infty} f(\tau) f(t + \tau) d\tau$
Stochastic	$E[X_t Y_{t-\tau}^*]$	$E[X_t Y_{t+\tau}^*]$	$E[X_t X_{t+\tau}^*]$

For the convolution theorem as it applies to discrete signals via the Z-transform and to continuous signals via the Fourier transform and Laplace transform, see Section 3.4.

5.4.3. Discrete Multidimensional Convolution

Let \mathbf{X} be an $M \times N$ matrix ($x(i, j) = X_{ij}$) and \mathbf{H} be a $K \times L$ matrix ($h(u, v) = H_{uv}$).

The convolution $\mathbf{Y} = \mathbf{H} * \mathbf{X}$ ($y = x * h$) is a $(M - K + 1) \times (N - L + 1)$ matrix, where

$$y(i, j) = \sum_{k=0}^{K-1} \sum_{l=0}^{L-1} h(k, l) x(i - k, j - l), \quad \text{for } 1 \leq i \leq M - K + 1, 1 \leq j \leq N - L + 1.$$

Typical application: \mathbf{X} is a general input, \mathbf{H} is the impulse response of a filter, \mathbf{Y} is the output.

5.4.3. Correlation Theorems

Let $X(\omega)$ and $Y(\omega)$ be the Fourier transforms of the time-domain signals $x(t)$ and $y(t)$.

Cross-Correlation Theorem: the cross-spectral density $S_{XY}(\omega) = |X(\omega) Y(\omega)|$ of two signals $x(t)$ and $y(t)$, and the cross-correlation of $x(t)$ and $y(t)$, form a Fourier inverse pair:

$$\int_{-\infty}^{\infty} (x \star y)(t) e^{-j\omega t} dt = S_{XY}(\omega) \qquad \int_{-\infty}^{\infty} S_{XY}(\omega) e^{j\omega t} d\omega = (x \star y)(t)$$

forward transform inverse transform

Wiener-Khinchin Theorem: the power spectral density $S_{XX}(\omega) = |X(\omega)|^2$ of a signal $x(t)$ and the autocorrelation of $x(t)$ form a Fourier-inverse pair:

$$\int_{-\infty}^{\infty} (x \star x)(t) e^{-j\omega t} dt = S_{XX}(\omega) \qquad \int_{-\infty}^{\infty} S_{XX}(\omega) e^{j\omega t} d\omega = (x \star x)(t)$$

forward transform inverse transform

The Wiener-Khinchin theorem is a special case of the cross-correlation theorem with $x = y$.

5.4.4. FIR and IIR Filters

The pulse function is defined as $\delta_k = \{1 \text{ if } k = 0 \text{ else } 0\} = \{1, 0, 0, 0, \dots\}$
(Kronecker Delta; discrete version of Dirac Delta function).

A discrete-time system has a transfer function given by the Z -transform of its pulse response.
(Analogous to continuous-time transfer functions as the Laplace transform of the impulse response).

For a discrete-time system (digital filter) with transfer function $G(z)$, input u_k and output y_k :

- $G(z)$ is the Z -transform of g_k (the pulse response when $u_k = \delta_k$).
- For a general input, the output is given by $y_k = (g_k * u_k)$ (discrete convolution theorem).
- Causal system: if $g_k = 0$ for all $k < 0$. All physically realisable systems are causal, in which k represents a discretisation of real time.
- Finite impulse response (FIR): if $g_k = 0$ for all $k > n$ for some smallest finite n . Otherwise, it is an Infinite impulse response (IIR) filter.
- Stability: a system is BIBO stable if, for any bounded $\{u_k\}$, the output $\{y_k\}$ is bounded. (A signal $\{u_k\}$ is bounded such that $|u_k| < M$ for some positive M for all k .)
- Step response: if $u_k = 1$ for all k then $y_k \rightarrow G(1)$ as $k \rightarrow \infty$ (final value theorem).
- Frequency response: if $u_k = \cos k\theta$ then at steady state, $y_{ss}(k) = |G(e^{j\theta})| \cos(k\theta + \angle G(e^{j\theta}))$.
- Causal system transfer function: $G(z) = \sum_{k=0}^{\infty} g_k z^{-k}$ (For IIR, all poles are at $z = 0$.)


Stability criterion: a rational TF $G(z) = \frac{n(z)}{d(z)} = \frac{\sum_{k=0}^m b_k z^{m-k}}{\sum_{k=0}^n a_k z^{n-k}} = \frac{b_0 z^m + b_1 z^{m-1} + \dots + b_m}{z^n + a_1 z^{n-1} + \dots + a_n}$ must have $m \leq n$.

For such a system, $G(z)$ is stable, all of the roots p_i of $d(z)$ satisfy $|p_i| < 1$ and $\sum_{k=0}^{\infty} |g_k|$ is finite.

A stable filter has a decaying transient response, so that its steady state is independent of the initial conditions. Any linear filter can be written as $A(z) Y(z) = B(z) U(z) + C(z, y_i)$, where C accounts for the initial conditions.

5.4.6. Discrete-Time Markov Chains (DTMCs)

A Markov process is a stochastic process in which the distribution of the next state is a function of only the current state, and not the previous states: $p(X_{n+1} | X_1, X_2, \dots, X_n) = p(X_{n+1} | X_n)$, where $p(X)$ is a **row** vector of probabilities of the random variable X being in each state.



		to		
		State 1	State 2	State 3
from	State 1	p_{11}	p_{12}	p_{13}
	State 2	p_{21}	p_{22}	p_{23}
	State 3	p_{31}	p_{32}	p_{33}

Definitions

- State space: the enumeration of the different states i.e. the domain of $X_n \in S$
- Absorbing state: if $p_{ii} = 1$ (on the leading diagonal of \mathbf{M}) then state i is an absorbing state.
- Recurrent set / Equivalence class: a set of states within which any state can reach any other state
- Irreducible chain: if all states are recurrent (i.e. no absorbing states; one equivalence class)
- Periodicity: for a chain of period δ , the nonzero eigenvalues of \mathbf{M} are the δ th roots of unity.
- Regular ergodic: an irreducible and aperiodic chain, which has limiting (stationary) distribution.

Probability Relationships

- Transition matrix: $(\mathbf{M})_{ij} = p_{ij} = P(X_{n+1} = j | X_n = i)$
- Columns of \mathbf{M} sum to 1: $\sum_j p_{ij} = 1$. (right stochastic matrix)
- Joint distribution: $p(X_0 = i_0, X_1 = i_1, \dots, X_n = i_n) = P(X_0 = i_0) \times p_{i_0 i_1} p_{i_1 i_2} \dots p_{i_{n-1} i_n}$
- State transition probabilities: $p(X_{n+1}) = p(X_n) \mathbf{M}$ and $p(X_{n+k}) = p(X_n) \mathbf{M}^k$
- Higher order transition probability: $p_{ij}^{(m+n)} = (\mathbf{M}^{m+n})_{ij} = \sum_k p_{ik}^{(m)} p_{kj}^{(n)}$ (Chapman-Kolmogorov equation)
- Stationary (ergodic) state: if $p(X_n) = \boldsymbol{\pi}$ then $\boldsymbol{\pi} \mathbf{M} = \boldsymbol{\pi}$ ($\boldsymbol{\pi}$ is an eigenvector of \mathbf{M} with eigenvalue 1)
- Unconditional probability: $P(X_n = j) = \sum_i p_{ij}^{(n)} P(X_0 = i) = \text{mean value of column } j \text{ in } \mathbf{M}^n$

First Step Transition Analysis / Waiting Time Problems

- Transitions between neighbouring states $i \rightarrow i + 1$ occur with time $T = \text{Geo}(p_{i(i+1)})$, so $E[T] = 1 / p_{i(i+1)}$.
- Expected steps required for a transition from i to j : $\mu_{ij} = E[\min(n \geq 1: X_n = j) | X_0 = i] = 1 + \sum_{k \neq j} p_{ik} \mu_{kj}$.


5.4.7. Continuous-Time Markov Chains (CTMCs, Discrete State Space)

The Markov matrix \mathbf{Q} for a continuous-time process is infinite dimensional: $\mathbf{x}' = \mathbf{x} \mathbf{Q}$

- $x_n(t) = P_n(t) = P(X(t) = n)$, $\mathbf{x} = [x_1(t), x_2(t), \dots]$ (**row vector**) and $(\mathbf{Q})_{ij} = \frac{\partial(x_j')}{\partial x_i}$. Entries of \mathbf{x} sum to 1.
- Rows of \mathbf{Q} sum to zero: $\sum_j (\mathbf{Q})_{ij} = 0$ (probability mass conserved).
- Discrete state space: $X(t)$ takes discrete values; $X \in \{0, 1, 2, \dots, n, \dots\}$ (to infinity, in general)
- Continuous state space: $X(t)$ takes continuous values (a range); $X \in S$.

Note that in solving $\mathbf{x}\mathbf{Q} = 0$ for the stationary state, one equation obtained from the columns of \mathbf{Q} is degenerate, and should be replaced with $\sum_j x_j = 1$ to give a determinate system.

The Birth-Death Process: the state $X(t)$ represents the number n of some entity at time t



- Birth: transition from state n to state $n + 1$, occurring at a rate λ_n per unit time ($T_{n \rightarrow n+1} \sim \text{Exp}(\lambda_n)$).
- Death: transition from state n to state $n - 1$, occurring at a rate μ_n per unit time ($T_{n \rightarrow n-1} \sim \text{Exp}(\mu_n)$).
- Transitioning to the same state is also possible in general, with rate $1 - \lambda_n - \mu_n$ (allowable if $\lambda_n + \mu_n \neq 1$).
- Equation: $x_n(t + \Delta t) = x_n(t) (1 - \lambda_n \Delta t - \mu_n \Delta t) + x_{n-1}(t) (\lambda_{n-1} \Delta t) + x_{n+1}(t) (\mu_{n+1} \Delta t)$ for $n > 1$.
- Differential Equation: $\frac{dx_n}{dt} = \lambda_{n-1} x_{n-1} + \mu_{n+1} x_{n+1} - (\lambda_n + \mu_n) x_n$ for $n > 1$, and $\frac{dx_0}{dt} = \mu_1 x_1 - \lambda_0 x_0$.



Pure birth process (Yule-Furry process): $\mu_n = 0$ (no deaths) and $\lambda_n = n\lambda$ (proportional growth rate).

Solution: $x_n(t) = {}^{n-1}C_{n-n_0} e^{-\lambda n_0 t} (1 - e^{-\lambda t})^{n-n_0}$ if $x_{n_0}(0) = 1$ i.e. given initial state is $X(0) = n_0$.

Poisson process: $\mu_n = 0$, $\lambda_n = \lambda$ (constant birth rate) and $X(0) = 0$. Solution: $x_n(t) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$.

For results of other CTMCs which can be interpreted as FIFO queues, see Section 5.4.8.

The Renewal Process: how many observations of an RV in series before a given time t



Let $X_t = \max \left\{ n : \left(\sum_{i=1}^n S_i \right) \leq t \right\}$, where S is a random variable with known distribution

(with support $S > 0$ representing interval times), and S_i is the i th i.i.d. observation of S .

- X_t is a renewal process, a type of generalised Poisson process, representing the number of observations of S that can be made by time t . The values of t where X_t changes by 1 are the 'jumping times', J_i .

- Renewal function equation: $E[X_t] = F_S(t) + \int_0^t E[X_{t-\tau}] f_S(\tau) d\tau$

where $F_S(t)$ is the CDF of S and $f_S(t)$ is the PDF of S .

- Strong law of large numbers for limiting observations per unit time: $\lim_{t \rightarrow \infty} \frac{X_t}{t} = \frac{1}{E[S]}$

- Central limit theorem: for large t , X_t is asymptotically Gaussian: $X_t \sim N\left(\frac{t}{E[S]}, \frac{t \times \text{Var}[S]}{E[S]^3}\right)$

- Renewal-reward process: at each observation time J_i , a 'reward variable' W_i is observed, and


the accumulated reward up to time t is $Y_t = \sum_{i=1}^{X_t} W_i$. The strong law of large numbers for the

limiting reward accumulation rate is $\lim_{t \rightarrow \infty} \frac{E[Y_t]}{t} = \frac{E[W]}{E[S]}$ (S and W need not be independent).

- Wald equation: $E[J_{X_t}] =$

5.4.8. Queueing Theory

Kendall notation: a queue model is named $A/B/c/N/K$ where:



A: inter-arrival time distribution (M: Markovian (exponential), D: deterministic (constant), G: general)

B: service time distribution (M: Markovian (exponential), D: deterministic (constant), G: general)

c : number of parallel servers (each serves 1 at a time)

N : system size (maximum queue length + c) (assumed ∞ if omitted)

K : population size (absolute maximum system size) (assumed ∞ if omitted)

(L : number of customers in the system, T : time of a customer in the system,

L_Q : number of customers in the queue, T_Q : time of a customer in the queue,

$\rho = \frac{\lambda_{eff}}{c\mu}$: server utilisation (load factor) - the queue is ergodic (a stationary state exists) if $\rho < 1$.)

Little's law: $mean\ throughput = \frac{mean\ number\ in\ system}{mean\ service\ time} \Leftrightarrow \lambda_{eff} = \frac{E[L]}{E[T]} = \frac{E[L_Q]}{E[T_Q]}$, always.

Coefficient of variation: $(cv)^2 = \frac{Var[X]}{(E[X])^2}$ for a specified random variable X . If $X \sim \text{Exp}$, $cv = 1$.

M/M/c/N/N Queue: those who leave the service node immediately rejoin the population (closed queue)

$$\pi_0 = \left[\sum_{n=0}^{c-1} {}^N C_n (c\rho)^n + \sum_{n=c}^N \frac{N! (c\rho)^n}{(N-n)! c! c^{n-c}} \right]^{-1}, \quad \pi_n = \begin{cases} \pi_0 {}^N C_n (c\rho)^n & \text{if } 0 \leq n < c, \\ \frac{N! (c\rho)^n}{(N-n)! c! c^{n-c}} & \text{else} \end{cases}$$

$$E[L] = \sum_{n=0}^N n \pi_n, \quad \lambda_{eff} = \sum_{n=0}^N (N-n)\lambda \pi_n.$$

M/M/c/N Queue: capped system size. Functionally identical to an M/M/c/ ∞ /N queue.

$$\pi_0 = \left[1 + \sum_{n=1}^c \frac{(c\rho)^n}{n!} + \frac{(c\rho)^c}{c!} \sum_{n=c+1}^N \rho^{n-c} \right]^{-1}, \quad \pi_n = \begin{cases} \pi_0 \frac{(c\rho)^n}{n!} & \text{if } 0 \leq n < c, \\ \pi_0 \frac{(c\rho)^n c^{c-n}}{c!} & \text{else} \end{cases}, \quad \pi_N = \frac{(c\rho)^N}{c! c^{N-c}} \pi_0,$$

$$E[L_Q] = \frac{\pi_0 (c\rho)^c \rho}{c! (1-\rho)^2} (1 - \rho^{N-c} - (N-c)(1-\rho)\rho^{N-c}), \quad \lambda_{eff} = \lambda(1 - \pi_N) = \frac{E[L_Q]}{E[T_Q]} = \frac{E[L]}{E[T]}, \quad E[T] = E[T_Q] + \mu^{-1}$$

M/M/c Queue (Erlang-C Model): birth-death process with $\lambda_n = \lambda$ and $\mu_n = \min\{c, n\} \mu$.

$$\pi_0 = \left[\sum_{n=0}^{c-1} \frac{(c\rho)^n}{n!} + \frac{(c\rho)^c}{c! (1-\rho)} \right]^{-1}, \quad P(L \geq c) = \sum_{n=c}^{\infty} \pi_n = \frac{(c\rho)^c}{c! (1-\rho)} \pi_0, \quad E[L] = c\rho + \frac{\rho}{1-\rho} P(L \geq c),$$

$$E[L_Q] = \frac{\rho}{1-\rho} P(L \geq c), \quad E[L - L_Q] = c\rho, \quad \pi_n = \begin{cases} \pi_0 \frac{(c\rho)^n}{n!} & \text{if } 0 \leq n < c, \\ \pi_0 \frac{(c\rho)^n c^{c-n}}{c!} & \text{else} \end{cases}, \quad \lambda_{eff} = \lambda.$$

M/M/1 Queue: birth-death process with $\lambda_n = \lambda$ and $\mu_n = \mu$.

$$\pi_0 = 1 - \rho, \quad \pi_n = \pi_0 \rho^n, \quad E[L] = \frac{\rho}{1-\rho}, \quad E[T] = \frac{1}{\mu - \lambda}, \quad E[L_Q] = \frac{\rho^2}{1-\rho}, \quad E[T_Q] = \frac{\rho}{\mu(1-\rho)}, \quad \lambda_{eff} = \lambda.$$

M/G/1 Queue: service times are randomly distributed with mean time μ^{-1} and variance σ^2 .

$$\rho = \frac{\lambda}{\mu}, \quad \pi_0 = 1 - \rho, \quad E[L_Q] = \frac{\rho^2(1 + \sigma^2 \mu^2)}{2(1-\rho)}, \quad E[T_Q] = \frac{\lambda(\mu^{-2} + \sigma^2)}{2(1-\rho)}, \quad E[L] = \rho + E[L_Q], \quad E[T] = \mu^{-1} + E[T_Q].$$

5.4.9. Priority Queues and Networks

Service Disciplines


- First-in, first-out (FIFO) / first-come, first-served (FCFS): longest waiting served first.
- Last-in, first out (LIFO) / stack: shortest waiting served first.
- Processor sharing: service capacity is shared equally among customers.
- Priority queue: customers are assigned a priority and can jump to the front of the queue on arrival or even displace a customer being served (preemptive).

Customer Waiting Behaviour

- Balking: customers decide not to join the queue if it is too long, as if an M/M/c/N queue.
- Jockeying: when there are multiple queues, customers move to the shortest at any time.
- Reneging: customers will leave the queue if they have waited too long.

Queuing Networks

- A network of queues is represented as m nodes (queue + service) each with population x_i and black-box parameters i.e. $M/M/c_i, \lambda_i^{(eff)}, \mu_i$.



- Transition probability for leaving node i to enter node j : p_{ij} (can also leave system with p_i)
- Overall arrival rate into queue j : $\lambda_j = a_j + \sum_i \lambda_i p_{ij}$ (a_j : arrival rate from outside the system)
- Matrix / vector form: $\lambda = (\mathbf{I} - \mathbf{P}^T)^{-1} \mathbf{a}$
- Utilisation of queue j : $\rho_j = \lambda_j / (c_j \mu_j)$
- The flow of customers through a network at steady-state can often (not always) be thought of as a fluid flowing through 'pipes' (the fluid limit).

- Jackson network: an open network of m M/M/ c_i queues.

Joint steady-state distribution product form: $\pi(\mathbf{x}) = \prod_{i=1}^m \pi(x_i)$ (independence)

Discrete event simulation (DES) can be used to investigate complex queuing systems, such as the simple module in Python, the SimEvents MATLAB add-on or enterprise software.

Simulations of queueing are subject to initialisation bias, in which the initial state can influence the averaged statistics even after a long time has elapsed. To mitigate this, data collection can be delayed until a given 'warm-up period' has passed. Determining the ideal length of this period may be challenging, with several methods proposed (e.g. Welch plot). Alternatively, a known/theoretical average steady state can be set at the start (although this may affect measures of variability).

5.4.9. Continuous-Time, Continuous State Space Markov Chains

The Wiener Process: continuous random walk (Brownian motion).

Discrete Brownian motion: Let $X_n = \sum_{k=1}^n \zeta_k$ where $\zeta_k \in \{-\delta, \delta\}$ with $P(\zeta_k = -\delta) = P(\zeta_k = \delta) = 1/2$.

The limiting distribution of X_n as $n \rightarrow \infty$ where time $t = n\delta$ as $\delta \rightarrow 0$ (so t is finite) is $p(x, t) \sim N_x(0, t)$.

In general, the Fokker-Planck equation is $\frac{\partial p}{\partial t} = \frac{\partial^2}{\partial x^2} (\alpha p)$ (diffusion equation).

If $p(x, 0) = \delta(x - a)$ then $p(x, t) \sim N_x(a, 2\alpha t)$. The standard deviation is unbounded: $\sigma = \sqrt{2\alpha t}$.

The Ornstein-Uhlenbeck Process: continuous analogue of the AR(1) process.

