

Real Analysis

Course Notes
C. McMullen

Contents

1	Introduction	1
2	Set Theory and the Real Numbers	4
3	Lebesgue Measurable Sets	13
4	Measurable Functions	26
5	Integration	35
6	Differentiation and Integration	44
7	The Classical Banach Spaces	60
8	Baire Category	72
9	General Topology	81
10	Banach Spaces	97
11	Fourier Series	112
12	Harmonic Analysis on \mathbb{R} and S^2	126
13	General Measure Theory	131
A	Measurable A with $A - A$ nonmeasurable	136

1 Introduction

We begin by discussing the motivation for real analysis, and especially for the reconsideration of the notion of integral and the invention of Lebesgue integration, which goes beyond the Riemannian integral familiar from classical calculus.

1. *Usefulness of analysis.* As one of the oldest branches of mathematics, and one that includes calculus, analysis is hardly in need of justification. But just in case, we remark that its uses include:

1. The description of physical systems, such as planetary motion, by dynamical systems (ordinary differential equations);
2. The theory of partial differential equations, such as those describing heat flow or quantum particles;
3. Harmonic analysis on Lie groups, of which \mathbb{R} is a simple example;
4. Representation theory;

5. The description of optimal structures, from minimal surfaces to economic equilibria;
6. The foundations of probability theory;
7. Automorphic forms and analytic number theory; and
8. Dynamics and ergodic theory.

2. *Completeness.* We now motivate the need for a sophisticated theory of measure and integration, called the Lebesgue theory, which will form the first topic in this course.

In analysis it is necessary to take limits; thus one is naturally led to the construction of the real numbers, a system of numbers containing the rationals and closed under limits. When one considers functions it is again natural to work with spaces that are closed under suitable limits. For example, consider the space of continuous functions $C[0, 1]$. We might measure the size of a function here by

$$\|f\|_1 = \int_0^1 |f(x)| dx.$$

(There is no problem defining the integral, say using Riemann sums).

But we quickly see that there are Cauchy sequences of continuous functions whose limit, in this norm, are discontinuous. So we should extend $C[0, 1]$ to a space that is closed under limits. It is not at first even evident that the limiting objects should be *functions*. And if we try to include *all* functions, we are faced with the difficult problem of integrating a general function.

The modern solution to this natural issue is to introduce the idea of *measurable functions*, i.e. a space of functions that is closed under limits and tame enough to integrate. The Riemann integral turns out to be inadequate for these purposes, so a new notion of integration must be invented. In fact we must first examine carefully the idea of the mass or *measure* of a subset $A \subset \mathbb{R}$, which can be thought of as the integral of its indicator function $\chi_A(x) = 1$ if $x \in A$ and $= 0$ if $x \notin A$.

3. *Fourier series.* More classical motivation for the Lebesgue integral come from Fourier series.

Suppose $f : [0, \pi] \rightarrow \mathbb{R}$ is a reasonable function. We define the Fourier coefficients of f by

$$a_n = \frac{2}{\pi} \int_0^\pi f(x) \sin(nx) dx.$$

Here the factor of $2/\pi$ is chosen so that

$$\frac{2}{\pi} \int_0^\pi \sin(nx) \sin(mx) dx = \delta_{nm}.$$

We observe that if

$$f(x) = \sum_1^\infty b_n \sin(nx),$$

then at least formally $a_n = b_n$ (this is true, for example, for a finite sum).

This representation of $f(x)$ as a superposition of sines is very useful for applications. For example, $f(x)$ can be thought of as a sound wave, where a_n measures the strength of the frequency n .

Now what coefficients a_n can occur? The orthogonality relation implies that

$$\frac{2}{\pi} \int_0^\pi |f(x)|^2 dx = \sum_{-\infty}^\infty |a_n|^2.$$

This makes it natural to ask if, conversely, for any a_n such that $\sum |a_n|^2 < \infty$, there exists a function f with these Fourier coefficients. The natural function to try is $f(x) = \sum a_n \sin(nx)$.

But why should this sum even exist? The functions $\sin(nx)$ are only bounded by one, and $\sum |a_n|^2 < \infty$ is much weaker than $\sum |a_n| < \infty$.

One of the original motivations for the theory of Lebesgue measure and integration was to refine the notion of function so that this sum really does exist. The resulting function $f(x)$ however need to be Riemann integrable! To get a reasonable theory that includes such Fourier series, Cantor, Dedekind, Fourier, Lebesgue, etc. were led inexorably to a re-examination of the foundations of real analysis and of mathematics itself. The theory that emerged will be the subject of this course.

Here are a few additional points about this example.

First, we could try to define the required space of functions — called $L^2[0, \pi]$ — to simply be the metric completion of, say $C[0, \pi]$ with respect to $d(f, g) = \int |f - g|^2$. The reals are defined from the rationals in a similar fashion. But the question would still remain, can the limiting objects be thought of as functions?

Second, the set of point $E \subset \mathbb{R}$ where $\sum a_n \sin(nx)$ actually converges is liable to be a very complicated set — not closed or open, or even a countable union or intersection of sets of this form. Thus to even begin, we must have a good understanding of subsets of \mathbb{R} .

Finally, even if the limiting function $f(x)$ exists, it will generally not be Riemann integrable. Thus we must broaden our theory of integration to

deal with such functions. It turns out this is related to the second point — we must again find a good notion for the length or *measure* $m(E)$ of a fairly general subset $E \subset \mathbb{R}$, since $m(E) = \int \chi_E$.

2 Set Theory and the Real Numbers

The foundations of real analysis are given by set theory, and the notion of cardinality in set theory, as well as the axiom of choice, occur frequently in analysis. Thus we begin with a rapid review of this theory. For more details see, e.g. [Hal]. We then discuss the real numbers from both the axiomatic and constructive point of view. Finally we discuss open sets and Borel sets.

In some sense, real analysis is a pearl formed around the grain of sand provided by paradoxical sets. These paradoxical sets include sets that have no reasonable measure, which we will construct using the axiom of choice.

The axioms of set theory. Here is a brief account of the axioms.

- Axiom I. (Extension) A set is determined by its elements. That is, if $x \in A \implies x \in B$ and vice-versa, then $A = B$.
- Axiom II. (Specification) If A is a set then $\{x \in A : P(x)\}$ is also a set.
- Axiom III. (Pairs) If A and B are sets then so is $\{A, B\}$. From this axiom and $\emptyset = 0$, we can now form $\{0, 0\} = \{0\}$, which we call 1; and we can form $\{0, 1\}$, which we call 2; but we cannot yet form $\{0, 1, 2\}$.
- Axiom IV. (Unions) If A is a set, then $\bigcup A = \{x : \exists B, B \in A \ \& \ x \in B\}$ is also a set. From this axiom and that of pairs we can form $\bigcup\{A, B\} = A \cup B$. Thus we can define $x^+ = x + 1 = x \cup \{x\}$, and form, for example, $7 = \{0, 1, 2, 3, 4, 5, 6\}$.
- Axiom V. (Powers) If A is a set, then $\mathcal{P}(A) = \{B : B \subset A\}$ is also a set.
- Axiom VI. (Infinity) There exists a set A such that $0 \in A$ and $x+1 \in A$ whenever $x \in A$. The smallest such set is unique, and we call it $\mathbb{N} = \{0, 1, 2, 3, \dots\}$.
- Axiom VII (The Axiom of Choice): For any set A there is a function $c : \mathcal{P}(A) - \{\emptyset\} \rightarrow A$, such that $c(B) \in B$ for all $B \subset A$.

Cardinality. In set theory, the natural numbers \mathbb{N} are defined inductively by $0 = \emptyset$ and $n = \{0, 1, \dots, n - 1\}$. Thus n , as a set, consists of exactly n elements.

We write $|A| = |B|$ to mean there is a bijection between the sets A and B ; in other words, these sets have the same *cardinality*. A set A is *finite* if $|A| = n$ for some $n \in \mathbb{N}$; it is *countable* if A is finite or $|A| = |\mathbb{N}|$; otherwise, it is *uncountable*.

A countable set is simply one whose elements can be written down in a (possibly finite) list, (x_1, x_2, \dots) . When $|A| = |\mathbb{N}|$ we say A is *countably infinite*.

Inequalities. It is natural to write $|A| \leq |B|$ if there is an injective map $A \hookrightarrow B$. By the Schröder–Bernstein theorem (elementary but nontrivial), we have

$$|A| \leq |B| \quad \text{and} \quad |B| \leq |A| \implies |A| = |B|.$$

The power set. We let A^B denote the set of all maps $f : B \rightarrow A$. The *power set* $\mathcal{P}(A) \cong 2^A$ is the set of all subsets of A . A profound observation, due to Cantor, is that

$$|A| < |\mathcal{P}(A)|$$

for any set A . The proof is easy: if $f : A \rightarrow \mathcal{P}(A)$ were a bijection, we could then form the set

$$B = \{x \in A : x \notin f(x)\},$$

but then B cannot be in the image of f , for if $B = f(x)$, then $x \in B$ iff $x \notin B$.

Russel’s paradox. We remark that Cantor’s argument is closely related to Russell’s paradox: if $E = \{X : X \notin X\}$, then is $E \in E$? Note that the axioms of set theory do not allow us to form the set E !

Countable sets. It is not hard to show that $\mathbb{N} \times \mathbb{N}$ is countable, and consequently:

A countable union of countable sets is countable.

Thus \mathbb{Z} , \mathbb{Q} and the set of algebraic numbers in \mathbb{C} are all countable sets.

Remark: The Axiom of Choice. Recall this axiom states that for any set A , there is a map $c : \mathcal{P}(A) - \{\emptyset\} \rightarrow A$ such that $c(A) \in A$. This axiom is often useful and indeed necessary in proving very general theorems; for example, if there is a surjective map $f : A \rightarrow B$, then there is an injective map $g : B \rightarrow A$ (and thus $|B| \leq |A|$). (Proof: set $g(b) = c(f^{-1}(b))$.)

Another typical application of the axiom of choice is to show:

Every vector space has a basis.

To see this is nontrivial, consider the real numbers as a vector space over \mathbb{Q} ; can you find a basis?

The real numbers. In real analysis we need to deal with possibly wild functions on \mathbb{R} and fairly general subsets of \mathbb{R} , and as a result a firm grounding in basic set theory is helpful. We begin with the *definition* of the real numbers. There are at least 4 different reasonable approaches.

The axiomatic approach. As advocated by Hilbert, the real numbers can be approached axiomatically, like groups or plane geometry. Accordingly, the real numbers are *defined* as a *complete, ordered field*. Note that in a field, $0 \neq 1$ by definition.

A field K is *ordered* if it is equipped with a distinguished subset K_+ that is closed under addition and multiplication, such that

$$K = K_+ \sqcup \{0\} \sqcup (-K_+).$$

It is *complete* if every nonempty set $A \subset K$ that is bounded above has a *least upper bound*, which is denoted $\sup A \in K$.

Least upper bounds, limits and events. If we extend the real line by adding in $\pm\infty$, then *any* subset of \mathbb{R} has a natural supremum. For example, $\sup \mathbb{Z} = +\infty$ and $\sup \emptyset = -\infty$. The great lower bound for A is denoted by $\inf A$.

From these notions we can extract the usual notion of limit in calculus, together with some useful variants. We first note that monotone sequences always have limits, e.g.:

If x_n is an increasing sequence of real numbers, then $x_n \rightarrow \sup(x_n)$.

We then define the important notion of lim-sup by:

$$\limsup x_n = \lim_{N \rightarrow \infty} \sup_{n > N} x_n.$$

This is the limit of a *decreasing sequence*, so it always exists. The liminf is defined similarly, and finally we say x_n *converges* if

$$\limsup x_n = \liminf x_n,$$

in which case their common value is the usual limit, $\lim x_n$.

For example, $(x_n) = (2/1, -3/2, +4/3, -5/4, \dots)$ has $\limsup x_n = 1$ even though $\sup x_n = 2$.

The limsup and liminf of a sequence of 0's and 1's is again either 0 or 1. Thus given a sequence of sets $E_i \subset \mathbb{R}$, there is a unique sets $\limsup E_i$ such that

$$\chi_{\limsup E_i} = \limsup \chi_{E_i},$$

and similarly for $\liminf E_i$. In fact

$$\limsup E_i = \{x : x \in E_i \text{ for infinitely many } i\},$$

while

$$\limsup E_i = \{x : x \in E_i \text{ for all } i \text{ from some point on}\}.$$

These notions are particularly natural in probability theory, where we think of the sets E_i as *events*.

Consequences of the axioms. Here are some first consequences of the axioms.

1. The real numbers have characteristic zero. Indeed, $1 + 1 + \cdots + 1 = n > 0$ for all n , since \mathbb{R}_+ is closed under addition.
2. Given a real number x , there exists an integer n such that $n > x$. Proof: otherwise, we would have $\mathbb{Z} < x$ for some x . By completeness, this means we have a real number $x_0 = \sup \mathbb{Z}$. Then $x_0 - 1$ is *not* an upper bound for \mathbb{Z} , so $x_0 - 1 < n$ for some $n \in \mathbb{Z}$. But then $n + 1 > x_0$, a contradiction.
3. Corollary: If $\epsilon > 0$ then $\epsilon > 1/n > 0$ for some integer n .
4. Any interval (a, b) contains a rational number p/q . (In other words, \mathbb{Q} is dense in \mathbb{Q} .)

Constructions of \mathbb{R} . To show the real numbers exist, one must construct from first principles (i.e. from the axioms of a set theory) a field with the required properties. Here are 3 such constructions.

Dedekind cuts. One can visualize a real number x as a *cut* that partitions the rational numbers into 2 sets,

$$A = \{r \in \mathbb{Q} : r \leq x\} \quad \text{and} \quad B = \{r \in \mathbb{Q} : r > x\}.$$

Thus one can *define* \mathbb{R} to consists of the set of pairs (A, B) forming partitions of \mathbb{Q} into nonempty sets with $A < B$, *such that* B has no least element. The latter convention makes the cut produced by a rational number unique.

Dedekind cuts work well for addition: we define $(A, B) + (A', B') = (A + A', B + B')$. Multiplication is somewhat trickier, but completeness works fairly well. As a first approximation, one can define

$$\sup(A_\alpha, B_\alpha) = \left(\bigcup A_\alpha, \bigcap B_\alpha \right).$$

The problem here is that when the supremum is rational, the set $\bigcap B_\alpha$ might have a least element. (This suggests it might be better to introduce an equivalence relation on cuts, so that the ‘two versions’ of each rational number are identified.)

The *extended reals* $\mathbb{R} \cup \pm\infty$ are also nicely constructed using Dedekind cuts, by allowing A or B to be empty. We will often implicitly use the extended reals, e.g. by allowing the value of a sum of positive numbers to be infinite rather than simply undefined.

For more on the efficient construction of \mathbb{R} using Dedekind cuts, see [Con, p.25].

Remark: Ideals. Dedekind also proposed the notion of an *ideal* I in the ring of integers A in a number field K . The elements $n \in A$ give *principal ideals* $(n) \subset A$ consisting of all the elements that are divisible by n . Ideals which are not principal can be thought of as ‘ideal’ integers, which do not belong to A but which can be seen implicitly through the set of elements of A that they divide. In the same way a real number can be seen implicitly through the way it cuts \mathbb{Q} into two pieces.

Cauchy sequences. A more analytical approach to the real numbers is to define \mathbb{R} as the metric completion of \mathbb{Q} . Then a real number is represented by a Cauchy sequence $x_k \in \mathbb{Q}$. This means for all $n > 0$ there exists an $N > 0$ such that

$$|x_i - x_j| < 1/n \quad \forall i, j > N.$$

We consider two Cauchy sequences to be equivalent if $|x_i - y_i| \rightarrow 0$ as $i \rightarrow \infty$.

This definition works well with respect to the field operations, e.g. $(x_i) \cdot (y_i) = (x_i y_i)$. It is slightly awkward to prove completeness, since we have defined completeness in terms of upper bounds.

Decimals. A final, perfectly serviceable way to define the real numbers is in terms of decimals, such as $\pi = 3.14159265\dots$. As in the case of Dedekind cuts, one must introduce a convention for numbers of the form $p/10^n$, to deal with the fact that $0.9999\dots = 1.0$.

Other completions of \mathbb{Q} : One can also take the metric completion of \mathbb{Q} in other metrics, such as the d -adic norms where $|p/d^n| = d^n$ (assuming d does not divide p). These yield the rings \mathbb{Q}_d for each integer $d > 1$. All of these completions of \mathbb{Q} are *totally disconnected*.

The elements of \mathbb{Q}_{10} can be thought of as decimal numbers which are *finite* after the decimal point but *not* before it. This ring is not a field! If 5^n accumulates on x and 2^n accumulates on y , then $|x|_{10} = |y|_{10} = 1$ but $xy = 0$. One can make the solution canonical by asking that $x = (0, 1)$ and $y = (1, 0)$ in $\mathbb{Z}_{10} \cong \mathbb{Z}_2 \times \mathbb{Z}_5$; then $y = x + 1 = \dots 4106619977392256259918212890625$.)

On the other hand, \mathbb{Q}_p is a field for all primes p .

The size of the real numbers. It is easy to prove:

The real numbers \mathbb{R} are uncountable.

For example, if we had a list of all the real numbers x_1, x_2, \dots , we could then construct a new real number z whose i th decimal digit differs from the i th decimal digit of x_i , so that z is missing from the list.

A more precise statement is that $|\mathbb{R}| = |\mathcal{P}(\mathbb{N})|$. To see this, one can e.g. use decimals to show that $2^{\mathbb{N}} \hookrightarrow [0, 1]$, and use binary numbers to show that $2^{\mathbb{N}}$ maps onto $[0, 1]$, and finally show (by any number of arguments) that $|[0, 1]| = |\mathbb{R}|$.

The continuum. The real numbers have a natural topology, coming from the metric $d(x, y) = |x - y|$, with respect to which they are connected. In fact, classically the real numbers are sometimes called ‘the continuum’ (cf. Weyl), and its cardinality is denoted by \mathfrak{c} .

The *continuum hypothesis* states that any uncountable set $A \subset \mathbb{R}$ satisfies $|A| \geq |\mathbb{R}|$. This statement is *undecidable* in traditional set theory, ZFC.

The idea of the real numbers can be traced back to Euclid and plane geometry, where the real numbers appear as a *geometric line*. There is an interesting philosophical point here: classically, one can speak of a point on a line, but it is a major shift of viewpoint (from the synthetic to the analytical) to think of a line as simply a collection of points.

The modern perspective on \mathbb{R} , based on axioms and set theory, was not universally accepted at first (cf. Brouwer). And as we will discuss below, it is worth noting that most points in \mathbb{R} *have no names*, and it is these nameless points that form the glue holding the continuum together.

Intervals and open sets. We now return to a down-to-earth study of the real numbers. The simplest subsets of the real numbers are the *open intervals* (a, b) ; we allow $a = -\infty$ and/or $b = +\infty$. We can also form closed intervals $[a, b]$ or half-open intervals $[a, b)$, $(a, b]$.

Proposition 2.1 *Every open set $U \subset \mathbb{R}$ is a finite or countable union of disjoint open intervals, $U = \bigcup (a_i, b_i)$.*

Proof. The components U_α of U (the maximal open intervals it contains) are clearly disjoint and their union is U . They are countable in number because different U_α contain different rational numbers. ■

Note: in this proof we have implicitly used the axiom of choice to pick a rational number from each open interval. This can also be done explicitly.

Warning: the intervals forming U need not come *in order*, and in fact there exist examples (such as the complement of the Cantor set) where a third subinterval exists between any two subintervals of U .

Proposition 2.2 *The collection of all open subsets of \mathbb{R} has the same cardinality as \mathbb{R} itself.*

Proof. An open set is uniquely determined by the collection of open intervals with rational endpoints that it contains. ■

Remark: $\mathbb{N}^{\mathbb{N}}$ and the irrational numbers. The set of irrational numbers $\mathcal{I} \subset [0, 1]$ is isomorphic to $\mathbb{N}^{\mathbb{N}}$ by the continued fraction map

$$(a_0, a_1, \dots) \mapsto 1/(b_1 + 1/(b_2 + \dots)),$$

where $b_i = a_i + 1$. In fact this map is a homeomorphism.

Algebras of sets. It will turn out that there are some subsets of \mathbb{R} (constructed with the Axiom of Choice) that are so exotic, there is no reasonable way to assign them a measure. But for the purposes of analysis, we do not need to work with arbitrary subsets of \mathbb{R} , only a collection which is rich enough that it includes the open sets and is closed under basic set-theoretic operations and limits.

To be more precisely, we say a collection of sets $\mathcal{A} \subset \mathcal{P}(\mathbb{R})$ forms an *algebra* if

1. $\emptyset \in \mathcal{A}$,
2. $E \in \mathcal{A} \implies \tilde{E} = \mathbb{R} - E \in \mathcal{A}$; and
3. $E, F \in \mathcal{A} \implies E \cup F \in \mathcal{A}$.

This is equivalent to saying that the collection of indicator functions $\chi_E \in 2^{\mathbb{R}}$, $E \in \mathcal{A}$, form an algebra over the field with 2 elements. Note, for example, that

$$\chi_{E \cup F} = \chi_E + \chi_F - \chi_E \chi_F.$$

We say \mathcal{A} forms a σ -algebra if it is closed under countable unions; that is, if for any sequence (E_1, E_2, \dots) of elements of \mathcal{A} , we have

$$\bigcup E_i \in \mathcal{A}.$$

The Borel sets. The *Borel sets* $\mathcal{B} \subset \mathcal{P}(\mathbb{R})$ are the smallest σ -algebra containing the open sets. To see there is such a σ -algebra, simply take the intersection of *all* σ -algebras containing the open sets. This shows that \mathcal{B} is uniquely determined.

They are not the simplest σ -algebra though. The smallest single algebra containing the *singletons* is the algebra of countable and co-countable sets. This algebra makes no reference to topology.

The interval algebra. As a warm-up to the Borel sets, one can also consider the algebra \mathcal{A} generated by the open intervals (a, b) . It turns out the elements of \mathcal{A} are all sets of the form $E = \bigcup_1^n (a_i, b_i) \cup F$, where F is finite. If we consider single points as closed intervals, we can simply say that the elements of \mathcal{A} are *finite unions of intervals*.

The algebra \mathcal{A} can also be constructed from the outside, by taking the intersection of all algebras containing the open intervals. But it can also be constructed from the inside, inductively. We let \mathcal{A}_0 be the set of all open intervals, and define \mathcal{A}_{i+1} by adjoining to \mathcal{A}_i all finite unions and complements of elements in \mathcal{A}_i . It is then clear that $\mathcal{A} = \bigcup \mathcal{A}_i$ is an algebra, and that it is the smallest algebra containing the open intervals. Here we have used the fact that any 2 sets E and F are already present at some finite stage \mathcal{A}_i .

Transfinite induction. In a similar way, \mathcal{B} can be constructed by induction over the first uncountable ordinal Ω . The most important property of this well-ordered set is that any countable set $I \subset \Omega$ has an upper bound. (Compare this with the ordinal ω , which has the property that any *finite set* has an upper bound.)

We then define \mathcal{B}_0 to be the set of open sets (or even open intervals) in \mathbb{R} , define $\mathcal{B}_{\alpha+1}$ be adjoining to \mathcal{B}_α the complements and countable unions of the sets it contains, and setting $\mathcal{B}_\gamma = \bigcup_{\alpha < \gamma} \mathcal{B}_\alpha$ for limit ordinals in Ω . It is then readily verified that

$$\mathcal{B} = \bigcup_{\alpha < \Omega} \mathcal{B}_\alpha.$$

The important point here is that if $E_1, E_2, \dots \in \mathcal{B}$ then these sets all belong to some \mathcal{B}_α , and hence $\bigcup E_i \in \mathcal{B}_{\alpha+1} \subset \mathcal{B}$. Thus every Borel set is ‘born’ at some stage in this inductive process.

The Borel hierarchy. The early stages of the Borel hierarchy have standard names. We say E is a G_δ set if it is a countable intersection of open sets; and E is an F_σ set if it is a countable union of closed sets. A countable union of G_δ sets is a $G_\delta\sigma$ set, and so on.

Example. Let $\langle f_n \rangle$ be a sequence of positive continuous function on \mathbb{R} , and let

$$E = \{x : \langle f_n(x) \rangle \text{ is bounded}\}.$$

Then E is an F_σ set. Indeed, we can write

$$E = \bigcup_{M=1}^{\infty} \{x : f_n(x) \leq M \forall n\},$$

and each set appearing in this union is closed.

