Prof. Claude

Lessons from conversations with Claude

Greg Howell and Claude

Contents

Convolution: one operation wearing four costumes **3**
- Videos . . . 3
- Level 1: any group — the big picture . . . 3
- Level 2: $\mathbb{R}$ — where the minus sign comes from . . . 4
- Level 3: $\mathbb{Z}$ — convolution is multiplying polynomials . . . 5
- Level 4: $\mathbb{Z}/N$ — convolution becomes a matrix . . . 6
- The punchline: why Fourier diagonalizes all of it . . . 6
- FFT: exploiting the subgroup . . . 7
- The one-sentence summary . . . 8
- Fine print: why the trivial σ-algebra breaks "abelian ⟺ convolution commutative" 8
- The Gauss heat kernel . . . 9
- Haar Measure: the dy that makes everything else legal . . . 10
- Self-test: convolution from $\mathbb{R}$ to $\mathbb{R}^n$ to S^1 . . . 13

Euler Rotations, Spinors and Gimbal Lock **17**
- The big picture in one breath . . . 17
- Euler angles, Z-Y-Z flavor . . . 17
- Gimbal lock — where it breaks, and why it *must* . . . 18
- Quaternions from scratch . . . 18
- Rotation by conjugation, and the half-angle . . . 18
- Composition — same order as matrices . . . 19
- The same thing in physics clothes: $SU(2)$ and spinors . . . 20
- Euler Z-Y-Z, done in quaternions . . . 20
- Why the Euler rates blow up . . . 21
- Picturing the double cover . . . 23
- Why the flight computer carries q instead of (α, β, γ) . . . 24
- Summary . . . 25
- Problems . . . 25

Inequalities **27**
- Jensen's Inequality . . . 27
- Hölder . . . 31
- Minkowski . . . 33
- Young . . . 34
- Exercises . . . 36

Newton's Binomial Series **38**

What Newton actually did that everyone glosses over . . . 38
The numerics, done properly . . . 38
The coding application . . . 39
Generating functions — Catalan numbers fall out . . . 39
Physics: the correction terms are the physics . . . 40
Newton's own killer app: term-by-term integration . . . 40
The big-picture reveal . . . 40
Exercises . . . 41

Tensors **42**
A bridge from convolution . . . 42
Tensor product: the machine that eats bilinearity . . . 42
Symmetric and alternating: quotient the tensor algebra . . . 43
Wedge: why determinants exist . . . 43
Why forms are what you integrate . . . 44
Tensors on a smooth manifold . . . 45
Homology and cohomology . . . 47
Exercises . . . 48

Visualizing Grad, Div, and Curl **51**
The three pictures . . . 51
Is $\nabla \times \nabla f = 0$ visually obvious? . . . 52
And the determinant? . . . 52
One caveat to hold in reserve . . . 53
See it: the surface explorer . . . 53
Problems . . . 54

Convolution: one operation wearing four costumes

Videos

3Blue1Brown

Level 1: any group — the big picture

Here is the whole subject in a sentence:

> **A group has a multiplication. Functions on the group inherit it. That inherited multiplication is convolution.**

Everything else — the integrals, the flipping, the FFT — is bookkeeping around that one idea.

More precisely, the cleanest definition isn't the integral formula at all. It's this: let G be a group with a measurable operation, and let μ, ν be finite measures on G. Define

$$\mu * \nu := (\text{pushforward of } \mu \otimes \nu \text{ under the map } (x, y) \mapsto xy).$$

Equivalently, tested against a function,

$$\int_G f \, d(\mu * \nu) \;=\; \int_G \int_G f(xy) \, d\mu(x) \, d\nu(y).$$

That's it. No flipping, no minus signs, no $g(x - y)$. Those appear only when you insist on writing things in terms of densities.

Two sanity checks that make the definition feel inevitable:

- $\delta_a * \delta_b = \delta_{ab}$. The group sits inside its own measure algebra. Convolution *is* the group law, extended by linearity.
- $\mu * \nu$ is commutative iff G is abelian, associative always. You get these for free from G; you don't prove them, you inherit them.

And the probabilistic reading, which is the one to keep in your pocket: **if $X \sim \mu$ and $Y \sim \nu$ are independent, then $X + Y \sim \mu * \nu$.** Convolution is the arithmetic of independent randomness.

The rest of the chapter walks down a ladder of groups: $\mathbb{R}$, $\mathbb{Z}$, $\mathbb{Z}/N$, and then back up to Fourier analysis and the FFT, which work on all of them at once.

Level 2: $\mathbb{R}$ — where the minus sign comes from

Suppose $d\mu = f\,dx$ and $d\nu = g\,dx$. Feed that into the abstract definition:

$$\int h\,d(\mu * \nu) = \iint h(x+y) f(x) g(y)\,dx\,dy.$$

Substitute $u = x + y$ inside the x-integral, so $x = u - y$:

$$= \int h(u) \left[\int f(u-y) g(y)\,dy \right] du.$$

The bracket is the density of $\mu * \nu$. **The flip is a change of variables, not a definition.** Students find the flip mysterious because it is presented as a definition; nobody chose to reverse anything, the group law did it.

$$(f * g)(x) = \int_{\mathbb{R}} f(x-y)\, g(y)\,dy.$$

What convolution does to a function. Read it as: *place a copy of f, weighted by $g(y)$, at every location y, and add them up.* Convolution is superposition of translates. Which is exactly why it's the universal answer to "what does a linear time-invariant system do?" — a system that's linear (superposition) and commutes with translation must be convolution with its response to a delta. Feed a room an impulse (clap), record the response g; then the room's effect on *any* sound f is $f * g$. That's reverb, and it's why convolution reverb plugins are literally recordings of cathedrals.

Same theorem, other costumes: a camera's out-of-focus blur is convolution with a disk; a moving average is convolution with a box; diffusion for time t is convolution with the Gaussian heat kernel G_t, and the semigroup law $G_s * G_t = G_{s+t}$ says *diffusing for 3 seconds then 4 more is diffusing for 7.*

The three properties worth memorizing:

1. **Smoothing.** $(f * g)' = f' * g = f * g'$. The derivative can land on whichever factor can afford it. So $f * g$ is as smooth as the *smoothest* factor — convolve anything in L^1

with a C^∞ bump and you get C^∞. This one fact powers mollification, and hence half of distribution theory.

2. **Young's inequality.** $\|f*g\|_r \le \|f\|_p \|g\|_q$ when $\frac{1}{p}+\frac{1}{q} = 1+\frac{1}{r}$. The case $p=q=r=1$ says $L^1(\mathbb{R})$ is a Banach algebra.
3. **Support.** $\operatorname{supp}(f*g) \subseteq \overline{\operatorname{supp} f + \operatorname{supp} g}$. Convolution spreads.

A structural wrinkle worth noticing. $L^1(\mathbb{R})$ is a Banach algebra with *no identity*. The identity would have to be δ_0, and δ_0 is not a function. You get consolation prizes — approximate identities $\varphi_\varepsilon(x) = \varepsilon^{-1}\varphi(x/\varepsilon)$ with $f * \varphi_\varepsilon \to f$ — but no unit. Contrast this with $\ell^1(\mathbb{Z})$ below, which *does* have a unit. The difference is exactly discreteness: on a discrete group, counting measure is Haar measure, so the point mass *is* an honest function. Non-discrete groups pay for their continuum with a missing identity.

Fun fact to file away: **Titchmarsh's convolution theorem** says $\inf\operatorname{supp}(f*g) = \inf\operatorname{supp} f + \inf\operatorname{supp} g$. Consequence: $L^1[0,\infty)$ has no zero divisors. You cannot convolve two nonzero causal signals and get silence.

Level 3: $\mathbb{Z}$ — convolution is multiplying polynomials

Take $G = \mathbb{Z}$, Haar = counting measure. The integral becomes a sum:

$$(a*b)_n = \sum_{k\in\mathbb{Z}} a_{n-k}\, b_k.$$

Now stare at that and recall how you multiply polynomials:

$$\Big(\sum a_n x^n\Big)\Big(\sum b_n x^n\Big) = \sum_n \Big(\sum_k a_{n-k} b_k\Big) x^n.$$

Polynomial multiplication is convolution of coefficient sequences. Not "is like" — is. And since $x^i \cdot x^j = x^{i+j}$, this is just the $\delta_a * \delta_b = \delta_{a+b}$ rule again. The group algebra $\mathbb{C}[\mathbb{Z}]$ *is* the Laurent polynomial ring.

Two more instances of the same statement: long multiplication of integers is convolution of digit strings (with carries as an afterthought), and the product of two generating functions is the convolution of their sequences — which is why generating functions solve counting problems where you combine independent choices.

Here $\ell^1(\mathbb{Z})$ is a Banach algebra *with* unit $\delta_0 = (\dots, 0, 1, 0, \dots)$. Its Gelfand theory is gorgeous: every character of $\ell^1(\mathbb{Z})$ is evaluation of the Fourier series at a point of the circle, so the maximal ideal space is $\mathbb{T}$. Immediate corollary — **Wiener's lemma**: if $f(\theta) = \sum a_n e^{in\theta}$ has $\sum |a_n| < \infty$ and never vanishes, then $1/f$ also has an absolutely convergent Fourier

series. Wiener's original proof was hard; Gelfand's is three lines, because *nonvanishing = invertible in every character = invertible.*

Level 4: $\mathbb{Z}/N$ — convolution becomes a matrix

Finite cyclic group, indices mod N:

$$(a \circledast b)_n = \sum_{k=0}^{N-1} a_{(n-k) \bmod N}\, b_k.$$

Fix b and regard $a \mapsto a \circledast b$ as a linear map. Its matrix is

$$C_b = \begin{pmatrix} b_0 & b_{N-1} & \cdots & b_1 \\ b_1 & b_0 & \cdots & b_2 \\ \vdots & & \ddots & \vdots \\ b_{N-1} & b_{N-2} & \cdots & b_0 \end{pmatrix}$$

— a **circulant matrix**, each row the previous one shifted. So: *the circulant matrices are exactly the operators commuting with the cyclic shift*, and they form a commutative algebra isomorphic to $\mathbb{C}[\mathbb{Z}/N]$. Convolution algebra = commutant of translation. Hold that thought for ten seconds.

The punchline: why Fourier diagonalizes all of it

Convolution commutes with translation. So does every operator built from it. A commuting family of normal operators is simultaneously diagonalizable — so there ought to be a basis in which *every* convolution operator is diagonal. What are the eigenvectors of translation?

$$\chi(x+y) = \chi(x)\chi(y).$$

The characters. On $\mathbb{R}$: $e^{i\xi x}$. On $\mathbb{Z}/N$: $\omega^{kn}, \omega = e^{2\pi i/N}$. **The Fourier transform is nothing but the change of basis into the eigenbasis of translation.** That's the whole reason it exists.

And now the convolution theorem proves itself:

$$\widehat{\mu * \nu}(\chi) = \iint \overline{\chi(x+y)}\, d\mu\, d\nu = \int \overline{\chi(x)} d\mu \int \overline{\chi(y)} d\nu = \hat{\mu}(\chi)\, \hat{\nu}(\chi).$$

One line, and the only input was $\chi(x+y) = \chi(x)\chi(y)$. Every familiar version — $\widehat{f * g} = \hat{f}\hat{g}$ on $\mathbb{R}$, "the characteristic function of a sum is the product of characteristic functions" in probability, "DFT turns cyclic convolution into pointwise product" in signal processing — is this same line, read on a different group. Pontryagin duality is the statement that this picture is symmetric: $\widehat{\widehat{G}} \cong G$, and $\mathbb{R}$ is self-dual, $\widehat{\mathbb{Z}} = \mathbb{T}$, $\widehat{\mathbb{Z}/N} = \mathbb{Z}/N$.

Free bonus — the CLT, in three moves. Let X_i be iid, mean 0, variance 1. The density of the sum is μ^{*n}, so its characteristic function is $\hat{\mu}(\xi)^n$. Rescale by $\sqrt{n}$: $\left(\hat{\mu}(\xi/\sqrt{n})\right)^n$. Expand $\hat{\mu}(\xi) = 1 - \xi^2/2 + o(\xi^2)$, so this is $\left(1 - \frac{\xi^2}{2n} + o(1/n)\right)^n \to e^{-\xi^2/2}$. The Gaussian is the fixed point of convolution because it's the fixed point of squaring-after-rescaling on the Fourier side. Convolution is hard; multiplication is easy; that's the entire trick.

FFT: exploiting the subgroup

The DFT is a matrix–vector multiply, $O(N^2)$. Cooley–Tukey (1965; Gauss got there in 1805 and didn't publish, because of course he did) makes it $O(N \log N)$. The idea is pure group theory: $\mathbb{Z}/N$ for even N has a subgroup of index 2.

Split a into even- and odd-indexed halves. Then

$$\hat{a}_k = \underbrace{\hat{a}_k^{\text{even}}}_{\text{DFT of length } N/2} + \omega_N^k \underbrace{\hat{a}_k^{\text{odd}}}_{\text{DFT of length } N/2},$$

and because $\omega_N^{k+N/2} = -\omega_N^k$, the second half of the output is the same two pieces recombined with a minus sign. Two half-size transforms plus N multiply-adds. $T(N) = 2T(N/2) + O(N) \Rightarrow O(N \log N)$. That $\pm$ pair is the "butterfly," and the recursion is just Fourier analysis on $\mathbb{Z}/N$ relative to the filtration by subgroups.

The payoff:

$$a * b \;=\; \text{FFT}^{-1}\big(\text{FFT}(a) \cdot \text{FFT}(b)\big)$$

turns $O(N^2)$ convolution into $O(N \log N)$. Practical warning that bites everyone once: the FFT gives you *cyclic* convolution. If you want ordinary (linear) convolution of sequences of lengths m and n, zero-pad both to length $\geq m + n - 1$, or the tail wraps around and contaminates the head. This is called time-domain aliasing, and it is the source of many mysterious bugs.

Downstream consequences: multiplying two n-digit integers in $O(n \log n)$ (Schönhage–Strassen and descendants, which is how your computer multiplies big numbers), fast

polynomial multiplication, JPEG and MP3, MRI reconstruction, radio astronomy, and every convolutional neural net layer that's large enough to make the trade worthwhile.

The one-sentence summary

Convolution is multiplication in the group algebra; the Fourier transform is the change of basis to characters, which are the eigenvectors of translation; and the FFT is that change of basis performed cleverly by recursing through subgroups.

The next three sections fill in fine print: when the measure structure is too coarse for the abelian/commutative equivalence, the heat kernel as the model convolution semigroup, and Haar measure, the dy every formula above quietly depends on.

Fine print: why the trivial σ-algebra breaks "abelian ⟺ convolution commutative"

The *easy* direction is unconditional: G abelian $\Rightarrow$ convolution commutative, because $\int f\, d(\mu * \nu) = \iint f(xy)\, d\mu\, d\nu$ and you may swap $xy \rightsquigarrow yx$ under the double integral. Nothing to assume.

The *hard* direction — commutative convolution $\Rightarrow G$ abelian — is what needs the measure structure, and here is exactly where it leaks. The natural proof goes through point masses: $\delta_a * \delta_b = \delta_{ab}$, so if convolution commutes then $\delta_{ab} = \delta_{ba}$, and ***if point masses separate points*** you conclude $ab = ba$. That last clause is the load-bearing hypothesis.

Take $\mathcal{F} = \{\emptyset, G\}$. A measure is then determined by the single number $\mu(G)$, so $M(G) \cong [0, \infty)$ — scalars. Convolution collapses to

$$(\mu * \nu)(G) = \mu(G)\, \nu(G),$$

multiplication of total masses, which is commutative **no matter what G is**. Concretely: put $\mathcal{F} = \{\emptyset, S_3\}$ on the symmetric group. Convolution is commutative; S_3 is not. The iff is false.

What went wrong: with the trivial σ-algebra, $\delta_a = \delta_b$ for *all* a, b (every point mass is "mass 1 on the only measurable set that contains anything"). The point masses have been identified, so they can no longer witness $ab \neq ba$.

The fix is precisely what Borel measures on a Hausdorff group buy you. In a Hausdorff space singletons are closed and distinct points lie in disjoint Borel sets, so $\delta_a \neq \delta_b$ whenever

$a \neq b$ — the point masses separate points, and the hard direction runs. So the correct statement is:

> If the σ-algebra separates points (e.g. the Borel σ-algebra of a Hausdorff topological group, with $\cdot$ Borel-measurable), then convolution of measures is commutative $\iff$ G is abelian.

The trivial σ-algebra is the maximal violation of point-separation, which is why it's the canonical thing to exclude. (You also need $\cdot : G \times G \to G$ measurable for the pushforward to exist at all, and product-measurability of the diagonal for $\delta_a * \delta_b = \delta_{ab}$ — Borel + second-countable gives you all of it.)

The Gauss heat kernel

On $\mathbb{R}^n$,

$$G_t(x) = \frac{1}{(4\pi t)^{n/2}}\, e^{-|x|^2/4t}, \qquad t > 0,$$

is the fundamental solution of the heat equation $\partial_t u = \Delta u$: it's the unique solution with initial data δ_0. So the solution with *any* reasonable initial data f is

$$u(t, \cdot) = f * G_t.$$

The three facts that make it the star of the show:

- **Mass one, for all t:** $\int_{\mathbb{R}^n} G_t = 1$. It's a probability density — a Gaussian with covariance $2t\, I$.
- **Semigroup law:** $G_s * G_t = G_{s+t}$. Diffuse for s, then for t, same as diffusing for $s + t$. This is the whole content of "heat flow is a one-parameter semigroup $e^{t\Delta}$."
- **Fourier-diagonal:** $\widehat{G_t}(\xi) = e^{-t|\xi|^2}$. The semigroup law becomes $e^{-s|\xi|^2} e^{-t|\xi|^2} = e^{-(s+t)|\xi|^2}$ — trivial. Convolution is hard in x, a one-line product in ξ. This is *why* G_t is Gaussian: $e^{-t|\xi|^2}$ is the only thing whose logarithm is additive in t and radial.

