Course notes

GRE Subject Test · GRE Mathematics

Gre Mathematics

Computer-delivered undergraduate GRE Mathematics

Approximately 66 multiple-choice items in 170 minutes; no separately timed sections. Calculus about 50%, algebra 25%, additional topics 25%. This is graduate-admissions testing of undergraduate mathematics, not an A-level qualification.

Official scope: https://www.ets.org/gre/test-takers/subject-tests/about/content-structure.html

C.1 · Single-variable calculus and applications

- A limit describes values near a point, without requiring the function to be defined there. Continuity adds the requirement that the function value exists and agrees with the limit. Algebraic cancellation is valid only away from the cancelled zero, but can reveal a removable limit. One-sided limits must agree for a two-sided limit. For quotient limits, check the denominator and hypotheses before applying a rule; 0/0 is an indeterminate form, not an answer.
- A derivative is a limit of difference quotients and implies continuity; the converse fails, as |x| at zero shows. Product, quotient and chain rules describe different structures: the derivative of f(g(x)) is $f'(g(x))g'(x)$, not a product of unrelated values. Differentiating a composition a second time generally produces two terms. Check domains, nonzero denominators and differentiability assumptions before using a symbolic expression as a derivative.
- The fundamental theorem says that the derivative of $\int$ from a to x of a continuous integrand f(t) is f(x). With a variable upper bound h(x), multiply by $h'(x)$; with two variable bounds, subtract the corresponding lower-bound contribution. A definite integral also arises as a limit of Riemann sums. Rewrite the sum as $(1/n)\Sigma f(k/n)$ before identifying the integral on [0,1], rather than treating n-dependent terms as constants.
- For example Σ from k=1 to n of $n/(n^2+k^2)$ equals $(1/n)\Sigma 1/(1+(k/n)^2)$, tending to $\int_0^1 1/(1+t^2)dt=\pi/4$. This limit uses continuity and the partition width 1/n. A series over an unbounded number of terms is a different limit: terms tending to zero do not alone guarantee convergence. For $\Sigma 1/k$ the partial sums diverge; compare, estimate or apply a valid series test instead of using the necessary term condition as sufficient.

exercises.lol exercises.com.cn

■ Checked example

For F(x)=integral from 0 to x^2 of e^t dt, the fundamental theorem and chain rule give F′(x)=e^(x^2) · 2x. At x=1, F′(1)=2e. The upper limit is x^2, so omitting 2x misses its rate of change.

■ Common error

A-level calculus is useful prerequisite material but does not cover the undergraduate analysis and applications tested here.

C.2 · Multivariable calculus and vector analysis

- A partial derivative changes one coordinate while fixing the others. For f=x^2+3xy, f_x=2x+3y and f_y=3x. The gradient collects these derivatives; if f is differentiable, the directional derivative along a unit vector v is grad f · v. Normalise the direction before taking this dot product. A direction vector of length two would double the answer if used without normalisation.
- Differentiability means a valid linear approximation, not merely the existence of some partial derivatives at one point. Continuous first partials in a neighbourhood are a sufficient condition. For a composition f(x(t),y(t)), the chain rule gives f_x x′+f_y y′. Second derivatives can introduce both direct and mixed terms; a zero mixed partial alone places no restriction on the pure second partials.
- A multiple integral sums contributions over a specified region. Describe the region before choosing iterated limits; nonrectangular bounds may change when the order changes. In polar coordinates area is r dr dtheta, not just dr dtheta. A general coordinate change uses the absolute Jacobian determinant. Sign belongs to oriented vector quantities, while area and volume scaling use a nonnegative factor.
- If both first partials of a globally defined function on R^2 are constant, f_x=a and f_y=b, integrating successively gives f=ax+by+c, a plane. This cannot be inferred from parallel straight level sets alone: e^x has vertical parallel level lines but a curved graph. Nor do f_xy=f_yx=0 force a plane: x^2+y^2 is a counterexample. Distinguish first-derivative constancy from absent mixed dependence and from the shape of selected level sets.

■ Checked example

For f(x,y)=x^2+3y^2, gradient f=(2x,6y). At (1,1) this is (2,6). In direction (3,4), the unit vector is (3/5,4/5), giving directional derivative 2 · 3/5+6 · 4/5=6.

■ Common error

A non-unit direction vector gives a scaled directional rate, not the derivative per unit distance.

exercises.lol exercises.com.cn

A.1 · Linear algebra

- A real vector space is closed under its specified addition and scalar multiplication and satisfies the vector-space axioms. A basis is an independent spanning list: every vector has a unique coordinate representation in that list. To test independence of matrix columns, solve the homogeneous system Ac=0; only the zero coefficient vector means independence. A spanning set can contain redundant vectors, so spanning alone does not make a basis.
- Row reduction exposes pivots and free variables without changing the solution set of a linear system. Rank is the dimension of the column image, equivalently the number of pivots. Nullity is the kernel dimension: for a map from an n-dimensional domain, rank+nullity=n. In a nonhomogeneous system, a zero coefficient row with nonzero right side means inconsistency; free variables give infinitely many solutions only after consistency has been established. The codomain dimension need not equal rank.
- A square matrix is invertible exactly when its determinant is nonzero, its kernel is zero and its rank equals its size. These statements do not require each entry to be nonzero. Triangular determinants are products of diagonal entries, so a matrix depending on a complex variable can be singular at complex roots absent from a real-only calculation. Row swaps reverse determinant sign; adding a multiple of one row to another leaves it unchanged. Check singularity before applying an inverse formula.
- An eigenvector is nonzero and satisfies $Av=\lambda v$, so eigenvalues are roots of $\det(A-\lambda I)$. Diagonalisation requires a full independent eigenvector basis over the chosen field. Distinct eigenvalues give independent eigenvectors, while repeated eigenvalues may have too small an eigenspace. A real symmetric matrix has a real orthonormal eigenbasis. A real odd-dimensional matrix has at least one real eigenvalue because its real characteristic polynomial has odd degree; that alone does not imply diagonalisation or all eigenvalues real.

■ Checked example

A=[[1,1],[0,1]] has characteristic polynomial $(1-\lambda)^2$. Eigenvectors satisfy y=0, so its eigenspace has dimension one. Two independent eigenvectors are needed to diagonalise a 2×2 matrix; A is not diagonalizable.

■ Common error

The algebraic multiplicity of an eigenvalue is not automatically the dimension of its eigenspace.

exercises.lol exercises.com.cn

A.2 · Abstract algebra and number theory

- A group needs closure, associativity, identity and inverses. Commutativity is an additional condition.
- A homomorphism preserves the operation. Its kernel is a normal subgroup and identifies elements mapping to the identity.
- A field permits division by every nonzero element. Integers form a ring but not a field.
- Work with congruences modulo n. A residue a has a multiplicative inverse exactly when gcd(a,n)=1.

Checked example

Modulo 8, 3 has inverse 3 because $3 \cdot 3 = 9 \equiv 1$. But 2 has no inverse because gcd(2,8)=2. Modulo a prime, every nonzero residue has an inverse; this gives a finite field.

Common error

Do not cancel a factor in a modular equation without checking that it is invertible.

T.1 · Real analysis and topology

- An epsilon–delta statement controls all sufficiently close inputs. The quantifier order matters.
- In real Euclidean space, closed and bounded sets are compact. Do not apply this equivalence to every metric space.
- Continuous images of compact sets are compact, so a real continuous function on a compact domain attains extrema.
- Connectedness rules out a separation into disjoint nonempty open parts. Continuity preserves connectedness; completeness is a separate property.

Checked example

f(x)=x on (0,1) is continuous and bounded but never equals its supremum 1. On [0,1], the same function attains its maximum at 1. The missing endpoint explains why the compact-domain theorem does not apply to the first case.

