MATH 725 — APPLIED FUNCTIONAL ANALYSIS

LECTURE NOTES

ANDREAS BUTTENSCHOEN*

Andreas Buttenschoen.

University of Massachusetts Amherst

These are work in progress notes.

CONTENTS

1 Function Spaces and Foundations 6
1.1 Overview . . . 6
1.2 What You Will Learn . . . 6
1.3 Modern Applied Mathematics . . . 7
1.3.1 Differential equations . . . 7
1.3.2 The Diffusion Equation . . . 8
1.4 Spaces of Continuous and Differentiable Functions . . . 10
1.4.1 Learning Goals . . . 10
1.4.2 Banach Spaces . . . 10
1.4.3 Spaces of Continuous and Differentiable Functions . . . 10
1.5 Spaces of Integrable Functions . . . 12
1.5.1 Learning Goals . . . 12
1.5.2 Spaces of Integrable Functions . . . 12
1.5.3 What Do L^p Norms Measure? . . . 12
1.5.4 L^p is a Banach Space . . . 13
1.5.5 The Rainbow of Function Spaces . . . 13
1.6 Density and Approximation . . . 14
1.6.1 Dense Subsets . . . 14
1.6.2 Separability . . . 14
1.6.3 Mollification . . . 14
1.6.4 Density of Smooth Functions in L^p . . . 15
1.6.5 Approximation Beyond Polynomials . . . 15
1.7 Topological Equivalence of Norms . . . 16

For more visit www.buttenschoen.ca.

*andreas.buttenschoen@umass.edu

1.8 Compactness . . . 18
1.9 Compactness in Function Spaces . . . 20
1.9.1 Arzelà-Ascoli for C(K) . . . 20
1.9.2 Kolmogorov-Riesz for L^p . . . 20

2 Continuous Linear Operators 22
2.1 Overview . . . 22
2.2 What You Will Learn . . . 22
2.3 Bounded Operators . . . 23
2.4 The Baire Category Theorem . . . 24
2.4.1 Small sets in a Measure Space . . . 24
2.4.2 Small Sets in Topological Spaces . . . 24
2.4.3 Consequences of the Baire Category Theorem . . . 27
2.5 The uniform boundedness principle (Banach-Steinhaus) . . . 28
2.5.1 Why Does Banach-Steinhaus Need Completeness? . . . 28
2.5.2 Why This Matters in Practice . . . 28
2.6 The Open Mapping Theorem . . . 30
2.6.1 Why Does the Open Mapping Theorem Need Completeness? . . . 30
2.6.2 Quantitative Formulation . . . 30
2.7 Quotient Spaces and Codimension . . . 33
2.7.1 Quotient Spaces . . . 33
2.7.2 Codimension . . . 34
2.7.3 The First Isomorphism Theorem . . . 34
2.8 Compact Operators . . . 36
2.8.1 Definition and Basic Properties . . . 36
2.8.2 The Space of Compact Operators . . . 36
2.8.3 Finite-Rank Operators: Compact Operators as "Infinite-Dimensional Matrices" 37
2.8.4 The Canonical Example: Integral Operators . . . 39
2.8.5 Spectral Theory of Compact Operators . . . 40
2.8.6 Looking Ahead . . . 42
2.9 Exercises . . . 43
2.9.1 Linear Operators . . . 43
2.9.2 Banach-Steinhaus . . . 43
2.9.3 Banach-Steinhaus Applications . . . 43
2.9.4 Open Mapping Theorem . . . 44
2.9.5 Open Mapping Theorem Applications . . . 44
2.9.6 Compact Operators . . . 45

3 Duality: Hyperplanes, Functionals, and Hahn–Banach 46
3.1 Overview . . . 46
3.2 What You Will Learn . . . 46
3.3 The Dual Space . . . 47
3.3.1 Introduction . . . 47
3.3.2 Isomorphisms and Isometries . . . 47
3.3.3 The Finite-Dimensional Case . . . 47

3.3.4 What Goes Wrong in Infinite Dimensions? . . . 49
3.4 The Geometry of Linear Functionals . . . 51
3.4.1 Kernels and level sets . . . 51
3.4.2 The fundamental decomposition . . . 51
3.4.3 The foliation picture . . . 51
3.4.4 Functionals are determined by their kernels . . . 53
3.4.5 Characterizing continuity via hyperplanes . . . 54
3.4.6 A Worked Example: Hahn–Banach Extension in $\mathbb{R}^2$ . . . 56
3.5 The Hahn–Banach Theorem . . . 59
3.5.1 Geometric Consequences . . . 60
3.5.2 Geometric Hahn–Banach: Separation of Convex Sets . . . 62
3.6 Dual Spaces: From Self-Dual to Non-Reflexive . . . 64
3.6.1 Hilbert Spaces: $H^* \cong H$ (Self-Dual) . . . 64
3.6.2 Reflexive Spaces: $X^{**} \cong X$ . . . 65
3.6.3 Non-Reflexive Spaces: $X \subsetneq X^{**}$ . . . 68
3.6.4 Spaces with Trivial Dual . . . 70
3.7 Application: The Universal Approximation Theorem . . . 71
3.7.1 Single-hidden-layer networks . . . 71
3.7.2 Discriminatory activations . . . 71
3.7.3 Cybenko's theorem . . . 71
3.7.4 Which activations are discriminatory? . . . 72
3.7.5 A second strategy for density proofs . . . 73
3.8 The Weak Topology . . . 74
3.8.1 What we need from a topology . . . 74
3.8.2 From norm balls to weak neighborhoods . . . 75
3.8.3 Strong and weak convergence . . . 77
3.8.4 Weak compactness in reflexive spaces . . . 81
3.8.5 Separation and convexity in the weak topology . . . 83
3.9 The Weak* Topology and Banach-Alaoglu . . . 85
3.9.1 The weak* topology . . . 85
3.9.2 Where weak* sits: the hierarchy of topologies . . . 86
3.9.3 Weak* convergence . . . 87
3.9.4 The Banach-Alaoglu theorem . . . 89
3.10 Applications of Weak and Weak* Convergence . . . 92
3.10.1 Convergence in probability: all three in action . . . 92
3.10.2 The direct method: a first look . . . 93
3.11 The Adjoint (Transpose) Operator . . . 97
3.11.1 The finite-dimensional picture . . . 97
3.11.2 The Banach adjoint (transpose) . . . 98
3.11.3 The four subspace relations: general Banach spaces . . . 98
3.11.4 Reflexive spaces: full symmetry restored . . . 100
3.11.5 Hilbert spaces: orthogonal complements . . . 100
3.12 Application: Fredholm Operators . . . 102
3.12.1 Solvability conditions . . . 104
3.12.2 The bordered operator . . . 105

3.12.3 Solving the bordered system by block elimination . . . 107

4 Distributions and Sobolev Spaces 109
4.1 Why distributions? . . . 109
4.2 The two topologies . . . 109
4.3 The hierarchy of generalized functions . . . 110
4.4 What comes next . . . 110
4.5 The Space of Test Functions . . . 111
4.5.1 Definition . . . 111
4.5.2 Why C_c^∞ is rich enough . . . 111
4.5.3 The inductive limit topology . . . 112
4.5.4 Characterizing continuous functionals . . . 114
4.5.5 What transfers from earlier chapters? . . . 114
4.6 Distributions as Functionals . . . 117
4.6.1 Definition of Distributions . . . 117
4.6.2 Convergence of Distributions . . . 117
4.6.3 The Hierarchy of Distributions . . . 118
4.6.4 Topological Properties of $\mathcal{D}'(\Omega)$ . . . 121
4.7 The Calculus of Distributions . . . 125
4.7.1 Multiplication by Smooth Functions . . . 125
4.7.2 Differentiation of Distributions . . . 126
4.7.3 Weak Derivatives . . . 131
4.7.4 Translation and Reflection . . . 132
4.7.5 Convolution . . . 133
4.8 Support and Locality . . . 135
4.8.1 The Structure Theorem for Distributions . . . 135
4.9 Sobolev Spaces . . . 137
4.9.1 From weak derivatives to function spaces . . . 137
4.9.2 What does the Sobolev norm measure? . . . 138
4.9.3 Examples: what is and isn't in $W^{1,2}$ . . . 141
4.10 Sobolev Embeddings . . . 143
4.10.1 Motivation: when does a bounded sequence converge? . . . 143
4.10.2 From Sobolev regularity to integrability . . . 143
4.10.3 The uncertainty principle and the general theorem . . . 147
4.10.4 Compact embeddings: Rellich-Kondrachov . . . 151
4.11 Sobolev Spaces as Banach Algebras . . . 154
4.11.1 Why multiplication matters . . . 154
4.11.2 The algebra property from embeddings . . . 154
4.11.3 Multiplication in frequency space: why regularity implies bandedness . . . 156
4.12 Sobolev Duality and Structure . . . 159
4.12.1 Duality, reflexivity, and the Sobolev scale . . . 159
4.13 Application: Differential Equations . . . 163
4.13.1 From ODE Superposition to PDE Convolution . . . 163
4.13.2 Fundamental Solutions . . . 163
4.13.3 The Malgrange-Ehrenpreis Theorem . . . 165

4.13.4 Solving PDEs via Convolution 166
4.14 Application: Weak Formulation of PDEs 167
4.14.1 From classical to weak solutions 167
4.14.2 The Gelfand triple 170
4.14.3 Nonlinear PDE via Compactness 171

References 175

1 FUNCTION SPACES AND FOUNDATIONS

Big Idea

PDEs look like ODEs—if you squint. The key insight of functional analysis is to treat solutions as points in infinite-dimensional function spaces, where differential operators become linear maps between Banach spaces. This abstraction lets us port finite-dimensional intuition (eigenvalues, convergence, compactness) to the PDE world.

1.1 *Overview*

Why do we need functional analysis? Consider a reaction-diffusion equation:

$$\frac{\partial u}{\partial t} = d\Delta u + f(u) \tag{1}$$

If we set $\mathrm{A}[u] := d\Delta u$, this looks like the ODE $\dot{x} = \mathrm{A}x + f(x)$. For ODEs, solutions live in $\mathbb{R}^n$—a finite-dimensional space with a well-understood theory. But PDE solutions live in **infinite-dimensional function spaces**, and much of finite-dimensional intuition breaks down:

- Not all norms are equivalent.
- Bounded closed sets need not be compact.
- Convergence comes in many flavors (strong, weak, weak-*).

This chapter sets up the language and tools needed to work in these spaces: norms, completeness, density, separability, and compactness.

1.2 *What You Will Learn*

- How to define and compare norms on function spaces (L^p, C^k, Sobolev).
- Completeness and Banach spaces—why we need them and how to verify them.
- Density and approximation—mollifiers, Weierstrass approximation.
- Compactness in infinite dimensions—Arzelà–Ascoli and why bounded sequences don't always have convergent subsequences.

1.3 *Modern Applied Mathematics*

- Development of new mathematical models (modeling).
- **Theoretical analysis of mathematical models.**
- Numerical solutions of these models.
- Data inferences and statistical learning.

This course focuses on the second point.

1.3.1 Differential equations

We motivate the need for the development of functional analysis through the ubiquitous nature of differential equations. Similar functional analytical techniques are useful when studying stochastic differential equations, and many other applications.

- **ODEs** (ordinary differential equations) are defined by having a finite number of variables of a single independent variable usually t.
  For example the following is a linear ODE:

$$\frac{d\vec{x}}{dt} = \mathrm{A}\vec{x} \tag{2}$$

  where A is a $n \times n$ matrix. Then the solution is:

$$t \mapsto \vec{x}(t) \in \mathbb{R}^n \tag{3}$$

  The space in which the solution of a differential equation "lives" is usually called a **phase-space**.
- **PDEs** require new language.
  - State space is a Banach space of functions.
  - PDEs are operators between Banach spaces.
  - Solutions often arise as weak or weak* limits in these Banach spaces.
- The following reaction-diffusion is an example of a PDE:

$$\frac{\partial u}{\partial t} = d\Delta u + f(u) \tag{4}$$

  for an unknown function $u(x, t)$. Here Δ is the Laplacian:

$$\Delta := \sum_{i=1}^{n} \frac{\partial^2}{\partial x_i^2} \tag{5}$$

  If we formally introduce the linear operator:

$$\mathrm{A}[u] := d\Delta u \tag{6}$$

  then we can rewrite the PDE as:

$$\frac{\partial u}{\partial t} = \mathrm{A}[u] + f(u) \tag{7}$$

and it looks like an ODE.

- If the PDE was an ODE we can use the variation of constant formula and the matrix exponential to find a solution i.e.:

$$u(t) = e^{\mathrm{A}t}u_0 + \int_0^t e^{\mathrm{A}(t-s)}f(u(s))ds. \tag{8}$$

But here A is a differential operator and not a matrix. Can we define matrix exponentials for unbounded operators? The answer is yes as we will see in the chapter on semigroup theory.

- If $u(x;t)$ is a solution of the "ODE" then we have a mapping:

$$t \mapsto u(x;t) \in \text{Function Space}. \tag{9}$$

Thus the phase-space for a PDE is a function space, and the PDE becomes an operator between function spaces. Whether we can invert these operators we will discuss in our chapter on operator theory.

1.3.2 The Diffusion Equation

Let's discuss the heat / diffusion equation example in a bit more detail. Suppose we have the heat equation on the interval $[0, \mathrm{L}]$.

$$\begin{cases} u_t & = d\,u_{xx} \\ u(t,0) & = u(t,\mathrm{L}) = 0 \\ u(x,0) & = u_0(x). \end{cases} \tag{10}$$

Using separation of variables and the principles of superposition we find a solution as a Fourier sine series:

$$u(x,t) = \sum_{k=1}^{\infty} a_n e^{-\lambda_n d\,t} \sin\left(\frac{n\pi x}{\mathrm{L}}\right) \tag{11}$$

The coefficients a_n are determined from the initial condition i.e.:

$$u_0(x) = \sum_{k=1}^{\infty} a_n \sin\left(\frac{n\pi x}{\mathrm{L}}\right) \tag{12}$$

and we can use the Fourier trick to solve this equation for a_n. Now the family of functions:

$$\mathcal{L} := \left\{ \sin\left(\frac{n\pi x}{\mathrm{L}}\right) : n \in \mathbb{N} \right\} \tag{13}$$

forms an orthogonal set with the inner product:

$$(\phi_n, \phi_m) = \int_0^{\mathrm{L}} \phi_n(x)\phi_m(x)dx \tag{14}$$

Recall that this define a norm via $\|x\| = \sqrt{(x, x)}$. Then $\mathcal{L}$ can be seen as the basis of the function space $\mathrm{L}^2(0, \mathrm{L})$ the square integrable functions. Since $\mathcal{L}$ is infinite, $\mathrm{L}^2(0, \mathrm{L})$ is an infinite dimensional vector space. Thus we understand the heat equation as an infinite dimensional differential equation.

The final note is that the linear operator Δ is crucially associated with the boundary conditions. Different boundary conditions lead to different basis and thus different function spaces. Much of PDE theory is concerned with the identification of the correct basis functions and right function spaces. We will come back to this later when discussing operator theory.

1.4 *Spaces of Continuous and Differentiable Functions*

1.4.1 Learning Goals

1. What is a Banach space?
2. The spaces of continuous and differentiable functions.
3. What do C^k norms measure?

1.4.2 Banach Spaces

1.1 DEFINITION (Banach Space). A Banach space $(X, \|\cdot\|_X)$ is a complete normed vector space.

1.2 EXAMPLE (Basic Banach Space). $\mathbb{R}^n$ is a Banach space with any p-norm i.e. $\|x\|_p = \left(\sum_{i=1}^{n} |x_i|^p\right)^{1/p}$

1.4.3 Spaces of Continuous and Differentiable Functions

1.3 DEFINITION (— Continuous Functions). C^0

Let $\Omega \subset \mathbb{R}^n$ be open. The space of continuous functions on $\overline{\Omega}$ is

$$C^0(\overline{\Omega}) = C(\overline{\Omega}) = \{f : \overline{\Omega} \to \mathbb{R} : f \text{ is continuous}\} \tag{15}$$

equipped with the **sup-norm** (or uniform norm)

$$\|f\|_\infty = \sup_{x \in \overline{\Omega}} |f(x)|. \tag{16}$$

When $\overline{\Omega}$ is compact, every continuous function is bounded and attains its supremum, so $\|f\|_\infty < \infty$ for all $f \in C(\overline{\Omega})$.

1.4 DEFINITION (—). C^k k**-times Continuously Differentiable Functions**

For $k \geq 1$, define

$$C^k(\overline{\Omega}) = \{f \in C(\overline{\Omega}) : D^\alpha f \in C(\overline{\Omega}) \text{ for all } |\alpha| \leq k\} \tag{17}$$

where α is a multi-index and $D^\alpha f$ denotes the corresponding partial derivative. This space is equipped with the C^k **norm**

$$\|f\|_{C^k} = \sum_{|\alpha| \leq k} \|D^\alpha f\|_\infty = \|f\|_\infty + \sum_{|\alpha|=1} \|D^\alpha f\|_\infty + \cdots + \sum_{|\alpha|=k} \|D^\alpha f\|_\infty. \tag{18}$$

1.5 DEFINITION (— Smooth Functions). C^∞

The space of smooth (infinitely differentiable) functions is

$$C^\infty(\overline{\Omega}) = \bigcap_{k=0}^{\infty} C^k(\overline{\Omega}). \tag{19}$$

This space does not carry a single norm but is topologized by the family of C^k seminorms: a sequence $f_n \to f$ in C^∞ if and only if $\|f_n - f\|_{C^k} \to 0$ for every k.

WHAT DO C^k NORMS MEASURE? The C^k norm controls the function **and** its first k derivatives uniformly. This has a concrete geometric meaning:

- $\|f\|_{C^0} = \|f\|_\infty$ controls the **height**: two functions are C^0-close if their graphs are uniformly close.
- $\|f\|_{C^1} = \|f\|_\infty + \|f'\|_\infty$ additionally controls the **slope**: two functions are C^1-close if their graphs are close *and* their tangent lines are close at every point.
- $\|f\|_{C^k}$ controls curvature, jerk, and higher-order shape information. Closeness in C^k means the functions "look the same" up to k-th order magnification.

The key point: **higher-order C^k norms are strictly stronger**. A sequence can converge in C^0 (graphs converge) without converging in C^1 (slopes may oscillate wildly).

1.6 EXAMPLE (with the sup-norm is incomplete). C^1

The space $C^1([0,1])$ equipped with the **wrong** norm $\|\cdot\|_\infty$ (instead of $\|\cdot\|_{C^1}$) is **not** a Banach space.

Consider the sequence $f_n(x) = \sqrt{x + 1/n}$ on $[0,1]$. Each f_n is C^1, and $f_n \to f(x) = \sqrt{x}$ uniformly. But $f(x) = \sqrt{x}$ is not differentiable at $x = 0$, so $f \notin C^1([0,1])$.

The sup-norm cannot detect that the derivatives $f_n'(x) = \frac{1}{2\sqrt{x+1/n}}$ are blowing up near $x = 0$. The C^1 norm would catch this: $\|f_n\|_{C^1} \to \infty$, so the sequence is not Cauchy in the C^1 norm.

This is a concrete instance of a general principle: **a function space is only complete in a norm strong enough to detect all the structure the space requires**. The sup-norm is too weak for C^1—it measures height but not slope, so Cauchy sequences can converge to functions that are continuous but not differentiable.

1.7 PROPOSITION (is a Banach space). $C^k(\overline{\Omega})$

For $\overline{\Omega}$ compact and $k \geq 0$, the space $(C^k(\overline{\Omega}), \|\cdot\|_{C^k})$ is a Banach space.

1.5 *Spaces of Integrable Functions*

1.5.1 Learning Goals

1. The spaces of integrable functions.
2. What do L^p norms measure? Intuition via bump functions.
3. The Rainbow of function spaces.

1.5.2 Spaces of Integrable Functions

1.8 DEFINITION (Spaces). L^p

For $1 \leq p < \infty$ we define the spaces of p-integrable functions

$$L^p(\Omega) = \left\{ f : \Omega \to \mathbb{R} : \|f\|_p = \left(\int_\Omega |f(x)|^p \, d\mu \right)^{1/p} < \infty \right\} \Big/ \sim \tag{20}$$

where $f \sim g$ if $f = g$ **almost everywhere** (i.e., $\mu(\{x : f(x) \neq g(x)\}) = 0$).
For $p = \infty$ we define

$$L^\infty(\Omega) = \{ f : \Omega \to \mathbb{R} : \|f\|_\infty = \operatorname{ess\,sup}_{x\in\Omega} |f(x)| < \infty \} \Big/ \sim \tag{21}$$

1.5.3 What Do L^p Norms Measure?

Different norms capture different features of a function. To build intuition, consider a smooth bump function $\phi \in C_c^\infty(\mathbb{R})$ with $\phi \geq 0$, $\operatorname{supp}(\phi) \subset [-1, 1]$, and $\|\phi\|_\infty = 1$.

We construct two families of functions that independently control **height** and **width**.

Scaling the height. Define $f_A(x) = A\,\phi(x)$ for $A > 0$. This rescales the amplitude without changing the support. Then:

$$\|f_A\|_\infty = A, \qquad \|f_A\|_1 = A\|\phi\|_1, \qquad \|f_A\|_p = A\|\phi\|_p. \tag{22}$$

All norms see the height equally — no surprise, since we are simply rescaling the function values.

Scaling the width. Define $g_R(x) = \phi(x/R)$ for $R > 0$. This stretches the support to $[-R, R]$ without changing the peak value. Then:

$$\|g_R\|_\infty = \|\phi\|_\infty = 1 \tag{23}$$

is **independent of** R — the L^∞ norm is completely insensitive to how wide the function is. On the other hand, by a change of variables $y = x/R$:

$$\|g_R\|_p^p = \int |g_R(x)|^p \, dx = R \int |\phi(y)|^p \, dy = R\,\|\phi\|_p^p \tag{24}$$

so that $\|g_R\|_p = R^{1/p}\|\phi\|_p$, which **grows** as we widen the support.

1.9 REMARK (Interpretation of). L^p **Norms**

This analysis reveals a fundamental distinction:

- L^∞ measures only the **height** (peak value) of a function. It is completely insensitive to the width of the support.
- L^1 measures the **total mass** — it averages the function over the domain, so both height and width contribute equally.
- L^p for $1 < p < \infty$ **interpolates** between these extremes: larger p gives more weight to peak values and less to the spread.

This is why L^∞ convergence (uniform convergence) is a much stronger requirement than L^1 convergence: a sequence of tall, narrow spikes can have L^1 norm tending to zero while the L^∞ norm remains constant.

1.5.4 L^p is a Banach Space

1.10 THEOREM (is a Banach Space). L^p

L^p *is a Banach space.*

1.5.5 The Rainbow of Function Spaces

1.11 REMARK (The Rainbow of Function Spaces). A key result from this lecture is the "rainbow" of function spaces. The spaces of continuous and differentiable functions naturally are nested. Similarly, the spaces of integrable functions are nested by Hölder's inequality for bounded domains.

In particular, spaces of integrable functions are only nestable when the measure of the domain is finite or the counting measure (in this case the role of p and q reverse). Thus crucially spaces of integrable functions with unbounded domains are **not** nested.

1.6 *Density and Approximation*

A central theme in functional analysis is **approximation**: given a function in some large space, can we approximate it arbitrarily well by functions from a nicer, more structured class? This question is formalized through the notions of density and separability.

1.6.1 Dense Subsets

1.12 DEFINITION (Dense Subset). A subset $Y \subset X$ is dense in X if $\bar{Y} = X$. Equivalently, given $x \in X$, $\forall \varepsilon > 0\ \exists y \in Y$ such that $\|x - y\| < \varepsilon$.

1.13 EXAMPLE (Rationals in Reals). $\bar{\mathbb{Q}} = \mathbb{R}$. Let $r \in \mathbb{R}$ and $\varepsilon > 0$ then there is n such that $\frac{1}{n} < \varepsilon$. Define $q_n = \frac{\lfloor nr \rfloor}{n}$ which implies that $\|r - \frac{nr}{n}\| < \frac{1}{n} < \varepsilon$

In other words we can approximate any real number r arbitrarily well by elements in the rationals e.g. the rationals are dense in the reals. This property of the rationals is crucially important for finite arithmetic approximations of the reals on computers.

1.14 EXAMPLE (Weierstrass Theorem). Let $f \in C^0([-1,1])$ and let $\varepsilon > 0$ then there exists a polynomial p such that $\|f - p\|_\infty < \varepsilon$.

1.6.2 Separability

1.15 DEFINITION (Separable Space). A Banach space which contains a dense **countable** subset is called separable.

1.16 EXAMPLE (Separability of Real Numbers). Since the rationals are countable and dense in $\mathbb{R}$, thus $\mathbb{R}$ is separable.

1.17 EXAMPLE (Separability of Continuous Functions). $C^0([-1,1])$ has the dense subset $\{ax^n : a \in \mathbb{Q},\ n \in \mathbb{N} \cup \{0\}\}$

This set is countable and dense in $C^0([-1,1])$: we can approximate any continuous function using a polynomial having rational coefficients. Thus $C^0([-1,1])$ is a separable Banach space.

1.18 EXAMPLE (Separability of Square Integrable Functions). The space of square integrable functions $L^2([-1,1])$ has basis set $\{1,\ \cos(nx),\ \sin(nx)\}$

which is countable and hence it is a separable space.

1.19 EXAMPLE (Non-separability of L-infinity). $L^\infty([-1,1])$ is not separable (see homework).

1.6.3 Mollification

The key technique for approximating rough functions by smooth ones is **mollification** — convolution with a smooth bump function that averages a function over a small neighborhood.

1.20 THEOREM (Mollifiers Theorem). *Given* $f \in C_c^0(\Omega)$, *for each* $\varepsilon > 0\ \exists \phi \in C_c^\infty(\Omega)$ *such that* $\|f - \phi\|_\infty < \varepsilon$

This is a powerful result: it tells us that any continuous function with compact support can be uniformly approximated by smooth, compactly supported functions. The idea is to convolve f with a smooth bump (a mollifier) at a sufficiently small scale.

1.6.4 Density of Smooth Functions in L^p

Mollification is the engine behind the fundamental density results for L^p spaces.

1.21 THEOREM (Density of). $\mathcal{C}^0$ ***in*** L^p

Let Ω be bounded, then $\mathcal{C}^0(\Omega)$ is dense in $L^p(\Omega)$ i.e.

$$L^p(\Omega) = \overline{\mathcal{C}_c^0(\Omega)} \tag{25}$$

where the closure is with respect to the p-norm, and L^p is separable.

Combining with the mollifiers theorem, we obtain the density of smooth functions.

1.22 COROLLARY (Density of). $\mathcal{C}_c^\infty$ ***in*** L^p

Let Ω be bounded. Then $\mathcal{C}_c^\infty(\Omega)$ is dense in $L^p(\Omega)$:

$$L^p(\Omega) = \overline{\mathcal{C}_c^\infty(\Omega)} \tag{26}$$

1.6.5 Approximation Beyond Polynomials

The classical approximation results above — Weierstrass, mollification, density of $\mathcal{C}^\infty$ in L^p — all share a common structure: they identify a "nice" class of functions (polynomials, smooth functions) that can approximate arbitrary elements of a larger space.

1.23 REMARK (Toward Neural Network Approximation). A natural question is: what other classes of functions are dense in common function spaces? It turns out that **neural networks** provide a modern and remarkably powerful answer. The Universal Approximation Theorem (Cybenko, 1989; Hornik, 1991) shows that single-hidden-layer neural networks with sufficiently many neurons are dense in $\mathcal{C}^0(K)$ for any compact K, and consequently in L^p.

From the functional-analytic viewpoint, this is a density result: the set of functions representable by a neural network architecture is dense in the spaces we care about. The proof, which we will see later in the course, is a beautiful application of the Hahn-Banach theorem — one of the central results in duality theory. We will develop this connection in detail in the chapter on Neural Network Connections.

1.7 *Topological Equivalence of Norms*

A single vector space can carry many different norms. Two norms are *equivalent* when they induce the same topology — the same open sets, the same convergent sequences, the same continuous functions. In finite dimensions this never matters: all norms on $\mathbb{R}^n$ are equivalent. In infinite dimensions it matters enormously, and the choice of norm determines which operators are bounded and which sequences converge.

1.24 DEFINITION (Equivalent Norms). Two norms $\|\cdot\|^1$ and $\|\cdot\|^2$ on a vector space X are **equivalent** if there exist constants $a, b > 0$ such that

$$a\|x\|^1 \le \|x\|^2 \le b\|x\|^1 \quad \forall x \in \mathrm{X}. \tag{27}$$

Equivalent norms induce the same underlying topology on X: they generate the same open sets and hence the same notion of "close."

Geometric intuition. The condition $a\|x\|^1 \le \|x\|^2 \le b\|x\|^1$ says the unit ball of one norm can be scaled to sit inside the unit ball of the other, and vice versa. In $\mathbb{R}^2$ the ℓ^1 ball (a diamond), the ℓ^2 ball (a circle), and the ℓ^∞ ball (a square) all have different shapes — yet each can be scaled to contain, or be contained in, any of the others. This mutual containment of balls is exactly equivalence of topologies.

1.25 EXAMPLE (Unit balls in). $\mathbb{R}^n$ **fit inside each other**

For $x \in \mathbb{R}^n$ the ℓ^p norms satisfy the pointwise inequalities

$$\|x\|_\infty \le \|x\|_2 \le \|x\|_1 \le n\,\|x\|_\infty, \qquad \|x\|_2 \le \sqrt{n}\,\|x\|_\infty, \qquad \|x\|_1 \le \sqrt{n}\,\|x\|_2. \tag{28}$$

Writing $\mathrm{B}_p = \{x : \|x\|_p \le 1\}$ for the closed unit ball of $\|\cdot\|_p$, the smaller the norm, the larger the ball. The inequalities above translate directly into the nested inclusions

$$\mathrm{B}_1 \subset \mathrm{B}_2 \subset \mathrm{B}_\infty \subset \sqrt{n}\,\mathrm{B}_2 \subset n\,\mathrm{B}_1. \tag{29}$$

So the diamond, the disk, and the square each fit inside a rescaled copy of the others. Pictorially in $\mathbb{R}^2$: the diamond B_1 is inscribed in the disk B_2, which is inscribed in the square B_∞; conversely, the square shrunk by $\sqrt{2}$ fits back inside the disk, and the disk shrunk by 2 fits back inside the diamond. Because the scaling factors $1, \sqrt{n}, n$ are *finite*, each pair of ℓ^p norms on $\mathbb{R}^n$ is equivalent.

1.26 THEOREM (Equivalence of Norms in). $\mathbb{R}^n$

On $\mathbb{R}^n$ all norms are equivalent.

Proof. Let $\|\cdot\|$ be any norm on $\mathbb{R}^n$ and $\|\cdot\|_2$ the Euclidean norm. Writing $x = \sum x_i e_i$ and using the triangle inequality, $\|x\| \le \sum |x_i|\,\|e_i\| \le b\,\|x\|_2$ where $b = (\sum \|e_i\|^2)^{1/2}$. In particular $\|\cdot\|$ is continuous with respect to $\|\cdot\|_2$. The Euclidean unit sphere $\mathrm{S} = \{x : \|x\|_2 = 1\}$ is compact (Heine-Borel), so the continuous function $x \mapsto \|x\|$ attains a positive minimum $a > 0$ on S. Homogeneity gives $a\|x\|_2 \le \|x\| \le b\|x\|_2$ for all x, so $\|\cdot\| \sim \|\cdot\|_2$, and hence any two norms on $\mathbb{R}^n$ are equivalent.

□

1.27 REMARK (Failure in Infinite Dimensions). This theorem fails in infinite dimensions. For instance, on $\mathcal{C}^0([0,1])$ the norms $\|\cdot\|_\infty$ and $\|\cdot\|_1$ are **not** equivalent: a sequence of tall, narrow spikes can converge in L^1 without converging uniformly. The choice of norm — and hence the topology — matters fundamentally in infinite-dimensional analysis and determines which operators are bounded, which functionals are continuous, and which sequences converge.

The proof above breaks down exactly because the unit sphere of one norm need not be compact in the topology of another. Equivalence of norms on $\mathbb{R}^n$ and compactness of the unit ball are two sides of the same finite-dimensional coin.

1.8 *Compactness*

Almost every existence argument in analysis follows the same three-step template:

1. **Construct an approximate sequence** of "almost-solutions" (minimizing sequences, Galerkin approximations, Euler polygonal approximations, maximizers of a Rayleigh quotient, etc.).
2. **Extract a convergent subsequence.** This is where compactness enters.
3. **Pass to the limit.** Show the limit actually solves the original problem.

In finite dimensions step 2 is free: Bolzano-Weierstrass guarantees that every bounded sequence in $\mathbb{R}^n$ has a convergent subsequence. This is why eigenvalue decompositions, the SVD, the direct method of the calculus of variations, and Peano's ODE existence theorem all "just work" in $\mathbb{R}^n$. Each reduces to building an approximate sequence on a compact set and extracting a convergent subsequence. In infinite dimensions, closed bounded sets are no longer compact, and step 2 fails without additional structure.

1.28 DEFINITION (Compact, precompact, and totally bounded sets). Let (M, d) be a metric space and $A \subset M$.

1. A is **compact** if every open cover of A has a finite subcover, or equivalently, if every sequence in A has a subsequence converging to a point in A.
2. A is **precompact** (or **relatively compact**) if its closure $\overline{A}$ is compact, or equivalently, if every sequence in A has a subsequence that is Cauchy.
3. A is **totally bounded** if for every $\varepsilon > 0$, A can be covered by finitely many balls of radius ε.

The key equivalence is:

> A **is compact** $\iff$ A **is complete and totally bounded.**

In a complete metric space (such as a Banach space), precompactness and total boundedness coincide: A is precompact if and only if it is totally bounded.

Why both conditions are needed.

- *Total boundedness without completeness fails.* The open interval $(0, 1)$ is totally bounded in $\mathbb{R}$, but the sequence $1/n$ converges to $0 \notin (0, 1)$. The limit escapes through a "hole."
- *Completeness without total boundedness fails.* The unit ball of ℓ^2 is complete, but the orthonormal sequence (e_n) has all elements distance $\sqrt{2}$ apart, so no finite ε-cover works for $\varepsilon < \sqrt{2}/2$. There are "too many directions."

Total boundedness is the condition that makes a set **approximately finite-dimensional**: at every resolution ε, the set is indistinguishable from a finite set. Completeness then ensures that Cauchy subsequences (built by the pigeonhole principle at finer and finer scales) actually converge.

In finite dimensions, compactness is characterized by the Heine-Borel theorem: a set is compact if and only if it is closed and bounded.

1.29 EXAMPLE (Compactness of Unit Ball in). $\mathbb{R}^n$

The closed unit ball $B_1(0) = \{x \in \mathbb{R}^n : \|x\|_p \leq 1\}$

is compact by Heine-Borel. Equivalently, we can take any sequence $\{x_n\} \subset B_1(0)$. Since that sequence is bounded Bolzano-Weierstrass implies that it has a convergent subsequence.

This picture changes dramatically in infinite dimensions.

1.30 EXAMPLE (The Unit Ball in). ℓ^2 **is Not Compact**

Consider the Hilbert space ℓ^2 of square-summable sequences with the standard orthonormal basis $\{e_n\}_{n=1}^{\infty}$ where $e_n = (0, \ldots, 0, 1, 0, \ldots)$ has a 1 in position n and zeros elsewhere.

Each e_n lies in the closed unit ball since $\|e_n\|_2 = 1$. However, for any $n \neq m$:

$$\|e_n - e_m\|_2 = \sqrt{\|e_n\|_2^2 + \|e_m\|_2^2} = \sqrt{2} \tag{30}$$

Every pair of distinct elements is at distance $\sqrt{2}$ apart. Therefore no subsequence of $\{e_n\}$ can be Cauchy, and hence no subsequence converges. The closed unit ball in ℓ^2 is **not** compact.

1.31 REMARK (Compactness in Infinite Dimensions). The failure of the Heine-Borel theorem in infinite dimensions is one of the most important distinctions between finite- and infinite-dimensional analysis. Closed and bounded sets need not be compact, and establishing compactness requires additional structure (e.g. equicontinuity in the Arzelà-Ascoli theorem, or tightness in probability theory). For this reason, compactness arguments play a crucial and often delicate role throughout functional analysis.

1.32 REMARK (Compact sets in infinite dimensions). Boundedness alone is not enough. The unit ball of ℓ^2 is closed and bounded, yet the orthonormal sequence (e_n) has every pair distance $\sqrt{2}$ apart, so no subsequence can be Cauchy. You need a constraint that suppresses this "too many directions" problem, i.e. you need total boundedness (Definition 1.28).

A concrete and important example: the unit ball of $H^1(\Omega)$, viewed inside $L^2(\Omega)$, is precompact (Rellich-Kondrachov). The reason is visible in Fourier space. An H^1 bound forces $|\hat{u}(k)|^2 \leq C/(1+|k|^2)$, so high-frequency modes are uniformly suppressed. For any $\varepsilon > 0$, all the L^2 energy beyond a sufficiently large frequency N is less than ε, uniformly over the entire H^1 ball. The remaining low-frequency part lives in a finite-dimensional space. This is total boundedness: the derivative bound kills oscillations that would otherwise prevent clustering, and the bounded domain prevents mass from escaping to spatial infinity.

1.9 *Compactness in Function Spaces*

In infinite dimensions, "bounded" does not imply "precompact" (Example 1.30). To pin down compactness in the two function spaces we care about ($C(K)$ and L^p) we need additional structure. Two classical theorems give the right structure: Arzelà-Ascoli for continuous functions, and its L^p analog Kolmogorov-Riesz for integrable functions. Both will be cited repeatedly throughout the course.

1.9.1 Arzelà-Ascoli for C(K)

1.33 THEOREM (Arzelà-Ascoli). *Let K be a compact metric space and $\mathcal{F} \subset C(K)$ a family of continuous functions. Then $\mathcal{F}$ is precompact (in the sup norm) if and only if*

1. ***(pointwise bounded)*** $\sup_{f\in\mathcal{F}} |f(x)| < \infty$ *for every* $x \in K$,
2. ***(equicontinuous)*** *for every* $\varepsilon > 0$ *there exists* $\delta > 0$ *such that* $|f(x) - f(y)| < \varepsilon$ *whenever* $d(x, y) < \delta$, *for all* $f \in \mathcal{F}$ *simultaneously.*

The two conditions prevent the two ways a sequence of continuous functions can fail to have a uniformly convergent subsequence: the values cannot blow up (boundedness), and they cannot oscillate arbitrarily fast (equicontinuity). A typical use: bounded sequences of C^1 functions on a compact interval are equicontinuous because $|f(x) - f(y)| \le \|f'\|_\infty |x - y|$, hence precompact in C^0 by Arzelà-Ascoli.

1.9.2 Kolmogorov-Riesz for L^p

The same question arises one floor down. Nesting (Remark 1.11) tells us $L^p(\Omega) \subset L^q(\Omega)$ continuously when $p \ge q$ and $|\Omega| < \infty$. Is this embedding *compact*?

The answer is **no**. On $\Omega = (0, 2\pi)$, the sequence $u_n(x) = \sin(nx)$ has $\|u_n\|_{L^p} \le C$ for every $p \in [1, \infty]$, so it is bounded in every L^p. Yet any two distinct terms satisfy $\|u_n - u_m\|_{L^q} \not\to 0$: the sequence has no convergent subsequence in L^q, for any q. Integrability alone does not suppress oscillation.

The L^p analog of Arzelà-Ascoli gives a complete answer.

1.34 THEOREM (Kolmogorov-Riesz-Fréchet). *Let $1 \le p < \infty$. A bounded set $K \subset L^p(\mathbb{R}^d)$ is precompact if and only if*

1. ***(bounded)*** $\sup_{u\in K} \|u\|_{L^p} < \infty$,
2. ***(equicontinuous under translation)*** $\sup_{u\in K} \|\tau_h u - u\|_{L^p} \to 0$ *as* $h \to 0$, *where* $\tau_h u(x) = u(x - h)$,
3. ***(tight)*** $\sup_{u\in K} \|u\|_{L^p(|x|>R)} \to 0$ *as* $R \to \infty$.

On a bounded domain Ω, condition (3) is automatic, so precompactness reduces to (1) and (2).

Each condition prevents a specific way mass can escape:

- (1) prevents mass from blowing up in size,
- (2) prevents mass from escaping into high frequencies, since a function whose translations converge uniformly cannot oscillate arbitrarily fast,
- (3) prevents mass from escaping to spatial infinity.

This is the L^p mirror of Arzelà-Ascoli: "bounded + equicontinuous + stays in a compact set" gives precompactness, with L^p-equicontinuity-under- translation playing the role of pointwise equicontinuity.

In later chapters, this theorem is the structural reason Sobolev embeddings produce compactness. A gradient bound of the form $\|\nabla u\|_{L^p} \leq M$ supplies condition (2) automatically: by the fundamental theorem of calculus, $\|\tau_h u - u\|_{L^p} \leq |h| \, \|\nabla u\|_{L^p}$. So *the Rellich-Kondrachov theorem is Kolmogorov-Riesz applied to the Sobolev ball*, with the equicontinuity condition upgraded from a hypothesis to a consequence of gradient control.

2 CONTINUOUS LINEAR OPERATORS

Big Idea

Linear operators are the morphisms of functional analysis—they connect spaces, encode differential equations, and determine when problems are well-posed. The three pillars (Banach–Steinhaus, Open Mapping, Closed Graph) give powerful tools for establishing boundedness and invertibility of operators without explicit computation.

2.1 *Overview*

Once we have Banach spaces, the next step is to study **maps between them**. A bounded linear operator $A : X \to Y$ satisfies $\|Ax\|_Y \leq M\|x\|_X$, and the space $\mathcal{L}(X, Y)$ of all such operators is itself a Banach space.

The deep results of this chapter all rely on the **Baire Category Theorem**—a topological result about complete metric spaces. From it we derive:

- **Banach–Steinhaus (Uniform Boundedness):** A pointwise bounded family of operators is uniformly bounded.
- **Open Mapping Theorem:** A surjective bounded operator between Banach spaces is an open map.
- **Closed Graph Theorem:** An operator with a closed graph is automatically bounded.

We also study **compact operators**, which map bounded sets to precompact sets. These operators behave much like matrices in finite dimensions, and their spectral theory is the gateway to Fredholm theory and applications in PDEs.

2.2 *What You Will Learn*

- Bounded vs. continuous operators and the operator norm.
- The Baire Category Theorem and its consequences.
- Banach–Steinhaus, Open Mapping, and Closed Graph Theorems.
- Compact operators and their spectral properties.

2.3 *Bounded Operators*

Linear operators between Banach spaces are key in functional analysis. They connect function spaces with each other, describe physical systems, cost functions, solve PDEs, and tell us when Banach spaces are isomorphic.

2.1 DEFINITION (Linear operator). Let $(X, \|\cdot\|_X)$ and $(Y, \|\cdot\|_Y)$ be two Banach spaces, and let $A : X \mapsto Y$ be a map:

1. A is linear if $A(\alpha x_1 + \beta x_2) = \alpha A x_1 + \beta A x_2$.
2. A is bounded if $\|Ax\|_Y \leq M\|x\|_X \ \forall x \in X$, $M > 0$.
3. We denote the set of all bounded operators by

$$\mathcal{L}(X, Y) = \{A : X \mapsto Y, \ A \text{ linear and bounded}\} \tag{31}$$

This space is equipped with a norm

$$\|A\|_{op} := \sup_{x \neq 0} \frac{\|Ax\|_Y}{\|x\|_X} = \sup_{\|x\|=1} \|Ax\|_Y \tag{32}$$

2.2 PROPOSITION. $\mathcal{L}(X, Y)$ *is a Banach Space*

2.3 PROPOSITION (Qualitative vs Quantitative properties). *Let* $A : X \mapsto Y$ *be linear, then* A *is bounded if and only if* A *is continuous.*

2.4 REMARK. This is the first result of the qualitative vs. quantitative theme we encounter in functional analysis. With qualitative we here mean that A is continuous i.e. well behaved, and that this is related to a quantitative bound on the operator which quantifies how well behaved the operator is. In terms of continuity it tells us how much the output can vary as we change the input.

2.5 EXAMPLE. Insert example of integral operator.

2.6 EXAMPLE. Insert example of second order differential operator.

2.7 DEFINITION. Let $A : X \mapsto Y$ we define

1. The **range** of the operator as $R[A] = \{g \in Y : g = A[f], \ f \in D(A)\}$
2. The **nullspace** or **kernel** of A is $N[A] = \{f \in D(A) : A[f] = 0\}$.

2.4 *The Baire Category Theorem*

In this section we introduce the Baire category theorem. Ultimately the Baire category theorem is used to answer the question when is a subset $E \subset X$ in a topological space small, which turns out to be a fundamental tools in the study of linear bounded operators. This section and in particular the qualiatative and quantitative discussion is inspired by T. Tao [4].

2.4.1 Small sets in a Measure Space

In a measure theory setting i.e. $X = (X, \mathcal{F}, \mu)$ the "small" sets are the nullsets i.e. $E \subset X$ such that $\mu(E) = 0$ or subsets of nullsets. Countable additivity implies that countable unions of nullsets are of measure zero. This gives us results like this:

2.8 LEMMA. *Let* $E_1, E_2, \ldots$ *be an at most countable sequence of measurable sets of a measure* X. *If* $\mu\left(\bigcup_n E_n\right) > 0$ *then at least one of the* E_n *has positive measure.*

The next proposition bridges between measure theory and topology. To motivate it, recall that the Lebesgue differentiation theorem tells us that for almost every point $x \in E$, the *density* of E at x,

$$\lim_{r \to 0} \frac{\mu(E \cap B(x, r))}{\mu(B(x, r))} = 1, \tag{33}$$

equals one. Such points are called *density points* of E — points where E fills up nearly all of every sufficiently small ball. One can think of density points as measure-theoretic interior points: an interior point of E is one where E contains an entire ball, while a density point is one where E nearly fills every small ball, even if it doesn't contain one outright. The next proposition globalizes this local statement: if E has positive measure, then E must be *dense in some ball* in the measure-theoretic sense that it occupies an arbitrarily large fraction of that ball's volume.

2.9 PROPOSITION. *Let* E *be a measurable subset of* $\mathbb{R}^d$. *Then the following are equivalent*

1. $\mu(E) > 0$
2. $\forall \varepsilon > 0$ *there exists a ball* B *such that* $\mu(E \cap B) \geq (1 - \varepsilon)\mu(B)$

Reading the equivalence from both sides:

- **(Forward):** If $\mu(E) > 0$, then E is measure-theoretically dense in some ball — it nearly fills that ball.
- **(Reverse):** If E is not measure-theoretically dense in any ball — that is, in every ball, E always fails to fill a definite fraction — then $\mu(E) = 0$.

A set of measure zero is therefore *nowhere dense in the measure-theoretic sense*: it is never concentrated in any ball. This suggests that to formalize "smallness" in a purely topological setting, we should look for the topological analogue of this property.

2.4.2 Small Sets in Topological Spaces

To make the analogy with the measure-theoretic setting precise, we first recall what *dense* means in a topological space.

2.10 DEFINITION (Dense in an open set). A set E in a topological space is *dense in an open set* U if every nonempty open subset of U intersects E, or equivalently, if the closure of E contains U.

2.11 REMARK. Compare this with the measure-theoretic notion from the previous section: there, E is "dense in B" if $\mu(E \cap B) \approx \mu(B)$, i.e. E fills almost all the volume of B. Here, E is dense in B if E meets every open subset of B, i.e. E is topologically everywhere in B. Volume has been replaced by topology, but the intuition is the same: E is thoroughly present throughout B, with no region where E is absent.

2.12 DEFINITION (Nowhere dense). A set E is *nowhere dense* if it is not dense in any ball, or equivalently, if its closure has empty interior.

This is the topological counterpart of a measure-zero set: just as a null set is never measure-theoretically dense in any ball (it never fills a large fraction of any ball), a nowhere dense set is never topologically dense in any ball (its closure contains no open set).

2.13 THEOREM (Baire Category Theorem). *Let* $E_1, E_2, \ldots$ *be an at most countable sequence of subsets of a complete metric space* X. *If* $\bigcup_n E_n$ *contains a ball* B, *then at least one of the* E_n *is dense in a sub-ball* $B' \subset B$ *(and in particular not nowhere dense). To put it in the contrapositive: the countable union of nowhere dense sets cannot contain a ball.*

Note that completeness is essential: on an incomplete metric space, the Baire property can fail, and the space itself may be a countable union of nowhere dense sets. This is why completeness is such an important hypothesis throughout the theory of bounded operators. For a proof, see [3].

2.14 DEFINITION (Meagre (First Category)). A set E in a topological space is *meagre* (or *of the first category*) if it can be written as a countable union of nowhere dense sets. A set that is not meagre is called *non-meagre* (or *of the second category*).

2.15 REMARK. A meagre set need not be nowhere dense itself. For example, $\mathbb{Q} \subset \mathbb{R}$ is meagre — it is the countable union $\bigcup_{q \in \mathbb{Q}} \{q\}$ of nowhere dense singletons — but $\mathbb{Q}$ is *not* nowhere dense, since $\overline{\mathbb{Q}} = \mathbb{R}$ has nonempty interior. Thus "meagre" is a strictly weaker notion than "nowhere dense": every nowhere dense set is meagre (a union of one nowhere dense set), but not conversely.

2.16 COROLLARY (A complete metric space is not meagre). *Let* X *be a complete metric space and* $\{F_j\}$ *a countable collection of nowhere dense sets. Then* $\bigcup_j F_j \neq X$. *Equivalently, a complete metric space is not meagre (i.e. not a countable union of nowhere dense sets).*

2.17 EXAMPLE (The rationals are meagre but not nowhere dense). On $X = \mathbb{R}$ (which is complete), the rationals $\mathbb{Q} = \bigcup_{q \in \mathbb{Q}} \{q\}$ are a countable union of nowhere dense sets, hence meagre. But $\mathbb{Q}$ is *not* nowhere dense: $\overline{\mathbb{Q}} = \mathbb{R}$, which has nonempty interior. Corollary 2.16 tells us $\mathbb{Q} \neq \mathbb{R}$, which we know. More importantly: the irrationals $\mathbb{R} \setminus \mathbb{Q}$ are a dense G_δ set (a countable intersection of open dense sets), and $\mathbb{Q}$ cannot be G_δ.

2.18 EXAMPLE (The Cantor set is meagre and nowhere dense). The middle-thirds Cantor set $C \subset [0, 1]$ is closed and has empty interior (it contains no interval), so it is nowhere dense. Since every nowhere dense set is trivially meagre (a "union" of one nowhere dense set), C is also meagre. Yet C is uncountable

— so a set can be topologically small (nowhere dense, meagre) while being large in cardinality. Contrast this with $\mathbb{Q}$, which is meagre but *not* nowhere dense.

2.19 EXAMPLE (Proper subspaces). **Why do we care?**

The consequence of this theorem is that $\mathbb{R}^d$ or $\mathbb{C}^d$ cannot be covered by a countable union of proper subspaces.

2.20 DEFINITION. $W \subset V$ is a proper subspace of V such that

1. A subset V i.e. $W \subset V$.
2. Strictly smaller than V.
3. Still a vector space.

In $\mathbb{R}^3$ examples of proper subspaces are:

1. The xy-plane.
2. A line through the origin.

Baire says that we can't write $\mathbb{R}^3$ as a countable union of these nowhere dense proper subspaces.

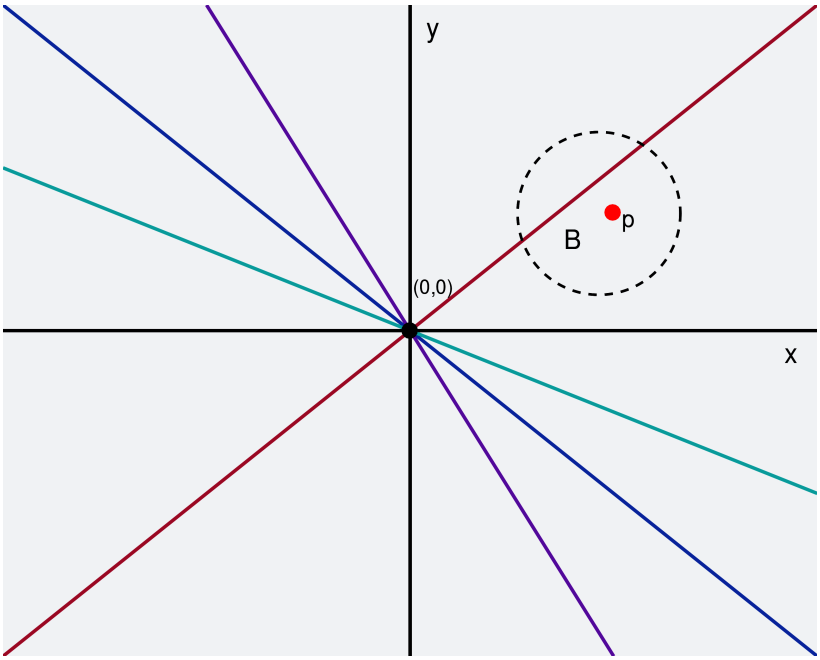


Figure 1: $\mathbb{R}^2$ cannot be covered by a countable union of proper subspaces (lines through the origin).

2.21 REMARK. Note that of course we could also prove these results using measure theory in finite dimensional spaces by using countable additivity of countable unions. The *advantage* of Baire's theorem is that it easily extends to infinite dimensional vectorspaces.

2.22 REMARK. The Baire Category Theorem holds in two important classes of topological spaces:

1. **Complete metric spaces** (and hence Banach spaces, Hilbert spaces, etc.).

2. **Locally compact Hausdorff spaces** (and hence $\mathbb{R}^d$, $\mathbb{C}^d$, and all finite-dimensional normed spaces).

In finite dimensions these overlap: $\mathbb{R}^d$ is both complete and locally compact. In infinite dimensions, however, a Banach space is *never* locally compact (the closed unit ball is not compact), so completeness is the only route to the Baire property.

2.4.3 Consequences of the Baire Category Theorem

The Baire Category Theorem leads to three fundamental equivalences between the qualitative theory of continuous linear operators on Banach spaces e.g. finiteness, surjectivity to the quantitative theory e.g. estimates [4]. We have already seen the first of these relationships. Recall that for a linear operator $A : X \mapsto Y$ we have that A is continuous $\iff$ A is bounded.

As we saw with the examples this allows us to prove the continuity of an operator A by simply establishing that it is bounded (which usually is easier).

1. *Uniform boundedness principle* that equates the qualitative operators with their quantitative boundedness.

2. *Open mapping theorem* equates qualitative solvability of a linear problem $Lu = f$ with quantitative solvability.

3. *Closed graph theorem* equates the qualitative regularity of a (weakly continuous) operator T with the quantitative regularity of the operator.

2.23 REMARK. Note that these results are not used much in practice, because one usually works in the reverse direction i.e. first prove bounds, and then derive the operator's qualitative properties. But crucially these results provide the motivation or justification of why we approach qualitative problems in functional analysis via their quantitative counter parts.

Next we will explore the uniform boundedness principle, and the open mapping theorem while we leave the closed graph theorem for later.

2.5 *The uniform boundedness principle (Banach-Steinhaus)*

2.24 THEOREM (Banach-Steinhaus). *Let* X *be a Banach space, and* Y *be a normed space and let* $(T_\alpha)_{\alpha\in A}$ *be a family of bounded linear operators* $T_\alpha : X \mapsto Y$. *Then the following are equivalent:*

1. *(Pointwise boundedness)* $\forall x \in X$ *the set* $\{T_\alpha x : \alpha \in A\}$ *is bounded.*
2. *(Uniform boundedness) The operator norms* $\{\|T_\alpha\|_{op}\}$ *are bounded.*

2.25 REMARK (Significance of Banach-Steinhaus).

1. The condition that $\{\|T_\alpha\| : \alpha \in A\}$ is bounded is equivalent to the family $\{T_\alpha\}$ being equicontinuous.
2. The theorem shows that pointwise control implies uniform control.
3. From the perspective of qualitative vs. quantitative: the qualitative statement says the family of operators is somehow well-behaved, and the quantitative perspective puts a particular number on the behaviour of the operator.
4. A version of this theorem with specific rates allows us to establish and control the convergence of sequences of approximation operators e.g. quadrature approximations to integrals, or finite difference approximations to derivatives.
5. From a metamathematical view it explains why sequential arguments often work in functional analysis. In other words, if the uniform limit didn't exist we could construct sequences that behave pathologically. For more details, see the subsequent corollary.

2.5.1 Why Does Banach-Steinhaus Need Completeness?

In finite dimensions, Banach-Steinhaus is trivial: it needs no topology, no Baire, not even continuity. In infinite dimensions completeness is essential.

2.5.2 Why This Matters in Practice

The contrapositive is often more useful: if the family is uniformly unbounded, then there exists a concrete point x where $\sup_\alpha \|T_\alpha x\| = \infty$. The blowup cannot remain invisible, it must concentrate at an actual point in the space. Moreover, the set of such witness points is **comeager** (residual) by Baire, meaning witnesses are generic and the well-behaved points are the topologically rare ones.

Banach-Steinhaus turns abstract collective pathology into concrete individual pathology. An incomplete space is itself meagre, so it can be exhausted by the nowhere dense sets F_j and no witness needs to exist. A complete space is not meagre, so Baire forces some F_j to have interior, and the witness is not only guaranteed but *generic*.

A common application is to converging sequences of bounded linear operators.

2.26 COROLLARY (Limits of sequences of bounded linear operators). *Let* X, Y *be Banach spaces, and let* $(T_n)_{n=1}^\infty$ *be a family of bounded linear operators* $T_n : X \mapsto Y$. *Then the following are equivalent:*

1. *(Pointwise convergence).* $\forall x \in X$, $T_n x \to y$ *in* Y.
2. *(Pointwise convergence to a continuous limit). There exists a bounded linear operator* $T : X \mapsto Y$ *such that* $\forall x \in X$, $T_n x \to Tx$ *in* Y.

3. *(Uniform boundedness + Dense subclass convergence). The operator norms* $\{\|\mathrm{T}_n\| : n = 1, 2, \dots\}$ *are bounded, and for a dense subset of* X, $\mathrm{T}_n x$ *converges in* Y.

2.6 *The Open Mapping Theorem*

The open-mapping theorem is another of the big theorems of functional analysis. Its significance is that it equates qualitative solvability of a linear problem $Lx = y$ with quantitative solvability. We will begin by stating the qualitative version and explore its consequences on the solvability of linear operators, followed by the quantitative version.

2.27 DEFINITION (Open map). A map $f : X \to Y$ between topological spaces is called **open** if it maps open sets to open sets, i.e. $f(U)$ is open in Y whenever U is open in X.

Geometrically, openness means **the image of every ball contains a ball**: sets with interior keep their interior, nothing gets crushed to zero volume. A map that is not open flattens the unit ball into something with empty interior, and any attempt to invert such a map is necessarily discontinuous since small perturbations in Y can jump to distant points in X.

2.28 THEOREM (Open Mapping Theorem). *Let* U, V *be the open unit balls of the Banach spaces* X, Y *respectively. To every bounded linear operator* L *of* X ***onto*** Y *there corresponds* $\delta > 0$ *so that* $L(U) \supset \delta V = \{y \in Y : \|y\| < \delta\}$.

The quantitative content is that there is a *uniform* lower bound $B_\delta(0) \subset L(B_1(0))$ on how much volume survives. In the bijective case, $\delta = 1/\|L^{-1}\|_{\text{op}}$. By the linearity of L this extends to all open sets: L maps every open set to an open set, i.e. L is an open mapping.

2.29 COROLLARY (Banach's Bounded Inverse Theorem). *Let* L *be a linear operator between two Banach spaces* X *and* Y*, that is* ***onto*** *and one-to-one, then* L^{-1} *is a bounded linear operator.*

The previous statement is remarkable in the sense that it lifts our intuition about linear operators in finite dimensions (i.e. that bijectiveness is equivalent to invertibility) to the infinite dimensional case.

2.30 COROLLARY (Equivalence of norms). *If* $\|\cdot\|_1$ *and* $\|\cdot\|_2$ *are two norms on a vector space* X*, both making it a Banach space, and* $\|x\|_2 \leq C\|x\|_1$ *for all* x*, then the norms are equivalent: there exists* $C' > 0$ *such that* $\|x\|_1 \leq C'\|x\|_2$ *for all* x*.*

2.31 REMARK. Normally, proving norm equivalence requires two inequalities. This corollary says that if both norms make the space complete, one inequality gives you the other for free. In particular, you cannot have two genuinely different Banach space topologies on the same space where one is strictly finer than the other: if a stronger norm and a weaker norm both give completeness, they must be equivalent.

2.6.1 Why Does the Open Mapping Theorem Need Completeness?

In finite dimensions the Open Mapping Theorem is almost trivial: compactness of the unit ball does all the work. In infinite dimensions compactness fails, and completeness must substitute via Baire's category theorem. The details below make this precise.

2.6.2 Quantitative Formulation

Surjectivity alone is qualitative: "solutions exist." Openness adds quantitative content: "solutions exist *with controlled norm*," i.e. there is a constant C such that every y has a preimage x with $\|x\| \leq C\|y\|$. The

Bounded Inverse Theorem sharpens this further when L is also injective: the preimage is unique and $\|x\| \leq \|L^{-1}\|_{op}\|y\|$.

In finite dimensions, a bijective linear map is automatically invertible with a bounded inverse. In infinite dimensions this fails: the c_{00} counterexample above is a bounded bijection whose inverse is unbounded. "Invertible" should really mean *stably* invertible, i.e. L^{-1} is also bounded. The Open Mapping Theorem is the bridge: for Banach spaces, surjectivity + boundedness automatically yields openness, and openness + injectivity gives a bounded inverse. So the algebraic notion of invertibility (bijection) coincides with the analytic notion (bounded inverse), but only in the Banach space setting.