Exercise: what is E for the sequence of functions

$$f_n(x) = \sum_{k=1}^n |\sin(\pi k!x)|^{1/n}?$$

In fact, E consists exactly of the rational numbers.

How many open sets are there? It is useful know that, while the number of subsets of \mathbb{R} is greater than \mathbb{R} , the number of tame subsets tends to be less. For example we have:

Theorem 2.3 *The set of all open subsets of \mathbb{R} is of the same cardinality as \mathbb{R} itself.*

Proof. Let Q denote the countable set of intervals with rational endpoints. An open set U is uniquely determined by the element $I \in Q$ that it contains, and thus the collection of all open sets is no larger than $|\mathcal{P}(Q)| = \mathfrak{c}$. ■

Corollary 2.4 *The number of closed subsets of \mathbb{R} is the same as the number of points in \mathbb{R} .*

Remark: the number of Borel sets. If we examine the inductive construction of the Borel sets, we find similarly that $|\mathcal{B}_\alpha| = \mathfrak{c}$ for all $\alpha < \Omega$. But it is easy to see that $|\Omega| < \mathfrak{c}$ and the union of a continuum number of copies of the continuum still has cardinality \mathfrak{c} (i.e. $|\mathbb{R}^2| = |\mathbb{R}|$), and thus the number of Borel sets is also equal to \mathfrak{c} .

As a corollary, *most* subsets of \mathbb{R} are not Borel sets, even though the vast collection of Borel sets is more than enough for many purposes in analysis.

Note: it is a general theorem in cardinal arithmetic that $\kappa^2 = \kappa$ if κ is an infinite cardinal.

3 Lebesgue Measurable Sets

Imagine the real line as a long, even strand of copper wire, weigh 1 (gram) per unit (centimeter, say).

A subset $E \subset \mathbb{R}$ gives us a piece of the real line which we can weigh — or does it? If so, what would its weight be? The theory of Lebesgue measure provides us with a large collection of *measurable sets*, that can be weighed, and tells us their properties. It forms the basis of integration, since $\int \chi_E = m(E)$ and the indicator functions come close to spanning all the measurable functions (those which can be integrated).

Goal. On \mathbb{R} we will construct:

- A σ -algebra \mathcal{M} containing the Borel sets, and
- A **measure** $m : \mathcal{M} \rightarrow [0, \infty]$, such that
- The measure of any interval has the expected value, $m([a, b]) = b - a$;
- The measure is *countably* additive: if the sets $E_i \in \mathcal{M}$ are disjoint, then

$$m\left(\bigcup E_i\right) = \sum m(E_i); \quad \text{and}$$

- The measure is translation invariant: $m(E + t) = m(E)$.

Outer measure. We begin by defining, for an arbitrary set $A \subset \mathbb{R}$, its *outer measure* $m^*(A)$. This is given by

$$m^*(A) = \inf \left\{ \sum \ell(I_i) : A \subset \bigcup_1^\infty I_i \right\},$$

where (I_i) is a collection of intervals (a_i, b_i) , and where $\ell(a, b) = b - a$.

Here are some of its basic properties:

Monotonicity. If $A \subset B$, then $m^*(A) \leq m^*(B)$.

Subadditivity. For any sequence of sets A_i , $m^*(\bigcup A_i) \leq \sum m^*(A_i)$.

Normalization: $m^*[a, b] = b - a = \ell([a, b])$.

Proof. Clearly $m^*[a, b] \leq b - a$. But if $[a, b]$ is covered by $\bigcup I_k$, by compactness we can assume the union is finite, and then

$$b - a = \int \chi_{[a, b]} \leq \int \sum \chi_{I_k} = \sum |I_k|,$$

so we also have $b - a \leq m^*[a, b]$. ■

In the foregoing proof we have used the Riemann integral, which is fine for bounded functions with finitely many discontinuities. An elementary argument could also be given.

Example. The outer measure of a single point is zero. By countable subadditivity, the same is true for any countable set; in particular,

$$m^*(\mathbb{Q}) = 0.$$

Measurable sets. It is a remarkable fact that the *measurable sets* \mathcal{M} form a σ -algebra which admits a ‘direct’ definition, i.e. rather than giving its generators we can be give a characterization of which sets belong to \mathcal{M} .

A set $E \subset \mathbb{R}$ is *measurable* if

$$m^*(E \cap A) + m^*(\tilde{E} \cap A) = m^*(A)$$

for all sets $A \subset \mathbb{R}$. This means E *cuts* any set A cleanly into two pieces whose outer measures add back up to the outer measure of A .

Because of subadditivity, only one direction needs to be checked: to show E is measurable we must show

$$m^*(E \cap A) + m^*(\tilde{E} \cap A) \leq m^*(A)$$

for all A . In particular we can always assume $m^*(A)$ is finite.

Examples of measurable sets. The simplest point is that *sets of measure zero* are measurable. This is because $m^*(E \cap A) = 0$, and of course $m^*(\tilde{E} \cap A) \leq m^*(A)$ by monotonicity.

Theorem 3.1 $E = [a, \infty)$ is measurable.

Proof. Given $\epsilon > 0$, pick a covering $\bigcup I_i$ for A such that $\sum \ell(I_i) \leq m^*(A) + \epsilon$. By intersecting I_i with E and \tilde{E} , we obtain covers I'_i for $E \cap A$ and I''_i for $\tilde{E} \cap A$ which show

$$m^*(E \cap A) + m^*(\tilde{E} \cap A) \leq \sum \ell(I'_i) + \ell(I''_i) = \sum \ell(I_i) \leq m^*(A) + \epsilon.$$

Since $\epsilon > 0$ was arbitrary, this shows E is measurable. ■

Here is one tricky point in the development of measurable sets.

Theorem 3.2 *The measurable sets form an algebra.*

Proof. Closure under complements is by definition. Now suppose E and F are measurable, and we want to show $E \cap F$ is. By the definition of measurability, E cuts A into two sets whose outer measures add up to the measure of A . Now F cuts $E \cap A$ into two sets whose outer measures add up, and similarly for the complements. Thus E and F cut A into 4 sets whose measures add up to the outer measure of A . Assembling 3 of these to form $A \cap (E \cup F)$ and the remaining one to form $A \cap \widetilde{E \cup F}$, we see $E \cup F$ is measurable.

Theorem 3.3 *If E_i are disjoint and measurable, $i = 1, 2, \dots, N$, then $\sum m^*(E_i \cap A) = m^*(A \cap \bigcup E_i)$.*

Proof. Let $A' = A \cap (E_1 \cup E_2)$. Then since E_1 is measurable,

$$m^*(A') = m^*(A' \cap E_1) + m^*(A' \cap \widetilde{E_1}) = m^*(A \cap E_1) + m^*(A \cap E_2).$$

This proves the theorem for $N = 2$. The general case follows by induction. ■

Theorem 3.4 *The measurable sets form a σ -algebra.*

Proof. Suppose E_i is a sequence of measurable sets; we want to show $\bigcup E_i$ is measurable. Since we already have an algebra, we can assume the E_i are disjoint. By the preceding lemma, we have for any finite N ,

$$\sum_1^N m^*(E_i \cap A) + m^*\left(A \cap \bigcap_1^N \widetilde{E_i}\right) = m^*(A).$$

The second term is only smaller for an infinite intersection, so letting $N \rightarrow \infty$ we get

$$\sum_1^\infty m^*(E_i \cap A) + m^*\left(A \cap \bigcap_1^\infty \widetilde{E_i}\right) \leq m^*(A).$$

By countable subadditivity of outer measure, the first term dominates $m^*(A \cap \bigcup E_i)$, so we are done. ■

Corollary 3.5 *All Borel sets are measurable.*

Definition. The Lebesgue measure of $E \in \mathcal{M}$ is defined by $m(E) = m^*(E)$.

Theorem 3.6 *Lebesgue measure is countably additive. That is, if E_i is a sequence of disjoint measurable sets, then $m(\bigcup E_i) = \sum m(E_i)$.*

Proof. This follows from the preceding proof. Explicitly, we first have finite additivity, and then for every N ,

$$\sum_1^\infty m(E_i) \geq m\left(\bigcup E_i\right) \geq m\left(\bigcup_1^N E_i\right) = \sum_1^N m(E_i),$$

which gives the desired result by taking N to infinity. ■

Continuity of measure. Here is a useful way to think of measure as a ‘continuous’ function of E , at least for monotone limits.

Theorem 3.7 *If $m(E_1)$ is finite and $E_1 \supset E_2 \supset E_3 \dots$, then $m(\bigcap E_i) = \lim m(E_i)$.*

Proof. Let $F = \bigcap E_i$ and write $E_1 = F \cup (E_1 - E_2) \cup (E_2 - E_3) \cup \dots$. Then we have

$$m(E_1) = m(F) + \sum_1^\infty m(E_i - E_{i+1}) = m(F) + m(E_1) - \lim m(E_i),$$

which gives the desired result. ■

Note that the indicator function of $\bigcap E_i$ is the pointwise limit of the functions χ_{E_i} , so it is reasonably to rewrite the conclusion as:

$$m(\lim E_i) = \lim m(E_i),$$

which looks more like ‘continuity’.

Borel–Cantelli and probability theory. Here is a result from probability theory that can be phrased in terms of measure. It has two parts.

Theorem 3.8 *Suppose $\sum m(E_i) < \infty$. Then the set of points that belong to infinitely many E_i has measure zero.*

Proof. Let A be the points that belong to infinitely many E_i . Then

$$m(A) \leq m\left(\bigcup_N^\infty E_i\right) \leq \sum_N^\infty m(E_i) \rightarrow 0$$

as $N \rightarrow \infty$; thus $m(A) = 0$. ■

In probability theory, an *event* is given by a measurable subset $A \subset [0, 1]$, and its *probability* is simply its measure: $P(A) = m(A)$. The result above says that if $\sum P(A_i)$ is finite, then almost surely only finitely many of these events occur.

This theorem has a converse if we assume the events are *independent*. This means that $P(A \cap B) = P(A)P(B)$, $P(A \cap B \cap C) = P(A)P(B)P(C)$, etc. for distinct events A, B, C .

Theorem 3.9 *If the events A_i are independent and $\sum P(A_i) = \infty$, then almost surely infinitely many events occur.*

Proof. The probability that no more than the first $N - 1$ events occur is given by

$$P\left(\bigcap_N \widetilde{A}_i\right) = \prod_N (1 - P(A_i)) = 0$$

because $\sum P(A_i)$ diverges (this is a general fact about infinite products). Thus with probability one, infinitely many events occur. ■

Example: normal numbers. Let us say a number $x \in [0, 1]$ is (weakly) *normal* if any finite sequence of digits occurs infinitely often in its decimal expansion. We claim the set E of all such numbers has $m(E) = 1$. To see this, let E_i be the set of $x \in [0, 1]$ such that the i th digit of x is 1. Then the E_i are independent, and $m(E_i) = 1/10$, so $\sum m(E_i) = \infty$. Thus the digit 1 occurs infinitely often for *almost every* $x \in [0, 1]$. The same reasoning applies to any finite sequence of digits. Intersecting these countably many sets of measure one again yields a set of full measure.

Hamlet. If you take x at random and convert it to a binary number, then to text, you will find infinitely many copies of Hamlet. It is widely believed, but not known, if numbers like π or $\sqrt{2}$ are normal. On the other hand one can give specific examples of normal numbers, such as $x = 0.1234567891011121314151617\dots$

A stronger notion of normality is that the digit 1 occurs with density 1/10th, and the same for any other finite sequence. It is also known that almost all numbers are normal in this sense.

Littlewood's principles. Littlewood remarked that the Lebesgue theory is actually fairly simple to understand intuitively, if phrased somewhat informally; namely:

1. A measurable set is nearly a finite union of intervals;

2. A measurable function is nearly continuous; and
3. A pointwise convergent sequence of measurable functions is nearly uniformly convergent.

We will give a precise quote from Littlewood when we consider measurable functions.

The first principle. We can now make the first principle precise, and prove it.

Theorem 3.10 *Suppose $m(E)$ is finite. Then for any $\epsilon > 0$ we can find a finite union of intervals J such that*

$$m(J\Delta E) < \epsilon.$$

Here the symmetric difference is defined by

$$A\Delta B = (A - B) \cup (B - A).$$

The quantity $d(A, B) = m(A\Delta B)$ is a good way to measure the ‘distance’ between measurable sets.

Here is a complement to the result above, which will be used in its proof.

Theorem 3.11 *Let $E \subset \mathbb{R}$ be a measurable set. Then we can find:*

1. *Closed and open sets $F \subset E \subset U$ such that $m(E - F) < \epsilon$ and $m(F - U) < \epsilon$; and*
2. *F_σ and G_δ sets $F' \subset E \subset U'$ such that $m(E - F') = m(U' - E) = 0$.*

Simplifying Borel sets. As a Corollary, every measurable set is the union of an F_σ and a set of measure zero. In particular, every *Borel set* is just an F_σ , if we are willing to neglect sets of measure zero.

Proof of both results. We treat the case where $E \subset [0, 1]$; the general case is similar. Since E has finite measure, there are open intervals such that $E \subset \bigcup I_i$ and

$$\sum \ell(I_i) = m\left(\bigcup I_i\right) < m(E) + \epsilon.$$

It follows that $m(U - E) < \epsilon$ for $U = \bigcup I_i$. If we take a finite union $J_n = \bigcup_1^n I_i$, then $m(J_n) \rightarrow m(U)$, and hence for n large enough we have $m(J_n\Delta E) < 2\epsilon$. This is Littlewood’s first principle. Moreover, we can find

open sets with $E \subset U_n$ such that $m(U_n - E) < 2^{-n}$; then $U' = \bigcap U_n$ gives a G_δ containing E and differing from it by a set of measure zero.

To show E can be approximated from the inside by closed sets and by an F_σ , just take complements in the argument above. For example, note that $A - B = \tilde{B} - \tilde{A}$; thus an open set $U \supset \tilde{E}$ with $m(U - \tilde{E}) < \epsilon$ yields a closed set $\tilde{U} \subset E$ with $m(E - \tilde{U}) = m(U - \tilde{E}) < \epsilon$. ■

Another approach to \mathcal{M} . One could turn these results around and use them to *define* the measurable sets \mathcal{M} . Namely we could say a set is measurable if it has the form $E = U \triangle A$ where $m^*(A) = 0$ and U is a G_δ set. This property is *clearly* closed under countable intersections, but some work is required to show it is closed under complements. That is, one must show that every G_δ set agrees with an F_σ set up to a set of measure zero.

Here is a useful ‘density result’ for sets of positive measure (we will later prove a more precise result in the same direction).

Corollary 3.12 *Let $E \subset \mathbb{R}$ have positive measure. Then for any $\epsilon > 0$ there is an open interval I such that $m(I \cap E)/m(I) > 1 - \epsilon$.*

Proof. We may assume $m(E)$ is finite, and that $U \supset E$ is an open set which closely approximates E . By rescaling, we may assume $m(U) = 1$ and $m(E) = 1 - \epsilon$. Write $U = \bigcup I_i$ as a union of disjoint open intervals. We then have

$$m(E) = 1 - \epsilon = \sum m(E \cap I_i) = \sum \frac{m(E \cap I_i)}{m(I_i)} m(I_i).$$

Since $\sum m(I_i) = 1$, this says the weighted average density of E in I_i is $1 - \epsilon$. Thus for some particular I_i , we have density at least $1 - \epsilon$, as desired. ■

Nonmeasurable sets. We will now justify the complicated definition of measurable sets by showing there *exists* a non-measurable set. Indeed, we will show there is a set for which no reasonable measure can be defined, if we require translation invariance and countable additivity.

Consider \mathbb{Q} as a subgroup of the additive group \mathbb{R} , and let $A \subset \mathbb{R}$ be a set of coset representatives for \mathbb{R}/\mathbb{Q} . That is, we choose A such that the cosets $a + \mathbb{Q}$ with $a \in A$ are disjoint and cover \mathbb{R} . How is A chosen? By the Axiom of Choice: we let H range over the cosets of \mathbb{Q} , and we apply a choice function for $\mathcal{P}(\mathbb{R})$ to each coset to get A . Note that A is uncountable.

The key point of A is *every* real number can be written *uniquely* as $x = a + q$ with $a \in A$ and $q \in \mathbb{Q}$.

Theorem 3.13 *The set A is nonmeasurable.*

Proof. Suppose A is measurable, and let $S = [0, 1] \cap \mathbb{Q}$. The countably many sets $s + (A \cap [a, b])$ with $s \in S$ are pairwise disjoint and contained in $[a, b + 1]$, and they all have the same measure; thus they all have measure zero. Therefore $m(A \cap [a, b]) = 0$ for any interval $[a, b]$, and hence $m(A) = 0$. But then $\mathbb{R} = \bigcup_{\mathbb{Q}}(q + A)$ has measure zero, a contradiction. ■

By similar reasoning one can show:

Theorem 3.14 *Any set of positive measure E contains a non-measurable set.*

Proof. We may assume E is bounded. Consider the equivalence relation on E given by $e \sim e'$ if $e - e' \in \mathbb{Q}$. Let $A \subset E$ be a set with exactly one element from each equivalence class. Suppose A is measurable. Then as above, the sets $s + A$ with $s \in \mathbb{Q} \cap [0, 1]$ are disjoint and contained in a bounded set. Thus $m(A) = 0$. But $E \subset \bigcup_{\mathbb{Q}}(q + A)$, contradicting the fact that $m(E) > 0$. ■

Remark: basis for \mathbb{R} as a vector space. With some more care, one can similarly show that if B is a basis for \mathbb{R} as a vector space over \mathbb{Q} , then B is nonmeasurable.

The first step is to observe that if B is measurable, then $m(B) > 0$. This is because every $x \in \mathbb{R}$ can be uniquely expressed, for some N , as $\sum_1^N q_i b_i$ with $(b_i) \in B^N$ and $q_i \in \mathbb{Q}$. If B has measure zero, then so does this rational projection of B^N , by an easy argument; but then \mathbb{R} has measure zero.

To finish the proof, we use the fact that the sets qB with $q \in \mathbb{Q}^*$ are all disjoint, and mimic the argument above. Of course here translation invariance must be replaced by the fact that $m(sB) = |s| \cdot m(B)$, but this makes little difference when s ranges in a bounded set.

A strange subset of the plane. Assume the Continuum Hypothesis. Then we can well-order $[0, 1]$ such that each initial segment is countable. Set $R = \{(x, y) : x < y\}$ in this ordering. Then horizontal slices (fixing y) have measure zero, while all vertical slices (fixing x) have measure one.

The Cantor set. Given the existence of nonmeasurable sets, it is useful to keep in mind that any set with $m^*(A) = 0$ is measurable. And since every measurable set looks fairly simple up to a set of zero, it is good to have examples of sets of measure zero.

An important example in topology and measure theory is the classical *Cantor set* $K \subset [0, 1]$. It shows, for example, that a set of measure zero need not be countable.

We can also use K together with the *Cantor function* to show that there are measurable sets that are not Borel sets. The Cantor function itself is a nice (paradoxical?) example of a nonconstant, continuous, monotone increasing function with $f'(x) = 0$ a.e.

Middle thirds. The Cantor set ‘middle third’ set can be defined in several ways. One way is as follows. Given an interval $I = [a, b]$, we define its middle third as the centered open subinterval U of length $L = (b - a)/3$. By cutting this out, we obtain 2 intervals of equal length,

$$I' \cup I'' = [a, a + L] \cup [a + 2L, b].$$

Now let $K_0 = [0, 1]$, let $K_1 = [0, 1/3] \cup [2/3, 1]$ be the result of removing the middle third from K_0 , and let K_{n+1} be the result of removing the middle $(1/3)$ from each of the 2^n intervals that make up K_n . Then the Cantor set is defined by

$$K = \bigcap_0^\infty K_n.$$

Theorem 3.15 *The middle third Cantor set has measure zero.*

Proof 1. By induction we have $m(K_n) = (2/3)^n$ — we remove $1/3$ of what remains at each stage — so $m(K) = \lim m(K_n) = 0$.

Proof 2. The total length of the intervals removed is given by

$$\frac{1}{3} + 2\frac{1}{9} + 4\frac{1}{27} + \cdots = (1/3) \sum_0^n (2/3)^n = 1.$$

■

Hausdorff dimension. The Cantor set is an example of a fractal. By construction, K is made of 2 copies of K , each scaled down by $1/3$. One can argue from this that the Hausdorff dimension δ of K is the solution to $1 = 2(1/3)^\delta$, which gives $\delta = \log 2 / \log 3 = 0.6309 \dots$

One can similarly construct a Cantor middle- α set K_α of measure zero, for any $0 < \alpha < 1$. Its dimension ranges in $(0, 1)$ as α varies.

Base three. Alternatively, we can define K as the set of all numbers in $[0, 1]$ which can be written in base 3 as $x = 0.x_1x_2x_3 \dots$ with each x_i equal

to 0 or 2. For example, the point $1 = 0.2222\dots$ is included in K , as is the point $x = 0.020202\dots = 1/4$.

The first ‘middle third’ consists of numbers which require $x_1 = 1$; the second pair of middle thirds, those which require $x_2 = 1$; and so on.

The Cantor function. One could also say that K consists of all numbers of the form

$$x = \sum_1^{\infty} 2y_i/3^i,$$

where (y_i) is a sequence of zeros and ones. This expression for x is unique. The *Cantor function*

$$f : K \rightarrow [0, 1]$$

is defined by $f(x) = \sum y_i/2^i$. In other words, it converts a base-3 expression of 0’s and 2’s into a base-2 expression of 0’s and 1’s. Since we can start with a base 2 expression and work backwards, the map f is surjective, and it is easily seen to be monotone and continuous. This shows:

Theorem 3.16 *The Cantor set is uncountable; indeed, $|K| = |\mathbb{R}|$.*

The Cantor function on K is not quite 1 – 1; the two endpoints of any complementary interval in $[0, 1] - K$ are identified by f . For example, $f(1/3) = f(2/3) = 1/2$, since

$$f(1/3) = f(0.02222\dots_3) = 0.01111\dots_2 = 0.1_2 = 1/2 = f(0.2_3) = f(2/3).$$

In fact, f has a unique extension to a monotone function $f : [0, 1] \rightarrow [0, 1]$, which is constant on each interval in the complement of K .

Probabilistic interpretation. There is an obvious way to choose a point $X \in K$ at random: construct the digits of X by flipping a coin countably many times. We can then say, in probability language:

$$f(x) = P(X < x),$$

i.e. $f(x)$ is the probability that a randomly constructed point in the Cantor set is less than x . It is then obvious that f is locally constant outside of K .

The devil’s staircase. Since the complementary intervals have full measure, this function has the amazing property that it climbs from 0 to 1 but for any randomly chosen point (i.e. for a set of full measure) we have $f'(x) = 0$. It is as if a strange particle travels between 2 points in space, but whenever it is observed, it is stationary.

Measurable but not Borel. Of course non-measurable sets are not Borel sets, but we can still ask:




Figure 1. The ends of this tree are the Cantor set $K \subset [0, 1]$.




Figure 2. Cantor's function: the devil's staircase.

Q. Is there a measurable set which is not a Borel set?

The answer is *yes*. For example, *any* set $E \subset K$ is measurable, since $m(E) = m(K) = 0$. But $|K| = |\mathbb{R}|$, so there are way too many elements in $\mathcal{P}(K)$ for all of them to be Borel.

For a concrete example, it is useful to turn the Cantor function into a homeomorphism by making it climb on the complementary intervals. This is done by setting

$$h(x) = x + f(x).$$

Then $h : [0, 1] \rightarrow [0, 2]$ is a homeomorphism with many interesting properties. For example, we have:

$$m(h(K)) = 1,$$

since $m(h([0, 1] - K)) = 1$. One can guess this equation from the fact that h doubles the size of $[0, 1]$, but preserves measure outside of K . In any case, since h sends K to $K' = h(K)$ with $m(K) = 0$ and $m(K') = 1$, we can conclude:

Theorem 3.17 *There exist a homeomorphisms $h : [a, b] \rightarrow [c, d]$ that sends a set of measure zero to a set of positive measure, and vice-versa.*

(For the vice-versa part, consider h^{-1} .)

Now let $A \subset K' = h(K)$ be a non-measurable set — which exists since $m(K') > 0$. Then $B = h^{-1}(A) \subset K$ is a subset of the Cantor set that *is not Borel*. For if B were Borel, then $A = h(B)$ would also be Borel, and hence measurable.