Common error

A theorem's conclusion cannot be used before its hypotheses have been checked.

exercises.lol exercises.com.cn

T.2 · Discrete mathematics, probability and numerical methods

- Choose the counting model first: ordered selections of r distinct objects use $n(n-1)\cdots(n-r+1)$, while unordered subsets use $C(n,r)$. A complete simple graph on n vertices has one edge per unordered vertex pair, giving $n(n-1)/2$ edges; loops and multiple edges would change the model. For n lines in general position in a plane, the kth line crosses the prior $k-1$ lines in distinct points and adds k regions. The total is $1+n(n+1)/2$; parallels or triple concurrence invalidate that count.
- Use complements to count at least one occurrence, and condition on the actual remaining population after a draw without replacement. For r iid outcomes chosen from n equally likely values, the probability that all are distinct is $n(n-1)\cdots(n-r+1)/n^r$ when $r\le n$. Its complement counts repeated outcomes. For a binomial count with N independent trials and fixed success probability p, mean is Np and variance $Np(1-p)$. Identical probabilities alone do not establish independence.
- A recurrence describes later values from earlier ones and needs enough initial data to determine a sequence. Separate the index from the value: $a_n=2a_{(n-1)}$ with $a_0=3$ gives $a_n=3\cdot 2^n$. Graph and algorithm arguments often establish a recurrence by identifying what a new vertex or step adds. A closed formula should satisfy both the recurrence and its initial conditions; fitting a few observed terms does not prove it for every index.
- Numerical approximation needs an error argument. Bisection preserves a sign-changing bracket for a continuous function and halves its width at each step; a zero may be absent if continuity fails. Newton's update is $x_{new}=x-f(x)/f'(x)$, requiring a nonzero derivative at the current point; convergence is not automatic from every starting value. An approximation's residual and its error in x are different. State the method's assumptions and a stopping criterion, rather than treating extra displayed decimals as accuracy.

■ Checked example

For $f(x)=x^2-2$ and $x_0=1$, Newton's step gives $1-(1-2)/2=1.5$. The next value is $1.5-0.25/3=1.4167$. These approximate sqrt(2), but $f'(0)=0$ makes zero an invalid starting point.

■ Common error

Events with positive probabilities cannot be both independent and mutually exclusive.

exercises.lol exercises.com.cn

T.3 · Complex analysis and residues

- Write f(z)=u(x,y)+iv(x,y). Complex differentiability imposes u_x=v_y and u_y=−v_x; with continuous first partials locally, the Cauchy–Riemann equations establish analyticity there. They differ from real differentiability of a two-coordinate map. The conjugate function x−iy fails these equations on every open neighbourhood, although real partial derivatives exist. State the region being checked, not only a convenient point.
- An analytic function is complex differentiable throughout a neighbourhood and has a local convergent power series. A removable singularity can be filled analytically when the function is bounded near the missing point. A pole has a finite principal part in its Laurent series; an essential singularity has infinitely many negative-power terms. The residue is the coefficient of $(z-a)^{-1}$, not necessarily the leading or largest negative-power term.
- For an isolated pole of order m, write f(z)=g(z)/(z−a)^m with g analytic at a. The residue is g^(m−1)(a)/(m−1)!. A simple pole uses g(a); a double pole uses $g'(a)$. This follows by expanding g into its Taylor series. Thus e^z/(z−a)^2 has residue e^a, while a constant numerator over a pure double pole has zero residue. The pole order alone does not determine the contour integral.
- For a positively oriented contour enclosing isolated singularities, the residue theorem gives $\oint f(z)dz = 2\pi i$ times the sum of enclosed residues, provided the function is analytic on the contour and elsewhere in the required interior. Reversing orientation changes the sign. Cauchy's derivative formula is $\oint g(z)/(z-a)^{(m+1)}dz = 2\pi i\, g^{(m)}(a)/m!$ under its analytic-domain hypotheses. A singularity on the contour prevents direct application; a singularity outside contributes nothing to this contour.

■ Checked example

For f(z)=conjugate(z), u=x and v=−y. Then u_x=1 but v_y=−1, so f is not complex differentiable. For 1/(z−2), a positively oriented circle |z−2|=1 encloses a simple pole of residue 1, hence the integral is $2\pi i$.

■ Common error

Continuity or real differentiability alone does not imply complex analyticity.

C.3 · Differential equations and initial conditions

- For $y'=ky$, separation and integration give y=Ce^(kt); recover the zero solution if division by y was used. An initial value determines C.
- A linear second-order constant-coefficient equation uses a characteristic polynomial. Distinct roots yield exponentials; a repeated root requires (C1+C2t)e^(rt).
- Complex roots $\alpha \pm i\beta$ give e^(αt)(C1 cos βt+C2 sin βt). The real motion includes both amplitude and phase information.

exercises.lol exercises.com.cn

- Existence and uniqueness depend on hypotheses near the initial point. A singular coefficient or a failure of local Lipschitz behaviour can defeat the familiar uniqueness conclusion.

■ Checked example

Solve y″+4y=0 with y(0)=3 and y′(0)=4. The characteristic roots are ±2i, so y=A cos2t+B sin2t. The first condition gives A=3; differentiating gives y′(0)=2B=4, so B=2.

■ Common error

A repeated characteristic root needs the factor t; two copies of the same exponential are not independent solutions.

A.3 · Groups, cosets and quotient maps

- Before computing an order, check closure, associativity, an identity and an inverse for every element under the stated operation. A subset can inherit associativity yet fail closure or omit the identity. In a finite group, the order of an element is the least positive power giving the identity. In the additive group Z/nZ, it is the least positive multiple giving zero; the order of residue a is n/gcd(a,n). Lagrange's theorem says subgroup orders divide the group order. The converse is not a general existence theorem, and the order of a group is not the order of each element.
- A left coset gH is a translate of a subgroup H. Cosets have equal size and partition the group, so the index is |G|/|H| in a finite group. In an additive group write g+H. Membership in the same coset means the difference lies in H. A coset usually is not itself a subgroup because it may omit the identity.
- A homomorphism preserves the operation. Its kernel consists of elements sent to the identity, and its image consists of values actually reached. Every kernel is normal. The first isomorphism theorem identifies G/ker(phi) with im(phi); do not replace the image with the whole codomain unless the map is onto.
- Quotient multiplication is well defined only when H is normal. All subgroups of an abelian group are normal. In a nonabelian group test gHg^−1=H; a subgroup of index two is normal. For permutations compose in the stated convention, here rightmost first. Disjoint cycle lengths give the permutation order by their least common multiple. Conjugation $h\sigma h^{-1}$ relabels the elements in σ's cycles, so it preserves cycle lengths; conversely permutations with the same cycle lengths can be related by a relabelling. Thus conjugacy classes in S_n correspond to partitions of n, including fixed-point cycles. In S4 the types are 1+1+1+1, 2+1+1, 2+2, 3+1 and 4: five classes, not one class for each possible element order. The types 2+1+1 and 2+2 both have order two but are not conjugate.

exercises.lol exercises.com.cn

■ Checked example

Define phi from Z/12Z to Z/3Z by reducing residues modulo 3. It is onto and preserves addition. Its kernel H is {0,3,6,9}, so |H|=4 and the index is 12/4=3. The other cosets are {1,4,7,10} and {2,5,8,11}. Thus (Z/12Z)/H is isomorphic to Z/3Z. The element 3 in the original group has order 12/gcd(3,12)=4, not 3.

■ Common error

A quotient has one element per coset, not one per element of its kernel. A homomorphism need not be onto its stated codomain.

A.4 · Rings, ideals and modules

- For the domain, ideal and quotient examples below, use a commutative ring with identity 1 distinct from 0. A source problem may specify a general ring instead; do not assume its multiplication commutes unless stated or proved. A unit has a multiplicative inverse. A nonzero zero divisor multiplies some nonzero element to zero. An integral domain has no such zero divisors; cancellation of a nonzero factor then works. A field is a domain in which every nonzero element is a unit. Z is a domain but not a field; Z/6Z is neither.
- An ideal I is an additive subgroup that absorbs multiplication by every ring element. This is stronger than being a subring. In Z, nZ is an ideal; quotient elements are integer residue classes modulo n. In a commutative ring, R/I is a field exactly when I is maximal, and it is a domain exactly when I is prime. The ideal must be proper in both statements.
- A module allows scalars from a ring instead of requiring a field. Every abelian group is a Z-module by repeated addition, but it need not have a vector-space basis. In Z/6Z as a Z-module, 6 times the nonzero residue 1 is zero; this is torsion. For a vector space over a field, a nonzero scalar is invertible and cannot annihilate a nonzero vector.
- A submodule is closed under addition and all permitted scalar actions. A linear map of modules preserves both. The kernel and image are submodules, and the quotient by the kernel is isomorphic to the image. Do not apply finite-dimensional rank-nullity to an arbitrary module without establishing an appropriate free-module setting; integer row operations and field row operations permit different divisions. In a general Boolean ring, every a satisfies a^2=a. Do not assume commutativity to prove it: idempotence of a+a gives 4a=2a, hence 2a=0. Expanding $(a+b)^2$=a+b gives ab+ba=0, and characteristic two makes −ba=ba; therefore ab=ba. Idempotence does not imply nilpotence: in F2, the nonzero element 1 satisfies 1^n=1 for every positive n.

exercises.lol exercises.com.cn

■ Checked example

In Z/6Z, 2 · 3=0, so 2 and 3 are zero divisors. The units are 1 and 5 because their gcd with 6 is 1. In Z, the ideal 5Z gives the field Z/5Z, while 6Z gives a quotient with zero divisors. As Z-modules, the map from Z to Z/6Z has kernel 6Z; the quotient identifies integers differing by a multiple of 6, rather than producing a real vector space.