Probabilist's convention warning: with $\partial_t = \frac{1}{2}\Delta$ you instead get $(2\pi t)^{-n/2} e^{-|x|^2/2t}$, covariance $t\, I$ — this is the transition density of standard Brownian motion, $B_{s+t} - B_s \sim N(0, tI)$, and the semigroup law is exactly "independent increments add their variances." Same object, factor-of-2 in the clock.

Two quick examples.

1. *Semigroup law by hand.* In 1-D, $G_t(x) = (4\pi t)^{-1/2} e^{-x^2/4t}$ is $N(0, 2t)$. Convolving $N(0, 2s) * N(0, 2t) = N(0, 2s + 2t) = G_{s+t}$ — the Gaussian-adds-variances fact, which is itself the cleanest instance of the convolution theorem.

2. *Instant smoothing.* Start with a discontinuous temperature, $f = \mathbf{1}_{[0,\infty)}$. Then $u(t,x) = (f * G_t)(x) = \frac{1}{2}(1 + \text{erf}(x/2\sqrt{t}))$ — the step is C^∞ for every $t > 0$. Heat flow has infinite propagation speed and kills every corner instantly; that's mollification wearing a physics hat.

Haar Measure: the dy that makes everything else legal

The big picture

Every convolution formula in this chapter has an invisible load-bearing symbol: the dy in $\int_G f(xy^{-1})\, g(y)\, dy$. On $\mathbb{R}$ it's Lebesgue, on $\mathbb{Z}$ it's counting, on S^1 it's $d\theta/2\pi$. The whole theory would collapse if these were *ad hoc* choices — but they aren't. Each is the same object wearing local clothes:

> **Haar measure is the essentially unique measure on a group that doesn't notice translation.**

"Doesn't notice translation" is the entire definition. If you slide a set over by multiplying every point by a fixed g, its measure must not change. That one demand — on a locally compact group — pins the measure down up to an overall scale. The reason convolution *works at all* is that this translation-blind measure exists; the reason the Fourier transform diagonalizes it is that the characters are the eigenfunctions of the translations the measure is blind to. Haar measure is the floor the whole building stands on.

The definition

Let G be a locally compact Hausdorff topological group. A **left Haar measure** is a nonzero Radon measure μ with

$$\mu(gE) = \mu(E) \qquad \text{for every } g \in G \text{ and every Borel } E,$$

where $gE = \{gx : x \in E\}$. Equivalently, on functions,

$$\int_G f(g^{-1}x)\, d\mu(x) = \int_G f(x)\, d\mu(x) \qquad \text{for all } f \in C_c(G).$$

Replace "left" with "right" ($\mu(Eg) = \mu(E)$) and you get a right Haar measure. The two can genuinely differ — that's the modular function below, and it's where the geometry gets interesting.

Theorem (Haar, 1933; Weil). Existence and uniqueness.

Every locally compact group carries a left Haar measure, and it is unique up to a positive multiplicative constant.

That uniqueness is the miracle. It means "translation-invariant measure" is not a *choice* you make but a *property the group has* — there's essentially only one, so Lebesgue measure isn't a definition we impose on $\mathbb{R}$, it's the thing $\mathbb{R}$ forces on us once we ask for shift-invariance.

The idea of existence (the tile-counting picture)

Haar's original construction, stripped to its essentials: To measure a compact set E, pick a tiny open set V and ask:

$$(E : V) = \text{the fewest left-translates } g_1V, \dots, g_kV \text{ needed to cover } E.$$

This count is *automatically left-invariant* — covering gE needs exactly as many tiles as covering E, because you just shift every tile by g. The raw count blows up as V shrinks, so normalize against a fixed reference set K_0:

$$\mu_V(E) = \frac{(E : V)}{(K_0 : V)}.$$

Now let V shrink to the identity. These ratios live in a compact product of intervals, so by Tychonoff a limit exists along an ultrafilter; the limit is finitely additive, left-invariant, and (this is the work) actually countably additive after invoking Riesz representation. That's it. **The measure of a set is "how many copies of a shrinking tile cover it," and invariance is free because shifting the set just shifts the tiles.** No cleverness about the specific group is needed — which is exactly why it works for *every* locally compact group at once.

The idea of uniqueness (a Fubini swap)

Suppose μ and ν are both left Haar. Take $f, g \in C_c(G)$ with $\int g\,d\mu \neq 0$ and look at the ratio $\int f\,d\nu \Big/ \int g\,d\mu$. The claim is it doesn't depend on g — which forces ν to be a constant multiple of μ.

The engine is a two-variable integral evaluated two ways:

$$\iint f(x)\, g(yx)\, d\nu(y)\, d\mu(x).$$

Integrate x first and use left-invariance of μ (substitute $x \rightsquigarrow y^{-1}x$) — the y drops out, giving $\int f\,d\nu \cdot \frac{1}{(\cdots)} \int g\,d\mu$-shaped terms. Integrate y first and use left-invariance of ν — you get the symmetric expression with roles swapped. Setting the two evaluations equal

shows $\int f\,d\nu / \int f\,d\mu$ is a constant independent of f. **Invariance under the group, pushed through Fubini, leaves no room for two independent measures.** (This is the skeleton; exercise 35 asks for every substitution written out.)

Left vs. right: the modular function

Fix a left Haar μ. For a fixed g, the measure $E \mapsto \mu(Eg)$ is *still* left-invariant (right and left multiplication commute). By uniqueness it must be a scalar multiple of μ:

$$\mu(Eg) = \Delta(g)\,\mu(E).$$

That scalar $\Delta(g) > 0$ is the **modular function**. Three facts fall out immediately:

- $\Delta : G \to (0, \infty)$ is a continuous **homomorphism** ($\Delta(gh) = \Delta(g)\Delta(h)$ — apply the definition twice).
- G is **unimodular** (left Haar = right Haar) exactly when $\Delta \equiv 1$.
- Right Haar measure is recovered as $d\mu_{\text{right}} = \Delta(g)^{-1}\,d\mu_{\text{left}}$.

And two beautiful sufficient conditions for unimodularity, each a one-line proof:

- **Abelian ⇒ unimodular.** $Eg = gE$, nothing to prove.
- **Compact ⇒ unimodular.** $\Delta(G)$ is a compact subgroup of $(0, \infty)$; the only one is $\{1\}$. (So all the finite groups, S^1, $SO(n)$, any compact Lie group — automatically unimodular, which is why Peter–Weyl works so cleanly.)

The groups where $\Delta \neq 1$ are the ones with genuine "stretch," and there's a canonical smallest example.

The worked examples

Group	Haar measure	Δ
$\mathbb{R}$ (add)	dx	1
S^1	$d\theta/2\pi$	1
$\mathbb{Z}$, finite groups	counting measure	1
$\mathbb{R}^\times$ (mult)	$dx/\lvert x\rvert$	1
$\mathbb{R}_{>0}$ (mult)	dx/x	1
$GL_n(\mathbb{R})$	$dX/\lvert\det X\rvert^n$	1
$ax+b$ group	left: $\frac{da\,db}{a^2}$	$\Delta = 1/a$

Why dx/x on $\mathbb{R}_{>0}$. The group law is multiplication, so "translation by a" is $x \mapsto ax$. Under it, $d(ax)/(ax) = dx/x$ — the a cancels. Lebesgue dx would *not* be invariant ($d(ax) = a\,dx$); dividing by x fixes it. The logarithm is secretly the statement that $(\mathbb{R}_{>0}, \times) \cong (\mathbb{R}, +)$ carries dx/x to dt.

Why $dX/|\det X|^n$ on GL_n. Left-multiplying $X \mapsto AX$ acts on the n^2 matrix entries as a linear map whose Jacobian is $(\det A)^n$ (it's A acting on each of the n columns). So Lebesgue measure dX scales by $|\det A|^n$; dividing by $|\det X|^n$ cancels it exactly. The same computation on the right gives the same factor — hence GL_n is unimodular even though it's wildly non-abelian and non-compact.

The $ax+b$ group — the canonical non-unimodular example. Elements (a, b) with $a > 0$, acting on $\mathbb{R}$ as $x \mapsto ax+b$; multiplication $(a, b)(a', b') = (aa',\ ab'+b)$. A direct Jacobian computation gives

$$d\mu_{\text{left}} = \frac{da\,db}{a^2}, \qquad d\mu_{\text{right}} = \frac{da\,db}{a}, \qquad \Delta(a, b) = \frac{1}{a}.$$

These really are different measures — the left one weights small-a (big-contraction) elements more heavily. This asymmetry isn't pathology; it's the simplest shadow of the fact that non-abelian groups can stretch volume differently depending on which side you push from. Everything downstream — why non-unimodular groups make representation theory harder, why the Plancherel formula needs correction factors — starts here.

The thread back to convolution

Now the payoff, and the reason this section belongs in a convolution book. Define $(f * g)(x) = \int_G f(xy^{-1})g(y)\,d\mu(y)$ with μ **left** Haar. Convolution is associative *because* μ is left-invariant — the associativity proof is one substitution $y \rightsquigarrow$ (shifted y) that is only legal under invariance. Pick a non-invariant measure and $f * (g * h) \neq (f * g) * h$; the algebra evaporates. And on a *non-unimodular* group, left and right convolution genuinely differ, with Δ appearing as the correction between them. So the modular function isn't a curiosity off to the side — it's the exact bookkeeping that tells you how much convolution cares which hand you use.

Haar measure is the quiet hypothesis under every "dy" in this chapter.

Self-test: convolution from $\mathbb{R}$ to $\mathbb{R}^n$ to S^1

Difficulty: ∘ easy, ∘∘ medium, ● hard.

A. Warm-ups on $\mathbb{R}$

1. ∘ Let $b = \mathbf{1}_{[-\frac{1}{2},\frac{1}{2}]}$. Compute $b*b$ explicitly and sketch it. What's the smoothest derivative that still exists? Then compute $b * b * b$ and say its smoothness class.
2. ∘ Compute $\mathbf{1}_{[0,\infty)} * g$ for $g \in L^1$. (You've just discovered an antiderivative operator — which one?)

3. ∘ Show $(f * g)' = f' * g$ directly from the integral, stating the hypothesis you need to differentiate under the sign.
4. ∘∘ Let $e_\lambda(x) = e^{-\lambda x}\mathbf{1}_{[0,\infty)}$. Compute $e_\lambda * e_\lambda$ and $e_\lambda * e_\mu$ ($\lambda \neq \mu$). Identify both as named probability densities.
5. ∘∘ Prove $\operatorname{supp}(f * g) \subseteq \overline{\operatorname{supp} f + \operatorname{supp} g}$, and give an example where the inclusion is strict.
6. ∘∘ Fix a bump $\varphi \geq 0$, $\int \varphi = 1$, and set $\varphi_\varepsilon(x) = \varepsilon^{-1}\varphi(x/\varepsilon)$. Show $f * \varphi_\varepsilon \to f$ in L^1 as $\varepsilon \to 0$. Where exactly do you use $\int \varphi = 1$?

B. Up to $\mathbb{R}^n$

7. ∘ Verify $\int_{\mathbb{R}^n} G_t = 1$ by reducing to the 1-D Gaussian integral.
8. ∘∘ Prove $\widehat{G_t}(\xi) = e^{-t|\xi|^2}$ (your convention), then re-derive $G_s * G_t = G_{s+t}$ in one line on the Fourier side.
9. ∘∘ Show $u(t,x) = (f * G_t)(x)$ solves $\partial_t u = \Delta u$ with $u(0,\cdot) = f$, for f bounded continuous. Which property of G_t gives you the initial condition?
10. • Young's inequality: prove $\|f * g\|_r \leq \|f\|_p \|g\|_q$ for $\frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r}$. (Interpolation or a direct Hölder-with-split-exponents argument — either is fair.)

C. On the circle $S^1 = \mathbb{R}/2\pi\mathbb{Z}$

11. ∘ For 2π-periodic f, g, define $(f * g)(\theta) = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(\theta - \phi)g(\phi)\,d\phi$. Show $\widehat{(f * g)}(n) = \hat{f}(n)\hat{g}(n)$ where $\hat{f}(n) = \frac{1}{2\pi}\int f e^{-in\theta}$. (Same theorem, group S^1, dual $\mathbb{Z}$.)
12. ∘∘ The Dirichlet kernel $D_N = \sum_{|n|\leq N} e^{in\theta}$ reproduces partial Fourier sums as $S_N f = f * D_N$. Show $\|D_N\|_{L^1} \to \infty$. What convergence failure does this predict?
13. ∘∘ The Fejér kernel $K_N = \frac{1}{N+1}\sum_{k=0}^{N} D_k$ is ≥ 0 with $\frac{1}{2\pi}\int K_N = 1$. Show K_N is an approximate identity, hence $f * K_N \to f$ uniformly for continuous f. Why does Fejér succeed where Dirichlet fails?
14. ∘∘ Heat kernel on the circle: $p_t(\theta) = \sum_{n\in\mathbb{Z}} e^{-n^2 t}e^{in\theta}$. Verify it solves $\partial_t u = \partial_\theta^2 u$ on S^1 and that $p_s * p_t = p_{s+t}$. (You're doing convolution on S^1, so the Fourier variable is discrete — the semigroup law is a statement about $e^{-n^2 t}$.)

D. Differential-geometry flavored

15. ∘∘ Explain, precisely, why convolution needs a *group* and not just a manifold: identify the step in $(f * g)(x) = \int f(xy^{-1})g(y)\,dy$ that fails on a bare smooth manifold, and what structure a homogeneous space must supply to rescue a version of it.
16. ∘∘ Mollify to a partition of unity: given an open cover of a compact manifold, use convolution with a bump (in charts) to build a smooth partition of unity subordinate to it. State where "convolution smooths" is doing the work.

17. • On a compact Lie group G with bi-invariant Haar measure, show $L^2(G)$ convolution is diagonalized by matrix coefficients of irreducible reps (Peter–Weyl), and that for *non-abelian* G the convolution theorem reads $\widehat{f * g}(\pi) = \hat{g}(\pi)\hat{f}(\pi)$ — a *product of matrices*, order-reversed. Explain why the order reversal is forced.
18. • Short-time heat asymptotics on S^1: using Poisson summation on $p_t(0) = \sum_n e^{-n^2 t}$, show $p_t(0) \sim \frac{1}{\sqrt{4\pi t}} \cdot 2\pi \cdot \frac{1}{2\pi} = \frac{1}{\sqrt{4\pi t}}$ as $t \to 0^+$, and read off the length of the circle from the leading coefficient. (This is the baby case of "the heat trace sees the volume": Weyl's law in one dimension, and the first step toward heat-kernel proofs of index theorems.)

E. Probability

19. ◦ X, Y independent uniform on $[0, 1]$. Find the density of $X + Y$ by convolution and name it.
20. ◦ Sum of independent Poisson(λ) and Poisson(μ): do the convolution on $\mathbb{Z}_{\geq 0}$ and show it's Poisson$(\lambda + \mu)$. (Convolution of sequences = multiplying the generating functions $e^{\lambda(z-1)}$.)
21. ◦◦ Via characteristic functions $\hat{\mu}(\xi) = \mathbb{E}[e^{i\xi X}]$, show the Gaussian family is closed under convolution and identify G_t as a transition density. Which convention (Δ vs $\frac{1}{2}\Delta$) matches standard Brownian motion?
22. • Prove the CLT for iid mean-0 variance-1 X_i: characteristic function of $S_n/\sqrt{n}$ is $\hat{\mu}(\xi/\sqrt{n})^n$; expand and take the limit. State the one regularity fact (Lévy continuity) that upgrades pointwise convergence of $\hat{\mu}^n$ to convergence in distribution.

F. Boss level (pick one)

23. • Prove $L^1(\mathbb{R})$ has no convolution identity: assume $e * f = f$ for all $f \in L^1$, test against an approximate identity, derive a contradiction.
24. • Wiener's lemma via Gelfand: show every character of the Banach algebra $\ell^1(\mathbb{Z})$ is evaluation of $\sum a_n e^{in\theta}$ at some θ, and deduce that a nowhere-vanishing absolutely-convergent Fourier series has an absolutely-convergent reciprocal.
25. • Poisson summation as a convolution/duality statement: prove $\sum_n f(n) = \sum_k \hat{f}(2\pi k)$ for Schwartz f, then specialize to $f = G_t$ to recover the functional equation relating p_t on S^1 to G_t on $\mathbb{R}$ — i.e. the Jacobi theta transformation $\vartheta(1/t) = \sqrt{t}\,\vartheta(t)$ up to constants.

G. Haar measure

26. ◦ Verify directly that $d\theta/2\pi$ on S^1 and counting measure on $\mathbb{Z}/N$ are left-invariant, and compute the total mass of each.
27. ◦ Show dx/x is left-invariant for $(\mathbb{R}_{>0}, \times)$ and that pushing it through log gives Lebesgue dt on $\mathbb{R}$. What's the Haar measure of the interval $[a, b] \subset \mathbb{R}_{>0}$?