This is why the Open Mapping Theorem matters so much for PDEs: it is not enough to know $Lu = f$ has a solution. You need to know $\|u\| \leq C\|f\|$, that the solution depends continuously on the data. Open Mapping says surjectivity *automatically* gives you this stability, as long as you are working in Banach spaces.

Now to the version of the open-mapping theorem by T. Tao [4].

2.32 THEOREM (Quantitative and Qualitative version of the Open Mapping Theorem). *Let* $L : X \mapsto Y$ *be a bounded linear operator between two Banach spaces, then the following are equivalent:*

1. L *is onto.*
2. L *is open.*
3. *(Qualitative solvability). For every* $y \in Y$ *there exists (not necessarily unique) a solution* $x \in X$ *to the equation* $Lx = y$.
4. *(Quantitative solvability).* $\exists C > 0 : \forall y \in Y,\ \exists x \in X$ *to* $Lx = y$ *such that* $\|x\| \leq C\|y\|$.
5. *(Quantitative solvability for dense subsets). There exists a* $C > 0$ *such that for a dense set of* y *in* Y *there exists solutions* $x \in X$ *to* $Lx = y$ *which obeys* $\|x\| \leq C\|y\|$.

2.33 REMARK. Make sure to take a close look at the proof of the open-mapping theorem. There are three key techniques on display: (1) **linearity**, (2) **series approximations**, and (3) the **Baire category theorem**.

1. The frequent dilations and translations in the above are allowed by linearity.
2. The series approximation technique has profound significance beyond proving results in functional analysis. Take a look at numerical methods such as gradient descent, conjugate gradient, or even Newton's method. These are all examples of methods that construct solution approximations by constructing approximating series.
3. Completeness enters twice. **Completeness of** Y feeds Baire's theorem: surjectivity gives $Y = \bigcup_n \overline{T(nB_X)}$, and Baire forces at least one of these closed sets to have nonempty interior, so $\overline{T(B_X)}$ contains a ball. **Completeness of** X then promotes from the closure to the set itself: approximate preimages are assembled into a geometrically decreasing series whose partial sums are Cauchy, and completeness guarantees convergence to an exact preimage with controlled norm. In the c_{00} counterexample, each $\overline{T(nB)} = \{y : |y_k| \leq n/k\}$ is nowhere dense (perturb the k-th coordinate for large k to escape), so c_{00} is covered by nowhere dense sets. Baire forbids this in a complete space, but c_{00} is incomplete, so the argument has no traction.

2.34 REMARK (Condition numbers). In numerical linear algebra, the central quantity controlling stability of solving $Lx = y$ is the **condition number** $\kappa(L) = \|L\| \cdot \|L^{-1}\|$. If the right-hand side y is perturbed by relative error ε, the solution can be perturbed by up to $\kappa(L) \cdot \varepsilon$. For finite-dimensional invertible matrices, $\kappa(A) < \infty$ automatically, and the whole game is about *how large* κ is. In infinite dimensions, the Bounded Inverse Theorem is what guarantees $\|L^{-1}\|_{\mathrm{op}} < \infty$ (and hence $\kappa(L) < \infty$) *at all*. The Open Mapping Theorem is the qualitative floor beneath condition numbers: it says the problem is well-conditioned before you even start worrying about how well-conditioned it is.

This has far-reaching consequences for numerical methods. When we discretize $Lu = f$ by a sequence of finite-rank operators L_n and solve $L_n u_n = f_n$, the Open Mapping Theorem guarantees well-posedness of the continuous problem, Banach-Steinhaus guarantees stability of convergent schemes (this is the content of the **Lax equivalence theorem**: convergence $\iff$ stability), and the Neumann series controls condition numbers of the discrete problems. We will return to these connections in the applications.

2.7 *Quotient Spaces and Codimension*

When we study a linear operator $T : X \to Y$, its kernel $\ker(T)$ captures what T "forgets." The quotient space $X/\ker(T)$ is the space that remains after collapsing everything T cannot distinguish.

2.7.1 Quotient Spaces

2.35 DEFINITION (Quotient space). Let X be a vector space and $M \subset X$ a subspace. The **quotient space** X/M is the set of equivalence classes

$$[x] = x + M = \{x + m : m \in M\} \tag{34}$$

under the equivalence relation $x \sim y \iff x - y \in M$. This is a vector space with operations

$$[x] + [y] = [x + y], \qquad \lambda[x] = [\lambda x]. \tag{35}$$

The **canonical projection** $\pi : X \to X/M$ defined by $\pi(x) = [x]$ is a surjective linear map with $\ker(\pi) = M$.

The elements of X/M are the "cosets" or translates $x + M$. Two vectors x and y represent the same coset precisely when $x - y \in M$ — that is, when they differ by something in M.

2.36 EXAMPLE (Quotient in). $\mathbb{R}^3$

Let $X = \mathbb{R}^3$ and $M = \{(x, y, 0) : x, y \in \mathbb{R}\}$ (the xy-plane). Then $X/M \cong \mathbb{R}$: each coset is a horizontal plane at height z, and the quotient map sends $(x, y, z) \mapsto z$. All the x- and y-information is "collapsed" — only the height survives.

2.37 EXAMPLE (Quotient by a kernel). Let $f : \mathbb{R}^3 \to \mathbb{R}$ be defined by $f(x, y, z) = 2x + 3y - z$. Then $\ker(f) = \{2x + 3y - z = 0\}$ is a plane through the origin. The cosets of $\ker(f)$ are the level sets $f^{-1}(c) = \{2x + 3y - z = c\}$ — parallel planes, each labeled by the value c. The quotient $\mathbb{R}^3/\ker(f)$ is one-dimensional, reflecting the fact that f maps onto $\mathbb{R}$.

THE QUOTIENT NORM When X is a normed space and M is a **closed** subspace, the quotient inherits a natural norm.

2.38 DEFINITION (Quotient norm). Let X be a normed space and $M \subset X$ a closed subspace. The **quotient norm** on X/M is

$$\|[x]\|_{X/M} = \inf_{m \in M} \|x - m\| = \operatorname{dist}(x, M). \tag{36}$$

The closedness of M ensures this is a genuine norm (not just a seminorm): if $\|[x]\|_{X/M} = 0$, then $\operatorname{dist}(x, M) = 0$, so $x \in \overline{M} = M$, hence $[x] = [0]$.

2.39 PROPOSITION. *If* X *is a Banach space and* $M \subset X$ *is a closed subspace, then* X/M *is a Banach space.*

Proof. Let $([x_n])$ be a Cauchy sequence in X/M. Passing to a subsequence, we may assume $\|[x_{n+1}] - [x_n]\|_{X/M} < 2^{-n}$. By definition of the infimum, choose representatives: pick $y_1 = x_1$, and for each n pick $m_n \in M$ such that $\|x_{n+1} - x_n - m_n\| < 2^{-n}$. Define $y_{n+1} = x_{n+1} - (m_1 + \cdots + m_n)$; then $[y_n] = [x_n]$ and $\|y_{n+1} - y_n\| < 2^{-n}$, so (y_n) is Cauchy in X. Since X is complete, $y_n \to y$ for some $y \in X$. Then $[x_n] = [y_n] \to [y]$ in X/M.

□

2.7.2 Codimension

2.40 DEFINITION (Codimension). Let M be a subspace of a vector space X. The **codimension** of M in X is

$$\operatorname{codim}(M) = \dim(X/M). \tag{37}$$

Equivalently, $\operatorname{codim}(M) = k$ means there exist vectors $e_1, \ldots, e_k \in X$ such that

$$X = M \oplus \operatorname{span}(e_1, \ldots, e_k), \tag{38}$$

and k is the smallest number for which such a decomposition exists.

Codimension measures how many dimensions are "missing" from M. A codimension-1 subspace is as large as possible while still being proper — it is a **hyperplane** through the origin.

2.41 EXAMPLE. In $\mathbb{R}^n$:

- A line through the origin has codimension $n - 1$.
- A plane through the origin in $\mathbb{R}^3$ has codimension 1.
- The subspace $\{0\}$ has codimension n.

2.7.3 The First Isomorphism Theorem

2.42 THEOREM (First Isomorphism Theorem). *Let* $T : X \to Y$ *be a linear map. Then* T *induces an isomorphism*

$$\widetilde{T} : X/\ker(T) \xrightarrow{\cong} \operatorname{Range}(T) \tag{39}$$

defined by $\widetilde{T}([x]) = Tx$. *This map is well-defined, linear, injective, and surjective onto* Range(T).

If X *and* Y *are Banach spaces,* T *is bounded, and* Range(T) *is closed, then* $\widetilde{T}$ *is a topological isomorphism (i.e.,* $\widetilde{T}$ *and* $\widetilde{T}^{-1}$ *are both bounded). This follows from the Theorem 2.28 applied to the bijective bounded operator* $\widetilde{T}$.

Proof. **Well-defined:** If $[x_1] = [x_2]$, then $x_1 - x_2 \in \ker(T)$, so $Tx_1 = Tx_2$.

Linear: $\widetilde{T}(\alpha[x] + \beta[y]) = \widetilde{T}([\alpha x + \beta y]) = T(\alpha x + \beta y) = \alpha Tx + \beta Ty = \alpha\widetilde{T}([x]) + \beta\widetilde{T}([y])$.

Injective: $\widetilde{T}([x]) = 0 \implies Tx = 0 \implies x \in \ker(T) \implies [x] = [0]$.

Surjective onto Range(T)**:** For any $y = Tx \in \operatorname{Range}(T)$, we have $\widetilde{T}([x]) = y$.

Bounded (when T **is bounded):** For any $m \in \ker(T)$, $\|Tx\| = \|T(x - m)\| \le \|T\|\|x - m\|$. Taking the infimum over m: $\|\widetilde{T}([x])\| \le \|T\|\,\|[x]\|$. So $\|\widetilde{T}\| \le \|T\|$.

□

The first isomorphism theorem immediately answers the question: *why does a nonzero linear functional have a codimension-1 kernel?*

2.43 COROLLARY. *Let* $f : X \to \mathbb{R}$ *be a nonzero linear functional. Then* $\operatorname{codim}(\ker(f)) = 1$.

Proof. Since f is nonzero, there exists $x_0 \in X$ with $f(x_0) \neq 0$. For any $c \in \mathbb{R}$, we have $f(cx_0/f(x_0)) = c$, so f is surjective. By the first isomorphism theorem (Theorem 2.42),

$$X/\ker(f) \cong \operatorname{Range}(f) = \mathbb{R}. \tag{40}$$

Therefore $\operatorname{codim}(\ker(f)) = \dim(X/\ker(f)) = \dim(\mathbb{R}) = 1$.

□

The key insight: the codimension equals 1 because the **target** is $\mathbb{R}$, which is one-dimensional. The first isomorphism theorem tells us that the quotient $X/\ker(f)$ is isomorphic to the range, and since f maps onto all of $\mathbb{R}$, the quotient is one-dimensional — regardless of whether X itself is finite- or infinite-dimensional.

2.8 *Compact Operators*

In infinite dimensions, closed bounded sets are no longer compact (Example 1.30), so bounded operators do not in general send bounded sequences to sequences with convergent subsequences. Compact operators are the class for which that *does* hold: they restore Bolzano-Weierstrass as a property of the operator, which is why they play the role of "infinite-dimensional matrices" throughout PDEs, spectral theory, and the calculus of variations. For the underlying metric-space notions of compactness, precompactness, and total boundedness, see Definition 1.28 and Remark 1.32.

2.8.1 Definition and Basic Properties

2.44 DEFINITION (Compact operator). An operator $K : X \to Y$ between normed spaces is **compact** if for each bounded set $U \subset X$, the image $K(U)$ is relatively compact in Y (i.e., $\overline{K(U)}$ is compact).

Equivalently, K is compact if and only if every bounded sequence (x_n) in X has a subsequence (x_{n_k}) such that (Kx_{n_k}) converges in Y.

The sequential characterization is the one we use most often in practice: **bounded sequence in, convergent subsequence out**. Compare this with a general bounded operator, which only guarantees bounded sequence in, bounded sequence out.

2.45 PROPOSITION (Compact operators are bounded). *Every compact operator is bounded.*

Proof. The unit ball $B_1(0) \subset X$ is bounded, so by compactness of the operator, $\overline{K(B_1(0))}$ is compact and hence bounded in Y. Thus

$$\sup_{\|x\| \le 1} \|Kx\|_Y \le M \tag{41}$$

for some $M > 0$, which gives $\|K\|_{\mathrm{op}} \le M$.

□

The converse is false in infinite dimensions: the identity operator $I : \ell^2 \to \ell^2$ is bounded but not compact, since the orthonormal sequence (e_n) is bounded but has no convergent subsequence.

2.8.2 The Space of Compact Operators

2.46 PROPOSITION (Compact operators form a closed ideal). *Let X, Y, Z be normed spaces with Y a Banach space.*

1. *If $K : X \to Y$ is compact and $T : Y \to Z$ is bounded, then TK is compact.*
2. *If $T : Z \to X$ is bounded and $K : X \to Y$ is compact, then KT is compact.*
3. *The set $\mathcal{K}(X, Y)$ of compact operators is a **closed subspace** of $\mathcal{L}(X, Y)$ in the operator norm.*

In algebraic language, $\mathcal{K}(X, Y)$ is a **closed two-sided ideal** in $\mathcal{L}(X)$ when $X = Y$. Composing a compact operator with any bounded operator (on either side) produces another compact operator.

The closedness statement is the most important: it says that the operator-norm limit of compact operators is again compact. This is the key tool for proving compactness of specific operators.

2.47 THEOREM (Limits of compact operators are compact). *Let* X *be a normed space and* Y *a Banach space. If* (K_n) *is a sequence of compact operators in* $\mathcal{L}(X, Y)$ *with* $K_n \to K$ *in the operator norm, then* K *is compact.*

Proof. Let (x_j) be a bounded sequence in X. We use a **diagonal argument**.

Since K_1 is compact, $(K_1 x_j)$ has a convergent subsequence $K_1 x_j^{(1)}$.

Since K_2 is compact, $(K_2 x_j^{(1)})$ has a convergent subsequence $K_2 x_j^{(2)}$.

Continuing, at stage l the sequence $(K_l x_j^{(l)})$ converges. The diagonal sequence $y_j := x_j^{(j)}$ satisfies: $(K_l y_j)_j$ converges for every l.

Now estimate:

$$\|Ky_i - Ky_j\| \le \underbrace{\|Ky_i - K_n y_i\|}_{\le \|K-K_n\|\cdot C} + \underbrace{\|K_n y_i - K_n y_j\|}_{\to 0 \text{ as } i,j\to\infty} + \underbrace{\|K_n y_j - Ky_j\|}_{\le \|K-K_n\|\cdot C} \tag{42}$$

where $C = \sup_j \|y_j\|$. For any $\varepsilon > 0$, first choose n so that $\|K - K_n\| < \varepsilon/3C$, then choose i, j large enough so that $\|K_n y_i - K_n y_j\| < \varepsilon/3$. Thus (Ky_j) is Cauchy in Y and converges since Y is a Banach space. □

2.48 COROLLARY (K(X,Y) is a Banach space). *The space* $\mathcal{K}(X, Y)$ *of compact operators from a normed space* X *into a Banach space* Y *is itself a Banach space.*

2.8.3 Finite-Rank Operators: Compact Operators as "Infinite-Dimensional Matrices"

The connection between compact operators and matrices runs through finite-rank operators.

2.49 DEFINITION (Finite-rank operator). A bounded linear operator $A : X \to Y$ has **finite rank** if its range R(A) is finite-dimensional. We write rank(A) = dim R(A).

2.50 PROPOSITION (Finite-rank operators are compact). *Every bounded operator with finite-dimensional range is compact.*

Proof. Let (x_n) be a bounded sequence in X. Then (Ax_n) is a bounded sequence in the finite-dimensional space R(A). By Bolzano-Weierstrass, it has a convergent subsequence.

□

Finite-rank operators are literally matrices: a matrix $A \in \mathbb{R}^{m\times n}$ has rank at most $\min(m, n)$. On Hilbert spaces, compact operators turn out to be exactly the norm limits of such operators:

2.51 PROPOSITION (Compact operators on Hilbert spaces are limits of finite-rank operators). *Let* H *be a Hilbert space and* $K : H \to H$ *a compact operator. Then there exists a sequence of finite-rank operators* (K_n) *such that* $\|K_n - K\| \to 0$.

Proof. Let $\{e_j\}_{j=1}^{\infty}$ be an orthonormal basis for H, and let $P_n : H \to H$ denote the orthogonal projection onto $V_n := \operatorname{span}\{e_1, \dots, e_n\}$, so that

$$P_n x = \sum_{j=1}^{n} \langle x, e_j \rangle\, e_j. \tag{43}$$

Define $K_n := P_n K$. Since $R(K_n) \subseteq V_n$, each K_n has rank at most n and is therefore compact by Proposition 2.50.

We claim $\|K_n - K\| \to 0$. Suppose for contradiction that this fails. Then there exists $\varepsilon > 0$ and a subsequence (still denoted n) such that $\|K_n - K\| > \varepsilon$ for all n. By definition of the operator norm, for each n there exists $x_n \in H$ with $\|x_n\| \leq 1$ and

$$\|(K - P_n K)x_n\| = \|(I - P_n)Kx_n\| > \varepsilon. \tag{44}$$

Since K is compact and (x_n) is bounded, there exists a subsequence (x_{n_k}) such that $Kx_{n_k} \to y$ for some $y \in H$. Now estimate:

$$\|(I - P_{n_k})Kx_{n_k}\| \leq \|(I - P_{n_k})(Kx_{n_k} - y)\| + \|(I - P_{n_k})y\|. \tag{45}$$

The first term satisfies $\|(I - P_{n_k})(Kx_{n_k} - y)\| \leq \|Kx_{n_k} - y\| \to 0$ since $\|I - P_n\| \leq 1$. The second term satisfies $\|(I - P_{n_k})y\| \to 0$ since partial sums of the Fourier expansion converge: $P_n y \to y$ for every $y \in H$. This gives $\|(I - P_{n_k})Kx_{n_k}\| \to 0$, contradicting the assumption that it stays above ε.

□

Combined with Theorem 2.47, this tells us that on Hilbert spaces, **compact operators are precisely the norm limits of finite-rank operators**. They are the closest thing to matrices in infinite dimensions.

2.52 REMARK. The result above extends beyond Hilbert spaces. A Banach space Y is said to have the **approximation property** if for every compact set $C \subset Y$ and every $\varepsilon > 0$, there exists a finite-rank operator $F : Y \to Y$ with $\|Fx - x\| < \varepsilon$ for all $x \in C$. When Y has the approximation property, every compact operator $K : X \to Y$ is the norm limit of finite-rank operators. All Hilbert spaces, L^p spaces, and $C(K)$ spaces have the approximation property. However, Enflo (1973) constructed a Banach space that fails the approximation property, so the result does not hold in full generality.

2.53 REMARK. On a Hilbert space, the analogy between compact operators and matrices is precise:

Finite dimensions ($\mathbb{R}^n \to \mathbb{R}^m$)	Compact operators ($H \to H$)
Every operator has finite rank	Norm limit of finite-rank operators
Spectrum = eigenvalues	Spectrum = eigenvalues $\cup \{0\}$
Each eigenvalue has finite multiplicity	Each nonzero eigenvalue has finite multiplicity
SVD: $A = \sum_{i=1}^{r} \sigma_i u_i v_i^{\mathrm{T}}$	SVD: $K = \sum_{i=1}^{\infty} \sigma_i \langle \cdot, v_i \rangle u_i$
Symmetric: $A = Q\Lambda Q^{\mathrm{T}}$	Self-adjoint: $K = \sum_{i=1}^{\infty} \lambda_i \langle \cdot, \psi_i \rangle \psi_i$
Fredholm alternative holds	Fredholm alternative holds

The SVD holds for every compact operator between Hilbert spaces. The eigenvalue decomposition into an orthonormal eigenbasis requires the additional assumption that the operator is **normal** ($K^*K = KK^*$), which includes self-adjoint operators as a special case (just as $A = Q\Lambda Q^{\mathrm{T}}$ requires symmetry in finite dimensions). Without normality, a compact operator's eigenvectors need not form a basis.

2.8.4 The Canonical Example: Integral Operators

Our main source of compact operators is integral operators with square-integrable kernels.

2.54 THEOREM (Hilbert-Schmidt integral operators are compact). *Let $\Omega \subset \mathbb{R}^d$ be bounded and let $k \in \mathrm{L}^2(\Omega \times \Omega)$. Then the integral operator*

$$\mathrm{K}f(x) := \int_\Omega k(x,y) f(y)\, dy \tag{46}$$

defines a compact operator $\mathrm{K} : \mathrm{L}^2(\Omega) \to \mathrm{L}^2(\Omega)$.

Proof. Choose an orthonormal basis $\{\varphi_j\}$ of $\mathrm{L}^2(\Omega)$. Then $\{\varphi_i(x)\varphi_j(y)\}$ is an ONB of $\mathrm{L}^2(\Omega \times \Omega)$, so

$$k(x,y) = \sum_{i,j=1}^{\infty} k_{ij}\, \varphi_i(x)\varphi_j(y), \qquad \|k\|^2_{\mathrm{L}^2(\Omega\times\Omega)} = \sum_{i,j=1}^{\infty} |k_{ij}|^2. \tag{47}$$

Define truncations $k_n(x,y) := \sum_{i,j=1}^{n} k_{ij}\, \varphi_i(x)\varphi_j(y)$ and corresponding operators K_n. For $f = \sum_{l=1}^{\infty} c_l \varphi_l$:

$$\mathrm{K}_n f = \sum_{i,j=1}^{n} k_{ij} c_j\, \varphi_i(x). \tag{48}$$

Each K_n has finite-dimensional range (rank $\le n$), hence is compact by Proposition 2.50. The approximation error satisfies

$$\|\mathrm{K} - \mathrm{K}_n\|^2 \le \int_\Omega \int_\Omega |k(x,y) - k_n(x,y)|^2\, dx\, dy = \sum_{i,j=n+1}^{\infty} |k_{ij}|^2 \to 0. \tag{49}$$

By Theorem 2.47, K is compact.

□

2.55 REMARK. The proof above is really a **finite-rank approximation** argument. The kernel $k(x,y)$ is expanded in an ONB for $\mathrm{L}^2(\Omega \times \Omega)$, and truncating to n terms gives a rank-n operator that converges to K in the operator norm. This is analogous to truncating the SVD of a matrix to r singular values to obtain a best rank-r approximation (the Eckart-Young theorem). However, the expansion here uses an *arbitrary* ONB, not the singular functions of K. The true SVD of the integral operator would use the eigenbases of $\mathrm{K}^*\mathrm{K}$ and $\mathrm{K}\mathrm{K}^*$.

2.56 EXAMPLE (The solution operator for Poisson's equation is compact). Consider the Poisson equation on $[0, \mathrm{L}]$ with Neumann boundary conditions:

$$-u'' = f, \quad u'(0) = u'(\mathrm{L}) = 0, \quad \int_0^{\mathrm{L}} u(x)\, dx = \mathrm{M}. \tag{50}$$

Integrating twice and applying Fubini gives the solution

$$u(x) = \int_0^{\mathrm{L}} g(x,s) f(s)\, ds + c \tag{51}$$

where g is the Green's function (an L^2 kernel). The solution operator $\mathrm{A}^{-1} : f \mapsto u$ is an integral operator with an L^2 kernel, hence compact by Theorem 2.54.

This is a fundamental pattern: **differential operators are unbounded, but their inverses (solution operators) are often compact.** The compactness of A^{-1} is what gives the Laplacian a discrete spectrum of eigenvalues.

2.57 EXAMPLE (Non-example: unbounded operators are not compact). The second derivative operator $\mathrm{A} = -\frac{d^2}{dx^2} : \mathrm{D}(\mathrm{A}) \to \mathrm{L}^2(0,1)$ is unbounded, and therefore not compact. Compact operators are always bounded (Proposition 2.45).

2.8.5 Spectral Theory of Compact Operators

The spectral theory of compact operators is the payoff: it tells us that compact operators have a spectrum that looks just like the spectrum of a matrix, up to a possible accumulation point at zero.

The Fredholm Alternative

2.58 THEOREM (Spectral Fredholm Alternative). *Let* $\mathrm{A} : \mathrm{H} \to \mathrm{H}$ *be a compact linear operator on a Hilbert space* H. *Then:*

1. *The spectrum* $\sigma(\mathrm{A})$ *is a compact subset of* $\mathbb{C}$ *whose only possible accumulation point is* $\lambda = 0$.
2. *For each* $\lambda \in \mathbb{C} \setminus \{0\}$, *exactly one of the following holds:*
   - ***(Alternative 1):*** $\lambda \in \rho(\mathrm{A})$, *i.e.,* $(\mathrm{A} - \lambda \mathrm{I})^{-1}$ *exists and is bounded. The equation* $\mathrm{A}x - \lambda x = y$ *has a unique solution for every* y.
   - ***(Alternative 2):*** $\lambda \in \sigma_p(\mathrm{A})$ *is an eigenvalue of* ***finite multiplicity****. The equation* $\mathrm{A}x - \lambda x = 0$ *has a nontrivial finite-dimensional solution space.*

2.59 REMARK. Compare this with finite dimensions: for a matrix $\mathrm{A} \in \mathbb{R}^{n \times n}$ and $\lambda \neq 0$, either $\det(\mathrm{A} - \lambda \mathrm{I}) \neq 0$ (invertible) or $\det(\mathrm{A} - \lambda \mathrm{I}) = 0$ (eigenvalue with finite-dimensional eigenspace). The Fredholm alternative is exactly this dichotomy, extended to infinite dimensions via compactness. Without compactness, the spectrum can contain continuous and residual parts that have no finite-dimensional analogue.

The Hilbert-Schmidt Theorem

When the compact operator is additionally self-adjoint, we get a complete spectral decomposition—the infinite-dimensional analogue of the eigenvalue decomposition of a symmetric matrix.

2.60 THEOREM (Hilbert-Schmidt Spectral Theorem). *Let* $\mathrm{A} : \mathrm{H} \to \mathrm{H}$ *be a linear, compact, and self-adjoint operator on a Hilbert space* H. *Then:*

1. *All eigenvalues* λ_i *of* A *are* ***real****.*

2. *Eigenvalues can accumulate only at 0.*

3. *There exists an orthonormal set of eigenfunctions* $\{\psi_i\}$ *such that* A *has the* ***spectral representation***

$$A\varphi = \sum_{i=1}^{\infty} \lambda_i \langle \varphi, \psi_i \rangle \, \psi_i. \tag{52}$$

This is the **infinite-dimensional eigendecomposition** for self-adjoint operators: the operator A is completely determined by its eigenvalues and eigenfunctions, just as a symmetric matrix is determined by its eigenvalues and eigenvectors. The spectral representation $A = \sum \lambda_i \langle \cdot, \psi_i \rangle \psi_i$ is the direct analogue of the matrix diagonalization $A = Q\Lambda Q^T$. Note that this is **not** the SVD—it is the eigenvalue decomposition, which requires self-adjointness. The SVD $K = \sum \sigma_i \langle \cdot, v_i \rangle u_i$ is a separate factorization that works for all compact operators (see Remark 2.53).

2.61 REMARK (Why eigenvalues accumulate at zero). On a Hilbert space, compact operators are norm limits of finite-rank operators (Proposition 2.51). A finite-rank operator is a matrix: it has finitely many nonzero eigenvalues. As the finite-rank approximations $K_n \to K$, the eigenvalues of K_n converge to those of K, but each K_n can only contribute finitely many. The only place infinitely many eigenvalues can pile up is at the shared limit 0. Compactness forces the spectrum to be "almost finite" at every scale, and 0 is the residue of that approximation.

Contrast this with a general bounded operator that is not compact: multiplication by x on $L^2([0,1])$, defined by $(Mf)(x) = xf(x)$, has $\|M\| = 1$ but no eigenvalues at all. If $Mf = \lambda f$, then $(x - \lambda)f(x) = 0$ a.e., forcing $f = 0$ a.e. The spectrum is the purely continuous set $[0,1]$. There is no eigenfunction expansion, no discrete decomposition. This operator preserves the full infinite-dimensional structure of L^2, which is exactly what compact operators suppress.

2.62 REMARK. The spectral representation also provides a natural way to define **fractional powers** of operators. If A is a positive operator with $A\varphi = \sum \lambda_i \langle \varphi, \psi_i \rangle \psi_i$, then we define

$$A^{\alpha}\varphi := \sum_{i=1}^{\infty} \lambda_i^{\alpha} \langle \varphi, \psi_i \rangle \, \psi_i \tag{53}$$

for $\alpha \geq 0$. This is used, for example, to define $\sqrt{-\Delta}$, which appears in the study of fractional diffusion and Levy flights.

Application: Spectral Theory of the Laplacian

2.63 THEOREM (Spectral theorem for differential operators with compact inverse). *Let* $A : D(A) \to H$ *be a symmetric, linear, unbounded operator with* $R(A) = H$, *and suppose* A^{-1} *exists and is compact. Then:*

1. *There exists an infinite sequence of real eigenvalues* $\{\lambda_n\}$ *with* $\lim_{n\to\infty} |\lambda_n| = +\infty$.

2. *The eigenvectors* $\{w_j\}$ *can be chosen to form an orthonormal basis, and*

$$Au = \sum_{j=1}^{\infty} \lambda_j \langle u, w_j \rangle \, w_j. \tag{54}$$

2.64 REMARK (Connection to the Hilbert-Schmidt theorem). If A^{-1} is compact and self-adjoint, then A^{-1} satisfies the Hilbert-Schmidt theorem with eigenvalues $\mu_j \to 0$. The eigenvalues of A are $\lambda_j = 1/\mu_j \to \infty$.

This is the typical situation for Laplacian-type operators: $A = -\Delta$ with suitable boundary conditions is unbounded, but the solution operator A^{-1} (given by a Green's function) is compact (Example 2.56). Hence $-\Delta$ has a discrete spectrum of eigenvalues tending to infinity, with eigenfunctions forming an ONB. This is why Fourier series work: the eigenfunctions of the Laplacian on $[0, 2\pi]$ are precisely $\{e^{inx}\}$.

The bounded domain is essential. On $\mathbb{R}^n$ the "cage" preventing translation is lost: a fixed bump $\varphi(x - n)$ slides off to infinity, so the solution operator is no longer compact. The spectrum of $-\Delta$ on $\mathbb{R}^n$ is the continuous set $[0, \infty)$, and the clean eigenfunction expansion breaks down.

2.8.6 Looking Ahead

Compact operators are the bridge between the abstract operator theory of this chapter and several later topics in the course:

- **Weak convergence + compact operator $\Rightarrow$ strong convergence.** If $x_n \rightharpoonup x$ weakly and K is compact, then $Kx_n \to Kx$ strongly. This is a key tool in the calculus of variations for passing to the limit in nonlinear problems.
- **Rellich-Kondrachov compactness.** The Sobolev embedding $H^1(\Omega) \hookrightarrow L^2(\Omega)$ is compact for bounded Ω. This is the source of compactness in elliptic PDE theory and the reason the direct method of the calculus of variations works.
- **Fixed point theory.** Compact operators are the setting for Schauder's fixed point theorem and the Leray-Schauder degree, which extend Brouwer's fixed point theorem to infinite dimensions.

2.9 *Exercises*

2.9.1 Linear Operators

Let X be a Banach space and $A : X \mapsto Y$ a bounded linear map. Show that

- The null-space is closed.
- If there exists $K > 0$ such that $\|x\| \leq K\|Ax\|$ then the range is closed.

2.9.2 Banach-Steinhaus

Prove the above corollary. For (3) $\implies$ (2) pick a dense set $D \subset X$ and define $T(x) = \lim_n T_n(x)$ for $x \in D$. Then show that $T_n(x) \to T(x)$ for all $x \in X$.

Let $S \subset \mathcal{L}(X, Y)$ be a family of linear operators such that $\sup_{T \in S} \|T\| < \infty$. Show that this means that S is equi-continuous.

2.9.3 Banach-Steinhaus Applications

Divergence of Lagrange interpolation.

Let L_n be the Lagrange interpolation operator at Chebyshev nodes on $[-1, 1]$.

1. Show that $\|L_n\| = \Lambda_n$ where Λ_n is the Lebesgue constant.
2. Using $\Lambda_n \sim \frac{2}{\pi} \log n$, explain why Banach-Steinhaus does **not** give divergence.
3. Prove that for any $f \in C[-1, 1]$, $\|f - L_n f\|_\infty \to 0$ (use Weierstrass + stability).

Quadrature stability.

Consider the midpoint rule $M_n f = \frac{1}{n} \sum_{k=1}^{n} f\left(\frac{2k-1}{2n}\right)$ on $[0, 1]$.

1. Compute $\|M_n\|$.
2. Show that $M_n p \to \int_0^1 p$ for any polynomial.
3. Conclude that $M_n f \to \int_0^1 f$ for all $f \in C[0, 1]$.

Failure of pointwise convergence.

Let $X = C[0, 1]$ and define $T_n : X \to \mathbb{R}$ by $T_n f = n \int_0^{1/n} f(t)\, dt$.

1. Show each T_n is bounded with $\|T_n\| = 1$.
2. Find $\lim_{n \to \infty} T_n f$ for continuous f.
3. This is an example where uniform boundedness holds and pointwise limits exist. What is the limit operator?

Unbounded inverse.
Let $X = \{f \in C^1[0,1] : f(0) = 0\}$ with norm $\|f\| = \|f'\|_\infty$, and $Y = C[0,1]$ with sup norm.
Define $T : X \to Y$ by $Tf = f$ (inclusion).

1. Show T is bounded, linear, and bijective onto its range.
2. Compute $\|T^{-1}\|$ (the inverse is differentiation). Is it bounded?
3. Why doesn't this contradict Banach's bounded inverse theorem?

2.9.4 Open Mapping Theorem

Show that all linear transformations from a finite-dimensional space to a normed vector space are continuous. Conclude that all norms on a finite-dimensional space are equivalent.

Hint: You can use either the open-mapping theorem or Banach-Steinhaus to prove this by considering the identity map.

2.9.5 Open Mapping Theorem Applications

Condition number of differentiation.
Consider $D : (C^1[0,1], \|\cdot\|_{C^1}) \to (C[0,1], \|\cdot\|_\infty)$ where $\|f\|_{C^1} = \|f\|_\infty + \|f'\|_\infty$.

1. Show D is bounded with $\|D\| \leq 1$.
2. Find $\|D\|$ exactly.
3. Describe the "inverse" D^{-1} (integration). On what domain is it defined?
4. Compute $\|D^{-1}\|$ and the condition number.

Inf-sup for the divergence.
Let $\Omega = (0,1)^2$, $V = H_0^1(\Omega)^2$, $Q = L_0^2(\Omega) = \{q : \int_\Omega q = 0\}$.
Define $b(\mathbf{v}, q) = \int_\Omega q \nabla \cdot \mathbf{v}$.

1. Show that b is continuous.
2. The inf-sup condition states $\inf_q \sup_{\mathbf{v}} \frac{b(\mathbf{v},q)}{\|\mathbf{v}\|_{H^1}\|q\|_{L^2}} \geq \beta > 0$. Interpret this as an Open Mapping statement.
3. Why does this imply existence of $\mathbf{v} \in V$ with $\nabla \cdot \mathbf{v} = q$ and $\|\mathbf{v}\|_{H^1} \leq C\|q\|_{L^2}$?

Closed Graph for integral operators.
Let $K : L^2[0,1] \to L^2[0,1]$ be defined by $(Kf)(x) = \int_0^x f(t)dt$.

1. Prove K has closed graph directly (without using continuity).
2. Show K is bounded by computing $\|K\|$.

3. Is K^{-1} bounded? (Consider the domain carefully.)

Perturbation theory.
Let $A : X \to X$ be invertible with $\|A^{-1}\| = M$.

1. Show that if $\|B\| < 1/M$, then $A + B$ is invertible.
2. Derive the bound $\|(A + B)^{-1} - A^{-1}\| \leq \frac{M^2\|B\|}{1-M\|B\|}$.
3. Apply this to $A = I$ to get the Neumann series.

2.9.6 Compact Operators

Show that if $K : X \to Y$ is compact and $T : Y \to Z$ is bounded, then TK is compact. Similarly, show that if $T : Z \to X$ is bounded and $K : X \to Y$ is compact, then KT is compact.

Let H be a Hilbert space and $K : H \to H$ a compact operator. Show that if (x_n) converges weakly to x (i.e. $\langle x_n, y\rangle \to \langle x, y\rangle$ for all $y \in H$), then $Kx_n \to Kx$ in norm. That is, **compact operators map weakly convergent sequences to strongly convergent sequences**.

Hint: Suppose for contradiction that $\|Kx_n - Kx\| \not\to 0$. Pass to a subsequence where $\|Kx_{n_k} - Kx\| \geq \varepsilon$, then use compactness to extract a further subsequence where $Kx_{n_{k_j}}$ converges. What must the limit be?

Let $K : L^2(0,1) \to L^2(0,1)$ be the integral operator

$$Kf(x) = \int_0^1 \min(x, y)\, f(y)\, dy. \tag{55}$$

1. Show that K is compact by verifying that $k(x, y) = \min(x, y) \in L^2((0,1) \times (0,1))$.
2. Show that K is self-adjoint: $\langle Kf, g\rangle = \langle f, Kg\rangle$.
3. Find the eigenvalues and eigenfunctions of K. *(Hint: differentiate $Kf = \lambda f$ twice to convert to an ODE.)*

Give an example of a bounded linear operator $T : \ell^2 \to \ell^2$ that is **not** compact. Prove that your operator fails the sequential characterization of compactness.

3 DUALITY: HYPERPLANES, FUNCTIONALS, AND HAHN–BANACH

Big Idea

In a general Banach space there is no inner product, so orthogonality is lost. What survives is **hyperplane geometry**: every continuous linear functional slices the space into parallel hyperplanes, and the Hahn–Banach theorem guarantees there are enough such slices to recover norms, separate points, and separate convex sets.

3.1 *Overview*

There are two ways to study a mathematical structure: look at the objects **in** the space, or look at the functions **on** the space. In functional analysis, the second viewpoint is remarkably powerful, but it rests on a geometric foundation that is worth making explicit.

In Hilbert spaces, we have orthogonality: the inner product lets us decompose vectors into components, project onto subspaces, and find nearest points. But most Banach spaces have no inner product, so **orthogonality is lost**. What *does* survive is **hyperplane geometry**. Every continuous linear functional $f : X \to \mathbb{R}$ slices the space into parallel hyperplanes (the level sets $f^{-1}(c)$), and its kernel is a codimension-1 subspace. This structure is purely algebraic and topological; it requires no inner product. The kernel determines the functional up to scaling, level sets foliate the space, and closed kernels correspond to continuous functionals. All of this works identically in $\mathbb{R}^3$ and in $L^p(\Omega)$.

Duality is the systematic study of this surviving geometry. The dual space $X^* = \mathcal{L}(X, \mathbb{R})$ collects all bounded linear functionals on X, encoding all possible ways to slice the space with hyperplanes. But is X^* rich enough? Could all functionals miss some structure in X? The **Hahn–Banach Theorem** answers this: functionals defined on subspaces can always be extended to the full space, guaranteeing that X^* is large enough to separate points, recover norms, and separate convex sets. From this we get:

- **Separation theorems** for convex sets, putting the hyperplane geometry to work.
- The **Riesz Representation Theorem** identifying dual spaces (e.g., $H^* \cong H$ for Hilbert spaces), where orthogonality is recovered as a special case.
- **Weak and weak-$*$ convergence**, providing compactness where norm convergence fails by testing against functionals.
- **Banach–Alaoglu**, giving weak-$*$ compactness of the unit ball in X^*.

3.2 *What You Will Learn*

- How functionals define hyperplanes, and why kernels have codimension 1.
- The geometry of foliations and what distinguishes continuous from discontinuous functionals.
- The Hahn–Banach Theorem and its geometric consequences: separation of convex sets and the sup formula for the norm.
- Dual spaces in practice: Riesz representation for Hilbert spaces, reflexive and non-reflexive spaces.
- Weak and weak-$*$ convergence, and the Banach–Alaoglu compactness theorem.

3.3 *The Dual Space*

There are two methods to studying mathematical structures.

1. By looking at the objects **in** the space.
2. By looking at functions **on** the space.

Let's see how the second viewpoint looks in our Banach space setting.

3.3.1 Introduction

In the Banach space setting the natural mappings to consider are the linear continuous functions between two Banach spaces X and Y. Now consider $\mathcal{L}(X, \mathbb{R})$ the space of bounded linear continuous functionals on X. This is a Banach space, and we will call it the **dual** of X denoted by X^*.

The intuitive view is that **functionals** are measurements of an object in the space X.

3.1 EXAMPLE. We encounter many functionals in daily life.

1. A scale measures the weight of objects — it assigns a real number to each object.
2. A ruler measures length.
3. CT-scanners measure line integrals of an unknown tissue density.
4. MRI-scanners measure proton densities along lines.

3.3.2 Isomorphisms and Isometries

First let's explore how we show that two Banach spaces have the same structure. In normed vector spaces the **natural map** is $T : X \to Y$. X and Y are equivalent if there exists an invertible continuous linear map $T : X \to Y$ i.e.

$$c\|x\|_X \leq \|Tx\|_Y \leq C\|x\|_X \tag{56}$$

If such a map T exists we say that X and Y are **isomorphic**. If $c = C = 1$ then T is an **isometry**.

3.2 REMARK. The upper bound $\|Tx\|_Y \leq C\|x\|_X$ is simply the statement that T is bounded (continuous). The lower bound $c\|x\|_X \leq \|Tx\|_Y$ is equivalent to T^{-1} being bounded. By Corollary 2.29 (Banach's Bounded Inverse Theorem), if $T : X \to Y$ is a **bijective** bounded linear map between Banach spaces, then T^{-1} is automatically bounded, so the lower bound comes **for free**. In other words, to show two Banach spaces are isomorphic it suffices to exhibit a bijective continuous linear map; one does not need to verify the lower bound separately.

3.3.3 The Finite-Dimensional Case

SETUP Let $V = \mathbb{R}^n$ with the Euclidean norm $\|x\| = \sqrt{x_1^2 + \cdots + x_n^2}$.

Question: What are the linear functionals $f : \mathbb{R}^n \to \mathbb{R}$?

Every such f is of the form $f(y) = a \cdot y$ for some fixed vector $a \in \mathbb{R}^n$.

Proof. If $e_1, \dots, e_n$ is the standard basis, set $a_i = f(e_i)$. Then by linearity $f(y) = \sum a_i y_i = a \cdot y$. □

THE DUAL ELEMENT OF A VECTOR Given $x \in \mathbb{R}^n$, define the associated linear functional:

$$f_x(y) = x \cdot y \tag{57}$$

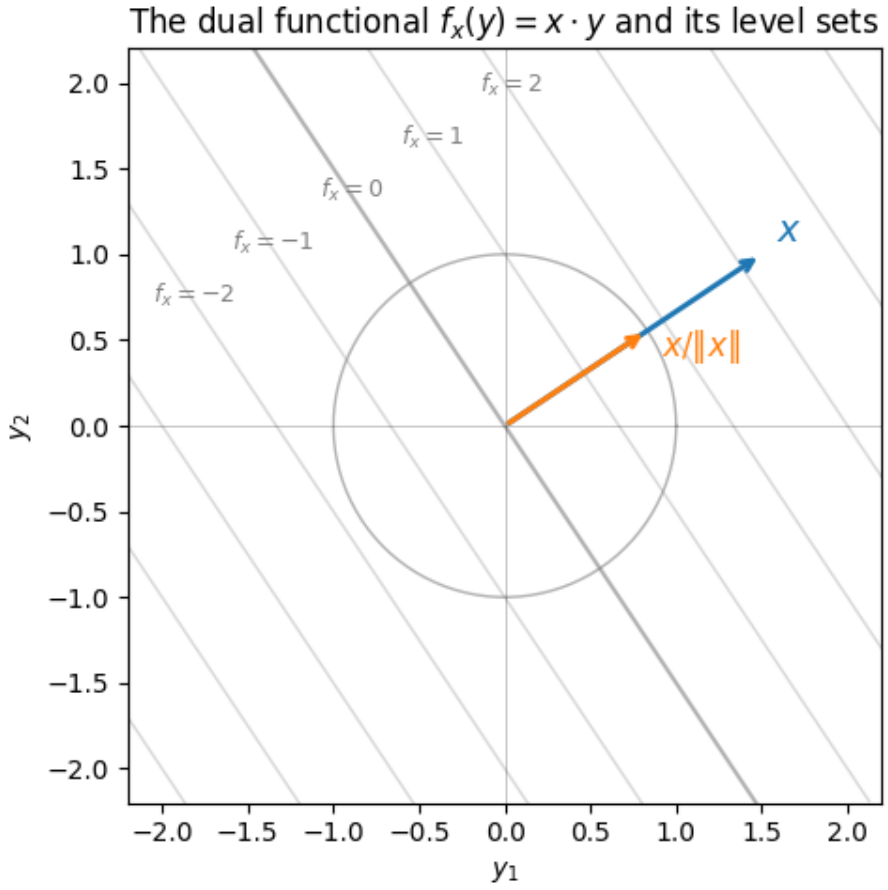
In matrix language: if vectors are columns, then f_x acts as the row vector x^{T}, so $f_x(y) = x^{\mathrm{T}} y$.
Key observations:

- $f_x(x) = x \cdot x = \|x\|^2$ (evaluating the functional at the vector itself)
- $|f_x(y)| \le \|x\|\|y\|$ by Cauchy–Schwarz
- The operator norm is $\|f_x\|_{\mathrm{op}} = \sup_{\|y\|=1} |f_x(y)| = \|x\|$

The upper bound $\|f_x\|_{\mathrm{op}} \le \|x\|$ is Cauchy–Schwarz. For the lower bound, plug in $y = x/\|x\|$:

$$f_x\left(\frac{x}{\|x\|}\right) = \frac{x \cdot x}{\|x\|} = \|x\| \tag{58}$$

So the map $x \mapsto f_x$ is an **isometric isomorphism** $\mathbb{R}^n \xrightarrow{\sim} (\mathbb{R}^n)^*$.



PROPERTIES OF THE FINITE-DIMENSIONAL DUAL

3.3 PROPOSITION. *In $\mathbb{R}^n$ the following hold:*

1. ***Every linear functional is continuous.*** *A linear map $f : \mathbb{R}^n \to \mathbb{R}$ is automatically bounded. There is no gap between the algebraic dual and the topological dual:* $V' = X^*$.

2. ***Extension is immediate.*** *If f is defined on a subspace $M \subset \mathbb{R}^n$, extend it to all of $\mathbb{R}^n$ by choosing a complementary basis. No axiom of choice required.*

3. ***The sup formula holds.*** *For any $x \in \mathbb{R}^n$:*

$$\|x\| = \max_{\|f\|_{op} \leq 1} |f(x)| \tag{59}$$

   Take $f = f_{x/\|x\|}$ to achieve the max. The norm is fully recovered from the dual.

4. X^* ***separates points.*** *If $f(x) = f(y)$ for all $f \in (\mathbb{R}^n)^*$, then $x = y$.*

Proof. We prove property 4. By linearity, $f(x - y) = 0$ for all f. Take $f = f_{x-y}$, the functional $f_{x-y}(v) = (x - y) \cdot v$. Then $0 = f_{x-y}(x - y) = (x - y) \cdot (x - y) = \|x - y\|^2$, so $x = y$.

□

The key observation is that given any vector we want to detect, the inner product lets us immediately build the right functional. In infinite dimensions we no longer have this luxury, and constructing functionals with prescribed behavior requires the **Hahn-Banach theorem**.

3.3.4 What Goes Wrong in Infinite Dimensions?

ALGEBRAIC VS. TOPOLOGICAL DUAL Let $X = (V, \|\cdot\|)$ be a normed space, with V a vector space over $\mathbb{R}$.

3.4 DEFINITION (Algebraic dual). V' is the set of *all* linear functionals $f : V \to \mathbb{R}$ — no continuity requirement. This is a purely algebraic object.

3.5 DEFINITION (Topological dual). $X^* \subset V'$ is the set of all *bounded* (equivalently, continuous) linear functionals $f : V \to \mathbb{R}$, equipped with the operator norm:

$$\|f\|_{X^*} = \|f\|_{\text{op}} = \sup_{\|x\| \leq 1} |f(x)| \tag{60}$$

X^* is always a Banach space (even if X itself is not complete).

In infinite dimensions, $V' \supsetneq X^*$ — there exist discontinuous linear functionals. These are constructed via a Hamel basis (using Zorn's lemma / axiom of choice), and they are thoroughly pathological: unbounded on every neighborhood of 0.

COULD THE TOPOLOGICAL DUAL BE TRIVIAL? Could X^* be trivially small? We know $V' \neq \{0\}$ (Hamel basis argument), but these algebraic functionals are useless for analysis. Could it happen that $X^* = \{0\}$ — that the *only* continuous linear functional is the zero functional?

If $X^* = \{0\}$, we cannot distinguish elements via functionals and there is no useful duality theory.

This happens! The spaces $L^p(\Omega)$ for $0 < p < 1$ have trivial dual: $(L^p)^* = \{0\}$. These are complete metric spaces (F-spaces) but not Banach spaces. The idea (due to Day [2]) is simple: one shows that L^p for $0 < p < 1$ has **no proper convex open subsets**. But any nonzero continuous linear functional f would make $\{x : |f(x)| < 1\}$ exactly such a set, so no such f can exist.

This motivates why **Banach spaces** (with their convex unit balls) are the right setting for duality.

THE DUAL AS A MEASUREMENT SYSTEM Here is the motivating picture for the rest of this chapter.

Suppose you have an element $x \in X$ and you want to know its norm $\|x\|$. But you're not allowed to compute $\|x\|$ directly — instead, you can only **probe** x using continuous linear functionals. Each $f \in X^*$ with $\|f\| \leq 1$ gives you a **reading** $|f(x)|$. Think of these as measurement instruments: each one returns a number, and you take the supremum over all available instruments:

$$\text{best reading} = \sup\{|f(x)| : f \in B_{X^*}\} \tag{61}$$

One inequality is trivial: since $|f(x)| \leq \|f\|\|x\| \leq \|x\|$ for any f in the unit ball of X^*, no instrument ever reads **too high**:

$$\sup\{|f(x)| : f \in B_{X^*}\} \leq \|x\| \tag{62}$$

But the real worry is the other direction: **could all your instruments read too low?** Could there be some "hidden size" in x that no continuous functional detects? If so, there is a gap between what the dual can see and what is actually there.

In finite dimensions this can't happen: the functional $f_{x/\|x\|}$ gives a reading of exactly $\|x\|$. But in infinite dimensions, we don't know *a priori* that such an optimal instrument exists.

If $X^* = \{0\}$ (as for L^p with $0 < p < 1$), the situation is as bad as possible: the supremum is 0 for every nonzero x. The dual sees nothing.

The question Hahn–Banach answers: Is X^* a **complete measurement system**, one with no information loss? That is, do we have

$$\|x\| = \max\{|f(x)| : f \in B_{X^*}\}? \tag{63}$$

The answer is **yes** for all Banach spaces: the sup is attained, and the dual recovers the full norm.

3.4 *The Geometry of Linear Functionals*

Before proving Hahn–Banach, it helps to understand what a linear functional *looks like* geometrically. This perspective clarifies what duality is really doing.

3.4.1 Kernels and level sets

For a nonzero linear functional $f \in X'$ (where X' denotes the algebraic dual), define two canonical sets:

- $K(f) = \ker(f) = f^{-1}(0) = \{x \in X : f(x) = 0\}$ — the **kernel**
- $I(f) = f^{-1}(1) = \{x \in X : f(x) = 1\}$ — the **unit level set**

Both are level sets of f: $K(f)$ is where f vanishes, and $I(f)$ is where f equals 1.

Note that $0 \in K(f)$ always (since $f(0) = 0$ by linearity), while $0 \notin I(f)$ (since $f(0) = 0 \neq 1$). So the kernel passes through the origin; the unit level set does not.

3.4.2 The fundamental decomposition

3.6 PROPOSITION (Codimension-1 decomposition). *Let $f \in X'$ be a nonzero linear functional, and let $x_0 \in X$ with $f(x_0) \neq 0$. Then:*

1. *$K(f)$ is a subspace of **codimension 1** (see Corollary 2.43) — it is "missing exactly one dimension."*
2. *Every vector $x \in X$ has a **unique** representation*

$$x = y + \lambda x_0, \qquad y \in K(f), \quad \lambda \in \mathbb{R} \tag{64}$$

3. *The scalar λ is determined: $\lambda = \dfrac{f(x)}{f(x_0)}$.*

Proof. Given any $x \in X$, set $\lambda = f(x)/f(x_0)$ and $y = x - \lambda x_0$. Then $f(y) = f(x) - \lambda f(x_0) = f(x) - \frac{f(x)}{f(x_0)} f(x_0) = 0$, so $y \in K(f)$ and $x = y + \lambda x_0$ by construction.

Uniqueness: If $x = y_1 + \lambda_1 x_0 = y_2 + \lambda_2 x_0$ with $y_1, y_2 \in K(f)$, then $(\lambda_1 - \lambda_2)x_0 = y_2 - y_1 \in K(f)$, so $(\lambda_1 - \lambda_2)f(x_0) = 0$. Since $f(x_0) \neq 0$, we get $\lambda_1 = \lambda_2$, hence $y_1 = y_2$.

Codimension 1: The decomposition gives $X = K(f) \oplus \operatorname{span}(x_0)$. Since $\operatorname{span}(x_0)$ is one-dimensional, $K(f)$ has codimension 1 — it is a maximal proper subspace, i.e., a **hyperplane through the origin**.

□

3.4.3 The foliation picture

The decomposition tells us exactly how f organizes the space geometrically.

Every $x \in X$ is uniquely $x = y + \lambda x_0$ with $y \in K(f)$, and $f(x) = \lambda f(x_0)$. So the level sets of f are:

$$f^{-1}(c) = \left\{ y + \frac{c}{f(x_0)}\, x_0 : y \in K(f) \right\} = K(f) + \frac{c}{f(x_0)}\, x_0 \tag{65}$$

These are parallel copies of $K(f)$, shifted along the x_0-direction. They **foliate** X — every point lies on exactly one level set — and f assigns a "height" c to each slice.

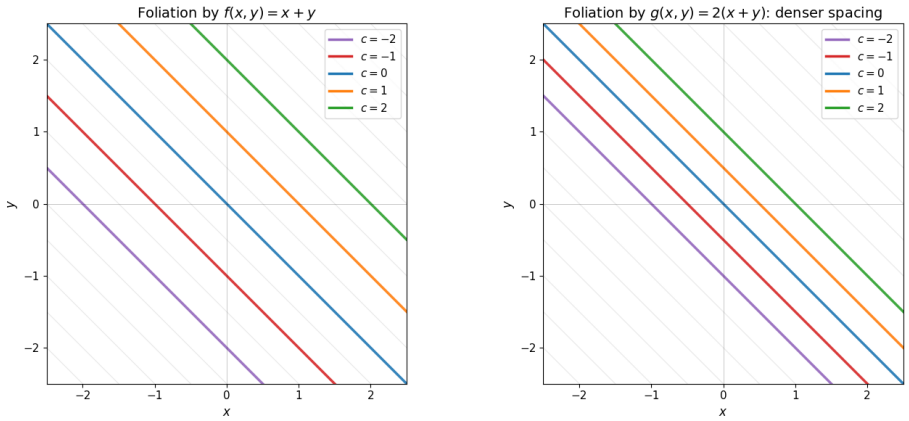
If we normalize so that $f(x_0) = 1$:

- $\mathrm{K}(f) = f^{-1}(0)$ is the slice through the origin.
- $\mathrm{I}(f) = f^{-1}(1) = \mathrm{K}(f) + x_0$ is the slice one step along x_0.
- $f^{-1}(n) = \mathrm{K}(f) + n\,x_0$ for any integer n — evenly spaced slices.

3.7 EXAMPLE (Foliation in $\mathbb{R}^3$).

Take $f(x, y, z) = 2x + 3y - z$.

- $\mathrm{K}(f) = \{2x + 3y - z = 0\}$ — a plane through the origin.
- Pick $x_0 = (1, 0, 2)$; then $f(x_0) = 2$.
- The level sets $f^{-1}(c) = \{2x + 3y - z = c\}$ are parallel planes.
- Moving along x_0 increases f by 2 per unit step.



Level sets of $f(x, y) = x + y$ (left) and $g(x, y) = 2(x + y)$ (right). Both functionals share the same kernel $\mathrm{K}(f) = \mathrm{K}(g) = \{x + y = 0\}$, but g has twice the spacing density: the $c = 1$ level set of g sits where the $c = 1/2$ level set of f would be.

WHAT THE KERNEL AND SCALE EACH DETERMINE A nonzero functional is determined by two independent pieces of geometric data:

1. The kernel determines the orientation of the slices. Two functionals with the same kernel produce exactly the same family of parallel hyperplanes. The kernel tells you *which directions are horizontal* (directions lying in $\mathrm{K}(f)$) and *which direction is transverse*.

2. The scale determines the spacing between slices. If $g = 2f$, then $\ker(g) = \ker(f)$ — same orientation, same family of parallel hyperplanes. But $g^{-1}(1) = f^{-1}(1/2)$, which sits halfway between $f^{-1}(0)$ and $f^{-1}(1)$. Rescaling f by λ doesn't rotate the slices — it compresses or stretches the spacing between them.

MULTIPLE FUNCTIONALS: ADDING VS. INTERSECTING A common source of confusion is the distinction between *adding* two functionals (an operation in X^*) and *intersecting* their kernels (an operation on subspaces of X). These produce very different geometric objects.

3.8 EXAMPLE (Adding vs. intersecting functionals in $\mathbb{R}^3$).

Work in $\mathbb{R}^3$ with the three coordinate functionals:

$$f_1(x, y, z) = x, \qquad f_2(x, y, z) = y. \tag{66}$$

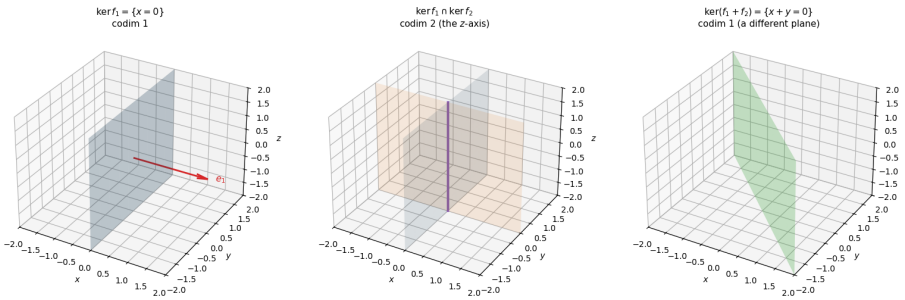
One functional. The kernel $\ker f_1 = \{x = 0\}$ is the yz-plane, a codimension-1 subspace. The decomposition gives $\mathbb{R}^3 = \ker f_1 \oplus \operatorname{span}\{e_1\}$: every vector is uniquely a point in the yz-plane plus a multiple of e_1.

Intersecting two kernels. The intersection $\ker f_1 \cap \ker f_2 = \{x = 0\} \cap \{y = 0\} = \{(0, 0, z) : z \in \mathbb{R}\}$ is the z-axis, a codimension-**2** subspace. Each independent functional removes one degree of freedom. The decomposition extends: $\mathbb{R}^3 = (\ker f_1 \cap \ker f_2) \oplus \operatorname{span}\{e_1, e_2\}$. Imposing n independent constraints carves out a codimension-n subspace.

Adding two functionals. The sum $f_1 + f_2$ is the functional $(f_1 + f_2)(x, y, z) = x + y$. This is a *single* functional, so its kernel $\ker(f_1 + f_2) = \{x + y = 0\}$ is a plane: still codimension **1**. It is a *different* plane from either $\ker f_1$ or $\ker f_2$, tilted at 45 between them. Adding functionals combines two measurements into one new measurement; it does not impose two constraints simultaneously.

Dependent functionals. If $f_2 = 2f_1$, then $\ker f_1 \cap \ker f_2 = \ker f_1$ — still the yz-plane, still codimension 1. The second functional measures the same direction as the first (just with a different scale), so it adds no new geometric information.

Cancellation. If $f_2 = -f_1$, then $f_1 + f_2 = 0$, the zero functional. Its kernel is all of $\mathbb{R}^3$: the two measurements cancel perfectly, and the constraint vanishes.



Left: the kernel of a single functional $f_1(x, y, z) = x$ is the yz-plane (codimension 1), with the transverse direction e_1 shown in red. Center: intersecting the kernels of two independent functionals f_1 and f_2 gives the z-axis (codimension 2), shown as the thick purple line where the two faint planes meet. Right: the sum $f_1 + f_2$ is a single functional with kernel $\{x + y = 0\}$, a different plane (codimension 1, green). Adding functionals produces one new measurement; intersecting kernels imposes multiple constraints.

3.9 REMARK (From intersecting kernels to weak neighborhoods). The intersection $\ker f_1 \cap \cdots \cap \ker f_n$ for n linearly independent functionals is a subspace of codimension n. If we thicken each kernel into a slab $\{y : |f_i(y)| < \varepsilon\}$ and intersect, we get a **tube**: bounded cross-section in n directions, unbounded in the remaining (codimension-n) directions. These tubes are exactly the basic neighborhoods of the weak topology (see Definition 3.41). Each additional independent functional constrains one more direction but still leaves infinitely many unconstrained in infinite dimensions.

3.4.4 Functionals are determined by their kernels

The decomposition immediately yields a powerful uniqueness result.

3.10 PROPOSITION. *If $f, g : X \to \mathbb{R}$ are nonzero linear functionals with $\ker(f) = \ker(g)$, then $f = \lambda g$ for some scalar $\lambda \neq 0$.*

Proof. Pick any $x_0 \notin \ker(f) = \ker(g)$. By the decomposition, every $x \in X$ can be written as $x = y + \frac{f(x)}{f(x_0)}\, x_0$

with $y \in \ker(f) = \ker(g)$. Applying g:

$$g(x) = g(y) + \frac{f(x)}{f(x_0)}\, g(x_0) = 0 + \frac{g(x_0)}{f(x_0)}\, f(x) \tag{67}$$

since $g(y) = 0$ (because $y \in \ker(g)$). Setting $\lambda = g(x_0)/f(x_0)$ gives $g = \lambda f$.