■ Common error

A set can be a submodule without being a vector space over the rationals. Scalar division is valid only when the scalar inverse belongs to the structure.

A.5 · Polynomials and field extensions

- A polynomial is irreducible over a field if it has positive degree and no factorisation into polynomials of smaller positive degrees there. A quadratic or cubic is irreducible exactly when it has no root in that field. This root test alone fails for degree four or higher; for instance $(x^2+1)(x^2+2)$ has no real root but is reducible over R. Always name the base field.
- The quotient F[x]/(p) is a field when p is irreducible. Reduce powers using p(alpha)=0, where alpha is the residue class of x. If F has q elements and p has degree d, the quotient has q^d elements represented by polynomials of degree below d. Z/4Z has four elements but has zero divisors, so it is not the field with four elements.
- Over F2, $p(x)=x^2+x+1$ has values 1 at both 0 and 1 and is irreducible. In its quotient $\text{alpha}^2=\text{alpha}+1$ because subtraction equals addition in characteristic two. The four elements are 0,1,alpha,alpha+1. All three nonzero elements must be units; compute their products rather than treating alpha as an ordinary real number.
- For nested finite-degree fields K inside L inside M, the tower law gives [M:K]=[M:L][L:K]. The degree of an algebraic element is the degree of its minimal polynomial. A finite extension of degree two does not contain an element of degree three over the base field. Over Q, sqrt(2) has degree two; adjoining sqrt(3) as well produces a degree-four extension, since sqrt(3) is not in Q(sqrt(2)). Over C, primitive nth roots of unity have exact order n and are the roots of the cyclotomic polynomial Φ_n. For n=10, divide x^5+1 by x+1 to exclude the order-two root −1: $\Phi_{10}=x^4-x^3+x^2-x+1$. Vieta's formulas give sum 1 and product 1 of its four primitive roots. Do not sum all tenth roots, or assume every nontrivial tenth root is primitive; a root's order must be checked. For a monic degree d polynomial, product of its roots is (−1)^d times the constant coefficient.

exercises.lol exercises.com.cn

■ Checked example

In F2[alpha] with alpha^2+alpha+1=0, multiply alpha(alpha+1) = alpha^2 + alpha = (alpha+1) + alpha = 1. Thus alpha^−1=alpha+1. Also alpha^3=1 and alpha is not 1, so its multiplicative order is 3. This produces a field of four elements. By contrast 2 · 2=0 in Z/4Z, proving that quotient is not a field.

■ Common error

Polynomial reducibility changes with the coefficient field. The polynomial x^2−2 is irreducible over Q but splits over R.

A.6 · Congruences, divisibility and arithmetic functions

- The equation ax≡b modulo n has a solution exactly when d=gcd(a,n) divides b. If it does, divide a,b,n by d to obtain an equation with an invertible coefficient. It has one residue solution modulo n/d and d distinct solutions modulo n. Dividing the coefficient but leaving the original modulus generally loses solutions. Prime factorisation gives another divisibility tool: if n^k must be divisible by a product of prime powers p^a, each prime exponent in n must be at least ceil(a/k). These minimum exponents independently produce the least positive admissible n. For n^4 divisible by $2^7 \cdot 3^5$, n must contain 2^2 and 3^2, so the least value is 36. This is a prime-exponent condition, not a congruence solved by modular division.
- The extended Euclidean algorithm expresses gcd(a,n) as ua+vn. When the gcd is 1, u is an inverse of a modulo n. Choose a representative in the required range after reduction. For 7 and 26, $1=15 \cdot 7-4 \cdot 26$, so 15 is the inverse of 7 modulo 26; verify the product to catch a sign error.
- For coprime positive moduli m,n, the Chinese remainder theorem gives exactly one solution modulo mn for each pair of residue conditions. Substitute x=r+mk into the second congruence and solve for k. If the moduli are not coprime, their residue values must agree modulo gcd(m,n); when consistent, uniqueness is modulo the least common multiple, not the product.
- Euler's totient phi(n) counts residues coprime to n. For distinct prime divisors p, phi(n)=n times the product of (1−1/p). Euler's theorem gives a^phi(n)≡1 only when gcd(a,n)=1. The prime case is Fermat's little theorem. Reduce exponents only after checking this condition; a nonunit can become zero under repeated powers instead.

■ Checked example

Solve 6x≡9 modulo 15. The gcd is 3 and divides 9. Divide all three quantities to get 2x≡3 modulo 5; the inverse of 2 is 3, so x≡9≡4 modulo 5. The original solutions modulo 15 are 4,9,14. For x≡2 modulo 4 and x≡3 modulo 7, write x=2+4k; then 4k≡1 modulo 7, giving k≡2. Hence x≡10 modulo 28.

exercises.lol exercises.com.cn

■ **Common error**

The expression a^phi(n)≡1 is false for arbitrary a. For example 2 is not invertible modulo 8 and 2^4 is zero modulo 8.

T.4 · Sequences, series and uniform convergence

- A convergent real sequence is Cauchy, and every real Cauchy sequence converges because R is complete. In Q a Cauchy sequence can approach an irrational number and fail to converge within Q. A bounded monotone real sequence converges. A bounded sequence need not converge, but Bolzano–Weierstrass guarantees a convergent subsequence; (-1)^n has two different subsequential limits.
- For a numerical series, absolute convergence implies convergence; conditional convergence does not allow arbitrary rearrangement without affecting the sum. The alternating harmonic series converges but its absolute-value series diverges. In a ratio test, a limit below 1 proves absolute convergence and one above 1 proves divergence; a limit equal to 1 is inconclusive, as both sum 1/n and sum $1/n^2$ demonstrate.
- Pointwise convergence chooses an index N separately for each x and tolerance. Uniform convergence chooses one N that works for every x in the domain. For real-valued functions, check the supremum of |f_n−f| over the whole domain. A continuous pointwise limit does not by itself prove uniform convergence. Domain endpoints and shrinking peaks often distinguish the two notions.
- A uniform limit of continuous functions is continuous. On a closed bounded interval, uniform convergence of Riemann-integrable functions permits exchanging limit and integral. Exchanging derivatives needs extra hypotheses; uniform convergence of the functions alone is insufficient. The Weierstrass M-test establishes uniform absolute convergence of a series of functions if each term is bounded by M_n on the whole domain and sum M_n converges.

■ **Checked example**

On [0,1], f_n(x)=x^n tends to 0 for x<1 and to 1 at x=1. The limit is discontinuous, so convergence cannot be uniform; directly, the supremum error is 1, approached below 1. On [0,a] with 0≤a<1, the limit is zero and the supremum is a^n, which tends to zero, so convergence is uniform. Changing the domain changes the conclusion.

■ **Common error**

Checking several fixed x values proves neither a supremum bound nor uniform convergence. The point producing the largest error may move with n.

exercises.lol exercises.com.cn

T.5 · Open sets, compactness and connectedness

- In a metric space, an open set contains a small ball around each of its points. The interior consists of such points; the closure includes all limit points; the boundary is closure minus interior. Open and closed are not mutually exclusive labels; the empty set and the whole space are both. In R, the set [0,1) has interior (0,1), closure [0,1] and boundary {0,1}.
- In a subspace X, an open set has the form X intersected with an ambient open set. Thus [0,1) is open relative to [0,2], using intersection with (-1,1), although it is not open in R. A set may also be relatively closed without being closed in the ambient space. Always state which space defines neighbourhoods and which metric is used.
- Compactness means every open cover has a finite subcover. In Euclidean R^n, Heine–Borel makes this equivalent to closed and bounded. In a metric space, compactness is equivalent to sequential compactness; completeness and boundedness alone do not suffice in arbitrary metric spaces. A continuous image of a compact set is compact, giving attained maxima and minima for real continuous functions on a nonempty compact domain.
- Connected sets cannot be separated into two disjoint nonempty relatively open parts; connected subsets of R are precisely intervals. A continuous image of a connected set is connected, which yields the intermediate value theorem. Path connectedness implies connectedness, but not conversely in every space. A compact set need not be connected, and a connected set need not be compact; a finite two-point set and an open interval supply the contrasting cases.