28. ∘∘ Prove $\mathbb{R}^\times = \mathbb{R} \setminus \{0\}$ under multiplication has Haar measure $dx/|x|$. Why the absolute value, and what goes wrong without it?
29. ∘∘ Prove $\Delta : G \to (0, \infty)$ is a homomorphism straight from $\mu(Eg) = \Delta(g)\mu(E)$, and that it's trivial when G is abelian.
30. ∘∘ Prove: **compact ⇒ unimodular.** (Image of Δ is a compact subgroup of the positive reals — finish it.)
31. ∘∘ For $GL_n(\mathbb{R})$, carry out the Jacobian computation showing $dX/|\det X|^n$ is left-invariant, and check it's also right-invariant (so GL_n is unimodular).
32. • The $ax + b$ group. From the multiplication law $(a, b)(a', b') = (aa', ab' + b)$, compute the left and right Haar measures by Jacobians, confirm $\Delta(a, b) = 1/a$, and verify $d\mu_{\text{right}} = \Delta^{-1} d\mu_{\text{left}}$.
33. • The 3-dimensional Heisenberg group (upper-triangular unipotent 3×3 matrices): show Lebesgue measure on the three strictly-upper entries is *both* left and right Haar, so the Heisenberg group is unimodular despite being non-abelian and non-compact. Contrast with #32 — what structural difference makes one unimodular and the other not?
34. ∘∘ Show convolution $(f * g)(x) = \int f(xy^{-1})g(y)\, d\mu(y)$ is associative, and pin down the exact line where left-invariance of μ is used. Then exhibit, on the $ax + b$ group, that using right Haar instead changes the product.
35. • (Uniqueness, full rigor — the boss.) Write out the Fubini-swap uniqueness proof completely: two left Haar measures μ, ν, the two-variable integral, both evaluations, and the conclusion that $\nu = c\mu$. The text above only sketches it; doing it yourself is the test of whether the idea landed.

Short paths. Problems 1–25 are tiered: the easy ones build fluency, the • ones are workouts. For one pass through every setting ($\mathbb{R}$, $\mathbb{R}^n$, S^1, differential geometry, probability) covering both the mechanics and the big theorems, do **1, 4, 8, 13, 18, 19, 22**.

For one pass through Haar measure, do **26, 27, 29, 30, 32, 34**. That walks from "verify the easy ones" through the homomorphism and compactness theorems to the non-unimodular example and the tie-back to convolution.

Euler Rotations, Spinors and Gimbal Lock

This is one of those places where a concrete engineering headache (why does a rocket's guidance system freak out?) turns out to be the same fact as a deep piece of topology (why is $\pi_1(SO(3)) = \mathbb{Z}_2$?).

The big picture in one breath

A rotation of 3-space has **3 degrees of freedom**. Euler's trick parametrizes them with three angles — cheap and intuitive, but it has singularities baked in that you *cannot* remove, for topological reasons. Spinors (equivalently unit quaternions, equivalently $SU(2)$) parametrize the *same* rotations with 4 numbers on a constraint surface, and they have **no** singularities — at the cost of being a *double* cover. Engineers pay that small tax (carrying a redundant number, and a sign ambiguity) to buy a representation of attitude that never blows up. That's the whole story; now the details.

Euler angles, Z-Y-Z flavor

Write $R = R_z(\gamma)\, R_y(\beta)\, R_z(\alpha)$, with

$$R_z(\alpha) = \begin{pmatrix} \cos\alpha & -\sin\alpha & 0 \\ \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad R_y(\beta) = \begin{pmatrix} \cos\beta & 0 & \sin\beta \\ 0 & 1 & 0 \\ -\sin\beta & 0 & \cos\beta \end{pmatrix}.$$

Read as **active rotations in a fixed frame**, you apply them right-to-left: spin by α about z, tilt by β about y, then spin by γ about the *fixed* z. (An *extrinsic* Z-Y-Z sequence equals the *intrinsic* Z-Y'-Z" sequence read in the opposite order: same matrix, two stories.) Conventional ranges: $\alpha \in [0, 2\pi)$, $\beta \in [0, \pi]$, $\gamma \in [0, 2\pi)$ — and β only needs half the circle precisely because the two z-spins can cover the rest.

Gimbal lock — where it breaks, and why it *must*

Set $\beta = 0$. Then $R_y(0) = I$ and the product collapses to $R_z(\gamma)R_z(\alpha) = R_z(\alpha+\gamma)$. The first and third gimbals now spin about the *same* axis: a degree of freedom is lost, and only the sum $\alpha + \gamma$ means anything. Two knobs, one effect. That's gimbal lock. Near $\beta = 0$ the map $(\alpha, \beta, \gamma) \mapsto SO(3)$ has a rank-deficient Jacobian, and if you try to *hold attitude* while passing through it, the required rates $\dot\alpha, \dot\gamma$ shoot to infinity (there's a $1/\sin\beta$ lurking in the kinematics, worked out below).

Why you can't just be cleverer about it: **you're trying to wrap a box around a sphere.** $SO(3)$ is a compact 3-manifold — in fact it's $\mathbb{RP}^3$ — and there is *no* smooth bijection from a coordinate box in $\mathbb{R}^3$ onto it. Any minimal 3-angle chart must degenerate *somewhere*. Euler angles aren't badly designed; they're paying an unavoidable topological tax. You can move the singularity around by choosing X-Y-Z or Z-X-Z conventions, but you can never delete it — the same reason one chart never covers a nontrivial compact manifold.

Quaternions from scratch

A quaternion is $q = a + b\,i + c\,j + d\,k$ with the multiplication rules $i^2 = j^2 = k^2 = ijk = -1$, which force $ij = k,\ jk = i,\ ki = j$ and $ji = -k$ etc. (noncommutative — that's the whole point). Write $q = (a, \vec{u})$ splitting scalar and vector parts; then the product is one formula worth memorizing:

$$(a, \vec{u})(b, \vec{w}) = \big(ab - \vec{u}\cdot\vec{w},\ \ a\vec{w} + b\vec{u} + \vec{u}\times\vec{w}\big).$$

The dot product and cross product of $\mathbb{R}^3$ are *both* hiding inside quaternion multiplication — Hamilton got there before Gibbs. Conjugate $\bar q = a - bi - cj - dk$, norm $|q|^2 = q\bar q = a^2+b^2+c^2+d^2$, inverse $q^{-1} = \bar q/|q|^2$. The norm is multiplicative, $|pq| = |p||q|$, so the **unit quaternions form a group** — and as a set they're exactly the sphere $S^3 \subset \mathbb{R}^4$.

A **pure** quaternion has $a = 0$; identify it with a vector $\vec v \in \mathbb{R}^3$. Key fact: if $\hat n$ is a *unit* pure quaternion then $\hat n^2 = -|\hat n|^2 = -1$. So $\hat n$ behaves exactly like i, which means the exponential series collapses Euler-style:

$$e^{\theta\hat n} = \cos\theta + \hat n \sin\theta.$$

Every unit quaternion is $e^{\theta\hat n}$ for some axis $\hat n$ and angle θ. That's the object we'll rotate with.

Rotation by conjugation, and the half-angle

For a unit quaternion q and a vector $\vec v$ (as a pure quaternion), define

$$\vec{v} \;\longmapsto\; q\,\vec{v}\,\bar{q}.$$

Two quick facts make this a rotation: it sends pure quaternions to pure quaternions (so the output is genuinely a vector), and it preserves length (since $|q\vec{v}\bar{q}| = |q||\vec{v}||\bar{q}| = |\vec{v}|$). A little more work shows it preserves orientation, so it lands in $SO(3)$, not just $O(3)$.

The claim is that

$$q = \cos\tfrac{\theta}{2} + \hat{n}\,\sin\tfrac{\theta}{2} \;=\; e^{(\theta/2)\hat{n}}$$

rotates by angle θ about $\hat{n}$. Watch *why the half appears*: q multiplies $\vec{v}$ on the left and $\bar{q}$ on the right. Those two "halves" each contribute $\theta/2$, and they add to θ in the vector's actual motion. (Concretely: split $\vec{v}$ into the piece along $\hat{n}$, which commutes with q so the q and $\bar{q}$ cancel and it's untouched, and the piece perpendicular to $\hat{n}$, which *anticommutes* so the two halves reinforce into a full θ.) The half-angle isn't a convention — it's the arithmetic of hitting $\vec{v}$ from both sides.

And there's the sign for free: q and $-q$ give the *same* rotation, because $(-q)\vec{v}\,\overline{(-q)} = q\vec{v}\bar{q}$ — the two minus signs annihilate. So the map

$$S^3 \longrightarrow SO(3), \qquad q \mapsto (\vec{v} \mapsto q\vec{v}\bar{q})$$

is **two-to-one**, with q and $-q$ collapsing to one rotation. Its kernel is $\{\pm 1\}$, and

$$SO(3) \;=\; S^3/\{\pm 1\} \;=\; \mathbb{RP}^3.$$

Since S^3 is simply connected, this *is* the universal double cover, and $\pi_1(SO(3)) = \mathbb{Z}_2$. A 2π rotation is a loop in $SO(3)$ that cannot be contracted: the path $q(t) = \cos\frac{t}{2} + \hat{n}\sin\frac{t}{2}$ for $t : 0 \to 2\pi$ runs from $q = 1$ to $q = -1$ — a genuine path in S^3, but its two endpoints project to the *same* rotation (the identity). Run t to 4π and you're back at $+1$; that loop contracts. This is Dirac's belt trick: give a belt a full 2π twist and you can't undo it without moving the ends, but a 4π twist shakes right out. Vectors can't tell 2π from 4π; quaternions can, and that is exactly what makes them the natural home for spin-$\frac{1}{2}$ particles.

Composition — same order as matrices

If rotation $R_1 \leftrightarrow q_1$ and $R_2 \leftrightarrow q_2$, then doing R_1 then R_2 is

$$q_2(q_1\vec{v}\bar{q}_1)\bar{q}_2 = (q_2q_1)\,\vec{v}\,\overline{(q_2q_1)},$$

so $R_2 R_1 \leftrightarrow q_2 q_1$ — the quaternion product respects the *same* order as the matrix product. No bookkeeping surprises.

The same thing in physics clothes: $SU(2)$ and spinors

Physicists write the same double cover with Pauli matrices. Encode a 3-vector $\vec{v}$ as a traceless Hermitian matrix $V = \vec{v}\cdot\vec{\sigma} = \begin{pmatrix} v_z & v_x - iv_y \\ v_x + iv_y & -v_z \end{pmatrix}$. A rotation acts by **conjugation**,

$$V \longmapsto UVU^\dagger, \qquad U \in SU(2),$$

and the rotation by θ about $\hat{n}$ is

$$U = \cos\tfrac{\theta}{2}\, I - i \sin\tfrac{\theta}{2}\,(\hat{n}\cdot\vec{\sigma}) \;=\; \exp\!\big(-i\tfrac{\theta}{2}\,\hat{n}\cdot\vec{\sigma}\big).$$

Match $-i\sigma_x, -i\sigma_y, -i\sigma_z$ with i, j, k and this *is* the unit quaternion $\cos\frac{\theta}{2} + \sin\frac{\theta}{2}\,\hat{n}$: $SU(2)$, the unit quaternions and S^3 are one object. U enters twice, so U and $-U$ give the same conjugation, and $SO(3) = SU(2)/\{\pm I\}$ again. The objects $SU(2)$ acts on directly — column vectors in $\mathbb{C}^2$ that pick up the sign under a 2π rotation — are **spinors**.

Euler Z-Y-Z, done in quaternions

Elemental rotations become one-liners:

$$q_z(\phi) = \cos\tfrac{\phi}{2} + k\sin\tfrac{\phi}{2}, \qquad q_y(\phi) = \cos\tfrac{\phi}{2} + j\sin\tfrac{\phi}{2}.$$

So $R_z(\gamma)R_y(\beta)R_z(\alpha)$ corresponds to $q = q_z(\gamma)\, q_y(\beta)\, q_z(\alpha)$. Grinding the two products (nothing but the multiplication rule, and half-angle sum identities) gives a clean closed form:

$$q = \cos\tfrac{\beta}{2}\cos\tfrac{\alpha+\gamma}{2} \;+\; \Big(\sin\tfrac{\beta}{2}\sin\tfrac{\alpha-\gamma}{2}\Big)i \;+\; \Big(\sin\tfrac{\beta}{2}\cos\tfrac{\alpha-\gamma}{2}\Big)j \;+\; \Big(\cos\tfrac{\beta}{2}\sin\tfrac{\alpha+\gamma}{2}\Big)k.$$

Two things to savor. The angles enter only through the combinations $\frac{\alpha+\gamma}{2}$ (in the scalar and k parts) and $\frac{\alpha-\gamma}{2}$ (in the i and j parts). In $SU(2)$ language the same product is the matrix

$$U = e^{-i\gamma\sigma_z/2}\, e^{-i\beta\sigma_y/2}\, e^{-i\alpha\sigma_z/2} = \begin{pmatrix} e^{-i(\alpha+\gamma)/2}\cos\frac{\beta}{2} & -e^{\,i(\alpha-\gamma)/2}\sin\frac{\beta}{2} \\ e^{-i(\alpha-\gamma)/2}\sin\frac{\beta}{2} & e^{\,i(\alpha+\gamma)/2}\cos\frac{\beta}{2} \end{pmatrix},$$

which is the **Wigner D-matrix** $D^{1/2}_{m'm}(\alpha,\beta,\gamma)$: the same Euler angles that torment gimbals

are the natural coordinates for the *representation theory* of the rotation group, and the sum and difference combinations are the seeds of the quantum numbers m, m'.

And **gimbal lock is now algebraically naked**: set $\beta = 0$ and the i, j parts vanish while the survivors depend *only* on $\alpha + \gamma$:

$$q\big|_{\beta=0} = \cos\tfrac{\alpha+\gamma}{2} + k\sin\tfrac{\alpha+\gamma}{2} = q_z(\alpha+\gamma).$$

Two knobs, one effect — the lost degree of freedom can be read straight off the formula, no Jacobian required. Quaternions don't *fix* the Euler chart's singularity; they show exactly what it is, and then offer a different chart (S^3) that doesn't have one.

Why the Euler rates blow up

Differentiate $R = R_z(\gamma)R_y(\beta)R_z(\alpha)$. The product rule for $\dot{R}R^{-1}$ with $R = ABC$ gives $\dot{A}A^{-1} + A\dot{B}B^{-1}A^{-1} + AB\dot{C}C^{-1}B^{-1}A^{-1}$, and since $\frac{d}{dt}R_{\hat{e}}(\phi)\cdot R_{\hat{e}}(\phi)^{-1} = \dot{\phi}\,[\hat{e}]_\times$, each term is just "rate times axis, carried by whatever rotations sit to its left":

$$\vec{\omega} \;=\; \dot{\gamma}\,\hat{z} \;+\; \dot{\beta}\,(R_z(\gamma)\hat{y}) \;+\; \dot{\alpha}\,(R_z(\gamma)R_y(\beta)\hat{z}).$$

That sentence is the whole geometry of a gimbal stack. The outer gimbal spins about the lab $\hat{z}$; the middle gimbal's axis has been carried around by the outer one; the inner gimbal's axis has been carried by both. Writing the three carried axes out,

$$\vec{\omega} = \underbrace{\begin{pmatrix} \sin\beta\cos\gamma & -\sin\gamma & 0 \\ \sin\beta\sin\gamma & \cos\gamma & 0 \\ \cos\beta & 0 & 1 \end{pmatrix}}_{A(\beta,\gamma)} \begin{pmatrix} \dot{\alpha} \\ \dot{\beta} \\ \dot{\gamma} \end{pmatrix}, \qquad \det A = \sin\beta.$$

The determinant is the punchline (expand along the last column; the 2×2 minor is $\sin\beta(\cos^2\gamma + \sin^2\gamma)$): the map from Euler rates to physical angular velocity **loses rank exactly at** $\beta = 0, \pi$. Inverting where you can:

$$\dot{\alpha} = \frac{\omega_x\cos\gamma + \omega_y\sin\gamma}{\sin\beta}, \qquad \dot{\beta} = -\omega_x\sin\gamma + \omega_y\cos\gamma, \qquad \dot{\gamma} = \omega_z - \cot\beta\,(\omega_x\cos\gamma + \omega_y\sin\gamma).$$

Now look at the third row of A: $\dot{\alpha}\cos\beta + \dot{\gamma} = \omega_z$. As $\beta \to 0$, $\cos\beta \to 1$, so

$$\dot{\alpha} + \dot{\gamma} \;\longrightarrow\; \omega_z \quad \text{(perfectly finite).}$$

The two divergences are equal and opposite. $\dot{\alpha} \to +\infty$ and $\dot{\gamma} \to -\infty$ with their sum staying bounded. That's the dynamical face of the algebraic fact in the quaternion formula: at $\beta = 0$ only $\alpha + \gamma$ is meaningful. Two gimbals counter-spinning like mad to produce a gentle finite motion.

Flying over the pole

Compute where the body's $\hat{z}$ (the boresight — the nose of the rocket, the telescope's pointing axis) actually goes:

$$R\,\hat{z} = (\sin\beta\cos\gamma,\ \sin\beta\sin\gamma,\ \cos\beta).$$

So β and γ **are literally the colatitude and longitude of the pointing direction**, and α is roll about that direction. Gimbal lock isn't analogous to the pole singularity of spherical coordinates — it *is* the pole singularity of spherical coordinates, with roll bolted on.