□

Consequence for the dual space: Classifying linear functionals up to scaling is the same as classifying codimension-1 subspaces. The projective algebraic dual $X'/\sim$ (where $f \sim \lambda f$ for $\lambda \neq 0$) is in bijection with the set of hyperplanes through the origin.

3.4.5 Characterizing continuity via hyperplanes

Here is where the algebraic vs. topological dual distinction becomes geometric. The key fact is surprisingly sharp: **a hyperplane is either closed (and nowhere dense) or dense in X — there is nothing in between.**

3.11 LEMMA. *A proper closed subspace of a normed space has empty interior (hence is nowhere dense).*

Proof. Let $M \subsetneq X$ be a proper closed subspace. Suppose for contradiction that M contains an open ball $B(x, r)$ for some $x \in M$, $r > 0$. Since M is a subspace and $x \in M$, for any $v \in B(0, r)$ we have $v = (x + v) - x$, and $x + v \in B(x, r) \subset M$, so $v \in M$. Thus $B(0, r) \subset M$. But then for any $y \in X$ with $y \neq 0$, the vector $\frac{r}{2\|y\|}\, y$ lies in $B(0, r) \subset M$, so $y \in M$ (since M is a subspace). This gives $M = X$, contradicting properness.

□

Why "empty interior" sounds wrong at first

Interior here means in the **full topology of** X, not the relative topology on M. The x-axis in $\mathbb{R}^2$ looks solid from the inside, but any open ball in $\mathbb{R}^2$ escapes into the y-direction — the subspace is "infinitely thin" when viewed from the ambient space.

So: **a continuous linear functional has a closed kernel, and this kernel is a "thin" hyperplane** — a codimension-1 subspace with empty interior in X.

3.12 LEMMA. *A codimension-1 subspace of a normed space is either closed or dense.*

Proof. Let $K(f)$ be a codimension-1 subspace (the kernel of some nonzero f). Its closure $\overline{K(f)}$ is a closed subspace containing $K(f)$. Since $K(f)$ has codimension 1, there are only two possibilities:

- $\overline{K(f)} = K(f)$ — in which case $K(f)$ is closed.
- $\overline{K(f)} \supsetneq K(f)$ — but then $\overline{K(f)}$ is a closed subspace strictly containing a codimension-1 subspace, so it must be all of X. Hence $K(f)$ is dense in X.

There is no intermediate option.

□

3.13 PROPOSITION (Continuity via the kernel). *A nonzero linear functional $f : X \to \mathbb{R}$ is continuous if and only if* $\ker(f)$ *is closed.*

Proof. **($\Rightarrow$):** If f is continuous, then $\ker(f) = f^{-1}(\{0\})$ is the preimage of a closed set, hence closed.

($\Leftarrow$): Suppose $\ker(f)$ is closed and $f \neq 0$. Pick x_0 with $f(x_0) \neq 0$; by the decomposition, $X = \ker(f) \oplus \operatorname{span}(x_0)$. Since $\ker(f)$ is closed, x_0 has positive distance from $\ker(f)$: set $\delta = \operatorname{dist}(x_0, \ker(f)) > 0$. For any $x = y + \lambda x_0$ with $y \in \ker(f)$:

$$\|x\| = \|y + \lambda x_0\| \geq |\lambda| \cdot \operatorname{dist}(x_0, \ker(f)) = |\lambda|\delta \tag{68}$$

(since $-y/\lambda \in \ker(f)$ when $\lambda \neq 0$, so $\|x_0 + y/\lambda\| \geq \delta$, and multiplying by $|\lambda|$ gives $\|x\| \geq |\lambda|\delta$). Therefore:

$$|f(x)| = |\lambda| \cdot |f(x_0)| \leq \frac{|f(x_0)|}{\delta} \|x\| \tag{69}$$

So f is bounded with $\|f\| \leq |f(x_0)|/\delta$.

□

3.14 REMARK (Connection to the quotient norm). The quantity $\delta = \operatorname{dist}(x_0, \ker(f))$ appearing in the proof above is the quotient norm $\|[x_0]\|_{X/\ker(f)}$ (see Definition 2.38). The closedness of $\ker(f)$ is essential: it guarantees that the quotient norm is a genuine norm, so that $\|[x_0]\| > 0$ for $x_0 \notin \ker(f)$.

In fact, the entire proof is the quotient perspective in disguise. The functional f factors through the quotient as

$$X \xrightarrow{\pi} X/\ker(f) \xrightarrow{\widetilde{f}} \mathbb{R} \tag{70}$$

where π is the canonical projection and $\widetilde{f}$ is the isomorphism from the Theorem 2.42. Since $X/\ker(f)$ is one-dimensional, $\widetilde{f}$ is automatically bounded (all norms on $\mathbb{R}$ are equivalent), and π is bounded with $\|\pi\| \leq 1$. The composite $f = \widetilde{f} \circ \pi$ is therefore bounded.

Combining everything:

- $f \in X^*$ (continuous) $\iff$ $K(f)$ is **closed** $\implies$ $K(f)$ is **nowhere dense**: a thin, clean hyperplane with empty interior. The space is sliced into well-separated parallel copies of $K(f)$, and the height function varies continuously.
- $f \in V' \setminus X^*$ (discontinuous) $\iff$ $K(f)$ is **not closed** $\implies$ $K(f)$ is **dense**: the hyperplane comes arbitrarily close to every point in X.

There is nothing in between: every hyperplane is either cleanly closed or pathologically dense.

3.15 REMARK (Dense hyperplanes are pathological). If $K(f)$ is dense, then *every* level set $f^{-1}(c)$ is also dense. Points at height 0, height 1, and height 1,000,000 are all interleaved at arbitrarily fine scales throughout the space. With a continuous functional, you'd see clean parallel bands of color, smoothly transitioning. With a discontinuous functional, every color appears in every tiny ball.

Concretely: if f is discontinuous, then for any $\varepsilon > 0$ and any $M > 0$, there exists x with $\|x\| < \varepsilon$ but $|f(x)| > M$. The slicing exists algebraically, but the height function is wildly discontinuous. This is why continuous functionals — the ones with closed, nowhere-dense kernels — are the only ones useful for analysis.

3.4.6 A Worked Example: Hahn–Banach Extension in $\mathbb{R}^2$

The Hahn–Banach theorem, stated in the hyperplane language, says: **if you know the heights of points along a subspace, you can extend the height function to the entire space by choosing a closed hyperplane.**

We make this explicit with a worked example. We use $\mathbb{R}^2$ with the ℓ^∞ norm $\|(x, y)\|_\infty = \max(|x|, |y|)$, so the unit ball is a **square**.

> **Why the ℓ^∞ norm and not Euclidean?**
>
> With the Euclidean norm (round unit ball), the Hahn–Banach extension turns out to be **unique** — the roundness of the ball forces a single choice. This is a general fact about Hilbert spaces: the Riesz representation theorem pins down the unique extension.
>
> With the ℓ^∞ norm (square unit ball), the extension is **non-unique** — there is a whole interval of valid choices, each giving a different kernel. This is the generic situation in Banach spaces.

Step 1: Define f on a subspace.

Let $\mathrm{M} = \mathrm{span}((1, 1))$. Define f on M by $f(t(1, 1)) = t$. Check the norm: $|f(t(1, 1))| = |t|$ and $\|t(1, 1)\|_\infty = |t| \cdot \|(1, 1)\|_\infty = |t|$. So $\|f\| = 1$ on M.

We know the heights along the diagonal line M: the origin has height 0, the point $(1/2, 1/2)$ has height $1/2$, the point $(1, 1)$ has height 1.

Step 2: Decompose.

Pick $z = (1, -1) \notin \mathrm{M}$ as the new direction. Every $(x, y) \in \mathbb{R}^2$ decomposes as:

$$(x, y) = \underbrace{\frac{x+y}{2}(1, 1)}_{\text{component in M}} + \underbrace{\frac{x-y}{2}}_{\text{coefficient } \alpha} (1, -1) \tag{71}$$

Step 3: The extension is determined by one number.

Set $c = \mathrm{F}(z) = \mathrm{F}(1, -1)$ — the height we assign to the new direction. By linearity:

$$\mathrm{F}(x, y) = \frac{x+y}{2} + \frac{x-y}{2} \cdot c = \frac{1+c}{2}\, x + \frac{1-c}{2}\, y \tag{72}$$

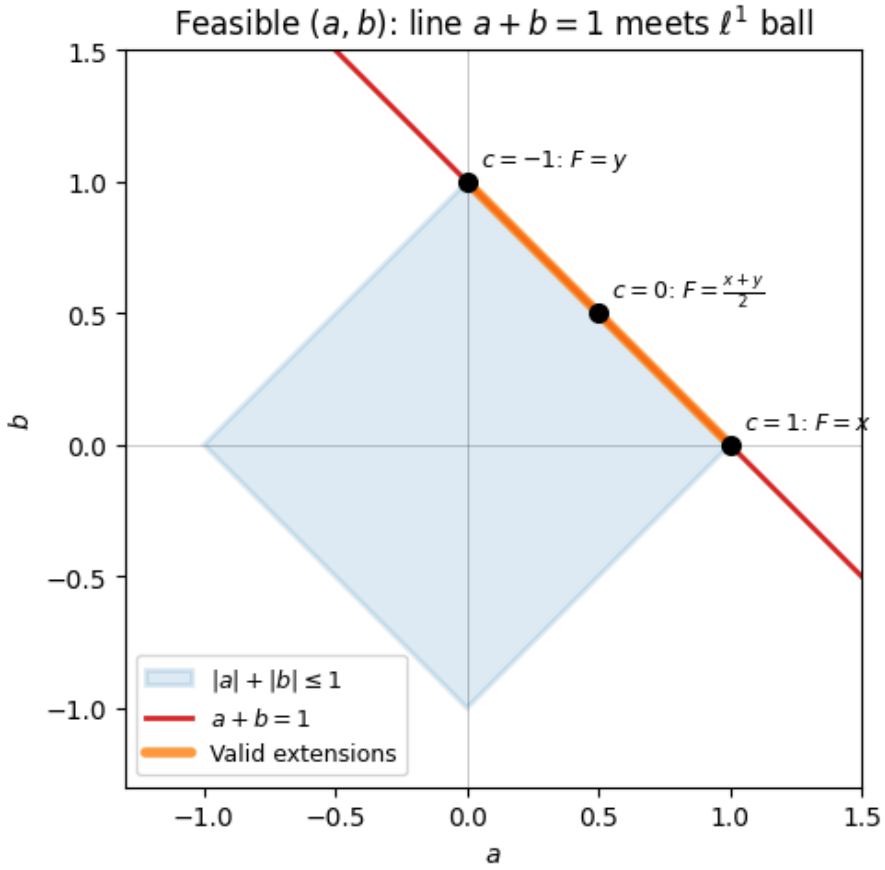
Note $a + b = 1$ where $a = (1 + c)/2$ and $b = (1 - c)/2$, ensuring $\mathrm{F}(1, 1) = 1 = f(1, 1)$.

Step 4: Find which values of c are valid.

We need $\|\mathrm{F}\|_{\mathrm{op}} \leq 1$ in the ℓ^∞ norm. For a linear functional $\mathrm{F}(x, y) = ax + by$ on $(\mathbb{R}^2, \|\cdot\|_\infty)$, the operator norm is $\|\mathrm{F}\| = |a| + |b|$. The constraint $|a| + |b| \leq 1$ combined with $a + b = 1$ forces $a, b \geq 0$, giving:

$$\boxed{c \in [-1, 1]} \tag{73}$$

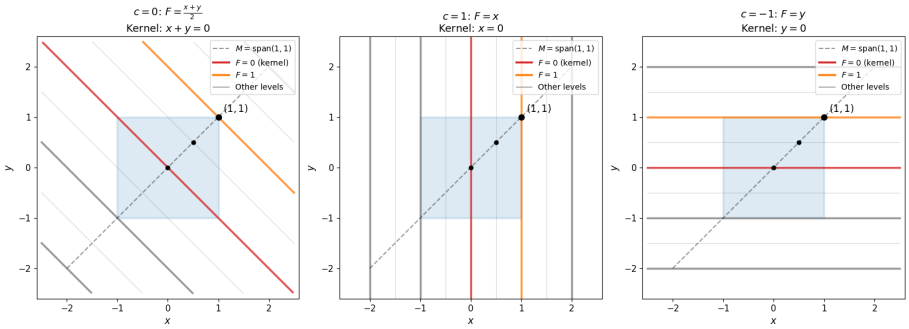
Every c in this interval gives a valid norm-preserving extension. The constraint $|a| + |b| \leq 1$ is the ℓ^1 unit ball in the (a, b)-plane, and $a + b = 1$ is a line cutting through it. The valid extensions correspond to the segment where the line meets the ball.



The ℓ^1 unit ball $|a| + |b| \le 1$ (diamond) is the dual ball of $(\mathbb{R}^2, \|\cdot\|_\infty)$, so the operator norm constraint $\|\mathrm{F}\| \le 1$ requires (a, b) to lie inside it. The extension constraint $a + b = 1$ (ensuring $\mathrm{F} = f$ on M) is a line. Valid extensions are the segment where the line meets the diamond — parametrized by $c \in [-1, 1]$.

Step 5: Each c gives a different kernel.

- $c = 0$: $\mathrm{F}(x, y) = \frac{x+y}{2}$. Kernel: $x + y = 0$ (the anti-diagonal).
- $c = 1$: $\mathrm{F}(x, y) = x$. Kernel: $x = 0$ (the y-axis).
- $c = -1$: $\mathrm{F}(x, y) = y$. Kernel: $y = 0$ (the x-axis).

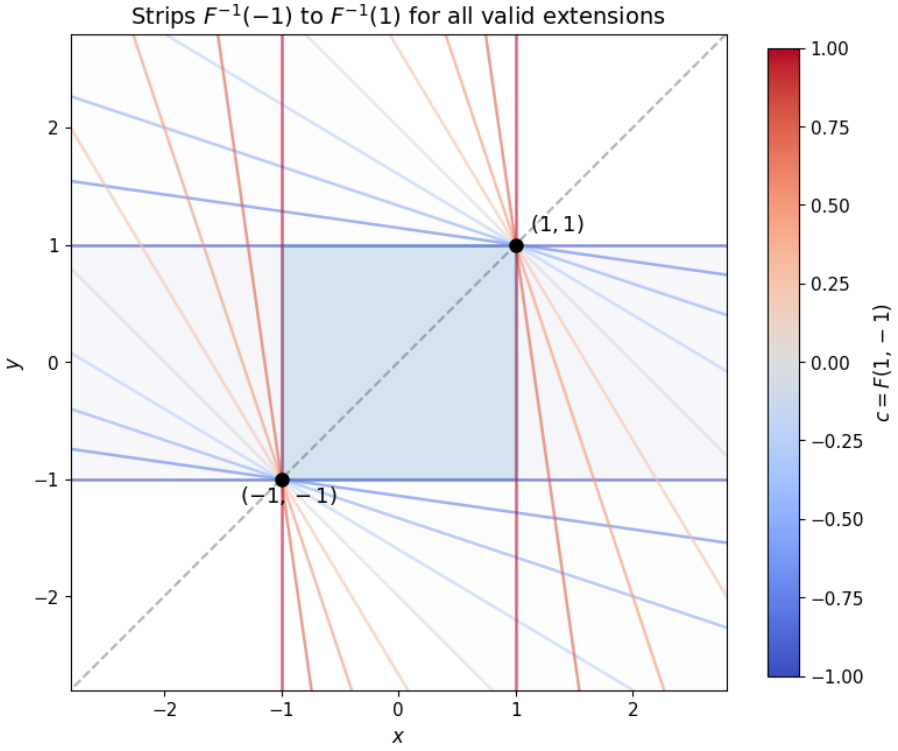


Three norm-preserving extensions of f from $\mathrm{M} = \mathrm{span}(1, 1)$ to all of $\mathbb{R}^2$, corresponding to $c = 0, 1, -1$. In each panel, the ℓ^∞ unit square (shaded) fits between the $\mathrm{F} = -1$ and $\mathrm{F} = 1$ level sets, confirming $\|\mathrm{F}\| = 1$. The kernel (red) and the subspace M (dashed) differ in each case, but $\mathrm{F}(1, 1) = 1$ in all three.

Takeaway

Same functional on M. Three different extensions. Three different kernels. In each case the unit square fits between the level sets $F^{-1}(-1)$ and $F^{-1}(1)$ — the norm constraint $\|F\| \leq 1$ is satisfied because the flat sides of the square allow the kernel to tilt freely. With a round ball, only one tilt would work. **Hahn–Banach says: in any Banach space, no matter how large, at least one valid c always exists.**

VISUALIZING THE FULL FAMILY OF EXTENSIONS Every $c \in [-1, 1]$ gives a different valid extension. Each extension defines a **strip** between $F^{-1}(-1)$ and $F^{-1}(1)$, and the norm constraint $\|F\| = 1$ means the unit square must fit inside that strip. The unit ball is the intersection of all such strips.



Each valid extension F *(parametrized by* $c \in [-1, 1]$*) defines a strip between* $F^{-1}(-1)$ *and* $F^{-1}(1)$*, shown as a pair of lines in matching color. The* ℓ^∞ *unit square (shaded) fits inside every strip. All* $F^{-1}(1)$ *lines pass through* $(1, 1)$ *and all* $F^{-1}(-1)$ *lines pass through* $(-1, -1)$*. The unit ball is exactly the intersection of all these strips, which is the geometric content of the sup formula* $\|x\| = \sup_{\|F\| \leq 1} |F(x)|$.

3.5 *The Hahn–Banach Theorem*

We saw that X^* could be trivial (L^p for $0 < p < 1$), and asked whether the dual is always a complete measurement system for Banach spaces. The following theorem answers yes: every bounded linear functional on a subspace extends to the whole space with the same norm. This guarantees that X^* is rich enough to separate points, recover the norm, and support the weak topology.

3.16 THEOREM (Hahn–Banach theorem). *Let* X *be a normed space and* $M \subset X$ *a linear subspace. Let* $f \in M^*$ *with* $|f(x)| \leq k\|x\|$ *for all* $x \in M$. *Then there exists an extension* $F \in X^*$ *of* f *to all of* X *with* $|F(x)| \leq k\|x\|$ *for all* $x \in X$.

Note *that* F *is not necessarily unique.*

Proof. The proof has three ingredients. The first two are structural; the third is where the analysis lives.

Ingredient 1: Zorn's lemma setup.

Define the poset:

$$E = \{(g, G) : G \supseteq M \text{ is a subspace, } g \text{ extends } f \text{ on } G,\ |g(x)| \leq k\|x\|\ \forall x \in G\} \tag{74}$$

Order by extension: $(h, H) \leq (g, G)$ if $H \subseteq G$ and $g|_H = h$. This is nonempty ($(f, M) \in E$), chains have upper bounds (take the union), so Zorn gives a maximal element (F, W).

Ingredient 2: Maximal implies total.

Suppose $W \neq X$. Pick $z \in X \setminus W$ and set $Z = W + \operatorname{span}(z)$. By the Proposition 3.6, every element of Z can be written uniquely as $w + \alpha z$ with $w \in W$, $\alpha \in \mathbb{R}$. We want to define:

$$\tilde{F}(w + \alpha z) = F(w) + \alpha c \tag{75}$$

for some constant c. If we can find *any* valid c, then $(\tilde{F}, Z)$ strictly extends (F, W), contradicting maximality.

Ingredient 3: The constant c exists.

The requirement $|\tilde{F}(w + \alpha z)| \leq k\|w + \alpha z\|$ constrains c to lie in an interval. The existing bound $|F(\cdot)| \leq k\|\cdot\|$ on W, combined with the **triangle inequality**, guarantees that the lower bound on c is $\leq$ the upper bound. The interval is nonempty.

Since c exists, we can extend, contradicting maximality. Therefore $W = X$.

□

What to take away

- ***Ingredients 1 and 2*** *follow the same Zorn's lemma template as the Hamel basis proof, combined with the codimension-1 decomposition. No new ideas are needed.*
- ***Ingredient 3*** *is the only piece that requires actual work, and the only tool it uses is the* ***triangle inequality****. This is where the assumption that we're in a normed space matters — and it's exactly the property that fails in* L^p *for* $0 < p < 1$, *which is why the dual collapses there.*

3.5.1 Geometric Consequences

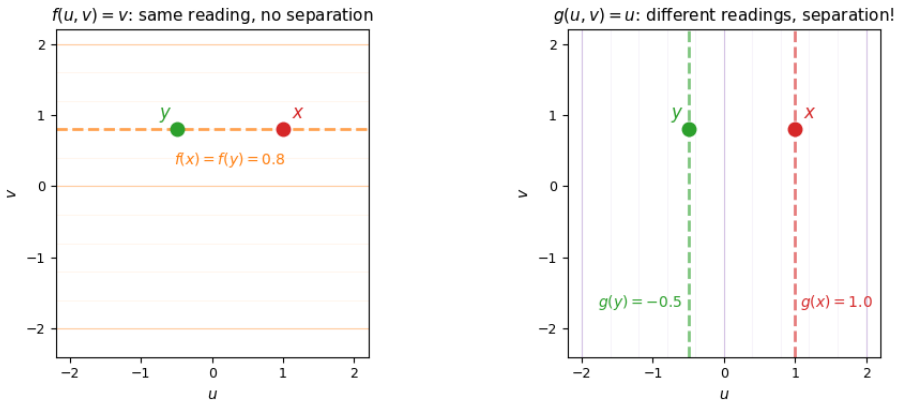
The corollaries of Hahn–Banach are all **separation results** — they say you can always find a hyperplane that separates things.

3.17 COROLLARY (The distinguishing property). *Let $x, y \in X$. If $f(x) = f(y)$ for all $f \in X^*$ then $x = y$.*

Proof. If $x \neq y$, define $\varphi(\alpha(x-y)) = \alpha\|x-y\|$ on $\operatorname{span}(x-y)$. Then $\|\varphi\| = 1$ and Hahn–Banach extends φ to $F \in X^*$ with $F(x) - F(y) = F(x-y) = \|x-y\| \neq 0$, contradicting the assumption.

□

The contrapositive gives the intuition: if two points are on the *same* level set for *every possible* foliation, no matter how you orient the hyperplanes, then they must be the same point.

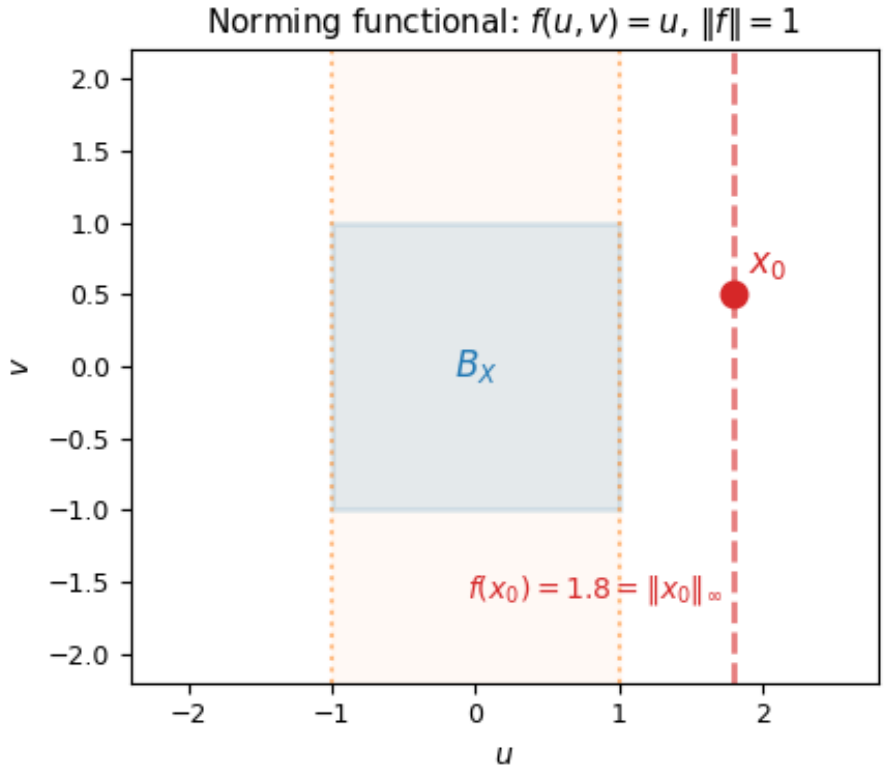


Left: the functional $f(u,v) = v$ gives the same reading for x and y (they lie on the same horizontal level set), so f alone cannot distinguish them. Right: the functional $g(u,v) = u$ gives different readings (different vertical level sets) and separates the two points. No single functional suffices; we need the entire dual X^* to guarantee that distinct points are always distinguishable.

3.18 COROLLARY (Norming property). *For each $x \in X$, there exists $f \in X^*$ with $f(x) = \|x\|$ and $\|f\|_{op} = 1$.*

Proof. Define $f(\alpha x) = \alpha\|x\|$ on $\operatorname{span}(x)$. Then $|f(\alpha x)| = |\alpha|\|x\| = \|\alpha x\|$, so $\|f\| = 1$ on $\operatorname{span}(x)$. Hahn–Banach extends to $F \in X^*$ with $\|F\| = 1$ and $F(x) = \|x\|$.

□



The functional $f(u, v) = u$ has $\|f\| = 1$ (the unit ball fits inside the strip $|f| \leq 1$) and achieves $f(x_0) = \|x_0\|_\infty = 1.8$. The reading equals the norm exactly.

3.19 COROLLARY (The sup formula). *Let X be a normed space and $B_{X^*} = \{f \in X^* : \|f\| \leq 1\}$ the closed unit ball of the dual. Then for all $x \in X$:*

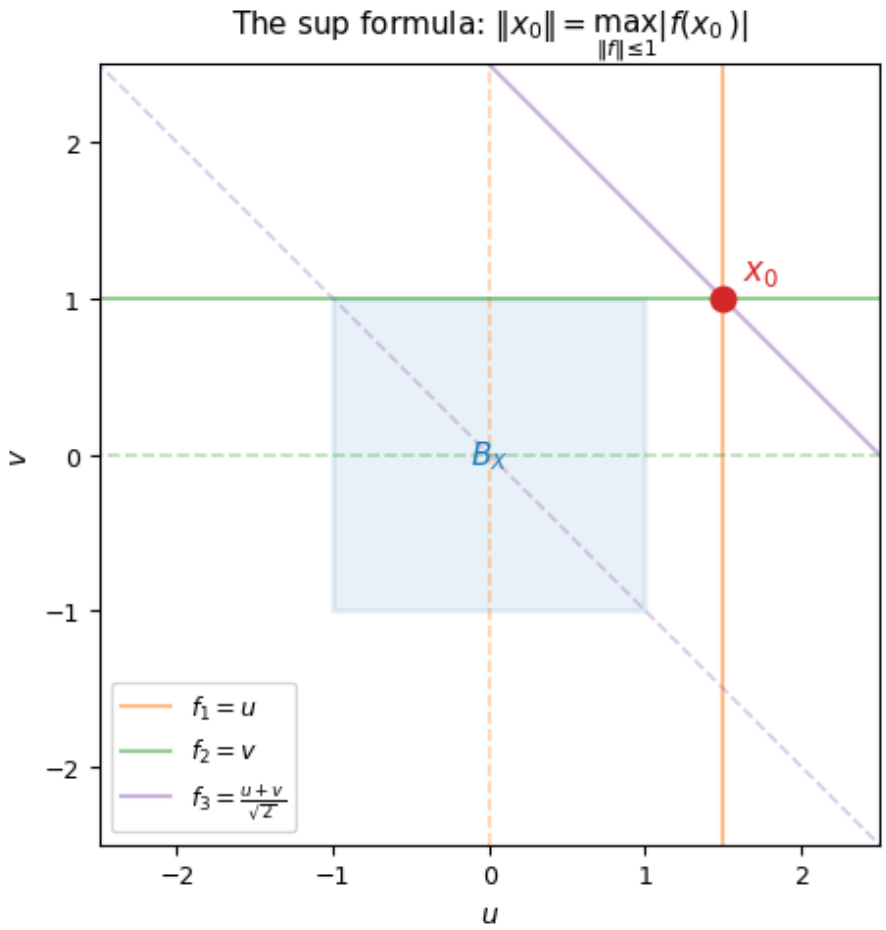
$$\|x\| = \sup\{|f(x)| : f \in B_{X^*}\} = \max\{|f(x)| : f \in B_{X^*}\} \tag{76}$$

*In particular, the supremum is **attained**: there exists $f \in B_{X^*}$ with $|f(x)| = \|x\|$. The dual X^* is a complete measurement system with no information loss.*

Proof. The inequality $\leq$: for any f with $\|f\| \leq 1$, we have $|f(x)| \leq \|f\|\|x\| \leq \|x\|$, so $\sup \leq \|x\|$.

The inequality $\geq$ (and that it is a max): by the Corollary 3.18, there exists $f \in X^*$ with $\|f\| = 1$ and $f(x) = \|x\|$. This f attains the supremum.

□



Three unit-norm functionals give different readings of $x_0 = (1.5, 1.0)$: $f_1(u, v) = u$ reads 1.5, $f_2(u, v) = v$ reads 1.0, and $f_3 = (u + v)/\sqrt{2}$ reads ≈ 1.25. The best reading equals $\|x_0\|_\infty = 1.5$. The norming property guarantees such an optimal functional always exists.

3.20 COROLLARY (Classification of closures). *Let* M *be a linear subspace of a normed space* X *and let $x_0 \in$ X. Then $x_0 \in \overline{\text{M}}$ if and only if there exists* ***no*** *bounded linear functional f such that $f(x) = 0$ for all $x \in$ M but $f(x_0) \neq 0$.*

Proof. ($\Rightarrow$) If $x_0 \in \overline{\text{M}}$, pick $x_n \in$ M with $x_n \to x_0$. For any $f \in \text{X}^*$ vanishing on M, continuity gives $f(x_0) = \lim f(x_n) = 0$.

($\Leftarrow$) Suppose $x_0 \notin \overline{\text{M}}$. Then $\overline{\text{M}}$ is a proper closed subspace and $d = \text{dist}(x_0, \overline{\text{M}}) > 0$. Define g on $\overline{\text{M}} + \text{span}(x_0)$ by $g(m + \alpha x_0) = \alpha$. Then $|g(m + \alpha x_0)| = |\alpha| \le \|m + \alpha x_0\|/d$ (since $\|m + \alpha x_0\| \ge |\alpha| d$). So $\|g\| \le 1/d$. Hahn-Banach extends g to $f \in \text{X}^*$. By construction $f = 0$ on M but $f(x_0) = 1 \neq 0$.

□

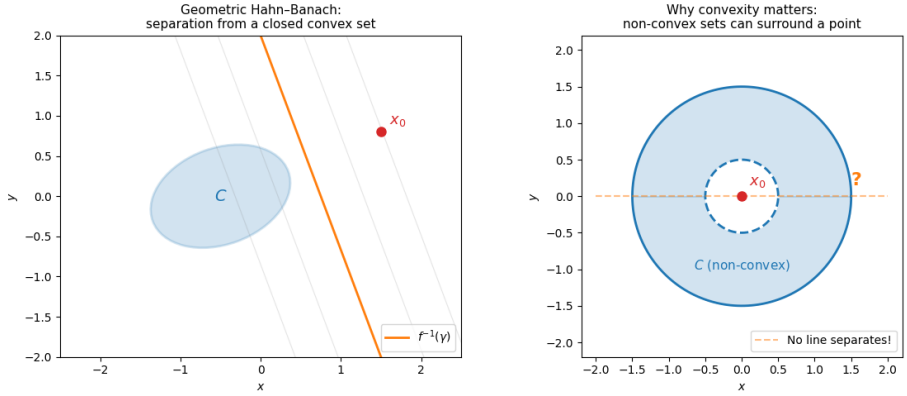
3.5.2 Geometric Hahn–Banach: Separation of Convex Sets

The corollaries above separate points from points and points from the unit ball. The **geometric form of Hahn–Banach** separates a point from any closed convex set.

3.21 THEOREM (Separation of point from closed convex set). *Let* $C \subset X$ *be a closed convex set and* $x_0 \notin C$. *Then there exists* $f \in X^*$ *and a constant* γ *such that*

$$f(x_0) > \gamma \geq f(c) \quad \text{for all } c \in C \tag{77}$$

That is, there is a closed hyperplane $f^{-1}(\gamma)$ *that* ***strictly separates*** x_0 *from* C.



WHY CLOSEDNESS AND CONVEXITY MATTER

- **Closedness:** If C is not closed, separation can fail. Consider $C = \{(x, y) : y > 0\} \cup \{(0, 0)\}$ and $x_0 = (1, 0)$. The point x_0 is not in C, but any separating line would need to separate from points arbitrarily close to the entire x-axis.
- **Convexity:** If C is not convex, a hyperplane can't separate — it's a flat cut, and non-convex sets can wrap around a point.

3.6 *Dual Spaces: From Self-Dual to Non-Reflexive*

The behavior of the dual space X^* varies dramatically across different Banach spaces. There are three main classes: self-dual (Hilbert), reflexive, and non-reflexive. The geometry of the unit ball is what drives the distinction.

3.6.1 Hilbert Spaces: $H^* \cong H$ (Self-Dual)

In a Hilbert space, every continuous linear functional comes from the inner product.

3.22 THEOREM (Riesz Representation Theorem). *Let* H *be a Hilbert space. For every* $f \in H^*$*, there exists a unique* $x_f \in H$ *such that*

$$f(y) = \langle x_f, y \rangle \quad \textit{for all } y \in H \tag{78}$$

and $\|f\| = \|x_f\|$*. The map* $f \mapsto x_f$ *is an isometric (conjugate-linear) isomorphism* $H^* \cong H$*.*

Proof. **Existence.** If $f = 0$, take $x_f = 0$. Otherwise $\ker(f)$ is a proper closed subspace of H. By the orthogonal decomposition, $H = \ker(f) \oplus \ker(f)^\perp$, and since $\ker(f) \neq H$, the orthogonal complement $\ker(f)^\perp$ is one-dimensional (because $\ker(f)$ has codimension 1). Let z be the unit vector in $\ker(f)^\perp$, so $\|z\| = 1$.

Now decompose any $y \in H$ using this splitting: write $y = y_0 + \alpha z$ where $y_0 \in \ker(f)$ and $\alpha \in \mathbb{R}$. Apply f to both sides: $f(y) = f(y_0) + \alpha f(z) = \alpha f(z)$, so $\alpha = f(y)/f(z)$.

Set $x_f = \overline{f(z)}\, z$. Then since $x_f \in \ker(f)^\perp$, we have $\langle x_f, y_0 \rangle = 0$, and:

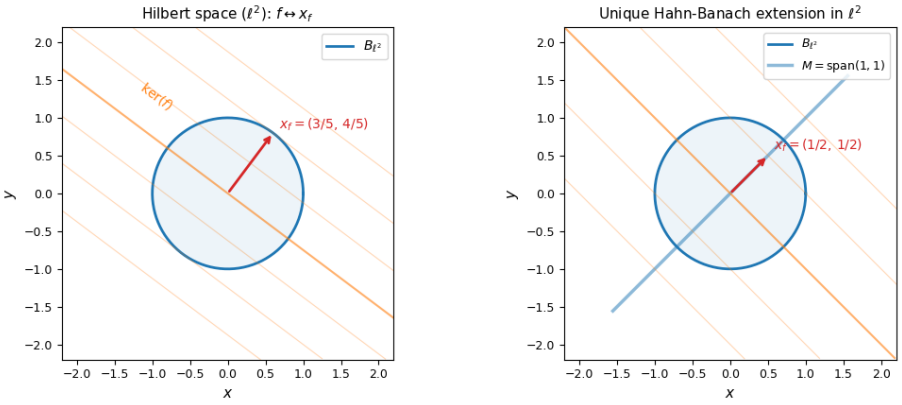
$$\langle x_f, y \rangle = \langle x_f, y_0 + \alpha z \rangle = \alpha \langle x_f, z \rangle = \frac{f(y)}{f(z)} \cdot \overline{f(z)} = f(y). \tag{79}$$

Uniqueness. If $\langle x_1, y \rangle = \langle x_2, y \rangle$ for all y, then $\langle x_1 - x_2, y \rangle = 0$ for all y. Taking $y = x_1 - x_2$ gives $x_1 = x_2$.

Isometry. $\|f\| = \sup_{\|y\| \leq 1} |\langle x_f, y \rangle| = \|x_f\|$ by Cauchy-Schwarz, with equality at $y = x_f/\|x_f\|$.

□

THE $\mathbb{R}^2$ PICTURE In $\mathbb{R}^2$ with the Euclidean norm, every functional $f(x, y) = ax + by$ is represented by the vector (a, b). The kernel $\ker(f)$ is the line perpendicular to (a, b), and the level sets are parallel lines spaced $1/\|(a, b)\|$ apart. The Riesz representative *is* the normal vector to the foliation.



Left: in $(\mathbb{R}^2, \|\cdot\|_2)$*, the functional* $f(x,y) = \frac{3}{5}x + \frac{4}{5}y$ *is represented by the vector* $x_f = (3/5, 4/5)$*, which is perpendicular to* $\ker(f)$*. Right: given* f *on* $\mathrm{M} = \mathrm{span}(1,1)$ *with* $f(1,1) = 1$*, the unique norm-preserving extension is* $\mathrm{F}(x,y) = (x+y)/2$*. The round ball forces a single choice of kernel, unlike the* ℓ^∞ *worked example in the previous section, where the square ball allowed a whole family of extensions.*

3.23 EXAMPLE (and). $\mathbb{R}^n \ell^2$

In $\mathbb{R}^n$ with the standard inner product, every functional is $f(y) = a \cdot y$ for some $a \in \mathbb{R}^n$. Thus $(\mathbb{R}^n)^* \cong \mathbb{R}^n$.

More generally, in $\ell^2 = \{(x_n) : \sum |x_n|^2 < \infty\}$, every functional is $f(x) = \sum a_n x_n$ for a unique $(a_n) \in \ell^2$, and $\|f\| = \|(a_n)\|_{\ell^2}$.

3.24 EXAMPLE. $\mathrm{L}^2(\Omega)$

In $\mathrm{L}^2(\Omega)$ every functional is given by

$$f[x] = \int_\Omega x(t)\, g(t)\, \mathrm{d}t \quad \text{for a unique } g \in \mathrm{L}^2(\Omega). \tag{80}$$

This is the Riesz Representation Theorem applied to L^2. The representing function g plays the role of the "normal vector" to $\ker(f)$.

The key property of Hilbert spaces: the inner product provides a *canonical* identification between H and H^*. Duality is not an abstract theorem but a concrete computation. Every closed hyperplane has a unique normal vector, and extending a functional from a subspace amounts to orthogonal projection.

3.6.2 Reflexive Spaces: $\mathrm{X}^{**} \cong \mathrm{X}$

3.25 DEFINITION (Reflexive space). A Banach space X is **reflexive** if the canonical embedding $\mathrm{J} : \mathrm{X} \to \mathrm{X}^{**}$ defined by

$$\mathrm{J}[x](f) = f(x) \quad \text{for all } f \in \mathrm{X}^* \tag{81}$$

is surjective (and hence an isometric isomorphism onto X^{**}).

3.26 LEMMA (The canonical embedding is an isometry). *For any normed space* X, *the canonical embedding* $J : X \to X^{**}$ *is an isometry:* $\|J[x]\|_{X^{**}} = \|x\|_X$ *for all* $x \in X$.

Proof. By definition, $\|J(x)\|_{X^{**}} = \sup_{\|f\|\le 1} |J[x](f)| = \sup_{\|f\|\le 1} |f(x)|$. The Corollary 3.19 gives $\sup_{\|f\|\le 1} |f(x)| = \|x\|$.

□

So X always embeds isometrically into X^{**}. Reflexivity asks: does every element of X^{**} come from a point in X? That is, does J miss anything?

All Hilbert spaces are reflexive (since $H^* \cong H$ implies $H^{**} \cong H^* \cong H$). But reflexivity is strictly weaker than self-duality.

The L^p family

3.27 EXAMPLE (spaces for). $L^p 1 < p < \infty$

The spaces $L^p(\Omega)$ for $1 < p < \infty$ are reflexive. The dual is $(L^p)^* \cong L^q$ where $\frac{1}{p} + \frac{1}{q} = 1$.

For $g \in L^q(\Omega)$, the functional

$$L_g(f) = \int_\Omega f(x)\, g(x)\, dx \tag{82}$$

satisfies $\|L_g\| = \|g\|_q$ by Hölder's inequality. The map $g \mapsto L_g$ is an isometric isomorphism.

Proof. **Step 1:** L_g **is bounded and** $\|L_g\| \le \|g\|_q$**.** For any $f \in L^p$, Hölder's inequality gives

$$|L_g(f)| = \left| \int_\Omega f\, g \right| \le \|f\|_p \, \|g\|_q, \tag{83}$$

so $L_g \in (L^p)^*$ with $\|L_g\| \le \|g\|_q$.

Step 2: $\|L_g\| = \|g\|_q$ **(the bound is sharp).** Assume $g \neq 0$ and define the test function

$$f_0(x) = \frac{|g(x)|^{q-1} \operatorname{sgn}(g(x))}{\|g\|_q^{q/p}}. \tag{84}$$

Then $|f_0|^p = |g|^{(q-1)p} / \|g\|_q^q = |g|^q / \|g\|_q^q$, so $\|f_0\|_p = 1$. Evaluating:

$$L_g(f_0) = \int_\Omega f_0\, g = \frac{1}{\|g\|_q^{q/p}} \int_\Omega |g|^q = \frac{\|g\|_q^q}{\|g\|_q^{q/p}} = \|g\|_q, \tag{85}$$

where we used $(q-1) + 1 = q$ and $q - q/p = q(1 - 1/p) = 1$. So $\|L_g\| \ge \|g\|_q$, and combined with Step 1, $\|L_g\| = \|g\|_q$. In particular $g \mapsto L_g$ is an isometry.

Step 3: Surjectivity. Let $\phi \in (L^p)^*$. Define a set function ν on measurable sets by $\nu(E) = \phi(\mathbf{1}_E)$. Since $|\nu(E)| = |\phi(\mathbf{1}_E)| \le \|\phi\| \, \|\mathbf{1}_E\|_p = \|\phi\| \, \mu(E)^{1/p}$, we have $\nu(E) \to 0$ whenever $\mu(E) \to 0$, so ν is absolutely

continuous with respect to μ. By the Radon-Nikodym theorem, there exists $g \in L^1(\Omega)$ with $\nu(E) = \int_E g \, d\mu$ for all measurable E. This gives $\phi(s) = \int_\Omega s\, g$ for all simple functions s. Since simple functions are dense in L^p and ϕ is continuous, $\phi(f) = \int_\Omega f\, g$ for all $f \in L^p$, i.e., $\phi = L_g$. The isometry from Step 2 forces $g \in L^q$ with $\|g\|_q = \|\phi\|$.

□

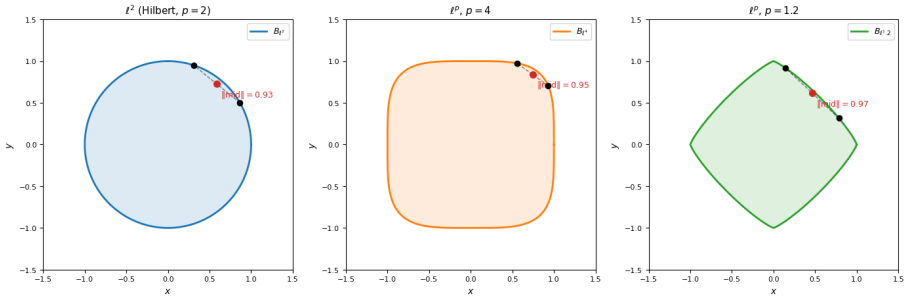
The duality pairing $(L^p)^* = L^q$ and $(L^q)^* = L^p$ closes the loop: $X^{**} = (L^q)^* = L^p = X$.

THE GEOMETRY OF REFLEXIVE SPACES What makes reflexive spaces special geometrically? The answer involves the shape of the unit ball.

3.28 DEFINITION (Uniform convexity). A Banach space X is **uniformly convex** if for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\|x\| = \|y\| = 1, \quad \|x - y\| \geq \varepsilon \implies \left\| \frac{x+y}{2} \right\| \leq 1 - \delta. \tag{86}$$

In words: if two points on the unit sphere are separated, their midpoint is strictly inside the ball. The boundary has no flat spots.



ℓ^p unit balls for three values of $p > 1$. In each case, the midpoint of two boundary points (red dot) lies strictly inside the ball. This is uniform convexity: the boundary curves inward, with no flat edges. Contrast with the ℓ^1 diamond and ℓ^∞ square, which have flat sides where the midpoint stays on the boundary.

FROM GEOMETRY TO REFLEXIVITY Recall what reflexivity asks: is every $\xi \in X^{**}$ of the form $\xi = J(x)$ for some $x \in X$? An element $\xi \in X^{**}$ is a linear functional on X^*, so it assigns a number $\xi(f)$ to every $f \in X^*$. If $\xi = J(x)$, then $\xi(f) = f(x)$, meaning the readings come from evaluating at a point. Non-reflexivity means there exist "phantom" evaluation functionals in X^{**} that do not come from any point in X.

The shape of the unit ball controls this through the **uniqueness of norm-attaining directions**:

- **Round ball** (uniformly convex): for any $\xi \in X^{**}$ with $\|\xi\| = 1$, the functional on X^* that ξ defines picks out a unique "direction" in X. The curvature of the ball means there is only one point on the unit sphere that could produce the readings $\xi(f) = f(x)$. This pins ξ down to an actual point in X.
- **Flat ball** (ℓ^1, ℓ^∞): the flat edges mean many points on the boundary give the same readings for large classes of functionals. This ambiguity creates room for elements of X^{**} that are consistent with the readings but do not correspond to any single point.

A concrete example: in c_0 (sequences converging to zero, sup norm), the canonical embedding $J(c_0) \subset \ell^\infty = c_0^{**}$. The constant sequence $(1, 1, 1, \dots) \in \ell^\infty$ defines a valid element of c_0^{**}, but it is not in c_0 since it does not converge to zero. The flat geometry of the ℓ^∞ ball accommodates this extra element.

3.29 THEOREM (Milman-Pettis). *Every uniformly convex Banach space is reflexive.*

Proof. **Idea.** We must show every $\xi \in X^{**}$ equals $J(x)$ for some $x \in X$. The key step is: if two points $x, y \in B_X$ produce *nearly the same readings* on all functionals (i.e., $|f(x) - f(y)|$ is small for all $f \in B_{X^*}$), then x and y must be *close in norm*.

Suppose not: $\|x - y\| \geq \varepsilon$ but the readings are close. By uniform convexity, $\|(x + y)/2\| \leq 1 - \delta$. But the readings of the midpoint are the averages of the readings of x and y, which are both close to $\xi(f)$. Choosing f that nearly norms $(x + y)/2$ gives a contradiction.

This means: if a sequence of "approximate representatives" $x_n \in B_X$ satisfies $f(x_n) \to \xi(f)$ for all f, then (x_n) is Cauchy in norm. Since X is complete, $x_n \to x$ for some $x \in X$, and $\xi = J(x)$.

□

The L^p spaces for $1 < p < \infty$ are uniformly convex (Clarkson's inequalities), so Milman-Pettis gives reflexivity as a consequence of their round unit balls.

Reflexivity is really about the *duality structure* (does J hit all of X^{**}?), not about any single geometric property. Uniform convexity is one sufficient condition, but spaces can be reflexive for other reasons, for instance because their dual pairing closes the loop as with $(L^p)^* = L^q$ and $(L^q)^* = L^p$.

3.6.3 Non-Reflexive Spaces: $X \subsetneq X^{**}$

When X is not reflexive, the canonical embedding $J : X \hookrightarrow X^{**}$ is a strict inclusion. The bidual X^{**} contains "phantom" elements that do not correspond to any point in X. Geometrically, the unit ball has flat spots or corners that create room for extra functionals.

L^1 AND L^∞

3.30 EXAMPLE (and). $L^1 L^\infty$

$(L^1(\Omega))^* \cong L^\infty(\Omega)$. The pairing is

$$L_g(f) = \int_\Omega f(x)\, g(x)\, dx, \quad g \in L^\infty. \tag{87}$$

But the duality chain does not close: $(L^\infty)^* \supsetneq L^1$. The dual of L^∞ consists of **finitely additive measures**, which include singular objects that cannot be represented by L^1 functions.

A **finitely additive measure** is a set function $\mu : \mathcal{A} \to \mathbb{R}$ satisfying $\mu(A \cup B) = \mu(A) + \mu(B)$ for disjoint A, B, but *not* necessarily for countable unions. Every countably additive (i.e., ordinary) measure is finitely additive, but not conversely. The extra elements of $(L^\infty)^*$ come from finitely additive measures that are *not* countably additive, such as **Banach limits** (which assign a "limit" to every bounded sequence, extending the usual limit).

The L^1 unit ball is not uniformly convex. In $\mathbb{R}^2$, the ℓ^1 ball (diamond) has flat edges connecting $(1, 0)$ to $(0, 1)$, where midpoints of boundary points remain on the boundary. This flat geometry is what allows non-reflexivity.

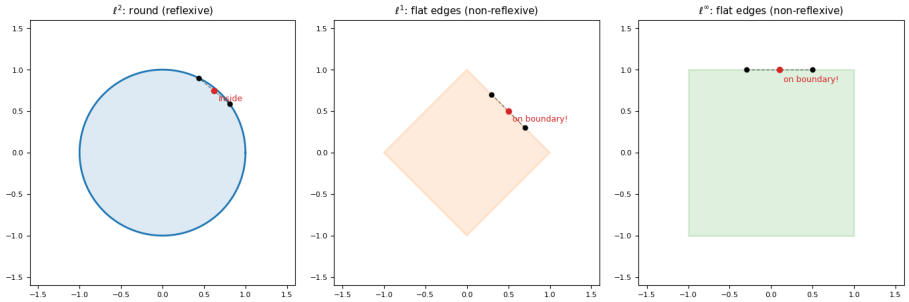
Sequence spaces: c_0 and ℓ^1

3.31 example (The). $c_0 \to \ell^1 \to \ell^\infty$ **chain**

$c_0 = \{(x_n) : x_n \to 0\}$ with the sup norm. The duality chain is:

$$c_0^* \cong \ell^1, \qquad (\ell^1)^* \cong \ell^\infty, \qquad (\ell^\infty)^* \supsetneq \ell^1. \tag{88}$$

Neither c_0 nor ℓ^1 is reflexive. The embedding $J : c_0 \hookrightarrow \ell^\infty = c_0^{**}$ misses sequences that do not converge to zero: for instance, the constant sequence $(1, 1, 1, \ldots) \in \ell^\infty$ is in c_0^{**} but not in $J(c_0)$.



Reflexive vs. non-reflexive geometry. In ℓ^2 (left), the midpoint of two boundary points is strictly inside the ball: uniform convexity. In ℓ^1 (center) and ℓ^∞ (right), flat edges allow midpoints to remain on the boundary. This failure of uniform convexity is the geometric signature of non-reflexivity.

C(K) and Radon measures

This is one of the most important dualities in analysis.

3.32 definition (Radon measure). Let K be a compact Hausdorff space. A **Radon measure** on K is a Borel measure μ that is:

- **Finite:** $|\mu|(K) < \infty$.
- **Inner regular:** $|\mu|(E) = \sup\{|\mu|(C) : C \subseteq E,\ C \text{ compact}\}$ for every Borel set E.
- **Outer regular:** $|\mu|(E) = \inf\{|\mu|(U) : U \supseteq E,\ U \text{ open}\}$ for every Borel set E.

A **signed Radon measure** is a difference $\mu = \mu^+ - \mu^-$ of two (positive) Radon measures with $|\mu|(K) = \mu^+(K) + \mu^-(K) < \infty$. The space $\mathcal{M}(K)$ of all finite signed Radon measures on K is a Banach space under the total variation norm $\|\mu\| = |\mu|(K)$.

Examples include absolutely continuous measures (with L^1 densities), point masses δ_x (where $\delta_x(f) = f(x)$), and singular measures such as the Cantor measure.

3.33 theorem (Riesz-Markov-Kakutani Representation Theorem). *Let K be a compact Hausdorff space. Then every bounded linear functional $\varphi \in C(K)^*$ is represented by a unique finite signed Radon measure μ on K:*

$$\varphi(f) = \int_K f \, d\mu \quad \textit{for all } f \in C(K) \tag{89}$$

and $\|\varphi\| = |\mu|(K)$ *(total variation). That is,* $C(K)^* \cong \mathcal{M}(K)$*, the space of finite signed Radon measures on* K.

Proof. The proof constructs the measure μ from the functional φ in several steps.

Positive case. If $\varphi \geq 0$ (i.e., $f \geq 0 \Rightarrow \varphi(f) \geq 0$), define for each open $U \subseteq K$:

$$\mu(U) = \sup\{\varphi(f) : 0 \leq f \leq 1,\ \operatorname{supp}(f) \subseteq U\} \tag{90}$$

This defines an outer-regular Borel measure. Urysohn's lemma (which requires K Hausdorff and compact) provides the continuous functions needed to approximate indicator functions, and the Riesz representation follows from monotone convergence arguments.

General case. Decompose $\varphi = \varphi^+ - \varphi^-$ into positive and negative parts (the Jordan decomposition of functionals), apply the positive case to each, and set $\mu = \mu^+ - \mu^-$.

□

3.34 EXAMPLE (Key instances).

- $C([0,1])^* = \mathcal{M}([0,1])$: functionals on continuous functions are signed measures. This includes absolutely continuous measures (represented by L^1 densities), singular measures, and point masses δ_x.
- $C(\mathbb{R}^n)$ (with compact support or vanishing at infinity): the dual captures probability distributions, empirical measures, and distributional objects.

The space C(K) is generally **not reflexive**. The dual $\mathcal{M}(K)$ is much larger than C(K): it contains point masses δ_x (which are not continuous functions) and singular measures. The bidual $C(K)^{**} = \mathcal{M}(K)^*$ is larger still.

This duality is central to probability theory (measures as linear functionals on observables), PDE (weak solutions via integration against test functions), and optimization (dual problems over measures). It also provides the setting for the weak* convergence of measures discussed in the next section.

3.6.4 Spaces with Trivial Dual

3.35 EXAMPLE (for). $L^p 0 < p < 1$

The spaces $L^p(\Omega)$ for $0 < p < 1$ have trivial dual: $(L^p)^* = \{0\}$.

These are complete metric spaces (F-spaces) but not Banach spaces. The "norm" $\|f\|_p = \int |f|^p$ fails the triangle inequality. The proof (due to Day, 1940) shows that the only convex open subsets of L^p for $0 < p < 1$ are $\emptyset$ and L^p itself. Since any nonzero continuous linear functional would have $f^{-1}((-1,1))$ as a proper convex open set, no such functional exists.

This is the extreme case where the dual sees nothing. It is precisely what the Hahn-Banach theorem prevents in Banach spaces: the convexity of the unit ball (guaranteed by the triangle inequality) is what keeps the dual nontrivial.

3.7 *Application: The Universal Approximation Theorem*

We give a clean application of the Hahn-Banach theorem to neural networks: Cybenko's proof [1] that single-hidden-layer networks are dense in $C(\Omega)$.

3.7.1 Single-hidden-layer networks

Fix $\Omega = [0,1]^d$ and an activation function $\sigma : \mathbb{R} \to \mathbb{R}$. A **single-hidden-layer network** with M neurons is a function of the form

$$f(x) = \sum_{j=1}^{M} a_j \, \sigma(w_j \cdot x + b_j), \qquad x \in \mathbb{R}^d \tag{91}$$

where $w_j \in \mathbb{R}^d$ are weights, $b_j \in \mathbb{R}$ are biases, and $a_j \in \mathbb{R}$ are output coefficients. The set of all such functions (for all M) forms a linear subspace of $C(\Omega)$:

$$\mathcal{N}_\sigma = \operatorname{span}\{\sigma(w \cdot x + b) : w \in \mathbb{R}^d,\ b \in \mathbb{R}\} \subset C(\Omega). \tag{92}$$

3.36 DEFINITION (Universal approximator). An activation function σ is a **universal approximator** if $\mathcal{N}_\sigma$ is dense in $C(\Omega)$, i.e. $\overline{\mathcal{N}_\sigma} = C(\Omega)$.

The question: Which activations σ are universal approximators?

3.7.2 Discriminatory activations

3.37 DEFINITION (Discriminatory function). A continuous function $\sigma : \mathbb{R} \to \mathbb{R}$ is **discriminatory** if, for any finite signed Borel measure μ on Ω,

$$\int_\Omega \sigma(w \cdot x + b)\, d\mu(x) = 0 \quad \text{for all } w \in \mathbb{R}^d,\ b \in \mathbb{R} \tag{93}$$

implies $\mu = 0$.

In other words, no nonzero measure can be invisible to every neuron $\sigma(w \cdot x + b)$.

3.7.3 Cybenko's theorem

3.38 THEOREM (Universal Approximation (Cybenko, 1989)). *Let $\sigma : \mathbb{R} \to \mathbb{R}$ be continuous and discriminatory. Then*

$$\overline{\mathcal{N}_\sigma} = C(\Omega). \tag{94}$$

Proof. Suppose for contradiction that $\overline{\mathcal{N}_\sigma} \neq C(\Omega)$.

Step 1 (Hahn-Banach). Since $\overline{\mathcal{N}_\sigma}$ is a proper closed subspace, the Corollary 3.20 gives a nonzero $F \in C(\Omega)^*$ with $F = 0$ on $\mathcal{N}_\sigma$.

Step 2 (Riesz-Markov). By the Theorem 3.33, there is a unique nonzero signed measure μ on Ω with

$$F(\phi) = \int_{\Omega} \phi \, d\mu \quad \text{for all } \phi \in C(\Omega). \tag{95}$$

Step 3 (Contradiction). Since F vanishes on $\mathcal{N}_\sigma$, we have

$$\int_{\Omega} \sigma(w \cdot x + b) \, d\mu(x) = 0 \quad \text{for all } w,\ b. \tag{96}$$

Because σ is discriminatory, $\mu = 0$, contradicting $F \neq 0$.

□

3.7.4 Which activations are discriminatory?

The definition asks that no nonzero measure hides from every neuron. The key observation is that steep rescalings of σ approximate indicator functions of half-spaces.

3.39 PROPOSITION (Sigmoidal activations are discriminatory). *Any bounded measurable function σ with*

$$\sigma(t) \to \begin{cases} 1 & t \to +\infty \\ 0 & t \to -\infty \end{cases} \tag{97}$$

is discriminatory.

Proof. For $\lambda > 0$, the rescaled neuron $\sigma(\lambda(w \cdot x + b))$ converges pointwise as $\lambda \to \infty$ to the half-space indicator

$$\mathbf{1}_{H_{w,b}}(x), \qquad H_{w,b} = \{x : w \cdot x + b > 0\}. \tag{98}$$

If $\int \sigma(w \cdot x + b) \, d\mu = 0$ for all w, b, then the same holds for $\sigma(\lambda(w \cdot x + b))$ (replace w by λw and b by λb). Dominated convergence gives

$$\mu(H_{w,b}) = \int \mathbf{1}_{H_{w,b}} \, d\mu = 0 \quad \text{for all half-spaces } H_{w,b}. \tag{99}$$

Half-spaces generate the Borel σ-algebra on Ω, so $\mu = 0$.

□

3.40 PROPOSITION (ReLU is discriminatory). *The ReLU activation $\sigma(t) = \max(0, t)$ is discriminatory.*

Proof. For any w, b and $\lambda > 0$, note that

$$\frac{1}{\lambda}\Big[\mathrm{ReLU}(\lambda(w \cdot x + b)) - \mathrm{ReLU}(\lambda(w \cdot x + b) - 1)\Big] \to \mathbf{1}_{H_{w,b}}(x) \tag{100}$$

pointwise as $\lambda \to \infty$. Since $\int \mathrm{ReLU}(w \cdot x + b)\, d\mu = 0$ for all w, b, the same dominated convergence argument gives $\mu(H_{w,b}) = 0$ for all half-spaces, hence $\mu = 0$.

□

Compare this with the Example 1.14 from our discussion of Density and Approximation: polynomials are dense in $C([a, b])$, and Cybenko's theorem says the same for neural networks in $C([0, 1]^d)$. Both are density results, but there is an important difference. Weierstrass approximation suffers from the **curse of dimensionality**: the number of monomials of degree $\leq n$ in d variables grows as $\binom{n+d}{d}$, which is exponential in d. Neural networks avoid this; the approximation is built from neurons $\sigma(w \cdot x + b)$ that each cut across all d dimensions simultaneously, so the number of terms needed can scale much more favourably with d.

On the other hand, the proof is non-constructive (it rests on Hahn-Banach, which uses Zorn's lemma), so it says nothing about **how many neurons** are needed for a given accuracy εor **how to find the weights**.

3.7.5 A second strategy for density proofs

In the Density and Approximation chapter, every density proof was **constructive**: we built an explicit approximating sequence (Bernstein polynomials for Weierstrass, mollifiers for $C_c^\infty \subset L^p$) and estimated the error directly. Cybenko's proof is fundamentally different. Instead of constructing an approximation, it argues by contradiction via duality:

1. **Assume** $\overline{M} \subsetneq X$.
2. **Hahn-Banach** produces a nonzero functional $F \in X^*$ vanishing on M.
3. A **representation theorem** identifies F concretely (here, as a measure via Riesz-Markov).
4. Problem-specific information (here, the discriminatory property) forces $F = 0$, a contradiction.

This is a general-purpose machine: the same scheme proves density of test functions in L^p, completeness of orthonormal systems, and many other approximation results. Only step 4 changes between applications. The price is that the proof is non-constructive and tells us nothing about rates of convergence, but when a direct construction is unavailable, the duality approach is often the only way in.