The stochastic differential equation is $dx = -\beta x dt + \sqrt{2\alpha} dW$ (W : Wiener process)

(α : diffusion term, β : drift term)

The Fokker-Planck equation is $\frac{\partial p}{\partial t} = \frac{\partial}{\partial x} (\beta x p) + \frac{\partial^2}{\partial x^2} (\alpha p)$.

If $p(x, 0) = \delta(x - a)$ then $p(x, t) \sim N_x(a e^{-\beta t}, \frac{\alpha}{\beta}(1 - e^{-2\beta t}))$. At steady state, $\lim_{t \rightarrow \infty} p(x, t) \sim N(0, \frac{\alpha}{\beta})$.

5.4.10. Sampling from Unnormalised Distributions

Consider a distribution $p(x) = \frac{f(x)}{C}$ where $f(x)$ is a known function but C is an unknown normalising constant. To approximate the drawing of samples from $p(x)$, we can use:

Rejection Sampling: Define a proposal distribution $q(x)$ and a constant k such that $f(x) \leq kq(x)$ for all x . Draw a candidate sample $Q \sim q(x)$, then draw another sample $U \sim \text{Uniform}(0, kq(Q))$. If $U \leq f(Q)$ then we accept the sample Q ; otherwise we reject. The accepted samples approach the distribution of $p(x)$.

Markov Chain Monte Carlo (MCMC): Construct a discrete-time Markov chain in which the states are the support (non-zero domain) of $f(x)$ and the 'time' is the sample number (index). The goal is to find the transition probabilities \mathbf{M} such that the stationary distribution is $p(x)$.

Detailed balance of equilibrium: $p(x)p(y|x) = p(y)p(x|y)$ equivalently $pM = p$.

Metropolis-Hastings Algorithm: Define a proposal distribution $q(x_{n+1} | x_n)$.

Generate samples of $x_{n+1} \sim q(x_{n+1} | x_n)$ and accept them with probability $A(x_{n+1} | x_n)$. (Note that $f = Aq$)

Detailed balance: $\frac{A(x_{n+1} | x_n)}{A(x_n | x_{n+1})} = \frac{f(x_{n+1})}{f(x_n)} \times \frac{q(x_n | x_{n+1})}{q(x_{n+1} | x_n)}$. Then $A(x_{n+1} | x_n) = \min \left\{ 1, \frac{f(x_{n+1})}{f(x_n)} \times \frac{q(x_n | x_{n+1})}{q(x_{n+1} | x_n)} \right\}$.

Gibbs Sampling: when $f = f(\mathbf{x})$ is a multivariable distribution and sampling from the joint pdf is difficult but from the conditional pdfs is easier. Choose an initial \mathbf{x}_0 . For each variable $x^{(i)}$, sample $x^{(i)}_{n+1} \sim p(x^{(i)}_{n+1} | \mathbf{x}^*_n)$ where \mathbf{x}^*_n is either \mathbf{x}_n or its partly/fully updated entries depending on implementation.

5.4.13. Wiener Deconvolution Filter

Given $y(t) = (h * x)(t) + n(t)$ (h : impulse response, n : noise), the goal is to find the MMSE estimate $\hat{x}(t) = (g * y)(t)$.

Signal-to-Noise ratio: $SNR(\omega) = \frac{S_{XX}(\omega)}{N(\omega)}$ ($S_{XX}(\omega)$: power spectrum of x)

Frequency spectrum of g : $G(\omega) = \frac{H^*(\omega) S_{XX}(\omega)}{|H(\omega)|^2 S_{XX}(\omega) + N(\omega)} = \frac{H^*(\omega) SNR(\omega)}{1 + |H(\omega)|^2 SNR(\omega)}$

Therefore $\hat{X}(\omega) = G(\omega) Y(\omega) \Rightarrow \hat{x}(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} G(\omega) Y(\omega) e^{i\omega t} d\omega$.

5.4.14. Stationary Time Series Analysis: AR, RA and ARMA Models

Gaussian noise: Let W_n be a sequence of random variables such that $E(W_n) = 0$ for all n and $E(W_i W_j) = \sigma^2$ if $i = j$ else 0.

Wide sense stationary (WSS): fixed mean, fixed correlation (independent of n).

The augmented Dickey-Fuller test (ADF test) can be used to test against the null hypothesis of a unit root (non-stationarity). Python: `statsmodels.tsa.stattools.adfuller(x)`

AR Model (Autoregressive Model)

AR(p) (order p) process: $X_n = \sum_{i=1}^p a_i X_{n-i} + W_n$

Correlation of AR(1): $R_{XX}(k) = E[X_n X_{n+k}] = a^k \sigma_X^2$

MA Model (Moving Average Model)

MA(p) (order p) process: $X_n = \mu + \sum_{i=1}^p a_i W_i$ where $\mu = E[X]$.

ARMA Model (Autoregressive Moving Average Model)

ARMA(p, q) process: $X_n = \mu + W_n + \sum_{i=1}^p a_i X_{n-i} + \sum_{j=1}^q b_j W_j$

Autocorrelation: $R_{XX}(k) = E[X_n X_{n+k}]$; Partial autocorrelation: $\phi(k) = E[(X_n - \overline{X_n})(X_{n+k} - \overline{X_{n+k}})]$

The ARMA process the white noise response of an IIR filter.

Model Fitting: estimating the hyperparameters of an ARMA process to fit observed data.

- The value of p is optimal at the elbow or peak of $|\phi(k)|$ (PACF).
- The value of q is optimal at the elbow or peak of $|R_{XX}(k)|$ (ACF, correlogram).

5.4.15. ARIMA Model (Autoregressive Integrated Moving Average Model)

ARIMA(p, d, q) process: $Y_n = (1 - L)^d X_n = \sum_{k=0}^d C_k (-1)^k X_{n-k}$ ($L^k X_n = X_{n-k}$: lag operator)

where X_n is an ARMA(p, q) process.

The ARIMA process Y_n is non-stationary, while the ARMA process X_n is stationary.

For time series analysis, the W_n (noise) terms are evaluated as the model error, which is assumed to be Gaussian. A time series can be made stationary by repeatedly differentiating (differencing) until sufficiently stationary. This is equivalent to removing polynomial terms from the Taylor series of a smoothed version of the series.

Akaike Information Criterion: $AIC = 2k - 2 \ln(\max\{\hat{L}\})$ (can be applied to any model)

(k : number of estimated parameters, L : likelihood function)

Further enhancements include the SARIMAX model (seasonal components and exogenous (externally supplied) variables). These are available in Python via the `statsmodels` module, with a similar interface to `tensorflow.keras` models.

5.4.16. Noise Response of an LTI System

Let $Hy = x$ where H : linear time invariant (LTI) system, y : system state, x : system input.

If H has impulse response h , the system response is given by $y = h * x$.

If $x = X(t)$ where X is a WSS stochastic process, then the response y is also WSS.

- Power spectral density (PSD) of y : $S_{YY}(\omega) = |H(\omega)|^2 S_{XX}(\omega)$ (H : FT of $h(t)$).
- Cross-spectral density (CSD): $S_{XY}(\omega) = S_{XX}(\omega) H(\omega)$

5.4.17. Multivariate Gaussian Distribution

A D -dimensional Gaussian random vector \mathbf{x} with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ has a joint pdf given by

$$N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^D |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right).$$

where $\boldsymbol{\mu}$ is a D -dimensional vector, $\boldsymbol{\Sigma}$ is a $D \times D$ positive definite symmetric matrix, and $|\boldsymbol{\Sigma}|$ its determinant.

- Distribution notation: writing $p(\mathbf{x}) = N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is equivalent to $\mathbf{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.
- If \mathbf{x} and \mathbf{y} are jointly Gaussian random vectors with marginal pdfs $p(\mathbf{x}) = N(\mathbf{x}; \mathbf{a}, \mathbf{A})$ and $p(\mathbf{y}) = N(\mathbf{y}; \mathbf{b}, \mathbf{B})$, with cross-covariance matrix $\mathbf{C} = \text{Cov}[\mathbf{x}, \mathbf{y}]$, then the joint pdf is

$$p\left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}\right) = N\left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}; \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}, \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ \mathbf{C}^\top & \mathbf{B} \end{bmatrix}\right),$$

and the conditional pdf is $p(\mathbf{x} | \mathbf{y}) = N(\mathbf{x}; \mathbf{a} + \mathbf{C}\mathbf{B}^{-1}(\mathbf{y} - \mathbf{b}), \mathbf{A} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^\top)$.

- Linear projection: if $p(\mathbf{x}) = N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{b}$ then $p(\mathbf{y}) = N(\mathbf{y}; \mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)$.
- The product of Gaussian densities is an unnormalised Gaussian:

$$N(\mathbf{x}; \mathbf{a}, \mathbf{A}) N(\mathbf{x}; \mathbf{b}, \mathbf{B}) = Z^{-1} N(\mathbf{x}; \mathbf{c}, \mathbf{C})$$

where $\mathbf{C} = (\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}$, $\mathbf{c} = \mathbf{C}(\mathbf{A}^{-1}\mathbf{a} + \mathbf{B}^{-1}\mathbf{b})$ and the normalising constant is Gaussian in both \mathbf{a} and \mathbf{b} : $Z^{-1} = (2\pi)^{-D/2} |\mathbf{A} + \mathbf{B}|^{-1/2} \exp(-(\mathbf{a} - \mathbf{b})^\top (\mathbf{A} + \mathbf{B})^{-1} (\mathbf{a} - \mathbf{b}) / 2)$.

- The (differential) entropy of a D -dimensional Gaussian random vector \mathbf{X} with pdf $p(\mathbf{x}) = N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is

$$h(\mathbf{X}) = \int p(\mathbf{x}) \log \frac{1}{p(\mathbf{x})} d\mathbf{x} = \frac{1}{2} \log((2\pi e)^D |\boldsymbol{\Sigma}|).$$

- KL (Kullback-Leibler) divergence between Gaussians:

If $p(\mathbf{x}) = N(\mathbf{x}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $q(\mathbf{x}) = N(\mathbf{x}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ then

$$KL(p, q) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} = \frac{1}{2} \left(\log \frac{|\boldsymbol{\Sigma}_2|}{|\boldsymbol{\Sigma}_1|} - D + \text{tr}(\boldsymbol{\Sigma}_2^{-1}\boldsymbol{\Sigma}_1) + (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^\top \boldsymbol{\Sigma}_2^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \right).$$

5.4.18. Gaussian Process

Formally, if $X: T \times \Omega \rightarrow \mathbb{R}$ is a Gaussian process with mean function m and covariance function K , then for every $t_1, \dots, t_k \in T$, the random vector $[X(t_1), \dots, X(t_k)]$ has a multivariate normal distribution with mean vector $\boldsymbol{\mu}_k = m(t_k)$ and covariance matrix $\boldsymbol{\Sigma}_{ij} = K(t_i, t_j)$.

(T : set of indices for Gaussian process, Ω : sample space of X)

Informally, X can be considered a ‘function’ that returns a distribution $N(x; m(t), K(t, t))$ for any given t , i.e. $X(t)$ has a univariate Normal distribution (with a specified covariance between any two different t). Function $X(t)$ has domain T and sample space Ω , returning a real-valued probability.

$X(t) \sim \text{GP}(m(t), K(t, t'))$ where $m(x)$ is the mean function, $K(t, t')$ is the covariance function with another given input x' .

It is useful to note that a Gaussian process can also be considered as an infinite-variable Gaussian distribution, with an infinite-length mean vector and infinity-by-infinity covariance matrix, since:

- An infinite-dimensional vector can be considered a scalar function of a single variable.
- An infinity-by-infinity matrix can be considered a scalar function of two variables:

Regression: let $y = X(t) + \varepsilon\sigma_y$ where $X(t)$ is a Gaussian process (the model \hat{y}), $\varepsilon \sim N(0, 1)$ and σ_y is the constant standard deviation of the model errors.

Assume a zero-mean Gaussian process prior distribution $f(x) | \theta \sim \text{GP}(0, K(x, x'))$ (θ : model hyperparameters used to define function K). Then $y | \theta \sim \text{GP}(0, K(x, x') + \mathbf{I}\sigma_y)$.

A common choice of K is $K(t, t') = \exp\left(-\frac{|t - t'|^2}{2\sigma}\right)$, where length-scale σ is a hyperparameter.

Prediction: let some given data be $\mathbf{y}_2 \sim N(\mathbf{b}, \mathbf{B})$ (finite dimensional) and data to predict $\mathbf{y}_1 \sim N(\mathbf{a}, \mathbf{A})$ (infinite dimensional). The joint distribution of \mathbf{y}_1 and \mathbf{y}_2 (Section 5.4.17) is a Gaussian process: $p(\mathbf{y}_1, \mathbf{y}_2) = N([\mathbf{y}_1; \mathbf{y}_2]; [\mathbf{a}; \mathbf{b}], [[\mathbf{A}, \mathbf{C}]; [\mathbf{C}^T, \mathbf{B}]])$.

Consider $p(\mathbf{y}_1 | \mathbf{y}_2) = \frac{p(\mathbf{y}_1, \mathbf{y}_2)}{p(\mathbf{y}_2)} = N(\mathbf{y}_1; \mathbf{a} + \mathbf{CB}^{-1}(\mathbf{y}_2 - \mathbf{b}), \mathbf{A} - \mathbf{CB}^{-1}\mathbf{C}^T)$.

The predictive mean is $E[y_1 | y_2] = \mathbf{a} + \mathbf{CB}^{-1}(y_2 - \mathbf{b}) = \mathbf{CB}^{-1}y_2$ ($\mathbf{a} = \mathbf{b} = \mathbf{0}$ for zero mean prior), which is linear in the given data y_2 . The predictive covariance is $\text{Cov}[y_1 | y_2] = \mathbf{A} - \mathbf{CB}^{-1}\mathbf{C}^T$ i.e. the uncertainty has been reduced from \mathbf{A} (prior uncertainty) by $\mathbf{CB}^{-1}\mathbf{C}^T$.

All Notes

5.4. Stochastic Processes, Signal Processing and Information Theory

5.4.19. Information Entropy

For a discrete random variable X with a probability mass function $P_X(x)$:

- **Shannon information content** ('surprise') of an outcome: $I_X(x) = -\log P_X(x)$
- The **entropy** of an r.v. X with pmf P is the expected information, $E[I_X(x)]$:

$$H(X) = \sum_x P(x) \log \frac{1}{P(x)} = -\mathbb{E}[\log(P(x))].$$

The entropy is measured in 'bits' if using log base 2, and 'nats' if using log base e.

- The **joint entropy** of random variables X_1, \dots, X_n with joint pmf $P_{X_1 \dots X_n}$ is

$$H(X_1, X_2, \dots, X_n) = \sum_{x_1, \dots, x_n} P_{X_1 \dots X_n}(x_1, \dots, x_n) \log \frac{1}{P_{X_1 \dots X_n}(x_1, \dots, x_n)}.$$

- The **conditional entropy** of Y given X is

$$H(Y|X) = \sum_x P_X(x) \underbrace{\sum_y P_{Y|X}(y|x) \log \frac{1}{P_{Y|X}(y|x)}}_{H(Y|X=x)} = \sum_x P_X(x) H(Y|X=x).$$

Note that a similar formula holds if we condition on a collection of random variables (X_1, \dots, X_n) instead of a single random variable X .

- **Chain rule for entropy:** The joint entropy of X_1, \dots, X_n can be written as

$$\begin{aligned} H(X_1, X_2, \dots, X_n) &= H(X_1) + H(X_2|X_1) + H(X_3|X_2, X_1) + \dots + H(X_n|X_{n-1}, \dots, X_1) \\ &= \sum_{i=1}^n H(X_i|X_{i-1}, \dots, X_1), \quad \text{where} \end{aligned}$$

$$H(X_i|X_{i-1}, \dots, X_1) = -\sum_{x_1, \dots, x_i} P_{X_1, \dots, X_i}(x_1, \dots, x_i) \log P_{X_i|X_1, \dots, X_{i-1}}(x_i|x_1, \dots, x_{i-1}).$$

- The **relative entropy or KL divergence** between two PMFs P and Q (defined on the same alphabet) is the information loss by using Q to represent P , given by

$$D(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)}.$$

- The **differential entropy** of a continuous random variable X with pdf p is

$$h(X) = \int_{-\infty}^{\infty} p(x) \log \frac{1}{p(x)} dx.$$

Joint differential entropy, conditional differential entropy, relative entropy/KL divergence, mutual information, chain rules for continuous random variables are all defined similarly to the discrete case with integrals replacing sums.

All Notes 5.4. Stochastic Processes, Signal Processing and Information Theory

5.4.20. Mutual Information

The mutual information between random variables X and Y represents the average reduction in uncertainty about X by knowing Y . For a joint pmf P_{XY} , mutual information is

$$\begin{aligned} I(X; Y) &= H(X) - H(X|Y) \\ &= H(Y) - H(Y|X) \\ &= H(X) + H(Y) - H(X, Y) \\ &= D(P_{XY} || P_X P_Y). \end{aligned}$$

Chain rule for mutual information:

$$\begin{aligned} I(X_1, X_2, \dots, X_n; Y) &= I(X_1; Y) + I(X_2; Y|X_1) + \dots + I(X_n; Y|X_{n-1}, \dots, X_1) \\ &= \sum_{i=1}^n I(X_i; Y|X_{i-1}, X_{i-2}, \dots, X_1). \end{aligned}$$

For maximum coding efficiency (optimal encoding), the mutual information / KL divergence $I(X; Y) = D_{\text{KL}}(P_{XY} || P_X P_Y)$ between an 'output' X and an 'input' Y must be maximised.

5.4.21. Decoding with Information Loss due to Noise

Data Processing Inequality: If X, Y, Z form a Markov chain, then

$$I(X; Y) \geq I(X; Z).$$

Discrete random variables X, Y, Z are said to form a Markov chain if their joint pmf can be written as $P_{XYZ} = P_X P_{Y|X} P_{Z|Y}$.

This is analogous to the second law of thermodynamics for statistical entropy (mutual information never increases with deterministic processing (conditioning)).

Fano's Inequality: Let X be a random variable taking values in a set χ with cardinality denoted by $|\chi|$. Let Y be a random variable jointly distributed with X , and $\hat{X} = f(Y)$ be any estimator of X from Y . Then the probability of error $P_e = P(\hat{X} \neq X)$ satisfies

$$1 + P_e \log |\mathcal{X}| \geq H(X | Y).$$

All Notes

5.4. Stochastic Processes, Signal Processing and Information Theory

5.4.22. Maximum Entropy Distributions

The entropy, representing the ‘surprise’ we get when making an observation. A random variable with a maximum entropy distribution (subject to given constraints) represents the most efficient coding of the information in the variable.

Constraint type	Constraint definition(s)	Maximum entropy distribution, $f_X(x)$	Distribution name
Limited range	$a \leq X \leq b$	$\frac{1}{b-a}$	Uniform
Nonnegative integers, limited range, fixed mean	$X \in \{0, 1, \dots, n\},$ $E[X] = \mu_x$	${}^n C_x \left(\frac{\mu_x}{n}\right)^x \left(1 - \frac{\mu_x}{n}\right)^{n-x}$	Binomial
Positive integers, fixed mean	$X \in \{1, 2, 3, \dots\},$ $E[X] = \mu_x$	$\frac{1}{\mu_x} \left(1 - \frac{1}{\mu_x}\right)^{x-1}$	Geometric
Nonnegative integers, fixed mean	$X \in \{0, 1, 2, \dots\},$ $E[X] = \mu_x$	$\frac{\mu_x^x e^{-\mu_x}}{x!}$	Poisson
Nonnegative, fixed mean	$X \geq 0, E[X] = \mu_x$	$\frac{1}{\mu_x} \exp \frac{-x}{\mu_x}$	Exponential
Fixed mean, fixed variance	$E[X] = \mu_x, Var[X] = \sigma_x^2$	$\frac{1}{\sqrt{2\pi\sigma_x^2}} \exp \frac{-(x-\mu_x)^2}{2\sigma_x^2}$	Normal

5.5. Machine Learning and Computational Statistics

5.5.1. Data Matrix and Notation for Datasets

The p ‘independent variables’ (features) are $\{x_1, x_2, \dots, x_p\}$. The single dependent variable is y . In a complete dataset, there are n recorded observations of each variable. A data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ and label column vector $\mathbf{y} \in \mathbb{R}^n$ is constructed:

Observation	Feature 1 (x_1)	Feature 2 (x_2)	...	Feature p (x_p)	Label ($y \in \mathbb{R}^n$)
1	X_{11}	X_{12}	...	X_{1p}	y_1
2	X_{21}	X_{22}	...	X_{2p}	y_2
...
n	X_{n1}	X_{n2}	...	X_{np}	y_n

X_{ij} is the i th observation of variable x_j .