The danger of differences. A more subtle phenomenon arises when we consider the difference set $A - A$, comprised of the numbers $x - y$ with $x, y \in A$. A standard exercise is to show that if $m(A) > 0$ then $A - A$ contains an interval. But the condition $m(A) > 0$ is not necessary; in fact $K - K = [-1, 1]$. The fact that a set of measure zero, like K , can have a difference set of positive measure, is a hint that the following ‘pathology’ holds: there exist measurable sets A such that $A - A$ is not measurable. The source of the problem is the following: while it is true that $A \times A$ is a measurable subset of \mathbb{R}^2 , it is not generally true that the projection of a measurable set from \mathbb{R}^2 to \mathbb{R} is measurable.

An open problem about the Cantor set. Is every $x \in K$ either rational or transcendental? As remarked above, all indications are that every

irrational algebraic number is normal. But in base 3 the elements of K are highly abnormal, so they should not be algebraic irrationals.

Appendix: Finitely-additive measures on \mathbb{N} . The natural numbers admit a finitely-additive measure defined on *all* subsets, and vanishing on finite sets. (Such a measure is cannot be countably additive.) This construction gives a ‘positive’ use of the Axiom of Choice, to construct a measure rather than to construct a non-measurable set.

A *filter* is a collections of sets $\mathcal{F} \subset \mathcal{P}(X)$ such that sets in \mathcal{F} are ‘big’:

- (1) $\emptyset \notin \mathcal{F}$,
- (2) $A \in \mathcal{F}, A \subset B \implies B \in \mathcal{F}$; and
- (3) $A, B \in \mathcal{F} \implies A \cap B \in \mathcal{F}$.

Example: the cofinite filter (if X is infinite).

Example: the ‘principal’ ultrafilter \mathcal{F}_x of all sets with $x \in F$. This is an *ultrafilter*: if $X = A \sqcup B$ then A or B is in \mathcal{F} .

Theorem 3.18 *Any filter is contained in an ultrafilter.*

Proof. Using Zorn’s lemma, take a maximal filter \mathcal{F} containing the given one. Suppose neither A nor $X - A$ is in \mathcal{F} . Adjoining to \mathcal{F} all sets of the form $F \cap A$, we obtain a larger filter \mathcal{F}' , a contradiction. (To check $\emptyset \notin \mathcal{F}'$: if $A \cap F = \emptyset$ then $X - A$ is a superset of F , so $X - A$ *was* in \mathcal{F} .)

Ideals and filters. In the ring $R = (\mathbb{Z}/2)^X$, ideals $I \neq R$ and filters are in bijection: $I = \{A : \widetilde{A} \in \mathcal{F}\}$. The ideal consists of ‘small’ sets, those whose complements are big.

(By (2), $A \in I \implies AB \in I$. By (3), $A, B \in I \implies A \cup B \in I \implies (A \cup B)(A \Delta B) = A + B \in I$.)

Lemma: if \mathcal{F} is an ultrafilter and $A \cup B = F \in \mathcal{F}$ then A or B is in \mathcal{F} .

Proof. We prove the contrapositive. If neither A nor B is in \mathcal{F} , then their complements satisfy $\widetilde{A}, \widetilde{B} \in \mathcal{F}$. Since \mathcal{F} is a filter,

$$\widetilde{A} \cap \widetilde{B} = \widetilde{A \cup B} \in \mathcal{F}$$

and thus $A \cup B \notin \mathcal{F}$.

Corollary: Ultrafilters correspond to prime ideals.

By Zorn’s Lemma, every ideal is contained in a maximal ideal; this gives another construction of ultrafilters.

Measures. Let \mathcal{F} be an ultrafilter. Then we get a finitely-additive measure on **all** subsets of X by setting $m(F) = 1$ or 0 according to $F \in \mathcal{F}$ or not.

Conversely, any 0/1-valued finitely additive measure on $\mathcal{P}(X)$ determines a filter.

Measures supported at infinity. The most interesting case is to take the cofinite filter, and extend it in some way to an ultrafilter. Then we obtain a finitely-additive measure on $\mathcal{P}(X)$ such that points have zero measure but $m(X) = 1$. When $X = \mathbb{N}$ such a measure cannot be countably additive.

4 Measurable Functions

In this section we begin to study the interaction of measure theory with functions on the real line.

Theorem 4.1 *Given $f : \mathbb{R} \rightarrow \mathbb{R}$, the following conditions are equivalent.*

1. $\{x : f(x) > a\}$ is measurable for all $a \in \mathbb{R}$.
2. $f^{-1}(U)$ is measurable for all open sets U .
3. $f^{-1}(B)$ is measurable for all open Borel sets B .

A function is *measurable* if any (and hence all) of these conditions hold. The first condition is the easiest to check.

Proof. Let $\mathcal{A} \subset \mathcal{P}(\mathbb{R})$ be the collection of sets $A \subset \mathbb{R}$ such that $f^{-1}(A)$ is measurable. Then \mathcal{A} forms a σ -algebra. Since the sets (a, ∞) generate the Borel sets as a σ -algebra, (1) \implies (3). The implications (3) \implies (2) \implies (1) are immediate. ■

First examples: continuous, monotone and indicator functions. Let $C(\mathbb{R})$ denote the space of all continuous functions on \mathbb{R} , and let $M(\mathbb{R})$ denote the set of all measurable functions on \mathbb{R} . Clearly we have $C(\mathbb{R}) \subset M(\mathbb{R})$, since open sets are measurable.

In addition $M(\mathbb{R})$ contains monotone functions, since for these the preimage of an interval is another interval. The indicator functions χ_E of any measurable set is also easily shown to be measurable.

Algebraic structure. We now examine which operations we can form to make new measurable functions out of existing ones. It is well-known that $C(\mathbb{R})$ is an *algebra*, meaning if $f, g \in C(\mathbb{R})$ then so are $f + g$, fg and αf , $\alpha \in \mathbb{R}$.

Theorem 4.2 *The space $M(\mathbb{R})$ is an algebra, containing the continuous functions.*

Proof. If f is continuous and U is open, then $f^{-1}(U)$ is open, and hence measurable. Thus $C(\mathbb{R}) \subset M(\mathbb{R})$. It is clear that $M(\mathbb{R})$ is closed under scalar multiplication.

The tricky part is addition. Suppose $f, g \in M(\mathbb{R})$ and $f(x) + g(x) > a$. Then we can find a rational number p/q such that $f(x) > p/q$ and $p/q + g(x) > a$. (Just take p/q between $f(x)$ and $a - g(x)$.) And of course, this condition implies $f(x) + g(x) > a$. Thus we have:

$$\{x : f(x) + g(x) > a\} = \bigcup_{p/q \in \mathbb{Q}} \{x : f(x) > p/q\} \cap \{x : p/q + g(x) > a\}.$$

This expresses the set on the left as a countable union of measurable sets, so it is measurable.

As for products, we note that $(f + g)^2 - f^2 - g^2 = 2fg$, so it suffices to show that $M(\mathbb{R})$ is closed under $f \mapsto f^2$. This follows from the fact that

$$\{x : f(x)^2 > c^2\} = \{x : f(x) > c\} \cup \{x : f(x) < -c\}.$$

■

Warning. Unlike the continuous functions, the algebra $M(\mathbb{R})$ is *not* closed under composition! However it is easy to show:

Proposition 4.3 *If $h : \mathbb{R} \rightarrow \mathbb{R}$ is continuous and f is measurable, then $h \circ f$ is also measurable.*

Analytic properties: limits. We say $f_n \rightarrow f$ *pointwise* if $f_n(x) \rightarrow f(x)$ for all $x \in \mathbb{R}$. The key advantage of the measurable functions over the continuous functions is the following:

Theorem 4.4 *The space $M(\mathbb{R})$ is closed under limits: if $f_n \in M(\mathbb{R})$ and $f_n \rightarrow f$ pointwise, then $f \in M(\mathbb{R})$.*

Proof. For any $c \in \mathbb{R}$ we have

$$\begin{aligned} \{x : f(x) > c\} &= \{x : \exists k \exists N \forall n \geq N f_n(x) > c + 1/k\} \\ &= \bigcup_k \bigcup_N \bigcap_{n \geq N} \{x : f_n(x) > c + 1/k\}, \end{aligned}$$

and the latter set is measurable because the functions f_n are measurable. ■

We remark that $M(\mathbb{R})$ is also closed under other limit such as $\liminf f_n$ and $\limsup f_n$, as well as $\sup f_n$ and $\inf f_n$.

Logic and set theory. In the proof above we have used two basic principles, common in real analysis: (i) the set of x satisfying a condition $Q(x)$, written as $\forall \exists \dots P(x)$, can always be described as $\bigcap \bigcup \dots P(x)$; i.e. quantifiers can be turned into unions and intersections; and (ii) conditions that range over uncountable sets, such as $\forall \epsilon > 0$, can often be replaced by conditions which range over countable sets, such as $\forall \textit{epsilon} = 1/n > 0$.

Conventions on domains of definitions. The *support* of a function $f \in M(\mathbb{R})$ is defined by

$$E = \text{supp}(f) = \{x \in \mathbb{R} : f(x) \neq 0\}.$$

Note: since we are doing measure theory and not topology, we do not form the closure! It may well be the case that $m(E) < \infty$ but E is dense in \mathbb{R} .

Whenever E is a measurable set, there is a natural notion of a measurable function $f : E \rightarrow \mathbb{R}$, exactly as above, and thus we can form the space $M(E)$ and prove analogous theorems here. Any such f can be extended by zero to yield a measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, and thus we have a natural inclusion $M(E) \subset M(\mathbb{R})$. Its image is exactly the space of functions with support *contained* in E .

Thus we can identify $M(E)$ with the space of $f \in M(\mathbb{R})$ supported on E .

One exception is that some results require that E has finite measure. For notational convenience, we sometimes consider the case where $f : [a, b] \rightarrow \mathbb{R}$ is defined on a finite interval, but most of the results for $M([a, b])$ only use the fact that $m([a, b]) < \infty$.

Measure functions defined a.e. It is common, and eventually necessary, to consider two measurable functions as being ‘the same’ if they agree *almost everywhere*, or a.e., which means outside a set of measure zero. Similarly it is usually acceptable for a measurable function to be *undefined* on a set of measure zero; thus we consider $f(x) = 1/x$ as a measurable function, with the value of $f(0)$ undefined or *irrelevant*.

The point of this convention is that when we *integrate* measurable functions, the integral will remain the same if we change the function on a set of measure zero. More on this later.

What do measurable functions look like? Littlewood writes (Lectures on the Theory of Functions, 1944, pp.26–27):

The extent of knowledge required is nothing like so great as is sometimes supposed. There are three principles, roughly expressible in the following terms: Every (measurable) set is nearly a finite sum of intervals; every function (of class L^λ) is nearly continuous; every convergent sequence of functions is nearly uniformly convergent. Most of the results of the present section are fairly intuitive applications of these ideas, and the student armed with them should be equal to most occasions where real variable theory is called for. If one of the principles would be the obvious means to settle a problem if it were “quite” true, it is natural to ask if the “nearly” is near enough, and for a problem that is actually soluble it generally is.

We have already proved Littlewood’s first principle. We now turn to the formulation of Littlewood’s second principle: measurable functions are nearly continuous. A rather strong version of this principle is:

Theorem 4.5 (Lusin’s theorem) *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function. Given $\epsilon > 0$ there exists a continuous function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $g(x) = f(x)$ outside a set of measure ϵ .*

Note: if $|f| \leq M$ then we can always assume $|g| \leq M$ by just cutting g off when it goes outside $[-M, M]$; that is, by replacing g with $\min(M, \max(-M, g))$.

To prove Lusin’s theorem we introduce some types of functions that provide intermediate steps between continuous functions and measurable functions.

Simple functions. We say f is a *simple function* if it is the linear span of the indicator functions of measurable sets. This means f can be written in the form

$$f(x) = \sum_1^n \alpha_i \chi_{E_i}(x),$$

with $\alpha_i \in \mathbb{R}$. We can make this expression canonical (up to the ordering of the terms) by requiring the E_i to be disjoint. Then f only assumes the values 0 and $(\alpha_1, \dots, \alpha_n)$. Conversely, a measurable function with finite range is simple. Since the indicator functions are measurable, so are the simple functions.

The simple functions form a subalgebra of $M(\mathbb{R})$.

Step functions. A special case of a simple function is a *step function*: this just means

$$f(x) = \sum_1^n \alpha_i \chi_{I_i}(x)$$

where the I_i are *intervals* (which can be taken to be disjoint).

Convergence in measure. Finally we introduce a new type of convergence, called convergence in measure. Namely we say $f_n \rightarrow f$ *in measure* if for each $\epsilon > 0$,

$$m\{x : |f_n(x) - f(x)| > \epsilon\} \rightarrow 0.$$

In other words, when n is large, f_n is quite close to f *outside* a set of small measure.

Example. Let $f(x) = 1/x^2$ and let $f_n(x) = \min(n^2, 1/x^2)$. Then $\sup |f_n - f| = \infty$, but actually $f_n = f$ outside the interval $[-1/n, 1/n]$. Thus $f_n \rightarrow f$ in measure.

Theorem 4.6 *Let $f : [a, b] \rightarrow \mathbb{R}$ be a measurable function. Then we can find a sequence $f_n \rightarrow f$ in measure such that the functions f_n are:*

- *Bounded and measurable; or*
- *Simple functions; or*
- *Step functions; or*
- *Continuous functions.*

If $|f| \leq M$ then we can also make $|f_n| \leq M$ for all n .

Proof. (1) Let $f(x)$ be measurable and let $f_n(x)$ be the truncation of f to a function with $|f_n(x)| \leq n$. Then $f_n = f$ outside

$$E_n = \{x : |f(x)| > n\}.$$

Since $\bigcap E_n = \emptyset$ and $E_1 \subset [a, b]$ has finite measure, we have $m(E_n) \rightarrow 0$ by continuity of measure. Thus $f_n \rightarrow f$ in measure.

(2) Because of (1), it now *suffices* to treat the case where f is bounded; say $|f| < M$. Chop the interval $[-M, M]$ into finitely many disjoint subintervals $I_i = [a_i, b_i]$ with $|I_i| < 1/n$, let

$$E_i = \{x : f(x) \in I_i\},$$

and define a simple function by

$$f_n(x) = \sum a_i \chi_{E_i}(x).$$

Then $|f_n - f| < 1/n$, so $f_n \rightarrow f$ in measure (even better, $f_n \rightarrow f$ uniformly).

(3) Because of (2), it now suffices to treat the case where f is a simple function. But a simple function is a finite linear combination of indicator functions, so it suffices to treat the case where $f(x) = \chi_E(x)$. Now by Littlewood's first principle, there exists a finite union of intervals J_n such that $m(E \Delta J_n) < 1/n$. Let $f_n = \chi_{J_n}$. Then

$$m(\{x : f_n(x) \neq f(x)\}) = m(E \Delta J_n) = 1/n \rightarrow 0,$$

so $f_n \rightarrow f$ in measure.

(4) Because of (3), it now suffices to treat the case where f is a step function, and this can be further reduced to the case where $f = \chi_I$ for a single interval $I = [c, d]$. Let $f_n(x)$ be a continuous function with $f_n(x) = 0$ outside $[c - 1/n, d + 1/n]$, $f_n(x) = 1$ on $[c, d]$, and f_n linear on the two small intervals that remain at the ends of I . Then $f_n = f$ outside a set of measure $2/n$, so $f_n \rightarrow f$ in measure. ■

We are now in a position to prove Lusin's theorem.

Proof of Theorem 4.5 (Lusin's theorem). We treat the case where $f : [a, b] \rightarrow \mathbb{R}$ is defined on a finite interval. The passage to the case where f is defined on all of \mathbb{R} is straightforward. We may also assume that f is bounded, say $|f| \leq M$, since f agrees with a bounded function outside a set of small measure.

By the preceding result, we can find a continuous function g_1 and write $f = g_1 + f_1$ where g_1 is continuous and the 'error' f_1 satisfies $|f_1| \leq 1/2$ outside a set E_1 of measure less than $\epsilon/2$. We can also set $f_1 = 0$ on E_1 , then $|f_1| \leq 1/2$ everywhere, but then $f = g_1 + f_1$ only outside E_1 .

Apply the same procedure to f_1 , we can find a continuous function g_2 such that

$$f = g_1 + g_2 + f_2$$

outside $E_1 \cup E_2$, with $m(E_2) \leq \epsilon/4$ and the error now reduced to $|f_2| \leq 1/4$. (Again, we set $f_2 = 0$ outside $E_1 \cup E_2$.) Continuing in this way we obtain

$$f = g_1 + g_2 + \cdots + g_n + f_n$$

with the equality holding outside $\bigcup_1^n E_i$ and $m(E_i) \leq \epsilon/2^i$, with $|f_n| \leq 2^{-n}$ on $[a, b]$, and with $f_n = 0$ on $\bigcup_1^n E_i$.

Since $|f_n| \leq 2^{-n}$ on $[a, b]$, at each stage we can choose the continuous function g_n that

$$\sup |g_n(x)| \leq \sup |f_{n-1}(x)| \leq 1/2^{n-1}.$$

Thus $g = \sum g_i$ converges uniformly to a continuous function, and $f = g$ outside the set $\bigcup_1^\infty E_i$ which has measure $\leq \epsilon$. ■

Warning. A measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$ is *not*, in general, a limit in measure of bounded functions. For example, $f(x) = x$ cannot be such a limit, since for any bounded function g we have $|f - g| > 1$ on a set of infinite measure.

However, Lusin's theorem gives:

Corollary 4.7 *Any measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$ is a limit in measure of continuous functions.*

Here we have leveraged the fact that continuous functions on \mathbb{R} , unlike step functions or simple functions, can be unbounded.

Interlude: pointwise convergence. We now have 2 notions of convergence: pointwise limit and limit in measure. How are they related?

Theorem 4.8 *If $f_n : [a, b] \rightarrow \mathbb{R}$ is a sequence of measurable functions converging to f pointwise, then $f_n \rightarrow f$ in measure.*

Proof. Given $\epsilon > 0$ let E_n be the set where $|f_n - f| > \epsilon$. Then $\bigcap E_n = \emptyset$, so $m(E_n) \rightarrow 0$, which is convergence in measure. ■

Remarks and warnings: the blimp function. Note that the same result holds when $f_n \rightarrow f$ a.e., meaning we have pointwise convergence outside a set of measure zero, and often this slightly weaker notion of convergence is what we need.

As usual, a result like that above cannot hold when the functions are defined on the whole real line. For example, $f_n = \sin(x)\chi_{[-n, n]}$ converges to $f(x) = \sin(x)$ pointwise, but not in measure.

Even on bounded sets, convergence in measure does *not* imply pointwise convergence a.e. For example, let $I_i \subset [0, 1]$ be a sequence of intervals with $m(I_i) \rightarrow 0$, but such that $\bigcup_N^\infty I_i = [0, 1]$ for all N . (In other words, these intervals should cover every point in $[0, 1]$ infinitely many times.) Then

$$f_i = \chi_{I_i} \rightarrow 0$$

in measure, but there is not a single value of x such that $\langle f_i(x) \rangle$ converges!

Nevertheless, we can obtain pointwise convergence by going to a subsequence.

Theorem 4.9 *Let $f_n : \mathbb{R} \rightarrow \mathbb{R}$ be a sequence of measurable functions converging to f in measure. Then there is a subsequence that converges to f pointwise a.e.*

Proof. By passing to a subsequence, we can arrange that $m(E_n) < 2^{-n}$ where E_n is the set where $|f_n - f| > 2^{-n}$. Then $f_n \rightarrow f$ outside the set $E = \limsup E_n$. Since $\sum m(E_n) < \infty$, we have $m(E) = 0$ by the easy Borel–Cantelli lemma. ■

Uniform convergence. Consider functions defined on an interval I . We say $f_n \rightarrow f$ *uniformly* if

$$\sup_I |f_n(x) - f(x)| \rightarrow 0$$

as $n \rightarrow \infty$. Intuitively, this says that the graph of f_n converges to the graph of f .

Uniform convergence is stronger than convergence in measure and pointwise convergence, and it confers many good properties. For example, a uniform limit of continuous functions is continuous, and if $f_n \rightarrow f$ uniformly then

$$\int_a^b f_n \rightarrow \int_a^b f$$

for any $[a, b] \subset I$.

A simple illustration of the difference between pointwise and uniform convergence is provided by the functions on $[0, 1]$ defined by $f_n(x) = n\chi_{(0, 1/n)}$. We have $f_n(x) \rightarrow f(x) = 0$ for all $x \in [0, 1]$, but $\sup |f_n - f| = n$; moreover,

$$\lim \int_0^1 f_n = 1 \neq \int_0^1 f = 0.$$

Littlewood’s third principle: Egoroff’s theorem. Nevertheless, it is not hard to show that a pointwise convergent sequence of *measurable* functions (on a domain of finite measure) is almost uniformly convergent. Here is a precise statement.

Theorem 4.10 (Egoroff) *Let $f_n \rightarrow f$ be a pointwise convergent sequence of measurable functions on $[a, b]$. Then for any $\epsilon > 0$, there is a set $E \subset [a, b]$ of measure less than ϵ such that $f_n \rightarrow f$ uniformly outside E .*

Proof. Consider the functions

$$\delta_n(x) = \sup_{i \geq n} |f_i(x) - f(x)|.$$

Note that $\delta_1 \geq \delta_2 \geq \delta_3 \cdots$, and $\delta_n \rightarrow 0$ pointwise since $f_n \rightarrow f$ pointwise.

But pointwise convergence implies convergence in measure, so we for each k there exists an index n_k such that $\delta_{n_k} < 1/k$ outside a set E_k with measure less than $2^{-k}\epsilon$. Then $\delta_{n_k} \rightarrow 0$ uniformly outside $E = \bigcup E_k$, and $m(E_k) < \epsilon$. But δ_n is a decreasing sequence of functions, so uniform convergence of a subsequence implies uniform convergence of the whole sequence. Finally $\delta_n(x) \geq |f_n(x) - f(x)|$, so $f_n \rightarrow f$ uniformly outside E . ■

Infinite borrowing with negligible debt. We remark on a basic principle has now been illustrated many times. Suppose we want to satisfy an infinite sequence of conditions, but for each one we must pay by excluding a set E_n of positive measure. On the other hand, $m(E_n)$ can be made as small as we like, so long as it is positive. Then we can arrange that our total payment, $m(\bigcup E_n)$, is also as small as we like, by taking $m(E_n) \leq \epsilon/2^n$.

The weak law of large numbers. Here an example showing that the notion of convergence in measure occurs naturally in probability theory.

We consider the binary digits $x = 0.x_1x_2x_3\cdots$ of a randomly chosen point in $[0, 1]$ as a model for an infinite sequence of coin flips. Define $f_i : [0, 1] \rightarrow \{-1, 1\}$ by $f_i(x) = 1$ if $x_i = 1$, and -1 otherwise. Then $\int f_i = 0$ for all i , and more important, $\int f_i f_j = 0$ if $i \neq j$, by independence.

Now let $S_n(x) = (1/n) \sum_1^n f_i(x)$. The *weak law of large numbers* says that there is a high probability that $S_n(x)$ is close to zero; in other words, there is a high probability that approximately half of the digits of x are 1s and half are 0s. More precisely, for any $\epsilon > 0$ we have

$$P(|S_n| > \epsilon) \rightarrow 0$$

as $n \rightarrow \infty$. In other words, we have:

Theorem 4.11 *The functions $S_n(x)$ converge to zero in measure.*

The proof can be based on *Chebyshev's inequality*, a simple estimate that is important in its own right: it says, for $f : \mathbb{R} \rightarrow \mathbb{R}$ a square-integrable function,

$$m(\{x : |f| > \epsilon\}) \leq \frac{1}{\epsilon^2} \int |f|^2.$$

In the case at hand, because of independence, we have

$$\int \left(\sum_1^n f_i \right)^2 = \sum_1^n \int f_i^2 = n.$$

Thus $\int S_n^2 = 1/n$, and so

$$m(\{x : |S_n(x)| > \epsilon\}) \leq \frac{1}{n\epsilon^2} \rightarrow 0$$

as $n \rightarrow \infty$.

The strong law of large numbers. The *strong law of large numbers*, which also holds here, states that $S_n(x) \rightarrow 0$ for almost every x . Its proof is more subtle. It would follow from the previous argument if $\sum 1/n$ were finite. The sum is just barely infinite, so it is not a big leap to expect the strong law to hold.