3.8 *The Weak Topology*

Big Idea

For applications in PDE and the calculus of variations, we need a topology that is **Hausdorff** (so limits are unique) and has **many compact sets** (so we can extract convergent subsequences). These demands pull in opposite directions. The **weak topology** is the sweet spot: coarse enough for compactness, fine enough for separation.

3.8.1 What we need from a topology

For the existence arguments that drive much of applied analysis, we need two things from a topology on a Banach space X.

Recall that convergence in a topological space is defined in terms of open sets: a sequence (x_n) in a topological space X **converges** to $x \in X$ if for every open set U containing x, there exists $n_0 \in \mathbb{N}$ such that $x_n \in U$ for all $n \geq n_0$. In a metric space, this reduces to the familiar ε-definition, but in a general topological space the open sets may look very different from metric balls.

Hausdorff (unique limits). We need limits to be unique: if $x_n \to x$ and $x_n \to y$, then $x = y$. A topology is **Hausdorff** if any two distinct points can be separated by disjoint open sets. This forces uniqueness: if $x \neq y$, choose disjoint open sets $U \ni x$ and $V \ni y$. Eventually $x_n \in U$ and eventually $x_n \in V$, but $U \cap V = \emptyset$, a contradiction. Without the Hausdorff property, a sequence can converge to multiple points simultaneously (in the indiscrete topology $\{\emptyset, X\}$, every sequence converges to every point), and limit arguments become meaningless.

Compact sets (convergent subsequences). We need to extract convergent subsequences from bounded families. A set K is **compact** if every open cover has a finite subcover. Why does this give subsequences? Take a sequence (x_n) in K. Any open cover of K has a finite subcover $U_1, \ldots, U_m$. Since infinitely many terms of the sequence are distributed among finitely many sets, at least one U_j contains infinitely many x_n. This U_j can be made arbitrarily small (by refining the cover), forcing the terms to cluster. Compactness guarantees that sequences cannot escape.

The tension. These two demands pull in opposite directions. The indiscrete topology $\{\emptyset, X\}$ makes every space trivially compact (there is only one open cover), but no two points can be separated. The discrete topology (every subset is open) makes every space Hausdorff, but only finite sets are compact: given an infinite set A, the cover $\{\{x\} : x \in A\}$ consists of open singletons and has no finite subcover. In general, making a topology coarser helps compactness but threatens separation; making it finer helps separation but destroys compactness.

In finite dimensions, the norm topology gives us both: it is Hausdorff and closed bounded sets are compact (Heine-Borel). In infinite dimensions, the norm topology remains Hausdorff but **loses compactness**: the closed unit ball is never norm-compact (Riesz's theorem). For instance, the standard basis vectors $e_1, e_2, e_3, \ldots$ in ℓ^2 all lie on the unit sphere with $\|e_m - e_n\| = \sqrt{2}$ for $m \neq n$, so no subsequence is Cauchy. The bumps "escape" into new dimensions and never cluster. The norm topology has too many open sets.

The idea is to **weaken the topology** just enough to recover compactness while keeping the Hausdorff property. The Hahn-Banach theorem makes this possible: it guarantees enough continuous linear functionals to separate points, so the topology they generate is Hausdorff. The resulting **weak topology** has fewer open sets (and therefore fewer open covers), making compactness easier to achieve.

3.8.2 From norm balls to weak neighborhoods

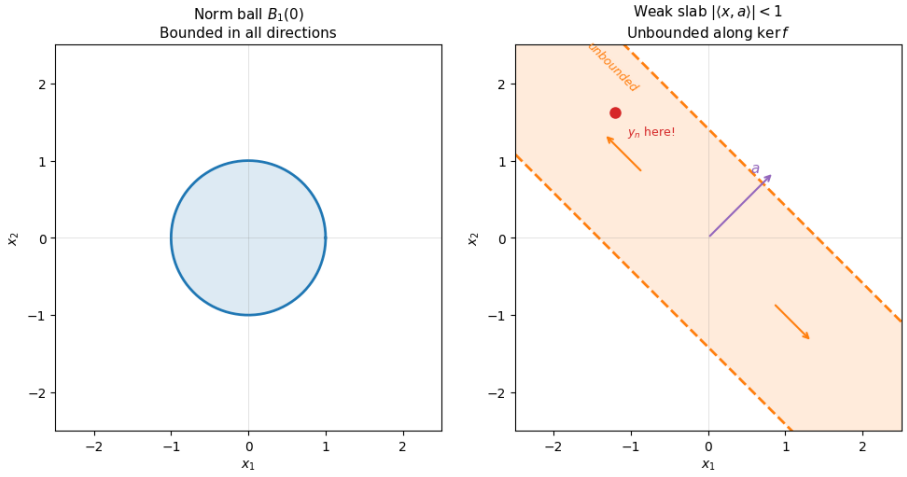
In a normed space X, the **norm topology** is generated by the open balls

$$B_r(x) = \{y \in X : \|y - x\| < r\}. \tag{101}$$

The idea behind the weak topology is to relax the notion of neighborhood. The building blocks are **slabs**: for a single functional $f \in X^*$ and $\varepsilon > 0$, define

$$S(f, x, \varepsilon) = \{y \in X : |f(y - x)| < \varepsilon\}. \tag{102}$$

Geometrically, f defines a family of parallel hyperplanes $H_c = \{x \in X : f(x) = c\}$, and the slab $S(f, x, \varepsilon)$ is the region between $H_{f(x)-\varepsilon}$ and $H_{f(x)+\varepsilon}$: a strip of width $2\varepsilon/\|f\|$ in the direction normal to the hyperplanes. It constrains y only in the direction that f measures, and is **unbounded** in every direction within $\ker f$.



Left: the norm ball is bounded in every direction. Right: the slab $|\langle x, a\rangle| < 1$ constrains only the direction normal to $\ker f$ and extends to infinity along it. The vector a is the normal direction. The point y_n lies inside the slab but far outside the norm ball.

A weak open set is built up in three layers. Recall that a **basis** for a topology is a collection $\mathcal{B}$ of open sets such that every open set in the topology is a union of members of $\mathcal{B}$. Equivalently, for every open set U and every $x \in U$, there exists $B \in \mathcal{B}$ with $x \in B \subseteq U$. In the norm topology, the open balls form a basis. In the weak topology, the role of open balls is played by **tubes**.

1. **Sub-basic open sets (slabs).** A single slab $S(f, x, \varepsilon)$ imposes one linear constraint. It is an infinite strip, constraining one direction and placing no restriction on any direction in $\ker f$.

2. **Basic open sets (finite intersections of slabs).** A finite intersection

$$U(x; f_1, \ldots, f_n, \varepsilon) = \bigcap_{i=1}^{n} S(f_i, x, \varepsilon) \tag{103}$$

constrains n directions simultaneously. Think of this as a **tube**: bounded cross-section in finitely many directions, infinite extent in all others. These tubes form a **basis** for the weak topology.

3. **General open sets (arbitrary unions of tubes).** Every weakly open set is a union of tubes, just as every norm-open set is a union of open balls.

The key point: in infinite dimensions, every tube is unbounded (since $n < \dim X$ always). Since every point in a weakly open set has a tube around it, and every tube is unbounded, **every nonempty weakly open set is unbounded** in infinite dimensions. No bounded set like a norm ball can be weakly open, and no bounded set can even *contain* a weakly open set.

3.41 DEFINITION (Weak topology). The **weak topology** on X, denoted $\sigma(X, X^*)$, is the coarsest topology making every $f \in X^*$ continuous. A neighborhood basis at $x \in X$ consists of the basic open sets:

$$U(x; f_1, \dots, f_n, \varepsilon) = \{y \in X : |f_i(y - x)| < \varepsilon \text{ for } i = 1, \dots, n\} \tag{104}$$

where $f_1, \dots, f_n \in X^*$ and $\varepsilon > 0$.

Continuity of f requires preimages of open sets to be open. Every open set in $\mathbb{R}$ is a union of open intervals, and $f^{-1}((a, b))$ is exactly a slab between two parallel hyperplanes. So any topology making all $f \in X^*$ continuous must contain all slabs. The coarsest such topology is the one generated by the slabs alone, with no extra open sets added.

The weak topology is **Hausdorff**: if $x \neq y$, then by the Hahn-Banach theorem there exists $f \in X^*$ with $f(x) \neq f(y)$, so x and y lie in disjoint slabs. Without Hahn-Banach, we would have no guarantee that X^* contains enough functionals to distinguish points, and the weak topology could fail to be Hausdorff.

3.42 EXAMPLE (The weak topology on). ℓ^2

Take $X = \ell^2$ with its standard orthonormal basis (e_n). By the Riesz representation theorem, every $f \in (\ell^2)^*$ is of the form $f(x) = \langle x, a\rangle$ for some $a \in \ell^2$. Consider the functional $f_1(x) = \langle x, e_1\rangle = x_1$. The slab

$$S(f_1, 0, 1) = \{x \in \ell^2 : |x_1| < 1\} \tag{105}$$

is the region between the hyperplanes $x_1 = -1$ and $x_1 = 1$. This slab contains vectors with arbitrarily large norm, as long as their first component is small. For instance, $y_n = \frac{1}{2}e_1 + n\, e_2$ satisfies $|f_1(y_n)| = \frac{1}{2} < 1$, so $y_n \in S(f_1, 0, 1)$ for every n, yet $\|y_n\| = \sqrt{1/4 + n^2} \to \infty$. The slab sees only the first coordinate and is blind to the e_2 direction.

Adding more functionals narrows the neighborhood but never makes it bounded. The set $U(0; f_1, \dots, f_n, \varepsilon) = \{x : |x_k| < \varepsilon \text{ for } k = 1, \dots, n\}$ constrains n coordinates but leaves infinitely many unconstrained. This is a tube: bounded cross-section in the first n coordinates, unbounded in all others.

3.43 REMARK (Open, closed, or neither). Every set in a topological space falls into exactly one of four categories: (1) open but not closed, (2) closed but not open, (3) both (clopen), or (4) neither. In any topology, the vast majority of sets are in category (4).

When passing from the norm topology to the weak topology, sets can only move **toward** category (4). A norm-open set may lose its openness; a norm-closed set may lose its closedness. No set gains either property, because every weakly open set is norm-open and every weakly closed set is norm-closed. The "neither" category grows at the expense of all three others.

For example, in ℓ^2 with the norm topology, the half-open shell $\{x : 1 \le \|x\| < 2\}$ is neither open nor closed. In the weak topology, the open ball $\{x : \|x\| < 1\}$ also joins the "neither" category: it is no longer weakly open (it is bounded), and it is not weakly closed either.

3.44 REMARK (Open and closed sets in the weak topology). The weak topology has fewer open sets than the norm topology. Two consequences:

- **Harder to be open.** Norm-open balls are generally *not* weakly open in infinite dimensions (they would need to contain an unbounded weak neighborhood, which they cannot).
- **Fewer closed sets, not more.** A common misconception is that fewer open sets means more closed sets. In fact, closed sets are complements of open sets, so **fewer open sets means fewer closed sets too**. A set that is "not open" in the weak topology does not become closed. Most sets are neither.
  However, **convex** norm-closed sets remain weakly closed — this is Theorem 3.57. The closed unit ball is weakly closed because it is convex, not because of a general principle about coarser topologies. Non-convex norm-closed sets can fail to be weakly closed. For example, the set of standard basis vectors $\{e_n\}$ in ℓ^2 is norm-closed (all pairwise distances are $\sqrt{2}$), but $e_n \rightharpoonup 0$ weakly, so 0 is in its weak closure.

3.45 REMARK (The unit ball as an intersection of slabs). The slab picture gives a clean characterization of the unit ball. Each $f \in X^*$ with $\|f\| \le 1$ defines a closed slab $\{x : |f(x)| \le 1\}$, and the unit ball fits inside every such slab. The Corollary 3.19 gives the converse:

$$B_X = \bigcap_{\|f\| \le 1} \{x : |f(x)| \le 1\}. \tag{106}$$

The unit ball is exactly the intersection of all dual slabs. This means B_X is **weakly closed** (it is an intersection of weakly closed sets, since each f is weakly continuous).

3.8.3 Strong and weak convergence

With the weak topology defined, we can now describe the corresponding notions of convergence. We begin with norm convergence for comparison, then turn to weak convergence.

3.46 DEFINITION (Strong convergence). Let X be a normed space and let $(x_n)_{n\ge1}$ be a sequence in X. We say $x_n \to x$ **strongly** (or **in norm**) if

$$\|x_n - x\| \to 0 \quad \text{as } n \to \infty. \tag{107}$$

This is the familiar notion: the points x_n physically move toward x, and the distance $\|x_n - x\|$ shrinks to zero. Since $|f(x_n) - f(x)| \le \|f\| \cdot \|x_n - x\|$, strong convergence forces all instrument readings to converge uniformly over $\|f\| \le 1$. But it is defined by the norm directly, not by the instruments.

3.47 DEFINITION (Weak convergence). Let X be a normed space and let $(x_n)_{n\geq 1}$ be a sequence in X. We say $x_n \rightharpoonup x$ **weakly** if

$$f(x_n) \to f(x) \quad \text{for every } f \in \mathrm{X}^*. \tag{108}$$

Each instrument f foliates X into level sets $\{x : f(x) = c\}$, the "isotherms" for that measurement. Weak convergence means: in every foliation, the readings $f(x_n)$ settle down to $f(x)$. The objects x_n need not become close to x; they can keep bouncing around, as long as every instrument eventually reads the same value as it does on x.

3.48 PROPOSITION (Strong convergence implies weak convergence). *Let* X *be a normed space. If* $x_n \to x$ *strongly, then* $x_n \rightharpoonup x$ *weakly.*

Proof. Let $f \in \mathrm{X}^*$. Since f is a bounded linear functional,

$$|f(x_n) - f(x)| = |f(x_n - x)| \leq \|f\| \cdot \|x_n - x\| \to 0. \tag{109}$$

Since f was arbitrary, $x_n \rightharpoonup x$.

□

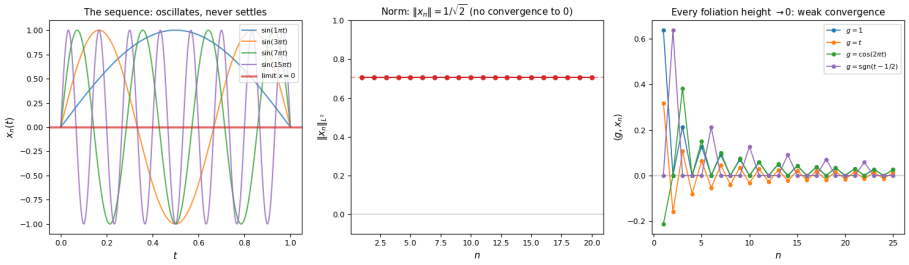
The converse is false in infinite dimensions:

3.49 EXAMPLE (Weak but not strong convergence in). L^2

In $\mathrm{L}^2([0,1])$, the sequence $x_n = \sin(n\pi t)$ converges weakly to 0, but **not** strongly. Each x_n has norm $\|x_n\|_{\mathrm{L}^2} = 1/\sqrt{2}$, so the sequence stays on a sphere. But for any $g \in \mathrm{L}^2$, the Riemann-Lebesgue lemma gives:

$$\int_0^1 g(t)\sin(n\pi t)\,dt \to 0 \quad \text{as } n \to \infty \tag{110}$$

so $f_g(x_n) \to 0$ for every functional $f_g \in (\mathrm{L}^2)^*$, i.e., $x_n \rightharpoonup 0$.



Left: the functions $\sin(n\pi t)$ *oscillate faster and faster, never settling down pointwise. Center: the* L^2 *norm stays at* $1/\sqrt{2}$ *for every* n*, so the sequence does not converge strongly. Right: the inner product* $\langle g, x_n\rangle$ *decays to 0 for every test function* $g \in \mathrm{L}^2$*, no matter its shape. Every "foliation height" converges, even though the functions themselves keep bouncing around. This is weak convergence without strong convergence.*

3.50 REMARK (Pointwise vs. uniform convergence over the dual). Why does weak convergence not imply strong convergence? After all, weak convergence requires $\langle g, x_n\rangle \to 0$ for **all** $g \in \mathrm{L}^2$ – isn't "all" a strong demand?

The key is that weak convergence is **pointwise in** g: we fix g first, then send $n \to \infty$. For any fixed g, the Fourier coefficients $\langle g, \sin(n\pi t)\rangle = \hat{g}(n) \to 0$ because the tail of a square-summable series vanishes. Each instrument eventually loses track of the oscillation.

But we are free to **change the instrument with** n. Choose $g_n = \sin(n\pi t)$, i.e., the instrument that is perfectly aligned with x_n. Then

$$\langle g_n, x_n\rangle = \|\sin(n\pi t)\|_{\mathrm{L}^2}^2 = \frac{1}{2} \tag{111}$$

for every n: this "adaptive" instrument always catches the oscillation. The supremum over the unit ball,

$$\sup_{\|g\|\le 1} |\langle g, x_n\rangle| = \|x_n\| = \frac{1}{\sqrt{2}}, \tag{112}$$

never decays. But this supremum **is** the norm of x_n, and driving it to zero would be exactly strong convergence.

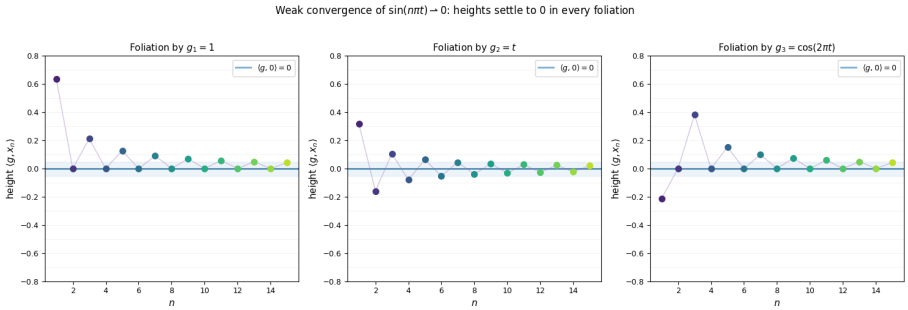
Why is this not a problem for weak convergence? Because the definition (Definition 3.47) quantifies as: *for every fixed* $g \in \mathrm{X}^*$, $\langle g, x_n\rangle \to 0$. The functional g is chosen **once and for all**, and then we ask whether the sequence of numbers $\langle g, x_n\rangle$ converges. An adaptive choice g_n that changes with n does not define a single sequence of real numbers: it defines a different measurement at each step. This is not what any individual instrument reads; it is a meta-observation assembled by switching instruments. No single linear measuring instrument in X^* witnesses the non-convergence.

So the gap between weak and strong convergence is precisely the gap between **pointwise** and **uniform** convergence over X^*: each fixed linear instrument eventually stops detecting the oscillation, but the worst-case instrument shifts with n to stay aligned with x_n.

The identity $\sup_{\|g\|\le 1} |\langle g, x_n\rangle| = \|x_n\|$ is a consequence of the Hahn-Banach theorem: for any $x \in \mathrm{X}$, there exists a **norming functional** $g_x \in \mathrm{X}^*$ with $\|g_x\| = 1$ and $g_x(x) = \|x\|$. So the supremum over the unit ball of X^* is always attained (or approached, in the non-reflexive case), and equals the norm. In other words, the norm of x is exactly the largest reading that any unit-norm linear instrument can produce on x.

This is also where the Banach-Steinhaus theorem (Theorem 2.24) enters: weak convergence $x_n \rightharpoonup x$ implies that $\langle g, x_n\rangle$ is bounded for each fixed g. Each x_n acts as a bounded linear functional on X^* via evaluation, $x_n(g) = g(x_n)$, and the family $\{x_n\}$ is pointwise bounded on X^*. The uniform boundedness principle then gives $\sup_n \|x_n\| < \infty$. So weak convergence automatically implies norm-boundedness, but not norm-convergence.

THE FOLIATION PICTURE OF WEAK CONVERGENCE Each functional $g \in \mathrm{L}^2$ defines a foliation of L^2 into parallel hyperplanes $\{x : \langle g, x\rangle = c\}$. Weak convergence $x_n \rightharpoonup 0$ means that in *every* foliation, the heights $\langle g, x_n\rangle$ converge to 0. The points x_n jump between different hyperplanes in each foliation, but eventually the heights settle near the origin's level set.



Each panel is a different foliation of L^2*, defined by a functional* g*. The vertical axis is the "height"* $\langle g, x_n \rangle$ *that the foliation assigns to* $x_n = \sin(n\pi t)$*. The points jump around between level sets, but in every foliation the heights converge to 0 (the blue line). The sequence never converges in norm, yet every foliation eventually reads near-zero heights. This is weak convergence.*

Strong vs. Weak

Strong convergence means the objects cluster in space. Weak convergence means every instrument reading converges, even though the objects may keep moving. The gap is real: $\sin(n\pi x) \rightharpoonup 0$ but $\|\sin(n\pi x)\| = 1/\sqrt{2} \not\to 0$.

In finite dimensions, this gap **disappears**: weak and strong convergence are equivalent (finitely many instruments suffice to control the norm). The gap is an essentially infinite-dimensional phenomenon.

Weak limits can lose mass Strong convergence preserves the norm: $\|x_n - x\| \to 0$ implies $\|x_n\| \to \|x\|$. Weak convergence does not. The norm can *drop* in the limit, but it cannot increase.

3.51 proposition (Weak lower semicontinuity of the norm). *Let* X *be a normed space. If* $x_n \rightharpoonup x$ *weakly, then*

$$\|x\| \le \liminf_{n\to\infty} \|x_n\|. \tag{113}$$

Proof. By the Corollary 3.19, for any $f \in \mathrm{X}^*$ with $\|f\| \le 1$, we have $|f(x_n)| \le \|f\| \cdot \|x_n\| \le \|x_n\|$. Since $x_n \rightharpoonup x$, the left side converges: $|f(x_n)| \to |f(x)|$. The right side $\|x_n\|$ need not converge (the norms may oscillate), but lim inf always exists. If $a_n \le b_n$ and $\lim a_n$ exists, then $\lim a_n \le \liminf b_n$. Therefore

$$|f(x)| = \lim_{n\to\infty} |f(x_n)| \le \liminf_{n\to\infty} \|x_n\|. \tag{114}$$

Taking the supremum over all f with $\|f\| \le 1$:

$$\|x\| = \sup_{\|f\|\le 1} |f(x)| \le \liminf_{n\to\infty} \|x_n\|. \tag{115}$$

□

This is the price of weak convergence: mass can escape to infinity. In the example $\sin(n\pi t) \rightharpoonup 0$, the norm stays at $1/\sqrt{2}$ while the weak limit has norm 0. The inequality is sharp.

Basic properties of weak convergence

3.52 PROPOSITION (Weak limits are unique). *Let* X *be a normed space. If* $x_n \rightharpoonup x$ *and* $x_n \rightharpoonup y$, *then* $x = y$.

Proof. For every $f \in X^*$, we have $f(x) = \lim f(x_n) = f(y)$, so $f(x - y) = 0$ for all $f \in X^*$. By the Corollary 3.19, $\|x - y\| = \sup_{\|f\| \leq 1} |f(x - y)| = 0$, so $x = y$.

□

3.53 PROPOSITION (Weakly convergent sequences are bounded). *Let* X *be a Banach space. If* $x_n \rightharpoonup x$, *then* $\sup_n \|x_n\| < \infty$.

Proof. Each x_n defines an element $J[x_n] \in X^{**}$ by $J[x_n](f) = f(x_n)$. Since $x_n \rightharpoonup x$, for each fixed $f \in X^*$ the sequence $J[x_n](f) = f(x_n)$ converges and is therefore bounded. So the family $\{J[x_n]\}_{n \geq 1} \subseteq X^{**}$ is pointwise bounded on X^*. Since X^* is a Banach space, the Banach-Steinhaus theorem (Theorem 2.24) gives $\sup_n \|J[x_n]\|_{X^{**}} < \infty$. By Lemma 3.26, $\|J[x_n]\|_{X^{**}} = \|x_n\|$, so $\sup_n \|x_n\| < \infty$.

□

3.54 PROPOSITION (Compact operators turn weak convergence into strong convergence). *Let* X, Y *be Banach spaces and* $A : X \to Y$ *a compact linear operator. If* $x_n \rightharpoonup x$ *in* X, *then* $Ax_n \to Ax$ *strongly in* Y.

Proof. Since $x_n \rightharpoonup x$, Proposition 3.53 gives $\sup_n \|x_n\| \leq C$. The sequence $(x_n - x)$ converges weakly to 0 and is bounded, so $(A(x_n - x))$ lies in the image of a bounded set under a compact operator, hence its closure is compact in Y.

It suffices to show every subsequence of (Ax_n) has a further subsequence converging to Ax. Let (x_{n_k}) be any subsequence. Since A is compact and (x_{n_k}) is bounded, there exists a further subsequence $(x_{n_{k_j}})$ such that $Ax_{n_{k_j}} \to z$ for some $z \in Y$. We identify $z = Ax$: for any $g \in Y^*$, the functional $f = g \circ A \in X^*$, so

$$g(Ax_{n_{k_j}}) = f(x_{n_{k_j}}) \to f(x) = g(Ax). \tag{116}$$

Since $Ax_{n_{k_j}} \to z$ strongly, also $g(Ax_{n_{k_j}}) \to g(z)$, giving $g(z) = g(Ax)$ for all $g \in Y^*$. By Hahn-Banach, $z = Ax$.

Since every subsequence of (Ax_n) has a further subsequence converging to Ax, the full sequence converges: $Ax_n \to Ax$.

□

3.8.4 Weak compactness in reflexive spaces

The whole point of weakening the topology was to gain compactness. With fewer open sets there are fewer open covers, so it becomes easier for a set to be compact. The following theorem makes this precise for reflexive spaces.

3.55 THEOREM (Weak compactness of the unit ball). *Let* X *be a* ***reflexive*** *Banach space. Then the closed unit ball* B_X *is compact in the weak topology. In particular, every bounded sequence in* X *has a weakly convergent subsequence.*

The proof uses the Banach-Alaoglu theorem (Theorem 3.67) and the canonical embedding; see Corollary 3.69.

This theorem is the reason reflexivity matters in applications. In reflexive spaces like L^p $(1 < p < \infty)$ and Sobolev spaces $W^{k,p}$ $(1 < p < \infty)$, every bounded sequence has a weakly convergent subsequence. This is the compactness step in the **direct method** of the calculus of variations: minimize an energy functional over a bounded set, extract a weakly convergent subsequence, and pass to the limit.

WHAT GOES WRONG WITHOUT REFLEXIVITY Why does the theorem require reflexivity? If weak compactness held for every Banach space, then every bounded sequence would have a weakly convergent subsequence. Consider $X = c_0$ with the duality chain $c_0^* = \ell^1$, $c_0^{**} = \ell^\infty$ (recall Example 3.31). The space c_0 is not reflexive: the canonical embedding $J : c_0 \hookrightarrow \ell^\infty$ misses elements like $(1, 1, 1, \ldots) \in \ell^\infty \setminus J(c_0)$.

Consider the bounded sequence $x_n = (\underbrace{1, 1, \ldots, 1}_{n}, 0, 0, \ldots) \in c_0$, with $\|x_n\|_\infty = 1$. Can we extract a weakly convergent subsequence? In fact the full sequence already has the property that the readings converge: for any $f = (f_k) \in \ell^1 = c_0^*$,

$$f(x_n) = \sum_{k=1}^{n} f_k \tag{117}$$

which is a partial sum of the absolutely convergent series $\sum f_k$ (since $f \in \ell^1$), so $f(x_n) \to \sum_{k=1}^{\infty} f_k$.

What we want: a subsequence (x_{n_j}) and some $x \in c_0$ such that $f(x_{n_j}) \to f(x)$ for every $f \in \ell^1$.

What this requires: we already know $f(x_n) \to \sum_{k=1}^{\infty} f_k$ for every f. Since a subsequence of a convergent sequence converges to the same limit, $f(x_{n_j}) \to \sum_{k=1}^{\infty} f_k$ as well. So the candidate x must satisfy $f(x) = \sum_{k=1}^{\infty} f_k$ for every $f \in \ell^1$.

Why the only candidate is not in c_0: choose $f = e_k = (0, \ldots, 0, 1, 0, \ldots) \in \ell^1$, which reads the k-th coordinate. On one hand, $e_k(x) = x_k$. On the other hand, $\sum_j (e_k)_j = 1$. So $x_k = 1$ for every k, forcing $x = (1, 1, 1, \ldots)$. But this constant sequence does not converge to zero, so $x \notin c_0$. No subsequence of (x_n) converges weakly in c_0.

Where does this candidate live? The canonical embedding $J : c_0 \hookrightarrow \ell^\infty = c_0^{**}$ identifies each element of c_0 with an evaluation functional on ℓ^1. The candidate $x = (1, 1, 1, \ldots)$ is a bounded sequence, so it defines a valid element of ℓ^∞. It acts on ℓ^1 by

$$J(x)(f) = \sum_{k=1}^{\infty} f_k \quad \text{for } f = (f_k) \in \ell^1. \tag{118}$$

This is a perfectly good element of the bidual $c_0^{**} = \ell^\infty$, but it is not in the range of J. Explicitly:

$$\mathrm{J}(c_0) = \{(y_k) \in \ell^\infty : y_k \to 0\}, \qquad (1,1,1,\ldots) \in \ell^\infty \setminus \mathrm{J}(c_0), \tag{119}$$

since the constant sequence does not converge to zero. So J is not surjective: $\mathrm{J}(c_0) \subsetneq \ell^\infty = c_0^{**}$. The sequence (x_n) has "escaped" into the bidual — the readings converge, but the limit lives in $c_0^{**} \setminus \mathrm{J}(c_0)$.

Reflexivity means $\mathrm{J}(\mathrm{X}) = \mathrm{X}^{**}$: there is no gap, so there is no room to escape. In a reflexive space, every candidate limit that is consistent with the readings already lives in X.

3.8.5 Separation and convexity in the weak topology

Optional Extension

This section explores the topological separation properties of the weak topology and their connection to Mazur's theorem. It is not needed for the main development but clarifies *why* convexity plays such a distinguished role.

Separation axioms: from T_2 to T_4 The separation axioms form a hierarchy of increasing strength:

$$\mathrm{T}_4 \implies \mathrm{T}_{3\frac{1}{2}} \implies \mathrm{T}_3 \implies \mathrm{T}_2 \implies \mathrm{T}_1. \tag{120}$$

- T_2 **(Hausdorff):** Any two distinct *points* can be separated by disjoint open sets.
- T_3 **(regular Hausdorff):** Any *point* and *closed set* not containing it can be separated by disjoint open sets.
- $\mathrm{T}_{3\frac{1}{2}}$ **(completely regular / Tychonoff):** Any *point* and *closed set* not containing it can be separated by a *continuous function* $f : \mathrm{X} \to [0,1]$.
- T_4 **(normal Hausdorff):** Any two *disjoint closed sets* can be separated by disjoint open sets.

Each level upgrades what can be separated: T_2 separates points from points, T_3 separates points from closed sets by open sets, $\mathrm{T}_{3\frac{1}{2}}$ strengthens this to separation by continuous functions, and T_4 separates closed sets from closed sets.

3.56 proposition (Metric spaces are). T_4

Every metric space (X, d) *is* T_4.

Proof. Given disjoint closed sets $\mathrm{A}, \mathrm{B} \subset \mathrm{X}$, the function

$$f(x) = \frac{d(x,\mathrm{A})}{d(x,\mathrm{A}) + d(x,\mathrm{B})} \tag{121}$$

is continuous with $f = 0$ on A and $f = 1$ on B (this works because $d(x,\mathrm{A}) + d(x,\mathrm{B}) > 0$ when $\mathrm{A} \cap \mathrm{B} = \emptyset$ and both are closed), so $\mathrm{U} = f^{-1}([0, \frac{1}{2}))$ and $\mathrm{V} = f^{-1}((\frac{1}{2}, 1])$ are disjoint open sets separating A and B.

□

Since every normed space is a metric space, every Banach space is T_4 in its norm topology. When we pass to the **weak topology**, we lose the metric but the topology is still Hausdorff (T_2), since X^* separates points.

MAZUR'S THEOREM: CONVEXITY RESTORES SEPARATION The natural question is: can we separate points from *closed sets*, not just from other points? The weak topology is defined so that every $f \in X^*$ is continuous, and the Hahn-Banach theorem produces such an f that strictly separates a point from a closed **convex** set. So for convex sets, the functionals in X^* play exactly the role of the separating continuous functions required by the $T_{3\frac{1}{2}}$ axiom. Mazur's theorem makes this precise: convexity is exactly the condition under which norm-closure and weak-closure agree.

3.57 THEOREM (Mazur's theorem). *Let* X *be a normed space and* $C \subseteq X$ *a convex set. Then* C *is norm-closed if and only if it is weakly closed.*

Proof. Since the weak topology is coarser than the norm topology, every weakly closed set is norm-closed. The content is the converse: a norm-closed convex set is weakly closed.

It suffices to show that if $x_0 \notin C$ and C is norm-closed and convex, then x_0 is not in the weak closure of C, i.e., there is a weakly open set containing x_0 that misses C.

Since C is norm-closed and $x_0 \notin C$, there exists $\delta > 0$ with $B_\delta(x_0) \cap C = \emptyset$. By the geometric form of the Hahn-Banach theorem (strict separation of a point from a closed convex set), there exists $f \in X^*$ and $\alpha \in \mathbb{R}$ such that

$$f(x_0) > \alpha \geq f(c) \quad \text{for all } c \in C. \tag{122}$$

The set $\{x \in X : f(x) > \alpha\}$ is a weakly open slab containing x_0 and disjoint from C. Therefore x_0 is not in the weak closure of C.

□

Mazur's theorem says that for convex sets, you cannot tell the difference between norm-closure and weak-closure. This is why the closed unit ball, closed convex hulls, and closed subspaces are all weakly closed. Non-convex sets do not enjoy this protection: the set $\{e_n\}$ in ℓ^2 is norm-closed but not weakly closed.

The connection to separation is now clear: the weak topology lacks T_4 (separation of arbitrary closed sets), but Hahn-Banach gives separation of points from **closed convex** sets. Mazur's theorem is the payoff — convexity is the precise condition under which the weak topology's separation power suffices.

3.9 *The Weak* Topology and Banach-Alaoglu*

Big Idea

The dual space X^* is itself a Banach space, so it has a weak topology. But that weak topology requires the bidual X^{**}, which can be enormous. The **weak* topology** sidesteps the bidual entirely by testing only against elements of X. This gives a coarser topology with fewer open sets, making compactness easier, and leads to the **Banach-Alaoglu theorem**.

We now want a weak-type topology on the dual space X^*, for the same reason we wanted one on X: to gain compactness. Since X^* is a Banach space in its own right, the most natural first attempt is to apply the same construction as before. Weak convergence of a sequence (f_n) in X^* would then mean

$$\phi(f_n) \to \phi(f) \quad \text{for every } \phi \in (X^*)^* = X^{**}. \tag{123}$$

This requires testing against every element of the **bidual** X^{**}.

3.58 EXAMPLE (The bidual can be much larger than). X

Recall the duality chain from Example 3.31: for $X = c_0$, $c_0^* = \ell^1$ and $c_0^{**} = \ell^\infty$. To check weak convergence of a sequence (f_n) in ℓ^1, we would need to test against every element of ℓ^∞.

Each $x \in c_0$ gives an element of ℓ^∞ via the canonical embedding J, and testing f_n against $J(x)$ just means evaluating $f_n(x) = \sum_k f_k x_k$, which is concrete: we point the instrument f_n at the object x and read off the result.

But ℓ^∞ also contains elements not in $J(c_0)$. Consider $\phi = (1, 1, 1, \ldots) \in \ell^\infty$. Its action on $f \in \ell^1$ is $\phi(f) = \sum_k f_k$. This is a perfectly valid bounded linear functional on ℓ^1, but it does not come from the canonical embedding: there is no $x \in c_0$ with $J(x) = \phi$, since $(1, 1, 1, \ldots)$ does not converge to zero. Requiring convergence against such "phantom" test functionals is a strictly stronger demand, and it makes compactness harder to achieve.

Now recall the canonical embedding $J : X \hookrightarrow X^{**}$, which identifies each $x \in X$ with the evaluation functional $J(x)(f) = f(x)$. Every element of X already lives inside X^{**}. So instead of testing against all of X^{**}, we could test against just J(X): ask only that

$$f_n(x) \to f(x) \quad \text{for every } x \in X. \tag{124}$$

This is weaker (we test against fewer functionals), but it has two advantages. First, it uses only elements of X, which we already understand. Second, because the topology is coarser (fewer open sets, fewer open covers), compactness becomes easier.

3.9.1 The weak* topology

3.59 DEFINITION (Weak* topology). The **weak* topology** on X^*, denoted $\sigma(X^*, X)$, is the coarsest topology making every evaluation map $\hat{x} : X^* \to \mathbb{R}$, $\hat{x}(f) = f(x)$, continuous. A basis of open neighborhoods of

$f \in X^*$ is:

$$V(f; x_1, \dots, x_n, \varepsilon) = \{g \in X^* : |g(x_i) - f(x_i)| < \varepsilon \text{ for } i = 1, \dots, n\} \tag{125}$$

where $x_1, \dots, x_n \in X$ and $\varepsilon > 0$.

Continuity of $\hat{x}$ requires preimages of open sets to be open. Every open set in $\mathbb{R}$ is a union of open intervals, and $\hat{x}^{-1}((a, b)) = \{g \in X^* : a < g(x) < b\}$ is a slab in X^*. So any topology making all evaluation maps continuous must contain all such slabs. The coarsest such topology is the one generated by the slabs alone. A weak* slab constrains the values of g on finitely many test objects.

3.60 EXAMPLE (The weak* topology on). $\ell^1 = (c_0)^*$

Take $X = c_0$ (sequences converging to 0 with sup norm). Its dual is $X^* = \ell^1$ with $f(x) = \sum_k f_k x_k$ for $f = (f_k) \in \ell^1$, $x = (x_k) \in c_0$.

The **norm ball** in ℓ^1 is $B_1 = \{f \in \ell^1 : \sum_k |f_k| \le 1\}$, the familiar cross-polytope (diamond) in finite-dimensional slices.

Now take the test object $x = e_1 = (1, 0, 0, \dots) \in c_0$. The evaluation $\hat{e}_1(f) = f_1$ defines a weak* slab

$$V_1 = \{f \in \ell^1 : |f_1| < 1\}. \tag{126}$$

This slab constrains only the first component of f. A sequence like $f^{(n)} = \frac{1}{2}e_1 + e_n$ has $|f_1^{(n)}| = \frac{1}{2} < 1$ (inside the slab) but $\|f^{(n)}\|_1 = \frac{3}{2}$ (outside the norm ball).

More generally, the weak* neighborhood $V(0; e_1, \dots, e_n, \varepsilon) = \{f \in \ell^1 : |f_k| < \varepsilon \text{ for } k = 1, \dots, n\}$ constrains the first n components but places no restriction on f_k for $k > n$. In infinite dimensions, every weak* neighborhood of 0 is unbounded in ℓ^1.

3.9.2 Where weak* sits: the hierarchy of topologies

The weak* topology is one of three natural topologies on X^*. The canonical embedding $J : X \hookrightarrow X^{**}$ mediates between them: each $x \in X$ defines an evaluation functional $J[x] \in X^{**}$ by $J[x](f) = f(x)$.

- The **weak* topology on** X^* is $\sigma(X^*, X)$: slabs come from $J[x]$ for $x \in X$.
- The **weak topology on** X^* is $\sigma(X^*, X^{**})$: slabs come from all $\phi \in X^{**}$, including those not in $J(X)$.
- The **norm topology on** X^* uses norm balls.

Since $J(X) \subseteq X^{**}$, every weak* open set is also weakly open, and every weakly open set is norm-open. More test functionals means more slabs, which means more open sets:

$$\text{weak* topology} \subseteq \text{weak topology on } X^* \subseteq \text{norm topology on } X^*. \tag{127}$$

Each inclusion means "coarser than or equal to." The first inclusion is strict when X is not reflexive (there exist elements of $X^{**} \setminus J(X)$ that generate extra slabs). When X is reflexive, $J(X) = X^{**}$ and the weak* and weak topologies on X^* coincide.

3.61 EXAMPLE. $c_0 \hookrightarrow \ell^1 \hookrightarrow \ell^\infty$

Take $X = c_0$. Then $X^* = \ell^1$ and $X^{**} = \ell^\infty$. The canonical embedding $J : c_0 \hookrightarrow \ell^\infty$ sends $x = (x_k) \in c_0$ to the same sequence viewed in ℓ^∞. Since $c_0 \subsetneq \ell^\infty$ (e.g., the constant sequence $(1, 1, 1, \ldots) \in \ell^\infty \setminus c_0$), the space c_0 is not reflexive.

The weak* topology on ℓ^1 uses slabs from c_0: neighborhoods constrain $f(x)$ for $x \in c_0$. The weak topology on ℓ^1 uses slabs from ℓ^∞: neighborhoods can additionally constrain f against sequences that do not converge to 0. The weak topology is strictly finer, with more open sets and therefore fewer compact sets.

3.62 REMARK (The compactness payoff). The coarser the topology, the easier it is for sets to be compact. This is why compactness improves as we move left in the chain:

- **Norm topology**: the closed unit ball of X^* is compact only if X^* is finite-dimensional.
- **Weak topology on** X^*: the closed unit ball is compact if and only if X^* is reflexive.
- **Weak* topology on** X^*: the closed unit ball is **always** compact (Banach-Alaoglu).

The weak* topology is the coarsest of the three, so it is the easiest setting for compactness. This is exactly why Banach-Alaoglu works without any reflexivity assumption.

3.9.3 Weak* convergence

3.63 DEFINITION (Weak* convergence). Let X be a normed space and let $(f_n)_{n\geq 1}$ be a sequence in X^*. We say $f_n \xrightarrow{w^*} f$ (**weak* convergence**) if

$$f_n(x) \to f(x) \quad \text{for every } x \in X. \tag{128}$$

Now the picture flips: each f_n *is* an instrument, and the sequence is a sequence of instruments, not objects. Think of replacing your entire measurement apparatus, swapping one thermometer for another, one scale for another. Fix any object x and read off $f_1(x), f_2(x), f_3(x), \ldots$. Weak* convergence means these readings stabilize to $f(x)$ for every fixed object. Geometrically, each f_n defines a different foliation (different isotherms), and these foliations rearrange from step to step, but at every fixed point the height reading converges.

3.64 PROPOSITION (Strong convergence in). X^* ***implies weak* convergence***

Let X *be a normed space. If* $f_n \to f$ *in the norm of* X^*, *then* $f_n \xrightarrow{w^*} f$.

Proof. For any fixed $x \in X$,

$$|f_n(x) - f(x)| = |(f_n - f)(x)| \leq \|f_n - f\| \cdot \|x\| \to 0. \tag{129}$$

Since x was arbitrary, $f_n \xrightarrow{w^*} f$.

□

3.65 PROPOSITION (Weak convergence in). X^* ***implies weak* convergence***

Let X *be a normed space. If* $f_n \rightharpoonup f$ *weakly in* X^*, *then* $f_n \xrightarrow{w^*} f$.

Proof. Weak convergence in X^* means $\phi(f_n) \to \phi(f)$ for every $\phi \in X^{**}$. For any $x \in X$, the canonical image $J(x) \in X^{**}$ acts by $J(x)(g) = g(x)$. In particular,

$$f_n(x) = J(x)(f_n) \to J(x)(f) = f(x). \tag{130}$$

Since $J(X) \subseteq X^{**}$, weak convergence tests against a larger set of functionals than weak* convergence, so the former implies the latter.

□

In summary:

$$\text{strong in } X^* \implies \text{weak in } X^* \implies \text{weak* in } X^*, \tag{131}$$

and neither arrow reverses in general. When X is reflexive (J is surjective), the last two notions coincide.

3.66 EXAMPLE (Weak* convergent but not weakly convergent in). ℓ^1

Take $X = c_0$, so $X^* = \ell^1$ and $X^{**} = \ell^\infty$. The standard basis vectors $e_n \in \ell^1$ converge weak* to 0: for any $x = (x_k) \in c_0$,

$$e_n(x) = x_n \to 0 \tag{132}$$

since $x \in c_0$ means $x_n \to 0$. But e_n does **not** converge weakly to 0 in ℓ^1. The element $\phi = (1, 1, 1, \ldots) \in \ell^\infty = (\ell^1)^*$ satisfies

$$\phi(e_n) = 1 \quad \text{for all } n, \tag{133}$$

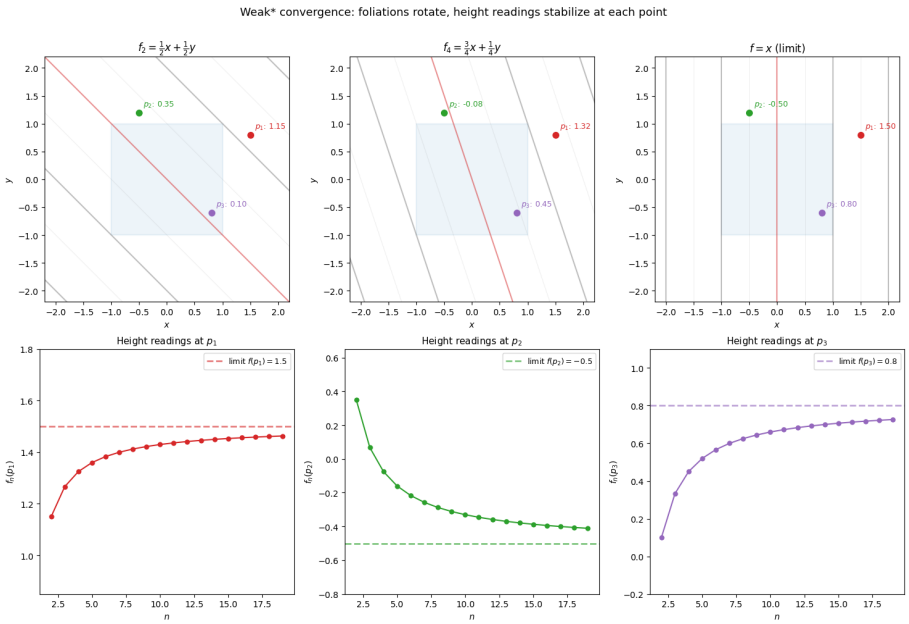
so $\phi(e_n) \not\to 0$. This ϕ is precisely the kind of "phantom" test functional in $X^{**} \setminus J(X)$ from Example 3.58: it does not correspond to any object in c_0, but it detects the non-convergence.

VISUALIZING WEAK* CONVERGENCE IN $\mathbb{R}^2$ Consider $f_n(x, y) = (1 - 1/n)\,x + (1/n)\,y$ on $(\mathbb{R}^2, \|\cdot\|_\infty)$. Each f_n has kernel line $(1 - 1/n)x + (1/n)y = 0$, which slowly rotates toward the y-axis as $n \to \infty$. For any fixed point (a, b):

$$f_n(a, b) = \left(1 - \frac{1}{n}\right) a + \frac{1}{n} b \to a = f(a, b) \tag{134}$$

where $f(x, y) = x$. The foliations converge pointwise: at each location, the height readings stabilize, even though the kernel lines are visibly rotating from picture to picture.

As with weak convergence, the $\mathbb{R}^2$ picture is a visual scaffold: in finite dimensions weak* = weak = strong, so the foliations converge uniformly, not just pointwise. The genuine weak* phenomenon, where pointwise convergence of height readings does not imply uniform convergence, requires infinite dimensions. The top row below shows the geometry (rotating level sets), but the real content is in the bottom row: height readings at fixed test points stabilizing one by one.



Top row: the level sets of f_n rotate as n increases. For $n = 2$ the level sets are diagonal; by $n = 4$ they are nearly vertical; at the limit $f(x, y) = x$ they are exactly vertical. Bottom row: for each fixed test point, the height reading $f_n(p)$ converges to the limiting value $f(p) = p_x$. This is weak convergence: the foliations rearrange, but at every fixed point the readings stabilize.*

3.9.4 The Banach-Alaoglu theorem

We claimed that the weak* topology makes compactness easier. The following theorem delivers on this promise: the closed unit ball of X^* is always weak* compact, for any normed space X.

The intuition is simple. Each $f \in B_{X^*}$ is determined by its readings on all objects, and since $\|f\| \leq 1$, each reading satisfies $|f(x)| \leq \|x\|$, so $f(x) \in [-\|x\|, \|x\|]$. An instrument is a choice of one bounded number per object. Given a sequence of instruments, focus on a single object x_1: the readings $f_n(x_1)$ are bounded, so Bolzano-Weierstrass gives a convergent subsequence. Extract a further subsequence to make the readings converge at x_2, then x_3, and so on.

This is the same diagonal argument behind Arzelà-Ascoli and compact operators: whenever there are **countably many** coordinates to control, Bolzano-Weierstrass plus diagonalization does the job. If X is separable, a countable dense subset provides the coordinates, and the argument goes through.

We give the full proof for separable spaces, then state the general result.

3.67 THEOREM (Banach-Alaoglu (separable case)). *Let X be a separable Banach space and $(f_n) \subset X^*$ a bounded sequence with $\|f_n\| \leq M$. Then (f_n) has a weak* convergent subsequence.*

Proof. **Step 1: set up the countable dense subset.** Since X is separable, there exists a countable dense subset $\{x_1, x_2, x_3, \ldots\} \subset X$.

Step 2: successive extraction. Evaluate the sequence (f_n) at x_1. Since $|f_n(x_1)| \leq M\|x_1\|$ for all n, the sequence $(f_n(x_1))$ is bounded in $\mathbb{R}$. By Bolzano-Weierstrass, extract a convergent subsequence:

$$(f_n) \supset (f_{n_1^{(j)}})_{j \geq 1} \quad \text{such that } f_{n_1^{(j)}}(x_1) \text{ converges.} \tag{135}$$

Now evaluate this subsequence at x_2. The numbers $(f_{n_1^{(j)}}(x_2))$ are again bounded, so extract a further subsequence:

$$(f_{n_1^{(j)}}) \supset (f_{n_2^{(j)}})_{j\geq 1} \quad \text{such that } f_{n_2^{(j)}}(x_1) \text{ and } f_{n_2^{(j)}}(x_2) \text{ both converge.} \tag{136}$$

The second subsequence still converges at x_1 because it is a subsequence of something that already converged there. Continue: at stage k, extract a subsequence $(f_{n_k^{(j)}})_{j\geq 1}$ that converges at $x_1, \dots, x_k$. This gives nested subsequences:

$$(f_n) \supset \underbrace{(f_{n_1^{(j)}})}_{\text{conv at } x_1} \supset \underbrace{(f_{n_2^{(j)}})}_{\text{conv at } x_1, x_2} \supset \cdots \supset \underbrace{(f_{n_k^{(j)}})}_{\text{conv at } x_1, \dots, x_k} \supset \cdots \tag{137}$$

Step 3: the diagonal trick. Define $g_j := f_{n_j^{(j)}}$, the j-th element of the j-th subsequence. For any fixed k and $j \geq k$, the element $g_j = f_{n_j^{(j)}}$ belongs to the j-th subsequence, which is a sub-subsequence of the k-th subsequence. So the tail $g_k, g_{k+1}, g_{k+2}, \dots$ is a subsequence of the k-th extracted sequence, which converges at x_k. Since finitely many initial terms do not affect convergence, $(g_j(x_k))_{j\geq 1}$ converges for every k. Call the limit L_k.

Step 4: extend from the dense subset to all of X**.** Take arbitrary $x \in X$. Given $\varepsilon > 0$, pick x_k from the dense subset with $\|x - x_k\| < \varepsilon$. Then:

$$|g_j(x) - g_l(x)| \leq |g_j(x) - g_j(x_k)| + |g_j(x_k) - g_l(x_k)| + |g_l(x_k) - g_l(x)|. \tag{138}$$

The first and third terms are bounded by $\|g_j\| \cdot \|x - x_k\| \leq M\varepsilon$ each. The middle term is less than ε for j, l large enough since $(g_j(x_k))$ converges. So $|g_j(x) - g_l(x)| \leq (2M + 1)\varepsilon$ for j, l sufficiently large. The sequence $(g_j(x))$ is Cauchy in $\mathbb{R}$, hence convergent. Define $g(x) := \lim_{j\to\infty} g_j(x)$.

Step 5: the limit is in X^***.** Linearity of g follows from linearity of each g_j and limits. Boundedness follows from $|g(x)| = \lim |g_j(x)| \leq M\|x\|$, so $\|g\| \leq M$ and $g \in X^*$. By construction $g_j(x) \to g(x)$ for all $x \in X$, which is exactly $g_j \xrightarrow{w^*} g$.

□

Leonidas Alaoglu extended this result to non-separable spaces in his 1938 PhD thesis at the University of Chicago, replacing the diagonal argument with Tychonoff's theorem.

3.68 REMARK (The general (non-separable) case). For arbitrary normed spaces, the theorem states that B_{X^*} is compact (not just sequentially compact) in the weak* topology. The proof embeds B_{X^*} into the product $P = \prod_{x\in X} [-\|x\|, \|x\|]$ via $\Phi(f) = (f(x))_{x\in X}$. By Tychonoff's theorem, P is compact in the product topology. The product topology is exactly the topology of pointwise convergence, which is the weak* topology. One checks that $\Phi(B_{X^*})$ is closed in P (limits of linear functions are linear, and the norm bound is preserved), so it is compact as a closed subset of a compact space.

3.69 COROLLARY (Weak sequential compactness in separable reflexive spaces). *Let* X *be a* ***separable reflexive*** *Banach space. Then every bounded sequence in* X *has a weakly convergent subsequence.*

Proof. Since X is reflexive, the canonical embedding $J : X \to X^{**}$ is surjective, so $J(B_X) = B_{X^{**}}$. Now apply Banach-Alaoglu to the space X^*: the unit ball $B_{(X^*)^*} = B_{X^{**}}$ is compact in the weak* topology of X^{**} (i.e., the topology of pointwise convergence on X^*). But under the identification $J : X \cong X^{**}$, the weak* topology on X^{**} corresponds exactly to the weak topology on X (both test against elements of X^*). Therefore B_X is weakly compact.

For sequences: since X is separable, X^* is separable in the weak* topology (a countable dense subset of X generates a countable family of weak* continuous functionals that separate points of X^*). The diagonal argument from the Banach-Alaoglu proof then applies directly: given (x_n) with $\|x_n\| \leq C$, pick a countable dense subset $\{f_1, f_2, \ldots\} \subset X^*$, diagonalize to extract a subsequence along which $f_k(x_{n_j})$ converges for every k, and extend to all of X^* by density and uniform boundedness.

□

This corollary is the reason reflexivity matters in applications. In a separable reflexive space (such as L^p or $W^{k,p}$ for $1 < p < \infty$), bounded sequences always have weakly convergent subsequences. Without reflexivity, Banach-Alaoglu still gives weak* compactness of B_{X^*}, but this is a statement about functionals on X, not about elements of X itself.

3.10 *Applications of Weak and Weak* Convergence*

Big Idea

The three notions of convergence (strong, weak, weak*) form a chain of implications. In probability, they correspond to convergence in L^1 norm, convergence in expectation, and convergence in distribution. The hierarchy of topologies explains when they coincide (reflexive spaces) and when they diverge. The **direct method** of the calculus of variations combines Banach-Alaoglu with weak lower semicontinuity to prove existence of minimizers.

3.10.1 Convergence in probability: all three in action

The three convergences appear naturally in probability, where they correspond to familiar notions.

Convergence in expectation as weak convergence A random variable X on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is an element of $L^1(\Omega, \mathbb{P})$ precisely when it has **finite expectation**: $\mathbb{E}[|X|] = \int_\Omega |X|\, d\mathbb{P} < \infty$. The L^1 norm *is* the expected absolute value: $\|X\|_{L^1} = \mathbb{E}[|X|]$.

3.70 example (Convergence in expectation in). L^1

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. The Banach space $L^1(\Omega, \mathbb{P})$ has dual $(L^1)^* = L^\infty(\Omega, \mathbb{P})$. Each bounded measurable function $h \in L^\infty$ defines a continuous linear functional on L^1 via

$$\varphi_h(X) = \int_\Omega X(\omega)\, h(\omega)\, d\mathbb{P}(\omega) = \mathbb{E}[Xh]. \tag{139}$$

Think of h as a **weight function** that selects which parts of X to measure. Typical choices:

- $h = \mathbf{1}_A$ (indicator of a measurable set): $\varphi_h(X) = \int_A X\, d\mathbb{P}$, the integral of X over A.
- $h = \operatorname{sgn}(X)$: $\varphi_h(X) = \mathbb{E}[|X|] = \|X\|_{L^1}$, recovering the norm.
- $h = e^{i\xi\cdot(\cdot)}$ (a character, if $\Omega \subseteq \mathbb{R}^d$): $\varphi_h(X) = \widehat{X}(\xi)$, a Fourier coefficient.

A sequence of random variables $X_n \in L^1$ converges **weakly** ($X_n \rightharpoonup X$) if and only if

$$\mathbb{E}[X_n h] \to \mathbb{E}[Xh] \quad \text{for every } h \in L^\infty. \tag{140}$$

In the foliation picture: each weight h defines a foliation of L^1, and weak convergence means the "weighted expectation" $\mathbb{E}[X_n h]$ stabilizes for every choice of weight. In particular, taking $h = \mathbf{1}_A$ for every measurable A requires $\int_A X_n\, d\mathbb{P} \to \int_A X\, d\mathbb{P}$: the integrals of X_n agree with those of X over every measurable set.

Strong convergence in L^1 is $\mathbb{E}[|X_n - X|] \to 0$, which implies weak convergence but not vice versa. A sequence can have all its weighted expectations converge without the random variables getting close in the L^1 norm.

MEASURES AND WEAK* CONVERGENCE The richest example of weak* convergence in practice is the convergence of measures. By the Riesz–Markov–Kakutani theorem, $C(K)^* \cong \mathcal{M}(K)$ (the space of finite signed Radon measures), and weak* convergence of measures means $\int g\, d\mu_n \to \int g\, d\mu$ for every continuous g. For probability measures, this is **convergence in distribution** — the convergence in the central limit theorem.

We develop this fully — including the total variation norm, Wasserstein distances, tightness, Prokhorov's theorem, and optimal transport — in the distributions chapter on measure norms. There, the duality framework from this chapter combines with the distributional framework to give a complete picture of what different norms on measures "see."

3.10.2 The direct method: a first look

The combination of Banach-Alaoglu and weak lower semicontinuity is the engine behind existence proofs in PDE and the calculus of variations. We give a brief preview here; we will develop this in detail in later sections.

3.71 DEFINITION (Weak lower semicontinuity). A functional $\mathcal{F}\colon X \to \mathbb{R} \cup \{+\infty\}$ is **weakly lower semicontinuous** if whenever $x_n \rightharpoonup x$ weakly in X,

$$\mathcal{F}(x) \le \liminf_{n\to\infty} \mathcal{F}(x_n). \tag{141}$$

The norm itself is the prototypical example: if $x_n \rightharpoonup x$, then $\|x\| \le \liminf \|x_n\|$ (this follows from $|f(x)| = \lim |f(x_n)| \le \|f\| \liminf \|x_n\|$ for any norming functional f). Many energy functionals in applications inherit this property from convexity.

The **direct method** proceeds in three steps:

1. **Boundedness.** Show that a minimizing sequence (u_n) for $\mathcal{F}$ is bounded: $\|u_n\| \le C$.

2. **Compactness.** Extract a weakly convergent subsequence $u_{n_k} \rightharpoonup u$. In a reflexive space this is Corollary 3.69; in a dual space $X = Y^*$, use Banach-Alaoglu (Theorem 3.67) for weak* compactness instead.

3. **Lower semicontinuity.** Conclude that

$$\mathcal{F}(u) \le \liminf_{k\to\infty} \mathcal{F}(u_{n_k}) = \inf \mathcal{F}, \tag{142}$$

so u is a minimizer.

Step 1 is problem-specific, typically a coercivity estimate. Step 2 is pure functional analysis, exactly what Banach-Alaoglu and Eberlein-Šmulian provide. Step 3 requires that $\mathcal{F}$ behaves well under weak limits, which is where convexity or compensated compactness arguments enter. We will return to this in detail when we study Sobolev spaces and variational problems.

THE DIRECT METHOD IN A REFLEXIVE SPACE

3.72 EXAMPLE (Minimizing the Dirichlet energy). Let $\Omega \subset \mathbb{R}^d$ be a bounded open set with smooth boundary, and fix a boundary datum $g \in \mathrm{H}^{1/2}(\partial\Omega)$. Consider the **Dirichlet energy**

$$\mathcal{E}(u) = \frac{1}{2} \int_\Omega |\nabla u|^2 \, dx \tag{143}$$

over the admissible set $\mathcal{A} = \{u \in \mathrm{H}^1(\Omega) : u|_{\partial\Omega} = g\}$. We claim that $\mathcal{E}$ attains its minimum on $\mathcal{A}$.

Step 1 (Boundedness). Let (u_n) be a minimizing sequence: $\mathcal{E}(u_n) \to \inf_{\mathcal{A}} \mathcal{E}$. Since $\mathcal{E}(u_n)$ is bounded, so is $\|\nabla u_n\|_{\mathrm{L}^2}$. The Poincaré inequality (applied to $u_n - \tilde{g}$ for any fixed extension $\tilde{g} \in \mathrm{H}^1(\Omega)$ of g) gives $\|u_n\|_{\mathrm{H}^1} \le C$.

Step 2 (Compactness). The Sobolev space $\mathrm{H}^1(\Omega) = \mathrm{W}^{1,2}(\Omega)$ is a Hilbert space, hence reflexive. By Theorem 3.55, the bounded sequence (u_n) has a weakly convergent subsequence $u_{n_k} \rightharpoonup u$ in $\mathrm{H}^1(\Omega)$. The trace operator is continuous in the weak topology, so $u|_{\partial\Omega} = g$ and $u \in \mathcal{A}$.