So fly over the pole. Let the vehicle slew at a constant, modest rate Ω about the fixed $\hat{y}$ axis, with the boresight on a great circle that misses the pole by angular distance ε:

$$\hat{n}(t) = (\cos\varepsilon\ \sin\Omega t,\ \sin\varepsilon,\ \cos\varepsilon\ \cos\Omega t), \qquad \vec{\omega} = \Omega\,\hat{y}\ \text{(constant, bounded, boring).}$$

Closest approach is at $t = 0$, where $\beta = \varepsilon$. The longitude is $\gamma = \operatorname{atan2}(\sin\varepsilon,\ \cos\varepsilon\sin\Omega t)$, and differentiating at $t = 0$:

$$\dot{\gamma}\big|_{t=0} = -\,\Omega\cot\varepsilon\ \approx\ -\frac{\Omega}{\varepsilon}, \qquad \text{and since } \omega_z = 0, \quad \dot{\alpha} \approx +\frac{\Omega}{\varepsilon}.$$

Put numbers on it. Slew at $\Omega = 1$ rad/s (about $57°$/s — brisk, not violent) and miss the pole by $\varepsilon = 0.01$ rad, i.e. **half a degree**. Then the guidance loop demands

$$|\dot{\gamma}| = |\dot{\alpha}| \approx 100\ \text{rad/s} \approx 5{,}700\ ^\circ\text{/s}$$

on two axes at once, in opposite directions, to accomplish a $57°$/s physical maneuver. Nothing physical is happening fast. The *coordinates* are hysterical. And if $\varepsilon = 0$ exactly, γ jumps by π instantaneously while α jumps by $-\pi$ — a genuine discontinuity in the state variables, with only $\alpha + \gamma$ surviving continuously.

The scalar version is familiar: walk one metre over the North Pole and your longitude changes by $180°$. Your body is fine; your *heading readout* is not. A control loop that closes on those numbers — commanding rates to null Euler-angle error, or numerically integrating them — will either saturate its actuators, or lag and lose track, or, in a fixed-step integrator, take a step of size $\Omega\,\Delta t/\varepsilon$ and lose the attitude entirely. The *hardware* version, physical

gimbals with finite torque, simply cannot slew fast enough; the platform gets dragged off its reference.

Contrast the quaternion. From $\dot{q} = \frac{1}{2}q\vec{\omega}$ (Problem 15), take norms — $|q| = 1$ and the quaternion norm is multiplicative:

$$|\dot{q}| = \tfrac{1}{2}|\vec{\omega}|.$$

Always. Everywhere. No exceptions, no orientations. Bounded physical rate ⇒ bounded state velocity. Said structurally: $S^3 \to SO(3)$ is a local *isometry* (up to the factor of 2), so it has zero metric distortion anywhere, while the Euler chart has unbounded distortion near the poles. That single line is the entire engineering case.

Picturing the double cover

The circle version is exact. Take the one-parameter subgroup of unit quaternions $q(\theta) = \cos\frac{\theta}{2} + k\sin\frac{\theta}{2}$; it's a circle in S^3, and it maps onto the circle of rotations about $\hat{z}$ in $SO(3)$. In the coordinate $\theta/2$ upstairs and θ downstairs, that map is exactly $e^{i\psi} \mapsto e^{2i\psi}$, i.e. $z \mapsto z^2$ on S^1. So $SU(2) \to SO(3)$ *restricted to a maximal torus is the squaring map*. Everything known about square roots — two branches, $\pm\sqrt{w}$, and going once around the origin downstairs takes you halfway around upstairs so you return on the *other* branch — is the 2π-versus-4π story verbatim. One caveat to keep the analogy honest: $\pi_1(\mathbb{C}^*) = \mathbb{Z}$, so you can keep climbing ($z \mapsto z^3$, z^4, the infinite spiral staircase of $\log z$). For $SO(3)$, $\pi_1 = \mathbb{Z}_2$ — the staircase has exactly two floors and then closes. 4π genuinely brings you home; there is no "8π sheet."

The picture to keep in your head is the solid ball. Represent a rotation by the vector $\vec{x} = \theta\hat{n}$ with $\theta \le \pi$: a point in the closed ball of radius π. The center is the identity. But rotation by π about $\hat{n}$ equals rotation by π about $-\hat{n}$, so **antipodal boundary points are glued**. That's $\mathbb{RP}^3$, and now the belt trick is something you can see.

A 2π rotation about $\hat{n}$ is the path that runs from the center out to the boundary at $\pi\hat{n}$, *reappears* at $-\pi\hat{n}$, and returns to the center. Closed loop. To contract it you'd have to pull those two boundary punctures together — but they're antipodal and stay antipodal, so you can't. Now do it **twice**: four boundary crossings, in two antipodal pairs. Slide one pair around the sphere until it meets the other pair; when they coincide, the two teleports cancel and the loop lifts off the boundary and shrinks to a point. That's a picture you can draw, and it's a real proof sketch that the generator has order exactly 2.

A warning about the Möbius strip. The right analogue is the *antipodal* double cover $S^n \to \mathbb{RP}^n$; $SO(3)$ is the $n = 3$ case. The Möbius band and $\mathbb{RP}^2$ carry the same $\mathbb{Z}_2$ covering structure, and "go around once, come back flipped; twice, restored" is exactly right. But **don't import the non-orientability**. $\mathbb{RP}^2$ is non-orientable because 2 is even; $\mathbb{RP}^3$ is perfectly orientable (every Lie group is parallelizable). The transferable content is the

antipodal gluing and the order-2 loop — not the one-sidedness.

And the version you can do right now. Hold a coffee cup on your flat palm, upright. Rotate it $360°$ (under the armpit) — cup is fine, arm is twisted. Rotate another $360°$ *the same way* (over the shoulder) — arm comes free. Your arm is a physical path-lift: the cup's orientation is the point in $SO(3)$, and your arm remembers which sheet of S^3 you're on. Vectors are the cup. Spinors are the arm.

Why the flight computer carries q instead of (α, β, γ)

No singularity. $S^3 \to SO(3)$ is a local diffeomorphism *everywhere*, so there's no attitude an integrator can't glide through. Contrast the Euler-rate kinematics above, with a $1/\sin\beta$ that detonates at gimbal lock.

Trivial kinematics. With body angular velocity $\vec{\omega}$ written as a pure quaternion, attitude evolves by

$$\dot{q} = \tfrac{1}{2}\, q\, \vec{\omega},$$

a bilinear expression — no trig, no division, cheap to integrate at kHz rates.

Cheap, stable housekeeping. Composing rotations is one 16-multiply quaternion product versus a 3×3 matrix product; correcting numerical drift is renormalizing one 4-vector, $q \to q/|q|$, versus re-orthonormalizing a matrix with Gram–Schmidt.

Geodesic interpolation (SLERP). "Blend between two orientations" becomes "walk the great circle on S^3": with $\cos\Omega = q_0 \cdot q_1$,

$$\mathrm{slerp}(q_0, q_1; t) = \frac{\sin((1-t)\Omega)}{\sin\Omega}\, q_0 + \frac{\sin(t\Omega)}{\sin\Omega}\, q_1,$$

constant angular speed, no wobble. This is why essentially every game engine stores rotations as quaternions internally and only *shows* you Euler angles in the inspector.

The famous real-world war story is **Apollo**. The inertial measurement unit tracked the spacecraft's attitude with a physical **three-gimbal** platform, and three gimbals means real, mechanical gimbal lock: certain orientations align the inner and outer gimbals and the platform loses a rotational axis, potentially tumbling the reference. A *fourth* gimbal would have prevented it (this is how you buy your way out of the topological corner in hardware), but it added weight, so NASA flew three plus software that warned the crew as they approached the forbidden cone. Michael Collins joked on Apollo 11 about wanting a fourth gimbal for Christmas. Modern spacecraft carry the attitude estimate as a quaternion precisely so that the *math* never locks even when maneuvering hard — the sign ambiguity $\pm q$ is a small bookkeeping cost next to never dividing by zero at the worst possible moment.

Summary

Euler angles are a cheap chart with an unavoidable hole; unit quaternions (spinors, $SU(2)$) are the double cover S^3 with no hole; and the half-angle and the $\pm q$ sign are the ledger entry recording that $SO(3)$ has a $\mathbb{Z}_2$ in its fundamental group.

Problems

Grouped by flavor, roughly increasing in difficulty.

Quaternion algebra (warm-ups)

1. Compute $(1+i)(1+j)$ and verify $|(1+i)(1+j)| = |1+i|\,|1+j|$.
2. Prove $\overline{pq} = \bar{q}\,\bar{p}$ (note the order-reversal). Then use it to re-derive $|pq| = |p||q|$ in one line.
3. Let $\hat{n}$ be a unit pure quaternion. Show $\hat{n}^2 = -1$, and hence that $e^{\theta\hat{n}} = \cos\theta + \hat{n}\sin\theta$.
4. Find the inverse of $q = 1+i+j+k$. What rotation (axis, angle) does $q/|q|$ represent?

Rotations by conjugation

5. Write the unit quaternion for a 90° rotation about $\hat{z}$. Apply $q\vec{v}\bar{q}$ to $\vec{v} = i$ (i.e. $\hat{x}$) and confirm you land on $\hat{y}$.
6. Show the rotation by 120° about $(1,1,1)/\sqrt{3}$ is $q = \frac{1}{2}(1+i+j+k)$. Compute q^3 — you should get -1, not $+1$. Explain why a rotation of order 3 gives a quaternion of order 6.
7. Prove that $\vec{v} \mapsto q\vec{v}\bar{q}$ sends pure quaternions to pure quaternions, for any unit q. (This is what makes it a well-defined map on $\mathbb{R}^3$.)
8. Extract the axis and angle of $q = \frac{1}{\sqrt{2}}(1+k)$ by hand, and again for $q = \frac{1}{2} + \frac{\sqrt{3}}{2}\,\frac{i+j}{\sqrt{2}}$.

Euler and gimbal lock

9. Verify the Z-Y-Z quaternion formula in two checks: set $\alpha = \gamma = 0$ and confirm you recover $q_y(\beta)$; then set $\beta = 0$ and confirm you recover $q_z(\alpha+\gamma)$.
10. Take $\alpha = \gamma = 0,\ \beta = \pi/2$ in the formula, convert the resulting quaternion back to axis-angle, and confirm it's 90° about $\hat{y}$.
11. Inverse problem: find Z-Y-Z Euler angles for a 180° rotation about $\hat{x}$. Show $\beta = \pi$ is forced and that α, γ are individually undetermined — only a combination matters. (A second face of gimbal lock.)

Double cover and topology

12. Directly from $q\vec{v}\bar{q}$, prove q and $-q$ give identical rotations. Then explain in one sentence why this makes $S^3 \to SO(3)$ exactly two-to-one rather than, say, three-to-one.
13. The belt trick as a computation: write the explicit path $q(t) = \cos\frac{t}{2} + k\sin\frac{t}{2}$, $t \in [0, 2\pi]$, list its endpoints in S^3, and explain why its image in $SO(3)$ is a loop that cannot be contracted — but the $t \in [0, 4\pi]$ version can.

14. Show the unit quaternions are closed under multiplication and inversion (hence a group), and that $\{\pm 1\}$ is its center. Why does "center" matter for the double cover being *normal*?

Kinematics and interpolation

15. For a constant body angular velocity $\vec{\omega} = \Omega\,\hat{\omega}$, the attitude is $q(t) = q_0\, e^{(\Omega t/2)\hat{\omega}}$. Differentiate to derive $\dot{q} = \frac{1}{2} q\,\vec{\omega}$. Where does the $\frac{1}{2}$ come from (same source as the half-angle)?
16. With $q_0 = 1$ and $q_1 = \frac{1}{\sqrt{2}}(1 + k)$, write out $\mathrm{slerp}(q_0, q_1; t)$ and evaluate at $t = \frac{1}{2}$. Confirm it's the 45°-about-$\hat{z}$ rotation, and check the angular speed is constant.
17. Explain why blindly SLERPing between q_0 and q_1 can take the "long way around," and how flipping to $-q_1$ when $q_0 \cdot q_1 < 0$ fixes it. What does this exploit about the double cover?

Harder / proof-flavored *(the main idea suffices unless you want full rigor)*

18. Prove $q \mapsto (\vec{v} \mapsto q\vec{v}\bar{q})$ is a group homomorphism $S^3 \to SO(3)$: check it preserves the quaternion product, preserves length, and lands in the *orientation-preserving* component (why SO, not O?).
19. Prove the map is surjective — every rotation arises. (Hint: you can hit every axis-angle pair; or use that the image is a subgroup containing rotations about all three coordinate axes.)
20. Assemble the topology: given $\pi_1(S^3) = 1$ and that $S^3 \to SO(3)$ is a 2-sheeted covering, deduce $\pi_1(SO(3)) = \mathbb{Z}_2$. Then say, in a sentence, why this $\mathbb{Z}_2$ is the reason spin-$\frac{1}{2}$ representations exist at all.

Problems 6, 13 and 20 are the ones to linger on — they're the same $\mathbb{Z}_2$ wearing three different costumes (algebra, geometry, topology). For 13 and 20, the solid-ball picture is the one to reach for.

Inequalities

Four inequalities that do most of the work in analysis, each with the picture that makes it true: Jensen (convexity), Hölder (duality), Minkowski (the shape of the unit ball) and Young (the Legendre transform). Exercises for all four are at the end.

Jensen's Inequality

The big picture in one sentence

Averaging and then bending is different from bending and then averaging — and for a convex function, bending-then-averaging always comes out bigger.

That's the whole theorem. Everything else is bookkeeping about what "averaging" means.

Picture the parabola $y = x^2$. Take two points on it and draw the chord. The chord sits *above* the curve. The midpoint of the chord is $\frac{f(a)+f(b)}{2}$; the point on the curve below it is $f\big(\frac{a+b}{2}\big)$. Jensen says this stays true for any weighted average of any number of points, and in the limit, for any probability distribution.

The statements

Finite form. Let φ be convex on an interval I, let $x_1, \dots, x_n \in I$, and let $\lambda_i \geq 0$ with $\sum \lambda_i = 1$. Then

$$\varphi\left(\sum \lambda_i x_i\right) \leq \sum \lambda_i \varphi(x_i).$$

Probabilistic form. X an integrable random variable taking values in I, φ convex:

$$\varphi(\mathbb{E}[X]) \leq \mathbb{E}[\varphi(X)].$$

Integral form. (Ω, μ) a probability space, $g \in L^1(\mu)$:

$$\varphi\left(\int_\Omega g \, d\mu\right) \leq \int_\Omega \varphi(g) \, d\mu.$$

Conditional form. For a sub-σ-algebra $\mathcal{G}$: $\varphi(\mathbb{E}[X \mid \mathcal{G}]) \leq \mathbb{E}[\varphi(X) \mid \mathcal{G}]$ a.s. — this is the version that makes martingale theory work.

Concave φ flips all inequalities. Equality holds iff X is a.s. constant, *or* φ is affine on the (essential) range of X.

The proof, in one move

A convex function lies above each of its supporting lines. Let $m = \mathbb{E}[X]$ and pick a supporting line at m: there's a slope c (any subgradient) with

$$\varphi(x) \geq \varphi(m) + c(x - m) \quad \text{for all } x \in I.$$

Substitute X and take expectations. The linear term dies because $\mathbb{E}[X - m] = 0$. Done.

That's it. The entire content of Jensen is "convex = supremum of affine functions, and expectation commutes with affine functions." Worth internalizing, because it tells you Jensen is *sharp* exactly when the supporting line is tight on the whole support.

Real-world instances

Volatility drag. An investment gains 50% one year, loses 50% the next. Arithmetic mean return: 0%. Actual outcome: $1.5 \times 0.5 = 0.75$, you lost 25%. Growth compounds multiplicatively, so what matters is $\mathbb{E}[\log R]$, and log is concave: $\mathbb{E}[\log R] \leq \log \mathbb{E}[R]$. Volatility is a pure tax on compounded growth. This one inequality explains why "average annual return" advertisements are misleading and why the Kelly criterion exists.

Why insurance exists. A risk-averse agent has concave utility u, so $\mathbb{E}[u(W)] \leq u(\mathbb{E}[W])$. The gap is exactly what she'll pay to eliminate the randomness. The insurer, holding thousands of independent policies, faces a nearly deterministic average — so the same gap is his profit. Both sides gain. Jensen's gap *is* the insurance industry.

Traffic jams and queues. In an M/M/1 queue, waiting time scales like $\frac{1}{1-\rho}$ in the utilization ρ — steeply convex. So a highway at fluctuating load 50%/90% is far worse than one steadily at 70%. Variability is expensive whenever the cost curve bends upward, which is why hospitals, servers, and airports obsess over smoothing demand rather than just reducing it.

Statistical mechanics (Gibbs–Bogoliubov). From convexity of e^x: $\langle e^{-\beta H} \rangle \geq e^{-\beta \langle H \rangle}$, which rearranges into

$$F \leq F_0 + \langle H - H_0 \rangle_0.$$

The true free energy is bounded above by what any trial Hamiltonian gives you. This is the foundation of variational mean-field theory — you get a *rigorous* bound, then minimize over trial states. Physicists have been eating off this inequality for a century.

Theoretical payoff: nonnegativity of relative entropy

One consequence deserves to be singled out, because a whole field rests on it.