The goal of an ML model is to find a mapping $f: \mathbf{X} \rightarrow \mathbf{y}$.

- Each row of \mathbf{X} can be considered a random vector.
- Each column of \mathbf{X} can be considered a set of observations from a random variable X_j .
- A ‘centred dataset’ contains features whose observations have zero mean: $\mathbf{x}_j' = \mathbf{x}_j - \overline{X_j}$.
- A ‘standardised dataset’ contains features whose observations have zero mean and unit standard deviation: $\mathbf{x}_j' = (\mathbf{x}_j - \overline{X_j}) / s_{X_j}$. The standardised X_j is comparable to (but not necessarily has) a standard normal (Gaussian) distribution.

5.5.2. Principal Component Analysis (PCA)


PCA is a method of constructing new features $\{x_1', x_2', \dots, x_p'\}$ that retain the information required to reconstruct y . Mathematically, it is a rotation of the coordinate axes used to specify the data into axes along which variance is maximised and covariance is minimised.

For a **standardised** ($\mu_i = 0, \sigma_i = 1$) dataset represented with data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ ($n \geq 2$):

- Singular value decomposition (SVD, Section 4.3.7): $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$
($\mathbf{U} \in \mathbb{R}^{n \times n}$, orthonormal; $\mathbf{\Sigma} \in \mathbb{R}^{n \times p}$, non-square diagonal; $\mathbf{V} \in \mathbb{R}^{p \times p}$, orthonormal, $\mathbf{V}^{-1} = \mathbf{V}^T$)
- Covariance matrix: $\mathbf{C} = \text{Cov}[\mathbf{X}, \mathbf{X}] = \frac{1}{n-1} \mathbf{X}^T \mathbf{X}$ such that $C_{ij} = \text{Cov}[X_i, X_j]$ ($\mathbf{C} \in \mathbb{R}^{p \times p}$, symmetric)
- Eigendecomposition (Section 4.3.6): $\mathbf{C} = \frac{1}{n-1} \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^T$ or $\mathbf{X}^T \mathbf{X} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^T$,
with eigenvalues in $\mathbf{\Sigma}^2 / (n - 1)$ or singular values in $\mathbf{\Sigma}$ are ordered descending.
- Data matrix in PC space: $\mathbf{X}' = \mathbf{X}\mathbf{V} = \mathbf{U}\mathbf{\Sigma}$. Each column of \mathbf{X}' is a principal component.
- The V_{ij} are the 'loadings' of X_i (coefficient in the expression for computing PC X'_j).
- The columns of \mathbf{V} are orthonormal vectors. This means that the PCs are independent i.e. uncorrelated: the new covariance matrix $\mathbf{C}' = \text{Cov}[\mathbf{X}', \mathbf{X}'] = \mathbf{V}^T \mathbf{C} \mathbf{V}$ is diagonal.
- Eigenvalues of \mathbf{C} are $\lambda_i = \sigma_i^2 / (n - 1)$, representing the variances of data in each PC.

PCA can be used for model dimensionality reduction by truncating $\mathbf{\Sigma}$ to retain only the $k \leq p$ largest singular values (projection: $\mathbb{R}^p \rightarrow \mathbb{R}^k$). The corresponding eigenvectors (rows of \mathbf{V}^T) are the dominant components. Then, $\mathbf{\Sigma}$ becomes a $n \times k$ matrix and \mathbf{V} becomes a $p \times k$ matrix. k can be chosen by plotting the eigenvalues in descending order and selecting the ones significantly larger than the rest (the 'elbow' of a 'scree plot'). For visualisation, $k = 2$ is often chosen (obtain PC1 and PC2).

For a multivariate distribution of zero-mean Normal variables, the principal components are along the axes of the p -dimensional hyper-ellipsoid formed by contours of the joint PDF. If the data has categorical labels, they can be colour-coded, around which 95% confidence ellipses can be drawn. This can sometimes help separate clusters before a clustering algorithm is applied (Section 5.5.6). A 'loading plot' shows the weights of the features (loadings) to a PC. A 'biplot' shows each loading as a vector to the point (PC1, PC2) in PC space, with a circle of radius 1 surrounding them, and the dataset in PC space optionally superimposed. Example (gene expression):



5.5.3. Scaling and Encoding (Pre-Processing)

Data can be normalised to remove potential variation due to physical units or magnitudes.

- Standard scaling: $x' = \frac{x - \mu}{\sigma}$ (sklearn.preprocessing.StandardScaler)
- Min-max scaling: $x' = \frac{x - \min(x)}{\max(x) - \min(x)}$ (sklearn.preprocessing.MinMaxScaler)

Non-numerical (categorical) data can be encoded into numerical data.

- Ordinal encoding: $x' \in \{0, 1, 2, \dots\}$ (sklearn.preprocessing.OrdinalEncoder)
- One-hot encoding: $x'_i \in \{0, 1\}$ (sklearn.preprocessing.OneHotEncoder)

where i is the number of distinct categories. Each one-hot feature is essentially a boolean for “is X equal to i ?”, where feature X takes values i .

5.5.4. Metrics for Evaluating Model Performance

Metrics for Evaluating Regression Models:

$$MSE = \frac{1}{n} \sum_i (y_i - \hat{y}_i)^2$$

Mean Squared Error (MSE)

$$MAE = \frac{1}{n} \sum_i |y_i - \hat{y}_i|$$

Mean Absolute Error (MAE)

$$MAPE = \frac{1}{n} \sum_i \left| \frac{y_i - \hat{y}_i}{y_i} \right|$$

Mean Absolute Percentage Error (MAPE)

$$R^2 = 1 - \frac{RSS}{TSS} = 1 - \frac{SSE}{SSE \text{ with } \hat{y}_i = \bar{y}} = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}$$

Coefficient of Determination (R^2)

(SSE (sum of squared errors) = RSS (residual sum of squares); SSE with $\hat{y} = \bar{y}$) = TSS (total sum of squares))

If the errors can be assumed to have zero mean (symmetric), then the MSE is equivalent to the variance of \hat{y} , and the RMSE = \sqrt{MSE} = std. dev.

Metrics for Evaluating Classification Models:

Confusion matrix: a contingency table showing frequency of classifications.

	Predicted Positive	Predicted Negative	Total
True Positive	TP True Positive	FN False Negative (Type II Error)	P
True Negative	FP False Positive (Type I Error)	TN True Negative	N
Total	PP	PN	$P + N = PP + PN$

$$\text{Accuracy} = \frac{TP + TN}{P + N}; \quad \text{Precision} = \frac{TP}{PP}; \quad \text{Sensitivity} = \frac{TP}{P}; \quad \text{Specificity} = \frac{TN}{N}; \quad F_1 \text{ Score} = \frac{2TP}{PP + P}$$

Precision is also known as Positive Predictive Value (PPV).

Sensitivity is also known as True Positive Rate (TPR), Recall, Hit Rate, or Power.

Specificity is also known as True Negative Rate (TNR) or Selectivity.

Metrics for Evaluating Clustering Models:

- Silhouette Score: $s_i = \frac{b_i - a_i}{\max\{a_i, b_i\}}$, compute for each point; measures separation.
(a_i : average distance between i and other points in same cluster; b_i : distance between i and centroid of nearest other cluster)
- Adjusted Rand Index: fraction of pairs of points in the correct same cluster, adjusted for randomness.
- Adjusted Mutual Information: measured between two given clusters; adjusted for the entropies of each cluster, see Section 5.5.3.
- Calinski-Harabasz Index (variance criterion): ratio of within-cluster dispersion to inter-cluster dispersion.

Many of these metrics are available in `scikit-learn`.

5.5.5. Traditional (Non Neural Network Based) Supervised Classification Algorithms

Naive Bayes Classifier

A Bayes classifier assumes conditional independence between each x given y . Classification labels \hat{y}_j are chosen to maximise the MAP, given by $p(y) \prod_i p(x_i | y)$ where $p(y) = 1 / (\# \text{ labels})$ is the prior relative frequency of y and $p(y | x_i)$ is calculated from a given distribution e.g. Gaussian.

Python (scikit-learn): `from sklearn.naive_bayes import GaussianNB`

Random Forest Classifier

A supervised classification method in which binary decision trees are created and optimised on their decision rule to minimise a loss function (often Gini impurity).

Python (scikit-learn): `from sklearn.ensemble import RandomForestClassifier`

XGBoost (extreme gradient boosting) is a powerful tree-based model implemented in C++.

Support Vector Classifier

For maximum margin classification, use a hinge loss:
$$L(\hat{y}, y) = \begin{cases} \max\{0, \varepsilon - \hat{y}\}, & \text{if } y > 0 \\ \max\{0, \varepsilon + \hat{y}\}, & \text{if } y < 0 \end{cases}$$

Python (scikit-learn): `from sklearn.svm import SVC`

K Nearest Neighbours

A simple classification algorithm in which data points are classified according to the most common label among that of the K nearest training data points (majority voting).

KD tree algorithm: constructs a binary tree in which each node represents a decision on the point coordinates (typically “above/below the median feature value?”), and the leaves of the tree are all the points within a class under these decision rules.

Python (scikit-learn): `from sklearn.neighbors import KNeighborsClassifier`

5.5.6. Clustering Algorithms (Unsupervised)

***K* Means Clustering**

For a dataset of n observations $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$, centroidal points \hat{y}_j ($0 \leq j < K$) are chosen such that $\sum_i \text{dist}(x_i, \hat{y}_{\text{closest}(i)})^2$ is minimised, where $\text{closest}(i)$ returns the centroidal point index which is closest to x_i . This is a nonlinear and non-smooth optimisation problem. The clusters are then given by $S_j = \{x_i: \text{closest}(i) = \hat{y}_j\}$.

Main limitation: cannot produce cluster boundaries whose centroids are necessarily far from any of their points (e.g. concentric rings).

Python (scikit-learn): `from sklearn.cluster import KMeans`

DBSCAN

DBSCAN (Density-Based Spatial Clustering of Applications with Noise) groups points of similar densities, without reference to any centroids. Can produce more complex cluster boundary topologies, including shapes with holes (annuli).

A modification is HDBSCAN (hierarchical DBSCAN) which has easier hyperparameter tuning.

Python (scikit-learn): `from sklearn.cluster import DBSCAN`

5.5.7. Generalised Linear Regression (Regression, Supervised)

Linear regression model: $\hat{y}(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0 = w_1 x_1 + w_2 x_2 + \dots + w_p x_p + w_0$. (linear combination of features)

The values of \mathbf{w} and b are chosen to minimise a particular cost function: $\mathbf{w}^*, w_0^* = \operatorname{argmin} C(\mathbf{w}, w_0)$.

- Ordinary least squares (OLS): $C(\mathbf{w}) = \sum_{i=1}^n (y_i - \hat{y}_i)^2$ (SSE: sum of squared errors)
- LASSO regression: $C(\mathbf{w}) = \sum_{i=1}^n (y_i - \hat{y}_i)^2 + \lambda \sum_{j=1}^p |w_j|$ (l^1 norm regularisation)
- Ridge regression: $C(\mathbf{w}) = \sum_{i=1}^n (y_i - \hat{y}_i)^2 + \lambda \sum_{j=1}^p w_j^2$ (l^2 norm regularisation)
- Support vector regression (SVR): $C(\mathbf{w}, b) = \lambda \sum_{i=1}^n \max\{0, |y_i - \hat{y}_i| - \varepsilon\} + \frac{1}{2} \sum_{j=1}^p w_j^2$ (hinge loss)

(n : number of observations, p : number of features; $\mathbf{X} \in \mathbb{R}^{n \times p}$; and λ, ε are hyperparameters.)

Regularisation terms penalise large components in \mathbf{w} , ensuring weight decay to prevent overfitting.

For polynomial regression up to order k , extend \mathbf{X} with new features given by

$x_1^{\alpha_1} x_2^{\alpha_2} x_3^{\alpha_3} \dots x_n^{\alpha_n}$ for all $0 \leq \alpha_i \leq k$ such that $2 \leq \sum_{i=1}^n \alpha_i \leq k$, then apply linear regression as

usual. E.g. for $k = 2$ and $n = 3$ (original features: $\mathbf{X} = \{x, y, z\}$), extend with $\{xy, yz, xz, x^2, y^2, z^2\}$.

Python (scikit-learn):

```
from sklearn.linear_model import LinearRegression # also: Ridge / Lasso / sklearn.svm.SVM
from sklearn.preprocessing import PolynomialFeatures
from sklearn.model_selection import train_test_split
from sklearn.metrics import mean_absolute_percentage_error

poly_features = PolynomialFeatures(degree=3, include_bias=False).fit_transform(X_data)
X_train, X_test, y_train, y_test = train_test_split(poly_features, y_data,
                                                    test_size=0.2)

poly_model = LinearRegression(fit_intercept=True).fit(X_train, y_train)
y_pred = poly_model.predict(X_test) # use the model to estimate y(X_test)
mpe = mean_absolute_percentage_error(y_test, y_pred) # quantify accuracy
```


5.5.8. Gaussian Process Regression (Uncertainty Regression, Supervised)

Gaussian process regression allows quantification of the uncertainty associated with a regression model about a mean predicted value. GPR is a useful form of nonlinear regression in which the computational complexity is independent of the form of the data: input x may be a scalar, a vector, a string, a graph, etc. It is suited when only a small training dataset is available.

- For a given training dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ such that all y_i are centralised (zero mean), the aim of the model is to produce predictive distributions on test points $\{(\mathbf{x}_i^*, y_i^*)\}_{i=1}^m$.
- Assume noisy observations of a Gaussian process (Section 5.4.18) f , i.e. $y_i = f(\mathbf{x}_i) + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$.
- Notation: \mathbf{X} = training data matrix, \mathbf{X}^* = testing data matrix, \mathbf{y} = training labels, $\mathbf{f} = \mathbf{f}(\mathbf{X})$, $\mathbf{f}^* = \mathbf{f}(\mathbf{X}^*)$.
- Kernel function (autocovariance): $\mathbf{K}_{\mathbf{X}\mathbf{X}}$ where element (i, j) is a function $K: \{\mathbb{R}^p \times \mathbb{R}^p\} \rightarrow \mathbb{R}$ evaluated as $K(\mathbf{x}_i, \mathbf{x}_j)$. A common choice of kernel for numeric \mathbf{x} is the radial basis function (RBF), $K(\mathbf{x}, \mathbf{x}' | \boldsymbol{\tau}) = \lambda \exp[-|\mathbf{x} - \mathbf{x}'| / 2\sigma]$ where $\boldsymbol{\tau} = \{\lambda: \text{output correlation scale}, \sigma: \text{input correlation scale}\}$ are hyperparameters. Kernel functions can be combined additively (and **approximately**, not formally, multiplicatively) to produce better fits to observed data. $\mathbf{K}_{\mathbf{X}\mathbf{X}^*}$, $\mathbf{K}_{\mathbf{X}^*\mathbf{X}}$ and $\mathbf{K}_{\mathbf{X}^*\mathbf{X}^*}$ are defined similarly. Note that $\mathbf{K}_{\mathbf{X}^*\mathbf{X}} = \mathbf{K}_{\mathbf{X}\mathbf{X}^*}^\top$ by the property in Section 5.4.17.
- Assume that \mathbf{y} and \mathbf{f}^* together form a joint $(n + m)$ -dimensional normal distribution:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}^* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \hat{\mathbf{K}}_{\mathbf{X}\mathbf{X}} & \mathbf{K}_{\mathbf{X}\mathbf{X}^*} \\ \mathbf{K}_{\mathbf{X}^*\mathbf{X}} & \mathbf{K}_{\mathbf{X}^*\mathbf{X}^*} \end{bmatrix} \right) \quad \text{where} \quad \hat{\mathbf{K}}_{\mathbf{X}\mathbf{X}} = \mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma_\varepsilon^2 \mathbf{I} \quad (\text{block matrices}).$$

The training dataset provides a marginalisation (conditioning) of one observation from this distribution, from which the distribution of the testing set can be inferred.

- The posterior distribution is then $\mathbf{f}^* | \mathbf{X}^*, \mathbf{X}, \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ where

$$\boldsymbol{\mu}_{\mathbf{f}^*} = \mathbf{K}_{\mathbf{X}^*\mathbf{X}} \hat{\mathbf{K}}_{\mathbf{X}\mathbf{X}}^{-1} \mathbf{y}, \quad \boldsymbol{\Sigma}_{\mathbf{f}^*} = \mathbf{K}_{\mathbf{X}^*\mathbf{X}^*} - \mathbf{K}_{\mathbf{X}^*\mathbf{X}} \hat{\mathbf{K}}_{\mathbf{X}\mathbf{X}}^{-1} \mathbf{K}_{\mathbf{X}\mathbf{X}^*}$$

- Individual posterior output distributions for y have 1D normal distributions with mean given by an entry in $\boldsymbol{\mu}_{\mathbf{f}^*}$ and variance given by the corresponding diagonal entry in $\boldsymbol{\Sigma}_{\mathbf{f}^*}$.
- Hyperparameter selection: choose $\boldsymbol{\theta} = \{\sigma_\varepsilon^2, \boldsymbol{\tau}\}$ such that $\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} \log p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta})$ (maximum log-likelihood prior). By algebra, $\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} \log p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta}) = \arg \max_{\boldsymbol{\theta}} \log \mathcal{N}(\mathbf{y} | \mathbf{0}, \hat{\mathbf{K}}_{\mathbf{X}\mathbf{X}})$

This is a differentiable function, allowing gradient-based optimisation techniques.

Python: (scikit-learn)

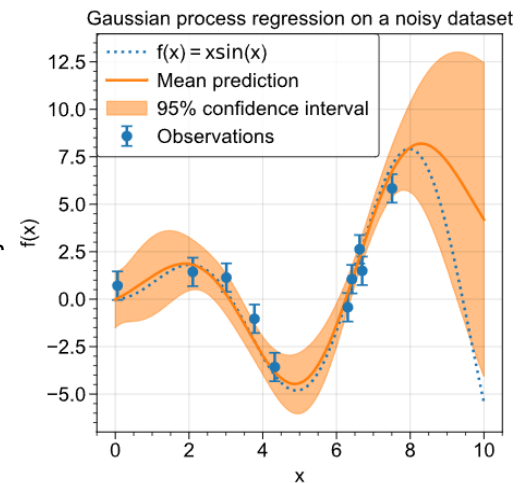
Gaussian process regression with one feature using the RBF kernel, showing the function $f(x)$ (mean) and $\pm 2\sigma$ range (95% confidence CI).

```
import numpy as np
import matplotlib.pyplot as plt
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF


X = np.linspace(start=0, stop=10, num=1_000).reshape(-1, 1)
y = np.squeeze(X * np.sin(X))

training_indices = np.random.choice(np.arange(y.size),
    size=10, replace=False)
X_train, y_train = X[training_indices], y[training_indices]
NOISE_STD = 0.75 # measurement noise  $\epsilon$ 
y_train_noisy = y_train + np.random.normal(loc=0.0,
    scale=noise_std, size=y_train.shape)
kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2)) # kernel function K
gaussian_process = GaussianProcessRegressor(kernel=kernel, alpha=NOISE_STD**2,
    n_restarts_optimizer=9)
gaussian_process.fit(X_train, y_train) # fit to observed data
print(gaussian_process.kernel_) # K multiplied by optimal  $\lambda$ 
mean_prediction, std_prediction = gaussian_process.predict(X, return_std=True)


plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.errorbar(X_train, y_train_noisy, NOISE_STD, linestyle="None", color="tab:blue",
    marker=".", markersize=10, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
plt.fill_between(X.ravel(), mean_prediction - 1.96 * std_prediction,
    mean_prediction + 1.96 * std_prediction, color="tab:orange",
    alpha=0.5, label=r"95% confidence interval")
plt.legend(); plt.xlabel("$x$"); plt.ylabel("$f(x)$")
plt.title("Gaussian process regression on a noisy dataset")
```




5.5.9. Logistic Regression (Classification, Supervised)



Schematic of a single neuron



Sigmoid: nonlinear activation function

Output for a binary classification task ($D=2$)

Activity (weighted sum): $u = w_1 x_1 + w_2 x_2 + \dots + w_D x_D + b$

Output: $z = \sigma(u) = \frac{1}{1 + \exp(-u)}$ (sigmoid: logistic curve)

The output of the network can be written as $y = f(\mathbf{x}; \mathbf{w})$. This can be interpreted as a (posterior) probability between 0 and 1 i.e. $y = p(z = 1 \mid \mathbf{x}, \mathbf{w})$ for a known label z .