In fact, if we pass to the subsequence $n(k) = [k^{1+\alpha}]$ with $0 < \alpha \ll 1$, then we get $\sum_k 1/n(k) \asymp \sum_k 1/k^{1+\alpha} < \infty$, so we can assert $S_{n(k)}(x) \rightarrow 0$ almost sure as $k \rightarrow \infty$. (This is called *sparsification*.) On the other hand, it is not hard to show that for $n(k) < i < n(k+1)$, the values of $S_i(x)$ are all almost the same when k is large. So convergence along the subsequence $n(k)$ implies convergence of the whole sequence, which is the strong law of large numbers.

5 Integration

In this section we introduce the Lebesgue integral of a measurable function on \mathbb{R} , written simply $\int f$.

Quick start. It is possible to sum up the definition and main properties of the Lebesgue integral fairly quickly. First, the integral will only be defined for certain measurable functions. Bending notation slightly, we say f is *positive* if $f \geq 0$. For a *positive* function, we define

$$\int f = \sup \left\{ \sum a_i m(E_i) : 0 \leq \sum_1^n a_i \chi_{E_i} \leq f \right\}.$$

Note that this sup can be infinite. We say f is *integrable* if $\int |f| < \infty$. Finally we define, for a general integrable function,

$$\int f = \int f_+ - \int f_-,$$

where we have written f as the difference of two positive functions — its positive part, and its negative part, defined by $f_+ = \max(0, f)$ and $f_- = \max(0, -f)$.

Here are some of the basic properties of the integral:

1. For any measurable set E , $\int \chi_E = m(E)$.
2. The integral is linear on the vector space of integrable functions; in particular,

$$\int f + g = \int f + \int g.$$

3. The integral is *monotone*: if $f \leq g$ then $\int f \leq \int g$.
4. The integral is countably additive for positive functions. That is, if $f_n \geq 0$ for all n , then

$$\int \sum f_n = \sum \int f_n.$$

This is called the *monotone convergence theorem*.

Integration over a region. It is often useful to integrate f along an interval, say from a to b , or more generally over a measurable set E . Extending by zero if necessary, we can always assume f itself is defined on the whole real line. We then define

$$\int_E f = \int f\chi_E \quad \text{and} \quad \int_a^b f = \int_{[a,b]} f.$$

From the properties above we have the useful estimates

$$\left| \int f \right| \leq \int |f|$$

and

$$\int_E |f| \leq m(E) \cdot \sup_E |f(x)|.$$

Alternating sums. We remark that a Lebesgue integrable function is like an *absolutely convergent* series. Although the sum $1 - 1/2 + 1/3 - 1/4$ converges, the corresponding function

$$f(x) = \sum_1^\infty \frac{(-1)^{n+1}}{n} \chi_{[n,n+1)}(x)$$

is *not* Lebesgue integrable. Similarly, $\sin(x)/x$ is not Lebesgue integrable. We can still ‘find’ the values of these integrals, but an explicit approximation scheme must be chosen, such as taking the limit of $\int_{-n}^n f$ as $n \rightarrow \infty$.

The problem with this quick start is that it is difficult to establish the main properties of the integral — especially linearity — directly from the definition above. Thus we will proceed by another route, where we gradually enlarge the domain of definition of the integral. This route also reveals more intuitively how Lebesgue integration works.

Integration of simple functions. For a simple function ϕ supported on a set of finite measure, we would like to define

$$\int \phi = \int \sum a_i \chi_{E_i} = \sum a_i m(E_i).$$

The problem is, a simple function can be written as a sum in many different ways, so the integral above is potentially ill-defined. To skirt this issue, recall that a simple function f has a *canonical* representation in the form above such that (i) the a_i are distinct and (ii) the E_i are disjoint. We then require that the canonical representation is used.

Example: $\int \chi_{\mathbb{Q}} = 0$.

Theorem 5.1 *Integration is linear on the vector space of simple functions.*

Proof. Clearly $\int a\phi = a \int \phi$. We must prove $\int \phi + \psi = \int \phi + \int \psi$.

First note that for any representation of ϕ as $\sum b_i \chi_{F_i}$ with the sets F_i disjoint, we have $\int \phi = \sum b_i m(F_i)$. Indeed,

$$\int \sum b_i \chi_{F_i} = \int \sum a_j \chi_{\bigcup_{b_i=a_j} F_i} = \sum a_j \sum_{b_i=a_j} m(F_i) = \sum b_i m(F_i).$$

Now take the finite collection of sets F_i on which ϕ and ψ are *both* constant, and write $\phi = \sum a_i \chi_{F_i}$ and $\psi = \sum b_i \chi_{F_i}$. Then

$$\int \phi + \psi = \sum (a_i + b_i) m(F_i) = \int \phi + \int \psi.$$

■

Integration of bounded functions on bounded sets. Now let $f : [a, b] \rightarrow \mathbb{R}$ be an *arbitrary* bounded function supported on a bounded interval. We can assume $|f| \leq M$. We define the *Lebesgue integral* of f by

$$\int_a^b f = \inf_{\psi \geq f} \int_a^b \psi = \sup_{f \geq \phi} \int_a^b \phi,$$

assuming sup and inf agree. Here ϕ and ψ range over all *simple functions* on $[a, b]$.

Theorem 5.2 *The two definitions of the integral of f above agree iff f is a measurable function.*

Proof. Suppose f is measurable. Since $\int \psi \geq \int \phi$, we just need to show the simple functions ϕ and ψ can be chosen such that their integrals are arbitrarily close. To this end, cut the interval $[-M, M]$ into N pieces $[a_i, a_{i+1})$ of length less than ϵ . Let E_i be the set on which $f(x)$ lies in $[a_i, a_{i+1})$. Then $\phi = \sum a_i \chi_{E_i}$ and $\psi = \sum a_{i+1} \chi_{E_i}$ satisfying $\phi \leq f \leq \psi$ and $\int(\psi - \phi) \leq \epsilon m(E)$, so we are done.

Conversely, if the sup and inf agree, then we can choose simple functions $\phi_n \leq f \leq \psi_n$ such that $\int(\psi_n - \phi_n) \rightarrow 0$. Let $\phi = \sup \phi_n$ and $\psi = \inf \psi_n$. Then ϕ and ψ are measurable, and $\phi \leq f \leq \psi$.

We claim $\phi = \psi$ a.e. (and thus f is measurable). Otherwise, there is a set of positive measure A and an $\epsilon > 0$ such that $\psi - \phi > \epsilon$ on A . But then $\epsilon \chi_A \leq \psi_n - \phi_n$ for all n , and thus $\int \psi_n - \phi_n \geq \epsilon m(A) > 0$. ■

Vertical versus horizontal. The method of approximation by simple function is the first cut the *range* of f into small intervals, then consider their preimages to decompose the domain. In Riemann integration one first cuts the domain into small pieces. Thus it is sometimes said that Riemann integration is based on vertical cuts, while Lebesgue integration is based on horizontal cuts.

Theorem 5.3 *Integration is linear on the space of bounded measurable functions on $[a, b]$.*

Proof. Clearly $\int \alpha f = \alpha \int f$. Using the sup definition of the integral gives

$$\int f + \int g = \sup_{\phi \leq f, \psi \leq g} \int \phi + \psi \leq \sup_{\alpha \leq f+g} \int \alpha = \int f + g.$$

From the inf definition, we obtain the reverse inequality. ■

Theorem 5.4 *Let f be a bounded function on an interval $[a, b]$, and suppose f is Riemann integrable. Then f is also Lebesgue integrable, and the two integrals agree.*

Proof. If f is Riemann integrable then there are *step functions* $\phi_n \leq f \leq \psi_n$ with $\int(\psi_n - \phi_n) \rightarrow 0$. Since step functions are special cases of simple functions, we see f is Lebesgue integrable. ■

Positive functions. For $f \geq 0$ we define $\int f = \sup_{0 \leq g \leq f} \int g$, where g ranges over bounded functions supported on sets of finite measure. Clearly this is the same as saying $\int f = \lim = \int f_M$, where $f_M = \min(f, M) \cdot \chi_{[-M, M]}$.

Theorem 5.5 *If $f_1, f_2 \geq 0$ are measurable functions, then $\int f_1 + \int f_2 = \int f$.*

Proof. From the definition it is immediate, as in the case of bounded functions, that $\int f_1 + \int f_2 \leq \int f$. To get the reverse inequality, we will split up g with $0 \leq g \leq f$ into $g = g_1 + g_2$ with $g_1 \leq f_1$ and $g_2 \leq f_2$. Namely we take $g_1 = \min(g, f_1)$ and $g_2 = g - g_1$.

We claim $g_2 \leq f_2$. Indeed, if $g_1(x) = g(x)$ then $g_2(x) = 0 \leq f_2(x)$, while if $g_1(x) = f_1(x)$ then $g_2(x) = g(x) - f_1(x) \leq f(x) - f_1(x) = f_2(x)$.

It follows that

$$\int g = \int g_1 + \int g_2 \leq \int f_1 + \int f_2.$$

Now taking the supremum over g gives

$$\int f \leq \int f_1 + \int f_2,$$

which is the desired reverse inequality. ■

The general Lebesgue integral. For general f , we require that $\int |f| < \infty$ before $\int f$ is defined. Then writing $f = f_+ - f_-$, we define $\int f = \int f_+ - \int f_-$.

Theorem 5.6 *The general Lebesgue integral is linear.*

Proof. Note that

$$(f + g)_+ - (f + g)_- = f + g = (f_+ - f_-) + (g_+ - g_-).$$

Rearranging terms so only positive functions appear, we get:

$$(f + g)_+ + f_- + g_- = (f + g)_- + f_+ + g_+.$$

By linearity for positive functions, we get

$$\int (f + g)_+ + \int f_- + \int g_- = \int (f + g)_- + \int f_+ + \int g_+.$$

Putting the terms back in their original order, this gives

$$\int f + g = \int (f + g)_+ - \int (f + g)_- = \int f_+ - \int f_- + \int g_+ - \int g_- = \int f + \int g.$$

■

Linearity, sup and inf. Note that the proofs of linearity for simple, bounded, positive and general functions are all different! It is worthwhile to think through the bases of these proofs, which can roughly be summarized as follows:

- Simple* — common refinement
- Bounded* — agreement of sup and inf
- Positive* — separation of integrand into 2 pieces
- General* — rearrangement.

Basic properties of integrals. It is easy to check that the integral just defined satisfies some natural properties, e.g.

$$\int_{A \sqcup B} f = \int_A f + \int_B f,$$

and $f \leq g$ implies $\int f \leq \int g$; in particular,

$$\left| \int_E f \right| \leq \int_E |f| \leq m(E) \cdot \sup_E |f(x)|.$$

Integrals and limits. Suppose $f_n \rightarrow f$ pointwise. When can we conclude that

$$\int f_n \rightarrow \int f?$$

There are 2 phenomena to be aware of: mass can escape to infinity ‘vertically’ or horizontally. For example, if $f_n(x) = n\chi_{(0,1/n)}$ then mass escapes vertically: we have $\int f_n = 1$ for all n but $f = 0$. If $f_n(x) = \chi_{[n,n+1]}$ then again $\int f_n = 1$ but $f = 0$; mass escapes horizontally.

Our first result says that if we prevent both types of escape, then in fact the limit of the integral is the integral of the limit. We say a sequence of functions is *uniformly bounded* if $\sup_n \sup_x |f_n(x)| = M < \infty$.

Theorem 5.7 (Bounded convergence) *Let f_n be a sequence of uniformly bounded measurable functions supported on $[a, b]$. Suppose $f_n \rightarrow f$ pointwise. Then $\int f_n \rightarrow \int f$.*

Proof. We will use Littlewood’s third principle. Outside a set A with $m(A) < \epsilon$, the convergence is uniform. Thus, letting $B = [a, b] - A$, we have

$$\left| \int_B f_n - f \right| \leq (b - a) \sup_B |f_n - f| \rightarrow 0.$$

On the other hand, outside B we still have $|f_n|, |f| \leq M$ and therefore

$$\limsup \left| \int f_n - f \right| = \limsup \left| \int_A f_n - f \right| \leq 2M \cdot m(A) \leq \epsilon m(A).$$

Since $\epsilon > 0$ was arbitrary, this shows $\int f_n \rightarrow \int f$. ■

Note: this result is a special case of the *dominated convergence theorem*, which we establish below.

Positive functions. We now focus on the case where $f_n \geq 0$ for all n , and $f_n \rightarrow f$ pointwise on \mathbb{R} . The first result says that mass can only escape, it cannot appear out of the blue.

Theorem 5.8 (Fatou's lemma) *If $f_n \geq 0$ is a sequence of measurable functions, and $f_n \rightarrow f$ pointwise, then*

$$\int f \leq \liminf \int f_n.$$

Proof. Consider any bounded measurable function g with bounded support such that $0 \leq g \leq f$. Let $g_n = \min(g, f_n)$. Then $g_n \rightarrow g$ as $n \rightarrow \infty$, so by bounded convergence we have

$$\int g_n \rightarrow \int g.$$

Since $\int g_n \leq \int f_n$, this gives

$$\int g \leq \liminf \int f_n.$$

Taking the supremum over all such $g \leq f$ gives

$$\int f \leq \liminf \int f_n.$$

■

Theorem 5.9 (Monotone convergence) *Let $0 \leq f_1 \leq f_2 \leq \dots$ be a monotone increasing sequence of positive functions, and let $f(x) = \lim f_n(x)$. Then*

$$\int f = \lim \int f_n.$$

Proof. Since $f_n \leq f$ we have $\limsup \int f_n \leq \int f$, and Fatou's Lemma gives the reverse inequality. ■

The space $L^1(\mathbb{R})$. As an application of monotone convergence, we will now prove a density result for the space of integrable functions $L^1(\mathbb{R})$.

Let us begin by introducing this space. If f is an integrable function on \mathbb{R} , a natural way to measure its size is by the norm

$$\|f\| = \int |f|.$$

This norm satisfies the expected properties, e.g.

$$\|f + g\| \leq \|f\| + \|g\|.$$

Given 2 integrable functions, we measure their distance by $d(f, g) = \|f - g\|$. Now note that $d(f, g) = 0$ does not quite imply that $f = g$; it only implies that these functions agree *almost everywhere*.

Thus we let $L^1(\mathbb{R})$ denote the normed vector space whose elements are *equivalence classes* of integrable functions, with $f \sim g$ iff $f = g$ a.e. (Usually we ignore this nuance.) Since $L^1(\mathbb{R})$ is a metric space, it makes sense to talk about dense sets of functions.

Theorem 5.10 *The following classes of functions are dense in $L^1(\mathbb{R})$:*

- *Bounded measurable functions with bounded support;*
- *Simple functions;*
- *Step functions; and*
- *Continuous functions with compact support.*

Proof. Given $f \in L^1(\mathbb{R})$, let f_M be the truncation of f to a function with $|f_M| \leq M$ supported on $[-M, M]$. Then $f_M \rightarrow f$ pointwise, and $|f_M| \rightarrow |f|$ monotonically, as $M \rightarrow \infty$. By monotone convergence, we have

$$\int |f_M| \rightarrow \int |f|.$$

Since f_M and f have the same sign, this implies

$$\|f - f_M\| = \int |f - f_M| = \int |f| - |f_M| \rightarrow 0.$$

Thus bounded functions with bounded support are dense.

To show simple functions are dense, it now suffices to show that they approximate any bounded function $f \in L^1(\mathbb{R})$ with bounded support. As

we have seen already, there are simple functions with the same bounds as f such that $\phi_n \rightarrow f$ in measure. Passing to a subsequence, we get pointwise convergence. Then $|f - \phi_n| \rightarrow 0$; and by the bounded convergence theorem, this implies $\|f - \phi_n\| \rightarrow 0$.

The same reasoning applies to step functions and continuous functions. ■

Dominated convergence. We now turn to the dominated convergence theorem. It begins with following ‘continuity’ property of integrable functions.

Theorem 5.11 (Modulus of integrability) *Let $f \geq 0$ be integrable. Then for any $\epsilon > 0$ there is a $\delta > 0$ such that $m(E) < \delta \implies \int_E f < \epsilon$.*

Corollary 5.12 *The function $F(t) = \int_{-\infty}^t f(x) dx$ is uniformly continuous on \mathbb{R} .*

Proof of the Theorem. As we have seen, the truncation f_M of f satisfies $\int |f - f_M| < \epsilon/2$ for M sufficiently large. Then for $m(E) < \delta = \epsilon/(2M)$, we have $\int_E f \leq \int_E (f - f_M) + Mm(E) \leq \epsilon$. ■

Dominated convergence. Let $f_n \rightarrow f$, with $|f_n|, |f| \leq g$ and $\int g < \infty$. Then $\int f_n \rightarrow \int f$.

Proof. Given $\epsilon > 0$ there is a $\delta > 0$ such that $\int_A g < \epsilon$ whenever $m(A) < \delta$. We can also choose M such that $\int_E g < \epsilon$ outside $[-M, M]$. Then by Littlewood’s third principle, there is a set $A \subset [-M, M]$ with $m(A) < \delta$ outside of which $f_n \rightarrow f$ uniformly. Thus

$$\limsup \left| \int f_n - f \right| \leq 2 \left(\int_{\mathbb{R} - [-M, M]} g + \int_A g \right) \leq 4\epsilon.$$

Since ϵ was arbitrary, $\int f_n \rightarrow \int f$. ■

Convergence in measure. All the theorems about interchanging limits and integration above also holds with *pointwise convergence* replaced by *convergence in measure*. This can be proved using the fact that if $f_n \rightarrow f$ in measure, then a subsequence converges pointwise.

6 Differentiation and Integration

We are now in a position to confront two of the pillars of calculus: the statements that

$$\int f' = f$$

and that

$$\frac{d}{dx} \int f = f.$$

Here is a summary of what we will find.

1. If f is *monotone*, or more generally if f has *bounded variation*, then $f'(x)$ exists a.e., and $\int_a^b |f'| < \infty$.
2. However, this is not enough to insure $f = \int f'$, as the Cantor function shows: it has bounded variation, $f'(x) = 0$ a.e. but f is not constant.
3. Instead, we find that if f has the stronger property of *absolutely continuity*, then $f = \int f'$.
4. Finally, if f is an integrable then $F = \int f$ is always absolutely continuous, so in fact $(d/dx) \int f = f$ holds (a.e.) whenever it makes sense (whenever f is integrable).

Applying the last result to the indicator function of a set E of finite measure, we obtain the *Lebesgue density theorem*: for almost every $x \in E$, we have

$$\lim_{r \rightarrow 0} \frac{m(E \cap B(x, r))}{m(B(x, r))} = 1.$$

The same limit is zero for a.e. $x \notin E$. This enhances our intuition about measurable sets.

We note that functions of bounded variation are of independent importance, as they correspond to the distribution functions of *signed Borel measures*, as we will sketch below.

This chapter will be in two sections. The first will treat the *existence* of derivatives, while the second will treat the *integration* of derivatives and the *differentiation* of the integral.

I. Differentiation

Functions that are differentiable everywhere. We say f is *differentiable* if $f'(x)$ exists for all x .

Even if $f'(x)$ exists everywhere, it does not have to be continuous. For example, if $|f(x)| \leq x^2$, then no matter how badly $f'(x)$ oscillates near $x = 0$, we have $f'(0) = 0$. It is difficult to describe all the functions that can arise as $f'(x)$.

A nowhere differentiable function. Next we show:

Theorem 6.1 *There exists a continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ that is nowhere differentiable.*

Proof. Let $f(x) = \sum_1^\infty a_n \sin(b_n x)$, where $\sum a_n$ converges quickly but $a_n b_n \rightarrow \infty$ rapidly. For concreteness, one might take $a_n = 10^{-n}$, $b_n = 10^{6n}$; then $a_{n-1} b_{n-1} / (a_n b_n) = 1/10^5$.

Now for any n and x , we can choose $\Delta x \approx 1/b_n$ such that $\Delta a_n \sin(b_n x) \asymp a_n$. (Here $\Delta h = h(x + \Delta x) - h(x)$.) For $k < n$, we have

$$\sum \Delta a_k \sin(b_k x) \leq \sum a_k b_k / b_n \asymp a_{n-1} b_{n-1} / b_n < a_{n-1} \ll a_n,$$

and for $k > n$ we have

$$\Delta a_k \sin(b_k x) \leq a_k \ll a_n$$

Thus $\Delta f / \Delta x \asymp a_n b_n \rightarrow \infty$, so $f'(x)$ does not exist. ■

The graph of f is easy to visualize to convince yourself that $f'(x)$ exists nowhere and that this is really no surprise. At any given scale, $y = f(x)$ looks like a sine wave. But on closer inspection, the graph has rapid wiggles at a much smaller scale. This behavior persists at every scale, so there is no reasonable value of the slope of the graph at x .

Riemann's 'example'. Riemann thought that the function

$$f(x) = \sum \exp(2\pi i n^2 x) / n^2$$

was nowhere differentiable. This is almost true, however it turns out that $f'(x)$ actually does exist at certain rational points.

This function has the advantage over the example above in that it comes up 'naturally' in mathematics; it defines an analytic function $f(z)$ on the upper half-plane, which is related to theta functions and automorphic forms.

Monotone functions. We say $f : [a, b] \rightarrow \mathbb{R}$ is *increasing* if $x \leq y \implies f(x) \leq f(y)$. If f or $-f$ is increasing then f is *monotone*.

Example: write $\mathbb{Q} = \{q_1, q_2, \dots\}$ and set

$$f(x) = \sum_{q_i < x} 2^{-i}$$

Then $f : \mathbb{R} \rightarrow \mathbb{R}$ is monotone increasing, and f has a dense set of points of discontinuity. One of the deepest theorems we will cover in these notes is:

Theorem 6.2 (Monotone differentiability) *A monotone function $f : [a, b] \rightarrow \mathbb{R}$ is differentiable almost everywhere.*

Thus the oscillations of the preceding example are *necessary* to produce nowhere differentiability. Gleason has remarked that this property of monotone functions helped lead him to his proof of Hilbert's 5th problem (which topological groups are Lie groups?).

Covering theorems. The proof of differentiability will use some covering lemmas which are of importance in their own right. In fact covering lemmas are pervasive in real analysis.

We will formulate these results for metric spaces. Note that if $B = B(x, r)$ in \mathbb{R}^n , then the *set* B determines x (its center) and r (its radius) uniquely. Not so in a general metric space. Thus a ball in a metric space should really be thought of as the *data* (x, r) , which determines the set $B(x, r)$. For example, x might be an isolated point, in which case $B(x, r) = \{x\}$ for all r sufficiently small.

The key idea comes from the following result. Given a ball $B = B(x, r)$, we let $3B = B(x, 3r)$.

Lemma 6.3 *Let K be a compact metric space covered by a collection of balls \mathcal{B} . Then we can find disjoint balls D_1, \dots, D_n within this covering, such that $K \subset \bigcup_1^n 3D_i$.*

Proof. We use the *greedy algorithm*. By compactness we can assume that \mathcal{B} consists of a finite collection of balls $B(x_i, r_i)$ with radii in decreasing order, $r_1 \geq r_2 \geq \dots \geq r_N$. Let $D_1 = B(x_1, r_1)$ and let D_{n+1} be the largest ball on the list that is disjoint from those already chosen. This procedure leads to a finite collection of balls D_1, \dots, D_n .

Now suppose $x \notin \bigcup_1^n D_j$. Then $x \in B(x_j, r_j)$ for some ball which was not chosen by the greedy algorithm. This means $B(x_j, r_j)$ must meet one of the balls $B(x_i, r_i)$ which *was* chosen. But then $r_i > r_j$, and since these balls meet we have $x \in B(x_j, r_j) \subset B(x_i, 3r_i)$. ■

Vitali coverings. This idea has many variations. A simple modification of the argument gives an infinite sequence of disjoint balls with a similar property.

To be more precise, we say a collection of balls \mathcal{B} gives a *Vitali covering* of a metric space X if for all $x \in X$ and $r > 0$, there is a ball $B(y, s)$ in the cover with $x \in B(y, s)$ and $s < r$. In other words, \mathcal{B} provides a small ball around every point of X ; equivalently, \mathcal{B} provides a *basis* for the metric topology on X .

Theorem 6.4 *Let \mathcal{B} be a Vitali covering of a compact set K . Then we can find a sequence of disjoint balls $D_i \in \mathcal{B}$ such that*

$$K \subset \bigcup_1^N \overline{D_i} \cup \bigcup_{N+1}^{\infty} 3D_i$$

for all $N > 0$.