■ Checked example

Let X=[0,1] with its usual relative topology. The set U=[0,0.5) equals X intersected with (-1,0.5), so U is open in X. Its closure in X is [0,0.5], and its boundary in X is {0.5}; zero is an interior point relative to X. For a continuous f on X with f(0)<0<f(1), connectedness ensures a zero, while compactness separately ensures attained extrema.

■ Common error

Heine–Borel's closed-and-bounded test is a Euclidean-space theorem. Connectedness and compactness answer different questions.

C.4 · Taylor expansions and power-series endpoints

- The Taylor polynomial of degree m at a is the sum of f^(k)(a)(x−a)^k/k! for k from zero to m. The factorial belongs to each coefficient. If the next derivative is bounded in magnitude by M between a and x, the Lagrange remainder has magnitude at most M|x−a|^(m+1)/(m+1)!. Smoothness alone does not guarantee that the infinite Taylor series equals the function everywhere.
- A power series sum c_n(x−a)^n converges absolutely inside its radius R and di-

exercises.lol exercises.com.cn

verges outside it. Ratio or root tests usually determine R, with possible values zero and infinity. At x=a−R and x=a+R the test often becomes inconclusive; substitute each endpoint into the original series. The two endpoint behaviours may differ.

- Within the open interval of convergence, termwise differentiation and integration preserve the radius. They can change whether endpoints are included. Start from the geometric series 1/(1−x)=sum x^n for |x|<1, then integrate from zero to x to obtain −ln(1−x)=sum x^n/n for n≥1. Check the integration constant and the real logarithm domain.
- Series also resolve removable limit forms. To evaluate (e^x−1−x)/x^2 near zero, retain the first surviving term x^2/2 rather than using only e^x≈1+x. An asymptotic truncation establishes the limit; a finite-interval inequality needs a separate remainder sign or magnitude argument. Do not substitute into a series outside its convergence interval.

■ Checked example

For sum x^n/n with n≥1, the ratio test gives radius 1. At x=1 it is the divergent harmonic series; at x=−1 it is an alternating convergent series. Thus its real convergence interval is [−1,1). Differentiating inside gives sum x^(n−1)=1/(1−x), which converges at neither endpoint. The radius stayed 1 while the endpoint inclusion changed.

■ Common error

A radius is not a complete interval of convergence. Endpoint tests and factorial coefficients must be checked explicitly.

C.5 · Improper integrals and geometric applications

- An infinite integration endpoint is replaced by a finite bound and a limit. At a singular point inside the interval, split the integral and require both one-sided integrals to converge separately. Symmetric cancellation can define a Cauchy principal value but does not prove convergence of the ordinary improper integral. For 1/x across zero, the two sides diverge even though symmetric cutoffs cancel.
- The integral of x^(−p) from 1 to infinity converges exactly when p>1. From zero to 1 it converges exactly when p<1. The same exponent behaves differently at the two boundaries. For positive integrands, comparison transfers convergence from a larger integrable function or divergence from a smaller nonintegrable one; keep the inequality direction correct. For an integral of a maximum or minimum, solve branch crossings and check which expression dominates each interval. For max(√(1−x^2),x+1) on [−1,1], the interior crossing is zero; use the semicircle on [−1,0] and the line on [0,1]. The area is π/4+3/2, not the integral of either branch over the entire interval.

exercises.lol exercises.com.cn

- For rotation around the x-axis, disks or washers integrate $\pi(R^2-r^2)$ dx. Cylindrical shells use 2π times radius times height and integrate in the matching variable. Select the method by the geometry and verify nonnegative radii. Area between curves integrates upper minus lower, splitting where their order changes. A signed integral is not automatically geometric area.
- For a differentiable plane curve y=f(x), arc length is the integral of sqrt(1+(f′(x))^2) dx. A surface formed by rotating a nonnegative f around the x-axis has area integral 2πf sqrt(1+(f′)^2) dx. These are different quantities from volume. If an interval is unbounded, the geometric formula still needs a convergence test.

■ Checked example

Rotate y=1/x for x≥1 around the x-axis. The disk volume is π times the integral of $1/x^2$, so V=π. The lateral surface integrand is $2\pi(1/x)$sqrt(1+$1/x^4$), at least $2\pi/x$; its improper integral diverges. Finite volume therefore does not imply finite surface area. This comparison avoids trying to find an unnecessary antiderivative.

■ Common error

Never count a principal value as a convergent improper integral, or infer one geometric quantity's finiteness from another.

C.6 · Multivariable extrema and constrained optimisation

- For a differentiable function on an open domain, an interior local extremum has zero gradient. This condition is necessary, not sufficient. For two variables at a critical point, let D=f_xx f_yy−(f_xy)^2. If D>0, f_xx>0 gives a strict local minimum and f_xx<0 gives a strict local maximum; D<0 gives a saddle. These tests require suitable second-derivative regularity near the point.
- When D=0 the second-derivative test is inconclusive, not evidence of a saddle. The functions x^4+y^4 and x^4-y^4 have the same zero Hessian at the origin, but one has a strict minimum and the other changes sign. Evaluate along contrasting directions or use a direct inequality. For a quadratic form, positive definiteness gives a more general Hessian interpretation.
- On a smooth equality constraint g(x,y)=c with nonzero gradient g, solve gradient f=lambda gradient g together with g=c. These equations generate candidates; they do not classify or guarantee a global extremum. If the constraint gradient vanishes, the regularity condition fails and the multiplier equations can miss a constrained extremum. Treat such points separately. For x^2+y^2=a and xy=b>0, $(x-y)^2$=a−2b shows necessity a≥2b. It is also sufficient: take s=√(a+2b), d=√(a−2b), x=(s+d)/2 and y=(s−d)/2. These real values have xy=b and x^2+y^2=a. A minimum condition alone needs this attainment check to establish solvability.
- A continuous function on a compact feasible set attains global extrema. To find them, compare all interior candidates and all boundary pieces, including corners

exercises.lol exercises.com.cn

and endpoints. For a rectangle, optimise the restrictions on each edge. For a disk, the circular boundary may use a parameter or multiplier. Solving only the unconstrained gradient ignores possible boundary winners.

Checked example

Minimise x^2+y^2 subject to $x+y=6$. The constraint gradient is $(1,1)$, so $2x=\lambda$ and $2y=\lambda$. Thus $x=y$ and the constraint gives $x=y=3$, with value 18. Completing the square on the line gives $x^2+(6-x)^2=2(x-3)^2+18$, proving a global minimum. There is no maximum on this unbounded line, despite the existence of a multiplier candidate.

Common error

A Lagrange multiplier solution is a candidate, not automatically a maximum or minimum. Check regularity and the complete feasible set.

C.7 · Coordinate changes and vector integral theorems

- Changing variables scales area or volume by the absolute Jacobian determinant. Polar coordinates use r dr dtheta; cylindrical coordinates use r dr dtheta dz; spherical coordinates with phi measured from the positive z-axis use $\rho^2 \sin(\varphi)\, d\rho\, d\varphi\, d\theta$. State angle conventions and transform both the integrand and the region. A missing Jacobian changes a uniform density integral into the wrong physical quantity.
- Green's theorem equates positively oriented planar boundary circulation integral $P\,dx+Q\,dy$ with the double integral of Q_x-P_y on the region. The usual hypotheses require first derivatives continuous on an open set containing the region. A hole needs its own negatively oriented inner boundary, or another valid treatment of the missing domain. The theorem cannot integrate across a field singularity.
- Stokes' theorem equates circulation on a surface boundary with the surface integral of curl F dot the oriented normal. The right-hand rule links boundary direction to the normal. The divergence theorem equates outward flux across a closed surface with the volume integral of div F. Circulation, flux, curl and divergence are distinct; a closed surface is required for the usual divergence theorem.
- A gradient field has path-independent line integrals, determined by endpoint potential differences. A continuously differentiable curl-free field on a simply connected open domain is conservative. Curl-free alone on a domain with a hole is insufficient. For $F=(-y/(x^2+y^2), x/(x^2+y^2))$, the origin is excluded; unit-circle circulation is 2π, although the curl is zero wherever the field is defined.

exercises.lol exercises.com.cn

■ Checked example

For F=(x,y,z), divergence is 3. The outward flux through the sphere of radius 2 is therefore 3 times its volume, or $3 \cdot (4\pi \cdot 2^3/3)=32\pi$. This avoids a surface parameterisation. For the planar field (−y,x), Green's theorem gives counterclockwise unit-circle circulation as integral of $1-(-1)=2$ over the disk, hence 2π. Reversing the orientation changes the circulation sign.