Step 3 (Lower semicontinuity). The functional $u \mapsto \int_\Omega |\nabla u|^2 \, dx$ is convex and continuous in the norm topology, so by Mazur's theorem (Theorem 3.57) its sublevel sets are weakly closed. Equivalently, $\mathcal{E}$ is weakly lower semicontinuous:

$$\mathcal{E}(u) \le \liminf_{k\to\infty} \mathcal{E}(u_{n_k}) = \inf_{\mathcal{A}} \mathcal{E}. \tag{144}$$

Therefore u is a minimizer. The minimizer is unique by strict convexity of $\mathcal{E}$, and satisfies the Euler-Lagrange equation $-\Delta u = 0$ in Ω (Laplace's equation).

The entire argument rests on two pillars: reflexivity of H^1 gives weak compactness, and convexity of $\mathcal{E}$ gives weak lower semicontinuity. Remove either and the argument collapses.

3.73 REMARK (Why convexity gives weak lower semicontinuity). Step 3 is the subtlest part of the direct method. Why can't the energy *increase* when we pass to a weak limit? There are three ways to see this.

Direct proof for the Dirichlet energy. Since $u_{n_k} \rightharpoonup u$ in H^1, we have $\nabla u_{n_k} \rightharpoonup \nabla u$ in L^2. Weak convergence means $\langle \nabla u_{n_k}, w\rangle \to \langle \nabla u, w\rangle$ for every $w \in \mathrm{L}^2$. Taking $w = \nabla u$:

$$\|\nabla u\|_{\mathrm{L}^2}^2 = \langle \nabla u, \nabla u\rangle = \lim_k \langle \nabla u_{n_k}, \nabla u\rangle \le \liminf_k \|\nabla u_{n_k}\|_{\mathrm{L}^2} \cdot \|\nabla u\|_{\mathrm{L}^2} \tag{145}$$

by Cauchy–Schwarz. Dividing by $\|\nabla u\|_{\mathrm{L}^2}$ gives $\|\nabla u\|_{\mathrm{L}^2} \le \liminf_k \|\nabla u_{n_k}\|_{\mathrm{L}^2}$, hence $\mathcal{E}(u) \le \liminf_k \mathcal{E}(u_{n_k})$. The same argument shows that **every norm is weakly lower semicontinuous**: the norm can only drop under weak limits, never increase.

Via Mazur's theorem. More generally, let $\mathcal{F}$ be any convex and norm-continuous functional. Its sublevel sets $\{u : \mathcal{F}(u) \le c\}$ are convex (by convexity of $\mathcal{F}$) and norm-closed (by continuity). By Mazur's theorem (Theorem 3.57), convex norm-closed sets are weakly closed. So the sublevel sets are weakly closed, which is equivalent to $\mathcal{F}$ being weakly lower semicontinuous.

The intuition. By Mazur's theorem, there exist convex combinations $v_k = \sum_j \lambda_j^{(k)} u_{n_j}$ converging *strongly* to u. For a convex functional, Jensen's inequality gives

$$\mathcal{F}(v_k) = \mathcal{F}\left(\sum_j \lambda_j^{(k)} u_{n_j}\right) \le \sum_j \lambda_j^{(k)} \mathcal{F}(u_{n_j}). \tag{146}$$

The right-hand side is at most $\max_j \mathcal{F}(u_{n_j})$, which is eventually close to $\liminf \mathcal{F}(u_{n_k})$. By norm continuity, $\mathcal{F}(v_k) \to \mathcal{F}(u)$, giving the inequality. In short: weak limits are limits of averages, and convex functionals can't be fooled by averaging.

What goes wrong without reflexivity

3.74 EXAMPLE (Failure of the direct method in). L^1

Consider the functional

$$\mathcal{F}(u) = \left| \int_0^1 u(x)\,dx - 1 \right| \tag{147}$$

over the constraint set $\mathcal{C} = \{u \in \mathrm{L}^1([0,1]) : u \ge 0,\ \|u\|_{\mathrm{L}^1} = 1\}$. The infimum is $\inf_{\mathcal{C}} \mathcal{F} = 0$, attained by the constant function $u = 1$.

Now replace the objective with something that rewards concentration. Define

$$\mathcal{G}(u) = -\int_0^1 u(x)^2\,dx \tag{148}$$

over the same constraint set $\mathcal{C}$. A minimizing sequence is given by the **approximations to the identity**:

$$u_n(x) = n\,\mathbf{1}_{[0,1/n]}(x). \tag{149}$$

Each $u_n \in \mathcal{C}$ with $\|u_n\|_{\mathrm{L}^1} = 1$ and $\mathcal{G}(u_n) = -n \to -\infty$, so $\inf_{\mathcal{C}} \mathcal{G} = -\infty$.

But more instructive is what happens to the sequence itself. As measures, $u_n\,dx \rightharpoonup^* \delta_0$ in $\mathcal{M}([0,1]) = \mathrm{C}([0,1])^*$: for any continuous test function φ,

$$\int_0^1 \varphi(x)\,u_n(x)\,dx = n\int_0^{1/n} \varphi(x)\,dx \to \varphi(0). \tag{150}$$

The weak-$*$ limit δ_0 is a perfectly good Radon measure, but it does not belong to L^1. The mass has **concentrated** at a single point. This is the failure mode: L^1 is not reflexive, so bounded sequences need not have weakly convergent subsequences in L^1. The sequence escapes to the larger space of measures, and the direct method cannot recover a minimizer in the original space.

The contrast between Example 3.72 and Example **??** illustrates why reflexivity is not a technical convenience but a structural requirement. In reflexive spaces, bounded sets are weakly compact and the direct method closes. In non-reflexive spaces like L^1, minimizing sequences can concentrate or oscillate their way out of the space, and one must either enlarge the space (to measures) or impose additional compactness conditions (such as equi-integrability via the Dunford-Pettis theorem).

3.11 *The Adjoint (Transpose) Operator*

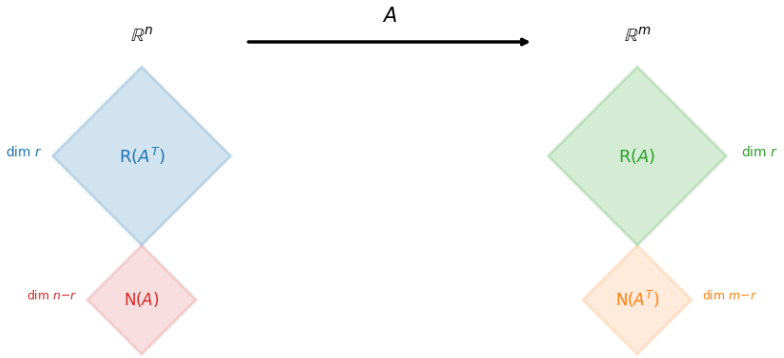
3.11.1 The finite-dimensional picture

In $\mathbb{R}^n$, a linear map $T : \mathbb{R}^n \to \mathbb{R}^m$ is represented by an $m \times n$ matrix A. The **transpose** A^T is the $n \times m$ matrix satisfying

$$\langle Ax, y\rangle = \langle x, A^T y\rangle \quad \text{for all } x \in \mathbb{R}^n,\ y \in \mathbb{R}^m. \tag{151}$$

The fundamental theorem of linear algebra decomposes $\mathbb{R}^n$ and $\mathbb{R}^m$ into four mutually orthogonal subspaces:

$$\mathbb{R}^n = \ker(A) \oplus \mathrm{Range}(A^T), \qquad \mathbb{R}^m = \ker(A^T) \oplus \mathrm{Range}(A). \tag{152}$$



The four fundamental subspaces of an $m \times n$ matrix A of rank r. The domain $\mathbb{R}^n$ decomposes into the row space $R(A^T)$ (dimension r) and the nullspace $N(A)$ (dimension $n - r$), which are orthogonal complements. The codomain $\mathbb{R}^m$ decomposes into the column space $R(A)$ (dimension r) and the left nullspace $N(A^T)$ (dimension $m - r$). The map A sends $R(A^T)$ isomorphically onto $R(A)$, and kills $N(A)$.

The four subspace picture says: A and A^T are perfectly coupled. The kernel of one is the orthogonal complement of the range of the other. This is the cleanest version of the story, and it relies on two features of $\mathbb{R}^n$:

1. **Inner product:** orthogonal complements make sense.

2. **Finite dimensions:** all subspaces are closed, and ranges are automatically closed.

In infinite dimensions, both features can fail. The adjoint theory generalizes this picture, but the four subspace relations require progressively more assumptions as we move away from the finite-dimensional setting.

3.11.2 The Banach adjoint (transpose)

3.75 DEFINITION (Banach adjoint). Let X, Y be normed spaces and $T : X \to Y$ a bounded linear operator. The **adjoint** (or **transpose**) $T^* : Y^* \to X^*$ is defined by

$$(T^* g)(x) = g(Tx) \quad \text{for all } g \in Y^*,\ x \in X. \tag{153}$$

The definition is simple: T^* takes a functional on Y and pulls it back to a functional on X by precomposing with T. If g defines a foliation of Y, then $T^* g$ defines a foliation of X whose level sets are the preimages under T of the level sets of g.

3.76 PROPOSITION (Basic properties). *Let* $T : X \to Y$ *be bounded. Then:*

1. $T^* : Y^* \to X^*$ *is bounded with* $\|T^*\| = \|T\|$.
2. $(S + T)^* = S^* + T^*$ *and* $(\alpha T)^* = \alpha T^*$.
3. $(ST)^* = T^* S^*$ *(reversal of order).*
4. T^* *is injective if and only if* $\overline{\mathrm{Range}(T)} = Y$.
5. T^* *is surjective only if* T *is bounded below (i.e.,* $\|Tx\| \geq c\|x\|$ *for some* $c > 0$*).*

Proof. **(1)** For $g \in Y^*$: $\|T^* g\| = \sup_{\|x\| \leq 1} |g(Tx)| \leq \|g\| \cdot \|T\|$, so $\|T^*\| \leq \|T\|$. For the reverse, the **??** gives $\|Tx\| = \sup_{\|g\| \leq 1} |g(Tx)| = \sup_{\|g\| \leq 1} |(T^* g)(x)| \leq \|T^*\| \cdot \|x\|$, so $\|T\| \leq \|T^*\|$.

(3) $((ST)^* g)(x) = g(STx) = (S^* g)(Tx) = (T^* S^* g)(x)$.

(4) $T^* g = 0$ means $g(Tx) = 0$ for all x, i.e., g vanishes on Range(T). By Hahn-Banach (**??**), g vanishes on $\overline{\mathrm{Range}(T)}$. So $\ker(T^*) = \{0\}$ iff the only functional vanishing on $\overline{\mathrm{Range}(T)}$ is zero, iff $\overline{\mathrm{Range}(T)} = Y$. □

3.11.3 The four subspace relations: general Banach spaces

To state the four subspace relations in Banach spaces, we need a substitute for orthogonal complements. Since we have no inner product, we use **annihilators**.

3.77 DEFINITION (Annihilator). Let X be a normed space.

- For $M \subseteq X$, the **annihilator** of M in X^* is

$$M^\perp = \{f \in X^* : f(x) = 0 \text{ for all } x \in M\}. \tag{154}$$

- For $N \subseteq X^*$, the **pre-annihilator** of N in X is

$${}^\perp N = \{x \in X : f(x) = 0 \text{ for all } f \in N\}. \tag{155}$$

Annihilators are the Banach-space replacement for orthogonal complements. In a Hilbert space with Riesz identification, $M^\perp$ reduces to the usual orthogonal complement. But in general, $M^\perp$ lives in the *dual* space, not in X itself.

3.78 THEOREM (Four subspace relations (general Banach spaces)). *Let* X, Y *be Banach spaces and* $T : X \to Y$ *a bounded linear operator. Then:*

1. $\ker(T^*) = \mathit{Range}(T)^\perp$
2. $\ker(T) = {}^\perp\mathit{Range}(T^*)$
3. $\overline{\mathit{Range}(T)} = {}^\perp \ker(T^*)$
4. $\overline{\mathit{Range}(T^*)} \subseteq \ker(T)^\perp$

Proof. **(1)** $g \in \ker(T^*)$ iff $T^*g = 0$ iff $g(Tx) = 0$ for all x iff $g \in \mathrm{Range}(T)^\perp$.

(2) $x \in {}^\perp\mathrm{Range}(T^*)$ iff $(T^*g)(x) = 0$ for all $g \in Y^*$ iff $g(Tx) = 0$ for all $g \in Y^*$. By the **??**, this holds iff $Tx = 0$ iff $x \in \ker(T)$.

(3) From (1), ${}^\perp\ker(T^*) = {}^\perp(\mathrm{Range}(T)^\perp)$. For any subspace M, the Corollary 3.20 gives ${}^\perp(M^\perp) = \overline{M}$. Applying this to $M = \mathrm{Range}(T)$ yields (3).

(4) If $f \in \overline{\mathrm{Range}(T^*)}$, then $f = \lim T^*g_n$ for some sequence in Y^*. For $x \in \ker(T)$: $f(x) = \lim(T^*g_n)(x) = \lim g_n(Tx) = 0$. So $f \in \ker(T)^\perp$.

□

3.79 REMARK (The asymmetry in relation (4)). Relation (4) is an **inclusion**, not an equality. This is the key difference from the finite-dimensional case, and it arises because:

- In (3), we used ${}^\perp(M^\perp) = \overline{M}$, which is a theorem about X (using Hahn-Banach for X).
- To upgrade (4) to equality, we would need $({}^\perp N)^\perp = \overline{N}$ for subspaces $N \subseteq X^*$. This requires Hahn-Banach applied to X^* with functionals from X^{**}, and it gives $({}^\perp N)^\perp = \overline{N}^{w^*}$ (weak* closure), which equals $\overline{N}$ (norm closure) only when X is reflexive.

So in a non-reflexive space, $\mathrm{Range}(T^*)$ can be a proper subset of $\ker(T)^\perp$ even after taking its closure. The "missing" elements of $\ker(T)^\perp$ are functionals that annihilate $\ker(T)$ but cannot be reached as limits of T^*g_n.

THE CLOSED RANGE THEOREM The most important application of the four subspace picture is the closed range theorem, which restores full symmetry when the range is closed.

3.80 THEOREM (Closed Range Theorem (Banach)). *Let* X, Y *be Banach spaces and* $T : X \to Y$ *bounded. Then*

$$\mathit{Range}(T) \text{ is closed} \iff \mathit{Range}(T^*) \text{ is closed.} \tag{156}$$

When this holds, relation (4) becomes an equality:

$$\mathit{Range}(T^*) = \ker(T)^\perp. \tag{157}$$

Proof. **Sketch.** ($\Rightarrow$) If $\mathrm{Range}(T)$ is closed, the open mapping theorem applied to $T : X/\ker(T) \to \mathrm{Range}(T)$ gives a bounded inverse. One then shows that $T^*g_n \to f$ implies g_n (restricted to $\mathrm{Range}(T)$) converge, producing g with $T^*g = f$.

($\Leftarrow$) Apply the same argument to T^*.

Equality in (4). If $\mathrm{Range}(T^*)$ is closed, then $\mathrm{Range}(T^*) = \overline{\mathrm{Range}(T^*)} \subseteq \ker(T)^\perp$. For the reverse inclusion: if $f \in \ker(T)^\perp$, then f defines a functional on $X/\ker(T)$. Since T induces an isomorphism $X/\ker(T) \cong \mathrm{Range}(T)$ (by the closed range), we can pull f back to a functional g on $\mathrm{Range}(T)$, extend by Hahn-Banach to $g \in Y^*$, and verify $T^*g = f$.

□

The closed range condition is a property of the **operator** T, not the ambient space. It is essential for solvability of equations: $Tx = y$ has a solution iff $y \in \mathrm{Range}(T)$, and the closed range theorem lets us check this via the adjoint: $y \in \overline{\mathrm{Range}(T)}$ iff $g(y) = 0$ for all $g \in \ker(T^*)$. An alternative route to full equality in relation (4) is to work in a reflexive space, where it holds for every bounded operator automatically. An important class of operators with closed range are the Fredholm operators, discussed in Application: Fredholm Operators.

3.11.4 Reflexive spaces: full symmetry restored

In a reflexive space ($X \cong X^{**}$ via the canonical embedding), the asymmetry in relation (4) disappears.

3.81 THEOREM (Four subspace relations (reflexive spaces)). *Let X, Y be **reflexive** Banach spaces and $T : X \to Y$ bounded. Then all four relations hold with equality:*

1. $\ker(T^*) = \mathit{Range}(T)^\perp$
2. $\ker(T) = {}^\perp\mathit{Range}(T^*)$
3. $\overline{\mathit{Range}(T)} = {}^\perp\ker(T^*)$
4. $\overline{\mathit{Range}(T^*)} = \ker(T)^\perp$

Proof. Relations (1)–(3) hold in any Banach space. For (4): in a reflexive space, every functional on X^* comes from an element of X, so the weak* closure and the norm closure of any convex subset of X^* coincide. Therefore $({}^\perp N)^\perp = \overline{N}$ for subspaces $N \subseteq X^*$, and applying this to $N = \mathrm{Range}(T^*)$ gives $\ker(T)^\perp = ({}^\perp\mathrm{Range}(T^*))^\perp = \overline{\mathrm{Range}(T^*)}$.

□

Reflexivity restores the full symmetry of the finite-dimensional picture: each kernel is the annihilator of the other's range, and vice versa. No closed range assumption is needed for the four relations to hold (though the ranges themselves may still not be closed).

3.11.5 Hilbert spaces: orthogonal complements

In a Hilbert space, the Riesz representation theorem identifies H^* with H, and annihilators become orthogonal complements. The Banach adjoint T^* becomes the **Hilbert adjoint**, defined by the familiar relation:

3.82 DEFINITION (Hilbert adjoint). Let H, K be Hilbert spaces and $T : H \to K$ a bounded linear operator. The **Hilbert adjoint** $T^* : K \to H$ is the unique operator satisfying

$$\langle Tx, y\rangle_K = \langle x, T^*y\rangle_H \quad \text{for all } x \in H,\ y \in K. \tag{158}$$

The Hilbert adjoint maps back into the *same type of space* (not the dual), which is why we can write the four subspace decompositions as orthogonal direct sums within H and K themselves.

3.83 COROLLARY (Four subspaces in Hilbert space). *Let* H, K *be Hilbert spaces and* $T : H \to K$ *bounded. Then:*

$$H = \ker(T) \oplus \overline{\mathit{Range}(T^*)}, \qquad K = \ker(T^*) \oplus \overline{\mathit{Range}(T)}. \tag{159}$$

The four subspaces satisfy:

1. $\ker(T^*) = \mathit{Range}(T)^\perp$
2. $\ker(T) = \mathit{Range}(T^*)^\perp$
3. $\overline{\mathit{Range}(T)} = \ker(T^*)^\perp$
4. $\overline{\mathit{Range}(T^*)} = \ker(T)^\perp$

All orthogonal complements are taken in the ambient Hilbert space (not the dual).

This is exactly the finite-dimensional picture, with closures added to account for the possibility that ranges are not closed. Every vector in H decomposes uniquely into a component killed by T and a component in the closure of what T^* produces.

See the four subspace picture above, which now carries over with $A \mapsto T$, $\mathbb{R}^n \mapsto H$, and $\mathbb{R}^m \mapsto K$, and closures added to account for possibly non-closed ranges.

3.12 *Application: Fredholm Operators*

The most important class of operators with closed range are the **Fredholm operators**, where the four subspace picture is as clean as in finite dimensions, even in non-reflexive spaces.

3.84 DEFINITION (Fredholm operator). Let X, Y be Banach spaces and $T : X \to Y$ a bounded linear operator. Then T is **Fredholm** if:

1. $\dim \ker(T) < \infty$,
2. $\operatorname{codim} \operatorname{Range}(T) < \infty$ (equivalently, $\dim(Y/\operatorname{Range}(T)) < \infty$).

The **Fredholm index** of T is

$$\operatorname{ind}(T) = \dim \ker(T) - \operatorname{codim} \operatorname{Range}(T). \tag{160}$$

The index measures the net dimensional defect of T: how many more dimensions are killed than are missed.

3.85 THEOREM (Properties of Fredholm operators). *Let $T : X \to Y$ be Fredholm. Then:*

1. $\operatorname{Range}(T)$ *is closed.*
2. $T^* : Y^* \to X^*$ *is Fredholm with* $\operatorname{ind}(T^*) = -\operatorname{ind}(T)$.
3. *The four subspace relations hold with full equality:*

$$\ker(T^*) = \operatorname{Range}(T)^\perp, \qquad \operatorname{Range}(T^*) = \ker(T)^\perp. \tag{161}$$

4. X *and* Y *decompose as topological direct sums:*

$$X = \ker(T) \oplus Z, \qquad Y = \operatorname{Range}(T) \oplus W, \tag{162}$$

where $T|_Z : Z \to \operatorname{Range}(T)$ *is an isomorphism and* $\dim W = \operatorname{codim} \operatorname{Range}(T)$.

Proof. **(1)** Since $\operatorname{codim} \operatorname{Range}(T) < \infty$, there exists a finite-dimensional subspace $W \subset Y$ with $Y = \operatorname{Range}(T) + W$. Define $\tilde{T} : X \times W \to Y$ by $\tilde{T}(x, w) = Tx + w$. This is bounded, linear, and surjective, so the open mapping theorem gives $\tilde{T}$ open. It follows that $\operatorname{Range}(T)$ is closed. (Alternatively: any subspace of finite codimension in a Banach space is closed.)

(2) From (1) and the closed range theorem, $\operatorname{Range}(T^*)$ is closed. We have $\ker(T^*) = \operatorname{Range}(T)^\perp$, so $\dim \ker(T^*) = \operatorname{codim} \operatorname{Range}(T) < \infty$. Similarly, $\operatorname{Range}(T^*) = \ker(T)^\perp$ (by the closed range theorem), so $\operatorname{codim} \operatorname{Range}(T^*) = \dim \ker(T) < \infty$. Therefore T^* is Fredholm and $\operatorname{ind}(T^*) = \operatorname{codim} \operatorname{Range}(T) - \dim \ker(T) = -\operatorname{ind}(T)$.

(3) Follows from (1) and the closed range theorem.

(4) Every finite-dimensional subspace of a Banach space is complemented. So $\ker(T)$ has a topological complement Z. The restriction $T|_Z : Z \to \operatorname{Range}(T)$ is bounded, bijective, and between Banach

spaces, so the open mapping theorem makes it an isomorphism. The complement W exists because Range(T) has finite codimension.

□

3.86 REMARK (Fredholm decomposition vs. Hilbert decomposition). The Fredholm decomposition $X = \ker(T) \oplus Z$ guarantees that *some* closed complement Z exists (since ker(T) is finite-dimensional), but there are infinitely many such Z. For example, if $\ker(T) = \operatorname{span}\{v\}$ in $\mathbb{R}^2$, then any line not parallel to v is a valid complement, and each gives a different projection onto ker(T): the same vector x decomposes differently depending on which Z we choose. Without an inner product, no choice is preferred.

In a Hilbert space, the inner product singles out the **orthogonal** complement $Z = \ker(T)^\perp$: the unique complement for which the projection is self-adjoint (equivalently, the decomposition $x = x_{\ker} + x_Z$ minimizes $\|x_Z\|$). The four subspace theorem identifies this canonical choice concretely as $\ker(T)^\perp = \overline{\text{Range}(T^*)}$ — the closure of what T^* produces.

Despite the non-uniqueness of Z in a general Banach space, the Fredholm decomposition still tells us that T is "almost invertible" up to a finite-dimensional correction, regardless of the geometry of the ambient space.

3.87 EXAMPLE (Shift operators on). ℓ^p

The shift operators on ℓ^p ($1 \le p < \infty$) illustrate the index:

- **Right shift** $R(x_1, x_2, \ldots) = (0, x_1, x_2, \ldots)$: injective, range has codimension 1. So $\operatorname{ind}(R) = 0 - 1 = -1$.
- **Left shift** $L(x_1, x_2, \ldots) = (x_2, x_3, \ldots)$: surjective, kernel is $\operatorname{span}\{e_1\}$. So $\operatorname{ind}(L) = 1 - 0 = 1$.

Note $L = R^*$ (under the natural pairing of ℓ^p and ℓ^q), confirming $\operatorname{ind}(L) = -\operatorname{ind}(R)$. These are Fredholm on every ℓ^p, including the non-reflexive cases $p = 1$ and $p = \infty$.

3.88 THEOREM (Stability of the Fredholm index). *Let* $T : X \to Y$ *be Fredholm.*

1. ***Compact perturbation:*** *If* $K : X \to Y$ *is compact, then* $T + K$ *is Fredholm with* $\operatorname{ind}(T + K) = \operatorname{ind}(T)$.
2. ***Small perturbation:*** *There exists* $\varepsilon > 0$ *such that if* $\|S\| < \varepsilon$, *then* $T + S$ *is Fredholm with* $\operatorname{ind}(T + S) = \operatorname{ind}(T)$.

In particular, the Fredholm index is a topological invariant: it is constant on connected components of the set of Fredholm operators.

The stability theorem is why the Fredholm index matters: it is computable (often by deformation to a simpler operator) and robust under perturbation. This makes it the central tool in index theory, where one relates the analytic index of a differential operator to topological data of the underlying manifold.

3.89 REMARK (The Fredholm alternative). For any Fredholm operator $T : X \to Y$, the solvability condition

$$Tx = y \text{ is solvable} \iff \psi(y) = 0 \text{ for all } \psi \in \ker(T^*) \tag{163}$$

holds, since Range(T) is closed and $\ker(T^*) = \text{Range}(T)^\perp$ by Theorem 3.85. However, the classical **Fredholm alternative** — the dichotomy that T is either bijective or has a nontrivial kernel — requires

$\operatorname{ind}(T) = 0$. When the index is zero, $\dim\ker(T) = \operatorname{codim}\operatorname{Range}(T)$, so injectivity and surjectivity are equivalent: either both hold (and T is bijective) or neither does.

For nonzero index this dichotomy fails. The right shift R on ℓ^p (Example 3.87) is injective but not surjective: $\ker(R) = \{0\}$ yet $\operatorname{codim}\operatorname{Range}(R) = 1$. There is no contradiction with the solvability condition — $Rx = y$ is solvable if and only if $y \perp \ker(R^*) = \operatorname{span}\{e_1^*\}$, i.e., $y_1 = 0$ — but the "either bijective or..." structure is lost.

3.12.1 Solvability conditions

The solvability condition from Remark 3.89 is one of the most useful consequences of Fredholm theory in PDE applications. We state it explicitly.

3.90 COROLLARY (Fredholm solvability condition). *Let* $T : X \to Y$ *be a Fredholm operator with* $\ker(T^*) = \operatorname{span}\{\psi_1, \dots, \psi_m\}$. *Then the equation* $Tx = f$ *has a solution if and only if*

$$\psi_j(f) = 0, \qquad j = 1, \dots, m. \tag{164}$$

When a solution exists, it is unique up to an element of $\ker(T)$*: the general solution is* $x = x_p + \sum_{i=1}^{n} \alpha_i v_i$*, where* x_p *is any particular solution,* $\{v_1, \dots, v_n\}$ *is a basis of* $\ker(T)$*, and the* α_i *are arbitrary.*

Proof. $Tx = f$ is solvable iff $f \in \operatorname{Range}(T)$. Since T is Fredholm, $\operatorname{Range}(T)$ is closed (Theorem 3.85), so $\operatorname{Range}(T) = {}^{\perp}\ker(T^*)$ by Theorem 3.78. Therefore $f \in \operatorname{Range}(T)$ iff $\psi(f) = 0$ for all $\psi \in \ker(T^*)$, which is equivalent to $\psi_j(f) = 0$ for each basis element. □

In PDE applications, the ψ_j are typically known explicitly (from the symmetry or structure of the problem), and the solvability conditions take the form of **integral constraints** on the data.

3.91 EXAMPLE (Solvability of the Neumann problem). Consider the Neumann problem for Poisson's equation on a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$:

$$-\Delta u = f \quad \text{in } \Omega, \qquad \frac{\partial u}{\partial n} = 0 \quad \text{on } \partial\Omega. \tag{165}$$

Formulated weakly, the operator $T = -\Delta_N : H^1(\Omega) \to H^1(\Omega)^*$ defined by

$$\langle Tu, v\rangle = \int_\Omega \nabla u \cdot \nabla v \, dx \tag{166}$$

is Fredholm with index zero:

- **Kernel.** $\ker(T) = \operatorname{span}\{1\}$: the constant functions, since $\int_\Omega |\nabla u|^2 \, dx = 0$ implies u is constant.
- **Cokernel.** By self-adjointness, $\ker(T^*) = \ker(T) = \operatorname{span}\{1\}$, so $\operatorname{codim}\operatorname{Range}(T) = 1$.

By Corollary 3.90, $-\Delta_N u = f$ has a solution if and only if $\psi(f) = 0$ for the single cokernel element $\psi = 1$:

$$\int_\Omega f\,dx = 0 \qquad \text{(the data must have zero mean).} \tag{167}$$

This is the classical compatibility condition: by the divergence theorem, $\int_\Omega \Delta u\,dx = \int_{\partial\Omega} \frac{\partial u}{\partial n}\,dS = 0$, so $\int_\Omega f\,dx = 0$ is necessary. Fredholm theory shows it is also sufficient. When satisfied, the solution is unique up to an additive constant.

The Fredholm index is $\operatorname{ind}(T) = 1 - 1 = 0$, but T is not invertible. To obtain a well-posed problem, we must simultaneously *impose a constraint* to select a unique solution (pin the constant) and *add a degree of freedom* to absorb right-hand sides that violate the compatibility condition.

3.12.2 The bordered operator

3.92 DEFINITION (Bordered operator). Let $T : X \to Y$ be Fredholm with $\operatorname{ind}(T) = 0$ and $n = \dim\ker(T)$. Let $\{\phi_1, \dots, \phi_n\} \subset Y$ and $\{\ell_1, \dots, \ell_n\} \subset X^*$. The **bordered operator** is

$$\widehat{T} = \begin{pmatrix} T & \Phi \\ L & 0 \end{pmatrix} : X \times \mathbb{R}^n \to Y \times \mathbb{R}^n, \tag{168}$$

where $\Phi : \mathbb{R}^n \to Y$ maps $\alpha \mapsto \sum_{j=1}^{n} \alpha_j \phi_j$ and $L : X \to \mathbb{R}^n$ maps $x \mapsto (\ell_1(x), \dots, \ell_n(x))$.

The bordered system $\widehat{T}(x, \alpha) = (y, c)$ reads:

$$Tx + \sum_{j=1}^{n} \alpha_j \phi_j = y, \qquad \ell_j(x) = c_j, \quad j = 1, \dots, n. \tag{169}$$

The idea is transparent: the functionals ℓ_j select a unique element of $\ker(T)$ (by prescribing the "coordinates" c_j), and the vectors ϕ_j span a complement to $\operatorname{Range}(T)$, so the extra unknowns α_j allow us to absorb the component of y outside the range.

3.93 THEOREM (Invertibility of the bordered operator). *Let $T : X \to Y$ be Fredholm with $\operatorname{ind}(T) = 0$ and $n = \dim\ker(T)$. Let $\{v_1, \dots, v_n\}$ be a basis of $\ker(T)$ and $\{\psi_1, \dots, \psi_n\} \subset Y^*$ a basis of $\ker(T^*)$. Then the bordered operator $\widehat{T}$ is boundedly invertible if and only if:*

1. ***The ℓ_j detect the kernel:*** *the $n \times n$ matrix $M_{ij} = \ell_i(v_j)$ is invertible.*
2. ***The ϕ_j complement the range:*** *the $n \times n$ matrix $N_{ij} = \psi_i(\phi_j)$ is invertible.*

Proof. By Theorem 3.85, the spaces decompose as

$$X = \ker(T) \oplus Z, \qquad Y = \operatorname{Range}(T) \oplus W, \tag{170}$$

where $T|_Z : Z \to \operatorname{Range}(T)$ is an isomorphism and W has dimension n.

Injectivity. Suppose $\widehat{\mathrm{T}}(x,\alpha) = (0,0)$. The second component gives $\ell_j(x) = 0$ for all j. The first gives $\mathrm{T}x = -\sum \alpha_j \phi_j$. Applying each $\psi_i \in \ker(\mathrm{T}^*) = \mathrm{Range}(\mathrm{T})^\perp$ to the first equation:

$$0 = \psi_i(\mathrm{T}x) = -\sum_{j=1}^{n} \alpha_j \psi_i(\phi_j) = -(\mathrm{N}\alpha)_i. \tag{171}$$

Since N is invertible by assumption (2), $\alpha = 0$. Therefore $\mathrm{T}x = 0$, so $x \in \ker(\mathrm{T})$. Writing $x = \sum \beta_j v_j$, we have $\ell_i(x) = \sum \beta_j \ell_i(v_j) = (\mathrm{M}\beta)_i = 0$. Since M is invertible by assumption (1), $\beta = 0$, hence $x = 0$.

Surjectivity. Given $(y, c) \in \mathrm{Y} \times \mathbb{R}^n$, we must find (x, α) such that $\mathrm{T}x + \sum \alpha_j \phi_j = y$ and $\ell_j(x) = c_j$.

Step 1. Decompose $y = y_{\mathrm{R}} + y_{\mathrm{W}}$ with $y_{\mathrm{R}} \in \mathrm{Range}(\mathrm{T})$ and $y_{\mathrm{W}} \in \mathrm{W}$. Since $\{\psi_i\}$ is a basis for $\mathrm{Range}(\mathrm{T})^\perp$, the coordinates of y_{W} relative to the ϕ_j are determined by $\psi_i(y) = \sum \alpha_j \psi_i(\phi_j)$, i.e., $\alpha = \mathrm{N}^{-1}(\psi_1(y), \dots, \psi_n(y))^{\mathrm{T}}$. With this choice, $y - \sum \alpha_j \phi_j \in \mathrm{Range}(\mathrm{T})$.

Step 2. Since $\mathrm{T}|_{\mathrm{Z}}$ is an isomorphism onto $\mathrm{Range}(\mathrm{T})$, there exists a unique $z \in \mathrm{Z}$ with $\mathrm{T}z = y - \sum \alpha_j \phi_j$.

Step 3. Set $x = z + \sum \beta_j v_j$ where $\beta = \mathrm{M}^{-1}(c_1 - \ell_1(z), \dots, c_n - \ell_n(z))^{\mathrm{T}}$. Then $\mathrm{T}x = \mathrm{T}z = y - \sum \alpha_j \phi_j$ and $\ell_i(x) = \ell_i(z) + (\mathrm{M}\beta)_i = c_i$.

Bounded inverse. $\widehat{\mathrm{T}}$ is a bounded bijection between Banach spaces, so the open mapping theorem gives a bounded inverse.

□

3.94 REMARK. The conditions are easy to check in practice. Condition (1) says: the constraints ℓ_j must be *linearly independent when restricted to the kernel*. Condition (2) says: the vectors ϕ_j must be *transverse to the range*, i.e., they span a complement. When T is self-adjoint (so $\ker(\mathrm{T}) = \ker(\mathrm{T}^*)$), a natural choice is $\phi_j = v_j$ — use the kernel elements themselves to span the cokernel complement.

3.95 EXAMPLE (Bordered Neumann Laplacian). Returning to the Neumann Laplacian $\mathrm{T} = -\Delta_{\mathrm{N}}$ with $n = 1$:

- **Kernel basis:** $v_1 = 1$ (the constant function).
- **Cokernel basis:** $\psi_1 = \frac{1}{|\Omega|}\int_\Omega (\cdot)\, dx$ (the mean functional), since T is self-adjoint.

Choose $\phi_1 = 1$ and $\ell_1(u) = \frac{1}{|\Omega|}\int_\Omega u\, dx$. Then:

- $\mathrm{M}_{11} = \ell_1(v_1) = \frac{1}{|\Omega|}\int_\Omega 1\, dx = 1 \neq 0$. ✓
- $\mathrm{N}_{11} = \psi_1(\phi_1) = \frac{1}{|\Omega|}\int_\Omega 1\, dx = 1 \neq 0$. ✓

By Theorem 3.93, the bordered operator

$$\widehat{\mathrm{T}} = \begin{pmatrix} -\Delta_{\mathrm{N}} & 1 \\ \dfrac{1}{|\Omega|}\displaystyle\int_\Omega (\cdot)\, dx & 0 \end{pmatrix} \tag{172}$$

is invertible. The bordered system

$$-\Delta u + \lambda = f, \qquad \frac{1}{|\Omega|}\int_\Omega u\, dx = 0, \tag{173}$$

has a unique solution (u, λ) for every $f \in \mathrm{H}^1(\Omega)^*$. The constraint pins the mean of u to zero, and the Lagrange multiplier $\lambda = \frac{1}{|\Omega|}\int_\Omega f\,dx$ absorbs the mean of f.

Proof. That $\lambda = \frac{1}{|\Omega|}\int_\Omega f\,dx$ follows by integrating the first equation: $-\int_\Omega \Delta u\,dx + \lambda|\Omega| = \int_\Omega f\,dx$. By the divergence theorem and the Neumann condition, $\int_\Omega \Delta u\,dx = \int_{\partial\Omega} \frac{\partial u}{\partial n}\,dS = 0$, so $\lambda = \frac{1}{|\Omega|}\int_\Omega f\,dx$.

□

3.12.3 Solving the bordered system by block elimination

The bordered system can be solved using T and T^* as **black boxes** — we never form the bordered matrix, only call the forward and adjoint solvers. Consider the general bordered system

$$\begin{pmatrix} \mathrm{T} & \phi \\ \ell & 0 \end{pmatrix}\begin{pmatrix} x \\ \lambda \end{pmatrix} = \begin{pmatrix} f \\ c \end{pmatrix}, \tag{174}$$

where $\mathrm{T} : \mathrm{X} \to \mathrm{Y}$ is Fredholm with index zero, $\dim\ker(\mathrm{T}) = 1$, $v \in \ker(\mathrm{T})$, and $\psi \in \ker(\mathrm{T}^*)$. The algorithm is a walk through the four subspace decomposition:

$$\mathrm{X} = \ker(\mathrm{T}) \oplus \mathrm{Z}, \qquad \mathrm{Y} = \mathrm{Range}(\mathrm{T}) \oplus \ker(\mathrm{T}^*). \tag{175}$$

Step 1: Project f onto $\ker(\mathrm{T}^*)$ to determine λ. The right-hand side $f \in \mathrm{Y}$ decomposes as $f = f_\mathrm{R} + f_\perp$ with $f_\mathrm{R} \in \mathrm{Range}(\mathrm{T})$ and $f_\perp \in \ker(\mathrm{T}^*)$. Apply $\psi \in \ker(\mathrm{T}^*) = \mathrm{Range}(\mathrm{T})^\perp$ to the first row: $\psi(\mathrm{T}x + \lambda\phi) = \psi(f)$. Since $\psi(\mathrm{T}x) = (\mathrm{T}^*\psi)(x) = 0$, this kills the $\mathrm{Range}(\mathrm{T})$ component entirely and reads off the $\ker(\mathrm{T}^*)$ component:

$$\lambda = \frac{\psi(f)}{\psi(\phi)}. \tag{176}$$

The denominator $\psi(\phi) \neq 0$ is the transversality condition from Theorem 3.93: ϕ is not in $\mathrm{Range}(\mathrm{T})$, so ψ sees it. This works for **any** f — the multiplier λ absorbs whatever component of f lies outside $\mathrm{Range}(\mathrm{T})$.

Step 2: Solve in $\mathrm{Range}(\mathrm{T})$ for a particular solution. By construction, $f - \lambda\phi$ has no $\ker(\mathrm{T}^*)$ component: Step 1 chose λ precisely so that $f - \lambda\phi \in \mathrm{Range}(\mathrm{T})$. The restriction $\mathrm{T}|_\mathrm{Z} : \mathrm{Z} \to \mathrm{Range}(\mathrm{T})$ is an isomorphism (Theorem 3.85), so the black-box solver returns the unique $x_p \in \mathrm{Z}$ with

$$\mathrm{T}x_p = f - \lambda\phi. \tag{177}$$

This is one call to the forward solver. The solution x_p lies in the complement Z of $\ker(\mathrm{T})$, but a general solver may return x_p only up to an element of $\ker(\mathrm{T})$ — the next step handles this.

Step 3: Project onto $\ker(\mathrm{T})$ using the constraint. The general solution of the first row is $x = x_p + \alpha v$, decomposed along $\mathrm{X} = \ker(\mathrm{T}) \oplus \mathrm{Z}$. The parameter α is the $\ker(\mathrm{T})$ component, which T cannot see. The second row $\ell(x) = c$ determines it:

$$\alpha = \frac{c - \ell(x_p)}{\ell(v)}. \tag{178}$$

The denominator $\ell(v) \neq 0$ is the detection condition from Theorem 3.93: ℓ is nontrivial on $\ker(\mathrm{T})$. The solution is $x = x_p + \alpha v$, λ from Step 1.

3.96 EXAMPLE (Block elimination for the Neumann Laplacian). For the Neumann Laplacian $\mathrm{T} = -\Delta_{\mathrm{N}}$ with $v = 1$, $\psi = \ell = \frac{1}{|\Omega|}\int_\Omega (\cdot)\,dx$, and $\phi = 1$:

- **Step 1.** Project onto $\ker(\mathrm{T}^*) = \operatorname{span}\{1\}$: $\lambda = \psi(f)/\psi(\phi) = \frac{1}{|\Omega|}\int_\Omega f\,dx$, the mean of f. This extracts the incompatible component of f.
- **Step 2.** Solve in $\mathrm{Range}(\mathrm{T})$: $-\Delta_{\mathrm{N}} u_p = f - \lambda$. The right-hand side has zero mean by construction, so this is a compatible Neumann problem for any f.
- **Step 3.** Project onto $\ker(\mathrm{T}) = \operatorname{span}\{1\}$: $\alpha = (c - \ell(u_p))/\ell(1) = c - \frac{1}{|\Omega|}\int_\Omega u_p\,dx$, and $u = u_p + \alpha$.

Setting $c = 0$ pins the mean of u to zero.

3.97 REMARK. The bordered operator technique is not limited to index-zero operators or to self-adjoint problems. For a general Fredholm operator with $\mathrm{ind}(\mathrm{T}) = k \neq 0$, one can still form a bordered system, but the sizes of Φ and L differ: if $\dim\ker(\mathrm{T}) = n$ and $\operatorname{codim}\mathrm{Range}(\mathrm{T}) = m$ (with $n - m = k$), then

$$\widehat{\mathrm{T}} = \begin{pmatrix} \mathrm{T} & \Phi \\ \mathrm{L} & 0 \end{pmatrix} : \mathrm{X} \times \mathbb{R}^m \to \mathrm{Y} \times \mathbb{R}^n \tag{179}$$

is invertible under the analogous transversality and detection conditions. The rectangular bordering reflects the dimensional mismatch measured by the index.

4 DISTRIBUTIONS AND SOBOLEV SPACES

Big Idea

Distributions are the **dual space** of smooth, compactly supported test functions. The topology on test functions is *not* a norm topology; it is an inductive limit of Fréchet spaces. But the topology on distributions themselves is the **weak-$*$ topology** from the duality chapter, making convergence of distributions exactly the pointwise convergence of functionals we have already studied.

4.1 *Why distributions?*

Classical analysis works with functions that assign values to points. This breaks down in at least three ways.

Differentiation destroys regularity. The absolute value function $|x|$ is continuous but not differentiable at the origin. Its "derivative" should be the sign function, and the "derivative" of the sign function should be $2\delta_0$, but neither of these operations makes sense in the classical framework. More dramatically, solutions to the wave equation develop shock discontinuities, and the fundamental solution of Laplace's equation has a point singularity. We need a derivative that always exists.

Point evaluation is too rigid. The Dirac delta "function" $\delta(x)$, which physicists write as $\int \delta(x)\varphi(x)\,dx = \varphi(0)$, is not a function: no locally integrable function can reproduce point evaluation via integration. Yet δ appears naturally as the identity for convolution, the fundamental solution of the identity operator, and the density of a point mass. We need a framework where δ is a legitimate object.

Measures are not enough. The dual of $C_c(\mathbb{R}^d)$ (continuous, compactly supported functions) is the space of Radon measures, by the Riesz representation theorem. Measures handle δ just fine (it is a point mass), but the derivative δ', defined by $\langle \delta', \varphi\rangle = -\varphi'(0)$, is *not* a measure. It is a genuinely new object that only exists when the test functions are differentiable. To get a framework closed under differentiation, we must test against C_c^∞ rather than just C_c. Write $\mathcal{D}(\Omega) = C_c^\infty(\Omega)$ for this space of test functions.

Distribution theory resolves all three issues in one stroke: define a **distribution** to be a continuous linear functional on $\mathcal{D}(\Omega)$. Every locally integrable function, every Radon measure, and every derivative of every distribution is again a distribution. The price is that we must carefully topologize $\mathcal{D}(\Omega)$.

4.2 *The two topologies*

The word "continuous" in the definition of a distribution requires a topology on $\mathcal{D}(\Omega)$. The right topology turns out to be surprisingly subtle: no single norm can capture both "all derivatives converge uniformly" and "supports stay bounded." The solution is an **inductive limit**: we decompose $\mathcal{D}(\Omega)$ into Fréchet spaces $\mathcal{D}_K$ (one for each compact $K \subset \Omega$), each equipped with the seminorms $\|\cdot\|_{C^k(K)}$, and take the finest locally convex topology making every inclusion continuous. The resulting convergence is very strong: a sequence $\varphi_n \to \varphi$ only if all supports sit in a single compact set and all derivatives converge uniformly. The details of this construction are in the next section.

Once we have $\mathcal{D}(\Omega)$ topologized, its dual $\mathcal{D}'(\Omega)$ carries the **weak-$*$ topology**, exactly the construction from the duality chapter. Convergence $T_n \to T$ in $\mathcal{D}'(\Omega)$ means $\langle T_n, \varphi\rangle \to \langle T, \varphi\rangle$ for every test function φ. This is **convergence in the sense of distributions**, one of the weakest convergence notions in analysis, yet limits are unique (the weak-$*$ topology is Hausdorff) and every Cauchy sequence converges.

4.3 *The hierarchy of generalized functions*

Distribution theory fits into a chain of increasing generality:

$$\mathrm{C}^\infty(\Omega) \subset \mathrm{C}(\Omega) \subset \mathrm{L}^1_{\mathrm{loc}}(\Omega) \subset \mathcal{M}(\Omega) \subset \mathcal{D}'(\Omega) \tag{180}$$

where $\mathcal{M}(\Omega)$ denotes Radon measures. Each inclusion is strict:

- The sign function is in $\mathrm{L}^1_{\mathrm{loc}}$ but not continuous.
- The Dirac delta δ is a measure but not in $\mathrm{L}^1_{\mathrm{loc}}$.
- The derivative δ' is a distribution but not a measure.

The embedding $\mathrm{L}^1_{\mathrm{loc}}(\Omega) \hookrightarrow \mathcal{D}'(\Omega)$ given by $f \mapsto \mathrm{T}_f$, where $\langle \mathrm{T}_f, \varphi \rangle = \int f \varphi \, dx$, is injective and continuous. This justifies identifying a function with the distribution it defines, so we write $\langle f, \varphi \rangle$ instead of $\langle \mathrm{T}_f, \varphi \rangle$ without ambiguity.

4.4 *What comes next*

With the overview in hand, we now develop the details:

1. **Test functions and their topology**: the inductive limit construction, sequential convergence, and key properties of $\mathcal{D}(\Omega)$.

2. **Distributions as continuous linear functionals**: the formal definition, the weak-$*$ topology, separation properties, and the fundamental examples.

3. **Operations on distributions**: differentiation, multiplication by smooth functions, and convolution, all defined by duality.

4. **Sobolev spaces** $\mathrm{W}^{k,p}(\Omega)$: where distributional derivatives meet L^p integrability, giving Banach spaces suited to PDE theory.

4.5 *The Space of Test Functions*

We construct the space $\mathcal{D}(\Omega)$ and equip it with the topology that determines which linear functionals qualify as distributions.

4.5.1 Definition

Throughout, $\Omega \subseteq \mathbb{R}^d$ is an open set. We use standard multi-index notation: for $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}_0^d$, we write $|\alpha| = \alpha_1 + \cdots + \alpha_d$ and $D^\alpha = \partial_1^{\alpha_1} \cdots \partial_d^{\alpha_d}$.

4.1 DEFINITION (Test functions). The space of **test functions** on Ω is

$$\mathcal{D}(\Omega) = C_c^\infty(\Omega) = \Big\{\varphi \in C^\infty(\Omega) : \operatorname{supp}(\varphi) \Subset \Omega\Big\}, \tag{181}$$

where $\operatorname{supp}(\varphi) = \overline{\{x \in \Omega : \varphi(x) \neq 0\}}$ and $\operatorname{supp}(\varphi) \Subset \Omega$ means the support is a compact subset of Ω.

Test functions are the "probes" we use to measure distributions. They are infinitely differentiable and vanish outside a bounded region, so integrals against them are always well-defined.

Notation. In the duality chapter, we wrote X' for the algebraic dual (all linear functionals) and X^* for the topological dual (continuous linear functionals). The standard convention in distribution theory, due to Schwartz, writes $\mathcal{D}'(\Omega)$ for the space of distributions, which is the **topological** dual of $\mathcal{D}(\Omega)$. This conflicts with our earlier notation, but is universal in the literature. Throughout this chapter, $\mathcal{D}'(\Omega)$ always means the topological dual, i.e. the space of continuous linear functionals on $\mathcal{D}(\Omega)$.

4.5.2 Why C_c^∞ is rich enough

Before diving into the topology of $\mathcal{D}(\Omega)$, we address the first natural question: is this space large enough to be useful?

$\mathcal{D}(\Omega)$ is dense in $L^p(\Omega)$ for $1 \le p < \infty$ and in $C_0(\Omega)$. This follows from the approximation results in the function spaces chapter: continuous compactly supported functions are dense in L^p, and mollification promotes them to smooth compactly supported functions.

This density is a special case of a general principle.

4.2 THEOREM (Stone-Weierstrass (locally compact version)). *Let* X *be a locally compact Hausdorff space. If* $\mathcal{A} \subset C_0(X)$ *is a subalgebra that* ***separates points*** *(for every* $x \neq y$*, some* $f \in \mathcal{A}$ *has* $f(x) \neq f(y)$*) and* ***vanishes nowhere*** *(for every* x*, some* $f \in \mathcal{A}$ *has* $f(x) \neq 0$*), then* $\mathcal{A}$ *is dense in* $C_0(X)$ *in the supremum norm.*

The space $\mathcal{D}(\Omega)$ satisfies both hypotheses on $X = \Omega$: smooth bump functions can separate any two points, and for every $x \in \Omega$ one can construct a bump with $\varphi(x) > 0$. Stone-Weierstrass therefore gives $\mathcal{D}(\Omega)$ dense in $C_0(\Omega)$, and since C_0 is dense in L^p for $1 \le p < \infty$, density in L^p follows by transitivity.

4.3 REMARK. The mollification proof from the function spaces chapter is more constructive: it produces an explicit approximating sequence $\rho_\epsilon * f \to f$. Stone-Weierstrass gives a stronger structural insight — $\mathcal{D}(\Omega)$ is dense because it is an algebra that separates points, not because of any particular approximation procedure. This matters when we move beyond $\mathbb{R}^d$: on a manifold or abstract locally compact space, mollification requires a convolution structure, but Stone-Weierstrass only requires an algebra of continuous functions.

More practically, density of $\mathcal{D}(\Omega)$ is what makes the choice of test functions *non-restrictive*. If T_1 and T_2 are distributions that agree on every test function, they agree as distributions — because $\mathcal{D}(\Omega)$ is rich enough to detect any difference.

$\mathcal{D}(\Omega)$ is also dense in the Sobolev spaces $W^{k,p}(\Omega)$ for $1 \le p < \infty$ (the **Meyers-Serrin theorem**, $H = W$). This is a deeper result that does not follow from Stone-Weierstrass alone — it requires the calculus of mollification to commute with weak derivatives.

4.5.3 The inductive limit topology

The solution is to build the topology in two stages tao$_d istributions.$

4.4 DEFINITION (Fréchet space). A **Fréchet space** is a topological vector space X whose topology is generated by a countable family of seminorms $\{p_k\}_{k=0}^{\infty}$ and which is complete with respect to the resulting metric

$$d(x,y) = \sum_{k=0}^{\infty} \frac{1}{2^k} \frac{p_k(x-y)}{1+p_k(x-y)}. \tag{182}$$

Every Banach space is a Fréchet space (take $p_0 = \|\cdot\|$ and $p_k = 0$ for $k \ge 1$). The difference is that a Fréchet space needs only seminorms, not a single norm. A seminorm p satisfies all the axioms of a norm except that $p(x) = 0$ does not imply $x = 0$. Consequently, a Fréchet space is metrizable and complete (so the Baire category theorem applies), but the metric is generally not translation-homogeneous: $d(\lambda x, 0) \neq |\lambda|\, d(x,0)$. This means there is no notion of "operator norm" for linear maps between Fréchet spaces, and bounded sets are defined by seminorms rather than by a single ball.

4.5 DEFINITION (The spaces). $\mathcal{D}_K$

For a compact set $K \subset \Omega$, define

$$\mathcal{D}_K(\Omega) = \{\varphi \in C_c^{\infty}(\Omega) : \operatorname{supp}(\varphi) \subseteq K\}. \tag{183}$$

Equip $\mathcal{D}_K$ with the countable family of seminorms

$$\|\varphi\|_{K,k} = \max_{|\alpha| \le k} \sup_{x \in K} |D^{\alpha}\varphi(x)|, \qquad k = 0, 1, 2, \ldots \tag{184}$$

The metric

$$d_K(\varphi,\psi) = \sum_{k=0}^{\infty} \frac{1}{2^k} \frac{\|\varphi-\psi\|_{K,k}}{1+\|\varphi-\psi\|_{K,k}} \tag{185}$$

makes $\mathcal{D}_K$ a **Fréchet space** (complete and metrizable).

4.6 DEFINITION (Inductive limit topology on). $\mathcal{D}(\Omega)$

Let $\mathrm{K}_1 \subset \mathrm{K}_2 \subset \cdots$ be a compact exhaustion of Ω, so that $\mathcal{D}(\Omega) = \bigcup_{j=1}^{\infty} \mathcal{D}_{\mathrm{K}_j}$. The **inductive limit topology** on $\mathcal{D}(\Omega)$ is the finest locally convex topology making every inclusion

$$\iota_j : \mathcal{D}_{\mathrm{K}_j}(\Omega) \hookrightarrow \mathcal{D}(\Omega) \tag{186}$$

continuous. Concretely, a convex set $\mathrm{U} \subseteq \mathcal{D}(\Omega)$ is open if and only if $\mathrm{U} \cap \mathcal{D}_{\mathrm{K}_j}$ is open in $\mathcal{D}_{\mathrm{K}_j}$ for every j.

Local convexity. A topological vector space is **locally convex** if every point has a neighborhood basis consisting of convex sets, or equivalently, if the topology is generated by a family of seminorms. Each $\mathcal{D}_{\mathrm{K}_j}$ is locally convex because its topology comes from the seminorms $\|\cdot\|_{\mathrm{K}_j,k}$. The inductive limit topology on $\mathcal{D}(\Omega)$ is locally convex *by construction*: it is defined as the finest **locally convex** topology making the inclusions ι_j continuous (Definition 4.6). Local convexity matters because it is the hypothesis of the geometric Hahn-Banach theorem (the separation theorem). Without it, the dual space $\mathcal{D}'(\Omega)$ might be too small to separate points of $\mathcal{D}(\Omega)$.

Why not a single norm? In a Banach space, the topology is determined by a single norm. In a Fréchet space, it comes from countably many seminorms. For $\mathcal{D}(\Omega)$, neither works: the topology must enforce both that all derivatives converge uniformly *and* that supports stay bounded. No single norm (nor even countably many seminorms) can do both. Each $\mathcal{D}_{\mathrm{K}}$ is Fréchet, with seminorms $\|\cdot\|_{\mathrm{K},k}$, but $\mathcal{D}(\Omega) = \bigcup_{\mathrm{K}} \mathcal{D}_{\mathrm{K}}$ is a union over all compact $\mathrm{K} \subset \Omega$, and no single compact set suffices. The inductive limit topology is precisely the tool that handles this union.

This is not an abstract curiosity; it has a clean sequential characterization.

4.7 PROPOSITION (Sequential convergence in). $\mathcal{D}(\Omega)$

A sequence $\varphi_n \to \varphi$ in $\mathcal{D}(\Omega)$ if and only if:

1. *There exists a compact set $\mathrm{K} \subset \Omega$ such that $\operatorname{supp}(\varphi_n) \subseteq \mathrm{K}$ and $\operatorname{supp}(\varphi) \subseteq \mathrm{K}$ for all n.*
2. *For every multi-index α, $\mathrm{D}^\alpha \varphi_n \to \mathrm{D}^\alpha \varphi$ uniformly on K.*

Proof. The backward direction is immediate: if both conditions hold then $\varphi_n \to \varphi$ in $\mathcal{D}_{\mathrm{K}}$, and since the inclusion $\iota_{\mathrm{K}} : \mathcal{D}_{\mathrm{K}} \hookrightarrow \mathcal{D}(\Omega)$ is continuous, convergence in $\mathcal{D}(\Omega)$ follows.

For the forward direction, suppose $\varphi_n \to \varphi$ in $\mathcal{D}(\Omega)$ but the supports are not contained in any single compact set. Then there exist $x_n \in \operatorname{supp}(\varphi_n - \varphi)$ with $|x_n| \to \infty$. One constructs a continuous seminorm p on $\mathcal{D}(\Omega)$ (using the inductive limit structure) that assigns large values near x_n, contradicting $p(\varphi_n - \varphi) \to 0$. Once the supports lie in a fixed K, convergence in $\mathcal{D}(\Omega)$ restricts to convergence in $\mathcal{D}_{\mathrm{K}}$, which is precisely uniform convergence of all derivatives on K.

□

This is a very **strong** notion of convergence. Both conditions must hold: derivatives of all orders must converge uniformly, *and* the supports cannot wander off to infinity. For comparison, convergence in L^2 only requires $\int |\varphi_n - \varphi|^2 \to 0$, which says nothing about derivatives and nothing about where the mass lives.

Not metrizable, but sequences suffice. A topological space is **first countable** if every point has a countable neighborhood basis (a sequence $\mathrm{U}_1 \supset \mathrm{U}_2 \supset \cdots$ such that every neighborhood contains

some U_n). Every metrizable space is first countable (take $U_n = B(x, 1/n)$). The space $\mathcal{D}(\Omega)$ is **not first countable**: a neighborhood of 0 must restrict to a neighborhood of 0 in every $\mathcal{D}_{K_j}$, and the "size" of these restrictions can be chosen independently, so no countable collection suffices. In general, sequential continuity does not imply continuity in spaces that are not first countable. However, the inductive limit structure gives a remarkable exception: for linear maps $T : \mathcal{D}(\Omega) \to X$ into a locally convex space, **continuity is equivalent to sequential continuity**. This is because T is continuous if and only if each restriction $T \circ \iota_j : \mathcal{D}_{K_j} \to X$ is continuous, and each $\mathcal{D}_{K_j}$ *is* metrizable. This is why we can work entirely with sequences throughout distribution theory.

4.5.4 Characterizing continuous functionals

The inductive limit topology determines exactly which linear functionals on $\mathcal{D}(\Omega)$ are continuous.

4.8 THEOREM (Continuity estimate). *A linear functional $T : \mathcal{D}(\Omega) \to \mathbb{R}$ is continuous if and only if: for every compact $K \subset \Omega$, there exist constants $C > 0$ and $k \geq 0$ such that*

$$|T(\varphi)| \leq C\|\varphi\|_{C^k(K)} \qquad \text{for all } \varphi \in \mathcal{D}_K(\Omega). \tag{187}$$

*Both C and k are allowed to **depend on** K; this is the crucial feature inherited from the inductive limit.*

Proof. By definition of the inductive limit topology, T is continuous on $\mathcal{D}(\Omega)$ if and only if each restriction $T|_{\mathcal{D}_K} : \mathcal{D}_K \to \mathbb{R}$ is continuous. Since $\mathcal{D}_K$ is a Fréchet space with seminorms $\|\cdot\|_{K,k}$, continuity of a linear functional on $\mathcal{D}_K$ is equivalent to boundedness by some finite combination of these seminorms, i.e. there exist $C > 0$ and $k \geq 0$ such that $|T(\varphi)| \leq C\|\varphi\|_{K,k}$ for all $\varphi \in \mathcal{D}_K$.

□

4.9 DEFINITION (Order of a distribution). The **order** of a distribution T is the smallest integer $k \geq 0$ such that for every compact $K \subset \Omega$, there exists $C > 0$ with

$$|T(\varphi)| \leq C\|\varphi\|_{C^k(K)} \qquad \text{for all } \varphi \in \mathcal{D}_K(\Omega). \tag{188}$$

A distribution of finite order k extends to a continuous functional on $C_c^k(\Omega)$. A distribution of infinite order genuinely requires test functions in C_c^∞.

4.5.5 What transfers from earlier chapters?

The spaces $\mathcal{D}_K$, $\mathcal{D}(\Omega)$, and $\mathcal{D}'(\Omega)$ are progressively further from Banach spaces:

Space	Type	Metrizable?	Complete?	Key property
$\mathcal{D}_K(\Omega)$	Fréchet space	Yes	Yes	Baire category holds
$\mathcal{D}(\Omega)$	Inductive limit of Fréchet spaces	**No**	Yes (sequentially)	Continuity = sequential continuity for linear maps
$\mathcal{D}'(\Omega)$	Dual of inductive limit, with weak-$*$ topology	**No**	Sequentially	Montel (every bounded set is relatively compact)

The key theorems from the linear operators chapter were proved using the Baire category theorem, which requires **completeness** and **metrizability**. Each $\mathcal{D}_K$ has both, so the theorems apply there. The full space $\mathcal{D}(\Omega)$ is not metrizable, so we cannot apply them directly, but we can apply them on each piece $\mathcal{D}_K$ and glue the results together.