The goodness-of-fit is measured by the log-likelihood objective function evaluated over N data points in a batch:

$$G(\mathbf{w}) = - \sum_{i=1}^N z^{(i)} \log y^{(i)} + (1 - z^{(i)}) \log (1 - y^{(i)}) \quad (\text{binary cross-entropy})$$

For multi-class classification i.e. $y = p(z = \mathbf{1} \mid \mathbf{x}, \mathbf{w})$, the softmax activation is applied before computing the loss (the categorical cross-entropy). (Note that if errors are Normally distributed i.e. in a regression / estimation problem, the ideal objective function is the mean-squared error (MSE):

$$G(\mathbf{w}) = \frac{1}{N} \sum (z^{(i)} - y^{(i)})^2$$

The aim of training is to find \mathbf{w} such that $G(\mathbf{w})$ is minimised i.e. $\mathbf{w}^* = \text{argmin} G(\mathbf{w})$. This is done by gradient descent i.e. $\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla G$ i.e. $w_i \leftarrow w_i - \eta \frac{\partial G}{\partial w_i}$ (η : learning rate)

Common improvements to the gradient descent method are:

- Stochastic gradient descent (SGD): compute ∇G over a smaller sample (mini-batch).
- Adaptive step: allow η to vary e.g. AdaGrad, RMSProp, Adam
- Regularisation: add $\alpha E(\mathbf{w}) = \frac{1}{2} \alpha \mathbf{w}^T \mathbf{w}$ to objective function, penalises extreme weights, preventing overfitting. α controls trade-off.


Single layer perceptrons generate linear decision regions, and can only be generalised by 1) manual 'feature engineering' or 2) adding more layers (MLPs, Section 5.5.10) to learn the feature transforms. The hidden layer activations represent the new features in latent space.

Python (scikit-learn):

```
from sklearn.linear_model import LogisticRegression
ans = LogisticRegression(penalty='l2', C=1.0).fit(X_train) # C = 1/alpha
y_pred = ans.predict(X_test)
```

5.5.10. Multilayer Perceptron (Feedforward / Fully Connected Neural Networks)

An MLP is a neural network which uses multiple neurons to allow multiple output dimensions. There are often 'hidden layers' of neurons which help in hierarchical pattern matching in the data.



- $a_j^{(L)}$: activation of neuron j in layer L of the network.
- $w_{j,i}^{(L)}$ ($i_{L-1} \rightarrow j_L$): weight from neuron i in layer $L - 1$ to neuron j in layer L .
- $b_j^{(L)}$ indicates the bias of neuron j in layer L .

These are typically compressed into vectors $\mathbf{a}^{(L)}$, $\mathbf{b}^{(L)}$ and matrices $\mathbf{W}^{(L)}$ for each layer.


- Weighted sum: $\mathbf{z}^{(L)} = \mathbf{W}^{(L)} \mathbf{a}^{(L-1)} + \mathbf{b}^{(L)}$ (in batch normalisation, \mathbf{a} is standard scaled)
- Feedforward activation: $\mathbf{a}^{(L)} = \mathbf{f}(\mathbf{z}^{(L)})$ (\mathbf{f} : nonlinear function e.g. sigmoid, ReLU, tanh)
- Output cost: $C(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{2} \sum_i (\hat{y}_i - y_i)^2$

Nonlinear Functions: $f(x)$ $f'(x)$ or Jacobian element

- Sigmoid: $\sigma(x) = (1 + e^{-x})^{-1}$, $\sigma'(x) = e^{-x} (1 + e^{-x})^{-2}$
- ReLU (rectified linear unit): $\text{ReLU}(x) = \max\{0, x\}$, $\text{ReLU}'(x) = \mathbf{I}_{\{x>0\}}$ ($x \neq 0$)
- Softmax: $(\mathbf{p}(\mathbf{x}))_i = e^{x_i} / \sum_j e^{x_j}$, $\left(\frac{\partial p}{\partial x_j}\right)_i = p_i(\mathbf{I}_{\{i=j\}} - p_j)$

Backpropagation: derivatives on computational graphs use the chain rule ($\dots \rightarrow h_{L-2} \rightarrow i_{L-1} \rightarrow j_L$)

To compute the gradient of C with respect to a weight in the **final** layer $w_{j,i}^{(L)}$:



$$\underbrace{\frac{\partial C}{\partial w_{j,i}^{(L)}}}_{\text{component of } \nabla C} = \underbrace{\frac{\partial C}{\partial a_j^{(L)}}}_{(a_j^{(L)} - y_j)} \times \underbrace{\frac{\partial a_j^{(L)}}{\partial z_j^{(L)}}}_{f'(z_j^{(L)})} \times \underbrace{\frac{\partial z_j^{(L)}}{\partial w_{j,i}^{(L)}}}_{a_i^{(L-1)}}$$

To compute the gradient of C with respect to a weight in a **further** layer $w_{i,h}^{(L-1)}$, sum over all final layer contributions when computing $\partial C / \partial a_i^{(L-1)}$:

$$\underbrace{\frac{\partial C}{\partial a_i^{(L-1)}}}_{\text{deeper layer}} = \sum_{j=1}^{n_L} \underbrace{\frac{\partial C}{\partial a_j^{(L)}}}_{(a_j^{(L)} - y_j)} \times \underbrace{\frac{\partial a_j^{(L)}}{\partial z_j^{(L)}}}_{f'(z_j^{(L)})} \times \underbrace{\frac{\partial z_j^{(L)}}{\partial a_i^{(L-1)}}}_{w_{j,i}^{(L)}} \quad \text{and} \quad \underbrace{\frac{\partial C}{\partial w_{i,h}^{(L-1)}}}_{\text{component of } \nabla C} = \underbrace{\frac{\partial C}{\partial a_i^{(L-1)}}}_{\text{deeper layer}} \times \underbrace{\frac{\partial a_i^{(L-1)}}{\partial z_i^{(L-1)}}}_{f'(z_i^{(L-1)})} \times \underbrace{\frac{\partial z_i^{(L-1)}}{\partial w_{i,h}^{(L-1)}}}_{a_h^{(L-2)}}$$

Gradient Descent: update weights by $\Delta \mathbf{W}_{n+1}^{(L)} = -\eta \nabla_{\mathbf{W}_n^{(L)}} C$. (C averaged over whole training set.)

Python Typical Implementation (TensorFlow with the Keras API):

```
from tensorflow.keras import Input, layers, models
model = models.Sequential(); model.add(Input(num_inputs, ...))
model.add(layers.Dense(num_neurons, ...)) # add as many hidden layers as desired
model.compile(optimizer=..., loss=..., metrics=...)
model.fit(X, y, epochs=..., batch_size=...) # training
model.evaluate(...); # check loss; model.predict(...) # testing
```

5.5.11. Techniques for Training Deep Neural Networks

Optimisations of the gradient descent algorithm:

- **Stochastic gradient descent (SGD):** compute ∇C per training example, rather than averaging over the whole training set (i.e. use batch size = 1). Decrease η over time to allow convergence.
- **Adam (adaptive moment estimation):** combines AdaGrad and RMSProp by keeping track of an exponentially decaying average (EMA) of past gradients and their squares.

Problems faced by deep neural networks and their solutions:

- **Unstable gradient problem** (vanishing gradient / exploding gradient): layers further from the output are harder to train in backpropagation. Resolved by using non-saturating activation functions (gradients do not fall to zero as $x \rightarrow \pm\infty$, e.g. leaky ReLU), using batch normalisation ($\mathbf{a}^{(L-1)}$ is standard scaled over a batch of values before evaluating the sum), and layer normalisation (or batch) can also be used in which $\mathbf{a}^{(L-1)}$ is standard scaled over the neurons in the layer.
- **Overfitting:** a model may perform well on training data but bad in unseen test data. Resolved by using a regularisation term in the loss function to prevent large weights from forming; using a recurrent dilution / dropout to randomly temporarily shrink / exclude neurons in training to force the network to generalise to wider patterns; and using pruning to drop unimportant weights, all of which reduce model complexity.
- **Internal covariate shift:** the erratic change in the distribution of neuron activations due to fluctuating data inputs and hence sharp changes to the weights. Mitigated by batch normalisation ($\mathbf{a}^{(L-1)}$ is standard scaled over a mini-batch of values before evaluating the sum).

Metric evaluation:

- **K-fold cross-validation:** Shuffle the data. Allocate a fixed proportion as the test data. Split the remaining data into K equal groups (folds). For each fold i in the folds, allocate fold i as the validation set and the remaining data as the training set, and train the data, and find performance on the validation set. Python: `from sklearn.model_selection import KFold`
- **Leave-one-out cross-validation (LOOCV):** K -fold CV but with $K =$ number of data points in training set i.e. use every data point once as the validation 'set'.

Hyperparameter optimisation: in Python, can use the keras-tuner library for TensorFlow models.

- **Grid search:** train models with all combinations of hyperparameters within a search space.
- **Random search:** train models with randomly chosen hyperparameters and locate good clusters.
- **Gradient-based optimisation** (hypernetworks): use gradient descent to find the optimal hyperparameters in the same way as a regular neural network finds its optimal weights.
- **Bayesian optimisation:** a probabilistic approach to estimate the optimal hyperparameters.
- **Evolutionary algorithm:** use a fitness function (e.g. CV) to rank performance, then select the best for crossover and mutation of (encoded) hyperparameters, run until desired performance is observed.


5.5.12. Convolutional Neural Networks (CNNs, ConvNets)

A CNN uses hidden layers consisting of convolution cells. CNNs are often used for computer vision, with each weight kernel representing a filter to identify a particular feature e.g. edges / blobs in early layers, more complex features in later layers. Each kernel is convolved against the input (Section 5.4.3) then adds an array-wide bias to produce a set of N activation maps, which are stacked into a 3 dimensional array and then subject to a pointwise nonlinear function.


The hyperparameters for a convolutional layer are the number of filters N (number of neurons), the dimension of the filters $K \times K$ (kernel size), the stride S (step size) and the padding P (extend with zeros).

A convolutional layer takes an input array of dimension $W_1 \times H_1 \times D_1$ and produces an output of $W_2 \times H_2 \times D_2$, where $W_2 = \frac{W_1 - K + 2P}{S} + 1$, $H_2 = \frac{H_1 - K + 2P}{S} + 1$ and $D_2 = N$. An optional pooling layer (using an operation such as max pooling or average pooling) reduces the output size by subdividing the output into square arrays (per layer) and choosing the maximum or mean value. The stride for a pooling layer is usually equal to the pool size, so that there is no overlap between subarrays. With parameter sharing, it introduces a total of $(K^2 D_1 + 1)N$ parameters (weights + biases) per layer.

CNNs can be represented visually as transforming an array into different dimensions as blocks:



Another common notation is to label each layer symbolically:



($C_n[a \times b @ s]$: convolutional layer # n using $a \times b$ kernel and stride s , $P_n[a \times b @ s]$: pooling layer # n (max / average) using $a \times b$ subsampling and stride s , $F_n[N]$: fully connected layer # n with N neurons, $_/$: ReLU activation, dashed border: applies dropout. The number of neurons is inferred from the last dimension $D = K$. Zero padding is assumed unless otherwise specified.)

The last layer of a convolutional neural network for classification is typically a fully connected (FC, dense) layer, with softmax activation, which takes in the flattened output array and produces a vector representing the classification probabilities.

1D convolutions can be used for short-term time series forecasting, for which they are faster to train than LSTMs, but are less capable of detecting longer term patterns. Hybrid CNN-LSTMs can be used to combine the strengths of both, and are useful for e.g. anomaly detection.

Data augmentation (image manipulation e.g. cropping, reflecting) can be used to artificially enlarge an image dataset to generate more images. This can also be used in the testing stage, where the testing set is augmented and predictions are based on the modification with the maximum output.


Python (Keras, TensorFlow): convolutional layers are added separately to pooling layers, e.g.

```
# for 32x32 pixel input RGB (3 dimensions) images
model.add(layers.Conv2D(32, (3, 3), activation='relu', input_shape=(32, 32, 3)))
model.add(layers.MaxPooling2D((2, 2)))
```

5.5.13. Recurrent Neural Networks

RNNs make predictions based on a sequence of inputs (typically in time) rather than a single input. Each input sequence can be represented by a data matrix \mathbf{X} (d variables \times n observations; transpose of convention used in Section 5.5.6). A batch of inputs can be represented as a rank 3 tensor (data matrices stacked along a third dimension: d variables \times n observations \times b sequences):

$$\mathbf{X} = \begin{bmatrix} x_0^{(0)} & x_0^{(1)} & \cdots & x_0^{(n)} \\ x_1^{(0)} & x_1^{(1)} & \cdots & x_1^{(n)} \\ \vdots & \vdots & \ddots & \vdots \\ x_d^{(0)} & x_d^{(1)} & & x_d^{(n)} \end{bmatrix} \quad \underline{\mathbf{B}} = [^{(0)}\mathbf{X} \quad ^{(1)}\mathbf{X} \quad \cdots \quad ^{(b)}\mathbf{X}]$$




Each recurrent cell takes both an input observation $\mathbf{x}^{(j)}$ as well as the output of the previous cell (as its 'hidden state') $\mathbf{h}^{(j-1)} = \mathbf{y}^{(j-1)}$ to produce an output $\mathbf{y}^{(j)}$, giving the network an ability to recall previous values: $\mathbf{y}^{(j)} = \sigma(\mathbf{W}_x \mathbf{x}^{(j)} + \mathbf{W}_h \mathbf{y}^{(j-1)} + \mathbf{b}) = \sigma(\underline{\mathbf{W}} \mathbf{X}^{(j)})$ (\mathbf{W}_x : weight matrix for input, \mathbf{W}_h : weight matrix for hidden input, \mathbf{b} : bias vector, all identical over j)

Backpropagation Through Time (BPTT): unroll the network through time (as shown above), then use regular NN backpropagation. The cost function is based on the T most recent outputs (where T is a hyperparameter in truncated BPTT).

Layer Normalisation: standardise inputs to (or outputs from) $\sigma(\cdot)$ by learning an offset (mean) and scale (std.dev) for each observation.

Long Short Term Memory (LSTM): memory cells can recall further back.



Layer outputs are: $\mathbf{f}, \mathbf{g}, \mathbf{i}, \mathbf{o} = \sigma(\mathbf{W}_x \mathbf{x} + \mathbf{W}_h \mathbf{h} + \mathbf{b})$ for different weights in each FC.

Then $\mathbf{c}^{(n)} = \mathbf{c}^{(n-1)} \otimes \mathbf{f}^{(n)} + \mathbf{g}^{(n)} \otimes \mathbf{i}^{(n)}$ and $\mathbf{y}^{(n)} = \mathbf{o}^{(n)} \otimes \tanh \mathbf{c}^{(n)}$.

(FC: fully-connected (dense) layer, \mathbf{f} : forget gate controller, \mathbf{g}, \mathbf{i} : input gate controllers, \mathbf{o} : output gate controllers, \mathbf{c} : cell state, $\mathbf{h} = \mathbf{y}$: hidden state (output), \otimes : Hadamard (simple elementwise) matrix product)

Gated Recurrent Unit (GRU): a simplified LSTM which maintains similar performance.

5.5.14. Transformer Networks

Embeddings

Each token i is assigned an embedding vector e_i , based on a lookup table. The token position is encoded by adding an orthogonal set of sinusoids.

Single Head Self Attention Mechanism: updates embeddings

- Compute the queries $\mathbf{Q} = \mathbf{W}_q \mathbf{E}$, keys $\mathbf{K} = \mathbf{W}_k \mathbf{E}$ and values $\mathbf{V} = \mathbf{W}_v \mathbf{E}$.
- Compute the attention pattern, $\mathbf{A} = \text{softmax}(\mathbf{Q}\mathbf{K}^T / d_k)$.
(input may be 'masked' setting all entries below the leading diagonal to $-\infty$, \mathbf{Q} : matrix of \mathbf{q}_i , \mathbf{K} : matrix of \mathbf{k}_i , d_k : dimension of query/key space, softmax is computed columnwise.)
- Compute the attention output $\Delta \mathbf{E} = \mathbf{A}\mathbf{V}$. The new embedding is then $\mathbf{E}' = \mathbf{E} + \Delta \mathbf{E}$.

The value weights \mathbf{W}_v is a low rank transformation matrix, implemented in non-square factorised form as $\mathbf{W}_v = \mathbf{W}_v^{(1)} \mathbf{W}_v^{(1)}$.

The attention pattern matrix size is $O(N^2)$ where N is the context size, making this a bottleneck in transformer architectures. Recent modifications have allowed for more scalable models (e.g. sparse attention mechanisms, blockwise attention, ring attention...).

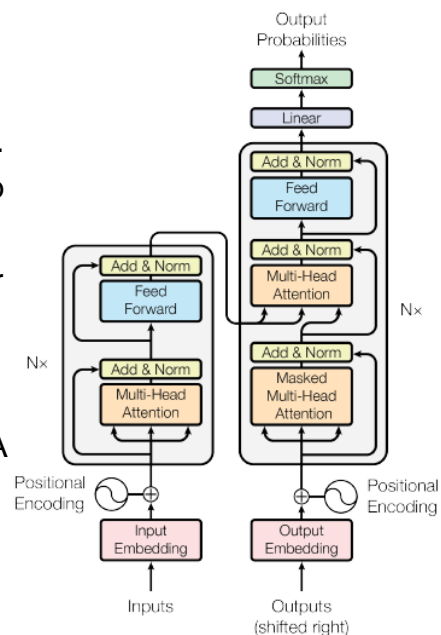
Cross Attention: the key and query matrices act on embeddings of two different token sets.

Multi Head Attention (MHA): run parallel single-head attentions, each with different parameters. Each head produces a change $\Delta \mathbf{E}^{(h)}$ which contributes to the overall change in \mathbf{E} . The computation $\Delta \mathbf{E} = \mathbf{A}\mathbf{V}$ is performed using only $\mathbf{W}_v^{(1)}$ in each head, then concatenating all $\mathbf{W}_v^{(1)}$ (output matrix) to compute the summed $\Delta \mathbf{E}$.

Transformer Architecture:

- The **encoder** block is an unmasked MHA followed by an MLP layer.
- The **decoder** block is a masked MHA, then an unmasked MHA also accepting inputs from the encoder block output, and an MLP layer.
- The transformer as a whole consists of a series of encoder-decoder blocks in series.

If the encoder is fed inputs, the decoder is fed outputs, and the decoder unmasked MHA takes the concatenation of the output MHA and encoder output (shown right), the transformer will be trained to predict the **next** token in the input (trained on shifted examples).



5.6.15. Modern Machine Learning Techniques

Large Datasets: used for Training in Machine Learning in Computer Vision

- **MNIST:** 70,000 28×28 grayscale images of 10 handwritten numerical digits (0-9).
- **CIFAR-10:** 60,000 32×32 RGB images of 10 object classes. (Also: CIFAR-100, 100 classes)
- **ImageNet:** 14 million images of 20,000 classes. (Also: ILSVRC 2012 subset for classification).
- **COCO:** 330,000 images of scenes segmented to 91 classes.
- **Labelled faces in the wild (LFW):** 13,000 images of labelled faces, for face recognition.

Successful Architectures in Computer Vision

- **VGG-16:** CNN, small kernels, frequent max pooling to reduce the number of parameters.
- **AlexNet:** 8-layer CNN, trained on ILSVRC 2012 on GPU.
- **ResNet:** deep CNN, residual blocks (identity layer/skip connections), trained on ILSVRC 2012.
- **Fully convolutional network (FCN):** replaces FC layers with 1×1 convolutional layers. Used for semantic segmentation (using interpolation, each pixel gets a prediction).
- **YOLO:** uses ResNet for instance segmentation (object detection). Can be run fast enough for real-time segmentation from live video on mid-range mobile GPUs.
- **Siamese network:** same weights applied to image pairs, trained with triplet/contrastive loss.
- **U-net:** downsampling (pooling) and upsampling (transposed convolution), extendable with e.g. spatial attention, skip connections in between. Used in biomedical image segmentation.
- **FaceNet:** face recognition, trained on LFW, generates embeddings per face, triplet loss.
- **Variational autoencoder:** input to encoder, output of decoder (with transposed convolution to scale up image). Uses 'evidence lower bound' (ELBO) loss function.
- **ViT (vision transformer).** Uses a transformer network with patches as tokens.