■ Common error

Zero curl is not enough when the domain has a hole. Outward flux and counterclockwise circulation use different theorems and orientation rules.

T.6 · Functions, inverse branches and composition

- A function assigns exactly one output to each allowed input. Injective means equal outputs force equal inputs; surjective means every element of the stated codomain is reached. The map $x\mapsto x^2$ from R to $[0,\infty)$ is surjective but not injective. From $[0,\infty)$ to $[0,\infty)$ it is both, hence bijective. From $[0,\infty)$ to R it remains injective but fails surjectivity. Keep domain, image and codomain separate when testing each claim.
- An inverse function reverses a bijection. Solve $y=f(x)$ for x, then use the original domain to choose a branch. For $f(x)=(x-2)^2-5$ on $x\geq 2$, the inverse is $2+\sqrt{(y+5)}$ on $y\geq -5$. The minus branch would return inputs outside the chosen domain. Check $f^{-1}(f(x))=x$ for allowed x and $f(f^{-1}(y))=y$ for allowed y; one unchecked composition can conceal a domain error.
- The inverse graph reflects the original graph across $y=x$; it does not take reciprocals of output values. A self-inverse function, or involution, satisfies $f(f(x))=x$ wherever the composition is defined. Both $-x$ on R and $1/x$ on R excluding zero are involutions. A strictly increasing involution on an interval must be the identity: if $f(x)>x$, increasingness gives $f(f(x))>f(x)>x$, and the analogous argument rules out $f(x)<x$.
- Iteration f^n means repeated composition, not the power $(f(x))^n$. Calculate the first few compositions and check their domains before looking for a period. For $f(x)=1/(1-x)$ on R excluding 0 and 1, $f^2(x)=(x-1)/x$ and $f^3(x)=x$. The image stays in the same allowed domain. Therefore reduce an iteration count modulo three. A displayed formula equal to x after cancellation does not restore forbidden inputs.

■ Checked example

For $f(x)=(x-2)^2-5$ with $x\geq 2$, $f(5)=4$ and $f^{-1}(4)=2+\sqrt{9}=5$. The unrestricted quadratic would have two inputs, −1 and 5, giving output 4. For $g(x)=1/(1-x)$, the orbit $3\to -1/2\to 2/3\to 3$ has period three, so $g^8(3)=2/3$; the count concerns compositions, not eighth powers.

exercises.lol exercises.com.cn

■ **Common error**

Changing a codomain changes surjectivity. Reflection across y=x is an inverse graph; reciprocating y-values is a different operation. Do not cancel away excluded inputs.

T.7 · Set images, equivalence relations and logical negation

- For f:X→Y and A⊆X, the image f(A) consists of all outputs f(a) with a in A. A point is in f(A∪B) exactly when it comes from A or B, so f(A∪B)=f(A)∪f(B). If A⊆B, then f(A)⊆f(B). For an intersection only f(A∩B)⊆f(A)∩f(B) is automatic: the same output may come from different inputs. Use $f(x)=x^2$, A={−1}, B={1}; the left image is empty while the right intersection is {1}.
- Preimages behave differently. For C⊆Y, $f^{-1}(C)$ here denotes the set of all inputs sent into C, even if f has no inverse function. Membership in two preimages means the very same input maps into both target sets. Consequently preimages preserve unions, intersections and complements relative to the stated domain/codomain. Do not transfer a theorem about preimages to images. If f is injective, image intersections do become equal, because equal outputs then force a shared input.
- A relation R on X is reflexive when xRx for every x, symmetric when xRy implies yRx, and transitive when xRy and yRz imply xRz. All three make an equivalence relation; its classes partition X. On integers, xRy when x and y have the same remainder modulo four gives four classes. The relation |x−y|≤1 is reflexive and symmetric but not transitive: 0R1 and 1R2 while 0 is not related to 2. Checking only a diagram or two properties is insufficient.
- An implication P⇒Q is false exactly when P is true and Q false. Thus the negation of P⇒(Q∧R) is P∧(¬Q∨¬R), not ¬P⇒(¬Q∧¬R). Negating "every x has property A" gives "there exists x without A"; negating "there exists x" gives "every x does not". A counterexample can disprove a universal claim, while examples cannot prove it. The quantifiers retain their order when individually negated: ¬(∀x∃y S(x,y)) is ∃x∀y ¬S(x,y).

■ **Checked example**

Let f map both a and b to label L. With A={a}, B={b}, f(A∩B)=∅ but f(A)∩f(B)={L}. For xRy defined by x=y or x=−y on R, reflexivity and symmetry follow directly, and two sign changes still give z=±x, proving transitivity. The classes are {x,−x}, with {0} a singleton. A false statement "every stored file has a hash" means at least one stored file has no hash; it does not mean all files lack hashes.

■ **Common error**

$f^{-1}(C)$ can mean a preimage set without an inverse function. A reflexive, symmetric relation may still fail transitivity. Negate the whole statement before simplifying its parts.

exercises.lol exercises.com.cn

T.8 · Conditioning, Bayesian inference and sampling error

- For P(B)>0, P(A|B)=P(A∩B)/P(B): restrict the population to B before computing the fraction in A. Independence means P(A∩B)=P(A)P(B), equivalently P(A|B)=P(A) when the denominator is nonzero. Mutually exclusive events with positive probabilities cannot be independent, because their joint probability is zero. Sampling without replacement usually changes later probabilities; a fixed denominator for successive draws would describe a different experiment.
- Build joint probabilities by multiplying a prior category proportion by the relevant conditional probability. If D is the condition and + a positive result, P(D∩+)=P(D)P(+|D). The total positive probability is P(D)P(+|D)+P(not D)P(+|not D), because these are disjoint and cover every positive result. Bayes divides the first joint probability by this total. Sensitivity is P(+|D), while specificity is P(−|not D); the false-positive rate is one minus specificity.
- For a hypothetical test with prevalence 2%, sensitivity 90% and specificity 95%, use a population of 10,000 for clarity. Of 200 people with the condition, 180 test positive. Of 9,800 without it, 490 test positive. Therefore 180 of 670 positive results have the condition: posterior 18/67, about 26.9%. Reversing the conditional would give 90%, answering a different question. This is a mathematical model, not a recommendation about clinical decisions.
- For independent identically distributed observations with finite variance σ^2, the sample mean has expectation μ and variance σ^2/n; its standard error is $\sigma/\sqrt{\mathrm{n}}$. This is spread of repeated sample means, not the spread of individual observations. Quadrupling n halves the standard error, rather than quartering it. Normal population data give an exactly normal mean; otherwise a central-limit approximation needs adequate conditions and sample size. Correlation invalidates the simple independent variance calculation.

■ Checked example

Let P(A)=0.4, P(B)=0.5 and P(A∩B)=0.3. Then P(A|B)=0.6, so A and B are not independent, since $0.3 \neq 0.4 \cdot 0.5$. If a variable has standard deviation 12, an independent sample of size 36 has mean standard error 2. Increasing the sample size to 144 gives 1; the individual-observation standard deviation stays 12.