Proof. By compactness we can assume \mathcal{B} consists of an sequence of balls of decreasing size, $B(x_i, r_i)$. The same greedy algorithm then produces either a finite set of disjoint balls which cover K — and then the result is trivial — or an infinite sequence of disjoint balls (D_i) . Assume we are in the latter case, and fix $N > 0$, and consider a point $x \notin \bigcup_1^N \overline{D_i}$. By the Vitali property, we have $x \in B(y_i, r_i) \in \mathcal{B}$ with r_i smaller than the radius of D_N . Thus $B(y_i, r_i)$ was in the running for being chosen after D_N . Since it wasn't, the ball $B(y_i, r_i)$ must meet a ball $B(y_j, r_j)$ which *was* chosen, with $j < i$. Then $B(y_j, r_j) = D_k$ for some $k > N$, and $x \in 3D_k$ by the same reasoning as above. ■

The Vitali lemma. We can now prove a very sharp form of Littlewood's first principle. We will continue to refer to coverings by balls, but of course since we are now working on the real line, each B_i is actually an interval.

Theorem 6.5 (Vitali's Lemma) *Let $E \subset \mathbb{R}$ be a set of finite, positive measure, and let \mathcal{B} be a Vitali covering of E . Then for any $\epsilon > 0$, there exists a finite collection of disjoint balls $D_i \in \mathcal{B}$ such that*

$$m(E \Delta \bigcup_1^N D_i) < \epsilon.$$

Proof. For any $\delta > 0$ we can choose a compact set K and an open set U such that $K \subset E \subset U$, and $m(U) - m(K) < \delta$. Then \mathcal{B} provides a Vitali

covering of K , and we can assume the balls in \mathcal{B} all lie in U . Now apply the preceding theorem to get a collection of disjoint disks $D_i \in \mathcal{B}$ such that

$$K \subset \bigcup_1^N \overline{D_i} \cup \bigcup_{N+1}^{\infty} 3D_i$$

for all $N > 0$. Since the D_i are disjoint and lie in U , we have $\sum m(U_i) < m(U) < \infty$. In particular, we can choose N large enough that $3 \sum_{N+1}^{\infty} m(D_i) < \delta$. We then get

$$m(U) - \delta \leq m(K) \leq \sum_1^N m(D_i) + \delta.$$

Since $\bigcup D_i \subset U$, we get $m(U \Delta \bigcup_1^N D_i) < 2\delta$. We already have $m(E \Delta U) < \delta$, and thus $m(E \Delta \bigcup_1^N D_i) < 3\delta$. Choosing $\delta = \epsilon/3$ gives the Theorem. \blacksquare

Application of coverings to density. To illustrate the use of Vitali coverings, we can now prove:

Theorem 6.6 (Lebesgue Density) *Let $E \subset \mathbb{R}$ be a measurable set. Then for almost every $x \in \mathbb{R}$,*

$$\lim_{r \rightarrow 0} \frac{m(E \cap B(x, r))}{m(B(x, r))} = \begin{cases} 1 & \text{if } x \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. It suffices to treat the case where $E \subset [a, b]$, and to show the limit above tends to 1 for almost every $x \in E$. (Applying the same reasoning to $[a, b] - E$ shows the limit tends to 0 elsewhere.)

Fix $n > 0$ and let $A \subset E$ denote the set where the limsup of the density of E in $B(x, r)$ is strictly less than $1 - 1/n$. Then for every $x \in A$ we can find arbitrarily small intervals I around x with the density of E in I less than $1 - 1/n$. Since A is contained in E , its density in A is also less than $1 - 1/n$. By the Vitali lemma, given $\epsilon > 0$ there is a disjoint set of such intervals such that $m(A \Delta \bigcup_1^n I_i) \leq \epsilon$. This gives

$$m(A) \leq \epsilon + \sum_1^n m(A \cap I_i) \leq \epsilon + (1 - 1/n)m(\bigcup_1^n I_i) \leq \epsilon + (1 - 1/n)(m(A) + \epsilon).$$

Letting $\epsilon \rightarrow 0$ we get $m(A) \leq (1 - 1/n)m(A)$ and hence $m(A) = 0$. Since n was arbitrary, this proves the density of E in $B(x, r)$ tends to one as $r \rightarrow 0$ for almost every $x \in E$. \blacksquare

A choice function. Suppose we declare two measurable sets to be equivalent if $m(E \Delta E') = 0$. Can we choose, concretely, a canonical set from each equivalence class? The answer is *yes*: we can just take the set of points of density 1. This set is the same for E and E' .

Ergodic theory on the circle. The Lebesgue density theorem has many basic applications in ergodic theory. Here is an example.

Theorem 6.7 *Let $\theta \in \mathbb{R} - \mathbb{Q}$ be an irrational number, and define $f : S^1 \rightarrow S^1$ by $f(x) = x + \theta \pmod{1}$. Then f is ergodic: if $E \subset S^1$ has positive measure, and $f(E) \subset E$, then $m(E) = 1$.*

Proof. Let $\delta_r(x) = m(E \cap B(x, r)) / m(B(x, r))$. Since E has positive measure, it has a point of density x_0 . This means $\delta_r(x_0) \rightarrow 1$ as $r \rightarrow 0$.

On the other hand, $\delta_r(x)$ is a continuous function of x , and (by invariance of E) it is constant on the orbits of f , which are dense. Thus $\delta_r(x)$ does not depend on x ! So once we know $\delta_r(x_0) \rightarrow 1$ as $r \rightarrow 0$, we know $\delta_r(x) \rightarrow 1$ for all $x \in S^1$. But $\delta_r \rightarrow \chi_E$ a.e., by the Lebesgue density theorem, so $m(E) = m(S^1) = 1$. ■

Corollary 6.8 *Any measurable function $h : S^1 \rightarrow \mathbb{R}$, invariant under the irrational rotational f , is constant a.e.*

Proof. For any partition of \mathbb{R} into disjoint intervals I_i of length ϵ , we have $m(f^{-1}(I_i)) = 1$ for exactly one i . As $\epsilon \rightarrow 0$, this distinguished interval shrinks down to the constant value assumed by f . ■

Monotone functions. We now return to the study of monotone functions, and complete the proof that $f'(x)$ exists a.e. It is similar in spirit to the proof of Lebesgue density.

Proof of Theorem 6.2 (Monotone differentiability). We may assume $f : [a, b] \rightarrow \mathbb{R}$ is monotone increasing. Let $|A| = \sup A - \inf A$ denote the length of the smallest interval containing $A \subset \mathbb{R}$. On any open interval $I \subset [a, b]$ we have an approximate value of f' which we denote by

$$D(I) = \frac{|f(I)|}{|I|}.$$

Let $E \subset [a, b]$ denote the set of points where $f'(x)$ does not exist. For convenience we exclude from E the countable set of points of discontinuity

of f , and the endpoints of $[a, b]$. We also exclude from E the set S where $f'(x) = +\infty$. It is easy to show that S has measure zero, using the Vitali covering lemma. (Similarly, the set where $f'(x) > M$ has measure at most $(b - a)/M$.)

Now if $x \in E$, then the values of $D(I)$ must certainly oscillate on arbitrarily small intervals with $x \in I$. Thus we can write $E = \bigcup_{rs} E_{rs}$, where the union is over all rationals with $r < s$, where E_{rs} consists of the points $x \in [a, b]$ such that there are arbitrarily small open intervals I and J containing x with

$$D(I) < r < s < D(J).$$

We will show that for fixed $r < s$, we have $m(E_{rs}) = 0$. For brevity of notation, let $A = E_{rs}$.

The idea of the proof is that along A , f' behaves as if its derivative is both less than r and bigger than s . Thus

$$sm(A) \leq m(f(A)) \leq rm(A),$$

which is a contradiction unless $m(A) = 0$.

By Vitali's lemma, given $\epsilon > 0$ we can find a disjoint intervals I_1, \dots, I_n such that $D(I_i) < r$ and $m(A \Delta \bigcup I_i) < \epsilon$. In particular, we have

$$m\left(\bigcup I_i\right) = \sum |I_i| < m(A) + \epsilon.$$

Applying Vitali's lemma to $A \cap \bigcup I_i$, we can find disjoint intervals J_1, \dots, J_m contained in $\bigcup I_i$ with $D(J_i) > s$ and

$$m\left(\bigcup J_i\right) = \sum |J_i| > m(A) - 2\epsilon.$$

But by monotonicity, we have

$$s \sum |J_i| \leq \sum |f(J_i)| \leq \sum |f(I_i)| \leq r \sum |I_i|$$

Since both $\sum |J_i|$ and $\sum |I_i|$ differ from $m(A)$ by a small amount, this gives

$$\frac{m(A) - 2\epsilon}{m(A) + \epsilon} \leq \frac{\sum |J_i|}{\sum |I_i|} \leq \frac{r}{s} < 1.$$

Letting $\epsilon \rightarrow 0$, we obtain a contradiction unless $m(A) = 0$. Thus E is a countable union of sets E_{rs} of measure zero, so $m(E) = 0$ and $f'(x)$ exists a.e. ■

Note: we have excluded points of discontinuity for the following reason: if $I = [x, x + t]$ is a *closed* interval with $D(I) > s$, and f is continuous at x , then there is also a slightly larger *open* interval J with $x \in J$ and $D(J) > s$; and similarly if $D(I) < r$.

Borel measures: the secret life of monotone functions. Here is an aside that provides more insight into the meaning of a general monotone function.

Let \mathcal{B} denote the Borel subsets of $[a, b]$. A finite *Borel measure* on $[a, b]$ is a map $\mu : \mathcal{B} \rightarrow [0, \infty)$ satisfying countable additivity.

Examples.

(1) The measure δ is given by $\delta(B) = 1$ if $0 \in B$, and $\delta(B) = 0$ otherwise. More generally, if $a_i > 0$, $\sum a_i < \infty$ and $x_i \in [a, b]$ are given, we can form the *atomic measure*

$$\mu = \sum a_i \delta(x - x_i),$$

which is characterized by

$$\mu(B) = \sum_{x_i \in B} a_i.$$

(2) An integrable function $f \geq 0$ on $[a, b]$ defines a measure by

$$\mu(B) = \int_B f.$$

Such a measure is said to be *absolutely continuous*; it has the property that $\mu(B) = 0$ whenever $m(B) = 0$.

(3) We have a measure on the Cantor set $K \subset [0, 1]$ obtained by choosing the digits 0 and 2 with equal probability to form the ternary expansion of $x \in K$. This measure has no atoms. It is characterized by

$$\mu(B) = m(f(B)),$$

where f is the Cantor function. Since K has measure zero, this measure is also not of type (2) above. A measure supported on a set of measure zero is said to be *singular*.

Every measure is a canonical sum of these three types:

$$\mu = (\text{atomic measure}) + (f(x) dx) + (\text{nonatomic singular measure}).$$

Relation to monotone functions. From μ we obtain a monotone increasing function

$$f(x) = \mu([a, x]).$$

Since these intervals generate \mathcal{B} , we can recover f from μ .

Then:

(1) The atomic part of μ comes from the jumps in f (its discontinuities). (2) The absolutely continuous part of μ comes from $f'(x) dx$. (3) The remainder of μ is a singular measure with no atoms. It corresponds to a continuous monotone function with $f'(x) = 0$ a.e.

Integrating the derivative of a monotone function. Now the total mass of the measure μ is $f(b) - f(a)$, while the integral of $f'(x)$ only picks up the mass of its absolutely continuous part, so it is no surprise that it should be bounded above by $f(b) - f(a)$. Let us prove this formally.

Theorem 6.9 *If $f : [a, b] \rightarrow \mathbb{R}$ is monotone increasing, then $\int_a^b f'(x) dx \leq f(b) - f(a)$.*

Proof. Define $f_n(x) = n(f(x + 1/n) - f(x)) \geq 0$. (For convenience we define $f(x) = f(b)$ for $x > b$.) Then $f_n(x) \rightarrow f'(x)$, so by Fatou's lemma we have $\int f' \leq \liminf f_n$. But $\int f_n$ is, for n large, the difference between the averages of f over two disjoint intervals, so it is less than or equal to the maximum variation $f(b) - f(a)$. ■

Since $f' \geq 0$ when f is increasing, $\int f' = \int |f'|$. Thus we can conclude:

Corollary 6.10 *If $f : [a, b] \rightarrow \mathbb{R}$ is monotone then $f' \in L^1([a, b])$.*

Distributions. Naturally if we insist on treat f' as a *traditional function*, it can never represent the singular part of μ . The theory of distributions or *generalized functions* allows one to recover the full measure μ from f as the *distributional derivative* df/dx . For example, $(d/dx)\chi_{[0, \infty)} = \delta$.

Topographical maps and critical points. There is a 2-variable version of the Cantor function, due to Whitney, which gives a function $f(x, y)$ on the plane whose derivatives exist everywhere, but which is not constant on its critical set (the place where both df/dx and df/dy are zero). In fact the critical set is connected.

This function describes the topography of a hill with a (fractal) road running from top to bottom passing only along the level or flat parts of the hillside.

Bounded variation. Note that if $f = g - h$ where g and h are both monotone, then $f'(x)$ also exists a.e. So it is desirable to characterize the full vector space of functions spanned by the monotone functions.

A function $f : [a, b] \rightarrow \mathbb{R}$ has **bounded variation** if

$$\|f\|_{BV} = \sup \sum_1^n |f(a_i) - f(a_{i-1})| < \infty.$$

Here the sup is over all finite dissections of $[a, b]$ into subintervals, $a = a_0 < a_1 < \dots < a_n = b$. This supremum is called the *total variation* of f over $[a, b]$.

Note: the expression above gives norm zero to the constant functions, so often we mod out by these.

Theorem 6.11 *A function f is of bounded variation iff $f(x) = g(x) - h(x)$ where g and h are monotone increasing.*

Proof. Clearly $\|f\|_{BV} = f(b) - f(a)$ if f is monotone increasing, and thus f has bounded variation if it is a difference of monotone functions.

For the converse, let us write $y = y_+ - y_-$, where $y = y_+$ if $y \geq 0$ and $y = -y_-$ if $y \leq 0$. Define

$$f_+(x) = \sup \sum_1^n [f(a_i) - f(a_{i-1})]_+$$

over all partitions $a = a_0 < \dots < a_n = x$, and similarly

$$f_-(x) = \sup \sum_1^n [f(a_i) - f(a_{i-1})]_-.$$

Clearly f_+ and f_- are monotone increasing, and they are bounded since the total variation of f is bounded.

We claim $f(x) - f(a) = f_+(x) - f_-(x)$. To see this, note that if we refine our dissection of $[a, b]$, then both f_+ and f_- increase. Thus for any $\epsilon > 0$, we can find a dissection of $[a, x]$ for which *both* sums are within ϵ of their supremums. But for a common partition of $[a, x]$, it is clear that

$$\begin{aligned} \sum_1^n [f(a_i) - f(a_{i-1})]_+ - [f(a_i) - f(a_{i-1})]_- &= \\ \sum_1^n f(a_i) - f(a_{i-1}) &= f(a_n) - f(a_0) = f(x) - f(a). \end{aligned}$$

Thus $f(x) - f(a) = f_+(x) - f_-(x)$. ■

Corollary 6.12 Any function $f : [a, b] \rightarrow \mathbb{R}$ of bounded variation is differentiable a.e., and its derivative f' lies in $L^1([a, b])$.

Relation to measures. Just as monotone increasing functions are related to measure, functions of bounded variations are related to *signed measures*. The canonical representation of f as a difference of monotone functions corresponds to the Hahn decomposition, $\mu = \mu_+ - \mu_-$, μ_+ and μ_- mutually singular positive measures.

II. Integration and Differentiation

In this section we introduce the notion of absolutely continuous functions, and establish the main result relating integration and differentiation:

Theorem 6.13 The operator $D(F) = F'$ gives a bijective map

$$D : \{ \text{absolute continuous functions on } [a, b] \} / \mathbb{R} \rightarrow L^1[a, b].$$

Its inverse is given by $I(f) = F(x) = \int_a^x f$.

Even better, inspection of the proof will show that

$$\|F\|_{BV} = \|D(F)\|_{L^1},$$

and indeed the decomposition of F into increasing and decreasing parts, $F = F_+ - F_-$, will correspond to the decomposition of $f = F'$ into its positive and negative parts, $f = f_+ - f_-$. We remark that the absolutely continuous functions form a *closed* subspace of the functions of bounded variation (exercise).

Note that we have modded out by the constant on one side (since the derivative of a constant is zero), and by the functions that vanish a.e. on the other (since their integral is zero).

Absolute continuity. A function $F : [a, b] \rightarrow \mathbb{R}$ is *absolutely continuous* if for any $\epsilon > 0$, there is a $\delta > 0$ such that for any finite set of non-overlapping intervals (a_i, b_i) , if $\sum_1^n |a_i - b_i| < \delta$ then $\sum_1^n |f(a_i) - f(b_i)| < \epsilon$.

Here is a motivation for the definition.

Theorem 6.14 If $f : [a, b] \rightarrow \mathbb{R}$ is integrable then $F(x) = \int_a^x f$ is absolutely continuous.

Proof. This follows from the fact that for any $\epsilon > 0$ there is a $\delta > 0$ such that $\int_A |f| < \epsilon$ whenever $m(A) < \delta$. ■

Non-example. Absolute continuity is designed exactly to *rule out* examples like the Cantor function $f : [0, 1] \rightarrow [0, 1]$. Indeed, the Cantor function has the property that we can find 2^n intervals each of length 3^{-n} such that $\sum |f(I_i)| = 1$, even though $\sum |I_i| = (2/3)^n$. This violates absolute continuity.

More generally, the result above shows that if F is *not* absolutely continuous, then there is *no chance* that $F = \int F'$.

Necessity of disjointness. We note that in the definition of absolute continuity, it is *critical* that the intervals $[a_i, b_i]$ are *disjoint* (or at least non-overlapping).

For example, the function $f(x) = \sqrt{x}$ is absolutely continuous on $[0, 1]$, since it is the integral of $1/(2\sqrt{x})$. But given $\delta > 0$ we can take N copies of the interval $[a_i, b_i] = [0, \delta/N]$ to get $\sum |f(a_i) - f(b_i)| = N\sqrt{\delta/N} = \sqrt{N}\delta$, which can be made as large as we like by taking N sufficiently large.

The same problem would arise for any absolutely continuous function $f(x)$ with $f'(x)$ unbounded.

Theorem 6.15 *An absolutely continuous function is continuous and of bounded variation.*

Proof. Continuity is clear. As for bounded variation, choose ϵ and δ as above; then over any interval of length δ , the total variation of f is at most ϵ , so over $[a, b]$ we have variation about $\epsilon(b - a)/\delta$. ■

Corollary 6.16 *If $f : [a, b] \rightarrow \mathbb{R}$ is absolutely continuous, then $f' \in L^1([a, b])$.*

Part I. The derivative of an integral. Now we show the derivative of an integral gives the expected result; in other words, that $D(I(f)) = f$. We have already proved this for the indicator function of a measurable set (Lebesgue density); the following argument gives a different proof.

Logically, the argument is to show (i) I is injective and (ii) $I(D(I(f))) = I(f)$. In other words, if we let $F = I(f)$, then $\int F' = \int f$, so $F' = f$.

Lemma 6.17 *If $f : [a, b] \rightarrow \mathbb{R}$ is integrable, and $F(x) = \int_a^x f(t) dt = 0$ for all x , then $f = 0$.*

Proof. (This shows injectivity of the map I .) Consider the collection of all sets over which the integral of f is zero. By assumption this contains all intervals in $[a, b]$, and it is closed under countable unions and complements.

Thus it contains all closed sets in $[a, b]$. But if $f \neq 0$, then either $\{f > 0\}$ or $\{f < 0\}$ contains a closed set F of positive measure. Then $\int_F f \neq 0$, contradiction. Thus $f = 0$. ■

The bounded / Lipschitz case. Recall that a function $F : [a, b] \rightarrow \mathbb{R}$ is Lipschitz if there exists an M such that

$$|F(x) - F(y)| \leq M|x - y|$$

for all $x, y \in [a, b]$.

Theorem 6.18 *If f is bounded, then $F = I(f)$ is Lipschitz and satisfies $F'(x) = f(x)$ a.e.*

Proof. Suppose $|f| \leq M$; then clearly $|F(x+t) - F(x)| \leq Mt$. We will show $\int_a^c F'(x) - f(x) dx = 0$ for all c . To this end, just note that $F'(x) = \lim F_n(x) = n(F(x+1/n) - F(x))$ satisfies $|F_n| \leq M$, so it is a pointwise limit of bounded functions. Thus

$$\begin{aligned} \int_a^c F'(x) dx &= \lim \int_a^c F_n(x) dx = \lim n \left(\int_c^{b+1/n} F - \int_a^{a+1/n} F \right) \\ &= F(c) - F(a) = \int_a^c f(x) dx, \end{aligned}$$

by continuity of F . ■

It is immediate that $F'(x)$ is bounded if F is Lipschitz. Thus we have actually shown that the derivative gives an isomorphism of normed spaces:

$$D : C^{Lip}[a, b]/\mathbb{R} \rightarrow L^\infty[a, b].$$

On the left the norm is given by the best Lipschitz constant; on the right, it is given by the *essential supremum* $M = \|f\|_\infty$, which is the smallest real number such that $|f| \leq M$ a.e.

The general case: f integrable. We can now take one more step and show:

Theorem 6.19 *For any $f \in L^1([a, b])$, we have*

$$\frac{d}{dx} \int f = D(I(f)) = f.$$

Proof. By linearity, it is enough to prove this assertion for positive f . Let $f_n = \min(n, f) \rightarrow f$. Then f_n increases to f . Let $F_n = I(f_n)$; then by monotone convergence, F_n increases to F pointwise. Moreover, we have $F'_n(x) = f_n(x)$ by the preceding result on bounded functions, and clearly

$$F'(x) \geq F'_n(x) = f_n(x)$$

since $F(x) - F_n(x) = \int f - f_n$ is monotone. Therefore, for any $c \in [a, b]$, we have

$$\int_a^c F' \geq \int_a^c f_n = F_n(c) \rightarrow F(c) = \int_a^c f.$$

On the other hand,

$$\int_a^c F' dx \leq F(c) - F(a) = \int_a^c f$$

by our general result on monotone functions such as F . Thus equality holds, and we have shown

$$\int_a^c F'(x) - f(x) dx = 0$$

for all c . Thus $F'(x) = f(x)$ a.e. ■

Part II. The integral of a derivative. Now we reverse the order of operations, and show that if F is absolutely continuous, then $I(D(F)) = F$.

Logically, since we know $DI = \text{id}$, to show $ID = \text{id}$ we need only show that D is injective. For then $D(I(D(F))) = D(F)$ since $DI = \text{id}$, and then $ID(F) = F$ since D is injective. Here is the injectivity of D :

Lemma 6.20 *If F is absolutely continuous and $F'(x) = 0$ a.e. then F is constant.*

Proof. Pick any $c \in [a, b]$. Using the Vitali lemma, cover $[a, c]$ with a finite number of intervals I_1, \dots, I_n such that $|\Delta F/\Delta t| < \epsilon$ over these intervals, and what's left over has total measure at most ϵ . Then by absolute continuity, the total variation of F over the complementary intervals is at most δ . Thus

$$|F(c) - F(a)| \leq \delta + \epsilon \sum m(I_i) \leq \delta + \epsilon(b - a),$$

and this can be made arbitrarily small so $F(c) = F(a)$. ■

Corollary 6.21 For any absolutely continuous function $F : [a, b] \rightarrow \mathbb{R}$, with $F(a) = 0$, we have

$$\int F' = I(D(F)) = F.$$

Proof. We have $D(I(D(F))) = D(F)$, and since D is injective, this gives $I(D(F)) = F$. ■

Summary. Letting M =monotone functions, we have

$$BV = M - M \supset AC = \{f : f = \int f'\}.$$

We will eventually see the differentiation form of this setup:

$$\{\text{signed } \mu\} = \{\mu - \nu : \mu, \nu \geq 0\} \supset \{\mu = f(x) dx\} \iff L^1(\mathbb{R}).$$

Convexity. A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is *convex* if for all $x, y \in \mathbb{R}$ and $t \geq 0$ we have

$$f(tx + (1 - t)y) \leq tf(x) + (1 - t)f(y).$$

In other words, the graph of the function f lies below every one of its chords. We say f is *strictly convex* if we have strict inequality above, aside from the trivial cases.

Theorem 6.22 The right and left derivatives of a convex function exist for all x , and agree outside a countable set.

Proof. The secant lines move monotonely. ■

We have yet to use the Monotone Convergence Theorem. When can we assert that the approximations to the derivative, $f_t(x) = (f(x+t) - f(x))/t$, converge to $f'(x)$ *monotonely* as t decreases to zero?

Answer: when f is a convex function!

Theorem 6.23 A function f is convex iff f is absolutely continuous and $f'(x)$ is increasing.