For probability distributions p, q on a finite set, the Kullback–Leibler divergence is

$$D(p\|q) = \sum_i p_i \log \frac{p_i}{q_i}.$$

Claim: $D(p\|q) \geq 0$, with equality iff $p = q$.

Proof. $-\log$ is strictly convex. Treat q_i/p_i as a random variable under p:

$$-D(p\|q) = \sum_i p_i \log \frac{q_i}{p_i} \leq \log \sum_i p_i \cdot \frac{q_i}{p_i} = \log \sum_i q_i = \log 1 = 0. \qquad \blacksquare$$

Three lines. And from it:

- **The maximum entropy principle.** Take q uniform: $D(p\|q) \geq 0$ instantly gives $H(p) \leq \log n$. Uniform maximizes entropy. Same trick with q Gaussian shows the Gaussian maximizes entropy at fixed variance — which is *why* Gaussians are everywhere.
- **Mutual information** ≥ 0, since $I(X;Y) = D(p_{XY}\|p_X p_Y)$. Independence is the extremal case.
- **Shannon's source coding theorem.** Any uniquely decodable code has expected length $\geq H(p)$: the gap between your code's length and the entropy is literally $D(p\|q)$ where q is the distribution your code is secretly optimal for. Using the wrong model costs you exactly the divergence.
- **The EM algorithm and variational inference.** The ELBO, $\log p(x) \geq \mathbb{E}_q[\log p(x,z)] + H(q)$, is Jensen applied to $\log \int q \cdot \frac{p}{q}$. Half of modern machine learning is this inequality plus gradient descent.

A second theoretical payoff worth noting: on a probability space, $\|f\|_p$ is nondecreasing in p (apply Jensen to $\varphi(t) = t^{q/p}$), and conditional Jensen makes $\varphi(M_n)$ a submartingale for any convex φ — which is the entry point to Doob's maximal inequalities and L^p convergence.

Worked competition problems

1. AM–GM, for free

log is concave, so for $x_i > 0$ and weights λ_i:

$$\log\left(\sum \lambda_i x_i\right) \geq \sum \lambda_i \log x_i = \log \prod x_i^{\lambda_i}.$$

Exponentiate: $\sum \lambda_i x_i \geq \prod x_i^{\lambda_i}$. You get the *weighted* AM–GM immediately, which the usual induction proofs make you sweat for.

2. Triangle warm-up

Show $\sin A + \sin B + \sin C \leq \frac{3\sqrt{3}}{2}$ *in any triangle.*

sin is concave on $[0, \pi]$, so $\frac{1}{3}\sum \sin A \leq \sin\left(\frac{A+B+C}{3}\right) = \sin(\pi/3) = \frac{\sqrt{3}}{2}$. Equality for the equilateral triangle. The moral: whenever the constraint is on a *sum* and the function is concave, Jensen is a reflex.

3. A clean one-liner

If $x_1 + \cdots + x_n = 1$ *with* $x_i > 0$*, show* $\sum_{i=1}^{n} \left(x_i + \frac{1}{x_i}\right)^2 \geq n\left(n + \frac{1}{n}\right)^2$.

Let $f(x) = (x + 1/x)^2 = x^2 + 2 + x^{-2}$. Then $f''(x) = 2 + 6x^{-4} > 0$, so f is convex on $(0, \infty)$. Jensen with equal weights:

$$\frac{1}{n}\sum f(x_i) \geq f\left(\frac{\sum x_i}{n}\right) = f(1/n) = \left(\tfrac{1}{n} + n\right)^2.$$

Multiply by n. Note that the naive move — expand and bound each piece separately — fails, because $\sum 1/x_i^2$ isn't controlled term-by-term. Convexity handles the whole thing at once.

4. IMO 2001, Problem 2 (the good one)

For positive reals a, b, c*:*

$$\frac{a}{\sqrt{a^2 + 8bc}} + \frac{b}{\sqrt{b^2 + 8ca}} + \frac{c}{\sqrt{c^2 + 8ab}} \geq 1.$$

The trick is choosing the *weights* rather than using uniform ones. Let $s = a + b + c$ and take $\lambda_a = a/s$, etc. Apply Jensen to the convex $f(t) = t^{-1/2}$:

$$\frac{1}{s}\sum \frac{a}{\sqrt{a^2 + 8bc}} = \sum \lambda_a f(a^2 + 8bc) \geq f\left(\sum \lambda_a (a^2 + 8bc)\right).$$

Now $\sum a(a^2 + 8bc) = a^3 + b^3 + c^3 + 24abc$, so the right side is $\left(\frac{a^3+b^3+c^3+24abc}{s}\right)^{-1/2}$, and the whole inequality reduces to

$$s^3 \geq a^3 + b^3 + c^3 + 24abc.$$

Expanding $(a + b + c)^3$, this is $3\sum_{\text{sym}} a^2 b + 6abc \geq 24abc$, i.e. $\sum_{\text{sym}} a^2 b \geq 6abc$ — six terms, AM–GM, done.

The lesson worth extracting: competitive Jensen is almost never "apply it with equal weights." It's *find the weighted average that turns your expression into* $\sum \lambda_i f(x_i)$. Here, pulling a/s out as a weight rather than as part of the function is the entire problem.

5. A cautionary counterexample

Prove or disprove: for $x_i > 0$ with $\sum x_i = 1$, $\sum \frac{1}{1+x_i} \le \frac{n^2}{n+1}$.

$f(x) = 1/(1+x)$ is *convex*, so Jensen gives $\sum f(x_i) \ge n f(1/n) = \frac{n^2}{n+1}$ — the inequality points the *wrong way*, and the claim is false (push one $x_i \to 1$ and the rest to 0: the sum tends to $n - \frac{1}{2}$, which exceeds $\frac{n^2}{n+1}$ for $n \ge 2$). Convexity gives you the *lower* bound; upper bounds under a sum constraint typically come from the boundary of the simplex, not the center. Always check which direction the curve bends before writing anything down.

Worth keeping in your pocket

- **Jensen's gap** $\mathbb{E}[\varphi(X)] - \varphi(\mathbb{E}[X])$ is roughly $\frac{1}{2}\varphi''(\mu)\,\mathrm{Var}(X)$ for small variance — a second-order Taylor heuristic that tells you *how much* you're losing, not just that you're losing.
- **The tangent line trick**: when f isn't convex but equality should occur at $x = t_0$, try proving $f(x) \ge f(t_0) + f'(t_0)(x - t_0)$ pointwise and summing. It's Jensen's proof without Jensen's hypothesis, and it rescues a lot of olympiad problems where convexity fails.
- **Hermite–Hadamard** gives a reverse: $\varphi\big(\frac{a+b}{2}\big) \le \frac{1}{b-a}\int_a^b \varphi \le \frac{\varphi(a)+\varphi(b)}{2}$ — both directions sandwiched.

The natural next stop is the martingale side of conditional Jensen, where the inequality stops being a computational tool and becomes structural (exercise 7).

Hölder

$$\sum_k |a_k b_k| \;\le\; \Big(\sum_k |a_k|^p\Big)^{1/p} \Big(\sum_k |b_k|^q\Big)^{1/q}, \qquad \tfrac{1}{p} + \tfrac{1}{q} = 1,\ p > 1$$

and the same for $\int fg$, and for $\mathbb{E}|XY|$. Cauchy–Schwarz is the case $p = q = 2$.

Big picture: it's a duality statement, not an inequality

A pairing $\langle a, b\rangle = \sum a_k b_k$ takes two objects and produces a number. Hölder says: **if you budget a in the p-norm and b in the q-norm, the pairing is bounded.** And crucially it's *sharp* — for any a you can find a b achieving equality, so

$$\|a\|_p = \sup_{\|b\|_q=1} \sum a_k b_k.$$

That's the real content: **the p-norm and the q-norm are dual measuring devices.** Hölder isn't "a bound"; it's the statement that $(\ell^p)^* = \ell^q$. Everything else is corollary.

Why does $\frac{1}{p} + \frac{1}{q} = 1$ show up? Because p and q trade *height* against *spread*. Large p punishes tall spikes and forgives long flat tails; large q does the reverse. Conjugacy is exactly the balance point where the two errors cancel. Concretely: take a to be a single spike of height H and b to be flat of height h on n slots. Then $\|a\|_p = H$ regardless of p, while $\|b\|_q = h\, n^{1/q}$. Their product is $Hh\, n^{1/q}$, and the actual pairing is Hh. The slack $n^{1/q}$ is the price of mismatched shapes.

An example: effort allocation

You have n days. Day k pays at rate b_k per unit of effort. You spend effort a_k. Total payoff $\sum a_k b_k$. Your stamina constraint is $\|a\|_p \le 1$ — and the exponent p encodes *what kind of person you are*:

- $p = 1$: total effort is what's limited. Optimal strategy — dump everything on the single best day. (Indeed $\|b\|_\infty = \max b_k$.)
- $p = 2$: fatigue is quadratic. Optimal — allocate proportional to b_k. Spread out, but favour good days.
- $p = \infty$: you can only work so hard per day, no matter what. Optimal — max effort every day, payoff $\sum b_k = \|b\|_1$.

The equality condition $|a_k|^p \propto |b_k|^q$ *is* the optimal allocation rule, and it interpolates continuously between "all eggs in one basket" and "perfectly uniform." That's a lot of economics for one line of algebra.

The physics twin: instantaneous power $p(t) = v(t)i(t)$, and $\int |vi| \le \|v\|_2 \|i\|_2$ says average power is bounded by RMS voltage times RMS current — with equality iff current is proportional to voltage, i.e. **a purely resistive load**. Power factor is literally the Cauchy–Schwarz defect.

Two workhorse corollaries

Spreading bound. Take $b_k = 1$ for $k \le n$: $\sum_{k \le n} a_k \le n^{1/q} \|a\|_p$. This is how you convert "I control the p-norm" into "I control the sum," and the $n^{1/q}$ is the exact tax. It's also the power-mean inequality in disguise.

Interpolation. If $\frac{1}{r} = \frac{\theta}{p} + \frac{1-\theta}{q}$ then $\|f\|_r \le \|f\|_p^\theta \|f\|_q^{1-\theta}$. In probability this is **Lyapunov's inequality**: $\log \mathbb{E}|X|^s$ is a convex function of s. So knowing two moments pins down everything between them. This is Hölder handing you back Jensen, which should make you suspicious in a productive way.

Main idea of the proof

Normalize so $\|a\|_p = \|b\|_q = 1$ (both sides are homogeneous, so this is free). Now you only need $\sum a_k b_k \leq 1$. Apply, term by term,

$$a_k b_k \leq \frac{a_k^p}{p} + \frac{b_k^q}{q}$$

and sum: you get $\frac{1}{p} + \frac{1}{q} = 1$. Done.

That one-line pointwise inequality is Young's, which is why the chain runs the way it does. **The whole of Hölder is: normalize, then be pointwise stupid.**

Minkowski

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p, \qquad p \geq 1.$$

Big picture: this is the sentence that makes L^p a *geometry*

Without Minkowski, $\|\cdot\|_p$ is just a number you compute. With it, you have distance, hence balls, hence convergence, completeness, Banach space, Hilbert space when $p = 2$, and the whole apparatus.

And the deep content is: **the triangle inequality is the convexity of the unit ball.** They're the same statement. If $\|x\|, \|y\| \leq 1$ then $\|\frac{1}{2}(x + y)\| \leq 1$ — the midpoint stays inside.

So picture the unit balls in $\mathbb{R}^2$:

- $p = 1$: diamond (corners on the axes)
- $p = 2$: circle
- $p = \infty$: square
- $p = 1/2$: a **four-pointed star** — pinched inward, *not convex*

And that's the real reason Minkowski fails for $p < 1$: the ball has dents. $(1, 0)$ and $(0, 1)$ both have $\|\cdot\|_{1/2} = 1$, but their sum $(1, 1)$ has $\|(1, 1)\|_{1/2} = (1 + 1)^2 = 4 > 2$. Nothing subtle went wrong; the shape is just wrong.

Real-world payoff of the ball shapes. In LASSO regression you minimize error subject to an ℓ^1 budget. The ℓ^1 ball is a diamond — it has *corners on the axes* — so the optimum tends to land on a corner, i.e. with coordinates exactly zero. That's why ℓ^1 produces sparse models and ℓ^2 doesn't: the geometry of the ball, nothing else. And compressed-sensing people flirt with $p < 1$ precisely because the star is even *pointier* — better sparsity — at the cost of losing convexity and hence losing every optimization guarantee. Minkowski is exactly the line they're stepping over.

Main idea of the proof

The slick one. Write

$$\int |f+g|^p = \int |f+g|^{p-1}\,|f+g| \le \int |f+g|^{p-1}|f| + \int |f+g|^{p-1}|g|,$$

then hit each piece with Hölder using exponent $q = \frac{p}{p-1}$ on the first factor. Note $(p-1)q = p$, so $\big\||f+g|^{p-1}\big\|_q = \|f+g\|_p^{p-1}$. You get

$$\|f+g\|_p^p \le \|f+g\|_p^{p-1}(\|f\|_p + \|g\|_p),$$

and divide. The trick is entirely "split off one power of $|f+g|$ so that Hölder's conjugate exponent lands back on p" — a conjugacy fixed point.

The *conceptual* proof is the convexity one: $t \mapsto t^p$ is convex, so by Jensen the unit ball is convex, so the triangle inequality holds. Same theorem, and it makes clear that Minkowski is Jensen with a coat on.

The version you'll actually use

Minkowski's integral inequality:

$$\Big\| \int f(\cdot, y)\, dy \Big\|_p \le \int \|f(\cdot, y)\|_p \, dy.$$

"The norm of a sum of a family is at most the sum of the norms" — infinitely many terms. This is the single most-used estimate in PDE: you write a solution as a superposition (Duhamel, convolution with a kernel, integral over characteristics) and this lets you bound the superposition by bounding each piece uniformly. It's also, once again, Jensen — for the convex functional $\|\cdot\|_p$.

Young

Two theorems share the name. Both matter.

3a. Young's product inequality

$$ab \le \frac{a^p}{p} + \frac{b^q}{q}, \qquad a, b \ge 0,\ \tfrac{1}{p} + \tfrac{1}{q} = 1.$$

At $p = q = 2$ this is $ab \le \frac{a^2+b^2}{2}$, i.e. AM–GM, i.e. $(a-b)^2 \ge 0$. So it's a *deformation of AM–GM* where you're allowed to weight the two terms unequally — as long as the weights are reciprocal.

The picture is the proof. Draw $y = x^{p-1}$ (increasing, through the origin). Its inverse is $x = y^{q-1}$, because $(p-1)(q-1) = 1$ — that identity is conjugacy, and it's the entire secret. Now:

Shade the region between the curve and the x-axis for $0 \le x \le a$: its area is $\int_0^a x^{p-1}dx = a^p/p$. Shade the region between the curve and the y-axis for $0 \le y \le b$: its area is $\int_0^b y^{q-1}dy = b^q/q$. The two regions don't overlap, and together they **cover** the $a \times b$ rectangle $[0,a] \times [0,b]$. Therefore $ab \le a^p/p + b^q/q$, with equality exactly when $b = a^{p-1}$, i.e. when the rectangle's far corner sits on the curve and nothing sticks out.

The generalization falls out for free: for *any* increasing f with $f(0) = 0$,

$$ab \le \int_0^a f(x)\,dx + \int_0^b f^{-1}(y)\,dy.$$

Power functions were never essential. That's the mark of a real proof.

The big picture: this is the Legendre transform

Let $\Phi(a) = a^p/p$. Then $\Phi^*(b) = b^q/q$ is precisely its **Legendre–Fenchel conjugate**, and Young's inequality is the general **Fenchel–Young inequality**

$$\langle x, y\rangle \le \Phi(x) + \Phi^*(y),$$

true for every convex Φ, with equality iff $y = \nabla\Phi(x)$.

The same object appears in mechanics. With $\Phi = L$ the Lagrangian in velocity and $\Phi^* = H$ the Hamiltonian in momentum:

$$p\,v \le L(v) + H(p),$$

equality exactly when $p = \partial L/\partial v$ — the physical relation between momentum and velocity. Young's inequality is the statement that **the Legendre transform is a duality, and the transform is only sharp along the physical trajectory.** Same structure in thermodynamics (U and F, entropy and temperature), same in convex optimization (weak duality; the gap is the part of the two shaded regions outside the rectangle).

So: Hölder is ℓ^p-ℓ^q duality, and it's proved by the pointwise inequality that expresses Legendre duality. Duality all the way down.

The version you'll use every day: ε-Young

$$ab \le \frac{\varepsilon a^2}{2} + \frac{b^2}{2\varepsilon} \qquad (\varepsilon > 0).$$

This is the single most-used inequality in PDE. You're estimating something and a bad cross-term $\int u \nabla u$ appears. You split it with ε-Young, choose ε small enough that the $\varepsilon\|\nabla u\|^2$ piece is *absorbed* into a good term you already have on the left-hand side, and you pay for it with a large-but-harmless $C_\varepsilon\|u\|^2$. That's how energy estimates, Gårding inequalities, and most a-priori bounds are actually made. The scaling $a \to \varepsilon a$, $b \to b/\varepsilon$ leaving ab invariant is the whole trick.

3b. Young's convolution inequality

$$\|f * g\|_r \le \|f\|_p \|g\|_q, \qquad \tfrac{1}{p} + \tfrac{1}{q} = 1 + \tfrac{1}{r}.$$

Big picture: blurring cannot concentrate. Convolution is smearing f with the profile g. The inequality says the smeared object is no more concentrated than its ingredients — and the bookkeeping exponent $1 + \frac{1}{r}$ tracks exactly how much smoothing you bought.