4.10 REMARK (Which theorems transfer). **On each $\mathcal{D}_K$ (Fréchet space), everything works:**

- **Baire category theorem**: $\mathcal{D}_K$ is a complete metric space, so it is a Baire space.
- **Banach–Steinhaus** (uniform boundedness): a pointwise bounded family of continuous linear functionals on $\mathcal{D}_K$ is equicontinuous. This is the key tool for proving that $\mathcal{D}'$ is sequentially complete and has the Montel property.
- **Open mapping theorem**: a surjective continuous linear map between Fréchet spaces is open.
- **Closed graph theorem**: a closed linear map between Fréchet spaces is continuous.

On $\mathcal{D}(\Omega)$ (inductive limit), partial transfer:

- **Banach–Steinhaus**: does not apply directly (not metrizable), but the **fix** K**, then vary** K strategy works. See Remark **??** below.
- **Open mapping / closed graph**: versions exist for inductive limits (the Grothendieck–de Wilde theorem), but we will not need them.
- **Hahn–Banach**: works on $\mathcal{D}(\Omega)$ because it is locally convex. This guarantees enough continuous functionals to separate points (the geometric version only needs local convexity, not metrizability).

On $\mathcal{D}'(\Omega)$ (weak-$*$ dual), the duality chapter transfers:

- **Weak-$*$ convergence**: convergence of distributions is exactly the weak-$*$ convergence from the duality chapter.
- **Banach–Alaoglu**: does not apply directly (no norm, so no "unit ball"). It is replaced by the **Montel property**: every bounded set in $\mathcal{D}'$ is relatively sequentially compact. The proof uses Banach–Steinhaus on each $\mathcal{D}_K$ plus a diagonal argument, the same structure as Banach–Alaoglu for separable preduals.
- **Separation**: test functions separate distributions and distributions separate test functions (Hahn–Banach on $\mathcal{D}$).

4.11 REMARK (The "fix). K**, then vary** K**" strategy**

This is the fundamental technique for working with the inductive limit. Suppose $\{T_n\}$ is a sequence of distributions that is **pointwise bounded**: $\sup_n |\langle T_n, \varphi\rangle| < \infty$ for every $\varphi \in \mathcal{D}(\Omega)$. We want to conclude

something uniform.

Step 1: Fix K**.** Pick any compact $K \subset \Omega$. The restrictions $T_n|_{\mathcal{D}_K}$ are a family of continuous linear functionals on the Fréchet space $\mathcal{D}_K$. They are pointwise bounded (by assumption). Apply Banach–Steinhaus on $\mathcal{D}_K$: the family is equicontinuous, meaning there exist $C_K > 0$ and $k_K \geq 0$ such that

$$|\langle T_n, \varphi \rangle| \leq C_K \|\varphi\|_{C^{k_K}(K)} \quad \text{for all } n \text{ and all } \varphi \in \mathcal{D}_K. \tag{189}$$

Step 2: Vary K**.** Repeat for every compact set in an exhaustion $K_1 \subset K_2 \subset \cdots$ of Ω. For each K_j, we get constants C_j and orders k_j, and these will generally **differ** from one K_j to the next. That is fine: any test function φ has compact support, so $\varphi \in \mathcal{D}_{K_j}$ for some j, and the estimate for that K_j applies.

The result is not a single uniform bound (which would require a single norm), but a **family of bounds, one per compact set**. This is exactly the structure of the continuity estimate in Theorem 4.8: the constants C and k depend on K because Banach–Steinhaus is applied separately on each $\mathcal{D}_K$.

This two-step strategy is used throughout distribution theory:

- **Sequential completeness** of $\mathcal{D}'$: fix K, get equicontinuity, then extend convergence from $\mathcal{D}_K$ to all of $\mathcal{D}$.
- **Montel property**: fix K, get equicontinuity, then use a diagonal argument over $K_1 \subset K_2 \subset \cdots$ to extract a convergent subsequence.
- **Convergent sequences are bounded**: fix K, apply Banach–Steinhaus to get uniform order on each compact set.

The recurring pattern: **do Fréchet space arguments on each $\mathcal{D}_K$, then glue via the inductive limit.** The inductive limit structure is the price we pay for the compact support condition; the reward is a dual space ($\mathcal{D}'$) where differentiation always makes sense.

4.6 *Distributions as Functionals*

With the inductive limit topology on $\mathcal{D}(\Omega)$ in hand, we can now define distributions and equip them with a topology.

4.6.1 Definition of Distributions

4.12 DEFINITION (Distributions). A **distribution** on Ω is a continuous linear functional $T : \mathcal{D}(\Omega) \to \mathbb{R}$. The space of all distributions is

$$\mathcal{D}'(\Omega) = \left(\mathcal{D}(\Omega)\right)^*, \tag{190}$$

the topological dual of $\mathcal{D}(\Omega)$. We write $\langle T, \varphi\rangle = T(\varphi)$ for the pairing.

By Theorem 4.8, T is a distribution if and only if for every compact $K \subset \Omega$, there exist $C > 0$ and $k \geq 0$ such that $|T(\varphi)| \leq C\|\varphi\|_{C^k(K)}$ for all $\varphi \in \mathcal{D}_K(\Omega)$.

Physical interpretation. No physical measurement captures a true point value. A thermometer averages molecular kinetic energy over a small volume. An electric field probe averages over the test charge's spatial extent. Position measurements in quantum mechanics are limited by the uncertainty principle. The pairing $\langle T, \varphi\rangle$ captures this: the test function φ is the sensitivity profile of a measurement device, with $\operatorname{supp}(\varphi)$ as the detection region and the shape of φ encoding how strongly each point contributes to the reading. A distribution T is then an object that can only be "observed" through such averaged measurements, never through point evaluation. This is not a limitation but a feature: it is precisely why distributions can accommodate objects like the Dirac delta that have no pointwise values.

4.6.2 Convergence of Distributions

4.13 DEFINITION (Weak-). $*$ **topology on distributions**

The **weak-$*$ topology** on $\mathcal{D}'(\Omega)$ is the coarsest topology making every evaluation functional $\Phi_\varphi(T) = \langle T, \varphi\rangle$ continuous. A sequence of distributions T_n converges to T in this topology if and only if

$$\langle T_n, \varphi\rangle \to \langle T, \varphi\rangle \quad \text{for all } \varphi \in \mathcal{D}(\Omega). \tag{191}$$

This is **convergence in the sense of distributions**, one of the weakest convergence notions in analysis. Yet limits are unique: the topology is **Hausdorff** because if $T_1 \neq T_2$, there exists φ with $\langle T_1, \varphi\rangle \neq \langle T_2, \varphi\rangle$.

4.14 REMARK. The parallel with Banach space duality is exact:

Banach space setting	Distribution setting
Banach space X	Test functions $\mathcal{D}(\Omega)$
Dual space X^*	Distributions $\mathcal{D}'(\Omega)$
Norm topology on X	Inductive limit topology on $\mathcal{D}(\Omega)$
Weak-$*$ topology on X^*	Standard topology on $\mathcal{D}'(\Omega)$
$f(x)$ for $f \in X^*, x \in X$	$\langle T, \varphi \rangle$ for $T \in \mathcal{D}', \varphi \in \mathcal{D}$
Weak-$*$ convergence: $f_n(x) \to f(x)$	Distributional convergence: $T_n(\varphi) \to T(\varphi)$
Banach–Alaoglu: closed balls are weak-$*$ compact	Distributions are weak-$*$ sequentially complete

The key difference is that $\mathcal{D}(\Omega)$ is **not** a Banach space, so results like Banach–Alaoglu do not apply directly. Compactness and completeness results for distributions rely on the specific structure of the inductive limit. But the *language* and *intuition* transfer directly from the duality chapter.

4.6.3 The Hierarchy of Distributions

The examples below are organized by the inclusion chain

$$C(\Omega) \subset L^1_{\text{loc}}(\Omega) \subset \mathcal{M}(\Omega) \subset \mathcal{D}'(\Omega), \tag{192}$$

with each level strictly enlarging the previous one. To understand *why* each inclusion is strict, we introduce a unified probe: **concentrating test functions**. Fix a point x_0 and a non-negative test function φ with $\int \varphi = 1$, and define

$$\varphi_n(x) = n^d \, \varphi(n(x - x_0)). \tag{193}$$

Then $\int \varphi_n = 1$ and $\operatorname{supp}(\varphi_n) \to \{x_0\}$. In the physical picture, φ_n is a measurement device focusing on x_0. The growth rate of $\langle T, \varphi_n \rangle$ as $n \to \infty$ is the signature that separates the levels.

Continuous functions: $C(\Omega) \hookrightarrow \mathcal{D}'$ A continuous function f defines a distribution T_f by integration. Probing with concentrating test functions recovers the point value:

$$\langle T_f, \varphi_n \rangle = \int f(x) \, \varphi_n(x) \, dx \to f(x_0) \tag{194}$$

since φ_n concentrates at x_0 and f is continuous there. The response is **bounded**.

Regular distributions: $L^1_{\text{loc}} \hookrightarrow \mathcal{D}'$

4.15 definition (Regular distributions). A function $f \in L^1_{\text{loc}}(\Omega)$ defines a distribution T_f by

$$\langle T_f, \varphi \rangle = \int_\Omega f(x) \, \varphi(x) \, dx \qquad \text{for all } \varphi \in \mathcal{D}(\Omega). \tag{195}$$

Distributions of this form are called **regular**.

The map $f \mapsto \mathrm{T}_f$ is injective (two functions that agree a.e. define the same distribution), so we often identify f with T_f and write $\langle f, \varphi\rangle$. Any measure absolutely continuous with respect to Lebesgue measure is representable by a locally integrable function, so regular distributions are exactly those generated by absolutely continuous measures. The embedding $f \mapsto \mathrm{T}_f$ gives an injection $\mathrm{L}^1_{\mathrm{loc}}(\Omega) \hookrightarrow \mathcal{D}'(\Omega)$.

4.16 EXAMPLE (Regular distributions).

1. Any $f \in \mathrm{L}^p(\Omega)$ with $1 \le p \le \infty$ is locally integrable, hence defines a regular distribution.
2. The **Heaviside function** $\mathrm{H}(x) = \mathbf{1}_{[0,\infty)}(x)$ is in $\mathrm{L}^1_{\mathrm{loc}}(\mathbb{R})$, so $\langle \mathrm{H}, \varphi\rangle = \int_0^\infty \varphi(x)\,dx$ defines a regular distribution.
3. The **sign function** $\operatorname{sgn}(x)$ is in $\mathrm{L}^1_{\mathrm{loc}}(\mathbb{R})$ but not continuous. It defines a regular distribution of order 0.

Probing. For $f \in \mathrm{L}^1_{\mathrm{loc}}$, the response $\langle \mathrm{T}_f, \varphi_n\rangle \to f(x_0)$ at every **Lebesgue point** x_0 of f (almost everywhere). The response is still **bounded**, just as for continuous functions. The inclusion $\mathrm{C}(\Omega) \subset \mathrm{L}^1_{\mathrm{loc}}(\Omega)$ is strict (the Heaviside function is locally integrable but not continuous), but both levels produce bounded responses.

Distributions from measures: $\mathcal{M} \hookrightarrow \mathcal{D}'$

4.17 DEFINITION (Distributions from measures). Any Radon measure μ on Ω defines a distribution T_μ by

$$\langle \mathrm{T}_\mu, \varphi\rangle = \int_\Omega \varphi(x)\,\mathrm{d}\mu(x). \tag{196}$$

This generalizes regular distributions: $\mathrm{T}_f = \mathrm{T}_{\mu_f}$ where $\mu_f = f(x)\,\mathrm{d}x$. The class of Radon measures also includes singular measures, so $\mu = \mu_{\mathrm{ac}} + \mu_{\mathrm{sing}}$. The singular part includes point masses and more exotic objects such as the Cantor measure. The embedding $\mu \mapsto \mathrm{T}_\mu$ gives an injection $\mathcal{M}(\Omega) \hookrightarrow \mathcal{D}'(\Omega)$.

4.18 EXAMPLE (Dirac delta distribution). For any $x_0 \in \Omega$, the **Dirac delta distribution** δ_{x_0} is defined by

$$\langle \delta_{x_0}, \varphi\rangle = \varphi(x_0). \tag{197}$$

This has order 0, since $|\langle \delta_{x_0}, \varphi\rangle| = |\varphi(x_0)| \le \|\varphi\|_{\mathrm{C}^0(\mathrm{K})}$. The Dirac delta is a singular distribution: it is a Radon measure (a point mass) but not representable by any locally integrable function. So $\delta_{x_0} \in \mathcal{M}(\Omega) \subset \mathcal{D}'(\Omega)$ but $\delta_{x_0} \notin \mathrm{L}^1_{\mathrm{loc}}(\Omega)$. This shows that the inclusion $\mathrm{L}^1_{\mathrm{loc}} \subset \mathcal{M}$ is strict.

Probing. For δ_{x_0}: $\langle \delta_{x_0}, \varphi_n\rangle = \varphi_n(x_0) = n^d\,\varphi(0)$. The response grows as $\mathcal{O}(n^d)$. More generally, if μ has a point mass of size c at x_0, then $\langle \mathrm{T}_\mu, \varphi_n\rangle \sim c\,n^d$. This blowup is the signature of a singular measure: no locally integrable function can produce $\mathcal{O}(n^d)$ growth, because $\int f\,\varphi_n \to f(x_0)$ stays bounded.

4.19 REMARK (Physical space embeds into distribution space). The map $x \mapsto \delta_x$ embeds $\mathbb{R}^d$ into $\mathcal{D}'(\mathbb{R}^d)$. This embedding is:

- **Injective**: if $x \neq y$, choose φ with $\varphi(x) = 1$ and $\varphi(y) = 0$; then $\langle \delta_x, \varphi \rangle \neq \langle \delta_y, \varphi \rangle$.
- **Continuous**: if $x_n \to x$ in $\mathbb{R}^d$, then $\langle \delta_{x_n}, \varphi \rangle = \varphi(x_n) \to \varphi(x) = \langle \delta_x, \varphi \rangle$ for every test function φ, so $\delta_{x_n} \to \delta_x$ in $\mathcal{D}'$.
- **A homeomorphism onto its image**: the weak-$*$ topology restricted to $\{\delta_x : x \in \mathbb{R}^d\}$ recovers the Euclidean topology, because test functions separate points of $\mathbb{R}^d$.
- **Closed**: the image $\{\delta_x : x \in \mathbb{R}^d\}$ is a closed subset of $\mathcal{D}'$. If $\delta_{x_n} \to \mathrm{T}$ in $\mathcal{D}'$ and $x_n \to \infty$, then $\langle \mathrm{T}, \varphi \rangle = 0$ for every test function (since eventually $x_n \notin \operatorname{supp}(\varphi)$), so $\mathrm{T} = 0$, which is not a δ_x. Any convergent subsequence must have x_n bounded, hence convergent, so $\mathrm{T} = \delta_{x_\infty}$.

Physical space $\mathbb{R}^d$ thus sits inside $\mathcal{D}'(\mathbb{R}^d)$ as a closed d-dimensional submanifold. The passage from classical to distributional analysis replaces "a function assigns values to points" with "a distribution assigns values to test functions." The points themselves are still present, embedded as the simplest distributions.

DISTRIBUTIONS BEYOND MEASURES: $\mathcal{D}' \setminus \mathcal{M}$ The full power of distribution theory appears at this level: objects that cannot be represented by any measure. A Radon measure μ acts on a function through integration, so its action depends only on the *values* of the function, controlled by

$$|\langle \mu, \varphi \rangle| = \left| \int \varphi \, d\mu \right| \leq |\mu|(\operatorname{supp} \varphi) \cdot \sup |\varphi|. \tag{198}$$

The right-hand side involves only $\sup|\varphi|$, never any derivatives. A distribution that "sees" derivatives of test functions cannot be a measure.

4.20 EXAMPLE (Derivative of the Dirac delta). The derivative δ' is defined by

$$\langle \delta', \varphi \rangle = -\varphi'(0). \tag{199}$$

More generally, for any multi-index α:

$$\langle \mathrm{D}^\alpha \delta, \varphi \rangle = (-1)^{|\alpha|} \mathrm{D}^\alpha \varphi(0). \tag{200}$$

The distribution $\mathrm{D}^\alpha \delta$ has order $|\alpha|$. The action of δ' depends on the derivative $\varphi'(0)$, and no bound of the form $|\varphi'(0)| \leq C \sup|\varphi|$ can hold: take $\varphi_n(x) = \frac{1}{n}\psi(nx)$ for a fixed bump function ψ with $\psi'(0) = 1$, then $\sup|\varphi_n| \to 0$ but $\varphi_n'(0) = 1$. So no measure can reproduce δ', and $\delta' \in \mathcal{D}'(\Omega) \setminus \mathcal{M}(\Omega)$. This shows that the inclusion $\mathcal{M} \subset \mathcal{D}'$ is strict.

4.21 EXAMPLE (Principal value of). $1/x$

The distribution p.v. $\frac{1}{x}$ is defined by

$$\left\langle \text{p.v.} \frac{1}{x}, \varphi \right\rangle = \lim_{\varepsilon \to 0} \int_{|x| > \varepsilon} \frac{\varphi(x)}{x} \, dx. \tag{201}$$

This limit exists for every $\varphi \in \mathcal{D}(\mathbb{R})$ because the odd part of the integrand cancels the singularity at

the origin. The principal value is not a regular distribution ($1/x$ is not locally integrable), but it is also not a measure. The same scaling argument applies: take $\varphi_n(x) = \frac{1}{n}\psi(nx)$ where ψ is an odd bump with $\int \psi(x)/x\,dx = c \neq 0$. Then $\sup|\varphi_n| \to 0$, but by the substitution $u = nx$,

$$\left\langle \text{p.v.}\,\frac{1}{x},\ \varphi_n \right\rangle = \int \frac{\psi(u)}{u}\,du = c \neq 0. \tag{202}$$

So no order-0 estimate holds, and p.v. $\frac{1}{x} \in \mathcal{D}'(\mathbb{R}) \setminus \mathcal{M}(\mathbb{R})$.

Probing. For a distribution of order m, the response grows as $\langle \mathrm{T}, \varphi_n\rangle = \mathcal{O}(n^m)$. For δ': $\langle \delta', \varphi_n\rangle = -\varphi_n'(0) = -n^{d+1}\varphi'(0)$, giving growth of order n^{d+1}. Measures (order 0) produce at most $\mathcal{O}(n^d)$ growth, so the faster rate is the quantitative reason δ' cannot be a measure.

The growth rate as a signature The concentrating probe φ_n acts as a diagnostic: bounded response means a function, $\mathcal{O}(n^d)$ growth means a singular measure, and $\mathcal{O}(n^m)$ growth for $m > d$ means a distribution that is not a measure. This connects back to Proposition 4.22: test functions separate distributions because different distributions respond differently to these probes. The concentrating sequence is the concrete mechanism by which test functions "see" the difference between a function, a measure, and a genuinely distributional object.

4.6.4 Topological Properties of $\mathcal{D}'(\Omega)$

With the examples in hand, we return to the topology of $\mathcal{D}'(\Omega)$ and establish the structural results that make distribution theory work: separation, completeness, density, and the remarkable collapse of the weak-$*$ and strong topologies.

Separation The weak-$*$ topology is Hausdorff: test functions separate distributions. The converse (that distributions separate test functions) is less obvious and requires the Hahn-Banach theorem.

4.22 proposition (Distributions separate test functions). *If $\varphi_1 \neq \varphi_2$ in $\mathcal{D}(\Omega)$, there exists a distribution $\mathrm{T} \in \mathcal{D}'(\Omega)$ such that $\langle \mathrm{T}, \varphi_1\rangle \neq \langle \mathrm{T}, \varphi_2\rangle$.*

Proof. The version of Hahn-Banach from the duality chapter assumes a normed space, and $\mathcal{D}(\Omega)$ is not normed. What applies here is the **geometric Hahn-Banach theorem** (the separation theorem), which holds on any **locally convex** topological vector space. Since $\mathcal{D}(\Omega)$ is an inductive limit of Fréchet spaces, it is locally convex, and the theorem guarantees a continuous linear functional separating φ_1 from φ_2. □

Sequential completeness

4.23 proposition (Sequential completeness). *Every weak-$*$ Cauchy sequence in $\mathcal{D}'(\Omega)$ converges to a distribution. That is, if $\langle \mathrm{T}_n, \varphi\rangle$ converges for every $\varphi \in \mathcal{D}(\Omega)$, then the limit $\mathrm{T}(\varphi) = \lim_n \langle \mathrm{T}_n, \varphi\rangle$ defines a distribution $\mathrm{T} \in \mathcal{D}'(\Omega)$.*

Proof. The map $\mathrm{T} : \mathcal{D}(\Omega) \to \mathbb{R}$ defined by $\mathrm{T}(\varphi) = \lim_n \mathrm{T}_n(\varphi)$ is linear (limits preserve linearity). It remains to show continuity. Fix a compact $K \subset \Omega$. The restrictions $\mathrm{T}_n|_{\mathcal{D}_K}$ are a pointwise bounded family of continuous linear functionals on the Fréchet space $\mathcal{D}_K$. By Banach-Steinhaus, they are

equicontinuous: there exist $C_K > 0$ and $k_K \geq 0$ such that $|T_n(\varphi)| \leq C_K \|\varphi\|_{K,k_K}$ for all n and all $\varphi \in \mathcal{D}_K$. Passing to the limit, $|T(\varphi)| \leq C_K \|\varphi\|_{K,k_K}$, so T is continuous on each $\mathcal{D}_K$, hence on $\mathcal{D}(\Omega)$.

□

Sequential density

4.24 PROPOSITION (Sequential density). *Every distribution is the weak-∗ limit of a sequence of test functions. That is, for every* $T \in \mathcal{D}'(\Omega)$, *there exists a sequence* $\{\psi_n\} \subset \mathcal{D}(\Omega)$ *such that* $\langle T, \varphi \rangle = \lim_n \int \psi_n \varphi \, dx$ *for all* $\varphi \in \mathcal{D}(\Omega)$.

As a concrete instance: if $\{\rho_n\}$ is an approximation to the identity (non-negative, $\int \rho_n = 1$, support shrinking to $\{0\}$), then $\rho_n \to \delta$ in the sense of distributions, since

$$\langle \rho_n, \varphi \rangle = \int \rho_n(x)\, \varphi(x)\, dx \to \varphi(0) = \langle \delta, \varphi \rangle \tag{203}$$

for every test function φ.

The Montel property: weak-∗ equals strong

4.25 REMARK (The Montel property). A **Montel space** is a locally convex topological vector space in which every bounded closed subset is compact. This is an infinite-dimensional analogue of the Heine-Borel theorem: in $\mathbb{R}^n$, bounded closed sets are compact, but in an infinite-dimensional Banach space they never are (by Riesz's lemma, the closed unit ball is not compact). Montel spaces recover this compactness despite being infinite-dimensional.

Why $\mathcal{D}_K$ is Montel. A bounded set in $\mathcal{D}_K$ is one where all seminorms $\|\cdot\|_{K,k}$ are uniformly bounded, i.e. all derivatives up to every order are uniformly bounded on K. By the Arzelà-Ascoli theorem, uniform boundedness of a family of functions and their first derivatives on a compact set implies a subsequence converging in $C^0(K)$. But the same family also has uniformly bounded second derivatives, so we can extract a further subsequence converging in $C^1(K)$, and so on. A diagonal argument produces a subsequence converging in every $C^k(K)$ norm, hence converging in $\mathcal{D}_K$. Since $\mathcal{D}(\Omega)$ is an inductive limit of the Montel spaces $\mathcal{D}_K$, it is itself Montel.

Consequence for $\mathcal{D}'(\Omega)$. On $\mathcal{D}'(\Omega)$, there are two natural topologies:

- The **weak-∗ topology** $\sigma(\mathcal{D}', \mathcal{D})$: pointwise convergence on test functions.
- The **strong dual topology** $\beta(\mathcal{D}', \mathcal{D})$: uniform convergence on bounded subsets of $\mathcal{D}(\Omega)$.

In a Banach space, these are genuinely different. The Montel property collapses the gap in four steps:

1. **Pointwise boundedness.** A weak-∗ convergent sequence $T_n \to T$ is pointwise bounded: $\sup_n |\langle T_n, \varphi \rangle| < \infty$ for each φ.

2. **Equicontinuity (Banach-Steinhaus on Fréchet spaces).** The key step. Fix a compact $K \subset \Omega$. The restrictions $T_n|_{\mathcal{D}_K}$ are a pointwise bounded family of continuous linear functionals on $\mathcal{D}_K$. This is where the Fréchet structure of $\mathcal{D}_K$ is essential: because $\mathcal{D}_K$ is a complete metrizable space (Definition **??**), the Baire category theorem applies, and Banach-Steinhaus gives equicontinuity – there exist $C_K > 0$ and $k_K \geq 0$ such that $|T_n(\varphi)| \leq C_K \|\varphi\|_{K,k_K}$ for *all* n and all $\varphi \in \mathcal{D}_K$, simultaneously.

(On $\mathcal{D}(\Omega)$ itself, which is not metrizable, Banach-Steinhaus does not apply directly; this is why we must descend to the Fréchet pieces $\mathcal{D}_K$ first.)

3. **Uniform convergence on precompact sets.** Equicontinuity + pointwise convergence $\Rightarrow$ uniform convergence on every precompact subset of $\mathcal{D}(\Omega)$. This is the Arzelà-Ascoli principle for functionals: an equicontinuous family that converges pointwise converges uniformly on any set where the arguments are "close together" (precompact sets are exactly such sets).

4. **Montel: bounded = precompact.** In a Banach space, bounded sets are generally *not* precompact (the unit ball is never compact in infinite dimensions). The Montel property changes this: every bounded set in $\mathcal{D}(\Omega)$ *is* precompact. So uniform convergence on precompact sets = uniform convergence on bounded sets = strong convergence.

Without Montel, step 4 fails. This is exactly what goes wrong in Banach spaces, where weak-$*$ and strong convergence are genuinely different. The chain of reasoning is:

$$\text{weak-}* \text{ conv.} \xrightarrow{1} \text{ptwise bounded} \xrightarrow{\text{2, Fréchet}} \text{equicontinuous} \xrightarrow{3} \text{unif. on precompact} \xrightarrow{\text{4, Montel}} \text{strong conv.} \tag{204}$$

Therefore **on $\mathcal{D}'(\Omega)$, the weak-$*$ and strong topologies have the same convergent sequences.** This is why many references simply say "the topology on $\mathcal{D}'$" without specifying which.

THE DOUBLE DUAL AND NON-REFLEXIVITY In the duality chapter, the canonical embedding $J : X \to X^{**}$, $J(x)(f) = f(x)$, embeds a Banach space into its double dual. A space is **reflexive** if J is surjective. What happens for test functions and distributions?

The canonical embedding $J : \mathcal{D}(\Omega) \to \mathcal{D}'(\Omega)^*$ sends each test function φ to the evaluation functional $J(\varphi)(T) = \langle T, \varphi\rangle$. This embedding is:

- **Injective**: if $J(\varphi_1) = J(\varphi_2)$, then $\langle T, \varphi_1\rangle = \langle T, \varphi_2\rangle$ for all distributions T. By Proposition 4.22, this forces $\varphi_1 = \varphi_2$.
- **Not surjective**: there exist continuous linear functionals on $\mathcal{D}'(\Omega)$ that do not come from any test function.

4.26 PROPOSITION (is not reflexive). $\mathcal{D}(\Omega)$

The canonical embedding $J : \mathcal{D}(\Omega) \to \mathcal{D}'(\Omega)^*$ is not surjective.

Proof. We construct an element of the double dual that is not in the image of J.

Let $\{\chi_n\}$ be a sequence of cutoff functions: $\chi_n \in \mathcal{D}(\Omega)$ with $0 \le \chi_n \le 1$, $\chi_n(x) = 1$ for $x \in K_n$, where $K_1 \subset K_2 \subset \cdots$ is an exhaustion of Ω by compact sets. Such cutoffs exist by Urysohn's lemma (or by explicit mollification of characteristic functions).

For any distribution $T \in \mathcal{D}'(\Omega)$ and any test function φ, the sequence $\langle T, \chi_n\varphi\rangle \to \langle T, \varphi\rangle$ since $\chi_n\varphi \to \varphi$ in $\mathcal{D}(\Omega)$.

Now define $\Lambda : \mathcal{D}'(\Omega) \to \mathbb{R}$ by $\Lambda(T) = \langle T, \mathbf{1}\rangle$, where the pairing is understood as $\lim_n \langle T, \chi_n\rangle$ whenever this limit exists. For every $T \in \mathcal{D}'(\Omega)$, the sequence $\langle T, \chi_n\rangle$ is eventually constant for distributions with compact support, and converges for tempered distributions.

On the subspace of compactly supported distributions $\mathcal{E}'(\Omega)$, Λ is a well-defined continuous linear functional. But Λ cannot equal $J(\varphi)$ for any $\varphi \in \mathcal{D}(\Omega)$: that would require $\langle T, \varphi\rangle = \langle T, \mathbf{1}\rangle$ for all

compactly supported T, hence $\varphi = \mathbf{1}$. But $\mathbf{1} \notin \mathcal{D}(\Omega)$, since the constant function 1 does not have compact support.

□

4.27 REMARK. Non-reflexivity means that the bidual $\mathcal{D}'(\Omega)^*$ is strictly larger than $\mathcal{D}(\Omega)$: there are continuous linear functionals on the space of distributions that cannot be represented by pairing with any test function. The constant function $\mathbf{1}$ is the simplest example: it "pairs" with every compactly supported distribution, but it is not itself a test function.

Note that this is *not* a statement about weak-$*$ versus strong convergence; we showed above (Remark 4.25) that the Montel property makes those agree on $\mathcal{D}'$. Rather, non-reflexivity is about the *size* of the dual: the space of test functions is too small to exhaust all continuous functionals on $\mathcal{D}'$.

Note the role of **Urysohn's lemma** (or equivalently, partitions of unity): it guarantees the existence of the cutoff functions χ_n that approximate $\mathbf{1}$ pointwise. The argument shows that the "function" $\mathbf{1}$ lives in the bidual but not in $\mathcal{D}$. It is a phantom test function that distributions can see but test functions cannot reach.

4.7 *The Calculus of Distributions*

Big Idea

Every operation on distributions is defined by **duality**: to apply an operation to a distribution T, transfer the adjoint operation to the test function. Differentiation becomes the integration-by-parts formula, and the payoff is immediate: every distribution is infinitely differentiable.

The key idea throughout this section is the same: if an operation A on smooth functions satisfies $\langle Af, \varphi\rangle = \langle f, A^\dagger\varphi\rangle$ for some adjoint $A^\dagger$ that maps test functions to test functions, then we *define* $\langle AT, \varphi\rangle = \langle T, A^\dagger\varphi\rangle$ for any distribution T.

A note on notation. We have a functional $T : \mathcal{D}(\Omega) \to \mathbb{R}$, and we want to define a new functional AT. The definition always acts on the test function: AT is the composition $T \circ A^\dagger$, and T itself is unchanged. Concretely:

- $D^\alpha T = (-1)^{|\alpha|}\, T \circ D^\alpha$, i.e. $\varphi \mapsto (-1)^{|\alpha|} T(D^\alpha \varphi)$,
- $f\,T = T \circ M_f$, i.e. $\varphi \mapsto T(f\varphi)$, where M_f is the multiplication operator,
- $\tau_h T = T \circ \tau_{-h}$, i.e. $\varphi \mapsto T(\tau_{-h}\varphi)$.

In each case the notation ($D^\alpha T$, $f\,T$, $\tau_h T$) suggests we are acting on T, but the operation is transferred to φ via the adjoint. This is the adjoint construction from the duality chapter (Definition 3.75) applied concretely: if $A : \mathcal{D}(\Omega) \to \mathcal{D}(\Omega)$ is a continuous linear map, its adjoint $A^* : \mathcal{D}'(\Omega) \to \mathcal{D}'(\Omega)$ acts by $A^*T = T \circ A$. Differentiation, multiplication, and translation on distributions are all instances of this. The notation is justified by consistency: when $T = T_f$ for a classical function f, the distribution AT agrees with the distribution T_{Af} defined by applying A to f directly.

4.7.1 Multiplication by Smooth Functions

Given $f \in C^\infty(\Omega)$ and $T \in \mathcal{D}'(\Omega)$, we want to define a new distribution that deserves the name "$f\,T$."

4.28 DEFINITION (Multiplication by a smooth function). Let $f \in C^\infty(\Omega)$ and $T \in \mathcal{D}'(\Omega)$. Define $U : \mathcal{D}(\Omega) \to \mathbb{R}$ by

$$U(\varphi) = T(f\varphi). \tag{205}$$

We write $f\,T := U$ and call it the **product** of f and T.

Note that U is well-defined as a map: $f\varphi \in \mathcal{D}(\Omega)$ whenever $\varphi \in \mathcal{D}(\Omega)$, because the product of a smooth function and a compactly supported smooth function is again compactly supported and smooth. So $T(f\varphi)$ makes sense.

4.29 PROPOSITION (Multiplication produces a distribution). $f\,T \in \mathcal{D}'(\Omega)$.

Proof. **Linearity.**

$$U(a\varphi + b\psi) = T(f(a\varphi + b\psi)) = T(af\varphi + bf\psi) = a\,T(f\varphi) + b\,T(f\psi) = a\,U(\varphi) + b\,U(\psi). \tag{206}$$

Continuity. Suppose $\varphi_n \to \varphi$ in $\mathcal{D}(\Omega)$ (supports in a fixed compact K, all derivatives converge uniformly). Then $f\varphi_n \to f\varphi$ in $\mathcal{D}(\Omega)$: the supports remain in K, and the Leibniz rule gives

$$\mathrm{D}^\alpha(f\varphi_n) = \sum_{\beta \le \alpha} \binom{\alpha}{\beta} \mathrm{D}^\beta f \cdot \mathrm{D}^{\alpha-\beta}\varphi_n \to \sum_{\beta \le \alpha} \binom{\alpha}{\beta} \mathrm{D}^\beta f \cdot \mathrm{D}^{\alpha-\beta}\varphi = \mathrm{D}^\alpha(f\varphi) \tag{207}$$

uniformly on K, since f and all its derivatives are bounded on K. Since T is continuous,

$$\mathrm{U}(\varphi_n) = \mathrm{T}(f\varphi_n) \to \mathrm{T}(f\varphi) = \mathrm{U}(\varphi). \tag{208}$$

□

4.30 REMARK. **Consistency with functions.** When $\mathrm{T} = \mathrm{T}_g$ for $g \in \mathrm{L}^1_{\mathrm{loc}}(\Omega)$, the distribution $f\mathrm{T}_g$ agrees with T_{fg}:

$$(f\mathrm{T}_g)(\varphi) = \mathrm{T}_g(f\varphi) = \int_\Omega g(x)\, f(x)\, \varphi(x)\, dx = \mathrm{T}_{fg}(\varphi). \tag{209}$$

So the distributional product extends pointwise multiplication of functions.

4.31 EXAMPLE (Multiplying the Dirac delta). For $f \in \mathrm{C}^\infty(\Omega)$: $(f\delta)(\varphi) = \delta(f\varphi) = f(0)\varphi(0) = f(0)\,\delta(\varphi)$. So $f\delta = f(0)\delta$: the delta "evaluates" the smooth factor at its support.

More generally, $f\delta_{x_0} = f(x_0)\delta_{x_0}$.

4.32 REMARK. We cannot multiply two arbitrary distributions. The construction above requires f to be smooth (or at least sufficiently regular) so that $\varphi \mapsto f\varphi$ maps $\mathcal{D}(\Omega)$ to itself. The difficulty of multiplying distributions is a fundamental limitation: it is related to the problem of renormalization in quantum field theory and was one of the motivations for Colombeau's theory of generalized functions.

4.7.2 Differentiation of Distributions

With the duality principle now familiar from translation and multiplication, we come to the main payoff: differentiation.

THE DEFINITION The definition is motivated by integration by parts. Suppose f is smooth. Then for every $\varphi \in \mathcal{D}(\Omega)$,

$$\int_\Omega f'(x)\, \varphi(x)\, dx = -\int_\Omega f(x)\, \varphi'(x)\, dx \tag{210}$$

since the boundary terms vanish (test functions have compact support in Ω). In the language of distributions: $\mathrm{T}_{f'}(\varphi) = -\mathrm{T}_f(\varphi')$. The right-hand side makes sense for *any* distribution T, not just those coming from smooth functions: given $\mathrm{T} \in \mathcal{D}'(\Omega)$, the map $\varphi \mapsto -\mathrm{T}(\varphi')$ is perfectly well-defined. This suggests a definition.

4.33 DEFINITION (Distributional derivative). Let $T \in \mathcal{D}'(\Omega)$ and α a multi-index. Define $U : \mathcal{D}(\Omega) \to \mathbb{R}$ by

$$U(\varphi) = (-1)^{|\alpha|}\, T(D^\alpha \varphi). \tag{211}$$

We write $D^\alpha T := U$ and call it the **distributional derivative** of T.

4.34 THEOREM (Every distribution is infinitely differentiable). *For any* $T \in \mathcal{D}'(\Omega)$ *and any multi-index* α, $D^\alpha T \in \mathcal{D}'(\Omega)$. *In particular, every distribution can be differentiated infinitely many times.*

Proof. Set $U = D^\alpha T$, i.e. $U(\varphi) = (-1)^{|\alpha|}\, T(D^\alpha \varphi)$.

Linearity.

$$U(a\varphi + b\psi) = (-1)^{|\alpha|}\, T(D^\alpha(a\varphi + b\psi)) = (-1)^{|\alpha|}\, T(a\, D^\alpha \varphi + b\, D^\alpha \psi) = a\, U(\varphi) + b\, U(\psi). \tag{212}$$

Continuity. If $\varphi_n \to \varphi$ in $\mathcal{D}(\Omega)$ (supports in a fixed compact K, all derivatives converge uniformly), then $D^\alpha \varphi_n \to D^\alpha \varphi$ in $\mathcal{D}(\Omega)$ (same compact K, all derivatives still converge uniformly — differentiating does not enlarge supports). Since T is continuous,

$$U(\varphi_n) = (-1)^{|\alpha|}\, T(D^\alpha \varphi_n) \to (-1)^{|\alpha|}\, T(D^\alpha \varphi) = U(\varphi). \tag{213}$$

Since U is linear and continuous, $U \in \mathcal{D}'(\Omega)$. As T and α were arbitrary, we can differentiate again: $D^\beta(D^\alpha T)$ is again a distribution, and so on indefinitely.

□

4.35 REMARK. **Consistency with functions.** When $T = T_f$ for $f \in C^{|\alpha|}(\Omega)$, integration by parts gives

$$(D^\alpha T_f)(\varphi) = (-1)^{|\alpha|}\, T_f(D^\alpha \varphi) = (-1)^{|\alpha|} \int_\Omega f\, D^\alpha \varphi \, dx = \int_\Omega D^\alpha f\, \varphi \, dx = T_{D^\alpha f}(\varphi). \tag{214}$$

So the distributional derivative extends the classical one.

This is the main payoff of the distributional framework: differentiation is always possible, and it always produces another distribution.

THE DIFFERENTIATION CASCADE The power of distributional derivatives is best seen through a chain of examples in one dimension ($\Omega = \mathbb{R}$).

4.36 EXAMPLE (Heaviside to Dirac). The **Heaviside function** $H(x) = \begin{cases} 1 & x > 0 \\ 0 & x < 0 \end{cases}$ is locally integrable and defines a regular distribution. Its distributional derivative is the Dirac delta:

$$\langle \mathrm{H}', \varphi \rangle = -\langle \mathrm{H}, \varphi' \rangle = -\int_0^\infty \varphi'(x)\,dx = \varphi(0) = \langle \delta, \varphi \rangle. \tag{215}$$

So $\mathrm{H}' = \delta$ in $\mathcal{D}'(\mathbb{R})$.

4.37 EXAMPLE (The). $|x|$ **cascade**

The function $f(x) = |x|$ is continuous but not differentiable at the origin. Its distributional derivative is

$$\langle f', \varphi \rangle = -\int_{-\infty}^{\infty} |x|\,\varphi'(x)\,dx = -\int_{-\infty}^{0} (-x)\varphi'(x)\,dx - \int_0^\infty x\,\varphi'(x)\,dx. \tag{216}$$

Integrating by parts on each half-line (boundary terms at $\pm\infty$ vanish because φ has compact support; at 0 the contributions from both sides cancel):

$$= -\int_{-\infty}^{0} \varphi(x)\,dx + \int_0^\infty \varphi(x)\,dx = \int_{-\infty}^{\infty} \operatorname{sgn}(x)\,\varphi(x)\,dx. \tag{217}$$

So $|x|' = \operatorname{sgn}(x)$ as distributions. Differentiating once more:

$$\langle \operatorname{sgn}', \varphi \rangle = -\int_{-\infty}^{\infty} \operatorname{sgn}(x)\,\varphi'(x)\,dx = \int_{-\infty}^{0} \varphi'(x)\,dx - \int_0^\infty \varphi'(x)\,dx = 2\varphi(0) = \langle 2\delta, \varphi \rangle. \tag{218}$$

The full cascade is:

$$|x| \xrightarrow{\mathrm{D}} \operatorname{sgn}(x) \xrightarrow{\mathrm{D}} 2\delta_0 \xrightarrow{\mathrm{D}} 2\delta_0' \xrightarrow{\mathrm{D}} \cdots \tag{219}$$

Each step is well-defined as a distributional derivative. The first step produces a discontinuous function, the second a measure, and the third an object that is not even a measure, but all are distributions.

4.38 EXAMPLE (Differentiating). $\log|x|$

The function $\log|x|$ is locally integrable on $\mathbb{R}$ (the singularity at the origin is integrable). Its distributional derivative is the principal value distribution:

$$(\log|x|)' = \text{p.v.}\,\frac{1}{x} \tag{220}$$

in $\mathcal{D}'(\mathbb{R})$. This can be verified by direct computation:

$$\langle (\log|x|)', \varphi\rangle = -\int_{-\infty}^{\infty} \log|x|\,\varphi'(x)\,dx = \lim_{\varepsilon\to 0}\int_{|x|>\varepsilon} \frac{\varphi(x)}{x}\,dx \tag{221}$$

where the last equality follows from integration by parts on $(-\infty,-\varepsilon)$ and (ε,∞) and observing that the boundary terms $\log\varepsilon\,[\varphi(\varepsilon)-\varphi(-\varepsilon)]\to 0$ since φ is smooth.

Properties of distributional differentiation

4.39 PROPOSITION (Basic properties). *Let* $S, T \in \mathcal{D}'(\Omega)$, α, β *multi-indices,* $c \in \mathbb{R}$.

1. ***Linearity:*** $D^\alpha(cS+T) = cD^\alpha S + D^\alpha T$.
2. ***Commutativity of mixed partials:*** $D^\alpha D^\beta T = D^\beta D^\alpha T = D^{\alpha+\beta}T$.
3. ***Consistency:*** *If* $f \in C^{|\alpha|}(\Omega)$, *the distributional derivative* $D^\alpha T_f$ *agrees with the classical derivative* $T_{D^\alpha f}$.
4. ***Continuity:*** *If* $T_n \to T$ *in* $\mathcal{D}'(\Omega)$, *then* $D^\alpha T_n \to D^\alpha T$ *in* $\mathcal{D}'(\Omega)$.

Proof. All four follow directly from the definition $\langle D^\alpha T, \varphi\rangle = (-1)^{|\alpha|}\langle T, D^\alpha\varphi\rangle$.

1. Linearity of T in the pairing.
2. $\langle D^\alpha D^\beta T, \varphi\rangle = (-1)^{|\alpha|}\langle D^\beta T, D^\alpha\varphi\rangle = (-1)^{|\alpha|+|\beta|}\langle T, D^{\alpha+\beta}\varphi\rangle$, which is symmetric in α,β since partial derivatives of smooth functions commute.
3. For $f \in C^{|\alpha|}$, integration by parts gives $(-1)^{|\alpha|}\int f\,D^\alpha\varphi = \int (D^\alpha f)\,\varphi$.
4. If $\langle T_n, \psi\rangle \to \langle T, \psi\rangle$ for all $\psi \in \mathcal{D}$, then in particular for $\psi = D^\alpha\varphi$: $\langle D^\alpha T_n, \varphi\rangle = (-1)^{|\alpha|}\langle T_n, D^\alpha\varphi\rangle \to (-1)^{|\alpha|}\langle T, D^\alpha\varphi\rangle = \langle D^\alpha T, \varphi\rangle$.

□

Property 4 is remarkable: **you can always interchange limits and derivatives in the sense of distributions.** In classical analysis, this requires uniform convergence of derivatives. In the distributional framework, pointwise convergence of the functionals is enough, because the derivative is transferred to the (fixed, smooth) test function.

4.40 PROPOSITION (Leibniz rule for distributions). *For* $f \in C^\infty(\Omega)$ *and* $T \in \mathcal{D}'(\Omega)$*:*

$$D(fT) = f'T + fT' \tag{222}$$

and more generally, for any multi-index α*:*

$$D^\alpha(fT) = \sum_{\beta\le\alpha}\binom{\alpha}{\beta} D^\beta f \cdot D^{\alpha-\beta}T. \tag{223}$$

Proof. For the first-order case in one dimension:

$$\langle \mathrm{D}(f\,\mathrm{T}), \varphi\rangle = -\langle f\,\mathrm{T}, \varphi'\rangle = -\langle \mathrm{T}, f\,\varphi'\rangle. \tag{224}$$

Now $f\,\varphi' = (f\,\varphi)' - f'\varphi$, so

$$= -\langle \mathrm{T}, (f\,\varphi)'\rangle + \langle \mathrm{T}, f'\varphi\rangle = \langle \mathrm{T}', f\,\varphi\rangle + \langle f'\mathrm{T}, \varphi\rangle = \langle f\,\mathrm{T}', \varphi\rangle + \langle f'\mathrm{T}, \varphi\rangle. \tag{225}$$

The general case follows by induction on $|\alpha|$.

□

TAYLOR'S THEOREM FOR DISTRIBUTIONS Now that we have both translation and differentiation, we can combine them into a distributional Taylor expansion. The idea is simple: to expand $\tau_h \mathrm{T}$ in powers of h, apply $\tau_h \mathrm{T}$ to a test function and use the classical Taylor expansion of $\varphi(x+h)$ in h.

4.41 THEOREM (Taylor's theorem for distributions). *Let* $\mathrm{T} \in \mathcal{D}'(\mathbb{R}^d)$ *and* $h \in \mathbb{R}^d$. *Then for every* $\mathrm{N} \geq 0$,

$$\tau_h \mathrm{T} = \sum_{|\alpha| \leq \mathrm{N}} \frac{(-h)^\alpha}{\alpha!}\, \mathrm{D}^\alpha \mathrm{T} + \mathrm{R}_\mathrm{N} \tag{226}$$

where the remainder $\mathrm{R}_\mathrm{N} \in \mathcal{D}'(\mathbb{R}^d)$ *satisfies*

$$\mathrm{R}_\mathrm{N}(\varphi) = \mathrm{T}\left(\sum_{|\alpha| = \mathrm{N}+1} \frac{h^\alpha}{\alpha!} \int_0^1 (1-t)^\mathrm{N}\, (\mathrm{D}^\alpha \varphi)(\cdot + th)\, (\mathrm{N}+1)\, dt \right) \tag{227}$$

for all $\varphi \in \mathcal{D}(\mathbb{R}^d)$.

Proof. By definition, $(\tau_h \mathrm{T})(\varphi) = \mathrm{T}(\tau_{-h}\varphi) = \mathrm{T}(\varphi(\cdot + h))$. Since φ is smooth, we apply the classical Taylor expansion in the variable h: for each fixed x,

$$\varphi(x+h) = \sum_{|\alpha| \leq \mathrm{N}} \frac{h^\alpha}{\alpha!}\, \mathrm{D}^\alpha \varphi(x) + \sum_{|\alpha| = \mathrm{N}+1} \frac{h^\alpha}{\alpha!} \int_0^1 (1-t)^\mathrm{N}\, (\mathrm{D}^\alpha \varphi)(x+th)\, (\mathrm{N}+1)\, dt. \tag{228}$$

This expansion, and all its x-derivatives, converge uniformly on compact sets (the remainder and its derivatives are controlled by finitely many derivatives of φ on a compact set containing $\operatorname{supp}(\varphi) + [0,1] \cdot h$). So the expansion holds in $\mathcal{D}(\mathbb{R}^d)$. Applying T:

$$(\tau_h \mathrm{T})(\varphi) = \sum_{|\alpha| \leq \mathrm{N}} \frac{h^\alpha}{\alpha!}\, \mathrm{T}(\mathrm{D}^\alpha \varphi) + \mathrm{T}(\text{remainder}). \tag{229}$$

Now $\mathrm{T}(\mathrm{D}^\alpha \varphi) = (-1)^{|\alpha|} (\mathrm{D}^\alpha \mathrm{T})(\varphi)$, so

$$(\tau_h \mathrm{T})(\varphi) = \sum_{|\alpha| \le \mathrm{N}} \frac{(-h)^\alpha}{\alpha!} (\mathrm{D}^\alpha \mathrm{T})(\varphi) + \mathrm{R}_\mathrm{N}(\varphi). \tag{230}$$

□

4.42 EXAMPLE (Taylor expansion of). δ_h

Taking $\mathrm{T} = \delta$ and $d = 1$: since $\tau_h \delta = \delta_h$ (Example 4.51), the Taylor formula gives

$$\delta_h = \sum_{n=0}^{\mathrm{N}} \frac{(-h)^n}{n!} \delta^{(n)} + \mathrm{R}_\mathrm{N}. \tag{231}$$

Let us verify this directly. Applying both sides to a test function φ:

$$\varphi(h) = \sum_{n=0}^{\mathrm{N}} \frac{(-h)^n}{n!} \delta^{(n)}(\varphi) + \mathrm{R}_\mathrm{N}(\varphi) = \sum_{n=0}^{\mathrm{N}} \frac{(-h)^n}{n!} (-1)^n \varphi^{(n)}(0) + \mathrm{R}_\mathrm{N}(\varphi) = \sum_{n=0}^{\mathrm{N}} \frac{h^n}{n!} \varphi^{(n)}(0) + \mathrm{R}_\mathrm{N}(\varphi). \tag{232}$$

This is exactly the classical Taylor expansion of φ at 0, evaluated at h. The distributional Taylor theorem for δ is nothing but the classical Taylor theorem for the test function, read through the duality.

4.7.3 Weak Derivatives

Distributional derivatives exist for every distribution, but they may not be representable by a function. When they are, we get the notion of a **weak derivative**, which bridges distribution theory and Sobolev spaces.

4.43 DEFINITION (Weak derivative). Let $u \in \mathrm{L}^1_{\mathrm{loc}}(\Omega)$. We say u has a **weak derivative** $\mathrm{D}^\alpha u = v$ if there exists $v \in \mathrm{L}^1_{\mathrm{loc}}(\Omega)$ such that

$$\int_\Omega u \, \mathrm{D}^\alpha \varphi \, dx = (-1)^{|\alpha|} \int_\Omega v \, \varphi \, dx \qquad \text{for all } \varphi \in \mathcal{D}(\Omega). \tag{233}$$

When it exists, the weak derivative is unique (up to a.e. equivalence).

In the language of adjoints: the distributional derivative $\mathrm{D}^\alpha \mathrm{T}_u$ is the functional $\varphi \mapsto (-1)^{|\alpha|} \int u \, \mathrm{D}^\alpha \varphi \, dx$. This always defines an element of $\mathcal{D}'(\Omega)$. The question is whether this functional can be *represented* by a function v via $\varphi \mapsto \int v \, \varphi \, dx$. When it can, $v = \mathrm{D}^\alpha u$ is the weak derivative of u. If moreover $u \in \mathrm{L}^p$ and $\mathrm{D}^\alpha u \in \mathrm{L}^p$, then u belongs to a **Sobolev space** $\mathrm{W}^{k,p}(\Omega)$: the space of L^p functions whose weak derivatives up to order k are also in L^p. Sobolev spaces are the natural home for PDE solutions that are not classically differentiable, and their theory rests entirely on the distinction between distributional and weak derivatives.

4.44 EXAMPLE (has a weak derivative). $|x|$

By Example 4.37, the distributional derivative of $|x|$ is $\operatorname{sgn}(x)$. Since $\operatorname{sgn} \in L^1_{\mathrm{loc}}(\mathbb{R})$, this is a genuine weak derivative: $|x|' = \operatorname{sgn}(x)$ weakly.

Note that $|x|$ is not classically differentiable at the origin, but the single point $\{0\}$ has measure zero and does not affect the L^1_{loc} function $\operatorname{sgn}(x)$.

4.45 EXAMPLE (The Heaviside function has no weak derivative in). L^1_{loc}

By Example 4.36, $H' = \delta$ as distributions. But δ is not a regular distribution: there is no $v \in L^1_{\mathrm{loc}}(\mathbb{R})$ with $\int v\,\varphi\,dx = \varphi(0)$ for all test functions φ. So H does **not** have a weak derivative.

This is why $H \in L^p(\Omega)$ for any bounded Ω but $H \notin W^{1,p}(\Omega)$ for any p: membership in the Sobolev space requires the weak derivative to exist *as a function in* L^p.

4.46 EXAMPLE (Characteristic function of an interval). The function $u = \mathbf{1}_{(a,b)}$ has distributional derivative $u' = \delta_a - \delta_b$ (a difference of point masses). This is a measure but not in L^1_{loc}, so u has no weak derivative. Like the Heaviside function, $u \in L^p$ but $u \notin W^{1,p}$ for any p.

4.47 REMARK. The weak derivative extends the classical derivative in the following sense:

- If f is classically differentiable (or even just absolutely continuous), its weak derivative exists and agrees with the classical derivative a.e.
- Weak differentiability allows corners and kinks (finitely many, or even countably many, as long as the derivative remains locally integrable).
- Weak differentiability does **not** allow jumps: a jump discontinuity produces a delta function in the derivative, which is not in L^1_{loc}.

The dividing line is **absolute continuity**: a function u on an interval has a weak derivative in L^1_{loc} if and only if u is (equivalent to) an absolutely continuous function.

4.7.4 Translation and Reflection

The simplest operations on distributions illustrate the duality principle with no calculus required.

Recall that for functions on $\mathbb{R}^d$, translation by $h \in \mathbb{R}^d$ is $(\tau_h\varphi)(x) = \varphi(x - h)$ and reflection is $\check{\varphi}(x) = \varphi(-x)$. Given a distribution $T \in \mathcal{D}'(\mathbb{R}^d)$, we want to define the translated distribution $\tau_h T$ and the reflected distribution $\check{T}$.

4.48 DEFINITION (Translation and reflection of distributions). Let $T \in \mathcal{D}'(\mathbb{R}^d)$.

1. Define $U_h : \mathcal{D}(\mathbb{R}^d) \to \mathbb{R}$ by $U_h(\varphi) = T(\tau_{-h}\varphi)$. We write $\tau_h T := U_h$ and call it the **translation** of T by h.

2. Define $V : \mathcal{D}(\mathbb{R}^d) \to \mathbb{R}$ by $V(\varphi) = T(\check{\varphi})$. We write $\check{T} := V$ and call it the **reflection** of T.

4.49 PROPOSITION (Translation and reflection produce distributions). $\tau_h T$ *and* $\check{T}$ *are distributions, i.e. they belong to* $\mathcal{D}'(\mathbb{R}^d)$.

Proof. We verify the two requirements for $U_h = \tau_h T$; the argument for $V = \check{T}$ is identical.

Linearity.

$$U_h(a\varphi + b\psi) = T(\tau_{-h}(a\varphi + b\psi)) = T(a\,\tau_{-h}\varphi + b\,\tau_{-h}\psi) = a\,T(\tau_{-h}\varphi) + b\,T(\tau_{-h}\psi) = a\,U_h(\varphi) + b\,U_h(\psi). \quad (234)$$

Continuity. If $\varphi_n \to \varphi$ in $\mathcal{D}(\mathbb{R}^d)$ (supports in a fixed compact K, all derivatives converge uniformly), then $\tau_{-h}\varphi_n \to \tau_{-h}\varphi$ in $\mathcal{D}(\mathbb{R}^d)$ (supports in the fixed compact $K + h$, all derivatives still converge uniformly). Since T is continuous,

$$U_h(\varphi_n) = T(\tau_{-h}\varphi_n) \to T(\tau_{-h}\varphi) = U_h(\varphi). \tag{235}$$

□

4.50 REMARK. **Consistency with functions.** When $T = T_f$ for $f \in L^1_{\mathrm{loc}}$, the distribution $\tau_h T_f$ agrees with $T_{\tau_h f}$:

$$(\tau_h T_f)(\varphi) = T_f(\tau_{-h}\varphi) = \int f(x)\,\varphi(x+h)\,dx = \int f(y-h)\,\varphi(y)\,dy = T_{\tau_h f}(\varphi). \tag{236}$$

So the distributional definition extends the classical one.

4.51 EXAMPLE (Translation of the Dirac delta). $(\tau_h\delta)(\varphi) = \delta(\tau_{-h}\varphi) = (\tau_{-h}\varphi)(0) = \varphi(h) = \delta_h(\varphi)$. So $\tau_h\delta = \delta_h$: translating the delta moves the point of evaluation.

4.7.5 Convolution

Convolution with a test function is a smoothing operation: it always produces a smooth function, even when applied to a distribution. The definition uses translation and reflection from the preceding section.

CONVOLUTION OF FUNCTIONS (REVIEW) For $f \in L^1_{\mathrm{loc}}(\mathbb{R}^d)$ and $\psi \in \mathcal{D}(\mathbb{R}^d)$:

$$(f * \psi)(x) = \int_{\mathbb{R}^d} f(y)\,\psi(x-y)\,dy. \tag{237}$$

The result $f * \psi$ is smooth (C^∞) and satisfies $D^\alpha(f * \psi) = f * D^\alpha\psi$.

CONVOLUTION OF A DISTRIBUTION WITH A TEST FUNCTION

4.52 DEFINITION (Convolution with a test function). Let $T \in \mathcal{D}'(\mathbb{R}^d)$ and $\psi \in \mathcal{D}(\mathbb{R}^d)$. The **convolution** $T * \psi$ is the function

$$(T * \psi)(x) = \langle T, \tau_x\check{\psi}\rangle = \langle T_y, \psi(x-y)\rangle \tag{238}$$

where τ_x is translation (Definition 4.48), $\check{\psi}$ is reflection, and T acts on the y variable.

When $T = T_f$ for a locally integrable function f, this recovers the usual convolution: $(T_f * \psi)(x) = \int f(y)\,\psi(x-y)\,dy = (f * \psi)(x)$.

4.53 THEOREM (Regularization by convolution). *Let* $T \in \mathcal{D}'(\mathbb{R}^d)$ *and* $\psi \in \mathcal{D}(\mathbb{R}^d)$. *Then:*

1. $T * \psi \in C^\infty(\mathbb{R}^d)$*: the result is always a smooth function.*
2. $D^\alpha(T * \psi) = (D^\alpha T) * \psi = T * (D^\alpha \psi)$.
3. $\operatorname{supp}(T * \psi) \subseteq \operatorname{supp}(T) + \operatorname{supp}(\psi)$ *(Minkowski sum).*

Proof. **Smoothness.** Fix $x_0 \in \mathbb{R}^d$ and consider the map $x \mapsto \tau_x \check{\psi}$ from $\mathbb{R}^d$ into $\mathcal{D}(\mathbb{R}^d)$. This map is smooth: for any direction e_i,

$$\lim_{h \to 0} \frac{\tau_{x_0 + h e_i} \check{\psi} - \tau_{x_0} \check{\psi}}{h} = \partial_i [\tau_{x_0} \check{\psi}] \tag{239}$$

and the convergence holds in $\mathcal{D}(\mathbb{R}^d)$ (all supports stay in a fixed compact set, all derivatives converge uniformly). Applying the continuous linear functional T:

$$\partial_i (T * \psi)(x_0) = \lim_{h \to 0} \frac{\langle T, \tau_{x_0 + h e_i} \check{\psi} \rangle - \langle T, \tau_{x_0} \check{\psi} \rangle}{h} = \langle T, \partial_i [\tau_{x_0} \check{\psi}] \rangle. \tag{240}$$

Iterating gives smoothness to all orders.

Interchange of derivative and convolution. From the proof above, $D^\alpha(T * \psi)(x) = \langle T, D_x^\alpha \psi(x - \cdot) \rangle = (T * D^\alpha \psi)(x)$. For the other identity, $\langle D^\alpha T, \psi(x - \cdot) \rangle = (-1)^{|\alpha|} \langle T, D_y^\alpha \psi(x - y) \rangle = \langle T, D_x^\alpha \psi(x - \cdot) \rangle$, using the chain rule $D_y^\alpha \psi(x - y) = (-1)^{|\alpha|} D_x^\alpha \psi(x - y)$.

□

Approximation to the identity

4.54 DEFINITION (Mollifier). A **mollifier** (or approximation to the identity) is a family $\{\rho_\varepsilon\}_{\varepsilon > 0}$ of test functions satisfying:

1. $\rho_\varepsilon \geq 0$,
2. $\int \rho_\varepsilon = 1$,
3. $\operatorname{supp}(\rho_\varepsilon) \subseteq B(0, \varepsilon)$.

The standard choice is $\rho_\varepsilon(x) = \varepsilon^{-d} \rho(x/\varepsilon)$ where ρ is a fixed non-negative test function with $\int \rho = 1$ supported in the unit ball.

4.55 THEOREM (Mollification approximates distributions). *Let* $T \in \mathcal{D}'(\mathbb{R}^d)$ *and* $\{\rho_\varepsilon\}$ *a mollifier. Then:*

1. *Each* $T * \rho_\varepsilon$ *is a smooth function.*
2. $T * \rho_\varepsilon \to T$ *in* $\mathcal{D}'(\mathbb{R}^d)$ *as* $\varepsilon \to 0$.
3. *If* $T = T_f$ *for* $f \in L^p(\mathbb{R}^d)$, $1 \leq p < \infty$, *then* $f * \rho_\varepsilon \to f$ *in* L^p.