Proof. Suppose f is convex. Then the slope of the secant line $f_t(x) = (f(x+t) - f(x))/t$ is an increasing function of t and of x . It follows that $f_t(x)$ is uniformly bounded on any compact interval $[c, d]$ in the domain of

f . Thus f is Lipschitz, which implies it is absolutely continuous. Finally $f'(x)$ is increasing since it is a limit of increasing functions.

For the converse, just note that when f is absolutely continuous, the secant slope

$$f_t(x) = \frac{1}{t} \int_x^{x+t} f'(s) dx$$

is just the average of f' . But the average of an increasing function is itself increasing, so the secants of the graph of f have increasing slope, which implies f is convex. ■

Corollary 6.24 *The convex functions are exactly the integrals of monotone increasing functions.*

Corollary 6.25 *If $f(x)$ is convex, then $f''(x)$ exists a.e. and $f''(x) \geq 0$.*

Optimization. Convex functions play a crucial role in practical applications because they are the easiest nonlinear functions to *optimize*. In this setting, optimizing f over $[a, b]$ means finding a point x where f achieves its minimum. By continuity, such a point exist; moreover, we have:

Proposition 6.26 *Any local minimum of a convex function is a global minimum. If f is strictly convex, the minimum is unique. In general, the set where a convex function f achieves its minimum is also convex.*

In several variables, when f is smooth, the minimum can be found by following $-\nabla f$. In other words, any local minimum is a global minimum.

Convexity and random variables. The next result is very useful for deriving inequalities. It also has interpretations in probability theory.

Theorem 6.27 (Jensen's inequality) *If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a convex function, and $X : [0, 1] \rightarrow \mathbb{R}$ is integrable, then*

$$f\left(\int X\right) \leq \int f(X) = \int_0^1 f(X(t)) dt.$$

Proof. First note that equality holds if f is a linear function. Also, both sides of the equation are linear functions of f (under pointwise addition). So it is enough to prove the result after modifying f by a linear function. To this end, let $m = \int X$ be the mean of X , take a linear supporting function $g(x) = ax + b$ with $g(m) = f(m)$ and $g(x) \leq f(x)$ otherwise; and replace f with $f - g$. Then $f(\int X) = 0$ but $\int f(X) \geq 0$. ■

Probabilistic interpretation. If f is convex, then $E(f(X)) \geq f(E(X))$ for any random variable X . Jensen's theorem is this statement where the distribution of the random variable is dictated by the function $X : [0, 1] \rightarrow \mathbb{R}$. **It includes** δ -masses as a special case, since these are obtained when X is a simple function.

The definition of convexity says the result holds for random variables assuming just two values x or y , with probabilities t and $(1-t)$ respectively.

A bettor's dilemma. You are about to gamble with \$100 on a fair game, where the payoff is proportional to your bet. A generous patron has offered to square your holdings. Do you ask for this boost before you start playing, to increase your stakes, or after you have gambled, to increase your payoff?

Answer: let X denote your payoff on a bet of \$100. To say the game is fair is to say $E(X) = 100$; on average you neither gain or lose. For squaring after the game, the expected earnings are $E(X^2)$; for squaring before, they are $E(100X) = E(X)^2$. Since x^2 is convex, $E(X^2) \geq (E(X))^2$, so squaring your payoff is better on average.

Example: in a game of double or nothing, squaring before betting gives a maximum payoff of 2×100^2 , while squaring after betting gives a maximum payoff of $(2 \times 100)^2 = 4 \times 100^2$.

Arithmetic/Geometric Mean. As is well-known, for $a, b > 0$, we have $\sqrt{ab} \leq (a+b)/2$, because:

$$0 \leq (\sqrt{a} - \sqrt{b})^2/2 = (a+b)/2 - \sqrt{ab}.$$

More generally, considering a random variable that assumes values a_1, \dots, a_n with equal likelihood, the concavity of the logarithm implies

$$E(\log X) = \frac{1}{n} \sum \log a_i \leq \log E(X) = \log \left(\frac{1}{n} \sum a_i \right)$$

and thus

$$\left(\prod a_i \right)^{1/n} \leq \frac{1}{n} \sum a_i.$$

Mnemonic: To remember the direction of this inequality, note that if $a_i \geq 0$ but $a_1 = 0$, then the geometric mean is zero but the arithmetic mean is not.

7 The Classical Banach Spaces

A *normed linear space* is a vector space V over \mathbb{R} or \mathbb{C} , equipped with a norm $\|v\| \geq 0$ defined for every vector, such that:

$\|v\| = 0 \implies v = 0$;
 $\|\alpha v\| = |\alpha| \cdot \|v\|$; and
 $\|v + w\| \leq \|v\| + \|w\|$. A norm is the marriage of metric and linear structures. It determines a distance by $d(v, w) = \|v - w\|$.

A *Banach space* is a complete normed linear space.

The unit ball. It is frequently useful to think of a norm in terms of its *closed unit ball*, $B = \{v : \|v\| \leq 1\}$. The conditions above insure:

B is convex; B is symmetric ($\alpha B = B$ if $|\alpha| = 1$); and B meets every line through the origin in a closed, nontrivial interval.

Conversely, if B satisfies these 3 properties, then it defines a norm on V by

$$\|v\| = \inf\{\alpha > 0 : v \in \alpha B\}.$$

The sub-additivity of the norm comes from convexity of B . Indeed, suppose B is convex and consider $x, y \in V$. Let $\alpha = \|x\|$, $\beta = \|y\|$. By scaling we can assume $\alpha + \beta = 1$. Write $x = \alpha x'$, $y = \beta y'$ with $\|x'\| = \|y'\| = 1$. Then by convexity of B , $\alpha x' + \beta y' \in B$, which implies $\|x + y\| \leq 1 = \|x\| + \|y\|$.

Theorem 7.1 (Verifying completeness) *A normed linear space is complete iff $\sum \|a_i\| < \infty \implies$ there is an $a \in V$ such that $\sum_1^N a_i \rightarrow a$.*

Proof. If a_i is a Cauchy sequence in V we can pass to a subsequence such that $d(a_i, a_{i+1}) < 2^{-i}$. Then $a_1 + (a_2 - a_1) + (a_3 - a_2) + \dots$ is absolutely summable, so it sums to some $s \in V$, and $a_i \rightarrow s$. The converse is obvious. ■

Example: $C(X)$. Let X be any compact Hausdorff space, and let $C(X)$ be the vector space of continuous functions $f : X \rightarrow \mathbb{R}$. Define $\|f\| = \sup_X |f|$. Then $\sum \|f_i\| < \infty$ implies the sum converges uniformly, and therefore $\sum f_i(x) = f(x)$ exists and is continuous; thus $C(X)$ is a Banach space.

Example: ℓ^p . Given $1 \leq p \leq \infty$, we let ℓ^p denote the space of sequences $a = (a_1, a_2, \dots)$ such that the corresponding norm is finite. The norms are defined by

$$\|a\|_p = \left(\sum |a_i|^p \right)^{1/p}, \quad \text{and}$$

$$\|a\|_\infty = \sup |a_i|.$$

The outer exponent is put in to give homogeneity of degree one.

Example: c and c_0 . Inside of ℓ^∞ we have the closed subspaces c and c_0 , consisting of the convergent sequences and the sequences such that $a_i \rightarrow 0$. Any closed subset of a Banach space is again a Banach space.

Example: $\ell^p(\mathbb{R}^n)$. It is useful to visualize the unit ball for the corresponding norm on the space of *finite* sequences, which we identify with \mathbb{R}^n . For $n = 2$, the unit ball defined by $x^p + y^p \leq 1$. Note that as p increases from 1 to ∞ , these balls are all convex, and they move steadily from a diamond through a circle to a square. In \mathbb{R}^3 they move from an octahedron through a sphere to a cube.

Example: Lipschitz, Hölder. For $0 < \alpha < 1$ we define the space $C^\alpha[a, b]$ to consist of the continuous functions satisfying a Hölder condition:

$$|f(x) - f(y)| \leq M \cdot |x - y|^\alpha.$$

Thus $C^1[a, b]$ consists of the *Lipschitz* functions (to avoid confusion with spaces of differentiable functions, this is sometimes denote $C^{Lip}[a, b]$). We define $\|f\|_{C^\alpha}$ to be the best possible value for M in the inequality above. Then $C^\alpha(\mathbb{R})/\mathbb{R}$ becomes a Banach space.

For example, the Cantor function belongs to $C^{2/3}$, since it sends intervals of length 3^{-n} to intervals of length $2^{-n} = (3^{-n})^{2/3}$.

Example: Bounded variation, absolutely continuous. As remarked earlier, $BV[a, b]/\mathbb{R}$ carries a natural norm coming from the total variation; this makes it into a Banach space. The absolutely continuous functions give a closed subspace $AC[a, b]/\mathbb{R}$.

The L^p spaces. For any measurable subset $E \subset \mathbb{R}$, and $1 \leq p \leq \infty$, we define $L^p(E)$ as the set of measurable functions $f : E \rightarrow \mathbb{R}$ such that $\int_E |f|^p < \infty$; and set

$$\begin{aligned} \|f\|_p &= \left(\int_E |f|^p \right)^{1/p}, \\ \|f\|_\infty &= \inf\{M \geq 0 : |f| \leq M \text{ a.e.}\}. \end{aligned}$$

Actually for the norm of f to vanish, it is only necessary for f to vanish a.e., so the elements of L^p are technically equivalence classes of functions defined up to agreement a.e.

Our first result verifies that we actually have a norm.

Theorem 7.2 (Minkowski's inequality) $\|f + g\|_p \leq \|f\|_p + \|g\|_p$. For $1 < p < \infty$, equality holds iff f and g lie on a line in L^p .

Proof. As mentioned above, it suffices to verify convexity of the unit ball; that is, assuming $\|f\| = \|g\| = 1$, we need only verify

$$\|tf + (1-t)g\| \leq 1$$

for $0 < t < 1$. In fact by convexity of the function x^p , $p > 1$, we have

$$\int |tf + (1-t)g|^p \leq \int t|f|^p + (1-t)|g|^p \leq t + (1-t) = 1.$$

This proves B is convex. For $1 < p$, the strict convexity of x^p gives strict convexity of B , furnishing strict inequality unless f and g lie on a line. ■

The scale of spaces. If $m(E) < \infty$ then we have

$$L^\infty(E) \subset L^p(E) \subset L^1(E),$$

i.e. the L^p spaces shrink as p rises. Precise inequalities between norms follow from Hölder's inequality below.

If $m(E) = \infty$, then there is no such comparison.

Continuity at $p = \infty$. If $m(A)$ is finite, then

$$\|\chi_A\|_p = (m(A))^{1/p} \rightarrow 1$$

as $p \rightarrow \infty$. More generally, if $f \in L^\infty[a, b]$, then

$$\|f\|_p \rightarrow \|f\|_\infty$$

as $p \rightarrow \infty$. (To see this, consider the case where $\|f\|_\infty = 1$.)

Completeness. We now come to one of the main motivations for the study of Lebesgue measurable functions.

Theorem 7.3 *For $1 \leq p \leq \infty$, the space $L^p(E)$ with the norm above is a Banach space.*

Proof. Suppose $\sum \|f_i\|_p < \infty$. Let $F(x) = \sum |f_i(x)|$. Then by monotone convergence, $\int F^p \leq (\sum_1^\infty \|f_i\|_p)^p$, so $F(x)$ is finite a.e. and it lies in L^p . Therefore the same is true for $f(x) = \sum f_i(x)$, since $|f(x)| \leq F(x)$; and we have $\|f\|_p \leq \sum \|f_i\|_p$. By virtue of the last inequality we also have

$$\|f - \sum_1^n f_i\|_p \leq \sum_{n+1}^\infty \|f_i\|_p \rightarrow 0,$$

and thus every absolutely summable sequence is summable, and L^p is complete. ■

Inequalities. The most important inequality in mathematics is the triangle inequality. Here is the runner-up.

Theorem 7.4 (Cauchy-Schwarz-Bunyakovskii inequality) *If f and g are in L^2 , then fg is in L^1 and*

$$\langle f, g \rangle = \int fg \leq \|f\|_2 \|g\|_2.$$

Proof 1. We can assume $\|f\|_2 = \|g\|_2 = 1$. Then

$$0 \leq \langle f \pm g, f \pm g \rangle = 2 \pm 2\langle f, g \rangle,$$

so $|\langle f, g \rangle| \leq 1$ as desired.

Proof 2. We can assume $f, g \geq 0$. For any $t > 0$ we have

$$\|f + tg\|^2 \leq (\|f\| + t\|g\|)^2 \leq \|f\|^2 + 2t\|f\|\|g\| + O(t^2),$$

while at the same time

$$\|f + tg\|^2 = \|f\|^2 + 2t\langle f, g \rangle + t^2\|g\|^2 \geq \|f\|^2 + 2t\langle f, g \rangle;$$

comparing terms of size $O(t)$, we find $\|f\|\|g\| \geq \langle f, g \rangle$. ■

Hilbert space. The *inner product* $\langle f, g \rangle$ is a symmetric, definite bilinear form making L^2 into a *Hilbert space*. It is an infinite-dimensional analogue of the inner product in \mathbb{R}^n . For example, if E and F are disjoint measurable sets, then $L^2(E)$ and $L^2(F)$ are orthogonal subspaces inside $L^2(\mathbb{R})$.

Hölder's inequality. A simpler (but still very useful) inequality is:

$$\left\| \int fg \right\| \leq \|f\|_1 \cdot \|g\|_{\infty}.$$

In fact, this inequality and the Cauchy-Schwarz inequality are two instances of a continuous family of similar inequalities.

Theorem 7.5 *Suppose $1/p + 1/q = 1$. Then for $f \in L^p(E)$ and $g \in L^q(E)$, we have*

$$\left\| \int fg \right\| \leq \|f\|_p \cdot \|g\|_q.$$

Young's inequality. One way to prove Hölder's inequality is to generalize the fact that $2|ab| \leq (a^2 + b^2)$ to an *inhomogeneous* inequality.

Lemma 7.6 For any $a, b \geq 0$ in \mathbb{R} , and $1/p + 1/q = 1$, we have

$$ab \leq a^p/p + b^q/q.$$

Proof 1 of Hölder's inequality. Draw the curve $y = x^{p-1}$, which is the same as the curve $x = y^{q-1}$. Then the area inside the rectangle $[0, a] \times [0, b]$ is bounded above by the sum of a^p/p , the area between the graph and $[0, a]$, and b^q/q , the area between the graph and $[0, b]$.)

Proof 2. We will mimic Proof 2 of Cauchy–Schwarz, and argue that Hölder's inequality follows from Minkowski's inequality by 'differentiation'.

We may assume $f, g \geq 0$, with $f \in L^p$ and $g \in L^q$. Then $f^{p/q} \in L^q$, and using the binomial expansion for $(a + b)^q$ we have:

$$\begin{aligned} \|f^{p/q} + tg\|_q^q &\leq \|f^{p/q}\|_q^q + qt\|f^{p/q}\|_q^{q-1}\|g\|_q + O(t^2) \\ &= \|f\|_p^p + qt\|f\|_p\|g\|_q + O(t^2) \end{aligned}$$

since $(q - 1)/q = 1/p$. On the other hand for $f, g \geq 0$ we have (by convexity of x^p),

$$\begin{aligned} \|f^{p/q} + tg\|_q^q &= \int |f^{p/q} + tg|^q \geq \int (f^{p/q})^q + qt(f^{p/q})^{q-1}g \\ &\geq \|f\|_p^p + qt \int fg. \end{aligned}$$

Putting these inequalities together gives the theorem. ■

Corollary 7.7 If $f : [a, b] \rightarrow \mathbb{R}$ is absolutely continuous, and $f' \in L^p[a, b]$, then f is Hölder continuous of exponent $1 - 1/p$.

If $p = 1$ we get no information. If $p = \infty$ we get Lipschitz continuity.

Proof. We have

$$|f(x) - f(y)| = \left| \int_x^y 1 \cdot f'(t) dt \right| \leq \|\chi_{[x,y]}\|_q \|f'\|_p \leq \|f'\|_p |x - y|^{1/q}.$$

■

Approximation. We now recast our earlier approximation theorems.

Theorem 7.8 (Density of simple functions) *For any p and E , simple functions are dense in $L^p(E)$. For $p \neq \infty$, step, continuous and smooth functions are dense in $L^p(\mathbb{R})$.*

Proof. First we treat $f \in L^\infty$. Then f is bounded, so it is a limit of simple functions in the usual way (cut the range into finitely many small intervals and round f down so it takes values in the endpoints of these intervals).

Now for $f \in L^p$, $p \neq \infty$, we can truncate f in the domain and range to obtain bounded functions with compact support, $f_M \rightarrow f$. Since $f - f_M \rightarrow 0$ pointwise and $|f - f_M|^p \leq |f|^p$, dominated convergence shows $\|f - f_M\|_p \rightarrow 0$. Finally we can find step, continuous or smooth functions $g_n \rightarrow f_M$ pointwise, and bounded in the same way. Then $\int |g_n - f_M|^p \rightarrow 0$ by bounded convergence, so such functions are dense. ■

Note! The step, continuous and smooth functions are **not** dense in L^∞ !

L^p as a completion. Given say $V = C_0^\infty(\mathbb{R})$ with the L^2 -norm, it is exceedingly natural to form the metric completion \bar{V} of V and obtain a Banach space. But what are the elements of this space? The virtue of measurable functions is that they do suffice to represent all elements of \bar{V} .

It is this completeness that makes measurable functions as important as real numbers.

Duality. Given a Banach space X , we let X^* denote the **dual space** of bounded linear functionals $\phi : X \rightarrow \mathbb{R}$, with the norm

$$\|\phi\| = \sup_{x \in X - \{0\}} \frac{|\phi(x)|}{\|x\|}.$$

There is a natural map $X \rightarrow X^{**}$. If $X = X^{**}$ then X is *reflexive*.

Theorem 7.9 *The dual space X^* is a Banach space.*

Proof. Suppose $\sum \|\phi_i\| < \infty$. Then for any $f \in X$, we have

$$\sum |\phi_i(f)| \leq \|f\| \cdot \sum \|\phi_i\|.$$

Thus we can define $\phi(f) = \sum \phi_i(f)$, since the latter sum is absolutely convergent. The resulting map ϕ is clearly linear, and it satisfies $\|\phi\| \leq \sum \|\phi_i\|$, so it belongs to X^* ; and evidently ϕ is the limit of the sums $\sum_1^n \phi_i$ as $n \rightarrow \infty$. ■

Uniform convergence of linear maps. A quick way to phrase this proof is that we have an *isometry*

$$X^* \subset C(B_X),$$

where B_X is the unit ball of X , and $C(B_X)$ is the Banach space of continuous functions on the ball in the sup norm; and the space of linear maps on the ball is closed under uniform limits.

The natural inclusion. By Hölder's inequality, we have a natural inclusion

$$L^q(\mathbb{R}) \subset L^p(\mathbb{R}),$$

In fact this map is an isometry, as can be seen by showing Hölder's inequality is sharp. Indeed, given $f \in L^p$ of norm one, we can set $g = \text{sign}(f)|f|^{p/q}$; then $g \in L^q$ also has norm 1, and we find

$$\int fg = \int |f|^{1+p/q} = \int |f|^p = 1,$$

since $(1 + p/q) = p(1/q + 1/p) = p$.

When is this map onto? The answer is provided by the *Riesz representation theorem* below. Note that the case $p = \infty$ is *excluded*.

Theorem 7.10 (Riesz representation) *If $p \neq \infty$, then $L^p(\mathbb{R})^* = L^q(\mathbb{R})$.*

Corollary 7.11 *For $1 < p < \infty$, L^p is reflexive.*

To give the main idea of the proof, we first treat the case of L^1 .

Theorem 7.12 *The dual of $L^1(\mathbb{R})$ is $L^\infty(\mathbb{R})$.*

Proof. Given $\phi \in L^1(\mathbb{R})^*$, we must find a $g \in L^\infty(\mathbb{R})$ such that

$$\phi(f) = \phi_g(f) = \int fg$$

for all $f \in L^1$. Now if this were to hold, we could obtain the integral of g over $[a, b]$ by simply evaluating $\phi(\chi_{[a,b]})$. With this in mind, we set

$$G(x) = \phi(\chi_{[0, x]})$$

for $x \geq 0$, and $G(x) = -\phi(\chi_{[x, 0]})$ for $x < 0$. Then

$$\phi(\chi_{[a,b]}) = G(b) - G(a)$$

for all $a < b$ in \mathbb{R} . Moreover, we have

$$|G(b) - G(a)| \leq \|\phi\| \cdot \|\chi_{[a,b]}\|_1 = \|\phi\| \cdot (b - a).$$

Thus G is Lipschitz, and hence $g(x) = G'(x)$ exists a.e. and lies in $L^\infty(\mathbb{R})$. Now by construction,

$$\phi_g(\chi_{[a,b]}) = \int_a^b G'(x) dx = G(b) - G(a) = \phi(\chi_{[a,b]}),$$

so ϕ and ϕ_g agree on indicator functions of intervals. By linearity, they agree on step functions. But step functions are dense in $L^1[a, b]$, so by continuity we have $\phi = \phi_g$. ■

Proof of the general Riesz representation theorem. Now consider any p with $1 \leq p < \infty$. Let $\phi : L^p(\mathbb{R}) \rightarrow \mathbb{R}$ be a bounded linear functional, with $M = \|\phi\|$. Define $G(x)$ as above by evaluating ϕ on indicator functions of intervals. Then for any collection of disjoint intervals (a_i, b_i) with $\sum |a_i - b_i| < \delta$, we have

$$\begin{aligned} \sum |G(a_i) - G(b_i)| &= \sum |\phi(\chi_{(a_i, b_i)})| = \phi\left(\sum \pm \chi_{(a_i, b_i)}\right) \\ &\leq M \left(\sum |a_i - b_i|\right)^{1/p} \leq M\delta^{1/p}. \end{aligned}$$

Thus $G(x)$ is absolutely continuous, and thus there is a locally integrable function $g(x) = G'(x)$ such that

$$\phi(\chi_I) = \phi_g(\chi_I)$$

for any interval I .

It is now routine to check that $\phi(f) = \phi_g(f)$ for all bounded functions with bounded support, by using the Bounded Convergence Theorem and density of step functions.

The main step is to prove that $g \in L^q(\mathbb{R})$. To this end, let g_n be the truncation of g to a function with $|g_n| \leq n$, supported on $[-n, n]$. We treat the case where $g \geq 0$; the general case is handled by just introducing an appropriate sign when needed. Then:

$$\begin{aligned} \int |g_n|^q &\leq \langle g_n^{q-1}, g \rangle = \phi(g_n^{q-1}) \leq M \|g_n^{q-1}\|_p \\ &= M \left(\int |g_n|^{(q-1)p} \right)^{1/p} = M \left(\int |g_n|^q \right)^{1/p}, \end{aligned}$$

since $pq = p + q$. Thus for every n we have

$$\int |g_n|^q \leq M^{1/(1-1/p)} = M^q.$$

Taking the limit as $n \rightarrow \infty$ and applying Fatou's lemma or monotone convergence, we have $\|g\|_q \leq M$.

Thus by Hölder's inequality, ϕ_g is a continuous linear functional that agrees with ϕ on a dense set in L^p , so $\phi = \phi_g$. ■

Duality for L^∞ . The dual of L^∞ is larger than L^1 . To indicate the proof, recall the analogous fact that we used an ultrafilter to construct a bounded linear function $L : \ell^\infty \rightarrow \mathbb{R}$, extending the usual function $L(a_n) = \lim a_n$ on $c \subset \ell^\infty$. On the other hand, for any $b \in \ell^1$ we can find $a \in c$ such that $L(a) = 1$ but $\langle a, b \rangle$ is as small as we like (slide the support of a off towards infinity.) Thus L is a linear functional that is not represented by any element of ℓ^1 .

A similar construction can be carried out by extending the point evaluations from $C[a, b]$ to $L^\infty[a, b]$.

Duality for general measure spaces. The fact that $L^p(E)^* = L^q(E)$ for $p \neq \infty$ also holds for general abstract measure spaces, where differentiation is not available. Instead one uses the definition $\mu(E) = \phi(\chi_E)$ to define a new *signed measure* on E , and then one takes the *Radon-Nikodym derivative* $g = d\mu/dm$, and shows that $g \in L^q(E)$ and $\phi = \phi_g$.

More on Hilbert spaces. For Hilbert space we have the important self-dual property, $L^2(\mathbb{R}) \cong L^2(\mathbb{R})^*$. The original Riesz representation theorem asserted exactly this equality.