Two limiting cases tell the story: - $q = 1, r = p$: convolving with a probability density (mass 1) doesn't increase any L^p norm at all. Averaging never sharpens. - $p = q = 2, r = \infty$: $|f * g| \le \|f\|_2 \|g\|_2$ pointwise — that's Cauchy–Schwarz.

Real-world: in probability, the density of $X + Y$ for independent X, Y is $f_X * f_Y$, and Young says the sum's density is flatter than either summand's. Add enough of them and you flatten your way to the CLT. In imaging, g is the point-spread function of your lens; Young quantifies why you can never recover a sharper image than the blur allows.

Exercises

Difficulty: ∘ easy, ∘∘ medium, ● hard.

Jensen

1. ∘ For $x_1, \dots, x_n > 0$, prove harmonic mean $\le$ geometric mean $\le$ arithmetic mean. (Jensen on log gives one inequality; apply it to $1/x_i$ for the other.)
2. ∘ On a probability space, show $\|f\|_p \le \|f\|_q$ for $1 \le p \le q$, using Jensen with $\varphi(t) = t^{q/p}$. Find a measure space where this fails.
3. ∘ If $\mathbb{E}[X] = 0$, show $\mathbb{E}[e^X] \ge 1$. When does equality hold?
4. ∘∘ For $X \sim N(\mu, \sigma^2)$, compute the Jensen gap $\mathbb{E}[e^X] - e^\mu$ exactly and compare it with the heuristic $\frac{1}{2}\varphi''(\mu)\,\mathrm{Var}(X)$. When is the heuristic good?

5. ∘∘ Volatility drag. Each period an investment returns $1+s$ or $1-s$ with probability $\frac{1}{2}$ each. Show the long-run growth rate per period is $\frac{1}{2}\log(1-s^2) \approx -s^2/2$, although the mean return is 0.
6. ∘∘ Prove $D(p\|q) \geq 0$ again, using only the tangent-line inequality $\log x \leq x-1$. Which step of the Jensen proof does that inequality replace?
7. • Let (M_n) be a martingale and φ convex with $\mathbb{E}|\varphi(M_n)| < \infty$. Use conditional Jensen to show $\varphi(M_n)$ is a submartingale. Deduce that $|M_n|$ and M_n^2 (when integrable) are submartingales.

Hölder

8. ∘ State the equality case of Hölder. For $a = (1,2,3)$ and $p = 3$, find the b with $\|b\|_{3/2} = 1$ that attains $\sum a_k b_k = \|a\|_3$.
9. ∘ Show $\left(\sum_{k=1}^n a_k\right)^2 \leq n\sum_{k=1}^n a_k^2$, and identify it as the spreading bound with $p = q = 2$.
10. ∘∘ Prove the interpolation inequality $\|f\|_r \leq \|f\|_p^{\theta}\|f\|_q^{1-\theta}$, where $\frac{1}{r} = \frac{\theta}{p} + \frac{1-\theta}{q}$. (Hint: write $|f|^r = |f|^{\theta r}|f|^{(1-\theta)r}$ and apply Hölder with exponents $\frac{p}{\theta r}$ and $\frac{q}{(1-\theta)r}$.)
11. ∘∘ Deduce Lyapunov's inequality from exercise 10: $s \mapsto \log \mathbb{E}|X|^s$ is convex on $[0,\infty)$.

Minkowski

12. ∘∘ For $0 < p < 1$, show $\|x+y\|_p \leq 2^{1/p-1}(\|x\|_p + \|y\|_p)$ on $\mathbb{R}^n$, and that $x = e_1$, $y = e_2$ attain the constant. (So $\|\cdot\|_p$ is a quasi-norm: Minkowski fails by a bounded factor, not arbitrarily.)
13. ∘∘ Use Minkowski's integral inequality to prove $\|f * g\|_p \leq \|f\|_p\|g\|_1$, the $q = 1$ case of Young's convolution inequality.

Young and Legendre

14. ∘ Prove $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$ a third way: from the weighted AM–GM inequality, i.e. concavity of log.
15. ∘ Compute the Legendre transforms of $\Phi(x) = x^p/p$ on $[0,\infty)$ and of $\Phi(x) = e^x$ on $\mathbb{R}$, and write out Fenchel–Young for each.
16. ∘∘ In ε-Young, $ab \leq \frac{\varepsilon a^2}{2} + \frac{b^2}{2\varepsilon}$, minimize the right side over $\varepsilon > 0$. Show the minimum is exactly ab, so the family of bounds is sharp.
17. • Relativistic Legendre transform. For $L(v) = -mc^2\sqrt{1-v^2/c^2}$ on $|v| < c$, compute $H(p) = \sup_v (pv - L(v))$, show $H(p) = \sqrt{p^2c^2 + m^2c^4}$, and find the equality case of $pv \leq L(v) + H(p)$. Which physical relation between momentum and velocity is it?

Newton's Binomial Series

Newton's form of the binomial theorem, for any rational exponent m/n:

$$(P+PQ)^{m/n} = P^{m/n} + \frac{m}{n}AQ + \frac{m-n}{2n}BQ + \frac{m-2n}{3n}CQ + \frac{m-3n}{4n}DQ + \cdots$$

where A is the first term, B the second, C the third — each letter is *whatever you just computed*.

What Newton actually did that everyone glosses over

Modern books write $\binom{\alpha}{k}$ and move on. Newton wrote a **recurrence**, and it's the better object. Set $\alpha = m/n$, $x = Q$, $T_0 = P^\alpha$:

$$\boxed{T_{k+1} = T_k \cdot x \cdot \frac{\alpha - k}{k+1}}$$

No factorials. No binomial coefficients. No catastrophic cancellation from computing huge numerator and denominator separately. **One multiply and one divide per term, forever.** And the $(P+PQ)$ packaging is Newton building in the normalization: you factor out the dominant part so that Q is the small quantity and convergence is fast.

That ratio $\frac{\alpha-k}{k+1}$ is the whole theorem. Everything below is a consequence of noticing it.

The numerics, done properly

$\sqrt{10}$: write $10 = 9 + 9\cdot\frac{1}{9}$, so $P = 9$, $Q = \frac{1}{9}$, $\alpha = \frac{1}{2}$.

k	T_k	partial sum
0	3	3
1	$3 \cdot \frac{1}{9} \cdot \frac{1}{2} = \frac{1}{6}$	3.1666667

k	T_k	partial sum
2	$\frac{1}{6}\cdot\frac{1}{9}\cdot\frac{-1/2}{2} = -\frac{1}{216}$	3.1620370
3	$-\frac{1}{216}\cdot\frac{1}{9}\cdot\frac{-3/2}{3} = \frac{1}{3888}$	3.1622942
4	$-\frac{5}{279936}$	3.1622763

True value 3.1622777. Four terms, six digits, and each term is *one keystroke sequence* from the last. This is why it was a calculating tool for 250 years.

The coding application

```
def newton_binomial(P, Q, alpha, terms=15):
    T = P ** alpha
    total, k = T, 0
    while k < terms and abs(T) > 1e-16 * abs(total):
        T *= Q * (alpha - k) / (k + 1)
        total += T
        k += 1
    return total
```

Notice what this *doesn't* do: it never forms a factorial, never overflows, never divides two large numbers to get a small one. A closed form built from factorials overflows double precision at $k = 171$ ($170! \approx 7 \times 10^{306}$), long after the *terms themselves* have become negligible — the recurrence never sees the intermediate monsters. This is the same reason you compute e^x by $T_{k+1} = T_k x/(k+1)$ rather than $x^k/k!$, and the same reason Horner beats naive polynomial evaluation. **Ratio recurrences are the numerically stable representation of a series; closed forms are the unstable one.**

For $x > 0$ the series alternates once $k > \alpha$, so once the terms are decreasing `abs(T)` is also a rigorous error bound (exercise 7). Free adaptive termination.

Generating functions — Catalan numbers fall out

Take $\alpha = \frac{1}{2}, Q = -4x$:

$$\sqrt{1-4x} = 1 - 2x - 2x^2 - 4x^3 - 10x^4 - 28x^5 - \cdots$$

Those coefficients are $-2C_n$ where $C_n = 1, 1, 2, 5, 14, 42$ are the **Catalan numbers**. So $C(x) = \frac{1-\sqrt{1-4x}}{2x}$, the generating function for binary trees, balanced parentheses, triangulations, Dyck paths — and Newton's recurrence generates them with three arithmetic operations apiece:

```
T, cat = 1.0, []
for k in range(10):
    T *= (-4) * (0.5 - k) / (k + 1)   # coefficient of x^(k+1) in sqrt(1-4x)
    cat.append(-T / 2)                 # -> 1, 1, 2, 5, 14, 42, ...
```

Use `Fraction` instead of floats and it's exact. The point is that the *combinatorial* recurrence and the *analytic* recurrence are the same recurrence.

Physics: the correction terms are the physics

Relativistic energy, $\gamma = (1 - v^2/c^2)^{-1/2}$, so $\alpha = -\frac{1}{2}$, $Q = -v^2/c^2$:

$$E = mc^2\left[1 + \tfrac{1}{2}\tfrac{v^2}{c^2} + \tfrac{3}{8}\tfrac{v^4}{c^4} + \tfrac{5}{16}\tfrac{v^6}{c^6} + \cdots\right] = mc^2 + \tfrac{1}{2}mv^2 + \tfrac{3}{8}\tfrac{mv^4}{c^2} + \cdots$$

Rest energy, then Newtonian kinetic energy, then the first relativistic correction — and the recurrence hands you as many corrections as you want without ever writing a binomial coefficient. Identically, $(1+r/R)^{-2}$ generates the tidal-force expansion, and $(1-2\frac{r}{R}\cos\theta+\frac{r^2}{R^2})^{-1/2}$ generates the multipole expansion: regroup Newton's terms by powers of r/R and the coefficients are the Legendre polynomials $P_\ell(\cos\theta)$ — monopole, dipole, quadrupole. (Not term by term: with $Q = -2\frac{r}{R}\cos\theta + \frac{r^2}{R^2}$, Newton's T_1 already contains an r^2/R^2 piece of the quadrupole.)

Newton's own killer app: term-by-term integration

This is what he was really after. Expand $(1 - x^2)^{1/2}$, integrate term by term, and you have the area of a circular segment — hence π. Newton got 15 decimal places this way and later wrote, a bit sheepishly, that he was ashamed to admit how far he'd carried the computation. Same move on $(1 + x)^{-1}$ gives $\log(1 + x)$; on $(1 - x^2)^{-1/2}$ gives arcsin. **The binomial series is what made calculus computational** — before it, you could define an integral but rarely evaluate one.

The big-picture reveal

Look again at $\dfrac{T_{k+1}}{T_k} = \dfrac{(\alpha - k)}{(k + 1)}\,x$. That's a **rational function of** k.

That property is the definition of a hypergeometric series. Newton's notation isn't a quaint 17th-century convenience — it's the ur-form of ${}_2F_1$, and it identifies the exact class of series that modern computer algebra can handle algorithmically: Gosper's algorithm, Zeilberger's creative telescoping, the whole "$A = B$" machinery for proving combinatorial identities automatically. All of it operates on the term ratio, never on the terms.

Newton wrote down the ratio because it was convenient for hand computation. Three centuries later it turned out that *the ratio is the invariant object* and the closed form is the derived one. That's the best kind of story.

Exercises

Difficulty: ◦ easy, ◦◦ medium, ● hard.

1. ◦ Show that Newton's recurrence $T_{k+1} = T_k \, x \, \frac{\alpha-k}{k+1}$ with $T_0 = P^\alpha$ gives $T_k = \binom{\alpha}{k} P^\alpha x^k$, where $\binom{\alpha}{k} = \frac{\alpha(\alpha-1)\cdots(\alpha-k+1)}{k!}$.
2. ◦ Show the series terminates exactly when α is a nonnegative integer, and point to the factor in the recurrence that makes it stop.
3. ◦ Compute $\sqrt[3]{30}$ to six digits by hand from $30 = 27\left(1+\frac{1}{9}\right)$. How many terms did you need?
4. ◦◦ Extend the relativistic energy expansion one term past $\frac{5}{16}\frac{v^6}{c^6}$ using only the recurrence. (You should get $\frac{35}{128}$.)
5. ◦◦ From $\sqrt{1-4x}$, show the Catalan numbers are $C_n = \frac{1}{n+1}\binom{2n}{n}$.
6. ◦◦ Integrate $(1-x^2)^{-1/2}$ term by term to get the series for $\arcsin x$, then use $\arcsin\frac{1}{2} = \frac{\pi}{6}$ to compute π to six digits. How many terms?
7. ◦◦ For $x > 0$ and $k > \alpha$ the terms alternate in sign. Prove that once $|T_k|$ is decreasing, the error in the partial sum is at most the first omitted term — so the `abs(T)` test in the code is a rigorous stopping rule.
8. ● Newton's own computation. On the circle of diameter 1, $y = \sqrt{x - x^2}$, show that

$$\pi = \frac{3\sqrt{3}}{4} + 24\int_0^{1/4} \sqrt{x-x^2}\,dx$$

   (draw the sector and the triangle). Expand $\sqrt{x}\,(1-x)^{1/2}$ with the binomial series, integrate term by term, and see how many terms give 10 digits of π.
9. ● A series $\sum t_k$ is hypergeometric when t_{k+1}/t_k is a rational function of k. Show that $(1+x)^\alpha = {}_1F_0(-\alpha;\,;-x)$, where ${}_1F_0(a;\,;z) = \sum_k \frac{(a)_k}{k!}z^k$ and $(a)_k = a(a+1)\cdots(a+k-1)$. Then write $\arcsin x$ (exercise 6) as $x \cdot {}_2F_1\left(\frac{1}{2},\frac{1}{2};\frac{3}{2};x^2\right)$.

Tensors

A bridge from convolution

The Convolution chapter defined $\mu * \nu$ as the pushforward of $\mu \otimes \nu$. That $\otimes$ is not an object on the dual group; it is the product measure on $G \times G$:

$$(u \otimes v)(x, y) = u(x)\, v(y).$$

That's a function of *two* variables. Convolution is then: form $u \otimes v$ on $G \times G$, then push it forward along the multiplication map $m : G \times G \to G$. All the "flip" business, $\int u(x - y)v(y)\, dy$, is just Fubini bookkeeping for that pushforward.

$$G \times G \xrightarrow{m} G, \qquad u * v = m_*(u \otimes v).$$

In probability language: $u \otimes v$ is the joint law of two **independent** variables (independence = the joint density factors = the joint is a pure tensor), and $u * v$ is the law of their sum.

A *different*, also-true fact does live on the dual: $\widehat{G \times H} \cong \widehat{G} \times \widehat{H}$, because every character of $G \times H$ factors as $\chi_1 \otimes \chi_2$. That's the one-line reason $\widehat{u * v} = \hat{u}\, \hat{v}$.

Tensor product: the machine that eats bilinearity

The big picture in one sentence: $V \otimes W$ **is the universal place where bilinear maps become linear.**

Bilinear maps are annoying — they're not linear, so all your linear algebra machinery stalls. So build a space $V \otimes W$ with a canonical bilinear map $(v, w) \mapsto v \otimes w$ such that every bilinear $B : V \times W \to Z$ factors uniquely:

$$B(v, w) = \tilde{B}(v \otimes w), \qquad \tilde{B} \text{ linear}.$$

Everything else is consequence. $\dim(V \otimes W) = mn$, basis $\{e_i \otimes f_j\}$, and the only rules are bilinearity: $(av + v') \otimes w = a(v \otimes w) + v' \otimes w$, etc.

The single most important fact: not every element of $V \otimes W$ is of the form $v \otimes w$. Generic elements are sums, $\sum_i v_i \otimes w_i$, and the minimal number of terms is the **rank**. This is not a technicality — it's the whole subject:

- $e_1 \otimes e_1 + e_2 \otimes e_2$ is rank 2. In quantum mechanics that's an **entangled state**; pure tensors are the separable ones. Entanglement *is* tensor rank > 1.
- In probability: pure tensor = independent. Rank > 1 = correlated.
- Matrices: $V^* \otimes W \cong \mathrm{Hom}(V, W)$, and tensor rank = matrix rank. Rank-1 matrices are exactly vw^T. SVD is the statement that you can always write a rank-r tensor as r orthogonal pure tensors.
- Functions: finite sums $\sum f_i(x)g_i(y)$ are dense in $L^2(X \times Y)$ — i.e. $L^2(X)\hat{\otimes}L^2(Y) \cong L^2(X \times Y)$. Separation of variables is a tensor-rank-1 ansatz.

Symmetric and alternating: quotient the tensor algebra

Form $T(V) = \bigoplus_k V^{\otimes k}$. Then:

$$\mathrm{Sym}^k V = V^{\otimes k}/\langle v \otimes w - w \otimes v \rangle, \qquad \Lambda^k V = V^{\otimes k}/\langle v \otimes v \rangle.$$

In characteristic 0 you can equivalently realize them as *subspaces*, via the projectors $\mathrm{Sym} = \frac{1}{k!}\sum_\sigma \sigma$ and $\mathrm{Alt} = \frac{1}{k!}\sum_\sigma \mathrm{sgn}(\sigma)\,\sigma$. That's the familiar symmetrizing and antisymmetrizing. The quotient definition is better because it works over any ring and doesn't hide factorials.