This is one of the most useful results in analysis: every distribution can be approximated by smooth functions. This is the distributional analogue of sequential density (Proposition 4.24), made constructive.

4.8 *Support and Locality*

Support of a Distribution A distribution T is said to be supported in a closed set K if $\langle T, \phi\rangle = 0$ for all $\phi \in \mathcal{C}_c^\infty(\mathbb{R}^d)$ with $\operatorname{supp}(\phi) \cap K = \emptyset$.

The support of T, denoted $\operatorname{supp}(T)$, is the smallest set K such that T is supported in K.

4.56 example.
- The support of the Dirac delta is $\operatorname{supp}(\delta) = \{0\}$.
- If f is a continuous function, $\operatorname{supp}(T_f) = \overline{\{x : f(x) \neq 0\}}$.

Local Structure Two distributions T_1 and T_2 are equal on an open set Ω if $\langle T_1, \phi\rangle = \langle T_2, \phi\rangle$ for all ϕ.

If T_f is a regular distribution and $T_f = 0$ on an open set Ω, then $f = 0$ almost everywhere on Ω.

Distributions with Point Support A fundamental result: If T is a distribution supported at a single point x_0, then T is a finite linear combination of derivatives of the delta function at x_0:

$$T = \sum_{|\alpha| \leq m} c_\alpha D^\alpha \delta_{x_0} \tag{241}$$

where m is the order of T.

4.8.1 The Structure Theorem for Distributions

The results above characterize distributions with *point* support. The structure theorem generalizes this to *all* distributions: every distribution is, locally, a finite-order derivative of a continuous function.

4.57 theorem (Structure theorem for distributions). *Let $T \in \mathcal{D}'(\Omega)$ and let $K \subset \Omega$ be compact. Then there exist an integer $N \geq 0$ and continuous functions $f_\alpha \in C(K)$, $|\alpha| \leq N$, such that*

$$T = \sum_{|\alpha| \leq N} D^\alpha f_\alpha \qquad \text{on a neighborhood of } K. \tag{242}$$

If T has finite order m globally (i.e., the seminorm bound $|\langle T, \varphi\rangle| \leq C_K \sum_{|\alpha|\leq m} \sup |D^\alpha \varphi|$ holds with the same m for all compact K), then we can take $N = m$ uniformly.

Proof. We sketch the argument in one dimension for clarity; the general case is identical in structure.

Step 1: From the order bound to a norm estimate. On a compact set K, the distribution T has some finite order m:

$$|\langle T, \varphi\rangle| \leq C_K \sum_{k=0}^{m} \|\varphi^{(k)}\|_\infty \qquad \text{for all } \varphi \in \mathcal{D}_K. \tag{243}$$

This means T is a continuous linear functional on $\mathcal{D}_K$ equipped with the C^m norm.

Step 2: Extension to a continuous linear functional. The space $\mathcal{D}_K$ is a subspace of $C^m(K)$. By the Hahn-Banach theorem, T extends to a continuous linear functional on $C^m(K)$.

Step 3: Representation via antiderivatives. A continuous linear functional on $C^m(K)$ can be written as a finite sum of integrals against Radon measures μ_α, $|\alpha| \le m$:

$$\langle T, \varphi \rangle = \sum_{|\alpha| \le m} \int D^\alpha \varphi \, d\mu_\alpha. \tag{244}$$

Each Radon measure μ_α on a compact subset of $\mathbb{R}^d$ has a continuous antiderivative: there exists $f_\alpha \in C(K)$ such that $\int \psi \, d\mu_\alpha = (-1)^{|\alpha|} \int f_\alpha \, D^\alpha \psi \, dx$ for test functions ψ. (In one dimension, this is just the fact that the antiderivative of a measure is a function of bounded variation, and one more antiderivative gives a continuous function.) Substituting:

$$\langle T, \varphi \rangle = \sum_{|\alpha| \le N} (-1)^{|\alpha|} \int f_\alpha \, D^\alpha \varphi \, dx = \sum_{|\alpha| \le N} \langle D^\alpha f_\alpha, \varphi \rangle \tag{245}$$

where $N = m + d$ accounts for the additional antiderivatives needed in d dimensions.

□

4.58 REMARK (Distributions are derivatives of continuous functions). The structure theorem demystifies distributions: no matter how singular a distribution may appear, it is always locally a finite sum of derivatives of continuous functions. The Dirac delta is H' (derivative of the Heaviside function, which is bounded and measurable — or the second derivative of the continuous function $\max(x, 0)$). The derivative δ' is the third derivative of $\max(x, 0)$. And so on.

The price of increased singularity is simply a higher derivative order, not a fundamentally different kind of object.

The structure theorem immediately implies the point support result stated above:

4.59 COROLLARY. *If* $T \in \mathcal{D}'(\mathbb{R}^d)$ *with* $\operatorname{supp}(T) = \{x_0\}$, *then* $T = \sum_{|\alpha| \le m} c_\alpha D^\alpha \delta_{x_0}$.

Proof. By the structure theorem, on a ball B around x_0, we have $T = \sum_{|\alpha| \le N} D^\alpha f_\alpha$ for continuous f_α. Since $\operatorname{supp}(T) = \{x_0\}$, the distribution T vanishes on $B \setminus \{x_0\}$, which means each f_α is a polynomial of degree at most $N - |\alpha|$ on $B \setminus \{x_0\}$, and by continuity on all of B. A polynomial supported at a point must be constant times a suitable derivative of δ_{x_0}. Collecting terms gives $T = \sum_{|\alpha| \le m} c_\alpha D^\alpha \delta_{x_0}$ where $m \le N$ is the order of T.

□

4.9 *Sobolev Spaces*

Big Idea

L^p norms measure height and width but are blind to oscillation. A function can oscillate faster and faster while its L^p norm stays constant. Sobolev norms add a third dimension, **frequency**, by penalizing derivatives. The Sobolev space $W^{k,p}(\Omega)$ consists of functions whose weak derivatives up to order k are in L^p: it sees not just how big a function is, but how fast it oscillates.

4.9.1 From weak derivatives to function spaces

In the previous sections we defined the weak derivative of a locally integrable function: $v = D^\alpha u$ weakly if

$$\int_\Omega u\, D^\alpha \varphi \, dx = (-1)^{|\alpha|} \int_\Omega v\, \varphi \, dx \quad \text{for all } \varphi \in C_c^\infty(\Omega). \tag{246}$$

This makes sense for any $u \in L^1_{\mathrm{loc}}(\Omega)$, but the derivative v might not be in L^p, or might not exist as an L^p function at all. Sobolev spaces are defined by requiring that weak derivatives exist **and** have controlled size.

4.60 DEFINITION (Sobolev space). $W^{k,p}(\Omega)$

Let $\Omega \subseteq \mathbb{R}^d$ be open, $k \geq 0$ an integer, and $1 \leq p \leq \infty$. The **Sobolev space** $W^{k,p}(\Omega)$ consists of all functions $u \in L^p(\Omega)$ whose weak derivatives $D^\alpha u$ exist and belong to $L^p(\Omega)$ for all multi-indices $|\alpha| \leq k$. The **Sobolev norm** is, for $1 \leq p < \infty$,

$$\|u\|_{W^{k,p}} = \left(\sum_{|\alpha| \leq k} \|D^\alpha u\|_{L^p(\Omega)}^p \right)^{1/p}, \tag{247}$$

and for $p = \infty$,

$$\|u\|_{W^{k,\infty}} = \max_{|\alpha| \leq k} \|D^\alpha u\|_{L^\infty(\Omega)}. \tag{248}$$

4.61 PROPOSITION (is a Banach space). $W^{k,p}(\Omega)$

For $1 \leq p \leq \infty$ and $k \geq 0$, the space $(W^{k,p}(\Omega), \|\cdot\|_{W^{k,p}})$ is a Banach space.

Proof. Let (u_n) be Cauchy in $W^{k,p}$. Then for each $|\alpha| \leq k$, the sequence $(D^\alpha u_n)$ is Cauchy in $L^p(\Omega)$. Since L^p is complete, $D^\alpha u_n \to v_\alpha$ in L^p for some v_α. In particular, $u_n \to v_0 =: u$ in L^p.

We verify that $v_\alpha = D^\alpha u$ weakly: for any $\varphi \in C_c^\infty(\Omega)$,

$$\int_\Omega v_\alpha\, \varphi \, dx = \lim_n \int_\Omega D^\alpha u_n\, \varphi \, dx = \lim_n (-1)^{|\alpha|} \int_\Omega u_n\, D^\alpha \varphi \, dx = (-1)^{|\alpha|} \int_\Omega u\, D^\alpha \varphi \, dx. \tag{249}$$

So $u \in \mathrm{W}^{k,p}$ and $u_n \to u$ in $\mathrm{W}^{k,p}$.

□

The completeness proof is simple because it reduces to the completeness of L^p: each derivative converges separately, and the weak derivative structure is preserved under L^p limits. This is a common pattern: Sobolev spaces inherit their analytical properties from L^p.

4.62 DEFINITION (The Hilbert-Sobolev spaces). $\mathrm{H}^k(\Omega)$

For $p = 2$, the Sobolev space $\mathrm{W}^{k,2}(\Omega)$ is denoted $\mathrm{H}^k(\Omega)$. It is a **Hilbert space** with inner product

$$\langle u, v \rangle_{\mathrm{H}^k} = \sum_{|\alpha| \le k} \int_\Omega \mathrm{D}^\alpha u \cdot \mathrm{D}^\alpha v \, dx. \tag{250}$$

The notation H^k reflects the Hilbert space structure. For PDE applications, $\mathrm{H}^1(\Omega)$ is the most important space: it controls the function and its first derivatives in L^2.

4.9.2 What does the Sobolev norm measure?

In the introduction (Remark 1.9), we asked "what do L^p norms measure?" and found that they capture two features of a function:

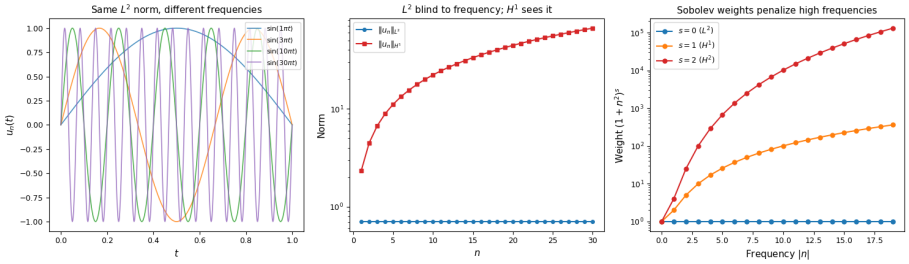
- **Height** (amplitude): the peak value, controlled by L^∞.
- **Width** (spatial extent): the size of the support, contributing to L^p for $p < \infty$.

But these two features miss something fundamental. Consider the sequence $u_n(t) = \sin(n\pi t)$ on $[0, 1]$:

$$\|u_n\|_{\mathrm{L}^2}^2 = \int_0^1 \sin^2(n\pi t)\, dt = \frac{1}{2} \quad \text{for all } n. \tag{251}$$

Every function in this sequence has the same height, the same support, and the same L^2 norm, yet they are clearly different: u_{100} oscillates 100 times faster than u_1. The L^2 norm is completely blind to this.

This is the same sequence that appeared in the weak convergence chapter (Example **??**): $\sin(n\pi t) \rightharpoonup 0$ weakly in L^2, but $\|\sin(n\pi t)\|_{\mathrm{L}^2} = 1/\sqrt{2}$ for all n. The L^2 norm cannot distinguish rapid oscillation from no oscillation at all.



Left: the functions $\sin(n\pi t)$ *oscillate faster but have constant* L^2 *norm. Center: the* L^2 *norm (blue) stays flat while the* H^1 *norm (red) grows linearly in n, showing that the derivative detects oscillation. Right: the*

Fourier weights $(1+n^2)^s$ *that define* H^s*; higher* s *penalizes high frequencies more aggressively.*

THE THIRD FEATURE: FREQUENCY The Sobolev norm adds the missing dimension. Computing the derivative:

$$u_n'(t) = n\pi\cos(n\pi t), \qquad \|u_n'\|_{\mathrm{L}^2}^2 = n^2\pi^2\int_0^1 \cos^2(n\pi t)\,dt = \frac{n^2\pi^2}{2}. \tag{252}$$

So the H^1 norm sees the oscillation:

$$\|u_n\|_{\mathrm{H}^1}^2 = \|u_n\|_{\mathrm{L}^2}^2 + \|u_n'\|_{\mathrm{L}^2}^2 = \frac{1}{2} + \frac{n^2\pi^2}{2} \to \infty. \tag{253}$$

The derivative term $\|u_n'\|_{\mathrm{L}^2}$ blows up like $n\pi$: differentiation converts oscillation frequency into amplitude. The H^1 norm detects the frequency because derivatives penalize oscillation.

THE BUMP FUNCTION SCALING HEURISTIC Following Tao, we can make this precise with a model function that independently controls all three features. Fix a bump function $\phi \in \mathrm{C}_c^\infty(\mathbb{R})$ with $\mathrm{supp}(\phi) \subset [-1,1]$, and consider

$$f_{\mathrm{A,R,N}}(x) = \mathrm{A}\,\phi(x/\mathrm{R})\,\sin(\mathrm{N}x) \tag{254}$$

where:

- $\mathrm{A} > 0$ controls the **height** (amplitude),
- $\mathrm{R} > 0$ controls the **width** (spatial extent: support $\approx [-\mathrm{R}, \mathrm{R}]$),
- $\mathrm{N} > 0$ controls the **frequency** (oscillation rate).

The L^p norm scales as (by the change of variables $y = x/\mathrm{R}$):

$$\|f_{\mathrm{A,R,N}}\|_{\mathrm{L}^p}^p = \int |\mathrm{A}\,\phi(x/\mathrm{R})\,\sin(\mathrm{N}x)|^p\,dx \approx \mathrm{A}^p\,\mathrm{R}\,\|\phi\cdot\sin\|_p^p \tag{255}$$

so $\|f_{\mathrm{A,R,N}}\|_{\mathrm{L}^p} \approx \mathrm{A}\,\mathrm{R}^{1/p}$. The frequency N does **not appear**: the L^p norm is blind to oscillation.
The $\mathrm{W}^{s,p}$ norm behaves differently. Each derivative of $\sin(\mathrm{N}x)$ pulls down a factor of N:

$$\mathrm{D}^s[\mathrm{A}\,\phi(x/\mathrm{R})\,\sin(\mathrm{N}x)] \approx \mathrm{A}\,\mathrm{N}^s\,\phi(x/\mathrm{R})\,\sin(\mathrm{N}x + s\pi/2) + \text{lower-order terms} \tag{256}$$

(the lower-order terms involve derivatives of ϕ, which contribute factors of R^{-1} instead of N, and are negligible when $\mathrm{N} \gg 1/\mathrm{R}$). Therefore:

$$\|f_{\mathrm{A,R,N}}\|_{\mathrm{W}^{s,p}} \approx \mathrm{A}\,\mathrm{R}^{1/p}\,\mathrm{N}^s. \tag{257}$$

The three parameters are now visible:

Norm	Scaling	What it sees
$\lvert f\rvert_{L^p}$	$A\,R^{1/p}$	Height × Width
$\lvert f\rvert_{W^{s,p}}$	$A\,R^{1/p}\,N^s$	Height × Width × Frequency^s

The extra factor N^s is the whole point: **the Sobolev norm penalizes high-frequency oscillation, with the penalty growing as** N^s. Higher Sobolev index s means stronger frequency penalty.

4.63 REMARK (The three dimensions of function behavior). A function's "size" in the broad sense has three independent components:

1. **Amplitude** A: how tall is the function? (Controlled by L^∞.)
2. **Spatial extent** R: how wide is the support? (Controlled by L^p for $p < \infty$.)
3. **Frequency scale** N: how fast does it oscillate? (Controlled by $W^{s,p}$.)

L^p norms see dimensions 1 and 2. Sobolev norms see all three. This is why Sobolev spaces are the natural setting for PDE: differential equations involve derivatives, and the Sobolev norm is the norm that controls derivatives.

REGULARITY AS FREQUENCY DECAY The bump function heuristic shows that the Sobolev norm penalizes high frequency. But how many degrees of regularity does a function have? To make this precise, consider the one-parameter family

$$G_{s,N}(x) = N^{-s}\,\phi(x)\,\sin(Nx), \qquad s > 0, \quad N \gg 1, \tag{258}$$

where $\phi \in C_c^\infty(\mathbb{R})$ is a fixed bump. The prefactor N^{-s} is chosen so that $\|G_{s,N}\|_{W^{s,p}}$ remains bounded as $N \to \infty$ (from the scaling $\|f_{A,R,N}\|_{W^{s,p}} \approx AR^{1/p}N^s$ with $A = N^{-s}$ and $R = 1$).

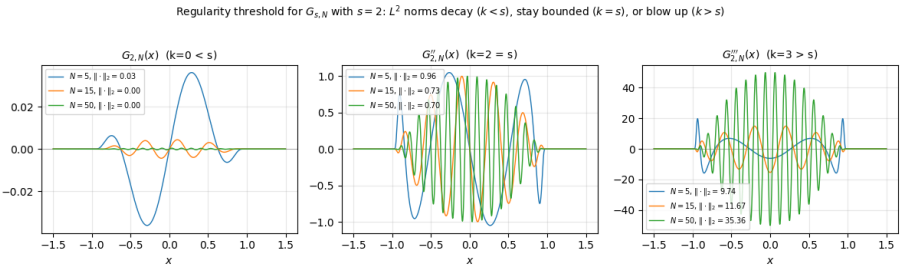
Computing derivatives (keeping only the dominant term when $N \gg 1$):

$$G_{s,N}^{(k)}(x) \approx N^{k-s}\,\phi(x)\,\sin(Nx + k\pi/2) + \text{lower-order terms.} \tag{259}$$

The behavior splits sharply at $k = s$:

Derivative order k	$\lvert G_{s,N}^{(k)}\rvert_{L^p}$ as $N \to \infty$	Interpretation
$k < s$	$\to 0$ (decays as N^{k-s})	Well within regularity budget
$k = s$	$O(1)$ (bounded)	Exactly at the regularity threshold
$k > s$	$\to \infty$ (grows as N^{k-s})	Beyond the regularity budget

The following plot shows $G_{s,N}$ for $s = 2$ and increasing N. The function itself shrinks (because of the N^{-s} prefactor), but its second derivative stays $O(1)$ and its third derivative grows.



4.64 REMARK (Sobolev regularity is sharp). The family $G_{s,N}$ demonstrates that s **is a sharp regularity threshold**: the function has exactly s derivatives in L^p and no more. This is the core idea of the Sobolev scale: the index s counts exactly how many derivatives a function can afford.

The mechanism is frequency: each derivative pulls down a factor of N, and the prefactor N^{-s} budgets exactly s such factors before the L^p norm diverges.

THE FOURIER PICTURE For $H^s = W^{s,2}$ on the torus $[0, 2\pi]$ (or $\mathbb{R}^d$ via the Fourier transform), the connection to frequency is exact. If $u(x) = \sum_n \hat{u}_n e^{inx}$, then $D^s u$ has Fourier coefficients $(in)^s \hat{u}_n$, and Parseval gives:

$$\|u\|_{H^s}^2 = \sum_n (1 + |n|^2)^s |\hat{u}_n|^2. \tag{260}$$

The weight $(1 + |n|^2)^s$ penalizes high frequencies $|n|$ by a factor of $|n|^{2s}$. This makes the frequency interpretation precise:

- L^2 **(**$s = 0$**):** all frequencies weighted equally. The norm is $\sum |\hat{u}_n|^2$: total energy, regardless of frequency.
- H^1 **(**$s = 1$**):** weight $(1 + n^2)$. High frequencies are penalized quadratically.
- H^k **(**$s = k$**):** weight $(1 + n^2)^k$. Higher k means stronger suppression of high frequencies.

A function is in H^s if and only if its Fourier coefficients decay fast enough to compensate for the weight $(1 + |n|^2)^s$. Smoother functions have faster-decaying Fourier coefficients, so they live in higher Sobolev spaces.

4.9.3 Examples: what is and isn't in $W^{1,2}$

4.65 EXAMPLE (Functions in and out of). H^1

In $H^1(0, 1)$**:**

- Any C^1 function on $[0, 1]$: classical derivatives are in L^2.
- $f(x) = |x - 1/2|$: the weak derivative is $f'(x) = \text{sgn}(x - 1/2)$, which is in L^2 (in fact in L^∞). The kink at $x = 1/2$ is invisible to the H^1 norm.

Not in $H^1(0, 1)$**:**

- The Heaviside function $H(x) = \mathbf{1}_{[1/2,1]}(x)$: its weak derivative is $\delta_{1/2}$, a measure and not an L^2 function. So $H \in L^2$ but $H \notin H^1$. The jump discontinuity is detected by the Sobolev norm.
- $f(x) = x^{-1/4}$ on $(0, 1)$: $f \in L^2(0, 1)$ (barely), but $f'(x) = -\frac{1}{4}x^{-5/4} \notin L^2(0, 1)$. The singularity at 0 is too strong for one derivative in L^2.

The pattern: H^1 tolerates kinks (Lipschitz singularities) but not jumps or blowup. This is because the weak derivative of a Lipschitz function is bounded (hence in every L^p), while the derivative of a discontinuous function is a measure.

4.10 *Sobolev Embeddings*

Big Idea

Bounded sequences in L^2 can fail to have convergent subsequences in two ways: **oscillation** (mass escaping to high frequencies) and **translation** (mass escaping to spatial infinity). Controlling a derivative kills oscillation; working on a bounded domain kills translation. Together they restore Bolzano-Weierstrass. This is **Rellich-Kondrachov compactness**.

Along the way, controlling derivatives also *upgrades* integrability. The **Sobolev embedding theorem** makes this quantitative: s derivatives in L^p buy integrability up to L^{p^*} (or continuity if $sp > d$), with exchange rate $1/p^* = 1/p - s/d$ set by an uncertainty principle.

4.10.1 Motivation: when does a bounded sequence converge?

In finite dimensions, Bolzano-Weierstrass guarantees that every bounded sequence has a convergent subsequence. In infinite-dimensional spaces like $L^2(\Omega)$ this fails dramatically, and understanding *how* it fails is the key to the whole chapter.

Consider a bounded sequence $\{u_n\} \subset L^2(\Omega)$, say $\|u_n\|_{L^2} \leq 1$. What can go wrong? There are essentially **two pathologies**:

(P1) Oscillation, where energy escapes to high frequencies. Take $\Omega = (0, 1)$ and $u_n(x) = \sqrt{2}\,\sin(n\pi x)$. Each u_n has $\|u_n\|_{L^2} = 1$, but any two distinct terms are orthogonal, so $\|u_n - u_m\|_{L^2} = \sqrt{2}$ for $n \neq m$. No subsequence is Cauchy. The mass does not escape in space; instead it escapes into ever higher frequencies.

(P2) Translation, where energy escapes to spatial infinity. On $\Omega = \mathbb{R}^d$, fix a bump $\varphi \in C_c^\infty$ with $\|\varphi\|_{L^2} = 1$ and set $u_n(x) = \varphi(x - ne_1)$. Each u_n has norm 1, but the supports become disjoint, so again $\|u_n - u_m\|_{L^2} = \sqrt{2}$. The mass slides off to infinity.

These two examples exhaust the ways compactness can fail in L^p: any non-precompact bounded sequence is, up to subsequence, a combination of oscillating modes and escaping bumps. The whole chapter can be read as a program to **defeat these two pathologies** and recover Bolzano-Weierstrass:

- *Kill oscillation* by controlling a derivative. If $\|\nabla u_n\|_{L^2}$ is also bounded, the high Fourier modes cannot carry the mass. The derivative of $\sin(n\pi x)$ has norm $\sim n$, which blows up, so controlling the gradient caps the usable frequency.
- *Kill translation* by working on a bounded domain, where there is nowhere for the mass to escape to.

Doing both yields **Rellich-Kondrachov compactness**: the embedding $H^1(\Omega) \hookrightarrow L^2(\Omega)$ is compact on bounded Ω. Along the way we will also see how controlling derivatives not only rescues compactness but *upgrades* integrability, climbing the ladder toward L^∞ and Hölder continuity. This is the Sobolev embedding theorem.

4.10.2 From Sobolev regularity to integrability

The Sobolev norm controls more than just derivatives; it forces the function to have better integrability and even continuity, depending on how many derivatives are controlled.

Before making this quantitative, it is worth asking *why we want integrability control at all*. The answer points to a single goal: **we are trying to climb the ladder**

$$L^p \subset L^q \subset \cdots \subset L^\infty \subset C^0 \subset C^{0,\alpha}, \tag{261}$$

because the top of the ladder is where classical tools live. Continuous and equicontinuous functions are precompact by Arzelà-Ascoli (Theorem 1.33); continuity lets us evaluate u at a point, impose classical boundary values, and make sense of nonlinear terms like u^q pointwise. The L^p rungs in between are what we settle for when we cannot reach the top. The Sobolev embedding theorem tells us *exactly* how high we can climb given a fixed budget of derivatives.

The simplest instance is the Poincaré inequality, which in one dimension turns out to climb all the way.

THE POINCARÉ INEQUALITY: THE PROTOTYPE EMBEDDING Everything in the Sobolev embedding theorem can be read off a single one-dimensional calculation. We do that calculation here, and the result is the **prototype**: a single pointwise estimate, obtained from the fundamental theorem of calculus and Cauchy-Schwarz, produces *three* regularity statements at once,

$$L^2 \text{ control} \;\to\; L^\infty \text{ control} \;\to\; \text{Hölder control.} \tag{262}$$

Every higher-dimensional Sobolev embedding is a way of running this same template in d directions.

The single estimate. Take $\Omega = (0, L)$ and $u \in H^1_0(0, L)$, so $u(0) = 0$. The fundamental theorem of calculus gives the representation $u(x) = \int_0^x u'(t)\,dt$, and Cauchy-Schwarz turns it into a *pointwise* bound:

$$|u(x)|^2 \;=\; \Big|\int_0^x 1 \cdot u'(t)\,dt\Big|^2 \;\le\; x \int_0^x |u'(t)|^2\,dt \;\le\; L\,\|u'\|_{L^2}^2. \tag{$\star$}$$

(Strictly speaking, the FTC representation holds pointwise for $u \in C^1$ with $u(0) = 0$; the estimate ($\star$) then extends to all of $H^1_0(0, L)$ by density, and simultaneously shows that each $u \in H^1_0$ has a continuous representative for which $u(x) = \int_0^x u'(t)\,dt$ holds.)

That is the engine. The three regularity statements all fall out by reading ($\star$) in different ways.

Rung 1: L^2 control (the Poincaré inequality). Integrate ($\star$) in x over $(0, L)$:

$$\|u\|_{L^2}^2 \le L^2\,\|u'\|_{L^2}^2, \qquad \|u\|_{L^2} \le L\,\|u'\|_{L^2}. \tag{263}$$

The gradient controls the function in mean square.

Rung 2: L^∞ control. Take the supremum over x in ($\star$):

$$\|u\|_{L^\infty(0,L)} \le L^{1/2}\,\|u'\|_{L^2}. \tag{264}$$

The gradient controls the function pointwise.

Rung 3: Hölder control. Apply the same Cauchy-Schwarz argument to the difference $u(x) - u(y) = \int_y^x u'(t)\,dt$:

$$|u(x) - u(y)| \le |x - y|^{1/2}\,\|u'\|_{L^2}, \tag{265}$$

so $u \in C^{0,1/2}(\overline{\Omega})$ with seminorm bounded by $\|u'\|_{L^2}$. The gradient controls the function's continuity.

Three regularity classes, one estimate. The classical Poincaré inequality is just rung 1, and we record it formally.

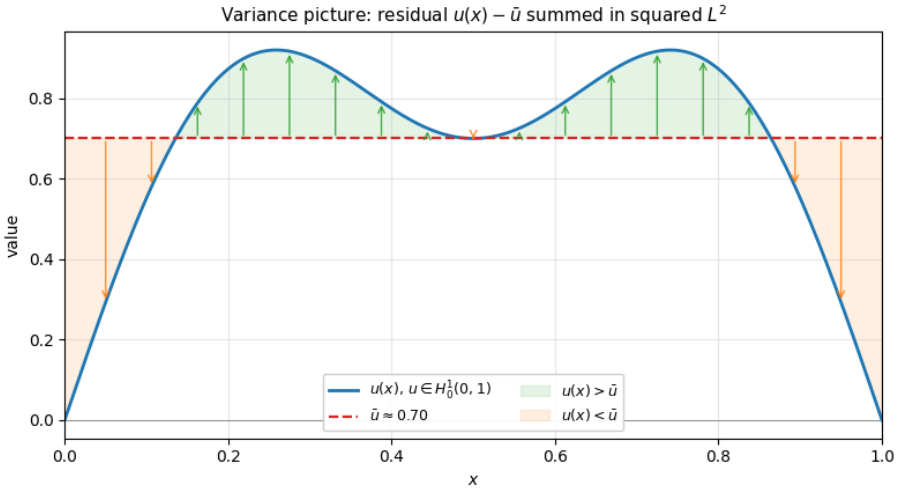
4.66 THEOREM (Poincaré inequality on). H_0^1

Let $\Omega \subset \mathbb{R}^d$ be a bounded, connected open set with Lipschitz boundary. There exists a constant $C = C(\Omega)$ such that for all $u \in H_0^1(\Omega)$,

$$\|u\|_{L^2(\Omega)} \le C\,\|\nabla u\|_{L^2(\Omega)}. \tag{266}$$

In one dimension this trade is *lossless*: a single derivative in L^2 buys all the way to Hölder continuity, the top of the integrability ladder. Rellich compactness then comes for free, since a bounded set in $H_0^1(0, L)$ is uniformly bounded and equicontinuous, and Arzelà-Ascoli (Theorem 1.33) delivers a C^0-convergent subsequence. The motivating oscillation pathology $u_n(x) = \sqrt{2}\sin(n\pi x)$ is automatically excluded: its H_0^1 norm $\sqrt{2}\,n\pi$ blows up, so $\{u_n\}$ never enters the bounded H_0^1 ball in the first place.

The variance picture The L^2 rung has a probabilistic flavor: $\|u - \bar{u}\|_{L^2}^2$ is the *variance* of u around its mean, summed in squared L^2 sense. The plot below shows a typical H_0^1 function on $(0, 1)$ with its mean line; each arrow has length $|u(x) - \bar{u}|$, and the squared sum of all arrows is $\|u - \bar{u}\|_{L^2}^2$. Poincaré says this total cannot exceed $C^2\|\nabla u\|_{L^2}^2$: a flatter function (small gradient) means short arrows, hence small variance.



The H_0^1 statement of Poincaré bounds $\|u\|_{L^2}$ rather than just $\|u - \bar{u}\|_{L^2}$, and the mean does *not* appear. This is sometimes called "the mean drops out," but it deserves a careful reading. The Pythagorean decomposition

$$\|u\|_{L^2}^2 \;=\; \underbrace{\|u-\bar{u}\|_{L^2}^2}_{\text{variance}} \;+\; \underbrace{|\bar{u}|^2\,|\Omega|}_{\text{mean energy}} \tag{267}$$

splits the squared norm into orthogonal pieces. On all of H^1, the gradient controls only the variance: a constant function has $\nabla u = 0$ but arbitrary mean. On H_0^1, the boundary condition forces u back to zero at $\partial\Omega$, so a small gradient also prevents the function from sitting at a high plateau; the mean energy $|\bar{u}|^2|\Omega|$ is itself bounded by the gradient. Indeed, FTC with $u(0) = 0$ and Cauchy-Schwarz give $|u(x)| \le L^{1/2}\|u'\|_{L^2}$ pointwise, so $|\bar{u}| \le L^{1/2}\|u'\|_{L^2}$ and thus $|\bar{u}|^2|\Omega| \le L^2\|u'\|_{L^2}^2$ — no appeal to Poincaré needed. Both orthogonal pieces are then controlled, and they combine into the cleaner $\|u\|_{L^2}^2 \le C^2\|\nabla u\|_{L^2}^2$. The mean is not made zero (a typical H_0^1 function has $\bar{u} \neq 0$); the boundary condition has just made it controllable in the same breath as the variance.

4.67 REMARK (The mean-free Poincaré-Wirtinger version on). H^1

Without a boundary condition, the gradient cannot detect constants ($\nabla(u + c) = \nabla u$), so any inequality of the form $\|u\|_{L^2} \le C\|\nabla u\|_{L^2}$ fails on H^1. The fix is to subtract the mean,

$$\|u-\bar{u}\|_{L^2(\Omega)} \le C\,\|\nabla u\|_{L^2(\Omega)}, \qquad \bar{u} = \frac{1}{|\Omega|}\int_\Omega u\,dx, \tag{268}$$

which projects u onto the orthogonal complement of constants where ∇ is injective. This is the Poincaré-Wirtinger inequality. Its proof requires a different technique (typically compactness and contradiction); the FTC argument above does not extend. We do not use this version in this chapter.

The frequency picture as confirmation The same ladder appears in Fourier coordinates. On $(0, L)$ with zero boundary, the eigenfunctions of $-\Delta$ are $\sin(n\pi x/L)$ with eigenvalues $(n\pi/L)^2$. These form an orthogonal basis of $L^2(0, L)$ and the same family, suitably normalized by $n\pi/L$, is an orthogonal basis of $H_0^1(0, L)$. So for $u \in H_0^1(0, L)$ we may write $u = \sum_n \hat{u}_n \sin(n\pi x/L)$ with convergence in both L^2 and H_0^1, and

$$\|u\|_{L^2}^2 = \sum_n |\hat{u}_n|^2, \qquad \|\nabla u\|_{L^2}^2 = \sum_n \tfrac{n^2\pi^2}{L^2}\,|\hat{u}_n|^2. \tag{269}$$

Rung 1: L^2 control (Poincaré). Term by term the coefficient inequality is immediate (since $n \ge 1$), and summing in n gives Poincaré:

$$|\hat{u}_n|^2 \le \tfrac{L^2}{\pi^2}\tfrac{n^2\pi^2}{L^2}|\hat{u}_n|^2 \quad\Longrightarrow\quad \|u\|_{L^2}^2 \le \tfrac{L^2}{\pi^2}\|\nabla u\|_{L^2}^2. \tag{270}$$

The optimal constant $C = L/\pi$ is attained by the lowest mode $\sin(\pi x/L)$, the slowest-oscillating eigenfunction.

Rung 2: L^∞ control. Bound the series pointwise by its coefficients and apply Cauchy-Schwarz on the index n:

$$|u(x)| = \Big|\sum_n \hat{u}_n \sin(n\pi x/\mathrm{L})\Big| \le \sum_n |\hat{u}_n| \le \Big(\sum_n \tfrac{\mathrm{L}^2}{n^2\pi^2}\Big)^{1/2}\Big(\sum_n \tfrac{n^2\pi^2}{\mathrm{L}^2}|\hat{u}_n|^2\Big)^{1/2} = \mathrm{C}_\infty \|\nabla u\|_{\mathrm{L}^2}. \tag{271}$$

The convergent sum $\sum n^{-2}$ is what makes the bound finite, and it is exactly the Fourier shadow of the $\int_0^x 1 \cdot u'\,dt$ step in ($\star$).

Rung 3: Hölder control. Expand the difference and use $|\sin a - \sin b| \le \min(2, |a-b|)$ termwise, then Cauchy-Schwarz:

$$|u(x)-u(y)| \le \sum_n |\hat{u}_n| \min\Big(2,\ \tfrac{n\pi|x-y|}{\mathrm{L}}\Big) \le \Big(\sum_n \tfrac{\min(2,\, n\pi|x-y|/\mathrm{L})^2}{n^2\pi^2/\mathrm{L}^2}\Big)^{1/2} \|\nabla u\|_{\mathrm{L}^2} \lesssim |x-y|^{1/2}\,\|\nabla u\|_{\mathrm{L}^2}. \tag{272}$$

Split the last sum at the crossover $n \sim \mathrm{L}/(\pi|x-y|)$ where the two arguments of min cross. *Low-frequency modes* (n below the crossover) are still in the smooth regime $|\sin a - \sin b| \approx |a-b|$, so they contribute the $|x-y|$-factor; *high-frequency modes* (n above the crossover) oscillate too fast to resolve the displacement $|x-y|$, so only the crude bound 2 is available. Balancing the two halves gives the $|x-y|^{1/2}$ exponent.

So all three rungs are visible in Fourier, each as a different way of weighting $\sum n^2|\hat{u}_n|^2$ against a convergent geometric factor.

L^2-compactness (Rellich-Kondrachov in 1D). The same expansion delivers L^2-compactness in one line. Let $\mathrm{B} = \{u : \|u\|_{\mathrm{H}^1_0} \le \mathrm{M}\}$, and let $\mathrm{P}_\mathrm{N} u = \sum_{n\le \mathrm{N}} \hat{u}_n \sin(n\pi x/\mathrm{L})$ be the projection onto the first N modes. Then $\mathrm{P}_\mathrm{N} u$ lives in the finite-dimensional (hence precompact) subspace $\mathrm{V}_\mathrm{N} = \mathrm{span}\{\sin(n\pi x/\mathrm{L}) : n \le \mathrm{N}\}$, while the tail is uniformly small in L^2:

$$\|u - \mathrm{P}_\mathrm{N} u\|_{\mathrm{L}^2}^2 \;=\; \sum_{n>\mathrm{N}} |\hat{u}_n|^2 \;\le\; \tfrac{\mathrm{L}^2}{\mathrm{N}^2\pi^2} \sum_{n>\mathrm{N}} \tfrac{n^2\pi^2}{\mathrm{L}^2}|\hat{u}_n|^2 \;\le\; \tfrac{\mathrm{L}^2\mathrm{M}^2}{\mathrm{N}^2\pi^2} \xrightarrow{\mathrm{N}\to\infty} 0. \tag{273}$$

So B is *totally bounded* in L^2 (covered by finitely many balls of any prescribed radius, using V_N's own compactness for the head), hence precompact. This is Rellich-Kondrachov in 1D.

The same ball B is also equicontinuous in $\mathrm{C}([0,\mathrm{L}])$ by the Hölder rung ($|u(x)-u(y)| \le \mathrm{M}|x-y|^{1/2}$ uniformly in $u \in \mathrm{B}$), so Arzelà-Ascoli gives uniform-norm precompactness — and $\mathrm{C} \hookrightarrow \mathrm{L}^2$ upgrades this to L^2-precompactness too. The two proofs are the same phenomenon seen twice: *high-frequency modes decay* and *uniform equicontinuity* are dual descriptions of "the bounded H^1_0 ball has no room to escape to infinity in L^2."

4.10.3 The uncertainty principle and the general theorem

The Sobolev embedding theorem says that controlling k derivatives in L^p forces membership in L^{p^*} with $1/p^* = 1/p - k/d$. Where does this critical exponent come from? The answer is an **uncertainty principle**: a function cannot be simultaneously localized in both space and frequency.

TWO REGIMES OF THE MODEL FUNCTION Fix a smooth bump ϕ on $\mathbb{R}^d$ with support in the unit ball, and a unit vector $e \in \mathbb{R}^d$. The model function is

$$f_{\mathrm{A,R,N}}(x) = \mathrm{A}\,\phi(x/\mathrm{R})\,\sin(\mathrm{N}\,e\cdot x), \qquad x \in \mathbb{R}^d, \tag{274}$$

a bump of height A, spatial scale R, oscillating at frequency N along direction e. Its L^p norms scale with the *volume* R^d of the support:

$$\|f\|_{\mathrm{L}^p} \;\sim\; \mathrm{A}\,\mathrm{R}^{d/p}. \tag{275}$$

Its gradient has two competing terms:

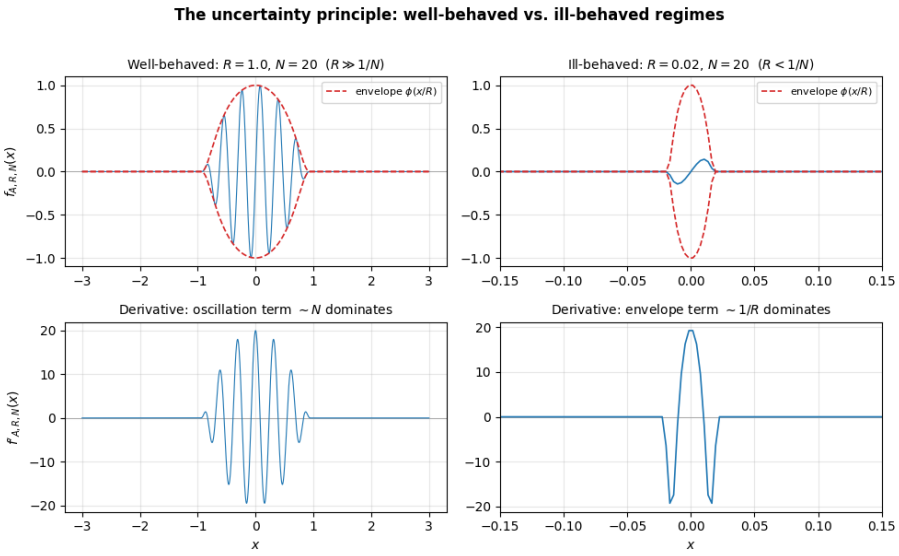
- The **envelope term** $\sim (\mathrm{A}/\mathrm{R})\,\phi'(x/\mathrm{R})\sin(\mathrm{N}\,e\cdot x)$, from differentiating the bump,
- The **oscillation term** $\sim \mathrm{AN}\,\phi(x/\mathrm{R})\cos(\mathrm{N}\,e\cdot x)$, from differentiating the sine.

Which dominates depends on the relationship between R and N:

Well-behaved regime: $\mathrm{RN} \geq 1$**.** The oscillation term dominates: $\|\nabla f\|_{\mathrm{L}^p} \sim \mathrm{AN}\,\mathrm{R}^{d/p}$. The function completes many oscillations within its support, and derivatives faithfully measure the frequency. The Sobolev norm behaves as the heuristic predicts: $\|f\|_{\mathrm{W}^{1,p}} \approx \mathrm{A}\,\mathrm{R}^{d/p}\,\mathrm{N}$.

Ill-behaved regime: $\mathrm{RN} < 1$**.** The envelope term dominates: $\|\nabla f\|_{\mathrm{L}^p} \sim \mathrm{A}\,\mathrm{R}^{d/p-1}$. The function is so narrow that it does not complete even one oscillation. The "frequency" N is invisible, and the gradient is controlled by the compression 1/R of the bump.

The following plot illustrates both regimes in the 1D slice along direction e (so $d = 1$ for the picture; the analysis above is in general $\mathbb{R}^d$).



The dividing line between these regimes is the **uncertainty principle for functions**:

$$\boxed{\mathrm{R}\cdot\mathrm{N} \gtrsim 1.} \tag{276}$$

A function oscillating at frequency N must spread over at least one wavelength in the oscillating direction: its spatial scale satisfies $\mathrm{R} \gtrsim 1/\mathrm{N}$. This is the same constraint in any dimension, since it

is a statement about a single oscillating direction, not about volume. The factor of d that enters the Sobolev exponent comes from a *different* place: it is the volume scaling $\|f\|_{L^p} \sim AR^{d/p}$ above, not from uncertainty itself.

4.68 REMARK (Classical uncertainty and Sobolev embedding). The classical Fourier uncertainty principle states: if f is concentrated in an interval of width Δx, then its Fourier transform $\hat{f}$ must be spread over a frequency range $\Delta\xi$ with $\Delta x \cdot \Delta\xi \geq 1$. A function cannot be simultaneously localized in both space and frequency.

The Sobolev embedding theorem is the *quantitative* form of this principle. Trading s derivatives of regularity (frequency control) for integrability (spatial behavior) is precisely the trade-off between frequency localization and spatial localization. The critical exponent $1/q = 1/p - s/d$ quantifies the exchange rate, combining the uncertainty bound $RN \gtrsim 1$ with volume scaling in d dimensions.

DERIVING THE CRITICAL EXPONENT FROM THE UNCERTAINTY PRINCIPLE Work in $\mathbb{R}^d$ with the model function above. For $W^{s,p}$, the highest-order term dominates in the well-behaved regime:

$$\|f\|_{W^{s,p}} \sim A\, R^{d/p}\, N^s. \tag{277}$$

Normalize $\|f\|_{W^{s,p}} = 1$, so $A \sim R^{-d/p} N^{-s}$. The L^q norm is then

$$\|f\|_{L^q} \sim A\, R^{d/q} = R^{d/q-d/p}\, N^{-s}. \tag{278}$$

We want to find the largest q such that $\|f\|_{L^q}$ stays bounded across *all* well-behaved functions (all R, N with $RN \geq 1$). The most dangerous case is the extremal one, where the function is as concentrated as uncertainty allows: $R = 1/N$. Substituting:

$$\|f\|_{L^q} \sim N^{d/p-d/q-s}. \tag{279}$$

This is bounded for all $N \geq 1$ iff $d/p - d/q - s \leq 0$, i.e. $1/q \geq 1/p - s/d$. The critical case gives exactly the Sobolev exponent:

$$\frac{1}{q} = \frac{1}{p} - \frac{s}{d} \qquad \Longleftrightarrow \qquad q = p^* = \frac{dp}{d - sp}. \tag{280}$$

Functions at the uncertainty boundary $R = 1/N$ are the **extremizers** of the embedding: as concentrated as the uncertainty principle allows, making the L^q norm as large as possible for a given $W^{s,p}$ norm. The Sobolev embedding theorem says that even these extremal functions have bounded L^{p^*} norm.

Notice what each ingredient contributed:

- **Uncertainty** ($RN \gtrsim 1$) picked out the extremizer $R = 1/N$.
- **Volume scaling** ($\|f\|_{L^p} \sim AR^{d/p}$) is where the factor d enters, yielding the exchange rate $1/d$ per derivative.

A purely dimensional derivation of the same Sobolev exponent, via rescaling $u_\lambda(x) = u(\lambda x)$, is worked through in **??**.

The general embedding theorem

4.69 THEOREM (Sobolev embedding theorem). *Let $\Omega \subseteq \mathbb{R}^d$ be open with suitable regularity (for instance Lipschitz boundary, or $\Omega = \mathbb{R}^d$).*

1. ***Subcritical case*** *($kp < d$):* $\mathrm{W}^{k,p}(\Omega) \hookrightarrow \mathrm{L}^{p^*}(\Omega)$, *where $p^* = dp/(d-kp)$ is the **Sobolev conjugate**. Controlling k derivatives in L^p gives integrability up to L^{p^*}.*

2. ***Critical case*** *($kp = d$):* $\mathrm{W}^{k,p}(\Omega) \hookrightarrow \mathrm{L}^q(\Omega)$ *for all $q < \infty$ (but not L^∞ in general). The function is "almost continuous" but may have logarithmic singularities.*

3. ***Supercritical case*** *($kp > d$):* $\mathrm{W}^{k,p}(\Omega) \hookrightarrow \mathrm{C}^{m,\alpha}(\overline{\Omega})$, *where the orders of smoothness and Hölder regularity are determined by the **Sobolev number***

$$s \;=\; k - \frac{d}{p} \;>\; 0, \tag{281}$$

which records "how much regularity is left over after spending d/p derivatives on dimension." The split is

- *If $d/p \notin \mathbb{Z}$: $m = \lfloor s \rfloor$ and $\alpha = s - \lfloor s \rfloor \in (0,1)$. Equivalently $m = k - \lfloor d/p \rfloor - 1$ and $\alpha = \lfloor d/p \rfloor + 1 - d/p$, so that $m + \alpha = s$.*
- *If $d/p \in \mathbb{Z}$ (boundary case): $m = k - d/p - 1$ and the embedding holds for every $\alpha \in (0,1)$, but in general not for $\alpha = 1$ (Lipschitz can fail).*

Controlling enough derivatives in L^p forces continuity, and even Hölder regularity, with the leftover budget s split between integer smoothness (m) and a fractional Hölder exponent (α).

Each clause is the intuition made precise. The subcritical exponent p^* is the one forced by scaling and saturated by the uncertainty extremizers. The supercritical continuity statement is the 1D Poincaré phenomenon extended to general d once enough derivatives are controlled to push p^* past infinity. The critical case is the knife edge between them.

Sketch. This is a sketch, not a full proof. Every case is the 1D Poincaré calculation re-used in d dimensions.

Recall the 1D identity $u(x) = \int_0^x u'(t)\,dt$, which fed either Cauchy-Schwarz (to get Hölder continuity) or integration in x (to get L^p control). In $\mathbb{R}^d$ we apply the same identity **along each coordinate axis**: for each direction i,

$$|u(x)| \;\le\; \int |\partial_i u|\,dy_i. \tag{282}$$

The two regimes then split as follows.

- **Supercritical ($kp > d$, Morrey).** Apply the 1D identity along a single line segment joining x and y, and estimate with Hölder. The result, $|u(x) - u(y)| \le \mathrm{C}\,|x-y|^{1-d/p}\,\|\nabla u\|_{\mathrm{L}^p}$, is verbatim the Poincaré Hölder estimate with the exponent adjusted for dimension.

- **Subcritical ($kp < d$, Gagliardo-Nirenberg-Sobolev).** Multiply the d axial bounds, take the $1/(d-1)$ power, and integrate. This is the same "integrate the pointwise bound" move from Poincaré, but done simultaneously in all d directions, and it delivers $\|u\|_{L^{d/(d-1)}} \le C\|\nabla u\|_{L^1}$. The case $p > 1$ follows by applying this to $|u|^\gamma$; iteration in k raises the exponent by $1/d$ each step.
- **Critical ($kp = d$).** The formula predicts $p^* = \infty$, but L^∞ narrowly fails (think $\log\log(1/|x|)$). Interpolating the subcritical estimates on either side recovers L^q for every finite q.

□

4.10.4 Compact embeddings: Rellich-Kondrachov

The Sobolev embedding theorem tells us that a bounded set in $W^{k,p}$ is bounded in L^{p^*}. However, **bounded** in an infinite-dimensional space does not mean **precompact** (Remark 1.32). The remarkable fact is that if we ask for slightly less integrability ($q < p^*$ instead of $q = p^*$), the embedding becomes **compact**: bounded sequences not only stay bounded but have convergent subsequences.

4.70 THEOREM (Rellich-Kondrachov compactness). *Let $\Omega \subset \mathbb{R}^d$ be bounded with Lipschitz boundary, and let $1 \le p < \infty$, $k \ge 1$.*

1. ***Subcritical ($kp < d$).** With $1/p^* = 1/p - k/d$, the embedding $W^{k,p}(\Omega) \hookrightarrow L^q(\Omega)$ is **compact** for every $1 \le q < p^*$. The endpoint $q = p^*$ is continuous but **not compact**.*
2. ***Critical ($kp = d$).** The embedding $W^{k,p}(\Omega) \hookrightarrow L^q(\Omega)$ is **compact** for every $1 \le q < \infty$. (Continuous into BMO, but not into L^∞.)*
3. ***Supercritical ($kp > d$).** Writing $\alpha = k - d/p \in (0, 1]$ (non-integer part), the embedding $W^{k,p}(\Omega) \hookrightarrow C^{0,\beta}(\overline{\Omega})$ is **compact** for every $0 \le \beta < \alpha$ (and in particular the embedding into $C(\overline{\Omega})$ is compact). The endpoint $\beta = \alpha$ is continuous but **not compact**.*

The pattern is uniform across regimes: **compactness holds strictly below the sharp embedding exponent, and fails at it.** The obstruction at the endpoint is always the same scaling/concentration phenomenon — the critical embedding is scale-invariant, so bumps $u_n(x) = \lambda_n^\alpha \varphi(\lambda_n x)$ keep a fixed norm while escaping every compact set. Giving up an arbitrarily small amount of the exponent breaks the scale invariance and restores compactness.

Proof. We prove the key case $H^1(\Omega) \hookrightarrow L^2(\Omega)$ using Fourier analysis; the general case follows the same logic.

What we are actually proving. You might expect a compactness proof to take a bounded sequence $\{u_n\}$ and extract a Cauchy subsequence by hand. We do not do that. Instead, we use the characterization

$$\text{precompact in a complete metric space} \iff \text{totally bounded} \tag{283}$$

(Definition 1.28), where *totally bounded* means: for every $\varepsilon > 0$, the set is covered by **finitely many** ε-balls. So we never mention sequences. Instead we show that the H^1-unit ball B, viewed inside L^2, is approximately finite-dimensional: at any resolution $\varepsilon > 0$, every element of B is within ε of a finite set.

This is a concrete instance of the general Kolmogorov-Riesz-Fréchet criterion (Theorem 1.34): a bounded set in L^p is precompact iff it is equicontinuous under translation and tight. The gradient bound in the H^1-norm is precisely what supplies the equicontinuity condition, since $\|\tau_h u - u\|_{L^2} \le |h|\,\|\nabla u\|_{L^2}$,

and the bounded domain supplies tightness for free. Seen this way, Rellich-Kondrachov is not a miracle: it is Kolmogorov-Riesz with the hypothesis (2) automatically verified by gradient control. The Fourier-cutoff proof below is one convenient way to package the argument.

Once total boundedness is in hand, subsequence extraction is automatic: any bounded sequence lands in B, and a totally bounded set in a complete space has convergent subsequences for free (this is the ($\Leftarrow$) direction of the equivalence above). So the proof below is shorter and more structural than any direct subsequence argument would be, and it explains *why* compactness holds: because the derivative bound crushes the infinite tail of Fourier modes down to a size we can ignore, leaving only a finite-dimensional piece.

The plan in two lines:

1. Split $u = u_{\le N} + u_{>N}$ at Fourier cutoff N.
2. The tail $u_{>N}$ is L^2-small uniformly in $u \in B$, and the head $u_{\le N}$ lives in a finite-dimensional subspace.

That is the whole argument. The rest is bookkeeping.

Let $B = \{u \in H^1(\Omega) : \|u\|_{H^1} \le 1\}$ and consider B as a subset of $L^2(\Omega)$. Working on a bounded domain (or torus) with Fourier basis $\{e_k\}$, every $u \in B$ has an expansion $u = \sum_k \hat{u}_k e_k$ with

$$\sum_k (1 + |k|^2)\,|\hat{u}_k|^2 = \|u\|_{H^1}^2 \le 1. \tag{284}$$

Split u into low and high frequencies at a cutoff N:

$$u = \underbrace{\sum_{|k|\le N} \hat{u}_k\, e_k}_{u_{\le N}} + \underbrace{\sum_{|k|>N} \hat{u}_k\, e_k}_{u_{>N}}. \tag{285}$$

The high-frequency tail is uniformly small. For every $u \in B$:

$$\|u_{>N}\|_{L^2}^2 = \sum_{|k|>N} |\hat{u}_k|^2 \le \frac{1}{1+N^2} \sum_{|k|>N} (1 + |k|^2)\,|\hat{u}_k|^2 \le \frac{1}{1+N^2}. \tag{286}$$

This bound is **uniform over all** $u \in B$: it depends only on N, not on the particular function. Given any $\varepsilon > 0$, choose N large enough that $1/(1+N^2) < \varepsilon^2$.

The low-frequency part is finite-dimensional. The projection $u \mapsto u_{\le N}$ maps B into the span of the finitely many modes $\{e_k : |k| \le N\}$, which is a finite-dimensional subspace of L^2. A bounded set in a finite-dimensional space is precompact (Bolzano-Weierstrass).

Combining: every element of B is within ε (in L^2) of the precompact set $\{u_{\le N} : u \in B\}$. A set that can be approximated to arbitrary accuracy by precompact sets is itself precompact. So B is precompact in L^2, i.e. the embedding is compact.

□

How each hypothesis defeats each pathology The proof makes explicit how the two hypotheses of Rellich-Kondrachov match the two pathologies (P1), (P2) identified at the start of the chapter:

- **The H^1 bound kills oscillation (P1).** The decay $|\hat{u}_k|^2 \lesssim 1/(1+|k|^2)$ forces the high-frequency tail to be uniformly small. Sequences like $\sin(n\pi x)$ are ruled out because their derivatives blow up.
- **The bounded domain kills translation (P2).** There is nowhere for the mass to escape to, and the Fourier basis $\{e_k\}$ is indexed by a discrete set, so "low frequency" means "finitely many modes."

Both hypotheses are essential. On $\mathbb{R}^d$, $H^1 \hookrightarrow L^2$ is **not** compact, because translation sequences defeat it. On a bounded domain without the derivative bound, the identity $L^2 \hookrightarrow L^2$ is not compact in infinite dimensions, because oscillations defeat it.

Rellich-Kondrachov is the Sobolev-space counterpart of the Arzelà-Ascoli theorem (Theorem 1.33). In Arzelà-Ascoli, equicontinuity prevents rapid spatial oscillation; here, the H^1 bound prevents rapid *frequency* oscillation. In both cases, a regularity condition suppresses one pathology and a bounded domain suppresses the other. Together they force approximate finite-dimensionality.

Connection to compact operators In the language of the compact operators chapter (Definition 2.44), the embedding $\iota : H^1(\Omega) \hookrightarrow L^2(\Omega)$ **is** a compact operator: it maps the bounded H^1 unit ball to a precompact subset of L^2. This is why compact embeddings and compact operators are interchangeable language.

The payoff for PDE is that compact embeddings upgrade weak convergence to strong convergence, which is essential for passing to the limit through nonlinearities. See the applications chapter for the full development: the weak formulation via Lax-Milgram (Theorem 4.96) and the nonlinear PDE example (**??**) both rely on Rellich-Kondrachov as a critical step.

4.11 *Sobolev Spaces as Banach Algebras*

Big Idea

The spaces $C^0(\overline{\Omega})$ and $C^\infty(\overline{\Omega})$ are algebras under pointwise multiplication: the product of two continuous (or smooth) functions is again continuous (or smooth). Sobolev spaces inherit this property precisely when they embed into C^0, which happens when $s > d/2$. The algebra structure connects the embedding theory from the previous chapter to the bandedness of multiplication operators in spectral methods.

4.11.1 Why multiplication matters

So far we have used Sobolev spaces as *linear* function spaces: we add functions, take limits, and apply linear operators. But functions can also be **multiplied**, and a natural question arises: when is the pointwise product uv again in the same Sobolev space? The answer turns out to connect the embedding theory we just developed to the structure of multiplication operators in spectral methods.

4.11.2 The algebra property from embeddings

An **algebra** over $\mathbb{R}$ (or $\mathbb{C}$) is a vector space A that also carries a **multiplication**, a bilinear map $A \times A \to A$ that is compatible with the linear structure. You already know several:

Space	Multiplication	Submultiplicative norm?
$\mathbb{R}^{n\times n}$	Matrix product	$\lvert AB\rvert \le \lvert A\rvert\lvert B\rvert$ (operator norm)
$C[0,1]$	Pointwise: $(fg)(x) = f(x)g(x)$	$\lvert fg\rvert_\infty \le \lvert f\rvert_\infty \lvert g\rvert_\infty$
$L^\infty(\Omega)$	Pointwise a.e.	$\lvert fg\rvert_\infty \le \lvert f\rvert_\infty \lvert g\rvert_\infty$
$\ell^1(\mathbb{Z})$	Convolution	$\lvert f * g\rvert_1 \le \lvert f\rvert_1 \lvert g\rvert_1$ (Young)

A **Banach algebra** is a Banach space whose norm is **submultiplicative**: $\|uv\| \le C\|u\|\|v\|$. This single inequality says that the multiplication is *continuous* as a bilinear map, or equivalently, that the space is **closed under multiplication** with quantitative control.

Not every function space is an algebra. The product of two L^2 functions is generally only in L^1 (by Cauchy–Schwarz, $\|fg\|_{L^1} \le \|f\|_{L^2}\|g\|_{L^2}$, but fg need not be in L^2). So L^2 with pointwise multiplication is *not* an algebra; multiplication takes you outside the space. The question for Sobolev spaces is: **does the extra regularity fix this?**

The key observation is that $C^0(\overline{\Omega})$ *is* a Banach algebra (the sup norm is submultiplicative), so any Banach space that embeds continuously into C^0 inherits the algebra property. If $X \hookrightarrow C^0$ with $\|u\|_\infty \le C\|u\|_X$, then

$$\|uv\|_X \le C'\|u\|_X\|v\|_X \tag{287}$$

whenever the product rule and embedding estimate can be combined. The Sobolev embedding theorem (Theorem 4.69) gives $H^s \hookrightarrow C^0$ exactly when $s > d/2$. This is the threshold.

4.71 THEOREM (Sobolev multiplication (Banach algebra property)). *Let $\Omega \subseteq \mathbb{R}^d$ be open with Lipschitz boundary (or $\Omega = \mathbb{R}^d$). If $s > d/2$, then $\mathrm{H}^s(\Omega)$ is a **Banach algebra** under pointwise multiplication: there exists a constant $\mathrm{C} > 0$ such that*

$$\|uv\|_{\mathrm{H}^s} \leq \mathrm{C}\,\|u\|_{\mathrm{H}^s}\,\|v\|_{\mathrm{H}^s} \quad \text{for all } u, v \in \mathrm{H}^s(\Omega). \tag{288}$$

In particular, $\mathrm{H}^s(\Omega)$ is closed under multiplication.

Proof. We prove the case $s = k$ a positive integer; the fractional case requires interpolation. The strategy is to apply the **Leibniz rule** (product rule for higher derivatives) and then split each term so that one factor is controlled in L^∞ via the Sobolev embedding.

Step 1: Leibniz rule. For any multi-index $|\alpha| \leq k$, the general Leibniz rule gives

$$\mathrm{D}^\alpha(uv) = \sum_{\beta \leq \alpha} \binom{\alpha}{\beta} \mathrm{D}^\beta u \cdot \mathrm{D}^{\alpha-\beta} v. \tag{289}$$

This is the higher-dimensional product rule: it distributes the α derivatives between u and v in all possible ways. (For $k = 1$ in one dimension, this is simply $(uv)' = u'v + uv'$.)