Now let X be an abstract Hilbert space. Then the inner product gives a natural map $X \rightarrow X^*$ sending f to $\phi_f(g) = \langle f, g \rangle$.

Theorem 7.13 (Abstract Riesz theorem) *The bracket on X gives an isomorphism between X and X^* .*

Orthonormal sets and bases. For the proof it is useful to introduce the notion of an orthonormal set and basis. First, a collection of vectors (e_i) , indexed by $i \in I$, is *orthonormal* if $\langle e_i, e_j \rangle = \delta_{ij}$.

Theorem 7.14 (Bessel's inequality) *For any orthonormal set e_i and any $x \in X$, if we set $a_i = \langle x, e_i \rangle$ then*

$$\sum |a_i|^2 \leq \|x\|^2.$$

Proof. It suffices to prove this for finite sums, and for this it follows by considering $\langle y, y \rangle$ where $y = x - \sum a_i e_i$. ■

Remark. In fact $\sum |a_i|^2$ gives $\|y\|^2$, where y is the orthogonal projection of x to the span of the (e_i) .

Bases. A *maximal* orthonormal set forms a *basis* for X . We also say an orthonormal set (e_i) is *complete* in this case. It has the property that whenever $\langle y, e_i \rangle = 0$ for all e_i , we have $y = 0$; otherwise we could add $y/\|y\|$ to the basis.

Theorem 7.15 *Every Hilbert space X has a basis. If X is separable, then the basis is countable.*

Proof. The existence of a basis comes from the Axiom of Choice. For the second statement, note that $\|e_i - e_j\| = \sqrt{2} > 1$, so the balls $B(e_i, 1/2)$ are disjoint. If X is separable (has a countable dense set), then there are only finitely many such balls. ■

Examples.

1. The easiest example of a Hilbert space is ℓ^2 , with the basis e_k being the sequence δ_{ik} .
2. Suitable multiples of 1, $\sin(nx)$ and $\cos(nx)$, $n \geq 1$, form an orthonormal basis for $L^2[0, 2\pi]$.
3. Similarly, $\langle z^n : n \in \mathbb{Z} \rangle$ gives an orthonormal basis for the complex Hilbert space $L^2(S^1)$, if we normalize so the total length of the circle $S^1 \subset \mathbb{C}$ is 1.
4. Another nice example is $L^2[0, \pi]$, with $e_n = \sin(nx)$ for $n \geq 1$, and with the revised norm:

$$\|f\|_2 = \frac{2}{\pi} \int_0^\pi f(x) dx.$$

Note that $\int_0^\pi \sin^2(nx) + \cos^2(nx) dx = \pi$, so the integral of $\sin^2(nx)$ gives half of this. We get a basis by (2) above, because f can be extended to an odd function on $[-\pi, \pi]$.

The Gram-Schmidt process. A more constructive proof of the existence of a basis, in the case of a separable Hilbert space X , is the following. As input data, suppose we are given an *algebraic basis* a_n for a *dense subspace* of X . (As a concrete example, one can take $a_n(x) = x^n$ for $X = L^2[-1, 1]$, for $n = 0, 1, 2, \dots$. Then a_n forms a basis for the space of polynomials.)

We can then obtain an orthogonal basis for X by setting $b_0 = a_0$, and then defining inductively

$$b_n = a_n - \sum_{i=0}^{n-1} \frac{\langle a_n, e_i \rangle e_i}{\langle e_i, e_i \rangle}.$$

The expression above simply removes from a_n its orthogonal projection to the linear span of (a_0, \dots, a_{n-1}) , which is the same as the linear span of (e_0, \dots, e_{n-1}) .

Finally to obtain an orthonormal basis, we set $e_n = b_n / \|b_n\|$.

In the concrete case of polynomials $a_n(x) = x^n$ on $[-1, 1]$, the result of this process is, up to constant multiples, the *Legendre polynomials* $P_n(x)$ of degree n . For example, $b_0(x) = 1$, $b_1(x) = x$, and $b_2(x) = x^2 - 1/3$.

Completeness. In all these cases, to check that we really have a basis one must verify completeness; that is, one must show that one of the following equivalent conditions holds:

1. (e_n) is a maximal orthonormal set.
2. The only $x \in X$ with $\langle x, e_n \rangle = 0$ for all n is $x = 0$.
3. The algebraic span of the vectors e_n is a dense subspace of X .

The last condition is the usually the easiest one to verify, and we will eventually do this using the Stone–Weierstrass theorem.

Applications of the basis.

Theorem 7.16 *Let (e_i) be a basis for X . Then for any $x \in X$ we have*

$$x = \sum a_i e_i \quad \text{and} \quad \|x\|^2 = \sum |a_i|^2,$$

where $a_i = \langle x, e_i \rangle$.

Proof. Define a_i as above. Then by Bessel's inequality, we have $\sum |a_i|^2 \leq \|x\|^2 < \infty$, so the sum converges. By continuity, $y = \sum a_i e_i$ satisfies $\langle y, e_i \rangle = a_i$, and hence $\langle y - x, e_i \rangle = 0$ for all i . Since (e_i) is maximal, this implies $y = x$. ■

Corollary 7.17 *Every Hilbert space is isomorphism to $\ell^2(I)$ for some index set I .*

Proof of Theorem 7.13 (Abstract Riesz representation) : Let (e_i) be an orthonormal basis for X . Given $\phi \in X^*$, let $a_i = \phi(e_i)$. For any finite set $J \subset I$, let $x_J = \sum_J a_i e_i$. Then we have

$$|\phi(x_J)| = \sum_J |a_i|^2 \leq \|\phi\| \cdot \|x_J\| = \|\phi\| \cdot \left(\sum_J |a_i|^2 \right)^{1/2}.$$

Thus

$$\sum_J |a_i|^2 \leq \|\phi\|^2$$

for all finite sets J . This implies that $y = \sum_I a_i e_i$ exists and gives an element of X such that $\langle y, e_i \rangle = a_i$ for all i . But this implies that ϕ is represented by y , and hence $X \cong X^*$. ■

8 Baire Category

In this section we discuss a topological notion of thin sets, analogous to sets of measure zero. The remarkable property that these thin sets share with sets of measure zero is that they are closed under countable unions.

The setting. Throughout, X will be a *nonempty, complete* metric space. (Note: a compact Hausdorff space would work just as well.)

Definitions. We say $F \subset X$ is *nowhere dense* if \overline{F} has empty interior. We say E is *meager* if it can be written as a countable union of nowhere dense sets $E = \bigcup_1^\infty F_i$. For example, any countable set (such as $\mathbb{Q} \subset \mathbb{R}$) is meager.

We say E is *residual* if its complement is meager.

Here are some readily verified properties.

1. Any subset of a meager set is meager.
2. The collection of meager sets is closed under countable unions.
3. The collection of residual sets is closed under countable intersections.
4. A countable intersection of dense open sets is residual.

The main result is that a residual set can never be empty, in fact it must always be dense.

Theorem 8.1 *Let X be a nonempty complete metric space, and let U_i be a sequence of dense open sets in X . Then $\bigcap U_i$ is dense.*

In particular, $\bigcap U_i$ is nonempty!

Proof. We will define a nested sequence of closed balls $B_0 \supset B_1 \supset \dots$ by induction. Let B_0 be arbitrary. Since U_n is dense, it meets the interior of B_n ; choose B_{n+1} to be any ball contained in $B_n \cap U_n$, with $\text{diam } B_{n+1} \leq 1/(n+1)$.

Then (if $X \neq \emptyset$), the centers of the balls B_n form a Cauchy sequence, so they converge to a limit $x \in X$. By construction, $x \in B_0 \cap \bigcap U_i$. Since B_0 was arbitrary, this shows $\bigcap U_i$ is dense. ■

Corollary 8.2 *A set is residual iff it contains a dense G_δ .*

Corollary 8.3 *A countable intersection of dense G_δ 's is again a dense G_δ .*

Corollary 8.4 *The complement of any meager set is dense.*

Measure and category. The sets of measure zero and the meager sets in \mathbb{R} both form σ -ideals (in the ring of all subsets of \mathbb{R}). That is, they are closed under taking subsets and countable unions.

Sometimes meager sets are called sets of *first category*. Even worse, a set of *second category* is not a residual set but rather a set that is not of first category. We will avoid this terminology!

Picture of Baire category. Consider a countable collection of embedded, possibly fractal arcs in the square, $\gamma_n \subset [0, 1]^2$. We can even have arcs of positive area. Nevertheless, each γ_n is nowhere dense, so $\bigcup \gamma_n$ is a *meager*; in particular, these countably many arcs cannot *cover* the square. A generic point of the square lies outside of all these arcs.

Application to continuous maps. Here is an interesting consequence of the Baire category theorem:

Theorem 8.5 *There is no function $f : [0, 1] \rightarrow \mathbb{R}$ that is continuous at the rational points and discontinuous at the irrational points.*

Proof. The set E of points of where f is continuous is a G_δ . If E contains the rational numbers then it is a dense G_δ , so E is a residual set. But if $E = \mathbb{Q} \cap [0, 1]$, then E is also meager — a contradiction. ■

The complementary formulation of Baire's theorem. The following reformulation of Baire's theorem is often useful.

Theorem 8.6 *Let $X = \bigcup_1^\infty F_i$ be a nonempty, complete metric space. Then the interior of $\overline{F_i}$ is nonempty for some i .*

Proof. Otherwise $\bigcup \overline{F_i}$ would express the entire space X as a meager set. ■

Applications. There are a wide variety of applications of the Baire category theorem. It is useful, especially in the complementary form, to deduce otherwise unexpected uniformity statements. It is also useful in the construction of examples and counterexamples, where a 'random' function or set or other object seems likely to 'work'. In function spaces choosing a point at 'random' with respect to a measure is problematic, but the notion of category applies so long as we have a complete metric.

We begin with an analytic application.

Theorem 8.7 (Uniform boundedness) *Let \mathcal{F} be a collection of continuous functions on a (nonempty) complete metric space X , such that for each x the functions are bounded — i.e. $\sup_{\mathcal{F}} f(x) \leq M_x$. Then there is a open set $U \neq \emptyset$ on which the functions are uniformly bounded: $\sup_U f(x) \leq M$ for all $f \in \mathcal{F}$.*

Proof. Let $F_n = \{x : f(x) \leq n \forall f \in \mathcal{F}\}$. Then F_n is closed and $\bigcup F_n = X$, so some F_M has nonempty interior U . ■

We will later use this result to deduce one of the three basic principles of functional analysis, the 'uniform boundedness principle'.

Diophantine approximation. We now turn to an application to number theory. Recall that by the pigeon hole principle, for any irrational number x we can find infinitely many rationals p/q such that

$$\left| x - \frac{p}{q} \right| \leq \frac{1}{q^2}.$$

(Proof: consider $x, 2x, \dots, nx \bmod 1$. Then $|n_1x - n_2x| \leq 1/n$ for some $0 \leq n_1 < n_2 \leq n$. Let $q = (n_2 - n_1)$. Then $0 \leq qx \bmod 1 \leq 1/n \leq 1/q$, which says $|p - qx| \leq 1/q$ for some integer p .)

A real number x is *Diophantine of exponent* α if there is a $C > 0$ such that

$$\left| x - \frac{p}{q} \right| > \frac{C}{q^\alpha}$$

for all rational numbers p/q . For example, if x is a *quadratic irrational*, then we also have a reverse inequality of the type above, saying x cannot be too well approximated by rationals. More generally, we have:

Theorem 8.8 *If x is algebraic of degree $d > 1$, then it is Diophantine of exponent d .*

Proof. Let $f(t) = a_0 t^d + \dots + a_d$ be an irreducible polynomial with integral coefficients satisfied by x . Then $|f(p/q)| \geq 1/q^d$. Since $|f'|$ is bounded, say by M , near x , we find

$$q^{-d} \leq |f(x) - f(p/q)| \leq M|x - p/q|$$

and thus $|x - p/q| \geq 1/(Mq^d)$. ■

Roth has proved the deep theorem that any algebraic number is Diophantine of exponent $2 + \epsilon$.

A number is Diophantine of exponent 2 iff the coefficients in its continued fraction expansion are bounded. For quadratic numbers, these coefficients are pre-periodic.

Liouville numbers. We say x is *Liouville* if x is irrational but for any $n > 0$ there exists a rational number with $|x - p/q| < q^{-n}$. Such a number is not Diophantine for any exponent, so it must be transcendental.

For example, $x = \sum 1/10^{n!}$ is an explicit and easy example of a transcendental number.

Theorem 8.9 (Measure vs. Category) *A random $x \in [0, 1]$ is Diophantine of exponent $2 + \epsilon$ for all $\epsilon > 0$. However a generic $x \in [0, 1]$ is Liouville.*

Proof. For the first part, fix $\epsilon > 0$, and let

$$E_q = \{x \in [0, 1] : \exists p, |x - p/q| < 1/q^{2+\epsilon}\}.$$

Since there are only q choices for p , we find $m(E_q) = O(1/q^{1+\epsilon})$, and thus $\sum m(E_q) < \infty$. Thus $m(\limsup E_q) = 0$ (by easy Borel-Cantelli). But this means that for almost every $x \in [0, 1]$, only finitely rationals approximate x to within $1/q^{2+\epsilon}$. Thus x is Diophantine of exponent $2 + \epsilon$. Taking a sequence

$\epsilon_n \rightarrow 0$ we conclude that almost every x is Diophantine of exponent $2 + \epsilon$ for all $\epsilon > 0$.

For the second part, just note that

$$E_n = \{x \in [0, 1] : \exists p, q, |x - p/q| < 1/q^n\}$$

contains the rationals and is open. Thus $\bigcap E_n = L$ is the set of Liouville numbers, and by construction it is a dense G_δ . ■

Sets with no category.

Lemma 8.10 *The set of closed subsets of \mathbb{R} has the same cardinality as \mathbb{R} itself.*

Lemma 8.11 *A closed subset of \mathbb{R} with no isolated points contains a Cantor set.*

Lemma 8.12 *Every uncountable closed set E in \mathbb{R} contains a Cantor set.*

Proof. Consider the subset F of $x \in E$ such that $F \cap B(x, r)$ is uncountable. It is easy to see that F is a nonempty, closed set, without isolated points, using the fact that countable unions preserve countable sets. Thus F contains a Cantor set. ■

Corollary 8.13 *Every uncountable closed set satisfies $|F| = |\mathbb{R}|$.*

Corollary 8.14 *Every set of positive measure contains a Cantor set.*

Proof. It contains a compact set of positive measure, which is necessarily uncountable. ■

By similar arguments, it is not hard to show:

Theorem 8.15 *Every dense G_δ set $X \subset [a, b]$ contains a Cantor set.*

Theorem 8.16 *There exists a set $X \subset \mathbb{R}$ such that X and X' both meet every uncountable closed set.*

Proof. Use transfinite induction on the smallest ordinal with $|\mathfrak{c}| = |\mathbb{R}|$, and the fact that the closed sets also have cardinality \mathfrak{c} . At each stage in the induction we have chosen less than \mathfrak{c} points, so we have not filled up any uncountable closed set. ■

Such a set X is called a *Bernstein set*.

Corollary 8.17 *If X is a Bernstein set, then for any interval $[a, b]$, neither $X \cap [a, b]$ nor $\tilde{X} \cap [a, b]$ is meager.*

Proof. If $X \cap [a, b]$ is meager, the complement of X contains a Cantor set K , contradicting the fact that X meets K . The same reasoning applies to \tilde{X} . ■

Thus X is an analogue, in the theory of category, of a non-measurable set. (One can think of a set X that meets **some** open set in a set of second category, as a set of positive measure).

Games and category. (Oxtoby, §6.) Let $X \subset [0, 1]$ be a set. Players A and B play the following game: they alternately choose intervals $A_1 \supset B_1 \supset A_2 \supset B_2 \cdots$ in $[0, 1]$, then form the intersection $Y = \bigcap A_i = \bigcap B_i$. Player A wins if Y meets X , otherwise player B wins.

Theorem 8.18 *There is a winning strategy for B iff X is meager.*

Proof. If X is meager then it is contained in a countable union of nowhere dense closed sets, $\bigcup F_n$. Player B simply chooses B_n so it is disjoint from F_n , and then $\bigcap B_n$ is disjoint from X .

Conversely, suppose B has a winning strategy. Then using this strategy, we can find a set of disjoint ‘first moves’ B_1^i that are dense in $[0, 1]$. To see this, let $B_1(A)$ be B ’s move if $A_1 = A$. Let J_1, J_2, \dots be a list of the intervals with rational endpoints in $[0, 1]$. Inductively define $B_1^1 = B(J_1)$ and $B_1^{i+1} = B(J_k)$ for the first k such that J_k is disjoint from B_1^1, \dots, B_1^i . Then every J_k meets some B_1^i so $\bigcup B_1^i$ is dense.

Similarly, we can find disjoint second moves that are dense in B_1^i for each i . Putting all these together, we obtain moves B_2^i , each contained in some B_1^i , that are also dense in $[0, 1]$.

Continuing in this way, we obtain a sequence B_k^i such that $U_k = \bigcup_i B_k^i$ is dense in $[0, 1]$. Let $Z = \bigcap U_k$. Any point $x \in Z$ is contained in a unique nested sequence $B_1^{i_1} \supset B_2^{i_2} \supset \dots$ obtained using B ’s winning strategy. Thus $x \notin X$. This shows X is disjoint from the dense G_δ Z , and thus X is meager. ■

By the same reasoning we have:

Theorem 8.19 *Player A has a winning strategy iff there is an interval A_1 such that $I \cap A_1$ has second category.*

Corollary 8.20 *There exists a set X such that neither A nor B has a winning strategy!*

One might try to take X equal to a non-measurable set $P \subset [0, 1) \cong S^1 = \mathbb{R}/\mathbb{Z}$ constructed so that $\mathbb{Q} + P = S^1$. By the Baire category theorem, P is not meager, but it is also not residual, since $P \cap P + 1/2 = \emptyset$.

However it might be the case that $P \cap I$ is small (even empty!) for some interval I . To remedy this, one considers instead a *Bernstein set*, i.e. a set X such that both X and its complement X' meet every uncountable closed subset of S^1 . Then, as we have seen above, $X \cap [a, b]$ has neither first nor second category.

Poincaré recurrence. Let X be a finite measure space, and let $T : X \rightarrow X$ be a measure-preserving automorphism. Then for any set A of positive measure, there exists an $n > 0$ such that $m(A \cap T^n(A)) > 0$.

Proof. Let $E = T(A) \cup T^2(A) \cup \dots$ be the strict forward orbits of the elements of A . Then, if A is disjoint from its forward orbit, we find A and E are disjoint sets and $T(A \cup E) = E$. Thus $m(A \cup E) = m(E) = m(E) + m(A)$, so $m(A) = 0$. ■

Recurrence and category. Now suppose X is also a compact metric space, $T : X \rightarrow X$ is a measure-preserving homeomorphism, and every nonempty open set has positive measure. We say $x \in X$ is *recurrent* if x is an accumulation point of the sequence $T^n(x), n > 0$.

Theorem 8.21 *The set of recurrent points is residual and of full measure.*

Proof. If x is not a recurrent point, then there is a positive distance from x to the closure of its forward orbit. That is, for some $r > 0$ we have

$$x \in E_r = \{y : d(y, T^n(y)) \geq r, \forall n > 0\}.$$

Note that E_r is closed, and hence compact. We claim $m(E_r) = 0$. If not, there is a ball such that $A = B(x, r/2) \cap E_r$ has positive measure. But then A is disjoint from its forward orbit, contrary to Poincaré recurrence.

Thus E_r is a closed set of measure zero, and hence nowhere dense. Since the non-recurrent points are exactly the set $\bigcup_i E_{1/i}$, we see the recurrent points are residual and of full measure. ■

The space of homeomorphisms. Let X be a compact metric space. Let us make the space $C(X, X)$ of all continuous maps $f : X \rightarrow X$ into a complete metric space by $d(f, g) = \sup d(f(x), g(x))$. What can we say about the subset $H(X)$ of homeomorphisms?

It is easy to see $H([0, 1])$ is already neither open nor closed. However it does consist exactly of the bijective maps in $C(X, X)$. Now surjectivity is a closed condition, and hence a G_δ -condition. What about injectivity? If f is not injective, then there are two points at definite distance, x and y , that are identified. Thus the non-injective maps are a union of closed sets,

$$\bigcup_n \{f : \exists x, y \in X, d(x, y) \geq 1/n, f(x) = f(y)\}.$$

(The closedness uses compactness of X). Putting these observations together we have:

Theorem 8.22 *For any compact metric space X , the homeomorphisms $H(X)$ are a G_δ subset of the complete metric space $C(X, X)$.*

The property of Baire. A topological space has the property of Baire if it satisfies the *conclusion* of the category theorem: namely if any intersection of dense G_δ 's is still dense.

Theorem 8.23 *If X is complete and $Y \subset X$ is a G_δ , then Y has the property of Baire.*

Proof. Apply the Baire category theorem to \bar{Y} , in which Y is a dense G_δ . ■

Alternatively, one can give $H(X)$ a different metric so its becomes complete: namely one can take $H(X)$ to be the subspace of $C(X) \times C(X)$ consisting of pairs (f, g) such that $f \circ g(x) = g \circ f(x) = x$.

Transitive maps of the square. Oxtoby and Ulam proved that a generic measure-preserving automorphism of any manifold is ergodic. We will prove a weaker result that illustrates the method.

Theorem 8.24 *There exists a homeomorphism of $[0, 1] \times [0, 1]$ with a dense orbit.*

Let $X = [0, 1] \times [0, 1]$. Since $H(X)$ has the property of Baire, it would suffice to show that a generic homeomorphism has a dense orbit. But this

is not true! Once there is a disk such that $f(D) \subset D$, any orbit that enters D can never escape. And in fact any homeomorphism can be perturbed slightly so that $f^n(D) \subset D$ for some disk D . The category method has failed!

Measure-preserving maps. What Oxtoby and Ulam proved is that the problem can be solved by making it harder.

Theorem 8.25 *A generic measure preserving homeomorphism of the square has a dense orbit.*

Proof. Let $M(X) \subset H(X)$ be the measure-preserving homeomorphisms. It also has the property of Baire, because it is a closed subset of $H(X)$.

Consider two balls, B_1 and B_2 , and let $U(B_1, B_2) \subset M(X)$ be the set of $f : X \rightarrow X$ such that $f^n(B_1)$ meets B_2 , for some $n > 0$. Clearly $U(B_1, B_2)$ is open; we will show it is dense.

To this end, fix $r > 0$, and consider any $f \in M(X)$. Choose a chain of nearly equally spaced points along a straight line in the square, x_0, \dots, x_n with $x_0 \in B_1$, $x_n \in B_2$, and $d(x_i, x_{i+1}) \asymp r$. Since a generic point is recurrent, after a slight perturbation we can also assume that x_i is recurrent. Thus we can also find high iterates $y_i = T^{n_i}(x_i)$ such that $d(x_i, y_i) \ll r$.

Now choose a sequence of disjoint short paths P_i (of length less than $2r$) from y_i to x_{i+1} , avoiding all other of the points we have considered; including $T^n(x_i)$ for $0 \leq n \leq n_i$. Construct a measure-preserving map within distance r of the identity, such that $g(y_i) = x_{i+1}$. This map g is supported close to $\bigcup P_i$.

Then $g \circ f$, under iteration, moves x_0 to y_{n_0-1} , then to $g(f(y_{n_0-1})) = g(y_{n_0}) = x_1$, and then x_1 to x_2 , etc.; so that ultimately x_n is in the forward orbit of x_0 , and hence f moves B_1 into B_2 .

Using a countable base for X , we can now conclude that a generic $f \in M(X)$ has the property that for any two nonempty open sets $U, V \subset X$, there exists an $n > 0$ such that $f^n(U) \cap V \neq \emptyset$.

We claim any such f has a dense orbit. Indeed, consider for any open ball B the set $U(B)$ of x such that $f^n(x) \in B$ for some $n > 0$. The set $U(B)$ is open, and it is dense by our assumption on f . Intersecting these $U(B)$ over a countable base for X , we find a *generic* $x \in X$ has a dense orbit. ■

Open problem. Does a *generic* C^1 diffeomorphism of a surface have a dense orbit? It is known that a sufficiently smooth diffeomorphism does *not* (KAM theory).

For more discussion, see [Ox, §18] and [Me, Thm. 4.3].

9 General Topology

Topological spaces. The collection of open sets \mathcal{T} satisfies: $X, \emptyset \in \mathcal{T}$; and finite intersections and arbitrary unions of open sets are open. Metric spaces give particular examples.