Some of these models are available pretrained in `tensorflow.keras`, others are open source elsewhere, trained on a particular dataset.

Successful Models in Natural Language Processing (NLP, LLMs) and Generative AI:

- **BERT:** first modern transformer used for various NLP purposes. Relatively hard to fine-tune.
- **GPT:** decoder-only transformer for text prediction. GPT-3 is single mode (text → text), while GPT-4 is multimodal (text/image → text). Has found commercial success (ChatGPT).
- **Gemini:** encoder-decoder transformer for (text → text), considered to outperform GPT-4.
- **DALL-E / Stable Diffusion / Midjourney:** generative AI using prompts (text → image).

These models are typically called from an API client side (MaaS: model as a service) rather than being embedded locally, as they contain billions to trillions of parameters (large file sizes).

Other notable leaps in AI technology have occurred in protein folding prediction (AlphaFold 3) and computational fluid dynamics. Generative models for (text → audio) and (text → video, e.g. Sora) have also emerged, with some debate as to their safety and practical usefulness, especially regarding the risks of distributing misleading or false content (e.g. deepfakes, misinformation, LLM hallucinations).

Adversarial Attacks: exploiting backpropagation (fast gradient sign method) to find the smallest possible change to an input image that would result in misclassification due to crossing decision boundaries in latent space. This is a serious concern for high dimensional networks ('the curse of dimensionality').

Transfer Learning: using a pretrained high-performing model as part of the architecture for another neural network with a different task, by freezing its weights and only training the additional layers. The base model acts as a feature extractor from which the additional layers complete the task. Fine tuning is achieved by unfreezing the base model weights.

Few-Shot Learning: learning with only a very small training dataset (few / one / zero per class).

Self-Supervised Learning (SSL): a method of training that can be applied to transformers (e.g. ViT). Models learn useful representations from unlabeled data by predicting parts of the data from other parts (e.g. predicting image patches or masked tokens in NLP).

Reinforcement Learning from Human Feedback (RLHF): the model learns a reward policy based on ranking feedback from human annotators, as a way to improve AI safety or content moderation.


5.5.16. Reinforcement Learning (RL, Unsupervised Learning)

An agent makes decisions (actions) A based on the state of the environment S , using a decision rule (policy) $\pi : S \mapsto A$ (the function $p_\pi(a_n | s_n)$). Each action results in a transition to a new state with probability $p(s_{n+1} | s_n, a_n)$, and a reward signal $r_n(s_n, a_n)$.

- Discounted return: $G_n = \sum_{r=n}^{\infty} \gamma^{r-n} R_{r+1}$ where $0 < \gamma < 1$ is the discount rate.
- State-value function: $V_\pi(s_n) = E[G_n | S_n = s_n]$
- Action-value function: $Q_\pi(s_n, a_n) = E[G_n | S_n = s_n, A_n = a_n]$
- TD(0) update rule at time $n + 1$: $V(s_n) \leftarrow V(s_n) + \alpha(r_{n+1} + \gamma V(s_{n+1}) - V(s_n))$
- ϵ -greedy policy: choose optimal (greedy) $a_n = \operatorname{argmax} Q(s_n, a_n)$ with probability ϵ and choose a random (exploratory) $a_n \in A$ with probability $1 - \epsilon$.
- Q -learning (off-policy TD): find $\pi \in \Pi$ such that $Q_\pi(s_0, a_0)$ is maximised (optimal policy π^*)

$$Q(s_n, a_n) \leftarrow Q(s_n, a_n) + \alpha \left(r_{n+1} + \gamma \max_{a_{n+1}} \{Q(s_{n+1}, a_{n+1})\} - Q(s_n, a_n) \right)$$

Deep Reinforcement Learning / Deep Q Networks (DQN): parameterise $V(s_n; \theta)$ and/or $Q(s_n, a_n; \theta)$ where θ are the weights and biases of a neural network used to estimate the value functions given the state and action, instead of calculating explicitly, which is infeasible for large search spaces.



- Experience replay: store the agent's observations of $(s_n = s, a_n, r_n, s_{n+1} = s')$ in a replay buffer until a given batch size. Train a copy of the network by sampling from the buffer using the original action-values as the truth.
- Gradient of loss function with respect to parameter θ_i :

$$\nabla_{\theta_i} L(\theta_i) = \mathbb{E}_{s, a, r, s'} \left[\left(r + \gamma \max_{a'} Q(s', a'; \theta_i^-) - Q(s, a; \theta_i) \right) \nabla_{\theta_i} Q(s, a; \theta_i) \right]$$

5.5.17. Evolutionary (Genetic) Algorithms for Optimisation

Evolutionary algorithms (EAs) mimic the biological concept of ‘natural selection’ in order to optimise an objective function of the state \mathbf{x} . They are useful when this objective is a ‘black box’ function of a high-dimensional \mathbf{x} .

- Individual: a particular candidate solution \mathbf{x} to the optimisation problem
- Population: a set of individuals
- Genes: encodes the state \mathbf{x} of an individual

Topology Optimisation

Evolutionary algorithms have been applied to topology optimisation in engineering design of load-efficient structures, using e.g. the solid isotropic material with penalisation (SIMP).

5.5.18. Python Examples of Various Machine Learning Tasks

For exploratory data analysis (EDA), Jupyter Notebooks (.ipynb files) can be used to execute one cell block of code at a time, view results step by step, and annotate code.

Loading and Cleaning a Dataset: in this example, from an Excel workbook (.xlsx file)

```
import pandas as pd
X = pd.read_excel('path/to/dataset.xlsx', sheet_name='InputData')
y = pd.read_excel('path/to/dataset.xlsx', sheet_name='OutputData')
null_indices = y[y.isnull().any(axis=1)].index # get rows with null values
X.drop(null_indices, inplace=True) # X: pd.DataFrame
y.drop(null_indices, inplace=True) # y: pd.DataFrame
```

Exploratory Data Analysis: display a report showing various useful metrics.

```
from ydata_profiling import ProfileReport
ProfileReport(df)
```

Standard Scaling and Exploratory Principal Component Analysis: show a scree plot.

```
from matplotlib import pyplot as plt
import numpy as np
from sklearn.preprocessing import StandardScaler
from sklearn.decomposition import PCA

X_std = StandardScaler().fit_transform(X)
pca = PCA(n_components=None)
df_pca = pd.DataFrame(pca.fit_transform(X_std))
explained_variance = pca.explained_variance_ratio_
plt.plot(range(1, len(explained_variance) + 1), explained_variance,
         label='Explained Variance')
plt.plot(range(1, len(explained_variance) + 1),
         np.cumsum(explained_variance), label='Cumulative Explained Variance')
plt.xlabel('Number of Components \n ($n$th largest eigenvalue, descending)')
plt.ylabel('Explained Variance \n (proportion of total variance)')
plt.legend()
plt.title('Principal Component Analysis: Scree plot')
plt.show()
```

Pipelines for Regression Algorithms: train models for 1) linear regression and 2) support vector regression including PCA and polynomial regression with regularisation (lasso). K -fold cross-validation is used and grid search is used for hyperparameter optimisation. The models are evaluated on various error metrics and then saved/loaded to/from the disk.

```

from sklearn.model_selection import train_test_split, GridSearchCV
from sklearn.preprocessing import StandardScaler, PolynomialFeatures
from sklearn.decomposition import PCA
from sklearn.linear_model import Lasso
from sklearn.pipeline import Pipeline
from sklearn.metrics import mean_absolute_error, mean_squared_error
from sklearn.svm import SVR
from sklearn.multioutput import MultiOutputRegressor
import joblib

pipeline_lasso = Pipeline([('scaler', StandardScaler()),
                           ('pca', PCA(n_components=5)), ('poly', PolynomialFeatures(degree=2)),
                           ('lasso', Lasso())])
pipeline_svr = Pipeline([('scaler', StandardScaler()),
                          ('svr', MultiOutputRegressor(SVR()))])

param_grid_lasso = {'lasso__alpha': [0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1.0]}
param_grid_svr = {'svr__estimator__C': [0.001, 0.01, 0.1],
                  'svr__estimator__epsilon': [0.001, 0.01, 0.1]}

def train_model(X: np.ndarray, y: np.ndarray, pipeline: Pipeline,
                param_grid: dict[str: list], model_name: str = '',
                test_size: float = 0.2, cv: int = 10,
                scoring: str = 'neg_mean_squared_error') -> Pipeline:

    mae = lambda y_test, y_pred: mean_absolute_error(y_test, y_pred)
    rmse = lambda y_test, y_pred: mean_squared_error(y_test, y_pred, squared=False)

    X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=test_size)
    search = GridSearchCV(estimator=pipeline, param_grid=param_grid, cv=cv,
                          scoring=scoring)
    search.fit(X_train, y_train)
    best_model = search.best_estimator_
    y_pred = best_model.predict(X_test)
    print(f'{model_name} - Best hyperparameters: {search.best_params_}')
    print(f'MAE: {mae(y_test, y_pred)}, RMSE: {rmse(y_test, y_pred)}')
    return best_model

lasso_model = train_model(X, y, pipeline_lasso, param_grid_lasso, model_name='Lasso')
svr_model = train_model(X, y, pipeline_svr, param_grid_svr, model_name='SVR')
joblib.dump(lasso_model, 'path/to/output/lasso_model.joblib')
lasso_model = joblib.load('path/to/output/lasso_model.joblib')

```

Pipelines for Neural Networks: train a CNN for multi-class image classification with one-hot encoding with pooling, dropout and batch normalisation layers, a validation set, and hyperparameter optimisation. Training neural networks can be computationally intensive, so it can help to use cloud computing (e.g. Google Colab/Cloud Platform, AWS) with access to hardware accelerators (GPUs and TPUs).

```

from matplotlib import pyplot as plt
import numpy as np; import pandas as pd; import cv2; import os
from sklearn.model_selection import train_test_split
from tensorflow.keras.models import Sequential, load_model
from tensorflow.keras.layers import Input, MaxPooling2D, \
    Reshape, Flatten, Dense, Dropout, Conv2D, BatchNormalization
from tensorflow.keras.callbacks import TensorBoard
from keras_tuner.tuners import BayesianOptimization
from keras_tuner.engine.hyperparameters import HyperParameters

tensorboard_callback = TensorBoard(log_dir='model_logs/fit', histogram_freq=1,
    write_graph=True, update_freq='epoch')
df_X = pd.DataFrame(columns=['img', 'gender'], index=None)
folder_names = ['imgs/men', 'imgs/women']
for folder_name in folder_names:
    for file in os.listdir(folder_name):
        img_arr = cv2.imread(os.path.join(folder_name, file), cv2.IMREAD_GRAYSCALE)
        img_arr = cv2.resize(img_arr, (96, 96))
        row = pd.DataFrame({'img': [img_arr], 'gender': [folder_name.split('/')[-1]]})
        df_X = pd.concat([df_X, row], ignore_index=True)
df_y = pd.get_dummies(df_X['gender'])
df_X.drop('gender', axis=1, inplace=True)
X, y = np.array(df_X['img'].tolist()), df_y.values.astype(np.float32)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=1/10)
X_train, X_val, y_train, y_val = train_test_split(X_train, y_train, test_size=1/9)

def build_convnet_model(hp: HyperParameters) -> Sequential:
    model_cnn = Sequential()
    model_cnn.add(Input(shape=(96, 96)))
    model_cnn.add(Reshape((96, 96, 1)))
    model_cnn.add(Conv2D(hp.Int('filters_1', min_value=64, max_value=256, step=32),
        kernel_size=hp.Int('size_1', min_value=2, max_value=4, step=1), activation='relu'))
    model_cnn.add(BatchNormalization())
    model_cnn.add(MaxPooling2D(pool_size=(2, 2)))
    model_cnn.add(Dropout(0.25))
    model_cnn.add(Conv2D(hp.Int('filters_2', min_value=8, max_value=32, step=8),
        kernel_size=hp.Int('size_2', min_value=4, max_value=8, step=2), activation='relu'))
    model_cnn.add(BatchNormalization())
    model_cnn.add(MaxPooling2D(pool_size=(2, 2)))
    model_cnn.add(Dropout(0.25))
    model_cnn.add(Flatten())
    model_cnn.add(Dense(hp.Int('nodes_fc', min_value=32, max_value=96, step=32), activation='relu'))
    model_cnn.add(Dense(2, activation='softmax'))
    model_cnn.compile(optimizer='adam', loss='categorical_crossentropy', metrics=['accuracy'])
    return model_cnn

tuner = BayesianOptimization(build_convnet_model, objective='val_loss', max_trials=10,
    directory='tuner_dir', project_name='model_tuner')
tuner.search(X_train, y_train, epochs=100, validation_data=(X_val, y_val))
best_model = tuner.get_best_models(num_models=1)[0]
best_hyperparameters = tuner.get_best_hyperparameters(num_trials=1)[0]
best_hist = best_model.fit(X_train, y_train, epochs=200, batch_size=128,
    validation_data=(X_val, y_val), callbacks=[tensorboard_callback])
plt.plot(best_hist.history['loss'], label='Training Loss')
plt.plot(best_hist.history['val_loss'], label='Validation Loss')
plt.xlabel('Epoch'); plt.ylabel('Loss'); plt.yscale('log'); plt.legend(loc='upper right'); plt.show()
best_loss = best_model.evaluate(X_test, y_test)
print(f'{best_hyperparameters.values}, loss: {best_loss}')
best_model.save('path/to/output/image_classifier.keras')
best_model = load_model('path/to/output/image_classifier.keras')

```


Pipelines for Neural Networks: train a hybrid CNN-LSTM for one-step-ahead multivariate time series forecasting, with pooling/dropout layers, a validation set, and hyperparameter optimisation.

```
import numpy as np; from matplotlib import pyplot as plt
from sklearn.preprocessing import StandardScaler; from sklearn.model_selection import train_test_split
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Input, MaxPooling1D, TimeDistributed,
    Reshape, Flatten, Dense, Dropout, LSTM, Conv1D
from tensorflow.keras.callbacks import TensorBoard
from keras_tuner.tuners import BayesianOptimization; from keras_tuner.engine.hyperparameters import HyperParameters

tensorboard_callback = TensorBoard(log_dir='model_logs/fit', histogram_freq=1,
    write_graph=True, update_freq='epoch')
df = pd.read_excel('time_series_data.xlsx', sheet_name='Daily')
LOOKBACK, FEATURES = 30, 2 # predict the 31st value of a given subsequence with 2 features
scaler = StandardScaler() # df: pd.DataFrame with features in columns
df_scaled = scaler.fit_transform(df[['adj_close_returns', 'volume_change']].values.reshape(-1, FEATURES))
# generate sliding window arrays (subsequences) from time series data
X = np.lib.stride_tricks.sliding_window_view(df_scaled, (LOOKBACK, FEATURES))
X = X.reshape((X.shape[0], X.shape[2], FEATURES)); y = X[1:, -1, :]; X = X[:-1, :, :]
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=1/10) # train:val:test = 80:10:10
X_train, X_val, y_train, y_val = train_test_split(X_train, y_train, test_size=1/9)

def build_model(hp: HyperParameters) -> Sequential:
    model = Sequential()
    model.add(Input(shape=(LOOKBACK, FEATURES)))
    model.add(Reshape((-1, LOOKBACK, FEATURES)))
    model.add(TimeDistributed(Conv1D(filters=hp.Int('filters_1', min_value=64, max_value=256, step=32),
        kernel_size=hp.Choice('kernel_1', values=[3, 5]), activation='relu')))
    model.add(TimeDistributed(Conv1D(filters=hp.Int('filters_2', min_value=32, max_value=128, step=32),
        kernel_size=hp.Choice('kernel_2', values=[3, 5]), activation='relu')))
    model.add(TimeDistributed(MaxPooling1D(pool_size=2)))
    model.add(TimeDistributed(Flatten()))
    model.add(LSTM(units=hp.Int('units_1', min_value=32, max_value=128, step=32), activation='relu',
        return_sequences=True))
    model.add(Dropout(rate=hp.Float('dropout_1', min_value=0.1, max_value=0.5, step=0.1)))
    model.add(LSTM(units=hp.Int('units_2', min_value=16, max_value=64, step=16), activation='relu'))
    model.add(Dropout(rate=hp.Float('dropout_2', min_value=0.1, max_value=0.5, step=0.1)))
    model.add(Dense(units=FEATURES, activation='linear'))
    model.compile(optimizer='adam', loss='mse')
    return model

tuner = BayesianOptimization(build_model, objective='val_loss', max_trials=10,
    directory='tuner_dir', project_name='model_tuner')
tuner.search(X_train, y_train, epochs=100, validation_data=(X_val, y_val))
best_model = tuner.get_best_models(num_models=1)[0]
best_hyperparameters = tuner.get_best_hyperparameters(num_trials=1)[0]
best_hist = best_model.fit(X_train, y_train, epochs=200, batch_size=128, validation_data=(X_val, y_val),
    callbacks=[tensorboard_callback])
plt.plot(best_history.history['loss'], label='Training Loss')
plt.plot(best_history.history['val_loss'], label='Validation Loss')
plt.xlabel('Epoch'); plt.ylabel('Loss'); plt.yscale('log'); plt.legend(loc='upper right'); plt.show()
best_loss = best_model.evaluate(X_test, y_test)
print(f'{best_hyperparameters.values}, MSE: {best_loss}')
best_model.save('path/to/output/time_series_forecasting.keras')
best_model = load_model('path/to/output/time_series_forecasting.keras')
```

5.6. Computer Vision and Computer Graphics

5.6.1. Representation and Processing of Digital Images

An image can be represented as a $w \times h \times c$ array (w : width in pixels, h : height in pixels, c : number of colour channels e.g. 3 for RGB, 1 for grayscale or monochrome). Each entry in the array is a binary number representing the intensity of the indicated pixel colour. The number of bits used per component channel is the 'bit depth'.

Common standard image dimensions are 640×480 (VGA), 1280×720 (HD), 4096×2160 (4K).

Common colour spaces are RGB(A) (red, green, blue, (alpha)), HSV (hue, saturation, value), CMYK (cyan, yellow, magenta, key/black), L*a*b* (lightness, green-red, blue-yellow), YUV / YCbCr.

Gamma correction: adjusts displayed intensities to account for the perceived nonlinearity of the different colour channels over the range of displayable colours (the gamut).

Image size [MB] = $\frac{\text{bit depth [b]}}{8} \times c \times h \times w \times 10^{-6}$ (+ storage of metadata, no coding/compression)

Conversion of the 3D World to a 2D Image: information is inevitably lost.

The intensity of a pixel $I(x, y)$ is dependent on the position/orientation of the camera, the geometry of the scene, the nature and distribution of light sources, the reflectance spectra of the surfaces (specular or diffuse/Lambertian), and the properties of the lens and CCD. Occlusion may obscure features of specific objects in a scene. In most practical situations, these factors do not affect the desired outcome, so data processing is required to prepare images with features independent of this 'noise'.

Image Processing in the Fourier Domain: useful for analysing filtering operations

Operations can be represented as convolution with a kernel, or multiplication in the Fourier domain.

1D Gaussian kernel: $g_{\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \Rightarrow G_{\sigma}(k) = e^{-\frac{k^2\sigma^2}{2}}$ (unnormalised Gaussian with std.dev $\frac{1}{\sigma}$)

2D Gaussian kernel: $g_{\sigma}(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \Rightarrow G_{\sigma}(k_x, k_y) = e^{-\frac{(k_x^2+k_y^2)\sigma^2}{2}}$, discretised: $(-n, -n) \leq (x, y) \leq (n, n)$

2D discrete convolution: $(w * I)(x, y) = \sum_{i=-n}^n \sum_{j=-n}^n w(i, j) I(x - i, y - j)$ (\mathbf{W} : $(2n + 1) \times (2n + 1)$ kernel array)

$O(n^2)$ 2D convolution as repeated $O(n)$ 1D convolution: $(w * I)(x, y) = w(x) * (w(y) * I(x, y))$

Differentiation as a convolution: $\frac{\partial S}{\partial x} = \frac{S(x+1, y) - S(x-y)}{2} = S(x, y) * [\frac{1}{2}, 0, \frac{-1}{2}]$

Directional derivative: $\nabla S \cdot \mathbf{n} = D_{\mathbf{n}} S(\mathbf{x}) = S_{\mathbf{n}}(\mathbf{x}) \approx S(\mathbf{x} + \mathbf{n}) - S(\mathbf{x})$.