Step 2: Estimating each term. Each term in the Leibniz sum has the form $\mathrm{D}^\beta u \cdot \mathrm{D}^{\alpha-\beta} v$ where $|\beta| + |\alpha - \beta| = |\alpha| \leq k$. We estimate in L^2 using Hölder's inequality with exponents 2 and ∞:

$$\|\mathrm{D}^\beta u \cdot \mathrm{D}^{\alpha-\beta} v\|_{\mathrm{L}^2} \leq \|\mathrm{D}^\beta u\|_{\mathrm{L}^2}\,\|\mathrm{D}^{\alpha-\beta} v\|_{\mathrm{L}^\infty}. \tag{290}$$

Since $|\alpha - \beta| \leq k$ and $k > d/2$, the function $\mathrm{D}^{\alpha-\beta} v$ has at least $k - |\alpha - \beta|$ remaining derivatives in L^2. When $|\alpha - \beta| \leq k$, we can apply the Sobolev embedding to get

$$\|\mathrm{D}^{\alpha-\beta} v\|_{\mathrm{L}^\infty} \leq \mathrm{C}\,\|v\|_{\mathrm{H}^k}. \tag{291}$$

(More precisely, $\mathrm{D}^{\alpha-\beta} v \in \mathrm{H}^{k-|\alpha-\beta|}$ and for the terms where $k - |\alpha - \beta| > d/2$ we embed directly into L^∞; for the remaining terms we swap the roles of u and v.)

Step 3: Combining. Summing over all $|\alpha| \leq k$ gives

$$\|uv\|_{\mathrm{H}^k} = \sum_{|\alpha| \leq k} \|\mathrm{D}^\alpha(uv)\|_{\mathrm{L}^2} \leq \mathrm{C} \sum_{|\alpha| \leq k} \sum_{\beta \leq \alpha} \|\mathrm{D}^\beta u\|_{\mathrm{L}^2}\,\|v\|_{\mathrm{H}^k} \leq \mathrm{C}'\,\|u\|_{\mathrm{H}^k}\,\|v\|_{\mathrm{H}^k}. \tag{292}$$

□

4.72 EXAMPLE (in one dimension). H^1

For $d = 1$ and $s = 1$, the condition $s > d/2$ becomes $1 > 1/2$, which holds. The proof is completely explicit:

$$\|uv\|_{\mathrm{H}^1} = \|uv\|_{\mathrm{L}^2} + \|(uv)'\|_{\mathrm{L}^2}. \tag{293}$$

For the first term: $\|uv\|_{L^2} \le \|u\|_{L^\infty}\|v\|_{L^2}$. For the second, the product rule gives $(uv)' = u'v + uv'$, so

$$\|(uv)'\|_{L^2} \le \|u'\|_{L^2}\|v\|_{L^\infty} + \|u\|_{L^\infty}\|v'\|_{L^2}. \tag{294}$$

By the Sobolev embedding $H^1(0,1) \hookrightarrow C[0,1]$, we have $\|u\|_{L^\infty} \le C\|u\|_{H^1}$ and similarly for v. Combining:

$$\|uv\|_{H^1} \le C\|u\|_{H^1}\|v\|_{H^1}. \tag{295}$$

Each term in the estimate pairs **derivatives on one factor** with L^∞ **control on the other**. The Sobolev embedding converts the L^∞ norms into H^1 norms, closing the estimate.

4.73 REMARK (Why). $s > d/2$ **is sharp**

The threshold $s > d/2$ cannot be improved. For $s = d/2$, the space $H^{d/2}(\mathbb{R}^d)$ does **not** embed into L^∞ (functions in $H^{d/2}$ can have logarithmic singularities), and multiplication fails to be continuous. For instance, in $d = 2$, the function $u(x) = \log\log(1/|x|)$ near the origin belongs to $H^1(\Omega)$ but u^2 does not.

4.11.3 Multiplication in frequency space: why regularity implies bandedness

The Banach algebra property tells us that multiplication by $u \in H^s$ is a bounded operator on H^s. But *what does this operator look like* when we expand in a spectral basis? The answer reveals a striking structural property: **regularity forces approximate bandedness**.

To see this most cleanly, work on the torus $\mathbb{T} = [0, 2\pi)$ with Fourier basis $\{e_k(x) = e^{ikx}\}_{k\in\mathbb{Z}}$. Expand $u = \sum_k \hat{u}_k e_k$. The multiplication operator $M_u : v \mapsto uv$ has the matrix representation

$$(M_u)_{jk} = \langle M_u e_k, e_j\rangle = \langle u \cdot e_k, e_j\rangle = \hat{u}_{j-k}. \tag{296}$$

4.74 OBSERVATION (Multiplication is convolution in frequency). The matrix of M_u in the Fourier basis is a **Toeplitz matrix**: the (j,k)-entry depends only on $j - k$, and equals the Fourier coefficient $\hat{u}_{j-k}$. Multiplying by u in physical space is convolution by $\hat{u}$ in frequency space.

Now the connection to regularity becomes immediate. If $u \in H^s$, then

$$\sum_k (1 + k^2)^s |\hat{u}_k|^2 < \infty \quad \Longrightarrow \quad |\hat{u}_k| \lesssim |k|^{-s-1/2} \text{ (by Cauchy–Schwarz).} \tag{297}$$

In particular, the off-diagonal entries of M_u decay:

$$|(M_u)_{jk}| = |\hat{u}_{j-k}| \lesssim |j - k|^{-s-1/2}. \tag{298}$$

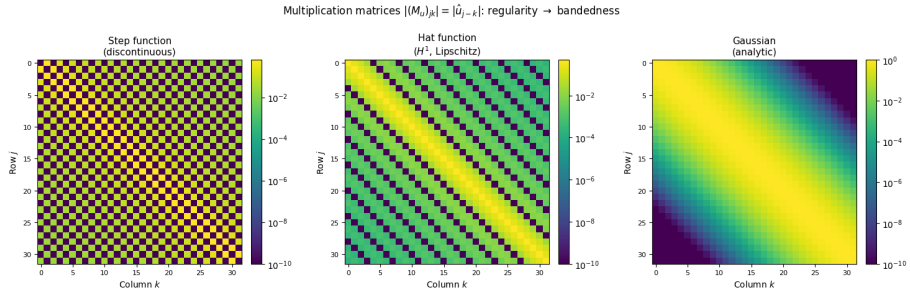
4.75 PROPOSITION (Off-diagonal decay of multiplication matrices). *Let $u \in H^s(\mathbb{T})$ with $s \ge 0$. The multiplication operator M_u in the Fourier basis satisfies*

$$|(\mathrm{M}_u)_{jk}| \leq \frac{\mathrm{C}\|u\|_{\mathrm{H}^s}}{(1+|j-k|)^{s+1/2}}. \tag{299}$$

In particular:

- $u \in \mathrm{H}^1$*: entries decay like* $|j-k|^{-3/2}$*, so off-diagonal entries are summable.*
- $u \in \mathrm{H}^2$*: decay like* $|j-k|^{-5/2}$*, giving rapid off-diagonal decay.*
- $u \in \mathrm{C}^\infty$*: superalgebraic decay, meaning the matrix is* ***numerically banded*** *(entries drop below machine precision within a few diagonals).*
- *u analytic:* ***exponential*** *off-diagonal decay, making it essentially a banded matrix.*

The following picture makes this concrete.



The multiplication matrix in the Fourier basis for three functions of increasing regularity. Left: a step function (discontinuous) produces a full, slowly-decaying matrix. Center: a hat function (H^1) *produces algebraic off-diagonal decay. Right: a Gaussian (analytic) produces exponential decay, so the matrix is effectively banded. The color scale is logarithmic.*

THE INTUITION: REGULARITY IS FREQUENCY LOCALIZATION Why should smoothness produce banded multiplication matrices? The intuition comes from **three equivalent ways to say the same thing**:

1. **Regularity = frequency concentration.** A function in H^s has Fourier coefficients decaying like $|k|^{-s-1/2}$. Higher regularity means the function's energy is *concentrated at low frequencies*.

2. **Multiplication = convolution in frequency.** The matrix M_u is Toeplitz with entries $\hat{u}_{j-k}$. Applying M_u to a vector of Fourier coefficients is **convolution** with the sequence $(\hat{u}_k)$.

3. **Convolving with something narrow produces something narrow.** If the sequence $(\hat{u}_k)$ is concentrated near $k = 0$ (because u is smooth), then convolution with it only couples mode k to nearby modes $k \pm \Delta k$, which is exactly the statement that M_u is approximately banded.

Putting these together: **smooth functions act locally in frequency space**. Multiplying by a smooth function is like a *short-range interaction* between Fourier modes. Multiplying by a rough function is a long-range interaction that couples all modes to all other modes.

4.76 REMARK (Analogy with integral operators). This pattern is familiar from integral operators. A kernel $\mathrm{K}(x,y)$ produces a banded matrix in some basis when the kernel is "close to diagonal." For multiplication, the kernel is $\mathrm{K}(x,y) = u(x)\delta(x-y)$, which is as diagonal as possible in physical space. But expanding in the Fourier basis, this perfectly diagonal operator becomes the Toeplitz matrix $\hat{u}_{j-k}$,

which is banded only when u is smooth. The same operator can be diagonal in one basis and banded in another; the bandedness in frequency reflects the smoothness in space.

CONNECTION TO SPECTRAL AND COLLOCATION METHODS This structure is exactly what makes **spectral methods** efficient for smooth problems.

In a pseudospectral (collocation) method, we want to compute the Fourier (or Chebyshev) coefficients of a product uv. The naive approach is to form M_u and multiply, an $\mathrm{O}(\mathrm{N}^2)$ operation. But the collocation approach exploits a factorization:

$$\mathrm{M}_u = \mathrm{F}^{-1}\ \mathrm{diag}(\mathbf{u})\,\mathrm{F} \tag{300}$$

where F is the DFT matrix and $\mathrm{diag}(\mathbf{u})$ contains the values of u at collocation points. **Multiplication is diagonal in physical space**, so we:

1. Transform $\hat{v} \to v$ at collocation points (inverse FFT, $\mathrm{O}(\mathrm{N}\log\mathrm{N})$),
2. Multiply pointwise: $w_j = u_j v_j$ ($\mathrm{O}(\mathrm{N})$),
3. Transform back $w \to \hat{w}$ (forward FFT, $\mathrm{O}(\mathrm{N}\log\mathrm{N})$).

The total cost is $\mathrm{O}(\mathrm{N}\log\mathrm{N})$ instead of $\mathrm{O}(\mathrm{N}^2)$. But here is the subtle point: this procedure is exact only for trigonometric polynomials of degree $\leq \mathrm{N}$. For general smooth functions, the product uv may have Fourier modes up to degree 2N, and the modes above N get **aliased** back into the $[-\mathrm{N}, \mathrm{N}]$ range.

The Banach algebra property tells us this aliasing is harmless for smooth functions: since $|\hat{u}_k| \lesssim |k|^{-s-1/2}$, the aliased modes are exponentially small (for analytic u) or at least rapidly decaying (for Sobolev u). The error from aliasing is controlled by the off-diagonal decay of M_u, precisely the entries we truncate by working with a finite matrix.

4.77 REMARK (Chebyshev vs. Fourier). For Chebyshev spectral methods, the multiplication matrix is not Toeplitz but **almost-banded**: the Chebyshev expansion $u = \sum_k a_k \mathrm{T}_k$ gives a multiplication matrix with entries involving the *linearization coefficients* of Chebyshev products ($\mathrm{T}_j\mathrm{T}_k = \frac{1}{2}(\mathrm{T}_{j+k} + \mathrm{T}_{|j-k|})$), and the matrix has the same off-diagonal decay governed by the rate of decay of the Chebyshev coefficients a_k. The story is identical: regularity of u controls the coefficient decay, which controls the bandwidth. This is the algebraic reason why collocation methods produce banded or nearly-banded discrete operators for smooth problems.

4.78 REMARK (The Banach algebra property as a spectral guarantee). From the numerical analyst's perspective, the Banach algebra property of H^s is a **stability guarantee for nonlinear spectral methods**: if $u, v \in \mathrm{H}^s$ and you form their product spectrally (with or without aliasing), the result stays in H^s with a controlled norm. This is essential for the convergence theory of spectral methods applied to nonlinear PDEs: it ensures that the nonlinearity does not create uncontrolled high-frequency content that would destabilize the computation.

4.12 *Sobolev Duality and Structure*

4.12.1 Duality, reflexivity, and the Sobolev scale

One of the main rewards of working with Sobolev spaces instead of the full distributional framework is that Sobolev spaces are **Banach spaces**, and for $1 < p < \infty$, they are **reflexive**. This means that the entire duality machinery from the duality chapter (weak topology, weak-$*$ topology, Banach–Alaoglu, weak compactness) applies directly, with no need for the inductive limit subtleties of $\mathcal{D}'(\Omega)$.

THE DUAL OF $W_0^{k,p}(\Omega)$

4.79 DEFINITION (and negative Sobolev spaces). $W_0^{k,p}(\Omega)$

Let $W_0^{k,p}(\Omega)$ denote the **closure of** $C_c^\infty(\Omega)$ in $W^{k,p}(\Omega)$. For $1 < p < \infty$ with conjugate exponent $p' = p/(p-1)$, the **negative Sobolev space** is defined as the dual:

$$W^{-k,p'}(\Omega) := \left(W_0^{k,p}(\Omega)\right)^*. \tag{301}$$

The duality pairing extends the distributional pairing: for $T \in W^{-k,p'}(\Omega)$ and $\varphi \in C_c^\infty(\Omega) \subset W_0^{k,p}(\Omega)$,

$$\langle T, \varphi \rangle_{W^{-k,p'} \times W_0^{k,p}} = \langle T, \varphi \rangle_{\mathcal{D}' \times \mathcal{D}}. \tag{302}$$

So every element of $W^{-k,p'}$ is a distribution, but a tame one: it extends continuously from test functions to all of $W_0^{k,p}$.

4.80 REMARK. The distinction between $W_0^{k,p}(\Omega)$ and $W^{k,p}(\Omega)$ matters on bounded domains: $W_0^{k,p}$ encodes zero boundary conditions (it is the closure of compactly supported functions), while $W^{k,p}$ does not. On $\mathbb{R}^d$, they coincide: $W_0^{k,p}(\mathbb{R}^d) = W^{k,p}(\mathbb{R}^d)$ because C_c^∞ is dense in $W^{k,p}$ on the whole space.

CONCRETE CHARACTERIZATION The dual characterization makes negative Sobolev spaces concrete. An element $T \in W^{-k,p'}(\Omega)$ can always be written as

$$\langle T, u \rangle = \sum_{|\alpha| \le k} \int_\Omega f_\alpha \, D^\alpha u \, dx \tag{303}$$

for some $f_\alpha \in L^{p'}(\Omega)$. Conversely, every such expression defines an element of $W^{-k,p'}$. In other words:

$$W^{-k,p'}(\Omega) = \left\{ \sum_{|\alpha| \le k} D^\alpha f_\alpha : f_\alpha \in L^{p'}(\Omega) \right\} \tag{304}$$

where the derivatives are taken in the distributional sense. This is a representation theorem analogous to the Riesz representation (Theorem 3.22): every continuous linear functional on $W_0^{k,p}$

comes from $L^{p'}$ data, with the pairing mediated by distributional derivatives.

THE SOBOLEV SCALE Sobolev spaces fit into the distributional framework as a graded scale of increasingly regular subspaces of $\mathcal{D}'(\Omega)$:

$$\cdots \subset H^{-2}(\Omega) \subset H^{-1}(\Omega) \subset L^2(\Omega) \subset H^1(\Omega) \subset H^2(\Omega) \subset \cdots \tag{305}$$

The index s measures regularity: positive s means the function has s derivatives in L^2; negative s means the distribution is "mildly singular," living $|s|$ levels below L^2.

4.81 EXAMPLE (The Dirac delta in the Sobolev scale). The Dirac delta $\delta \in \mathcal{D}'(\mathbb{R}^d)$ is not in L^p for any p, but it belongs to negative Sobolev spaces. In dimension d with $p = 2$, $\delta \in H^{-s}(\mathbb{R}^d)$ for $s > d/2$, because the Sobolev embedding theorem guarantees that pointwise evaluation is bounded on H^s when $s > d/2$:

$$|\langle \delta, \varphi \rangle| = |\varphi(0)| \leq C\|\varphi\|_{H^s}. \tag{306}$$

In one dimension, $\delta \in H^{-s}(\mathbb{R})$ for $s > 1/2$. In three dimensions, $\delta \in H^{-s}(\mathbb{R}^3)$ for $s > 3/2$. The more singular the object, the further down the scale it lives.

4.82 EXAMPLE (Derivatives shift the scale). Differentiation moves a distribution down the Sobolev scale: if $u \in H^s(\Omega)$, then $D^\alpha u \in H^{s-|\alpha|}(\Omega)$. For instance:

- If $u \in H^1$, then $u' \in L^2 = H^0$.
- If $u \in L^2$, then $u' \in H^{-1}$. The derivative exists as a distribution but is not a function.
- The Heaviside function $H \in L^2(0,1)$ has $H' = \delta_{1/2} \in H^{-1}(0,1)$.

The Sobolev index precisely tracks how many derivatives a function can "afford" before leaving L^2.

THE LIMITS OF THE SOBOLEV SCALE As k increases, H^{-k} grows larger, containing increasingly singular distributions. It is natural to ask whether the union exhausts all of $\mathcal{D}'(\Omega)$.

4.83 PROPOSITION. *The inclusion*

$$\bigcup_{k=0}^{\infty} H^{-k}(\Omega) \subsetneq \mathcal{D}'(\Omega) \tag{307}$$

*is strict. The Sobolev scale captures exactly the **finite-order** distributions. Distributions of **infinite order** (where the order of the continuity estimate grows without bound as the compact set K expands) escape every level of the scale.*

4.84 EXAMPLE (A distribution of infinite order). On $\mathbb{R}$, define

$$T = \sum_{n=1}^{\infty} \delta^{(n)}(x - n). \tag{308}$$

On a compact set containing $x = n$, the term $\delta^{(n)}(x-n)$ requires $k \geq n$ derivatives of the test function. No single k works for all compact sets, so $T \notin H^{-k}(\mathbb{R})$ for any k. But $T \in \mathcal{D}'(\mathbb{R})$: for any test function φ (which has compact support), only finitely many terms in the sum are nonzero, so $\langle T, \varphi \rangle$ is well-defined.

This is the precise sense in which $\mathcal{D}'(\Omega)$ is larger than the Sobolev scale: it accommodates distributions whose singularity **worsens without bound** across the domain. For such objects, the inductive limit topology of $\mathcal{D}$ is genuinely necessary; no single Banach space in the scale can hold them.

In practice, distributions arising from PDE theory almost always have finite order (often order ≤ 2), so the Sobolev scale is sufficient for most applications.

Reflexivity: weak and weak-* for free

4.85 THEOREM (Reflexivity of Sobolev spaces). *For $1 < p < \infty$ and $k \geq 0$, the spaces $W^{k,p}(\Omega)$ and $W_0^{k,p}(\Omega)$ are **reflexive** Banach spaces. In particular, $H^k(\Omega) = W^{k,2}(\Omega)$ is a reflexive Hilbert space.*

*For $p = 1$ or $p = \infty$, the Sobolev spaces are **not** reflexive.*

Proof. The map $\Phi : W^{k,p}(\Omega) \to \prod_{|\alpha| \leq k} L^p(\Omega)$ defined by $\Phi(u) = (D^\alpha u)_{|\alpha| \leq k}$ is an isometric embedding (by definition of the Sobolev norm). The product of reflexive spaces is reflexive, and L^p is reflexive for $1 < p < \infty$. A closed subspace of a reflexive space is reflexive, so $W^{k,p}$ is reflexive.

□

Reflexivity means the canonical embedding $J : W^{k,p} \to (W^{k,p})^{**}$ is surjective, so the double dual adds nothing new. This has immediate consequences from the duality chapter:

4.86 COROLLARY (Weak compactness in Sobolev spaces). *For $1 < p < \infty$:*

1. ***Banach–Alaoglu applies directly.*** *The closed unit ball of $W^{k,p}(\Omega)$ is weak-* compact in $(W^{-k,p'})^* = W^{k,p}$. But because $W^{k,p}$ is reflexive, the weak and weak-* topologies coincide, so the ball is **weakly compact**.*

2. ***Every bounded sequence has a weakly convergent subsequence.*** *If (u_n) is bounded in $W^{k,p}(\Omega)$, there exists a subsequence $u_{n_j} \rightharpoonup u$ in $W^{k,p}(\Omega)$.*

3. ***Weak sequential completeness.*** *Every weakly Cauchy sequence in $W^{k,p}(\Omega)$ converges weakly.*

This is the payoff: in $\mathcal{D}'(\Omega)$, we needed the Montel property and the non-trivial inductive limit machinery to extract convergent subsequences. In $W^{k,p}$ for $1 < p < \infty$, it is a direct consequence of reflexivity and Banach–Alaoglu, results we already proved in the duality chapter.

4.87 REMARK (The duality chapter toolkit, applied). Here is how the duality chapter results specialize to Sobolev spaces:

Duality chapter result	Sobolev space consequence
Weak topology $\sigma(X, X^*)$ is Hausdorff	Weak limits in $W^{k,p}$ are unique
Banach–Alaoglu: unit ball is weak-$*$ compact	Bounded sets in $W^{k,p}$ are weakly precompact ($1 < p < \infty$)
Reflexive $\Rightarrow$ weak = weak-$*$	No need to distinguish weak and weak-$*$ on $W^{k,p}$
Weak convergence + compact operator $\Rightarrow$ strong convergence	$u_n \rightharpoonup u$ in $H^1 \Rightarrow u_n \to u$ in L^2 (Rellich–Kondrachov)
Direct method: minimize over weakly compact sets	Existence of PDE solutions by energy minimization in H^1

The last row is the reason Sobolev spaces exist: they are the Banach spaces where the direct method works. Reflexivity gives weak compactness (subsequences exist), Rellich–Kondrachov gives compact embedding (the subsequences converge strongly in a weaker norm), and weak lower semicontinuity of the energy functional closes the argument.

WHY $p = 1$ AND $p = \infty$ ARE DIFFERENT For $p = 1$, the Sobolev space $W^{k,1}(\Omega)$ is not reflexive (just as L^1 is not reflexive). Its dual $(W_0^{k,1})^*$ contains measures and other singular objects. Bounded sequences in $W^{k,1}$ need not have weakly convergent subsequences. The direct method fails, and minimization problems over $W^{1,1}$ require the theory of functions of bounded variation (BV) instead.

For $p = \infty$, the space $W^{k,\infty}(\Omega)$ is essentially the Lipschitz functions (for $k = 1$), and its dual is again non-reflexive. The weak-$*$ topology on $(W^{k,\infty})^*$ does not coincide with the weak topology on $W^{k,\infty}$.

The range $1 < p < \infty$ is the sweet spot where duality theory works cleanly.

4.13 *Application: Differential Equations*

Big Idea

A fundamental solution E of a linear operator L solves $\mathrm{LE} = \delta$. Once you have it, convolution $u = \mathrm{E} * f$ solves $\mathrm{L}u = f$ for any source f. This is the continuous analogue of the superposition principle from ODEs: the integral replaces the sum, and δ replaces the basis vectors.

4.13.1 From ODE Superposition to PDE Convolution

In ODEs, if L is a linear operator and we can solve $\mathrm{L}y_i = f_i$ for individual sources f_i, then linearity gives $\mathrm{L}(\sum c_i y_i) = \sum c_i f_i$. This is the superposition principle for finite sums.

For PDEs, we replace the finite sum with an integral. Any source f can be decomposed as a continuous superposition of point sources:

$$f(x) = \int_{\mathbb{R}^d} f(y)\,\delta(x-y)\,dy. \tag{309}$$

This is just the sifting property of the delta, but read physically it says: f is a superposition of impulses $\delta(\cdot - y)$, each weighted by $f(y)$. If E is the response to a single impulse at the origin (i.e. $\mathrm{LE} = \delta$), then by translation $\mathrm{E}(\cdot - y)$ is the response to an impulse at y. Superposing all these responses:

$$u(x) = \int_{\mathbb{R}^d} f(y)\,\mathrm{E}(x-y)\,dy = (\mathrm{E} * f)(x). \tag{310}$$

Since L acts on x and the integral is over y, linearity gives

$$\mathrm{L}u(x) = \int_{\mathbb{R}^d} f(y)\,\mathrm{LE}(x-y)\,dy = \int_{\mathbb{R}^d} f(y)\,\delta(x-y)\,dy = f(x). \tag{311}$$

So $u = \mathrm{E} * f$ solves $\mathrm{L}u = f$. Convolution with the fundamental solution *is* the superposition principle, generalized from sums to integrals.

4.13.2 Fundamental Solutions

4.88 DEFINITION (Fundamental solution). Let $\mathrm{L} = \sum_{|\alpha| \le m} a_\alpha \mathrm{D}^\alpha$ be a linear partial differential operator with constant coefficients $a_\alpha \in \mathbb{R}$. A **fundamental solution** of L is a distribution $\mathrm{E} \in \mathcal{D}'(\mathbb{R}^d)$ satisfying

$$\mathrm{LE} = \delta. \tag{312}$$

The equation $\mathrm{LE} = \delta$ is understood in the distributional sense: $\langle \mathrm{LE}, \varphi \rangle = \langle \delta, \varphi \rangle = \varphi(0)$ for all $\varphi \in \mathcal{D}(\mathbb{R}^d)$. By the definition of distributional differentiation, this becomes $\langle \mathrm{E}, \mathrm{L}^*\varphi \rangle = \varphi(0)$, where L^* is the formal adjoint of L.

Why distributions are essential here. The equation $\mathrm{LE} = \delta$ has no meaning in classical analysis: the right-hand side is not a function. But even the left-hand side is misleading classically. Consider the Laplacian in $\mathbb{R}^3$: the function $\mathrm{E}(x) = -1/(4\pi|x|)$ satisfies $\Delta \mathrm{E}(x) = 0$ for every $x \neq 0$. Pointwise, the Laplacian cannot "see" the singularity at the origin. But as a distribution, E acts on test functions via integration, and the singularity contributes. When we compute $\langle \Delta \mathrm{E}, \varphi \rangle = \langle \mathrm{E}, \Delta \varphi \rangle$ and integrate by parts (excising a ball $\mathrm{B}_\varepsilon(0)$, applying Green's identity, taking $\varepsilon \to 0$), the boundary terms on the small sphere do not vanish: they converge to $\varphi(0)$. That is where the δ comes from. The distributional derivative detects a contribution at the singularity that pointwise differentiation misses entirely.

4.89 EXAMPLE (Fundamental solution of the Laplacian). The Laplacian $\Delta = \sum_{i=1}^{d} \partial_i^2$ in $\mathbb{R}^d$ has fundamental solution

$$\mathrm{E}(x) = \begin{cases} \dfrac{1}{2}|x| & d = 1, \\ \dfrac{1}{2\pi} \log |x| & d = 2, \\ \dfrac{-1}{d(d-2)\omega_d} |x|^{2-d} & d \geq 3, \end{cases} \tag{313}$$

where ω_d is the volume of the unit ball in $\mathbb{R}^d$. In three dimensions this gives the familiar $\mathrm{E}(x) = -\frac{1}{4\pi|x|}$.

Each of these is locally integrable (the singularity at the origin is integrable), so E defines a regular distribution. The equation $\Delta \mathrm{E} = \delta$ can be verified by integration by parts: for any $\varphi \in \mathcal{D}(\mathbb{R}^d)$,

$$\langle \Delta \mathrm{E}, \varphi \rangle = \langle \mathrm{E}, \Delta \varphi \rangle = \int_{\mathbb{R}^d} \mathrm{E}(x)\, \Delta \varphi(x)\, dx = \varphi(0). \tag{314}$$

The last equality uses Green's second identity on the domain $\{|x| > \varepsilon\}$ and a careful limit $\varepsilon \to 0$.

4.90 EXAMPLE (Fundamental solution of the heat operator). The heat operator $\mathrm{L} = \partial_t - \Delta$ in $\mathbb{R}^d \times (0, \infty)$ has fundamental solution

$$\mathrm{E}(x, t) = \frac{1}{(4\pi t)^{d/2}}\, e^{-|x|^2/4t}\, \mathrm{H}(t), \tag{315}$$

where $\mathrm{H}(t)$ is the Heaviside function. This is the **heat kernel**: a Gaussian that spreads over time. For any bounded continuous initial condition f, the convolution

$$u(x, t) = (\mathrm{E}(\cdot, t) * f)(x) = \int_{\mathbb{R}^d} \frac{1}{(4\pi t)^{d/2}}\, e^{-|x-y|^2/4t}\, f(y)\, dy \tag{316}$$

solves the initial value problem $\partial_t u - \Delta u = 0$, $u(\cdot, 0) = f$.

4.13.3 The Malgrange-Ehrenpreis Theorem

A natural question: does every linear constant-coefficient operator have a fundamental solution? The answer is yes.

4.91 THEOREM (Malgrange-Ehrenpreis). *Every non-zero linear partial differential operator with constant coefficients*

$$\mathrm{L} = \sum_{|\alpha| \le m} a_\alpha \mathrm{D}^\alpha, \qquad a_\alpha \in \mathbb{R}, \tag{317}$$

has a fundamental solution $\mathrm{E} \in \mathcal{D}'(\mathbb{R}^d)$.

This is a deep existence result, due independently to Malgrange and Ehrenpreis in the 1950s.

Proof. **(Sketch via Hahn-Banach.)** We need a distribution E satisfying $\langle \mathrm{E}, \mathrm{L}^*\varphi \rangle = \varphi(0)$ for all $\varphi \in \mathcal{D}(\mathbb{R}^d)$. Define the linear functional

$$\Phi : \mathrm{L}^*\Big(\mathcal{D}(\mathbb{R}^d)\Big) \to \mathbb{R}, \qquad \Phi(\mathrm{L}^*\varphi) = \varphi(0). \tag{318}$$

This is well-defined on the image of L^* (one must verify that $\mathrm{L}^*\varphi = 0$ implies $\varphi(0) = 0$, which is the hardest part of the proof and requires estimates on L^*). The key analytic step is showing that Φ is continuous with respect to the topology inherited from $\mathcal{D}(\mathbb{R}^d)$: there exist a compact set K and constants C, k such that

$$|\varphi(0)| \le \mathrm{C} \|\mathrm{L}^*\varphi\|_{\mathrm{C}^k(\mathrm{K})} \qquad \text{for all } \varphi \in \mathcal{D}(\mathbb{R}^d). \tag{319}$$

This estimate is the heart of the proof. Once it is established, the Hahn-Banach theorem extends Φ to a continuous linear functional E on all of $\mathcal{D}(\mathbb{R}^d)$. This E is a distribution satisfying $\mathrm{LE} = \delta$.

□

STRUCTURE OF THE GENERAL SOLUTION The fundamental solution E is a **particular solution** of $\mathrm{L}u = \delta$, just as in ODE theory. The general solution has the familiar form: general = homogeneous + particular.

4.92 PROPOSITION (General solution structure). *Let* L *be a constant-coefficient operator.*

1. ***Non-uniqueness of fundamental solutions.*** *If* E *is a fundamental solution of* L *and* $w \in \mathcal{D}'(\mathbb{R}^d)$ *satisfies* $\mathrm{L}w = 0$, *then* $\mathrm{E} + w$ *is also a fundamental solution.*

2. ***All fundamental solutions differ by homogeneous solutions.*** *If* E_1 *and* E_2 *are both fundamental solutions (*$\mathrm{LE}_1 = \mathrm{LE}_2 = \delta$*), then* $\mathrm{L}(\mathrm{E}_1 - \mathrm{E}_2) = 0$, *so* $\mathrm{E}_1 - \mathrm{E}_2$ *is a solution of the homogeneous equation.*

3. ***General solution of*** $\mathrm{L}u = f$***.*** *If* $u_p = \mathrm{E} * f$ *is a particular solution of* $\mathrm{L}u = f$, *then every distributional solution of* $\mathrm{L}u = f$ *has the form* $u = u_p + w$ *where* $\mathrm{L}w = 0$.

Proof. Parts 1 and 2 follow from linearity: $\mathrm{L}(\mathrm{E} + w) = \mathrm{LE} + \mathrm{L}w = \delta + 0 = \delta$, and $\mathrm{L}(\mathrm{E}_1 - \mathrm{E}_2) = \mathrm{LE}_1 - \mathrm{LE}_2 = \delta - \delta = 0$.

For part 3, if $Lu = f$ and $Lu_p = f$, then $L(u - u_p) = 0$, so $w = u - u_p$ solves the homogeneous equation.

□

This is the exact analogue of the ODE theorem: the general solution of a nonhomogeneous linear equation is a particular solution plus the general solution of the homogeneous equation. Distribution theory provides the particular solution via convolution; the homogeneous solutions depend on the specific operator L and on boundary or initial conditions.

4.13.4 Solving PDEs via Convolution

With a fundamental solution in hand, the superposition argument from the beginning of this section becomes rigorous.

4.93 THEOREM (Solution by convolution). *Let* L *be a constant-coefficient operator with fundamental solution* E. *If* $f \in \mathcal{D}(\mathbb{R}^d)$ *(or more generally, if* f *has compact support and the convolution* $E * f$ *is well-defined), then*

$$u = E * f \tag{320}$$

solves $Lu = f$ *in* $\mathcal{D}'(\mathbb{R}^d)$.

Proof. By the properties of convolution and distributional differentiation:

$$Lu = L(E * f) = (LE) * f = \delta * f = f, \tag{321}$$

where we used that L commutes with convolution (since L has constant coefficients), $LE = \delta$, and $\delta * f = f$ (the delta is the identity for convolution).

□

This is the distributional realization of the superposition principle: each point y contributes an impulse response $f(y)\,E(\cdot - y)$, and the integral over all y assembles the solution.

4.94 EXAMPLE (Solving Poisson's equation). Poisson's equation $-\Delta u = f$ in $\mathbb{R}^3$ has the solution

$$u(x) = (E * f)(x) = \frac{1}{4\pi} \int_{\mathbb{R}^3} \frac{f(y)}{|x - y|}\, dy, \tag{322}$$

which is the **Newton potential** of f. This is the gravitational (or electrostatic) potential generated by a mass (or charge) density f.

4.14 *Application: Weak Formulation of PDEs*

Big Idea

Fundamental solutions give explicit formulas but require constant coefficients and no boundary. The weak formulation replaces the PDE with an equation in a Sobolev dual space, reducing existence to the Riesz representation theorem (symmetric case) or Lax-Milgram (general case). The Gelfand triple $H_0^1 \hookrightarrow L^2 \hookrightarrow H^{-1}$ identifies where the solution, the data, and the bilinear form live.

4.14.1 From classical to weak solutions

Fundamental solutions give explicit formulas, but only for constant-coefficient operators on all of $\mathbb{R}^d$. For bounded domains with boundary conditions, or for variable-coefficient operators, we need a different approach. Sobolev spaces provide it: instead of finding a formula, we prove existence abstractly using duality.

The idea is to rewrite the PDE as a problem about functionals on a Hilbert space. Consider $-\Delta u = f$ on a bounded domain Ω with $u = 0$ on $\partial\Omega$. The boundary condition is encoded by working in $H_0^1(\Omega)$ (the closure of $C_c^\infty(\Omega)$ in H^1). Multiply the PDE by a test function $v \in H_0^1(\Omega)$ and integrate by parts:

$$\int_\Omega \nabla u \cdot \nabla v \, dx = \int_\Omega f \, v \, dx \qquad \text{for all } v \in H_0^1(\Omega). \tag{323}$$

The boundary terms vanish because $v \in H_0^1$. This is the **weak formulation**: find $u \in H_0^1(\Omega)$ satisfying the identity above.

Existence by the Riesz representation theorem **The PDE is an equation in the dual space.** The weak formulation is not an equation between functions in Ω; it is an equation between *elements of* $H^{-1}(\Omega) = H_0^1(\Omega)^*$. Read both sides of (weak) as functionals of the test direction v:

$$\underbrace{v \mapsto \int_\Omega \nabla u \cdot \nabla v \, dx}_{=: Au \in H^{-1}(\Omega)} \;=\; \underbrace{v \mapsto \int_\Omega f \, v \, dx}_{=: F \in H^{-1}(\Omega)} \qquad \text{in } H^{-1}(\Omega). \tag{324}$$

The left-hand functional Au is continuous on H_0^1 by Cauchy-Schwarz ($|Au(v)| \le \|\nabla u\|_{L^2}\|\nabla v\|_{L^2}$), so $A : H_0^1 \to H^{-1}$ is a bounded operator. The right-hand functional F is continuous on H_0^1 by Hölder plus Poincaré (for $f \in L^2$, or more generally for $f \in H^{-1}$ by definition). Solving the PDE means solving the linear equation $Au = F$ in H^{-1}.

This is the shift that distributions made possible. Classically $-\Delta u = f$ is an equation between functions at each point $x \in \Omega$; weakly, it is a single equation between two bounded linear functionals on H_0^1. The equality is tested by pairing against every $v \in H_0^1$, not by comparing pointwise values.

The left-hand side is simultaneously an inner product on $H_0^1(\Omega)$, $\langle u, v\rangle_{\dot{H}^1} = \int \nabla u \cdot \nabla v \, dx$ (using the Poincaré inequality to show $\|\nabla \cdot\|_{L^2}$ is a norm equivalent to $\|\cdot\|_{H^1}$ on H_0^1). That is what makes Riesz applicable: the operator A is the Riesz isomorphism $H_0^1 \xrightarrow{\sim} H^{-1}$ associated with this inner product.

By the Riesz representation theorem (Theorem 3.22), there exists a unique $u_f \in \mathrm{H}^1_0(\Omega)$ such that

$$\int_\Omega \nabla u_f \cdot \nabla v \, dx = f(v) \qquad \text{for all } v \in \mathrm{H}^1_0(\Omega). \tag{325}$$

This u_f is the weak solution. Existence and uniqueness are immediate corollaries of a theorem we already proved in the duality chapter.

4.95 REMARK. The Riesz map $(-\Delta)^{-1} : \mathrm{H}^{-1}(\Omega) \to \mathrm{H}^1_0(\Omega)$, $f \mapsto u_f$, is the **inverse Laplacian**. It gains two derivatives: the input f lives in H^{-1} (a distribution), but the output u_f lives in H^1 (a function with a derivative in L^2).

This two-derivative gain is not a coincidence. If $f \in \mathrm{H}^{-1}$, the pairing $\langle f, v\rangle$ is defined abstractly. But the Riesz map produces $u_f \in \mathrm{H}^1_0$, and now

$$\langle f, v\rangle_{\mathrm{H}^{-1}\times\mathrm{H}^1} = \int_\Omega \nabla u_f \cdot \nabla v \, dx. \tag{326}$$

The singular object f has been "absorbed" into $\nabla u_f \in \mathrm{L}^2$. The regularity of u_f compensates exactly for the singularity of f.

THE LAX-MILGRAM THEOREM The Riesz argument above works because the left-hand side $\int \nabla u \cdot \nabla v \, dx$ *is* an inner product on H^1_0. For more general operators (with lower-order terms, non-symmetric coefficients, or convection) the bilinear form $\mathrm{B}(u, v)$ may not be symmetric, so it is not an inner product. The Lax-Milgram theorem handles this case.

4.96 THEOREM (Lax-Milgram). *Let* V *be a real Hilbert space. Suppose* $\mathrm{B} : \mathrm{V} \times \mathrm{V} \to \mathbb{R}$ *is a bilinear form satisfying:*

1. ***Continuity:*** $|\mathrm{B}(u, v)| \le \mathrm{M}\|u\|_{\mathrm{V}}\|v\|_{\mathrm{V}}$ *for all* $u, v \in \mathrm{V}$,
2. ***Coercivity:*** $\mathrm{B}(u, u) \ge \alpha\|u\|^2_{\mathrm{V}}$ *for all* $u \in \mathrm{V}$,

for constants $\mathrm{M}, \alpha > 0$. *Then for every* $f \in \mathrm{V}^*$, *there exists a unique* $u \in \mathrm{V}$ *such that*

$$\mathrm{B}(u, v) = \langle f, v\rangle \qquad \textit{for all } v \in \mathrm{V}, \tag{327}$$

and $\|u\|_{\mathrm{V}} \le \frac{1}{\alpha}\|f\|_{\mathrm{V}^*}$.

Proof. For each fixed $u \in \mathrm{V}$, the map $v \mapsto \mathrm{B}(u, v)$ is a continuous linear functional on V (by the continuity bound). By the Riesz representation theorem, there exists a unique $\mathrm{A}u \in \mathrm{V}$ such that

$$\mathrm{B}(u, v) = \langle \mathrm{A}u, v\rangle_{\mathrm{V}} \qquad \text{for all } v \in \mathrm{V}. \tag{328}$$

This defines a linear operator $\mathrm{A} : \mathrm{V} \to \mathrm{V}$. Continuity of B gives $\|\mathrm{A}u\|_{\mathrm{V}} \le \mathrm{M}\|u\|_{\mathrm{V}}$. Coercivity gives

$$\alpha\|u\|_{\mathrm{V}}^2 \le \mathrm{B}(u,u) = \langle \mathrm{A}u, u\rangle_{\mathrm{V}} \le \|\mathrm{A}u\|_{\mathrm{V}}\|u\|_{\mathrm{V}}, \tag{329}$$

so $\|\mathrm{A}u\|_{\mathrm{V}} \ge \alpha\|u\|_{\mathrm{V}}$. This means A is injective and has closed range. To show surjectivity: if $w \perp \mathrm{range}(\mathrm{A})$, then $0 = \langle \mathrm{A}w, w\rangle = \mathrm{B}(w,w) \ge \alpha\|w\|^2$, so $w = 0$. Hence A is bijective.

Now represent f via Riesz: there exists $\tilde{f} \in \mathrm{V}$ with $\langle f, v\rangle = \langle \tilde{f}, v\rangle_{\mathrm{V}}$. Set $u = \mathrm{A}^{-1}\tilde{f}$. Then $\mathrm{B}(u,v) = \langle \mathrm{A}u, v\rangle_{\mathrm{V}} = \langle \tilde{f}, v\rangle_{\mathrm{V}} = \langle f, v\rangle$ for all v. The bound $\|u\|_{\mathrm{V}} \le \frac{1}{\alpha}\|f\|_{\mathrm{V}^*}$ follows from coercivity.

□

4.97 REMARK (Riesz as a special case). When $\mathrm{B}(u,v) = \langle u, v\rangle_{\mathrm{V}}$ is the inner product itself, Lax-Milgram reduces to the Riesz representation theorem with $\mathrm{M} = \alpha = 1$.

4.98 REMARK (Lax-Milgram and the direct method). The two hypotheses of Lax-Milgram play exactly the same roles as the two ingredients of the direct method for variational problems. Recall the direct method (Example 3.72): to minimize $\mathrm{E} : \mathrm{V} \to \mathbb{R}$, take a minimizing sequence u_n, use *coercivity* to force (u_n) to be bounded, use *Banach-Alaoglu* (Theorem 3.67) to extract a weakly convergent subsequence $u_{n_k} \rightharpoonup u$, and use *weak lower semicontinuity* to conclude $\mathrm{E}(u) \le \liminf \mathrm{E}(u_{n_k}) = \inf \mathrm{E}$. Two ingredients: coercivity to *get* a weak limit, weak l.s.c. to *pass to it*.

When B is symmetric, Lax-Milgram *is* the direct method applied to the energy

$$\mathrm{E}(v) = \tfrac{1}{2}\mathrm{B}(v,v) - \langle f, v\rangle, \tag{330}$$

whose Euler-Lagrange equation is $\mathrm{B}(u,v) = \langle f, v\rangle$. The two Lax-Milgram hypotheses translate into the two direct-method ingredients as follows.

Coercivity of B $\Longrightarrow$ **coercivity of** E**.** From $\mathrm{B}(v,v) \ge \alpha\|v\|^2$,

$$\mathrm{E}(v) \;\ge\; \tfrac{\alpha}{2}\|v\|^2 - \|f\|_{\mathrm{V}^*}\|v\| \;\longrightarrow\; \infty \quad \text{as } \|v\| \to \infty. \tag{331}$$

So minimizing sequences are bounded in V. This is the same mechanism that produces the a priori estimate $\|u\|_{\mathrm{V}} \le \frac{1}{\alpha}\|f\|_{\mathrm{V}^*}$ in the Lax-Milgram conclusion: coercivity controls the size of the solution.

Continuity of B **+ symmetry + coercivity** $\Longrightarrow$ **weak l.s.c. of** E**.** This step is subtler than it looks. Continuity of B makes $v \mapsto \frac{1}{2}\mathrm{B}(v,v)$ continuous in the *norm* topology (since $|\mathrm{B}(v,v)| \le \mathrm{M}\|v\|^2$), but norm continuity alone does **not** imply weak l.s.c. — $v \mapsto -\|v\|^2$ is norm-continuous yet fails to be weakly l.s.c. The bridge is **convexity**, supplied by Theorem 3.57:

> A convex, norm-continuous functional on a Banach space is weakly lower semicontinuous.

Symmetry together with coercivity makes $v \mapsto \frac{1}{2}\mathrm{B}(v,v)$ a positive-definite quadratic form, hence convex. Continuity then promotes convexity to weak l.s.c. via Mazur. So the honest reading of the Lax-Milgram hypotheses under symmetry is

$$\underbrace{\text{continuity}}_{\text{norm continuity of E}} \;+\; \underbrace{\text{coercivity} + \text{symmetry}}_{\text{convexity of E}} \xrightarrow{\text{Mazur}} \text{weak l.s.c. of E.} \tag{332}$$

The shorthand "continuity of B $\leftrightarrow$ weak l.s.c." is harmless in the symmetric setting because convexity is already on the table for free.

Non-symmetric case. When B is not symmetric there is no energy to minimize, and the direct method does not literally apply. The proof of Lax-Milgram instead routes through Riesz representation and the operator $A : V \to V$ with $B(u, v) = \langle Au, v\rangle_V$. But the *roles* of the two hypotheses are preserved:

- Coercivity gives the lower bound $\|Au\|_V \geq \alpha\|u\|_V$, which forces A to be injective with closed range — the "compactness / a priori bound" step.
- Continuity makes A a bounded operator in the first place — the "pass to the limit" step.

So the dichotomy *coercivity controls size*, *continuity controls limits* survives even when there is no energy functional in sight.

4.14.2 The Gelfand triple

The natural framework for weak formulations is a triple of spaces

$$V \hookrightarrow H \hookrightarrow V^* \tag{333}$$

where V is a Hilbert space (the "energy space"), H is a pivot space, and V^* is the dual. The embeddings are continuous and dense.

The prototypical example is

$$H_0^1(\Omega) \hookrightarrow L^2(\Omega) \hookrightarrow H^{-1}(\Omega). \tag{334}$$

An element $f \in H^{-1}$ acts on test functions in H_0^1 via the duality pairing. The PDE $Au = f$ is an equation *in* H^{-1}: we seek $u \in H_0^1$ such that $B(u, v) = \langle f, v\rangle_{H^{-1}\times H_0^1}$ for all $v \in H_0^1$. The Gelfand triple identifies which space each player lives in:

- The **solution** u lives in the energy space $V = H_0^1$.
- The **data** f lives in the dual $V^* = H^{-1}$.
- The **bilinear form** $B : V \times V \to \mathbb{R}$ connects them.
- The **pivot** $H = L^2$ sits in between, providing the L^2 inner product that mediates between function values and functionals.

4.99 EXAMPLE (General second-order elliptic operator). Consider the operator $Lu = -\operatorname{div}(A(x)\nabla u) + b(x)\cdot\nabla u + c(x)u$ on a bounded domain Ω with Dirichlet boundary conditions. The weak formulation is: find $u \in H_0^1(\Omega)$ such that

$$B(u, v) = \int_\Omega \Big[A(x)\nabla u \cdot \nabla v + (b(x)\cdot\nabla u)\, v + c(x)\, u\, v\Big]\, dx = \langle f, v\rangle \tag{335}$$

for all $v \in H_0^1(\Omega)$.

- **Continuity** of B follows from $A \in L^\infty$, $b \in L^\infty$, $c \in L^\infty$, and Cauchy-Schwarz.
- **Coercivity** requires uniform ellipticity ($A(x)\xi\cdot\xi \geq \lambda|\xi|^2$) and suitable bounds on b and c (or add a large enough multiple of u to absorb the lower-order terms via Poincaré).

By Lax-Milgram, there exists a unique weak solution $u \in H_0^1(\Omega)$.

WHAT SOBOLEV SPACES PROVIDE Compared to the fundamental solution approach, the weak formulation gains:

1. **Boundary conditions.** The space $H_0^1(\Omega)$ encodes $u = 0$ on $\partial\Omega$. No need to modify the fundamental solution.
2. **Variable coefficients.** Replace $\int \nabla u \cdot \nabla v$ by $\int A(x)\nabla u \cdot \nabla v$ for a matrix-valued coefficient $A(x)$. As long as A is uniformly elliptic, the same argument works (via Lax-Milgram instead of Riesz).
3. **Regularity from the Sobolev scale.** The solution u_f lives in H_0^1. If f is more regular ($f \in L^2$ rather than H^{-1}), elliptic regularity gives $u_f \in H^2$. The Sobolev scale tracks exactly how regular the solution is.
4. **Compactness.** The embedding $H_0^1(\Omega) \hookrightarrow L^2(\Omega)$ is compact (Rellich-Kondrachov), so the inverse Laplacian $(-\Delta)^{-1} : L^2 \to H_0^1 \hookrightarrow L^2$ is a compact operator. This gives eigenvalues, eigenfunctions, and the spectral theory from the linear operators chapter.

TWO APPROACHES, ONE SOLUTION On a bounded domain Ω, both approaches produce the same answer. The **Green's function** $G(x, y)$ is the fundamental solution modified to satisfy boundary conditions: $-\Delta_x G(x, y) = \delta(x - y)$ with $G = 0$ on $\partial\Omega$. The convolution formula becomes

$$u(x) = \int_\Omega G(x, y)\, f(y)\, dy, \tag{336}$$

and this u is also the Riesz representative in $H_0^1(\Omega)$. The fundamental solution gives the explicit kernel; the Sobolev approach gives existence without needing to find it.

4.14.3 Nonlinear PDE via Compactness

The Lax-Milgram theorem handles linear equations. For nonlinear problems, we need a different strategy: approximate, extract limits, and pass to the limit in the nonlinear terms. The tools are Banach-Alaoglu (weak compactness), Rellich-Kondrachov (compact embedding), and Sobolev embedding (controlling the nonlinearity). We illustrate with a model problem that ties together the entire theory.

THE MODEL PROBLEM Consider the semilinear elliptic equation

$$-\Delta u + u^3 = f \qquad \text{in } \Omega, \qquad u = 0 \text{ on } \partial\Omega, \tag{337}$$

where $\Omega \subset \mathbb{R}^d$ ($d \leq 3$) is bounded with Lipschitz boundary and $f \in H^{-1}(\Omega)$. The cubic nonlinearity u^3 models a restoring force that grows with amplitude.

The **weak formulation** is: find $u \in H_0^1(\Omega)$ such that

$$\int_\Omega \nabla u \cdot \nabla v \, dx + \int_\Omega u^3 v \, dx = \langle f, v \rangle \qquad \text{for all } v \in H_0^1(\Omega). \tag{338}$$

The nonlinear term $\int u^3 v\,dx$ is well-defined when $d \le 3$ because the Sobolev embedding $H_0^1(\Omega) \hookrightarrow L^6(\Omega)$ (for $d = 3$) gives $u^3 \in L^2$ and $u^3 v \in L^1$.

GALERKIN APPROXIMATION Choose finite-dimensional subspaces $V_n \subset H_0^1(\Omega)$ with $\bigcup_n V_n$ dense in $H_0^1(\Omega)$ (e.g., the span of the first n eigenfunctions of $-\Delta$). The **Galerkin problem** is: find $u_n \in V_n$ such that

$$\int_\Omega \nabla u_n \cdot \nabla v\,dx + \int_\Omega u_n^3 v\,dx = \langle f, v\rangle \qquad \text{for all } v \in V_n. \tag{339}$$

This is a finite-dimensional system, so Brouwer's fixed point theorem guarantees the existence of a Galerkin solution u_n.

A PRIORI BOUNDS Test with $v = u_n$ itself:

$$\|\nabla u_n\|_{L^2}^2 + \|u_n\|_{L^4}^4 = \langle f, u_n\rangle \le \|f\|_{H^{-1}}\|u_n\|_{H^1}. \tag{340}$$

By the Poincaré inequality (Theorem 4.66), $\|u_n\|_{H^1} \le C\|\nabla u_n\|_{L^2}$, so

$$\|\nabla u_n\|_{L^2}^2 \le \|f\|_{H^{-1}} \cdot C\|\nabla u_n\|_{L^2}. \tag{341}$$

Dividing: $\|\nabla u_n\|_{L^2} \le C\|f\|_{H^{-1}}$. The sequence (u_n) is **bounded in** $H_0^1(\Omega)$, uniformly in n.

EXTRACTING THE LIMIT The bounded sequence (u_n) in the reflexive space $H_0^1(\Omega)$ has a weakly convergent subsequence by the Banach-Alaoglu theorem (Theorem 3.67):

$$u_n \rightharpoonup u \quad \text{in } H_0^1(\Omega). \tag{342}$$

Weak convergence alone is not enough to pass to the limit in the nonlinear term u_n^3. The map $u \mapsto u^3$ is continuous (in the strong topology), but *continuous is not the same as weakly continuous*: continuity preserves *strong* limits, and weak convergence is in general only preserved by **linear** bounded operators. Nonlinear maps can destroy weak limits through oscillation. The textbook example is squaring: $u_n(x) = \sin(nx)$ on $(0, 2\pi)$ satisfies $u_n \rightharpoonup 0$ in L^2, but

$$u_n^2 = \tfrac{1}{2} - \tfrac{1}{2}\cos(2nx) \;\rightharpoonup\; \tfrac{1}{2} \;\neq\; 0 = u^2, \tag{343}$$

so $u \mapsto u^2$ is *not* weakly continuous. Cubing fails the same way: $u_n = 1 + \sin(nx)$ has $u_n \rightharpoonup 1$ but $u_n^3 \rightharpoonup \frac{5}{2} \neq 1 = u^3$ (the cross term $3\sin^2(nx) \rightharpoonup 3/2$ is what leaks out). So knowing only that $u_n \rightharpoonup u$ in H_0^1 is not enough to conclude $u_n^3 \to u^3$ in any sense; we need strong convergence.

The fix is to upgrade weak convergence to strong convergence in a space where the nonlinearity *is* continuous, and this is exactly what **Rellich-Kondrachov** provides: applying Proposition 3.54 with T the compact embedding $\iota : H_0^1 \hookrightarrow L^p$ (compact) for $p < 6$ (taking $d = 3$, so the Sobolev critical exponent is $2^* = 6$) upgrades the weak-H_0^1 convergence of the Galerkin sequence, along a further subsequence, to:

$$u_n \to u \quad \text{strongly in } \mathrm{L}^p(\Omega) \text{ for } p < 6. \tag{344}$$

In particular, $u_n \to u$ in $\mathrm{L}^4(\Omega)$. This strong convergence gives:

$$u_n^3 \to u^3 \quad \text{in } \mathrm{L}^{4/3}(\Omega) \tag{345}$$

(since $\|u_n^3 - u^3\|_{\mathrm{L}^{4/3}} \le \mathrm{C}\|u_n - u\|_{\mathrm{L}^4}(\|u_n\|_{\mathrm{L}^4}^2 + \|u\|_{\mathrm{L}^4}^2)$ by the algebraic identity $a^3 - b^3 = (a-b)(a^2+ab+b^2)$).

Passing to the limit Fix a test direction $v \in \mathrm{H}_0^1(\Omega)$ (not just $v \in \mathrm{V}_m$) and read each side of the Galerkin equation as a value of a functional on H_0^1. For every n, the maps

$$\Phi_n \;:\; v \;\mapsto\; \int_\Omega \nabla u_n \cdot \nabla v \, dx + \int_\Omega u_n^3 \, v \, dx, \qquad \mathrm{F} \;:\; v \;\mapsto\; \langle f, v \rangle \tag{346}$$

are continuous linear functionals on H_0^1, i.e. elements of $\mathrm{H}^{-1}(\Omega)$. The linear piece is continuous by Cauchy-Schwarz; the nonlinear piece is continuous because $u_n \in \mathrm{H}_0^1 \hookrightarrow \mathrm{L}^4$ implies $u_n^3 \in \mathrm{L}^{4/3} \hookrightarrow \mathrm{H}^{-1}$ (Sobolev embedding dualized). So the Galerkin identity is the statement $\Phi_n = \mathrm{F}$ in H^{-1}, but *restricted* to test directions $v \in \mathrm{V}_m$, i.e. tested only against a dense subset.

We want to upgrade this to $\Phi = \mathrm{F}$ in H^{-1} for a single limit object Φ. Both pieces of Φ_n converge, pointwise in v, to the corresponding pieces built from u:

- **Linear term.** $u_n \rightharpoonup u$ in H_0^1 is exactly the statement $\int \nabla u_n \cdot \nabla v \, dx \to \int \nabla u \cdot \nabla v \, dx$ for every $v \in \mathrm{H}_0^1$ (weak convergence *is* pointwise convergence of the Riesz functionals).
- **Nonlinear term.** Rellich upgraded weak-H_0^1 to strong-L^4 convergence, hence $u_n^3 \to u^3$ in $\mathrm{L}^{4/3}$. This upgrades to convergence in H^{-1} via the continuous embedding $\mathrm{L}^{4/3} \hookrightarrow \mathrm{H}^{-1}$: for any $g \in \mathrm{L}^{4/3}$, Hölder and Sobolev give

$$\left| \int g \, v \, dx \right| \le \|g\|_{\mathrm{L}^{4/3}} \, \|v\|_{\mathrm{L}^4} \le \mathrm{C} \, \|g\|_{\mathrm{L}^{4/3}} \, \|v\|_{\mathrm{H}_0^1}, \tag{347}$$

  so g defines a bounded linear functional on H_0^1 with $\|g\|_{\mathrm{H}^{-1}} \le \mathrm{C}\|g\|_{\mathrm{L}^{4/3}}$. This is exactly the dual of Sobolev $\mathrm{H}_0^1 \hookrightarrow \mathrm{L}^4$. Pairing with any fixed $v \in \mathrm{H}_0^1 \subset \mathrm{L}^4$ then gives $\int u_n^3 v \, dx \to \int u^3 v \, dx$.

So $\Phi_n(v) \to \Phi(v)$ for every $v \in \mathrm{H}_0^1$, where $\Phi(v) = \int \nabla u \cdot \nabla v + \int u^3 v$. For each $v \in \mathrm{V}_m$ (any m) the identity $\Phi_n(v) = \mathrm{F}(v)$ holds eventually, so passing $n \to \infty$ yields $\Phi(v) = \mathrm{F}(v)$ on the dense set $\bigcup_m \mathrm{V}_m$. Continuity of Φ and F on H_0^1 extends this to all $v \in \mathrm{H}_0^1$. The Galerkin equation has upgraded to the identity $\Phi = \mathrm{F}$ in H^{-1}, which is the weak form $-\Delta u + u^3 = f$ we set out to solve.

Linear vs. nonlinear: why compactness was needed Compare with the purely linear problem $-\Delta u = f$ from the start of the section. There, existence followed from Riesz representation alone: weak convergence in H_0^1 was *already* enough to pass to the limit, because the only operation applied to u_n was the continuous linear functional $v \mapsto \int \nabla u_n \cdot \nabla v$. Linear bounded operators preserve weak convergence by definition, so weak subsequential limits suffice and Rellich is never needed.

Nonlinearity changes this. The map $u \mapsto u^3$ is continuous but *not weakly continuous*, so weak convergence of u_n gives no information about u_n^3. The fix is to upgrade weak convergence to strong convergence in a space where the nonlinearity is continuous, and this upgrade is precisely what a compact embedding provides. This is the general shape of the direct method: for linear problems weak compactness (Banach-Alaoglu, Riesz) is enough; for nonlinear problems weak compactness must be combined with a compact embedding (Rellich-Kondrachov) to tame the nonlinear term.

Why multiplication is the hard operation. The obstruction is structural, and the Banach-algebra chapter already told us what it is: multiplication is a *frequency-convolution* operation ($\widehat{uv} = \hat{u} * \hat{v}$), so squaring or cubing a function *sums its frequencies*. Two modes at frequencies n and m produce a mode at frequency $n + m$. High-frequency content in u_n therefore breeds even higher frequencies in u_n^3, so u_n^3 can be strictly *worse behaved* than u_n, and may leave the Sobolev space that u_n lived in. That is why Sobolev spaces fail to be Banach algebras below the threshold $s > d/p$ (Theorem 4.71): squaring moves you out of the space, because the squared spectrum escapes the regularity bound. The compactness trick sidesteps the problem by reducing the question to strong convergence in a concrete L^q, where multiplication *is* continuous (Hölder), rather than working in the H^s where multiplication might blow up the regularity budget.

REFERENCES

[1] Cybenko, George: *Approximation by superpositions of a sigmoidal function*. Mathematics of Control, Signals and Systems, 2(4):303–314, 1989.

[2] Day, Mahlon M.: *The Spaces L^p with $0 < p < 1$*. Bulletin of the American Mathematical Society, 46(10):816–823, 1940.

[3] Hillen, Thomas: *Elements of Applied Functional Analysis*. June 2023. https://www.math.ualberta.ca/~thillen/FA-book-June2023.pdf.

[4] Tao, Terence: *245b, Notes 9: The Baire category theorem and its Banach space consequences*. https://terrytao.wordpress.com/2009/02/01/245b-notes-9-the-baire-category-theorem-and-its-banach-space-consequences/, feb 1 2009. https://terrytao.wordpress.com/2009/02/01/245b-notes-9-the-baire-category-theorem-and-its-banach-space-consequences/, [Online; accessed 2025-03-03].