Note $\langle v \otimes v \rangle$ is the right relation, not $\langle v \otimes w + w \otimes v \rangle$ — they agree away from characteristic 2, and the squares version is the honest one.

Dimensions: $\dim \Lambda^k V = \binom{n}{k}$, $\dim \mathrm{Sym}^k V = \binom{n+k-1}{k}$. The first is finite and vanishes for $k > n$; the second grows forever. Physics knows this: fermions live in Λ (Pauli exclusion *is literally* $v \wedge v = 0$ — you can't put two electrons in the same state), bosons live in Sym (pile in as many as you like, hence Bose–Einstein condensates and lasers).

Wedge: why determinants exist

$\Lambda^n V$ is **one-dimensional**. That's the whole theorem. Given $T : V \to V$, the induced map $\Lambda^n T$ on a 1-dimensional space is multiplication by a scalar — and that scalar is $\det T$. Multiplicativity $\det(ST) = \det S \det T$ is then functoriality, free. No cofactor expansions, no sign chasing.

Geometrically: $v_1 \wedge \cdots \wedge v_k$ is the **oriented k-volume** of the parallelepiped they span. Antisymmetry is exactly "swap two edges, flip orientation," and $v \wedge v = 0$ is "a degenerate box has no volume." Decomposable k-vectors, up to scale, correspond to k-dimensional subspaces (the Plücker embedding of the Grassmannian) — and, exactly as with tensor rank, most k-vectors are *not* decomposable once $2 \le k \le n-2$.

Convention warning: with the "determinant convention," $\alpha \wedge \beta = \frac{(k+l)!}{k!\,l!}\mathrm{Alt}(\alpha \otimes \beta)$, which makes $dx \wedge dy$ evaluate to the determinant $\det \begin{pmatrix} u_1 & v_1 \\ u_2 & v_2 \end{pmatrix}$. Spivak uses this; Kobayashi–Nomizu drops the factorial. Different books, different stray $1/k!$'s. Always check.

Graded commutativity: $\alpha \wedge \beta = (-1)^{kl}\beta \wedge \alpha$. So odd-degree forms anticommute, even-degree ones are central.

Why forms are what you integrate

Under a change of variables, a k-dimensional volume element picks up a Jacobian **determinant**. Alternating tensors are precisely the objects whose transformation law *is* that determinant. So k-forms are the natural integrands over k-dimensional surfaces — the change-of-variables formula is built into the algebra rather than bolted on.

- $\Lambda^n T^*M$ is a line bundle; a nowhere-zero section is a **volume form**, and it exists iff M is orientable. (Möbius strip: no.)
- On a Riemannian manifold, $\sqrt{\det g}\, dx^1 \wedge \cdots \wedge dx^n$ is the canonical one.
- On a Kähler manifold: ω is a 2-form, and $\omega^n/n!$ is the volume form. That $\int_M \omega^n > 0$ is what forces the Kähler class to be nontrivial — hence no Kähler structure on S^6, etc. It's all wedge products doing the work.
- Stokes: $\int_M d\omega = \int_{\partial M} \omega$. One line containing the fundamental theorem of calculus, Green, Gauss, and Stokes-classical.

The physics: the electromagnetic field is a 2-form $F = dA$, and Maxwell's four equations are $dF = 0$ and $d \star F = J$. The first is an identity ($d^2 = 0$); the second is the dynamics. Also: angular momentum is really the bivector $x \wedge p$, not a vector. The cross product is a 3-dimensional accident (where $\Lambda^2\mathbb{R}^3 \cong \mathbb{R}^3$); the wedge is what generalizes, which is why in 4D relativity you get an antisymmetric tensor $L^{\mu\nu}$ and not an "angular momentum vector."

Tensors on a smooth manifold

The tangent space

The modern definition avoids "arrows sticking out of the manifold," since there's no ambient space. Define

$$T_pM = \{\text{derivations at } p\} = \{X : C^\infty(M) \to \mathbb{R} \mid X(fg) = f(p)Xg + g(p)Xf\}.$$

A tangent vector *is* a directional derivative. In a chart $(x^1, \dots, x^n)$ this space has basis $\partial/\partial x^i|_p$, so $\dim T_pM = n$. The dual T_p^*M has basis $dx^i|_p$, characterized by $dx^i(\partial_j) = \delta^i_j$.

$T^{(p,q)}$

$$T^{(p,q)}(V) \;=\; \underbrace{V \otimes \cdots \otimes V}_{p} \otimes \underbrace{V^* \otimes \cdots \otimes V^*}_{q} \;\cong\; \{\text{multilinear } \underbrace{V^* \times \cdots \times V^*}_{p} \times \underbrace{V \times \cdots \times V}_{q} \to \mathbb{R}\}.$$

With $V = T_pM$: p is the **contravariant** rank, q the **covariant** rank. Components $T^{i_1 \cdots i_p}_{\;\;j_1 \cdots j_q}$ — upper indices contravariant, lower covariant. Dimension n^{p+q}.

The vocabulary is historical and feels backwards. Under a change of chart with Jacobian J, *basis vectors* ∂_i transform by J, while the *components* of a vector transform by J^{-1} — "contra" to the basis. Covector components go along with J, hence "co." Mnemonic that actually helps: **upper indices transform like** dx**, lower like** $\partial/\partial x$, and the Einstein summation convention only ever pairs one of each — because that pairing is the one that's coordinate-independent.

Familiar cases: - $(1,0)$: vector field - $(0,1)$: differential 1-form - $(0,2)$ symmetric positive-definite: **Riemannian metric** g - $(0,2)$ antisymmetric, closed, nondegenerate: **symplectic form** - $(1,1)$: an endomorphism field — e.g. the complex structure J with $J^2 = -\mathrm{id}$, which is where Kähler geometry starts (g, ω, J, any two determining the third) - $(1,3)$: the Riemann curvature tensor $R^i_{\;jkl}$ - $(0,k)$ totally antisymmetric: k**-form**, i.e. $\Lambda^k T^*M$

Bundles and sections

Glue the fibers: $T^{(p,q)}M = \bigsqcup_{x \in M} T^{(p,q)}(T_xM)$, a smooth vector bundle over M of rank n^{p+q}.

A **section** of a bundle $\pi : E \to M$ is a smooth $s : M \to E$ with $\pi \circ s = \mathrm{id}_M$ — a smoothly varying choice of one element from each fiber. Write $\Gamma(E)$. So:

> A **tensor field** of type (p,q) is an element of $\Gamma(T^{(p,q)}M)$.

$\Gamma(E)$ is a module over $C^\infty(M)$, not just a vector space, and that's the key structural fact. It gives you the **tensor characterization lemma**: a map $\Gamma(TM)^q \to C^\infty(M)$ comes from a tensor field iff it is $C^\infty(M)$-multilinear (not merely $\mathbb{R}$-multilinear). Being a tensor means

being *pointwise*. This is how you tell tensors from differential operators: the Lie bracket $[X, Y]$ and the connection $\nabla_X Y$ are not tensors in Y (Leibniz terms appear), but curvature $R(X, Y)Z$ miraculously *is* — all the derivative terms cancel. Sections also can't always be nonvanishing: TS^2 has no nowhere-zero section (hairy ball theorem), which is exactly why global sections detect topology.

Pushforward and pullback

Let $F : M \to N$ be smooth. Its differential at p,

$$dF_p : T_pM \to T_{F(p)}N, \qquad (dF_pX)(f) = X(f \circ F),$$

is the **pushforward** — the coordinate-free chain rule. Indeed $d(G \circ F)_p = dG_{F(p)} \circ dF_p$ *is* the chain rule, and its matrix in charts is the Jacobian.

Dualizing reverses arrows: the **pullback**

$$F^* : T^*_{F(p)}N \to T^*_pM, \qquad (F^*\alpha)(X) = \alpha(dF_pX).$$

Now the crucial asymmetry:

- **Covariant tensors always pull back.** $(F^*T)(X_1, \dots, X_q) = T(dFX_1, \dots, dFX_q)$, defined pointwise, no conditions on F whatsoever. So metrics, forms, and symplectic structures pull back freely — this is how you induce a metric on a submanifold.
- **Contravariant tensors do not push forward** in general. If F isn't injective, two points of M send vectors to the same point of N and disagree; if it isn't surjective, most of N gets nothing. You need F to be a **diffeomorphism** — and then mixed (p, q) tensors transport too, via dF and $(dF)^{-1}$.

That asymmetry is the entire reason differential forms are the well-behaved objects. And the two facts that make forms *the* right category:

$$F^*(\alpha \wedge \beta) = F^*\alpha \wedge F^*\beta, \qquad \boxed{F^*(d\omega) = d(F^*\omega)}.$$

d commutes with pullback. Combined with $(G \circ F)^* = F^* \circ G^*$, we get: $\Omega^\bullet$ is a **contravariant functor** from manifolds to cochain complexes. Everything below follows from that sentence.

Homology and cohomology

de Rham

$d : \Omega^k(M) \to \Omega^{k+1}(M)$ with $d \circ d = 0$, so $\operatorname{im} d \subseteq \ker d$, and

$$H^k_{dR}(M) = \frac{\ker(d : \Omega^k \to \Omega^{k+1})}{\operatorname{im}(d : \Omega^{k-1} \to \Omega^k)} = \frac{\text{closed}}{\text{exact}}.$$

The Poincaré lemma says that on a star-shaped open set, closed $\Rightarrow$ exact ($H^k = 0$ for $k > 0$). So de Rham cohomology is identically zero locally. Therefore *every class it produces is a global obstruction* — it measures holes, not local geometry. Exercise 8 is the prototype: ω is locally $d\theta$, but θ isn't globally defined, and $H^1(\mathbb{R}^2 \setminus 0) = \mathbb{R}$ detects the puncture.

Quick catalog: $H^0(M) = \mathbb{R}^{\#\text{components}}$ (closed 0-forms are locally constant). $H^k(S^n) = \mathbb{R}$ for $k = 0, n$, else 0. $H^1(T^2) = \mathbb{R}^2$, generated by $d\theta_1, d\theta_2$. And because F^* commutes with d, cohomology is a functor — moreover it's **homotopy invariant**: homotopy equivalent manifolds have isomorphic H^*_{dR}. A construction defined by calculus turns out to see only topology.

Homology, and why the two names

Dually, build **singular chains**: $C_k(M)$ = free abelian group on smooth maps $\sigma : \Delta^k \to M$, with boundary $\partial : C_k \to C_{k-1}$ the alternating sum of faces. Then $\partial^2 = 0$ (the boundary of a boundary is empty — a solid triangle's boundary is a loop, which has no endpoints), and

$$H_k(M) = \frac{\ker \partial}{\operatorname{im} \partial} = \frac{\text{cycles}}{\text{boundaries}}.$$

Cycles that don't bound = holes you can see directly.

Note the direction: ∂ lowers degree and pushes *forward* along maps (covariant); d raises degree and pulls *back* (contravariant). Hence homo- versus co-homology.

Stokes is the bridge

$$\int_M d\omega = \int_{\partial M} \omega.$$

Pair a closed k-form with a k-cycle by $\langle \omega, c \rangle = \int_c \omega$. Stokes makes this well-defined on classes, twice over: - c a cycle, $\omega = d\eta$ exact $\Rightarrow \int_c d\eta = \int_{\partial c} \eta = 0$. - ω closed, $c = \partial b$ a boundary $\Rightarrow \int_{\partial b} \omega = \int_b d\omega = 0$.

de Rham's theorem: this pairing is perfect, so $H^k_{dR}(M) \cong H^k_{\text{sing}}(M; \mathbb{R})$. An analytic invariant (solutions of PDEs, essentially) equals a combinatorial-topological one. The integral

of exercise 8, $\int_{S^1}\omega = 2\pi$, is literally this pairing evaluating the generator of H^1 against the generator of H_1, and getting a nonzero number is the proof both are nonzero.

Two more structures worth naming: - **Ring structure.** $\wedge$ descends to cohomology, making $H^*_{dR}(M)$ a graded-commutative ring (matching cup product). This distinguishes spaces that have the same Betti numbers — e.g. $S^2 \times S^4$ vs. $\mathbb{CP}^3$. - **Poincaré duality.** On a closed oriented n-manifold, $H^k \cong H^{n-k}$ via the pairing $(\omega, \eta) \mapsto \int_M \omega \wedge \eta$.

Hodge theory

Put a metric on the manifold. **Hodge theory** says each class has a unique harmonic representative: $H^k_{dR} \cong \mathcal{H}^k = \ker \Delta$, $\Delta = dd^* + d^*d$. Topology now has a *canonical* analytic representative — a PDE existence theorem in disguise.

On a **Kähler** manifold this explodes into structure: $\Delta_d = 2\Delta_{\bar\partial}$, which forces the Hodge decomposition

$$H^k(X;\mathbb{C}) = \bigoplus_{p+q=k} H^{p,q}(X), \qquad \overline{H^{p,q}} = H^{q,p}.$$

Consequences: $h^{p,q} = h^{q,p}$, so odd Betti numbers are even; ω^k is closed and nonzero in H^{2k}, so $b_{2k} \geq 1$; hence S^6 admits no Kähler structure. All of it is d, $\wedge$, and the tensor algebra above, wearing a complex structure.

Exercises

Ordered roughly by difficulty within each part.

Linear algebra and forms on $\mathbb{R}^n$

1. (Rank.) In $\mathbb{R}^2 \otimes \mathbb{R}^2$, show $t = e_1 \otimes e_2 + e_2 \otimes e_1$ is not a pure tensor $v \otimes w$. Then show every element of $\mathbb{R}^2 \otimes \mathbb{R}^2$ has rank ≤ 2. *(Hint: identify with 2×2 matrices and think about what "rank 1" means there.)*

2. (Grad, curl, div are one operator.) On $\mathbb{R}^3$ identify $\Lambda^0 \cong C^\infty$, $\Lambda^1 \cong \Lambda^2 \cong$ vector fields (via $dx \mapsto e_1$ and $dy \wedge dz \mapsto e_1$, cyclically), $\Lambda^3 \cong C^\infty$. Show d becomes grad, curl, div in degrees $0, 1, 2$. Conclude that $d^2 = 0$ is exactly curl $\circ$ grad $= 0$ and div $\circ$ curl $= 0$.

3. (Compute.) Let $\omega = xy\,dx + yz\,dy + zx\,dz$ and $\eta = z\,dx \wedge dy$ on $\mathbb{R}^3$. Compute $d\omega$, $d\eta$, $\omega \wedge \eta$, and verify $d(\omega \wedge \eta) = d\omega \wedge \eta - \omega \wedge d\eta$ by direct computation.

4. (Determinant, conceptually.) For linear $T : \mathbb{R}^n \to \mathbb{R}^n$, show the induced map $\Lambda^n T^*$ on the one-dimensional space $\Lambda^n(\mathbb{R}^n)^*$ is multiplication by $\det T$. Deduce $\det(ST) = \det S \det T$ in one line.

5. (Change of variables is a pullback.) Let $F(r,\theta) = (r\cos\theta,\ r\sin\theta)$. Compute $F^*(dx \wedge dy)$ directly and get $r\,dr \wedge d\theta$. Then explain in one sentence why the Jacobian factor in multivariable calculus is *automatic* in the language of forms.

6. (Contraction.) Define the interior product $\iota_X\omega(\cdot) = \omega(X,\cdot)$. Compute $\iota_{\partial_x}(dx\wedge dy\wedge dz)$ and $\iota_X(dx \wedge dy)$ for $X = y\partial_x - x\partial_y$. Show ι_X is an antiderivation: $\iota_X(\alpha\wedge\beta) = \iota_X\alpha\wedge\beta + (-1)^{|\alpha|}\alpha\wedge\iota_X\beta$.

7. (Decomposability.) In $\Lambda^2(\mathbb{R}^4)$, let $\omega = e_1\wedge e_2 + e_3\wedge e_4$. Compute $\omega\wedge\omega$ and conclude ω is **not** of the form $v\wedge w$. Then prove the general fact: a 2-form ω is decomposable iff $\omega\wedge\omega = 0$.

8. (The first cohomology class you'll ever meet.) On $\mathbb{R}^2\setminus\{0\}$ let

$$\omega = \frac{x\,dy - y\,dx}{x^2+y^2}.$$

(a) Show $d\omega = 0$. (b) Compute $\int_{S^1}\omega$. (c) Conclude ω is not exact, and explain why (b) doesn't contradict the fact that $\omega = d\theta$ locally.

9. (Leibniz, in general.) Prove that there is a *unique* operator $d:\Omega^k\to\Omega^{k+1}$ on $\mathbb{R}^n$ satisfying: d is $\mathbb{R}$-linear, df is the usual differential on functions, $d(\alpha\wedge\beta) = d\alpha\wedge\beta + (-1)^k\alpha\wedge d\beta$, and $d^2 = 0$. *(The point: d isn't a choice, it's forced.)*

10. (Where the decomposition breaks.) Verify $\dim \mathrm{Sym}^2(\mathbb{R}^n) + \dim\Lambda^2(\mathbb{R}^n) = n^2$. Now take $n=2$, $k=3$: compute $\dim(\mathbb{R}^2)^{\otimes 3}$, $\dim \mathrm{Sym}^3$, $\dim\Lambda^3$. What's missing, and how big is it? *(You've just discovered mixed-symmetry tensors — the beginning of representation theory of S_k and GL_n.)*

Tensors on a manifold

11. (Vectors are derivations.) Show a derivation X at p kills constants. Then, using Hadamard's lemma $f(x) = f(p) + \sum_i (x^i - p^i)g_i(x)$ with $g_i(p) = \partial_i f(p)$, show $X = \sum_i X(x^i)\,\partial_i|_p$, so $\dim T_pM = n$.