■ Common error

Do not reverse P(+|D) into P(D|+). Include false positives in the posterior denominator. Standard deviation of individual observations and standard error of their mean have different sample-size behaviour.

exercises.lol exercises.com.cn

T.9 · Similarity, scaling and conic distance loci

- Triangles are similar when their corresponding angles agree; their corresponding side ratios are then equal. Write the vertex correspondence explicitly. If triangle APQ has angle A equal to angle A of ABC and angle P equal to angle C, its ordered correspondence is A↔A, P↔C, Q↔B. Thus AP/AC=AQ/AB=PQ/CB. The side AP lies on AB but corresponds to AC; matching by where a side lies or by the letter P can give the wrong ratio.
- With a common length scale factor $k>0$, lengths multiply by k, areas by k^2 and volumes by k^3. A 30% length increase means $k=1.30$, hence an area increase of $1.30^2-1=0.69$, or 69%; it is not 60%. For similar solids with a volume ratio of 27, the length ratio is 3 and the surface-area ratio 9. Percentage changes and absolute differences are different quantities; write the new-to-old ratio before converting it to a percent.
- A circle fixes distance to one point. An ellipse fixes the sum of distances to two foci. A hyperbola fixes the absolute difference of those distances; a parabola fixes equality of distance to a focus and to a directrix. These are definitions of point sets, not sketches to memorise. For foci $(-c,0)$ and $(c,0)$, an ellipse uses a sum $2a$ with $a>c$. A nondegenerate hyperbola uses an absolute difference $2a$ with $0<a<c$; a signed difference distinguishes its two branches.
- For the hyperbola with horizontal transverse axis, $c^2=a^2+b^2$ and the equation is $x^2/a^2-y^2/b^2=1$. Let r_A be distance to $(-c,0)$ and r_B distance to $(c,0)$. The condition $r_A-r_B=2a>0$ selects the right branch $x\geq a$; the negative condition selects the left. Squaring can lose this sign, so check the original distance condition after obtaining the equation. Difference zero gives the perpendicular bisector rather than a hyperbola; a difference exceeding the focal separation is impossible by the triangle inequality.

■ Checked example

Let AB=10, AC=15 and BC=20. A point P on AB has AP=6, and Q on AC makes angle APQ equal to angle ACB. Then AP/AC=6/15=2/5, so PQ=8 and AQ=4. Separately, foci at $(-5,0),(5,0)$ and distance-to-left minus distance-to-right equal to 6 give $a=3$, $b^2=25-9=16$: $x^2/9-y^2/16=1$ with $x\geq 3$. The point $(-3,0)$ satisfies the squared equation but fails the signed condition.

■ Common error

Match angles before sides. Area percentages use the square of the length factor. A squared distance equation may describe both branches even when the original problem asks for only one.

exercises.lol exercises.com.cn

C.8 · Trigonometric phase and parametric curves

- For y=A cos(ωx−φ) with A>0 and ω>0, amplitude is A, period is 2π/ω, and horizontal shift is φ/ω. The phase angle φ is measured inside the cosine argument; it is not itself the horizontal shift unless ω=1. Start with peak-to-peak spacing for the period and midline-to-peak distance for amplitude. To find phase, substitute a known point and check the direction of motion there. Equivalent phases differ by 2π.
- For y=−2 cos(3x), rewrite the model with positive amplitude as 2 cos(3x−π). Its amplitude is 2, period 2π/3, and phase π; the equivalent right shift is π/3. Peaks occur where 3x−π is a multiple of 2π. A graph point at x=0 and y=0 alone cannot determine phase uniquely: slopes or another point distinguish the possible angles. Units and angle conventions matter; ordinary calculus trigonometric derivatives use radians.
- A parametric curve specifies x=x(t), y=y(t). Eliminating t describes a point set but can lose restrictions and direction. For $x=\cos^3 t$, $y=\sin^3 t$, real cube roots give $|x|^{(2/3)}+|y|^{(2/3)}=1$, an astroid with cusps on the axes. As t runs from 0 to 2π it starts at (1,0), passes (0,1) at π/2 and travels counterclockwise. Restricting t to [0,π/2] gives only the first-quadrant arc, not the entire implicit locus.
- When dx/dt≠0, dy/dx=(dy/dt)/(dx/dt). A horizontal tangent generally requires dy/dt=0 and dx/dt≠0; a vertical tangent generally reverses those conditions. If both vanish, inspect a limit or the local expansion rather than taking 0/0 as a slope. For $x=t^2$,$y=t^3$ at t=0, the quotient for t≠0 is 3t/2 and tends to zero: the cusp has a horizontal tangent. A zero parameter velocity is not by itself a local maximum or minimum of y as a function of x.

■ Checked example

For y=−2 cos(3x), minima include x=0 and maxima include x=π/3. The phase representation 2 cos(3x−π) gives the same values. For x=2 cos t,y=sin t, elimination gives $x^2/4+y^2=1$; at t=π/4, dx/dt=−√2 and dy/dt=√2/2, hence slope −1/2. The full parameter interval [0,2π] traces the ellipse once counterclockwise.

■ Common error

Phase angle and horizontal shift differ by the frequency factor. An implicit equation can add untraced parts of a curve. If both parameter derivatives vanish, use local reasoning instead of labelling every such point an extremum.

T.10 · Vector geometry, projections and oriented area

- For nonzero real vectors u,v, u · v=|u||v| cos θ. A positive, zero or negative dot product corresponds to an acute, right or obtuse smaller angle. The scalar projection of v along u is (v · u)/|u|; its vector projection is ((v · u)/(u · u))u. The difference from this projection is orthogonal to u. Do not confuse a projected vector

exercises.lol exercises.com.cn

with its signed scalar component. The zero vector is orthogonal to every vector but has no defined angle direction.

- For three-dimensional vectors, $u \times v$ is perpendicular to both with length $|u||v| \sin \theta$. Coordinate calculation uses $(u_2v_3 - u_3v_2, u_3v_1 - u_1v_3, u_1v_2 - u_2v_1)$. This length is the parallelogram area, so a triangle from two edge vectors has half that area. Reversing their order reverses the cross product but preserves area. Build both edge vectors from the same vertex; crossing two unrelated position vectors generally measures the wrong triangle.
- The plane through a point p with nonzero normal n has equation $n \cdot (x - p) = 0$. Distance from q to the plane is $|n \cdot (q - p)|/|n|$; the denominator normalises the scale of the equation. A scalar triple product $u \cdot (v \times w)$ gives signed parallelepiped volume; its absolute value is geometric volume. Zero triple product means dependence of the three edge vectors, not necessarily that each pair is perpendicular or parallel.
- For four planar vectors, there are six unordered dot products. The configuration $e_1, -e_1, e_2, -e_2$ has two negative products and four zeros; e_1, e_1, e_2, e_2 has two positive products and four zeros. Four nonzero vectors cannot have every pairwise dot product negative. Order their directions around the circle: each consecutive angular gap would have to exceed 90°, forcing the sum of four gaps above 360°. A zero vector cannot rescue a strict-negative requirement, because its dot products vanish.

Checked example

With $p = (0,0,0)$, $q = (2,0,0)$, $r = (0,3,0)$, the edges are $u = (2,0,0)$, $v = (0,3,0)$. Their cross product is $(0,0,6)$, giving triangle area 3. For $w = (3,4)$ and $u = (1,0)$, the vector projection is $(3,0)$ and the orthogonal remainder is $(0,4)$. The plane $2x - y + 2z = 6$ has normal length 3, so its distance from the origin is $6/3 = 2$.

Common error

A cross-product magnitude is parallelogram area, so halve it for a triangle. Dot-product zero is an algebraic orthogonality statement even for a zero vector. Plane distance must divide by normal length.

C.9 · Integrating factors and nonhomogeneous differential equations

- A first-order linear equation has the form $y' + p(x)y = q(x)$ on an interval where its coefficients are continuous. It need not be separable. Set $\mu(x) = \exp(\int p(x)dx)$. Since $\mu' = p\mu$, multiplying gives $(\mu y)' = \mu q$, hence $y = \mu^{-1}(\int \mu q \, dx + C)$. Choose any convenient antiderivative for p; its integration constant only rescales μ and cancels from the solution. Apply initial data after integration, and keep any singular coefficient points outside the chosen interval.

exercises.lol exercises.com.cn

- For y′+2xy=x with y(0)=1, μ=e^(x²). Then (e^(x²)y)′=x e^(x²), whose integral is e^(x²)/2+C. Thus y=1/2+C e^(−x²), and the initial condition gives C=1/2. Direct substitution checks y′+2xy=x. The integrating factor makes the left side a product derivative; it is not an extra factor that remains multiplying the original right side in the final answer.
- For ay″+by′+cy=r(x) with a nonzero, solve the characteristic equation aλ^2+bλ+c=0 for the homogeneous part. Distinct real roots give two exponentials; a repeated root λ gives (C_1+C_2x)e^(λx); roots $\alpha\pm i\beta$ give e^(αx)(C_1 cos βx+C_2 sin βx). The complete solution is y_h+y_p. A polynomial forcing suggests a polynomial particular trial; an exponential or sine/cosine forcing suggests the corresponding family, with enough coefficients to account for differentiation.
- If a particular trial duplicates a homogeneous solution, it cannot produce the forcing: multiply by x once for a simple root, twice for a repeated root. For y″−3y′+2y=e^x, the naive Ke^x is annihilated. Trying Kxe^x gives −Ke^x, so y_p=−xe^x. Initial conditions determine the homogeneous constants only after adding y_p. Substitute the final function into the original equation to check signs and forcing; an initial-value check alone cannot verify the differential equation.