Gradient of a convolution: $\nabla^n s_{\sigma}(x, y) = \nabla^n G_{\sigma}(x, y) * I(x, y)$

Laplacian is approximately a Difference of Gaussians (DoG): $g_{k\sigma}(x, y) - g_{\sigma}(x, y) \approx (k - 1)\sigma^2 \times \nabla^2 g_{\sigma}(x, y)$

Smoothing (low pass filter): $S(x, y) = (g_{\sigma} * I)(x, y) = \sum_{i=-n}^n \sum_{j=-n}^n g_{\sigma}(i, j) I(x - i, y - j)$

5.6.2. Feature Detection

Edge Detection: edges represent regions of sharp change in intensity (max gradient).

Canny edge detection algorithm: **1)** smooth: $s_\sigma = g_\sigma * I$, **2)** gradient: ∇s_σ , **3)** non-maximal suppression: place 'edgels' where $|\nabla s_\sigma|$ exceeds neighbours in directions $\pm \nabla s_\sigma$, **4)** threshold by $|\nabla s_\sigma|$.

Marr-Hildreth detection algorithm: **1)** Laplacian: $\nabla^2 g_\sigma * I$ by DoG approximation, **2)** find zeroes.

The motion of an edge cannot be inferred by looking at edges along (the aperture problem).

Corner Detection: corners have discontinuity in two separate directions (cross correlation)

$$C_n(x, y) = \frac{\mathbf{n}^T \mathbf{A} \mathbf{n}}{\mathbf{n}^T \mathbf{n}} \quad \mathbf{A} = \begin{bmatrix} \langle S_x^2 \rangle & \langle S_x S_y \rangle \\ \langle S_x S_y \rangle & \langle S_y^2 \rangle \end{bmatrix} \quad \lambda_1 \leq C_n(x, y) \leq \lambda_2$$

Smoothed directional derivative of $I(x, y)$ in the direction of \mathbf{n}

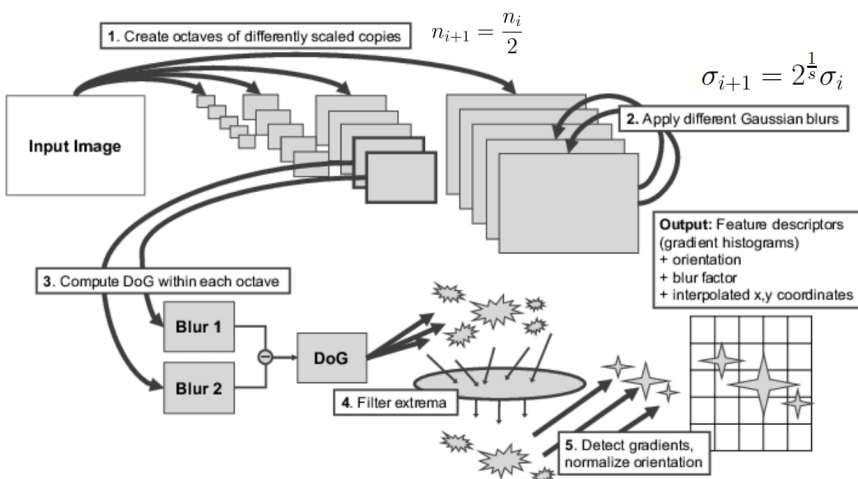
where $\langle S \rangle = s_\sigma = g_\sigma * I$ and $S_x = \partial S / \partial x$, etc. (λ_1, λ_2) : eigenvalues of \mathbf{A}

Harris-Stephens corner detection algorithm: **1)** cross-correlation: $\mathbf{A}(x, y)$, **2)** find $\lambda_1 \lambda_2 = \det \mathbf{A}$ and $\lambda_1 + \lambda_2 = \text{tr} \mathbf{A}$, **3)** threshold by $\lambda_1 \lambda_2 - \kappa(\lambda_1 + \lambda_2)^2$ for some small parameter κ .

If the eigenvalues are both large and distinct, then a corner is likely present.

Blob Detection: round regions enclosed by edges, indicative of keypoints

$|\sigma^2 \nabla^2 s_\sigma(x, y)|$ acts as a normalised 'blob detector' for features at length scale σ (it is a band pass filter). The circular blob diameter with maximum response is $d = 2\sqrt{2}\sigma$.





SIFT Feature Detection: **1)** create copies at sizes $n_{i+1} = n_i / 2$ by subsampling (i : octave number), **2)** apply sequential blurs in each octave, $\sigma_{j+1} = 2^{1/s} \sigma_j$ (j : index within octave), **3)** find $s_{j+1} - s_j$ in each octave (represents $\propto \nabla^2 s_j$ due to DoG approx), **4)** threshold to identify keypoints, **5)** sample Gaussian-weighted 16×16 patch at correct scale around keypoint, **6)** histogram of oriented gradients (HoG) in 4×4 subcells, **7)** concat to a vector, **8)** normalise, truncate outliers ($>0.2 \rightarrow 0.2$), renormalise. The output is a 128-vector per patch.

The SIFT descriptors for a given keypoint feature can be compared (for recognition tasks) by k nearest neighbours on a k -D tree (Section 5.5.16) or passed to a neural network for classification.

Image texture can be characterised by a repeated feature patch (textons).

5.6.3. Planar Perspective Projection

A planar perspective projection of a 3D object in the world is the enlargement about an optical centre point onto an image plane. Assumptions: pinhole camera, no nonlinear distortion.



$$\mathbf{X}_c = (X_c, Y_c, Z_c), \quad \mathbf{x} = (x, y) \text{ at } z = f$$

$$\text{Images are projected through } O_c: \tilde{\mathbf{x}} = \frac{f}{Z_c} \mathbf{X}_c$$

Homogeneous Coordinates: point $\mathbf{X} = (X, Y, Z)$ represented as $\tilde{\mathbf{X}} = [sX, sY, sZ, s]^T$. (WLOG $s = 1$).

If $\tilde{\mathbf{X}} = [X_1, X_2, X_3, X_4]^T$ then $\mathbf{X} = (X_1/X_4, X_2/X_4, X_3/X_4)$.

If $X_4 = 0$ then \mathbf{X} is the point at infinity in direction $[X_1, X_2, X_3]^T$.

1. Rotation in Homogeneous Coordinates:

where $\tilde{\mathbf{X}} = [X, Y, Z, 1]^T$ and $\tilde{\mathbf{X}}_c = [X_c, Y_c, Z_c, 1]^T$.

(\mathbf{R} : 3×3 rotation matrix, \mathbf{T} : 3×1 translation vector;

in Cartesian, $\mathbf{X}_c = \mathbf{R}\mathbf{X} + \mathbf{T}$.)

$$\tilde{\mathbf{X}}_c = \underbrace{\begin{bmatrix} r_{xx} & r_{xy} & r_{xz} & t_x \\ r_{yx} & r_{yy} & r_{yz} & t_y \\ r_{zx} & r_{zy} & r_{zz} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{affine rotation: } \mathbf{P}_r = [\mathbf{R}|\mathbf{T}]} \tilde{\mathbf{X}}$$

2. Projection in Homogeneous Coordinates:

for perspective projection, $3\text{D} \rightarrow 2\text{D}$.

(f : focal length for perspective projection)

$$\tilde{\mathbf{x}} = \underbrace{\begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\tilde{\mathbf{P}}_p: \text{projection}} \tilde{\mathbf{X}}_c = \begin{bmatrix} fX_c \\ fY_c \\ Z_c \end{bmatrix} \rightarrow \mathbf{x} = \begin{bmatrix} fX_c/Z_c \\ fY_c/Z_c \end{bmatrix}$$

3. CCD Imaging in Homogeneous Coordinates:

where $\tilde{\mathbf{w}} = [u, v, 1]^T$ and $\tilde{\mathbf{x}} = [x, y, 1]^T$

((k_u, k_v): pixel length scales, (u_0, v_0): optical centre offset)

$$\tilde{\mathbf{w}} = \underbrace{\begin{bmatrix} k_u & 0 & u_0 \\ 0 & k_v & v_0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{P}_c: \text{imaging}} \tilde{\mathbf{x}}$$

Overall, $\tilde{\mathbf{w}} = \mathbf{P}_c \mathbf{P}_p \tilde{\mathbf{X}}$. Intrinsic camera calibration matrix is $\mathbf{K} = \mathbf{P}_c \mathbf{P}_p$. Matrix $\mathbf{P}_r = [\mathbf{R}|\mathbf{T}]$ is extrinsic. Then, $\tilde{\mathbf{w}} = \mathbf{K}[\mathbf{R}|\mathbf{T}] \tilde{\mathbf{X}} = \mathbf{P}_{ps} \tilde{\mathbf{X}}$ where \mathbf{P}_{ps} is the camera projection matrix (3×4 , 10 dof).

A general 'projective camera model' refers to *any* 3×4 matrix \mathbf{P}_{ps} , which can have 11 dof (constrain by either enforcing $\|\mathbf{P}\| = 1$ or setting $P_{34} = 1$.)

All Notes

5.6. Computer Vision and Computer Graphics

Geometry in Homogeneous Coordinates

- Point in image: $\mathbf{w} = [u, v]^T \rightarrow \tilde{\mathbf{w}} = [u, v, 1]^T$, with reconstructed ray in camera coordinates $\mathbf{X}_c = \mathbf{0} + \lambda \tilde{\mathbf{w}} \rightarrow \mathbf{r} = \lambda [u, v, 1]^T$ (from the optical centre through the point in image plane)
- Line in image: $\mathbf{l}^T \tilde{\mathbf{w}} = 0 \rightarrow l_1 u + l_2 v + l_3 = 0$.
- Intersection of two image lines $\mathbf{l}_1^T \tilde{\mathbf{w}} = 0$ and $\mathbf{l}_2^T \tilde{\mathbf{w}} = 0$ occurs at $\tilde{\mathbf{w}} = \mathbf{l}_1 \times \mathbf{l}_2$.
- Vanishing point (VP) of a line: $\lim_{\lambda \rightarrow \infty} \mathbf{X} = \mathbf{a} + \lambda [a, b, c]^T$ projects to $\tilde{\mathbf{w}} = \mathbf{P}_{ps} [a, b, c, 0]^T$.
- Horizon of a plane $\mathbf{X} = \mathbf{a} + \lambda [a, b, c]^T + \mu [d, e, f]^T$ projects to the line connecting the vanishing points $\tilde{\mathbf{w}}_1 = \mathbf{P}_{ps} [a, b, c, 0]^T$ and $\tilde{\mathbf{w}}_2 = \mathbf{P}_{ps} [d, e, f, 0]^T$, equivalently $\mathbf{l} = \tilde{\mathbf{w}}_1 \times \tilde{\mathbf{w}}_2$.
- Parallel world **lines** have the **same vanishing point**.
- Parallel world **planes** have the **same horizon**.

Camera Calibration: find 3×4 projective camera matrix $\mathbf{w} = \mathbf{P}_{ps} \mathbf{X}$ (11 dof) from a known set $\{\mathbf{w}, \mathbf{X}\}_i$

Points $\mathbf{w}_i = [su_i, sv_i, s]^T$ and $\mathbf{X}_i = [X_i, Y_i, Z_i, 1]^T$ are given. The equations to solve are

$$u_i = \frac{su_i}{s} = \frac{p_{11}X_i + p_{12}Y_i + p_{13}Z_i + p_{14}}{p_{31}X_i + p_{32}Y_i + p_{33}Z_i + p_{34}} \quad \text{and} \quad v_i = \frac{sv_i}{s} = \frac{p_{21}X_i + p_{22}Y_i + p_{23}Z_i + p_{24}}{p_{31}X_i + p_{32}Y_i + p_{33}Z_i + p_{34}}$$

$$\begin{bmatrix} X_1 & Y_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & -u_1 X_1 & -u_1 Y_1 & -u_1 Z_1 & -u_1 \\ 0 & 0 & 0 & 0 & X_1 & Y_1 & Z_1 & 1 & -v_1 X_1 & -v_1 Y_1 & -v_1 Z_1 & -v_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ X_n & Y_n & Z_n & 1 & 0 & 0 & 0 & 0 & -u_n X_n & -u_n Y_n & -u_n Z_n & -u_n \\ 0 & 0 & 0 & 0 & X_n & Y_n & Z_n & 1 & -v_n X_n & -v_n Y_n & -v_n Z_n & -v_n \end{bmatrix} \begin{bmatrix} p_{11} \\ p_{12} \\ p_{13} \\ p_{14} \\ p_{21} \\ p_{22} \\ p_{23} \\ p_{24} \\ p_{31} \\ p_{32} \\ p_{33} \\ p_{34} \end{bmatrix} = \mathbf{0}$$

← n points yields the linear system of equations $\mathbf{A}\mathbf{p} = \mathbf{0}$. Solve by orthogonal least squares (\mathbf{p} = eigenvector of $\mathbf{A}^T \mathbf{A}$ corresponding to smallest eigenvalue, found via SVD: last column of \mathbf{V} in $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$)

If using the constraint $P_{34} = 1$, can write in the form $\mathbf{A}\mathbf{p} = \mathbf{b}$ and solve by (psuedo)inverse (i.e. ordinary least squares) $\mathbf{p} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$.

Refine \mathbf{p} to optimise the reprojection errors: $\mathbf{P}_{ps} = \operatorname{argmin} \left\{ \sum_{i=1}^N (u_i - \hat{u}_i)^2 + (v_i - \hat{v}_i)^2 \right\}$ where $[s\hat{u}_i, s\hat{v}_i, s]^T = \mathbf{P}_{ps} \mathbf{X}_i$.

The 'RQ' decomposition of top-left 3×3 submatrix of \mathbf{P}_{ps} yields $\mathbf{K}\mathbf{R}$, and $\mathbf{T} = \mathbf{K}^{-1} [P_{14}, P_{24}, P_{34}]^T$.

Image Mosaicing: any two images of a general scene with the same camera centre are related by a planar projective transformation (homography) given by $\tilde{\mathbf{w}}' = \mathbf{K}\mathbf{R}\mathbf{K}^{-1}\tilde{\mathbf{w}}$ (\mathbf{K} : camera calibration matrix, \mathbf{R} : rotation between views).

Given key points in an image (e.g. by SIFT), the RANSAC (random sample consensus) algorithm robustly fits keypoints to compute the homography.

All Notes

5.6. Computer Vision and Computer Graphics

5.6.4. Affine Projection

Weak Perspective Projection: when the whole scene has a similar Z_c , orthographic projection is a good approximation of perspective projection.

The weak perspective projection is $\tilde{\mathbf{x}} = \underbrace{\begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 0 & Z_{av} \end{bmatrix}}_{\mathbf{P}_{pl}: \text{weak perspective}} \tilde{\mathbf{X}}_c$ where Z_{av} is the average scene Z_c .

The overall image is then $\tilde{\mathbf{w}} = \mathbf{P}_c \mathbf{P}_{ppi} \mathbf{P}_r \tilde{\mathbf{X}} = \mathbf{P}_{aff} \tilde{\mathbf{X}}$, where $\mathbf{P}_{aff} = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ 0 & 0 & 0 & p_{34} \end{bmatrix}$ is the **affine camera** projection matrix (8 dof).

Planar Weak Perspective Projection: if Z_c is **exactly** constant across the scene, then the projection matrix can be simplified further. The third column can be removed, giving 6 dof.

Cross-Ratio (Section 2.2.10): for four ordered **collinear** points $\{A, B, C, D\} \rightarrow \{a, b, c, d\}$,

The cross ratio of $\{A, B, C, D\}$ is conserved under any perspective projection: $\frac{|AD||BC|}{|BD||AC|} = \frac{|ad||bc|}{|bd||ac|}$.


For five **coplanar** points $\{A, B, C, D, E\} \rightarrow \{a, b, c, d, e\}$, two conserved cross-ratios exist using specially constructed points: if $F = AB \cap CD$, $G = BC \cap AD$, $E_1 = EF \cap AG$, $E_2 = AF \cap EG$, then the cross-ratios of $\{A, E_2, B, F\}$ and $\{A, E_1, D, G\}$ are conserved under any perspective projection. (Notation: $AB \cap CD$ is the intersection of lines AB and CD .)

Another way to form projective invariant quantities is by ‘canonical views’, by constructing a canonical frame curve given two views of an object in perspective projection. A calibration operation aims to map four key points onto the corners of a unit square, giving an ‘invariant signature’.

5.6.5. Stereoscopic Vision (Stereo Vision)

Epipolar Geometry: relation between projections onto different image planes

Viewing a point X from two angles in perspective projection allows for triangulation of the point in 3D, by the intersection of the corresponding rays. **Epipolar lines** constrain the positions of a projected point in each image.



X from two views $\rightarrow x$ and x'
 Coords relate by $X_c' = R X_c + T$
 X_c' is p_1 , X_c is p_2

Epipolar plane Π contains X and the optical centres.

As X moves, the epipoles are invariant and Π rotates about OO' .

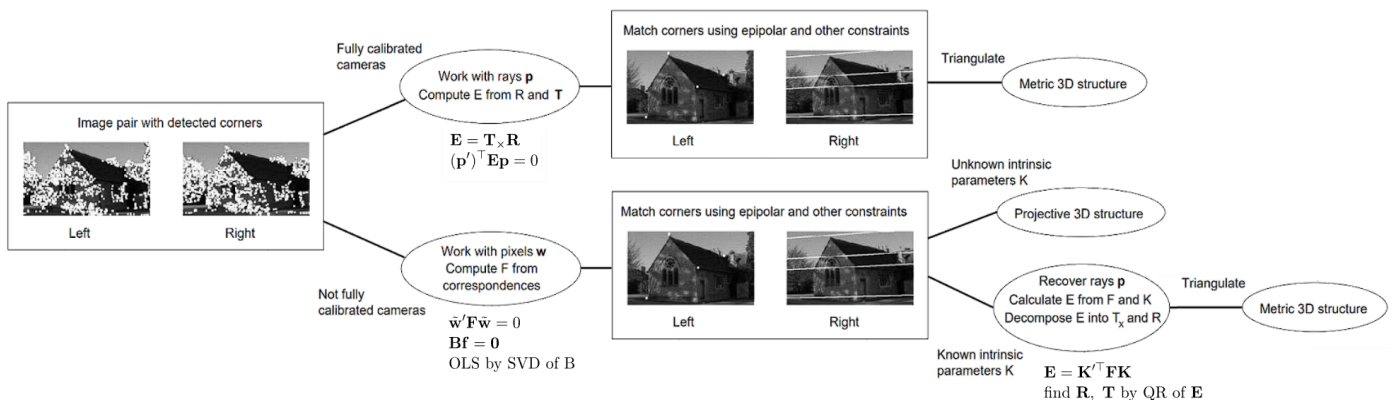
For two correspondences (1 and 2) of one world point X (as a projected ray in camera coordinates $p = \tilde{x} = [x, y, f]^T$) related by a **homography** (linear transformation R , translation T),

- **Epipolar constraint** (ray coordinates): $p_1^T E p_2 = 0$
 ($E = T_x R$: essential matrix, T_x : cross product matrix of T , Section 4.1.6)
 The **epipoles** lie in the **nullspace** of E ($Ee = 0$ and $E^T e' = 0$; and of F^T).

In the limit of parallel cameras, the epipoles approach infinity, allowing depth perception: if $T = [-d, 0, 0]^T$ and $R = I$ then $Z_c = df / (x - x')$.

- **Epipolar constraint** (pixel coordinates): $\tilde{w}_1^T F \tilde{w}_2 = 0$
 ($F = (K_1^{-1})^T E K_2^{-1}$: fundamental matrix, $\tilde{w} = Kp$: imaged ray, K : camera intrinsics).
 The equation of the epipolar line in the right image is $\tilde{l}'^T \tilde{w}' = 0$ where $\tilde{l}' = F \tilde{w}$.