12. (Transformation laws.) Under a change of chart $y = y(x)$, show vector components transform as $v'^i = \frac{\partial y^i}{\partial x^j}v^j$ and covector components as $\omega'_i = \frac{\partial x^j}{\partial y^i}\omega_j$. Check that $\omega_i v^i$ is chart-independent.

13. (Tensor or not?) Compute $[X, fY]$ and conclude the Lie bracket is not $C^\infty(M)$-linear in Y. Then show the torsion $T(X,Y) = \nabla_X Y - \nabla_Y X - [X,Y]$ of a connection *is* $C^\infty(M)$-bilinear, hence a tensor field. Which Leibniz terms cancel?

14. (Induced metric.) Parametrize the unit sphere by $(\theta,\varphi)\mapsto(\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$ and pull back $dx^2+dy^2+dz^2$. You should get $d\theta^2 + \sin^2\theta\,d\varphi^2$.

15. (Vectors don't push forward.) Let $F:\mathbb{R}\to S^1$, $F(t) = e^{it}$, and $X = t\,\partial_t$ on $\mathbb{R}$. Show

"$dF(X)$" is not a well-defined vector field on S^1. Which hypothesis on F would fix it?

16. (d commutes with pullback.) Prove $F^*(df) = d(F^*f)$ for functions (it's the chain rule), then extend to all forms using $F^*(\alpha \wedge \beta) = F^*\alpha \wedge F^*\beta$ and the uniqueness of d from exercise 9.

Homology and cohomology

17. (The circle.) Show a 1-form $f(\theta)\,d\theta$ on S^1 is exact iff $\int_0^{2\pi} f = 0$. Conclude $H^1_{dR}(S^1) \cong \mathbb{R}$, with the isomorphism given by integration.

18. (The pairing is well defined.) For $M = \mathbb{R}^2 \setminus \{0\}$, write out both Stokes arguments that make $\langle [\omega], [c] \rangle = \int_c \omega$ independent of the representatives, for ω a closed 1-form and c a closed curve.

19. (The torus.) Show $d\theta_1$ and $d\theta_2$ are closed but not exact on T^2 and form a basis of $H^1_{dR}(T^2)$, and that $d\theta_1 \wedge d\theta_2$ generates $H^2_{dR}(T^2)$.

20. (Rings see more than Betti numbers.) Granting $H^*(\mathbb{CP}^3) \cong \mathbb{R}[x]/(x^4)$ with $\deg x = 2$, show $S^2 \times S^4$ and $\mathbb{CP}^3$ have the same Betti numbers but non-isomorphic cohomology rings.

21. (Poincaré lemma, degree 1.) On a star-shaped open $U \subset \mathbb{R}^n$, let $\omega = \sum_i \omega_i\,dx^i$ be closed and set $f(x) = \int_0^1 \sum_i \omega_i(tx)\,x^i\,dt$. Show $df = \omega$. Where is $d\omega = 0$ used?

22. (No Kähler S^6.) On a compact Kähler manifold of complex dimension n, show $[\omega]^k \neq 0$ in H^{2k} for $1 \leq k \leq n$, using $\int_M \omega^n > 0$. Deduce that S^6 carries no Kähler structure, whatever the answer to the famous open question about complex structures on it.

Visualizing Grad, Div, and Curl

Three pictures, one for each operator, and then the question every student asks once the pictures are in place: is $\nabla \times \nabla f = 0$ *visibly* true, or just an algebraic accident?

The three pictures

Gradient — the topographic map. Not "a vector that points uphill" but rather: draw the level sets. ∇f is perpendicular to them, and its magnitude is the *reciprocal spacing* — closely bunched contours mean a steep gradient. The better mental object isn't the arrow at all, it's the stack of level surfaces (Misner–Thorne–Wheeler's "egg crate"). You measure $df(v)$ by counting how many surfaces v pierces. Keep that picture; it does all the work below.

Divergence — the dye blob. Drop a small blob of dye in the flow and watch its *volume*. div v is the fractional rate of volume change per unit time; equivalently, flux out of a shrinking box. Incompressible flow, div $v = 0$: blobs deform but never swell. $\nabla \cdot E = \rho/\varepsilon_0$ says field lines only begin and end on charge.

Curl — circulation per unit area, and here's where everyone's intuition is wrong. The paddlewheel picture is right, but the *phrase* people attach to it — "does the flow swirl?" — is wrong, and it's worth killing that instinct early. Two examples:

- $v = (y, 0, 0)$, a shear flow. Perfectly straight streamlines, no swirl anywhere — and yet $\nabla \times v = -\hat{z}$. A paddlewheel's top blade gets pushed harder than its bottom one, so the wheel turns.
- $v = \hat{\theta}/r$, the potential vortex (the bathtub drain). Everything *visibly* swirls, and the curl is zero everywhere off the axis. A paddlewheel is carried around the center but its arms keep pointing in fixed directions, because the inner blade is sped up exactly as much as the outer blade is slowed.

Curl is not "swirl." It's **circulation per unit area**, and it measures the local rotation of a *material element*, not the shape of the trajectory. (Precisely: $\frac{1}{2}\nabla \times v$ is the instantaneous angular velocity of the fluid element.)

Is $\nabla \times \nabla f = 0$ visually obvious?

Yes — but not through the paddlewheel. Through the loop.

Take curl in its integral definition, $(\nabla \times F) \cdot \hat{n} = \lim \frac{1}{A} \oint_{\partial A} F \cdot d\mathbf{r}$, and ask what the circulation of a gradient around a closed loop is:

$$\oint_C \nabla f \cdot d\mathbf{r} = \int_C df = f(\text{end}) - f(\text{start}) = 0.$$

You hiked a closed loop on the mountain and came back to the same altitude. That's it. Every loop, no matter how small, has zero circulation, so the circulation density is zero everywhere.

The contrapositive is where it becomes vivid: a field with nonzero curl would be an Escher staircase. A gradient field with a swirl in it is a loop on which you're always walking downhill and yet return to your starting elevation. The impossibility of the Escher waterfall *is* $\nabla \times \nabla f = 0$.

There's a second geometric reading worth having. A gradient field is everywhere orthogonal to a family of level surfaces that fit together globally — the contour map exists. Curl is precisely the obstruction to that: it's the "screw dislocation" in the stack of surfaces, the failure of the sheets to close up on themselves as you march around a loop. (This is Frobenius integrability: the weaker condition $F \cdot (\nabla \times F) = 0$ is what it takes for F to be normal to *some* family of surfaces, versus $\nabla \times F = 0$ for $F = \nabla f$. Problem 8 works it out.)

And the determinant?

The determinant version is bookkeeping, not intuition. Write the $\hat{z}$ component:

$$(\nabla \times \nabla f)_z = \partial_x(\partial_y f) - \partial_y(\partial_x f) = 0.$$

That's Clairaut — mixed partials commute. Full stop. The determinant just organizes three copies of that cancellation.

But there's a trap here, because it's the mnemonic everyone repeats: "$\nabla \times \nabla f = 0$ because the determinant has two identical rows — it's the cross product of a vector with itself." That reasoning is **bogus**, and bogus in an instructive way. ∇ is not a vector; its "components" are operators, and nothing says operators commute. The determinant collapses *only because* $\partial_i \partial_j = \partial_j \partial_i$. Replace the partials with covariant derivatives on a bundle and the identical rows do not cancel:

$$[D_i, D_j] \neq 0,$$

and what survives is the **curvature** — $F = dA + A \wedge A$, the field strength of a gauge theory, the Riemann tensor. So $\nabla \times \nabla = 0$ is not a formal triviality about determinants; it's the statement that flat space with ordinary derivatives has no curvature (Problem 9).

The same logic runs one level up: $\nabla \cdot (\nabla \times F) = 0$ because the flux of a curl through a closed surface equals the circulation around that surface's boundary, and a closed surface has no boundary. $d^2 = 0$ is $\partial^2 = 0$ — the boundary of a boundary is zero.

One caveat to hold in reserve

$\nabla \times F = 0$ does **not** give $F = \nabla f$ globally. The vortex above is the counterexample: it is curl-free on $\mathbb{R}^2 \setminus \{0\}$, but $\oint F \cdot d\mathbf{r} = 2\pi$ around the origin, so no potential exists. Locally exact, globally not. That gap is H^1_{dR}, and it's the reason de Rham cohomology exists at all (see the Tensors chapter). Poincaré's lemma buys the potential only on contractible domains.

In short: the loop argument gives the *conviction*; the mixed-partials argument gives the *mechanism* and the sign checks; Stokes' theorem is the bridge that says they're the same statement at two magnifications. If you only keep one, keep the mountain.

See it: the surface explorer

The pictures above are built into an interactive page, https://math.babelslibrary.com/surface-explorer.html: type any $f(x, y)$, fly around its graph, and drag a slider that highlights the level set $\{f = c\}$ on the surface and on a contour map projected onto the floor, with ∇f arrows on the map. Three things to try:

- **The fold.** On $f = \frac{y}{x^2} e^{-y/x^2}$ (close to Hubbard's example 1.5.24, the page's default), drag c up toward $1/e \approx 0.368$. The two parabolas of the level set converge and pinch off into the ridge $y = x^2$. Turn on the floor gradient: on either side of the ridge the arrows point *toward each other*. Writing $f = g(u)$ with $u = y/x^2$, that's $g'(u) = e^{-u}(1-u)$ changing sign at $u = 1$ — which is why the level set folds rather than passing through. $|\nabla f|$ vanishes on the whole ridge: a critical *curve*, not a critical point (Problem 5).
- **A degenerate critical point.** On the monkey saddle $x^3 - 3xy^2$, hide the surface and keep only the floor map and arrows: three ascending and three descending sectors meet at the origin, where every arrow dies.
- **The gradient is a covector.** The page squeezes the domain anisotropically into its display box. Tangent vectors scale with the box, $(t_x, t_y) \mapsto (s_x t_x, s_y t_y)$, so to keep $df(t)$ unchanged the gradient's components must scale the *other* way, $(g_x, g_y) \mapsto (g_x/s_x, g_y/s_y)$. Scale the arrows like vectors and they visibly stop meeting the contours at right angles — a nice-looking picture of a false statement (Problem 7).

The page draws graphs $z = f(x, y)$ only. Level *surfaces* of $F(x, y, z)$ need marching cubes or GPU raymarching instead of marching squares; once f has three arguments the graph lives in $\mathbb{R}^4$, and the stack of level surfaces stops being an aid and becomes the only picture

there is. Off-the-shelf tools that do go there: ParaView and PyVista (VTK) for isosurfaces with a slider, and CalcPlot3D or GeoGebra 3D with no installation at all.

Problems

Five simple, two medium, one hard, plus a stretch problem. Hints are given for the medium and hard problems only.

Problem 8(c) is the point of the whole set. $F = (y, 0, 1)$ isn't merely "not a gradient" — it's not normal to *any* family of surfaces, so there is no surface anywhere for it to be perpendicular to. That's the standard contact structure on $\mathbb{R}^3$, and it's the honest converse to everything drawn above: every picture in this chapter presupposes that a stack of level surfaces exists. Problem 8 asks what happens when it doesn't.

Simple

1. Let $f(x, y) = 3x - 4y$.

(a) Describe the level sets $\{f = c\}$.
(b) Compute ∇f and $|\nabla f|$.
(c) If you draw contours at increments $\Delta c = 0.5$, what is the perpendicular spacing between adjacent contours in the plane? Confirm it agrees with $|\nabla f| \approx \Delta c / \Delta s$.

---

2. Two fields on $\mathbb{R}^3$:

$$\mathbf{F} = (y,\, 0,\, 0), \qquad \mathbf{G} = \left(\frac{-y}{x^2 + y^2},\, \frac{x}{x^2 + y^2},\, 0 \right).$$

(a) Compute $\nabla \times \mathbf{F}$ and $\nabla \times \mathbf{G}$ (the latter away from the z-axis).
(b) Which of the two has *straight* streamlines? Which has a nonzero curl?
(c) State in one sentence why this pair shows that "curl measures swirl" is the wrong slogan.

---

3. Is $\mathbf{F} = (y,\, x,\, 0)$ a gradient field? If so, exhibit a potential. If not, produce the obstruction.

Then do the same for $\mathbf{H} = (y,\, -x,\, 0)$.

---

4. Let $\mathbf{r} = (x, y, z)$ and $r = |\mathbf{r}|$.

(a) Compute $\nabla \cdot \mathbf{r}$ and $\nabla \times \mathbf{r}$.
(b) Compute $\nabla \cdot \left(\mathbf{r} / r^3 \right)$ for $r \neq 0$. Interpret the answer physically, and say what is happening at the origin.

---

5. Return to $f(x,y) = \dfrac{y}{x^2}\, e^{-y/x^2}$ on $x > 0$.

(a) Verify $f(\lambda x, \lambda^2 y) = f(x,y)$ for all $\lambda > 0$.
(b) Deduce the shape of the level sets without computing anything further.
(c) Find every point where $\nabla f = 0$. Is the critical set a set of isolated points? What does that tell you about whether f is a Morse function?
(d) For which values of c does $\{f = c\}$ consist of exactly two parabolas, exactly one, or none?

Medium

6. Let **G** be the vortex field of problem 2, on the punctured plane $\Omega = \mathbb{R}^2 \setminus \{0\}$.

(a) Verify $\nabla \times \mathbf{G} = 0$ on Ω.
(b) Compute $\oint_C \mathbf{G} \cdot d\mathbf{r}$ around the unit circle.
(c) Reconcile (a) and (b) with Stokes' theorem — be precise about which hypothesis fails.
(d) Exhibit an open subset $U \subset \Omega$ on which **G** *is* a gradient, and give the potential explicitly. What is the largest such U you can find, and what goes wrong at its boundary?
(e) State what $b_1(\Omega)$ is and how (b) computes it.

Hint: the potential is an angle, and angles are only defined up to a choice of branch.

---

7. (A real bug, met while building the surface explorer.) Rescale the plane anisotropically by $\Phi(x,y) = (ax, by)$ with $a, b > 0$, $a \neq b$.

(a) A curve tangent $\mathbf{t} = (t_1, t_2)$ pushes forward to $\Phi_* \mathbf{t}$. Write it down.
(b) Suppose $\nabla f \cdot \mathbf{t} = 0$, i.e. $\mathbf{t}$ is tangent to a level set. Find the vector $\mathbf{n}$ such that $\mathbf{n}$ is perpendicular to $\Phi_* \mathbf{t}$ **for every** such $\mathbf{t}$. Express $\mathbf{n}$ in terms of ∇f, a, b.
(c) Conclude: does the gradient transform with the scale factors or their reciprocals? Which of the two objects — the level-set stack, or the arrow — is the one that survives the rescaling unchanged?
(d) For which (a,b) does naively pushing ∇f forward as a vector give the correct normal direction anyway? Give the general principle behind your answer.

Hint: (b) is asking you to preserve a pairing, not a length.

Hard

8. Let **F** be a smooth nonvanishing vector field on an open set $U \subset \mathbb{R}^3$. Call **F surface-orthogonal** if near each point there is a smooth function g with $\nabla g \neq 0$ such that **F** is

everywhere normal to the level sets of g — equivalently, if there is a nonvanishing scalar μ (an *integrating factor*) with $\mu\mathbf{F} = \nabla g$.

(a) Show that if $\mathbf{F}$ is surface-orthogonal then

$$\mathbf{F} \cdot (\nabla \times \mathbf{F}) = 0.$$

(b) Show that the converse holds locally. (You may cite Frobenius; if you do, say precisely which statement of it you are using and verify its hypothesis.)

(c) Let $\mathbf{F} = (y,\, 0,\, 1)$. Compute $\mathbf{F} \cdot (\nabla \times \mathbf{F})$ and conclude that **no** family of surfaces is orthogonal to $\mathbf{F}$ anywhere — not merely that $\mathbf{F}$ fails to be a gradient.

(d) Reconcile this with problem 1: there, the level sets existed and were obvious. What structural feature does $\mathbf{F} = (y, 0, 1)$ have that ∇f can never have?

(e) Translate the whole problem into the language of differential forms. Write $\omega = F_1\, dx + F_2\, dy + F_3\, dz$ and restate (a) and (b) as a single condition on ω and $d\omega$. Why is this restatement the *right* one — i.e. what does it make manifest that the vector formulation hides?

Hint for (a): differentiate $\mu\mathbf{F} = \nabla g$ and use $\nabla \times \nabla g = 0$ together with $\nabla \times (\mu\mathbf{F}) = \mu\nabla \times \mathbf{F} + \nabla\mu \times \mathbf{F}$. Then dot with $\mathbf{F}$ and notice that one term dies for free.

Hint for (e): the condition you want is three symbols long.

Stretch (optional, no credit, purely for fun)

9. The identity $\nabla \times \nabla f = 0$ rests on $\partial_i \partial_j = \partial_j \partial_i$. Replace the partials by covariant derivatives $D_i = \partial_i + A_i$, where the A_i are matrix-valued and need not commute.

Compute $[D_i, D_j]$ and identify what you get. In what precise sense is the failure of "curl grad $= 0$" the definition of curvature? Write the answer both in index form and as $F = dA + A \wedge A$.