■ Checked example

For y″−3y′+2y=4, the characteristic roots are 1 and 2. A constant particular solution y_p=2 gives 2y_p=4, so y=C_1e^x+C_2e^(2x)+2. For y′+2xy=x and y(0)=1, the separate first-order solution is y=(1+e^(−x²))/2. These illustrate distinct methods; neither equation becomes homogeneous just because its left side is linear.

■ Common error

Keep the particular solution. A resonant trial needs an x factor. The integrating-factor derivation applies on a valid coefficient interval and should not divide by y or silently discard zero solutions.

T.11 · Weak compositions, loop invariants and flowchart tracing

- The number of nonnegative integer solutions of $x_1+\cdots+x_k=n$ is C(n+k−1,k−1), for n≥0 and k≥1. Represent n identical stars separated into k named boxes by k−1 bars; adjacent bars and end bars allow zero entries. Boxes are distinguishable, objects are not. Counting distinct tokens would instead assign each token independently, giving k^n before other constraints. Decide the model from what the objects and recipients represent, not from a familiar-looking binomial coefficient.
- For positive allocations x_i≥1, put z_i=x_i−1; the remaining total is n−k, giving C(n−1,k−1) when n≥k. More general lower bounds x_i≥a_i are handled by subtracting each a_i. At least one zero is the complement of every entry positive: subtract C(n−1,k−1) from the nonnegative total. Overlapping cases such as

exercises.lol exercises.com.cn

exactly two empty boxes need a different count; subtracting each empty-box case independently double-counts their intersections.

- Trace a flowchart one executed statement at a time, recording variables in a table. Distinguish assignment from a mathematical equality test; the right side of an assignment uses the old value before replacement. Record which branch executes next and which variables reset. A new outer-loop candidate may reset an inner sum without resetting the candidate itself. For a process that adds 1,3,5,… to S from zero, after m additions $S=m^2$ and the next addend is $2m+1$.
- A loop invariant is true before and after each iteration. The square-sum invariant holds initially at $m=0$ and is preserved because $m^2+(2m+1)=(m+1)^2$. An invariant alone does not prove termination. To test whether an integer $N\geq 0$ is square by adding odd numbers until $S\geq N$, S grows without bound; equivalently m increases and must reach a bound such as N. At exit, equality means square and overshoot means nonsquare. A printed value must be reached on an actual execution path; a plausible numerical pattern alone is not a trace.

■ Checked example

Distribute 11 identical tokens among 4 named boxes. Nonnegative allocations number $C(14,3)=364$; all-positive allocations number $C(10,3)=120$. Therefore 244 allocations have at least one empty box. For the odd-sum test with $N=10$, S progresses 0,1,4,9,16; it overshoots and rejects 10. For $N=9$ it exits at $S=9$ after three additions. The next addends are 1,3,5,7,9; sum and next-addend columns must not be confused.

■ Common error

Identical objects and distinct objects use different models. At least one empty box is not the same as exactly one. In a flowchart, preserve update order and reset only variables that the executed path actually resets.

C.10 · Implicit differentiation and the inverse Jacobian

- For a differentiable map $F(u,v)=(x,y)=(f(u,v),g(u,v))$, its Jacobian is $J=[[f_u,f_v],[g_u,g_v]]$. Small changes satisfy $[dx,dy]^{\mathsf{T}}=J[du,dv]^{\mathsf{T}}$ to first order. To find u_x while holding y fixed, differentiate both defining equations with respect to x: f_u u_x+f_v v_x=1 and g_u u_x+g_v v_x=0. These are a coupled linear system, not two independent scalar inverse rules.
- If f and g are continuously differentiable near the point and $\det J=f_u\,g_v-f_v\,g_u$ is nonzero there, the inverse-function theorem supplies a differentiable local inverse. Its derivative is $J^{-1}=(1/\det J)[[g_v,-f_v],[-g_u,f_u]]$. Thus $u_x=g_v/\det J$, $u_y=-f_v/\det J$, $v_x=-g_u/\det J$ and $v_y=f_u/\det J$. Evaluate every derivative at the corresponding point. A local inverse need not extend to a global one.

exercises.lol exercises.com.cn

- For implicit equations H(u,v,x,y)=0 and K(u,v,x,y)=0, differentiate while fixing the requested independent coordinate. Solve [[H_u,H_v],[K_u,K_v]][u_x,v_x]$^{\mathsf{T}}$=−[H_x,K_x]$^{\mathsf{T}}$. The determinant in the unknown variables u,v must be nonzero to use the usual implicit-function theorem. Signs on the right come from moving known derivatives to the other side. If the determinant vanishes, this theorem is inconclusive; it does not by itself prove no inverse or no implicit solution exists.
- The scalar shortcut du/dx=1/(dx/du) holds for a one-variable inverse with nonzero derivative, but usually fails for a coupled system because v changes to keep y fixed. For x=u+v,y=u+2v, J=[[1,1],[1,2]] has determinant 1 and inverse [[2,−1],[−1,1]]. Therefore u_x=2 while 1/f_u=1. An inverse Jacobian transforms differential sensitivities; a change-of-variables integral instead uses the absolute determinant for area or volume scaling, not a selected inverse entry.

Checked example

Take x=u^2+v and y=u−v. At (u,v)=(1,0), the output is (1,1), and J=[[2,1],[1,−1]] has determinant −3. The inverse is [[1/3,1/3],[1/3,−2/3]], so u_x=1/3 and v_y=−2/3 there. Multiplying J by this inverse gives the identity. The reciprocal 1/f_u=1/2 is not u_x because changing u also requires changing v to keep y fixed.

Common error

Differentiate both equations and state which output is held fixed. Nonzero determinant guarantees a local inverse under the regularity hypotheses; zero determinant is not a proof of impossibility. Use an inverse entry for sensitivities and an absolute determinant for integration.

A.7 · LU factorisation and nullity of composed maps

- A factorisation A=LU expresses a square matrix as lower triangular L and upper triangular U. With unit diagonal L, forward substitution is especially simple. An invertible matrix need not admit this form without row exchanges: a zero leading pivot can require a permutation. For PA=LU, solve Ly=Pb, then Ux=y. The permutation acts on the right side as well as the coefficient matrix. Nonzero leading principal pivots justify the usual no-exchange elimination, not invertibility alone.
- In Ly=b, compute y from top to bottom, subtracting terms already known; divide by the current diagonal unless it is one. In Ux=y, compute x from bottom to top. For each equation substitute back into the original row as a check. Reusing LU for many right sides avoids repeating the elimination: dense factorisation takes cubic-order work in dimension, while each triangular solve takes quadratic-order work. Exact arithmetic can check small examples; numerical pivoting reduces some roundoff problems but cannot remove inherent ill-conditioning.
- For B:V→W and A:W→Z, ker B is contained in ker(A∘B), but the latter can be

exercises.lol exercises.com.cn

larger. Extra vectors are those mapped by B into ker A. Restrict B to ker(A∘B): its image is im B∩ker A and its kernel is ker B. Rank-nullity on this restricted map gives dim ker(A∘B)=dim ker B+dim(im B∩ker A). The intersection, not the whole kernel of A, determines the extra nullity.

- For endomorphisms of R^6 with nullity A=2 and nullity B=3, rank B=3 and the intersection dimension can range from 0 to 2. Thus nullity of A∘B ranges from 3 to 5. Bounds depend on the common intermediate space: for subspaces of dimensions r and s in an m-dimensional space, their intersection has dimension between max(0,r+s−m) and min(r,s). Reversing the composition can change its nullity, even though AB and BA are both defined. Choose compatible spaces before applying these formulas.