Other constraints to identify correspondences are uniqueness and ordering (for opaque surfaces). The metric structure information (R, T) up to a scale can be computed from the SVD of E .



Summary

Camera Calibration: robust method using RANSAC and nonlinear optimisation

1. Randomly sample 6 out of N image points and world points to form a set of $\{\tilde{\mathbf{w}}_i, \tilde{\mathbf{X}}_i\}$ ($1 \leq i \leq 6$)
2. Using $\tilde{\mathbf{w}} = \mathbf{P}\tilde{\mathbf{X}}$, write a system of 12 equations: $\mathbf{A}\mathbf{p} = \mathbf{0}$ (\mathbf{A} : 12×12 , \mathbf{p} : 12×1)
3. Calculate the SVD decomposition and take $\mathbf{p} = [\text{last row of } \mathbf{V}^T]$ where $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$
4. Calculate $|\mathbf{A}\mathbf{p}|$ for the full set of N points and count the 'inliers' (sufficiently close to zero)
5. Iterate from step 1 until the count of inliers is maximised. Use this \mathbf{P} as a seed for Step 6.
6. Minimise the reprojection errors: $\mathbf{P} = \operatorname{argmin} \left\{ \sum_{i=1}^N (u_i - \hat{u}_i)^2 + (v_i - \hat{v}_i)^2 \right\}$ where $[s\hat{u}_i, s\hat{v}_i, s]^T = \mathbf{P}\mathbf{X}_i$.
7. The 'RQ' decomposition of top-left 3×3 submatrix of \mathbf{P} yields \mathbf{KR} .
8. Calculate $\mathbf{T} = \mathbf{K}^{-1} [p_{14} \ p_{24} \ p_{34}]^T \rightarrow \mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{T}]$.

Stereoscopic Correspondences: recover 3D structure from two calibrated cameras

1. Given two images related by an unknown homography, identify keypoints.
2. Use SIFT to register invariant features $\mathbf{x}_i \in \mathbb{R}^{128}$ and $\mathbf{x}'_j \in \mathbb{R}^{128}$ in a k -D tree.
3. Find nearest neighbour correspondence estimates $\{\mathbf{x}_i, \mathbf{x}'_j\}$ from the images.
4. Refine the correspondence estimates using the RANSAC algorithm:
 - a. Randomly sample 8 pairs of correspondences $\{\tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}'_i\}$.
 - b. Using $\tilde{\mathbf{w}}'^T \mathbf{F} \tilde{\mathbf{w}} = 0$, write a system of 8 equations: $\mathbf{A}\mathbf{f} = \mathbf{0}$ (\mathbf{A} : 8×9 , \mathbf{f} : 9×1), solve with SVD as before to find 3×3 fundamental matrix \mathbf{F} with $F_{33} = f_9 = 1$.
 - c. Find $\tilde{\mathbf{w}}'^T \mathbf{F} \tilde{\mathbf{w}}$ using this estimate for \mathbf{F} on the whole dataset. Count the number of correspondences for which this value is sufficiently close to zero.
 - d. Iterate, finding new \mathbf{F} 's until the count of inliers is maximised. Take this as \mathbf{F} .
5. Enforce rank 2 constraint in \mathbf{F} by setting its smallest singular value to zero.
6. Compute $\mathbf{E} = \mathbf{K}'^T \mathbf{F} \mathbf{K}$.
7. Calculate the SVD decomposition $\mathbf{E} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. Then $\mathbf{R} = \mathbf{U}\mathbf{Y}\mathbf{V}^T$ and $\mathbf{T} = [\text{last column of } \mathbf{U}]$, where $\mathbf{Y} = [[0, -1, 0]; [1, 0, 0]; [0, 0, 1]] = \{\text{rotation about } V_z \text{ by } 90^\circ \text{ anticlockwise}\}$
8. There are 4 possible solutions for using $\pm\mathbf{T}$ and $\{\mathbf{R}, \mathbf{R}^T\}$: resolve ambiguity by ensuring all visible points lie in front of both cameras. Then $\mathbf{P}' = \mathbf{K}'[\mathbf{R}|\mathbf{T}]$ and $\mathbf{P} = \mathbf{K}[\mathbf{I}|\mathbf{0}]$.
9. To compute $\tilde{\mathbf{X}}$, solve $\tilde{\mathbf{w}}' = \mathbf{P}'\tilde{\mathbf{X}}$ and $\tilde{\mathbf{w}} = \mathbf{P}\tilde{\mathbf{X}}$ (4 equations in 3 unknowns: least squares).

5.6.6. Hough Transform and Radon Transform

Hough Transform: detects lines in 2D images.

Straight lines in images can be parameterised by (r, θ) , the polar coordinates of the closest point on the line to the origin of the image coordinate system. The equation of the line is then

$$r = x \cos \theta + y \sin \theta \quad (\text{Hough transform in } (r, \theta) \text{ space: Hesse normal form}).$$

A given (x, y) point corresponds to a sinusoid in (r, θ) space. Sinusoids are superposed for each pixel to build the transformed image in Hough (r, θ) space. The intersections of these sinusoids (represented as maxima in total amplitude) occur at values of (r, θ) where a line with these parameters exists in the original image.

The intersections are thresholded and returned as detected lines in the image.

Radon Transform:

An image $f(x, y)$ can be mapped to (r, θ) directly using the Radon transform:

$$Rf(r, \theta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \delta(r - x \cos \theta - y \sin \theta) dx dy$$

This is mathematically equivalent to the Hough transform in the continuous infinite size limit.

The **inverse Radon transform** is useful in tomography (reconstructing solid 3D objects from a stack of 2D image slices through), widely used in biomedical imaging (e.g. CT scans) to reconstruct from X-ray transmission intensity data. It is computed using the filtered backprojection formula.

5.7. Accounting, Finance and Business

5.7.1. Basics of Accounting for Businesses

Types of Entities

- **Non-Business Organisation:** exists to meet a societal need where making profit is not the goal.
- **Business Organisation:** exists to sell goods and/or services where making profit is a goal.
 - **Sole Trader / Proprietorship:** owned by one person, the same legal entity.
 - **Limited Company (LTD):** owned by one person, a different legal entity.
 - **Partnership:** owned by two or more associated people.
 - **Corporation:** owned by one or more shareholders.
 - **Publicly Accountable Enterprise:** sells shares on the stock exchange (after IPO).
 - **Private Enterprise:** shares owned by a group of associated people.
 - **Limited Liability (LLC):** shareholders are not personally responsible for debts.

Types of Accounting:

- **Financial Accounting:** reports the breakdown of financial decisions of a company, aimed at informing external stakeholders. Standardised according to GAAP / IFRS, enforced depending on locality (US: FASB, elsewhere: IASB).
- **Managerial Accounting:** reports the internal performance of the company, aimed at informing managers about decisions on projects and employees.
- **Tax Accounting:** reports the overall profits made by a company to facilitate collection of tax revenue by the government (US: IRS, UK: HMRC)

Incentives: actors may be incentivised to bias a report e.g. managers aim to minimise reported profits for a tax return (to pay less tax), but maximise profits when presenting to investors. This also applies for external bodies e.g. consultants, auditors and credit rating agencies which are paid by the company for evaluation, though these entities are typically more reputable as third parties.

Key Actors in Accounting

- **Financial Accountants:** hired by the firm to prepare their financial reports. Supervised by the Chief Financial Officer (CFO).
- **Auditors:** external agents hired by the firm to assess and verify the financial reporting quality of the firm. Typically works for an accounting firm (the 'Big 4': Deloitte, KPMG, EY, PwC). Auditors use statistical methods to determine the risk of a significant deviation from accounting standards.
- **Forensic Accountants:** investigates fraud and financial irregularities, and advises on financial disputes.
- **Investors:** sources of finance for developing companies, especially at the entrepreneurial stage. Seed funding may be received from angel investors, venture capitalists and banks (investment bankers) can provide sources of funding as well as advice on decision making, though VCs typically require a position on the company board (equity stake) for their services.
- **Standards Setters:** the entities making and enforcing the rules on accounting.

5.7.2. Financial Statements (Books)

Stakeholders may make major investment decisions based on financial reports:

- **Income Statement:** shows the breakdown of **revenue and expenses**, as summed flows over time.
- **Balance Sheet:** shows the total **assets, liabilities and equity**, at a point in time.
- **Cash Flow Statement:** shows the total sum of transactions of the company.
- **Environmental, Social and Governance (ESG Data):** currently being formalised, but is not yet universally mandatory as audit quality varies. Used to inform policymakers.

The ‘aggregation exercise’ is performed in each accounting cycle: ongoing accruals are recorded in a **journal**, noting the date, account debited and credited. To prepare the accounts, **T-accounts** are drawn up, grouping transactions by account. The balance for each account is found by summing and the necessary financial books are written.

Double Entry Bookkeeping: any one transaction affects two variables in the **accounting equation**, **Assets = Liabilities + Equity**. Financial statements are structured using ‘T-accounts’ to reflect this:

Ledger	General Ledger						Temporary Ledger			
Account Type	Assets Debits Credits		Liabilities Debits Credits		Equity (Shares) Debits Credits		Revenue Debits Credits		Expense Debits Credits	
Increase or Decrease										

Debits (Dr) increase the **assets**. For the temp ledger, **credits (Cr)** increase the **equity** (which decreases the assets). The temporary ledger balance (revenue - expense) is realised as ‘retained earnings’ in Equity.

Company Name Income Statement For the month or year ended xx/xx/xxxx	
Sales Revenue or Service Revenue	_____
Cost of Goods Sold (COGS)	_____
Gross Margin	_____
Operating Expenses	
e.g. Depreciation Expense	
e.g. Selling, General and Administrative Expense (SG&A)	
e.g. Research and Development Expense (R&D)	
Operating Income or Loss	_____
Other Revenues or Expenses	
e.g. Gain or Loss from disposal of long-term assets	
e.g. Impairment Loss	
Earnings Before Interest and Tax (EBIT)	_____
Interest Expense	
Tax Expense	
Net Income or Loss	_____

Company Name Statement of Cash Flow For the month or year ended xx/xx/xxxx	
Cash flow from operating activities:	
Net Income	
Adjust for non-cash items from the Income Statement:	
e.g. +depreciation/amortisation expense	
e.g. -gain/+loss from disposal of long-term assets	
e.g. + bad debt expense	
Adjust for changes in non-cash current assets and current liabilities:	
e.g. - increase in non-cash current assets	
e.g. + increase in current liabilities	
Total Cash flow from operating activities	_____
Cash flow from investing activities:	
Cash used to purchase long-term assets	
Cash received from disposal of long-term assets	
Total cash flow from investing activities	_____
Cash flow from financial activities:	
Cash payment of long-term liabilities	
Acquisition of long-term liabilities for cash	
Cash repurchase of share capital	
Issuance of new shares for cash	
Cash payment of dividends	
Total cash flow from financing activities	_____
Net Increase (Decrease) in cash	_____
Cash Balance at the beginning of the year	_____
Cash Balance at the end of the year	_____

Company Name Balance Sheet As at xx/xx/xxxx	
Assets	Liabilities and Shareholders' Equity
Current Assets	Liabilities
e.g. Cash	Current liabilities
e.g. Inventory	e.g. Accounts Payable (A/P)
e.g. Accounts Receivable (A/R)	e.g. Salaries Payable
Non-Current Assets	e.g. Interests Payable
e.g. Property, plant and equipment (PPE)	Non-Current liabilities
e.g. Research and Development	e.g. Bank Loans
e.g. Goodwill	Total Liabilities
Total Assets	_____
	Shareholders' Equities
	e.g. Share Capital
	e.g. Retained Earnings
	Total Equity

	Total Liability and Shareholders' Equity

5.7.3. Simple Metrics in Financial Accounting

Metrics of profitability:

- Gross Profit Ratio = Gross Profit / Sales Revenue
- Profit Margin Ratio = Net Income / Sales Revenue
- Return on Assets = Net Income / Total Assets

Metrics of liquidity:

- Working Capital = Current Assets - Current Liabilities
- Current Ratio = Current Assets / Current Liabilities.

Negative numbers on books are recorded in parentheses e.g. (100) is \$ -100.

Quantitative easing: a monetary policy in which a central bank buys government bonds, increasing the assets of commercial banks, stimulating economic activity. First used to mitigate the 2008 financial crisis.

5.7.4. Revenue Recognition and Bad Debt

- Revenue is earned when a product or service has been delivered and accepted by the customer.
- When money is received before a service is delivered, record the revenue as **'deferred revenue' (a liability)** until the service is delivered, since the company owes a refund until this time.
- When a service is delivered before money is received, record the value as **'accounts receivable' (A/R, an asset)** until the bill/invoice is fully paid. In B2B, many companies offer credit to customers when they make sales.
- If the customer seems unable to pay off their debt, the reported A/R value is now inaccurate, and an estimate of the loss is recorded in **'allowances for bad debt' (AFDA)**, a contra-asset valuation account). If the firm gives up on recovering this bad debt, it is **written off**, balancing the AFDA and crediting from the A/R account.
- To compute the bad debt expense, statistical methods based on an **ageing schedule (arrears)** are used. The longer the account remains overdue, the higher the probability it will default. Established firms have long histories of data to estimate probability of defaulting by lateness of the payment. Increment to ending AFDA = A/R amount due × Probability of default (summed over arrears buckets).

5.7.5. Statement of Cash Flows and Fraud in Accounting

The balance sheet and income statement are not sufficient to assess a company's financial health. Without cash, a company cannot operate and goes bankrupt (even if it owns assets). To predict this, financial analysts look at the cash inflows and outflows over a period, described in the cash flow statement.

The cash flow statement's net change in cash in year t gives the increase in balance sheet cash between years $t - 1$ and t .

Cash Flow Statement: divided into three sections

- Cash flow from **operations**: from income statement/balance sheet changes, also paying off interest on debts (as of US-GAAP ASC 230).
- Cash flow from **investing**: from investments and long-term assets.
- Cash flow from **financing**: obtaining long-term debt (borrowing), paying off principal of debt, issuing equity, paying dividends.

Operating cash flow is usually computed indirectly, subtracting **non-cash operations from 'net income'** on the income statement.

Gains and **Losses** are increases and decreases in equity (net assets), except those that result from revenues or investments by owners.

Fraud in Accounting: exploiting grey areas of accounting to misrepresent cash flow data.

- Channel stuffing: selling more products to distributors than they are capable of selling to the end customers in order to inflate sales.
- Underestimating default probabilities for a bad debt expense account to inflate net A/R income.

5.7.6. Financial Interest and the Time Value of Money

Interest: accrual of money per unit time

- Simple interest: $A = P(1 + rt)$ (interest = $A - P = Prt$)
- Compound interest: $A = P\left(1 + \frac{r}{n}\right)^{nt}$
- Continuous compound interest: $A = P e^{rt}$
- Annual percentage rate (APR): $APR = \frac{\text{interest } [\$] + \text{fees } [\$]}{\text{principal } [\$]} \times \frac{365 \text{ [days]}}{\text{loan term [days]}}$ (effective interest rate)
- Return on investment (RoI) = % increase in $P = \frac{A - P}{P} = \frac{\text{profits}}{P}$ ($\times 100\%$)

(A : accrued amount; future value FV, P : premium (present value PV), r : interest rate (as $\frac{r[\%]}{100}$) per year, t : time (years), n : number of compounding periods per year e.g. weekly $\rightarrow n = 52$, monthly $\rightarrow n = 12$. When working with r as an inflation rate, A in the future has the same purchasing power as P now.)

Time Value of Money: a unit of money is generally 'more valuable' now than in the future.

Cash amounts are only comparable when referring to the same point in time.

- Future Value (FV) = Present Value (PV) $\times (1 + r)^n$ (where $r > 0 \rightarrow FV > PV$)
- Present Value (PV) = $\frac{\text{Future Value (FV)}}{(1 + r)^n}$

If the net present value (NPV) of an investment is negative, it is likely not worthwhile, as the discounting exceeds the interest.

Perpetuities and Annuities: cash flows vary between times due to interest and inflation

- **Perpetuity:** stream of constant cash flows. $PV = \sum_{n=1}^{\infty} \frac{C}{(1+r)^n} = \frac{C_1}{r}$
- **Growing perpetuity:** stream of rising cash flows. $PV = \sum_{n=1}^{\infty} \frac{C_1(1+r_g)^{n-1}}{(1+r)^n} = \frac{C_1}{r-r_g}$ ($C_1 = C_0(1+r_g)$)
- **Annuity:** stream of N cash flows. $PV = \sum_{n=1}^N \frac{C}{(1+r)^n} = \frac{C}{r} \left(1 - \frac{1}{(1+r)^N}\right)$
- **Growing annuity:** stream of N rising cash flows. $PV = \sum_{n=1}^N \frac{C_1(1+r_g)^{n-1}}{(1+r)^n} = \frac{C_1}{r-r_g} \left(1 - \left(\frac{1+r_g}{1+r}\right)^N\right)$

(r : discount rate / hurdle rate / (opportunity) cost of capital, r_g : long-term growth rate)

5.7.7. Valuation of Stocks and Shares

Types of shares:

- Growth shares: investors expect to benefit from capital gains (future growth of earnings)
- Income shares: investors seek cash dividends

Metrics of share valuation:

- Book equity per share (BVPS) = $\frac{\text{value of common equity on balance sheet}}{\text{number of shares outstanding}}$
- Earnings per share (EPS) = $\frac{\text{net income} - \text{dividends to preferred shares}}{\text{number of shares outstanding}}$
- Return on equity (RoE) = $\frac{\text{net income}}{\text{shareholders' equity}} = \frac{\text{EPS}}{\text{BVPS}} = \frac{\text{net income} - \text{dividends to preferred shares}}{\text{value of common equity on balance sheet}}$
- Payout ratio = $\frac{\text{dividends}}{\text{EPS}}$
- Plowback / Retention ratio, $b = 1 - \text{Payout ratio}$
- Price per earning (P/E ratio) = $\frac{\text{market value per share (share price)}}{\text{EPS}}$

Valuation of Shares: common shares can be paid out as either dividends or capital gains.

(D : dividends, P : share price, discount rate: $r = \frac{D_1}{P_0} + r_g$, growth rate: $r_g = b$ (plowback) \times RoE.)

DCF share price is the sum of discounted future dividends (PV): $P_0 = \frac{P_1 + D_1}{1+r} = \sum_{n=1}^{\infty} \frac{D_n}{(1+r)^n}$.

Gordon-Shapiro growth model: $D_n = D_0(1+r_g)^n$ then $P_0 = \sum_{n=1}^{\infty} \frac{D_1(1+r_g)^{n-1}}{(1+r)^n} = \frac{D_1}{r-r_g}$ and so $P_n = \frac{D_{n+1}}{r-r_g}$.

For nonlinear growth then linear, use $P_0 = \frac{P_N}{(1+r)^N} + \sum_{n=1}^N \frac{D_n}{1+r}$ (where $P_N = D_N(1+r_g) / (r-r_g)$).

Valuation of Straight Bonds: investor (lender) buys the bond to receive fixed returns (like a loan)

- Maturity date: bond expiry date, issuer returns the face value amount to the lender
- Face value F / Par value: the amount promised to pay the investor on the maturity date
- Coupon C : periodic interest paid while the bond is held (before maturity):

$$\text{Coupon, } C \text{ [\$/]} = \text{Face Value, } F \text{ [\$/]} \times \text{Coupon rate, } r_C \text{ [\%/]}$$

- Yield to maturity r_{YTM} (YTM): annual expected return for the investor.
- Current yield: annual coupon divided by price

Fair value of bond is the sum of discounted debt cash flows: $P_0 = \frac{F}{(1+r_{\text{YTM}})^N} + \sum_{n=1}^N \frac{C}{(1+r_{\text{YTM}})^n}$.

Investors are at risks e.g. default risk, liquidity risk, regulatory risk, interest rate risk.

5.7.8. Cash Flow Analysis

The net present value (NPV) concept can be used to check feasibility of an investment (e.g. for a project):

$$\text{Net Present Value (NPV)} = \left(\sum_{n=1}^N \frac{C_n}{(1+r)^n} \right) - P_0 \text{ (PV of returns - } P_0 \text{: initial investment / cash outlay)}$$

Considerations in conducting a cash flow analysis:

- Use **incremental** cash flows: differences between revenue with and without the project.
- Account for **side effects** on other cash flows e.g. erosion (decrease) or synergy (increase).
- **Ignore sunk costs**: costs in the past cannot influence future decisions.
- Consider potential revenues excluded by the project as **opportunity costs**.
- Allocated cost should be included entirely when budgeting (as opposed to in accounting).