■ Checked example

Let L=[[1,0,0],[2,1,0],[−1,3,1]], U=[[2,1,−1],[0,3,2],[0,0,4]], and b=(−1,−1,12). Forward substitution gives y=(−1,1,8). Back substitution gives $x_3=2$, $x_2=(1-4)/3=-1$, and $x_1=(-1-(-1)+2)/2=1$. Therefore x=(1,−1,2). Multiplying Ux gives y and multiplying Ly gives b, independently checking the factor order.

■ Common error

Solve with L first and U second, and permute b if PA=LU. A composition's nullity is not automatically the sum of the two nullities; only the image-kernel intersection adds to nullity B.

C.11 · Integration by parts, order reversal and symmetry

- Integration by parts comes from the product rule: $\int u\,dv=uv-\int v\,du$. For definite integrals include the boundary term at both ends. With $I_n(x)=\int_1^x(\ln t)^n\,dt$, choose $u=(\ln t)^n$ and dv=dt. For n≥1 the lower boundary vanishes, giving $I_n=x(\ln x)^n-n\,I_{n-1}$. A plus sign would violate the differentiated identity. Keep the same lower limit in a recurrence; changing it changes the constants.
- For improper integrals first apply the identity on finite endpoints, then justify the limiting boundary and remaining integral. If $\int_{-\infty}^{\infty}e^{(-x^2)}dx=\sqrt{\pi}$, set u=x and $dv=x\,e^{(-x^2)}dx$ for $\int x^2e^{(-x^2)}dx$. Since $v=-e^{(-x^2)}/2$ and $x\,e^{(-x^2)}\to0$ at both infinities, the full second moment is $\sqrt{\pi}/2$. Evenness gives the positive-half-line moment $\sqrt{\pi}/4$. The vanished boundary is part of the argument, not an automatic property of every improper integral.
- To reverse $\int_0^1\int_x^1F(x,y)dy\,dx$, describe the triangle $0\le x\le y\le1$, then rewrite it as $\int_0^1\int_0^yF(x,y)dx\,dy$. Both the outer interval and inner bounds change; swapping symbols alone changes the region. For $F=e^{(y^2)}$, integrating over x first gives $\int_0^1y\,e^{(y^2)}dy=(e-1)/2$. Continuity on this compact triangle makes order reversal valid; singular or conditionally convergent cases require stronger care.
- If an integrable function on [a,b] satisfies f(a+b−x)=−f(x), reflection makes its

exercises.lol exercises.com.cn

integral equal to its negative, hence zero. For $\sin(2mx)/\sin x$ on $[0,\pi]$ with positive integer m, reflection $x\mapsto\pi-x$ changes the numerator sign and preserves the denominator. The apparent endpoint singularities are removable: limits are 2m at zero and $-2m$ at π. Check these limits before invoking symmetry; cancellation is not a substitute for integrability.

Checked example

For $I_2(x)=\int_1^x(\ln t)^2dt$, the recurrence gives $x(\ln x)^2-2I_1(x)$, with $I_1(x)=x\ln x-x+1$. Therefore $I_2=x[(\ln x)^2-2\ln x+2]-2$. The constant -2 ensures $I_2(1)=0$. Separately, reversing the triangular $e\hat{}(y^2)$ integral produces a factor y from the inner x-length; dropping that factor would restore the original difficulty and change the value.

Common error

Improper integration by parts needs endpoint limits. Reconstruct the region before reversing order. An odd-looking reflected integrand must be integrable, including any removable endpoints.

C.12 · Related rates and removable quotient limits

- For a quantity V depending on a changing depth h(t), the chain rule gives $dV/dt=V'(h)\,dh/dt$. A draining tank has $dh/dt<0$, so signed dV/dt is negative; an outflow magnitude is its negative. Draw the geometry and name the instantaneous depth before substituting numbers. Do not treat the rate as a static volume divided by elapsed time unless the situation actually specifies a constant average rate.
- A spherical tank of radius R has cross-section radius squared $R^2-(R-h)^2=2Rh-h^2$ at depth h from the bottom. Integrating these circular areas gives the cap volume $V(h)=\pi(Rh^2-h^3/3)$, for $0\le h\le 2R$. Its derivative $\pi(2Rh-h^2)$ is the current cross-sectional area. At $R=3,h=1,dh/dt=-1/4$, signed volume rate is $-5\pi/4$ and outflow magnitude $5\pi/4$. The formula works for both shallow and deep caps within the stated range.
- If f and g are continuously differentiable near zero, $f(0)=g(0)=0$ and $g'(0)\neq 0$, then $f(x)/g(x)$ tends to $f'(0)/g'(0)$. This follows from $f(x)=f'(0)x+o(x)$ and $g(x)=g'(0)x+o(x)$; continuity of g' keeps the quotient defined nearby except at zero. Filling in this limit gives a continuous extension. The argument does not require $f'(0)$ nonzero, and it does not prove the extended quotient differentiable.
- For $f(x)=x|x|$ and $g(x)=x$, both are continuously differentiable, $f(0)=g(0)=0$ and $g'(0)=1$. Their quotient for $x\neq 0$ is $|x|$, which extends continuously at zero but has unequal one-sided derivatives there. For $(f^2-f)/(2g-g^3)$, factor to $(f/g)(f-1)/(2-g^2)$; its removable limit is $-f'(0)/(2g'(0))$. Additional factors approach -1 and 2. Distinguish the limit of the quotient from the derivative of its extension.

exercises.lol exercises.com.cn

■ Checked example

For a sphere with R=3 and depth h=1, cross-sectional area is $\pi(6-1)=5\pi$. A depth decrease of 1/4 per time unit therefore removes volume at magnitude $5\pi/4$ per time unit. For f(x)=sin(2x),g(x)=3x, the quotient extends at zero with value 2/3. This limit says nothing about the size of a tank and should not be substituted as a derivative rate without identifying the dependent quantities.

■ Common error

Use instantaneous cross-section area, not total tank surface area. Report a signed volume change or a positive outflow as requested. A finite removable quotient limit establishes continuity, not differentiability of the filled-in quotient.

T.12 · Metric completeness and closure from a basis

- For an injective map h:X→R, d(x,y)=|h(x)−h(y)| is a metric: symmetry and the triangle inequality come from R, and injectivity ensures d(x,y)=0 only when x=y. The map h is an isometry onto its image h(X). If h is not injective, this construction may be only a pseudometric. Bounded distances do not prove a space complete, and a familiar set of points can have very different Cauchy behaviour under different metrics.
- The pullback metric is complete exactly when the image h(X) is complete in the ordinary real distance. A closed subset of R is complete. For h(x)=arctan x on R, the image is $(-\pi/2,\pi/2)$, which omits its endpoints. The sequence n is Cauchy in the pullback metric because arctan n→$\pi/2$, but no real x has arctan x=$\pi/2$, so the metric is incomplete. For $h(x)=x^3$, the image is all R and the pullback metric is complete. Topological equivalence alone does not preserve completeness.
- In a topology with a basis, x lies in the closure of A when every basic open neighbourhood of x meets A. This is a membership test, not automatically a Euclidean endpoint operation. The basis must cover the space, and for a point in two basis sets there must be a smaller basis set containing it within their intersection. A point can lie in the closure of a singleton even when it is not the singleton's own point; that possibility depends on the topology.
- On X={2,3,4,…}, take basic sets U_k={n in X: n divides k}, for k≥2. They cover X since x∈U_x; intersections are U_gcd(k,l) when the gcd is at least 2, or empty. A point x lies in closure of {n} exactly when every k divisible by x is also divisible by n: equivalently n divides x. More directly, the smallest basic neighbourhood U_x contains n precisely when n divides x, and every other neighbourhood containing x includes these divisors. Thus closure of {8} consists of positive multiples of 8, not the divisors of 8.

exercises.lol exercises.com.cn

■ Checked example

Under d(x,y)=|arctan x−arctan y|, the integer sequence goes toward a missing image endpoint rather than a point of R, so it is Cauchy without convergence in that space. In the divisor basis, 16 lies in closure of {8}, since any basic set containing 16 also contains 8. But 4 does not: U_4={2,4} is a neighbourhood of 4 missing 8. This explicitly separates multiples from divisors.

■ Common error

A Cauchy limit must belong to the same space. Completeness is a metric property, not just a topological one. For closure, test every neighbourhood of the candidate point, rather than every neighbourhood of the singleton value.

exercises.lol exercises.com.cn