Free cash flow to the firm (FCFF) method: typical tabular layout

- EBIT (earnings before interest and tax) = (sales revenue – costs) – depreciation [on income statement]
- NOPDAT (net operating profit after tax) = EBIT × (1 – r_{tax})
- CFO = NOPDAT + depreciation + amortisation
- Depreciation = Initial Outlay / Project Length (using ‘straight line’ depreciation method)

	Year 0	Year 1	Year 2 ...	Year N
Cash flow from operations (CFO)	0	CFO	CFO	CFO
Capital investment (CapEx)	(initial outlay) (-ve)	0	0	0
Investment in working capital (ΔNWC)	(WC) (-ve)	0	0	WC (+ve)
Free cash flow [sum down columns]	C_0	C_1	C_2	C_N
Discounted cash flow (DCF)	C_0	$C_1 / (1 + r)$	$C_2 / (1 + r)^2$	$C_N / (1 + r)^N$
Net Present Value (NPV) = sum of DCFs				

If $\text{NPV} > 0$ then the project is expected to be profitable. If $\text{NPV} < 0$ then the project is unprofitable.

5.7.9. Break-even Analysis

Internal rate of return (IRR): the discount rate r such that a cash flow analysis returns $\text{NPV} = 0$.

If $\text{IRR} > r_{\text{required}}$ then the project is profitable. If $\text{IRR} < r_{\text{required}}$ then the project is unprofitable.

- If the function $\text{NPV}(r)$ has multiple zeros, this method may be invalid (need to graph curve).
- When comparing two projects, higher IRR does not always imply higher NPV.

$$\text{Profitability Index} = \frac{\text{NPV [at fixed } r]}{\text{initial investment}}$$

Sensitivity analysis: study how NPV varies with inputs e.g. $\frac{\Delta \text{NPV}}{\Delta x}$ or $\frac{d(\text{NPV})}{dr}$.

5.7.10. International Business

An international business conducts some aspect of their business across borders e.g. products and services, capital (trading), people (employees) and information (digital companies).

Incentives for internationalisation of business:

- Market seeking: economic growth, infrastructure, less competition
- Efficiency seeking: regulatory arbitrage, lower labour costs, lower taxes / tax avoidance
- Innovation seeking: access to talent and specialist workforce
- Resource seeking: critical materials

Barriers to internationalisation of businesses:

- Trade tensions: politicians may legislate against trade with some countries
- Armed conflict: international business is deprioritised in times of war
- Disruptions: (e.g. Houthi attacks on ships in the Red Sea/Suez canal in 2024).
- Climate Change: shortages or failures of land-based raw materials

Strategy of internationalisation: experience or experimentation?

- Deliberate: top-down, global strategy with careful planning.
- Emergent: bottom-up, local strategy embracing learning the new cultural norms.

Modes of entry: resource commitment usually starts low and increases over time (Uppsala model).

- Non-equity modes: import/export, outsourcing, licensing, franchising.
- Joint venture / partial acquisition. Can help with gaining local knowledge.
- Equity modes: green field / full acquisition.

Increasing global supply chain resilience (reconfiguration):

- Shorter supply chains: onshoring / near-shoring
- Political alignment: friend-shoring
- Diversification: dual/multi-sourcing e.g. 'China Plus One' strategy for the West
- Bifurcations: separate supply chains for different markets

Examples of Internationalisation in Tech and Energy Industries

- Tesla expands into China in 2014: to take advantage of the EV market, manufacturing in China is necessary to avoid high tariffs.
- Many tech companies (Arm, Nvidia, etc) set up in high-talent areas e.g. Silicon Valley (California) or Cambridge.
- LG Chem builds a factory in the US in 2024 to take advantage of the IRA tax credits for making EV battery cathodes.

Factors affecting International Business for Energy and Tech Industry

- **US-China Trade War** (President Trump, 2019): nationalist protectionism increasing barriers to trade between the USA and China.
- **US Inflation Reduction Act** (President Biden, 2022): promotes US decarbonisation by subsidising domestic clean energy production with tax credits.
- **Critical minerals**: top producers (mining and refining) include China (Sb, Co, Ga, graphite, In, Mg, REE oxides, Si, Te, Sn, W, V), Vietnam (Bi), Australia (Li), DR Congo (Ta), South Africa (Pt), Brazil (Nb).
- **Artificial intelligence**: potential rise of LLMs and AGI in the near-to-medium future may force a restructuring of the software workforce.
- **Fragility in Semiconductor Industry**: impacts all technology and associated supply chains.
 - US firms design chips using software relying on IP licences from Europe.
 - Manufacturing equipment developed in the US, Japan and Europe. ASML (Europe) is the only producer of extreme UV lithography systems as of 2024: they sell to chip makers.
 - Silicon is mined and refined in the US, processed into wafers in Japan and South Korea.
 - Chips are manufactured and packaged in Taiwan and Malaysia. TSMC makes 92% of global advanced semiconductor chips.
 - Processors are assembled into electronic products in China.

5.7.11. Institutional Theory in International Business

Institution: a taken-for-granted set of organising principles ('rules of the game; social scripts').

Institutional theory was developed to criticise the 19th century 'economic man' theory, explaining how and why business and people behave irrationally. Modern (neo-institutionalism) theory explores why some organisations of the same type always have a very **similar structure** everywhere in the world, due to **regulative, normative and cognitive forces** driving uniformity.

Examples of institutions: democracy, marriage, banks, places of religious worship, schools...

Properties of institutions: persist for a long time, are collective, are mostly taken for granted, **guide and constrain social behaviour**, simplify decision-making, provide order, build trust and legitimacy, are hard to change (inertia).

5.7.12. Analysis of Institutional Distance for Internationalisation

Institutions in International Business:

- Institutional arrangements are highly context dependent. They may be literal (geographic), formal (legal, political, economic) or informal (cultural, religious, linguistic).
- Institutions are incredibly powerful structures: underestimating or misunderstanding 'the rules of the game' while entering a given country may lead to failure. However, it is hard to fully understand the institutional environment from the outside due to its taken-for-granted nature (liability of outsidership).
- If we need to change the institutional environment in order to enter a country, we need to understand it entirely, and will still be a very challenging task.

Analysing Potential for Internationalisation:

- Why is the company successful currently in the home country?
- What is their positioning (reputation, service type) in the market?
- Identify regulations responsible for forming institutions.
- Identify the cultural forces that drive these regulations.
- What are the implications of these differences on the current business model? Is adaptation feasible?

Types of Institutional Distance: may be literal (**geographic**), formal (**legal, political, economic**) or informal (**cultural, religious, linguistic**).

Approaches to Reduce Institutional Distance: adapt to the target, or influence the target?

- Cultural legitimisation ('glocalisation'): localisation, polycentric pricing, omnichannel marketing, workforce training, social responsibility. May require ethnographic studies for market research.
- Market to the diaspora of the target country, then attract locals (e.g. bubble tea in the West).
- Use cultural interests to promote associated products (e.g. K-pop/Korean wave → Asian food supermarket)
- Joint ventures with a local company to gain trust or avoid negative perception of host country.
- Institutional entrepreneurship: lobbying to relax regulations, marketing, PR, strategic collaboration with industry groups.

Example: commerce in the US (e.g. Walmart). Relevant institutional norms in the US include:

1. Economies of scale: a more capitalist free market allows monopolisation by rapid growth.
2. Driving culture: can have large stores, spaced apart.
3. Very high standards of customer service: requires more rigorous management of employees).
4. Low price guarantees: requires lower wages for employees.

These factors are unlikely to work in Europe (hence Walmart's failure in Germany) due to institutional distance: e.g. more government oversight, regulations on employee rights, salaries and environmentalism, limited store opening hours, stronger unions for employees, mandatory holidays, less emphasis on customer service and the appearance of friendliness.

5.7.13. Cross-Cultural Communication

Models for Understanding Culture:

- **Human mental programming hierarchy:** 1) human nature (universal; inherited), 2) culture (specific to group; learned), 3) personality (specific to person; inherited and learned).
- **Cultural iceberg:** only behaviours and practices are observable. The hidden factors informing them are the perceptions, attitudes, beliefs, values, which in turn are influenced by climate, geography, demographics, economics, media, education, ideology, religion.
- A nation is often not the best unit to study a culture: the cultural unit may be larger or smaller, and may not be geographically united.

Dimensions determined by culture (Hofstede and more recent critical work):

1. **Individualism vs Collectivism.** Integration of people into primary groups. Most significant.
2. **High vs Low power distance.** Solutions to basic inequality.
3. **Low vs High uncertainty avoidance.** Response to stress in the presence of unknowns.
4. **Motivation towards Achievement / “Masculinity vs Femininity”.** Division of emotional and gender roles in society, into competitive/tough/assertive vs cooperation/relationship building.
5. **Long vs Short term orientation.** Focus of people’s efforts in the present or the future.

Other important dimensions are **Indulgence vs Restraint** (Personal happiness, freedom of expression, and the importance of leisure) and **Perception of Time** (Sequential vs Overlapping).

Differences along any of these dimensions can lead to unexpected clashes in many social interactions while doing business.

5.7.14. Institutional Voids and Developmental Distances

Developing economies may have a lack of formal institutions (voids): within each dimension of institutional distance, there is also a developmental distance (informality). Often occurs in:

- Functioning political, economic and legal systems
- Hard infrastructure e.g. roads, rail, airports, seaports, telecomms, energy
- Soft infrastructure (business ecosystems) e.g. talent, logistics, information availability

Conglomerates (highly diversified, often family-founded corporations) tend to perform well in countries with voids as they provide all necessary services at once, becoming MNCs.

Network Effects: some services become intrinsically better when more people use them. Commonly exploited by digital companies to cement a monopoly in a void.

Examples: taxi apps (more drivers, more users), e-commerce (more buyers, more sellers), social media (more users, more engagement), shopping malls (more shops, more buyers).

'Super-apps' are all-in-one digital service apps, proving highly successful in 'recently' developed countries, combining the benefits of conglomerate-style void-filling and the network effects. Localisation helps to tailor services to the local market e.g. Grab in South-East Asia. Super apps are not popular in the West due to anti-competition laws, already matured institutions, data privacy concerns etc.

Technological leapfrogging: using voids as opportunities by building the service around what is already there on the ground, not worrying about the necessary infrastructure in the home country e.g. mobile internet in Africa, mobile fintech, renewable energy in Asia, electric mobility.

5.7.15. Business in the Anthropocene Epoch (Climate Change and Pollution)

The 'anthropocene epoch' refers to the observation that collective human activity is impacting the planet itself.

Climate risk: the negative financial effects of climate change on business e.g. supply chain disruption, higher costs, lower sales, transportation disruption, food shortages, regulatory risk.

Double materiality: the recognition that companies have a significant impact on the climate and should be held responsible for the waste they produce (rather than the consumer). Implementations include ESG data reporting (Section 5.7.2), extended producer responsibility (EPR) schemes, right to repair laws and results-based government funding.

Scopes of climate reporting (e.g. greenhouse gas emissions):

- Operational emissions: company-owned vehicles and facilities.
- Operational resources: purchased heating, cooling, energy, steam...
- Upstream activity (suppliers): assets, employee commuting, purchased goods/services, business travel, waste, fuel/energy, transport/distribution, capital goods
- Downstream activity (consumers): processing of sold product, use of product, leased facilities, investments, franchisees activity, end of life treatment (LCA)

Wastage per person is increasing over time due to increased consumerism and illegal business practices e.g. planned obsolescence, proprietary interfaces, fast fashion, as well as growth in emerging economies. In developed economies, 'degrowth' has become a viable strategy for decreasing waste without short-term economic loss. Companies producing in underdeveloped countries are being scrutinised for the ethics of their supply chains e.g. slavery, poor working conditions, sometimes resulting in supply-chain restructuring.

Recently, advanced technologies have been used to make deep supply chains more transparent (e.g. blockchain, molecular tagging in cotton supply chains), but this approach may not be locally suitable. Other methods include supplier code of conduct, monitoring the workplaces (auditing) of direct suppliers, publishing of supplier details for information transparency and industry alliances. However, the reputation of segments of the public remains almost the only incentive to comply, limiting full-scale adoption.

Companies are being pressured to take stances on these issues, as well as some social issues ('cancel culture': more controversial, less universally accepted). Companies must decide whether they want to risk alienating certain demographics with their choices.

5.7.16. Important Companies in the Semiconductor Industry

- **Arm** (UK): designs chips and licences the design as IP to manufacturers e.g. Nvidia, Intel, TSMC, Apple, Samsung.
- **Applied Materials** (USA): supplies equipment, services and software for manufacturing chips.
- **Nvidia** (USA): fabless tech company producing GPUs widely used to train large-scale machine learning models.
- **ASML** (Netherlands): the only producer of ‘extreme UV lithography’ machines, used to print ICs on silicon for high-performance chips. ASML was formed from the unification of several industry experts, and they sells their machines to chip makers, used to surpass the ‘7 nm’ process node from 2019. Its subsidiary, Cymer, is based in the US.
- **TSMC** (Taiwan): the world’s largest contract semiconductor manufacturer (foundry), selling to ‘fabless’ companies who rely on TSMC to produce their chips. TSMC does not try to compete with its customers. TSMC’s largest customer is Apple.
- **FoxConn** (China and Taiwan): a competitor to TSMC, which has much friendlier relations with the Chinese government, used to a smaller extent by Western companies.
- **Samsung** (South Korea): ‘chaebol’ (Korean conglomerate) with in-house manufacturing for their own electronics products, especially for memory chips.
- **Intel** (USA): manufactures its own high-end microprocessors and GPUs.

5.7.2. Income Tax (UK, FY2022):

Tax brackets vary over time and between different countries, as well as various special circumstances. The table is only useful as an example and should not be relied on.

The annual tax brackets for the fiscal year 2022 (6th April 2022 – 5th April 2023) in the UK using tax code 1257L (single source of fully taxable income) are:

Bracket	Taxable range	Progressive tax rate
Personal allowance	£ 0 – £ P	0%
Basic rate	£ $P+1$ – £ 50,270	20%
Higher rate	£ 50,271 – £ 150,000	40%
Additional rate	£ 150,001 and above	45%


The standard personal allowance $P = 12,570$, which holds when income is less than £ 100,000. If the income I is more than £ 100,000, then $P = \max\{12,570 - 0.5(I - 100,000), 0\}$.

Example calculation: if income before tax is £ 61,000, then the tax payable is

$$(12,570 - 0) * 0.00 + (50,270 - 12,571) * 0.20 + (61,000 - 50,271) * 0.40 = \mathbf{\pounds 11,831.40}.$$

Effective tax rate = (tax paid) / (income before tax) ($\times 100\%$).

5.7.3. Supply and Demand



In a perfectly competitive market, per-unit price of a good varies until quantity demanded equals quantity supplied (economic equilibrium).

(P : price of product, Q : quantity of product sold, S : supply curve, D : demand curve)

Isoelastic supply curve: $Q = kP^n$ (power law expression)

The graph on the left shows that an increase in demand ($D_1 \rightarrow D_2$) results in an increase in both price P and quantity Q .

Elasticity: sensitivity of demand to variations in environmental factors.

Liquidity: ease of buying/selling quickly and without affecting the price.

5.7.4. Inflation

Inflation is the reduction in value of money and the resulting decrease in consumer purchasing power.

Hyperinflation is caused primarily by excessive unsustainable growth in the money supply. Keynesian economics does not suggest moderate inflation directly results from growth, but rather that it is caused by excessive demand.

Real (inflation adjusted) rate of return = $1 + \frac{1 + \text{nominal rate}}{1 + \text{inflation rate}}$ (all expressed as decimals)

5.7.5. Pareto Principle

The Pareto principle (aka the 80-20 rule) is an empirical observation applicable to some scenarios. It states that 80% of consequences (data) come from 20% of the causes.

It is often used qualitatively, where the goal is to identify the dominant few actions (or problems to solve) that would generate the most results (profits).

5.7.6. Inventory Models

In operations research, an inventory model advises on the time and quantity of a supply to purchase in anticipation of a distribution of orders.

Let p : unit sale price, c : unit order cost, h : unit leftover holding cost, b : unit shortage penalty, D : random demand, x : initial inventory level, y : base stock level, $F_D(d) = P(D \leq d)$.

Assuming that $p + b > c$, the optimal order quantity q^* aims to maximise the expected profits:

$$q^* = \operatorname{argmax}_q \mathbb{E} \left[\underbrace{-cq}_{\text{order cost}} + \underbrace{p \min\{x + q, D\}}_{\text{products sold}} - \underbrace{h \max\{x + q - D, 0\}}_{\text{holding costs}} - \underbrace{b \max\{D - x - q, 0\}}_{\text{shortage penalties}} \right]$$

The solution (optimal order up to level) is $y^* = F_D^{-1}\left(\frac{p+b-c}{p+b-h}\right)$, where $y^* = \max\{q^*, 0\} + x$
 $\leftrightarrow q^* = \max\{y^* - x, 0\}$ and F_D^{-1} is the inverse cumulative distribution function (ICDF) of D .

Probability of running out of stock = $P(D > y^*) = 1 - \frac{p+b-c}{p+b-h}$.

If $D \sim N(\mu, \sigma^2)$ then $y^* = \mu + \sigma \times \Phi^{-1}\left(\frac{p+b-c}{p+b-h}\right)$ and the optimal (maximum) profit is $cx + (p-c)\mu - \sigma((h+c)z^* + (p+h+b)L(z^*))$ where $z^* = (y^* - \mu) / \sigma = \Phi^{-1}\left(\frac{p+b-c}{p+b-h}\right)$ and

$L(w)$ is the loss function defined as $L(w) = \frac{1}{\sqrt{2\pi}} \int_w^\infty (t-w) e^{-\frac{1}{2}t^2} dt = \phi(w) - w(1 - \Phi(w))$.

For an additional fixed (base) order cost K , the optimal reorder point s is defined as the value such that when $x < s$ we should order up to level ($q = q^*$) and when $x > s$ we do not order. The value of s is the smallest s such that $\text{Profit}(y^*) - \text{Profit}(s) = K$, which can be solved using the profit expression above.

5.7.7. Financial Instruments

Stocks (equities, shares): represent a fixed fraction of a single company's market capitalisation (value).

Bonds (fixed income securities): debt instruments with fixed interest rate returns and time to maturity (repay time to avoid defaulting), issued by governments (gilts / treasury notes) or corporations.

Commodities: raw materials such as fuels (e.g. crude oil, natural gas), agricultural produce (e.g. corn, sugar, live cattle), base metals (e.g. lead, copper), precious metals (e.g. gold, platinum), precious stones (e.g. diamond), lumber, rubber and water rights.

Currencies: the exchange rates of one currency relative to another, including cryptocurrencies, fluctuating due to national economics.

Derivatives: high-leverage instruments based on an underlying asset, used to hedge risk. May be contracts to purchase assets for a fixed price in the future (futures, options, forwards), or by exchanging loans with different interest rates (swaps).

Exchange-Traded Funds (ETFs): a collection of assets, whose price varies throughout the day as trading occurs.

5.7.7. Time Series Forecasting

A stock ticker represents time series data of the value of a stock over time. Various statistical and machine learning methods (e.g. SARIMAX, ES, LSTM, transformer etc) can be used to estimate (forecast) \hat{X}_{n+1} given $X_1 \dots X_n$.

The return of a stock $(\frac{X_n}{X_{n-1}} - 1)$ is typically more stationary (WSS) than the values X_n , and these returns are closer to a Normal or t -distribution, making use of standard scaling (Section 5.5.7) optimal in preprocessing.

Backtesting is used to test a trained model. For a large span of historic data \mathbf{b} , split \mathbf{b} into a set of adjacent sliding windows $\mathbf{X} = [\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots]$ and process each $\mathbf{x}^{(i)}$,

Processing involves checking data suitability (e.g. checking for stationarity in ARMA), computing the forecast, and evaluating the performance using a metric.

Stock tickers which perform well in backtesting can be ranked, with the top selection being used in real-time trading (diversification helps smooth out random variation).

White noise hypothesis: there is no actual pattern to stock data, and any model is as good as drawing from a Normal distribution with the same mean and variance as the training data. This is a good reference to see whether a model outperforms the 'naive' estimate.

~ LN, 2024