Compactness. A space is compact if every open cover has a finite sub-cover. Equivalent, any collection of closed sets with the finite intersection property has a nonempty intersection.

Theorem 9.1 . *A subset $K \subset \mathbb{R}^n$ is compact iff K is closed and bounded.*

Theorem 9.2 *A metric space (X, d) is compact iff every sequence has a convergent subsequence.*

The first result does not hold in general metric spaces: for example, the unit ball in $\ell^\infty(\mathbb{N})$ is closed and bounded but not compact. Similarly, the sequence of functions $f_n(x) = x^n$ is bounded in $C[0, 1]$, but has no convergent subsequence.

The second, we will also see, does not hold in general topological spaces. Nevertheless both results can be modified so they hold in a general setting.

Total boundedness. A metric space is *totally bounded* if for any $r > 0$, there exists a covering of X by a finite number of r -balls. In \mathbb{R}^n , boundedness and total boundedness are equivalent; but the latter notion is much stronger in infinite-dimensional spaces, and gives the correct generalization of Theorem 9.1.

Theorem 9.3 *A metric space (X, d) is compact iff X is complete and totally bounded.*

Arzela-Ascoli. Here is application of compactness to function spaces.

Let $C(X)$ be the Banach space of continuous functions on a compact metric space (X, d) . When does a set of functions $\mathcal{F} \subset C(X)$ have compact closure? That is, when can we assure that every sequence $f_n \in \mathcal{F}$ has a convergent subsequence (whose limit may or may not lie in \mathcal{F})?

Recall that $C(X)$ is complete, and that a metric space is compact iff it is complete and totally bounded. The latter property means that for any $r > 0$ there is a finite covering of X by r -balls.

The set \mathcal{F} is *equicontinuous* if all the functions satisfy the same modulus of continuity: that is, if there is a function $m(s) \rightarrow 0$ as $s \rightarrow 0$ such that $d(x, y) < s$ implies $|f(x) - f(y)| < m(s)$ for all $f \in \mathcal{F}$. Of course \mathcal{F} is *bounded* iff there is an M such that $|f(x)| \leq M$ for all $f \in \mathcal{F}$.

Theorem 9.4 $\mathcal{F} \subset C(X)$ has compact closure iff \mathcal{F} is bounded and equicontinuous.

Proof. First suppose $\overline{\mathcal{F}}$ is compact. Then clearly \mathcal{F} is bounded. Now take any $\epsilon > 0$, and cover \mathcal{F} by a finite collection of balls $B(f_i, \epsilon/3)$. Since X is compact, each f_i is uniformly continuous, so there is a δ such that

$$d(x, y) < \delta \implies \forall i, |f_i(x) - f_i(y)| < \epsilon/3.$$

Then for any $f \in \mathcal{F}$, we can find f_i with $d(f, f_i) < \epsilon/3$, and conclude that $|f(x) - f(y)| < \epsilon$ when $d(x, y) < \delta$. Thus \mathcal{F} is equicontinuous.

Now suppose \mathcal{F} is bounded by M , and equicontinuous. To show $\overline{\mathcal{F}}$ is compact, we need only show \mathcal{F} is totally bounded. To this end, fix $r > 0$, and by equicontinuity choose $\delta > 0$ such that $d(x, y) < \delta \implies |f(x) - f(y)| < r$. By compactness of X , we can find a finite set $E \subset X$ such that $B(E, \delta) = X$. Similarly we can pick a finite set $F \subset [-M, M]$ that comes within r of every point.

For each map $g : E \rightarrow F$, let

$$B(g) = \{f \in \mathcal{F} : \sup_E |g - f| < r\}.$$

Since there are only finitely many maps g , and every f is close to some g , these sets give a finite cover \mathcal{F} . Finally if $f_1, f_2 \in B(g)$, then for any $x \in X$, there is an $e \in E$ within δ of x . We then have

$$|f_1(x) - f_2(x)| \leq 2r + |f_1(e) - f_2(e)| \leq 4r,$$

so $\text{diam } B(g) \leq 4r$. It follows that \mathcal{F} is totally bounded, and thus $\overline{\mathcal{F}}$ is compact. ■

Example: Normal families. Let \mathcal{F} be the set of all analytic functions on an open set $\Omega \subset \mathbb{C}$ with $|f(z)| \leq M$. Then \mathcal{F} is compact in the topology of uniform convergence on compact sets.

Note: The functions $f_n(z) = z^n$ do *not* converge uniformly on the whole disk, so the restriction to compacta is necessary.

Proof. By Cauchy's theorem, if $d(z, \partial\Omega) > r$, then

$$|f'(z)| = \left| \frac{1}{2\pi i} \int_{S^1(p,r)} \frac{f(\zeta) d\zeta}{(\zeta - z)^2} \right| \leq \frac{2\pi r M}{2\pi r^2} = \frac{M}{r}.$$

Thus we can pass to a subsequence converging uniformly on the compact set $K_r = \{z \in \Omega : d(z, \partial\Omega) \geq r\}$. Diagonalizing, we get a subsequence converging uniformly on compact sets. Analyticity is preserved in the limit, so \mathcal{F} is a normal family. ■

Metrizability, Topology and Separation. Our next goal is to formulate purely topological versions of the best properties of metric spaces. This properties will help us recognize when a topological space (X, \mathcal{T}) is *metrizable*, i.e. when there is a metric d that determines the topology \mathcal{T} .

Given any collection \mathcal{C} of subsets of X , there is always a weakest topology \mathcal{T} containing that collection. We say \mathcal{C} *generates* \mathcal{T} .

A *base* \mathcal{B} for a topology is a collection of open sets such that for each $x \in U \in \mathcal{T}$, there is a $B \in \mathcal{B}$ with $x \in B \subset U$. Then U is the union of all the B it contains, so \mathcal{B} generates \mathcal{T} . Indeed \mathcal{T} is just the union of the empty set and all unions of subsets of \mathcal{B} .

If \mathcal{B} is given, it is a base for some topology iff for any $x \in B_1, B_2$ there is a B_3 with $x \in B_3 \subset B_1 \cap B_2$.

A *sub-base* \mathcal{B} for a topology \mathcal{T} is a collection of open sets such that for any $x \in U \in \mathcal{T}$, we have $x \in B_1 \cap \dots \cap B_n \subset U$ for some $B_1, \dots, B_n \in \mathcal{B}$. Any sub-base also generates \mathcal{T} . Conversely, any collection of set \mathcal{B} covering X forms a sub-base for the topology \mathcal{T} it generates.

Example: In \mathbb{R}^n , the open half-spaces $H = \phi^{-1}(a, \infty)$ for linear functions $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ form a sub-base for the topology. (By intersecting them we can make small cubes).

A *base at x* is a collection of open sets \mathcal{B}_x , all containing x , such that for any open U with $x \in U$, there is a $B \in \mathcal{B}_x$ with $x \in B \subset U$.

Example: in any metric space, the balls $B(x, 1/n)$ form a base at x .

Countability axioms. A topological space X is:

- *first countable* if every point has a countable base;
- *second countable* if there is a countable (sub-)base for the whole space; and
- *separable* if there is a countable dense set $S \subset X$.

Clearly a second countable space is separable: just choose one point from each open set.

Examples. Clearly any metric space is first countable.

A Euclidean spaces \mathbb{R}^n are first and second countable, and separable.

The space $\ell^\infty(\mathbb{N})$ is *not* separable or second countable. The uncountable collection of balls $B(\chi_A, 1/2)$, as A ranges over all subsets of \mathbb{N} , are disjoint. On the other hand, $\ell^p(\mathbb{N})$ is separable for $1 \leq p < \infty$.

Theorem 9.5 *Any separable metric space is second countable.*

Proof. Let (x_i) be a countable dense set of let $\mathcal{B} = \{B(x_i, 1/n)\}$. Then if $x \in U$, we have $x \in B(x, r) \subset U$, and hence $x \in B(x_i, 1/n) \subset U$ as soon as $d(x_i, x) < 1/n$ and $2/n < r$. ■

Theorem 9.6 *The number of open (or closed) sets in a separable metric space (like \mathbb{R}^n) is at most $|\mathbb{R}|$.*

Proof. $|\mathcal{T}| \leq |\mathcal{P}(\mathcal{B})| \leq |\mathcal{P}(\mathbb{N})| = |\mathbb{R}|$.

Corollary 9.7 *There are more subsets of \mathbb{R} than there are closed subsets.*

Question. Does first countable and separable imply second countable?
No!

Example. The half-open interval topology. Let

$$\mathcal{B} = \{[a, b) : a < b\};$$

this is a base for a topology \mathcal{T} on \mathbb{R} . In this topology, $x_n \rightarrow y$ iff x_n approaches y from above. Thus every strictly increasing sequence diverges.

This space is first countable and separable. (The rationals are dense.) But it is *not* second countable! If $a \in B \subset [a, b)$, then a must be the minimum of B . Thus for any base \mathcal{B} , the map $B \mapsto \inf B$ sends \mathcal{B} onto \mathbb{R} , and therefore $|\mathcal{B}| \geq |\mathbb{R}|$.

Cor: $(\mathbb{R}, \mathcal{T})$ is not metrizable.

This space is sometimes denoted \mathbb{R}_ℓ ; for an extended discussion, see Munkres, *Topology*.

The Lindelöf condition. A topological space is said to be *Lindelöf* if every open cover has a *countable* subcover. A second countable space is Lindelöf. The space \mathbb{R}_ℓ above is also Lindelöf, but not second countable. It is interesting to note that the Sorgenfrey (carefree?) plane, $\mathbb{R}_\ell \times \mathbb{R}_\ell$ is **not** Lindelöf (cf. Munkres, *Topology*, p. 193).

Separation axioms \mathbf{T}_i . (T for Tychonoff). Let us say disjoint subsets E, F of a topological space X can be *separated* if they lie in disjoint open sets. The separation axioms (or properties) are:

\mathbf{T}_1 (Tychonoff): Points are closed.

\mathbf{T}_2 (Hausdorff): Pairs of points x, y are separated.

\mathbf{T}_3 (regular): Points are separated from closed sets, and points are closed.

\mathbf{T}_4 (normal): Pairs of closed sets are separated, and points are closed.

Example: any metric space is normal. Given two closed sets A and B , they are separated by the open sets $U = \{x : d(x, A) < d(x, B)\}$ and $V = \{x : d(x, B) < d(x, A)\}$.

Zariski topology. Let k be a field. A natural example of a topology that is not Hausdorff is the Zariski topology on k^n . In this topology, a set is *F* closed if it is defined by system of polynomial equations: F is the zero set of a collection of polynomials $f_\alpha \in k[x_1, \dots, x_n]$.

A base for the topology consists of complements of hypersurfaces, $U_f = k^n - Z(f)$. Note that $U_f \cap U_g = U_{fg}$, so we indeed have a base.

By the Noetherian property, the ideal (f_α) is finitely generated, so only a finite number of polynomials are actually necessary to define F . Geometrically, this means any decreasing sequence of closed sets, $F_1 \supset F_2 \supset F_3 \dots$, eventually stabilizes. In particular, \mathbb{R}^2 is **compact**.

On \mathbb{R} , the Zariski topology is the cofinite topology. On \mathbb{R}^n , any two nonempty open sets meet; i.e. \mathbb{R}^n cannot be covered by a finite number of hypersurfaces. Thus the Zariski topology is T_1 but not T_2 .

The spectrum of a ring. Given a ring A , one also defines the Zariski topology on the set $\text{Spec } A$ of all prime ideals $p \subset A$, by taking the closed sets to have the form $V(a) = \{p : p \supset a\}$, where a ranges over all ideals in A . A point $p \in \text{Spec } A$ is closed iff p is a maximal ideal.

Thus $\text{Spec } A$ is usually not even Hausdorff. In fact, for any ring A , the ‘generic point’ p coming from the ideal (0) is dense; its closure is the whole space.

Theorem 9.8 *A compact Hausdorff space X is normal.*

Proof. We first show X is regular. Let p be a point outside a closed set F . Then for each $x \in F$ there are disjoint open sets $x \in U_x$ and $p \in V_x$. Passing to a finite subcover of F , we have $F \subset \bigcup_1^n U_i$ and $p \in \bigcap_1^n V_i$.

Now to prove normality, suppose E and F are disjoint closed sets. Then for each $x \in E$, there is are disjoint open sets with $x \in U_x$ and $F \subset V_x$. Passing to a finite subcover, we have $E \subset \bigcup_1^n U_i$ and $F \subset \bigcap_1^n V_i$. ■

Theorem 9.9 (Urysohn’s Lemma) *Let A, B be disjoint closed subsets of a normal space X . Then there exists a continuous function $f : X \rightarrow [0, 1]$ such that $f(A) = 0$ and $f(B) = 1$.*

Proof. Let $U_0 = A$ and let $U_1 = X$. The closed set A is a subset of the open set \tilde{B} . By normality, there exists an open set $U_{1/2}$ with $A \subset U_{1/2} \subset$

$\overline{U_{1/2}} \subset \widetilde{B}$. Iterating this construction, we obtain a family of open sets U_r indexed by the dyadic rationals in $[0, 1]$ such that $U_r \subset \overline{U_r} \subset U_s$ whenever $r < s$. Now let $f(x) = \inf\{r : x \in U_r\}$. Then $\{x : f(x) < s\} = \bigcup_{r < s} U_r$ is open, and $\{x : f(x) \leq s\} = \bigcap_{r > s} \overline{U_r}$ is closed, so f is continuous. ■

Corollary 9.10 *In a normal space, there are sufficiently many functions $f : X \rightarrow \mathbb{R}$ to generate the topology on X .*

Proof. We must show that every closed set A is the intersection of level sets of functions. But for any $p \notin A$ we can find a function with $f(A) = 0$, $f(p) = 1$, and so we are done. ■

Theorem 9.11 (Tietze Extension) *If X is normal and $A \subset X$ is closed, then every continuous function $f : A \rightarrow \mathbb{R}$ extends to a continuous function on X .*

Actually Tietze generalizes Urysohn, since the obvious function $f : A \sqcup B \rightarrow \{0, 1\}$ is continuous and $A \sqcup B$ is closed.

Approximating sets by submanifolds. For any compact set $X \subset \mathbb{R}^n$, and $r > 0$, there exists a smooth compact submanifold lying within $B(X, r)$ and separating X from ∞ .

Proof. Smooth the function given by Tietze and apply Sard's theorem.

Cor. Any compact set in \mathbb{R}^2 can be surrounded by a finite number of smooth loops. Any Cantor set in \mathbb{R}^3 can be surrounded by smooth closed surfaces; but their genus may tend to infinity! (Antoine's necklace).

Weak topology. Given a collection of functions \mathcal{F} on a set X , we can consider the weakest topology which makes all $f \in \mathcal{F}$ continuous. A base for this topology is given by the sets of the form

$$f_1^{-1}(\alpha_1, \beta_1) \cap f_2^{-1}(\alpha_2, \beta_2) \cap \dots \cap f_n^{-1}(\alpha_n, \beta_n),$$

where $f_i \in \mathcal{F}$. We have $x_n \rightarrow x$ iff $f(x_n) \rightarrow f(x)$ for all $f \in \mathcal{F}$.

The *weak topology* on a Banach space X is the weakest topology making all $\phi \in X^*$ continuous. For example, $f_n \rightarrow f$ weakly in L^1 iff

$$\int f_n g \rightarrow \int f g$$

for all $g \in L^\infty$. This topology is weaker than norm convergence; e.g. the functions $f_n(x) = \sin(nx)$ converge weakly to zero in $L^1[0, 1]$, but they do not converge at all in the norm topology.

Products. Given any collection of topological spaces $\{X_\alpha\}$, the product $X = \prod X_\alpha$ can be endowed with the Tychonoff topology, defined by the sub-basic sets $B(U, \alpha) = \{x \in X : x_\alpha \in U\}$ where $U \subset X_\alpha$ is open.

This is the weakest topology such that all the projections $f_\alpha : X \rightarrow X_\alpha$ are continuous.

Example: For any set A , \mathbb{R}^A is the set of all functions $f : A \rightarrow \mathbb{R}$, and $f_n \rightarrow f$ in the Tychonoff topology iff $f_n(a) \rightarrow f(a)$ for all $a \in A$. So the Tychonoff topology is sometimes called the topology of pointwise convergence.

Example: In $X = (\mathbb{Z}/2)^A \cong \mathcal{P}(A)$, we have $A_n \rightarrow A$ iff $(x \in A \text{ iff } x \in A_n \text{ for all } n \gg 0)$.

Theorem 9.12 *If X_i is metrizable for $i = 1, 2, 3, \dots$, then so is $\prod_1^\infty X_i$.*

Proof. Replacing the metric d_i by $\min(d_i, 1)$, we can assume each X_i has diameter at most 1. Then $d(x, y) = \sum 2^{-i}d(x_i, y_i)$ metrizes $\prod X_i$. ■

For example, $\mathbb{R}^{\mathbb{N}}$ is metrizable.

Theorem 9.13 (Urysohn metrization theorem) *A second countable topological space X is metrizable iff X is normal.*

Proof. Clearly a metric space is normal. For the converse, let (B_i) be a countable base for X . For each pair with $\overline{B_i} \subset B_j$, construct a continuous function $f_{ij} : X \rightarrow [0, 1]$ with $f_{ij} = 0$ on B_i and $f_{ij} = 1$ outside B_j . Let \mathcal{F} be the collection of all such functions, and consider the natural continuous map $f : X \rightarrow [0, 1]^{\mathcal{F}}$, sending x to $(f_{ij}(x))$. Since \mathcal{F} is countable, $f(X)$ is metrizable; we need only show that the inverse map $f(X) \rightarrow X$ is defined and continuous.

To see the map $f(X) \rightarrow X$ is defined, we must show f is injective. But given any points $x \neq y$, we can find open sets with $x \in \overline{B_i} \subset B_j$ and y outside B_j ; then f_{ij} separates x from y .

To see $f(X) \rightarrow X$ is continuous, we just need to show that the weakest topology \mathcal{T}' making all the functions f_{ij} continuous is the original topology \mathcal{T} on X . But if $x \in U \in \mathcal{T}$, then there are basis elements with $x \in \overline{B_i} \subset B_j \subset U$. Then $V = f_{ij}^{-1}[0, 1/2)$ is in \mathcal{T}' , and we have $x \in B_i \subset V \subset B_j \subset U$. Since this holds for every $x \in U$, we conclude that $U \in \mathcal{T}'$ and thus $\mathcal{T} = \mathcal{T}'$. ■

Regularity v. Normality. Tychonoff observed that Urysohn's metrization theorem also applies to regular spaces, since we have:

Theorem 9.14 *A regular space with a countable base is normal.*

Proof. Let A, B be disjoint closed sets in such a space. Then A is covered by a countable collection of open sets U_i whose closures are disjoint from B . There is a similar cover V_i of B by open sets whose closures are disjoint from A . Now set $U'_i = U_i - (\overline{V_1} \cup \dots \cup \overline{V_i})$, set $V'_i = V_i - (\overline{U_1} \cup \dots \cup \overline{U_i})$, and observe that $U = \bigcup U'_i$ and $V = \bigcup V'_i$ are disjoint open sets containing A and B . ■

A non-metrizable product. Example: $(\mathbb{Z}/2)^\mathbb{R} \cong \mathcal{P}(\mathbb{R})$ is not metrizable because it is not first countable.

A base at the set \mathbb{R} consists of the open sets $U(F)$, defined for each finite set $F \subset \mathbb{R}$ as

$$U(F) = \{A \subset \mathbb{R} : F \subset A\}.$$

Let \mathcal{F} be the set of finite subset $A \subset \mathbb{R}$. Then \mathcal{F} meets every $U(F)$ so $\mathbb{R} \in \overline{\mathcal{F}}$. But there is no sequence $A_n \in \mathcal{F}$ such that $A_n \rightarrow \mathbb{R}$! Indeed, if $A_n \in \mathcal{F}$ is given then we can pick $x \notin \bigcup A_n$, and A_n never enters the neighborhood $U(\{x\})$ of \mathbb{R} .

We will later see that $\mathcal{P}(\mathbb{R})$ is compact. But it has sequences with no convergent subsequences! To see this, let A_n be the set of real numbers $x = 0.x_1x_2x_3\dots$ such that $x_n = 1$. Given any subsequence n_k , we can find an x such that x_{n_k} alternates between 1 and 2 as $n \rightarrow \infty$. Suppose $A_{n_k} \rightarrow B$. If $x \in B$ then $x \in A_{n_k}$ for all $k \gg 0$, and if $x \notin B$ then $x \notin A_{n_k}$ for all $k \gg 0$. Either way we have a contradiction.

Nets. A *directed system* A is a partially ordered set so any two $\alpha, \beta \in A$ are dominated by some $\gamma \in A$: $\gamma \geq \alpha$ and $\gamma \geq \beta$.

A *net* x_α is a map $x : A \rightarrow X$ from a directed system into a topological space X .

Example: \mathbb{N} is a directed system, and a sequence x_n is a net.

Convergence. We say $x_\alpha \rightarrow x \in X$ iff for any neighborhood U of x there is an α such that $x_\beta \in U$ for all $\beta > \alpha$.

Theorem 9.15 *In any topological space, $x \in \overline{E}$ iff there is a net $x_\alpha \in E$ converging to x .*

Proof. Let $\alpha = \alpha(U)$ range over the directed set of neighborhoods of x in X , and for each U let x_α be an element of $U \cap E$. Then $x_\alpha \rightarrow x$.

Conversely, if $x_\alpha \in E$ converges to x , then every neighborhood of x meets E , so $x \in \overline{E}$. ■

Subnets. If B is also a directed system, a map $f : B \rightarrow A$ is *cofinal* if for any $\alpha_0 \in A$ there is a $\beta_0 \in B$ such that $f(\beta) \geq \alpha_0$ whenever $\beta \geq \beta_0$. Then $y_\beta = x_{f(\beta)}$ is a *subnet* of x_α .

Example: A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is cofinal iff $f(n) \rightarrow \infty$. So subsequences are special cases of subnets.

Theorem 9.16 *X is compact iff every net has a convergent subnet.*

Proof. Let \mathcal{F} be a collection of closed sets with the finite intersection property, and let α be the directed system of finite subsets of \mathcal{F} , and let x_α be a point lying in their common intersection. Then the limit point y of a convergent subnet of x_α will lie in every element of \mathcal{F} , so $\bigcap \mathcal{F} \neq \emptyset$.

Conversely, let x_α be a net in a compact space X . For every α let

$$F_\alpha = \{x_\beta : \beta \geq \alpha\}.$$

Since the index set is directed, any finite set of indices has an upper bound, and thus the $\overline{F_\alpha}$ have the finite intersection property. Therefore there is a y in $\bigcap \overline{F_\alpha}$.

Now let B be a base at Y ordered by inclusion, and let $C = A \times B$ with the product ordering. (This means $(a, b) < (a', b')$ iff $a < a'$ and $b < b'$.) Then the projection $A \times B \rightarrow A$ is cofinal.

For every pair $\gamma = (\alpha, \beta(U))$ there is an element $y_\gamma = x_{f(\gamma)} \in U \cap F_\alpha$. Then y_γ is a subnet converging to y . ■

Theorem 9.17 (Tychonoff) *A product $X = \prod_N X_n$ of compact sets is compact. (Here N is an arbitrary index set).*

Proof. By the Axiom of Choice we may assume the index set is an ordinal $N = \{0, 1, 2, \dots\}$. Given a net $x_\alpha \in X$, we will produce a convergent subnet y_α , by transfinite induction over N . In the process we will define nets x^n for each $n \in N$, with x^n a subnet of x^i for $i < n$, and with each coordinate $x^n(i)$ converging for $i < n$. We will have $f_{ij} : A_i \rightarrow A_j$ denote the re-indexing function for $i \geq j$.

Let $y^0 = x$. Passing to a subnet, we obtain a net x_α^0 indexed by $\alpha \in A_0$ and with $x_\alpha^0(0)$ converging in X_0 .

Given $n \in N+1$, let $B_n = \sqcup_{i < n} A_i$, and let $y^n(\alpha) = x_\alpha^i$ for $\alpha \in A_i$. Order A_n by $\alpha \leq \beta$ if α and β belong to A_i and A_j with $i \leq j$, and if $f_{ji}(\beta) \geq \alpha$. Finally to make y^n a subnet of x^i , let $g_{nj}(\alpha) = f_{ij}(\alpha)$ if $\alpha \in A_i$, $i \geq j$, and specify $g_{nj}(\alpha) \in A_j$ arbitrarily if $\alpha \in A_i$ for $i < j$.

(Check that this is a subnet: given $\alpha_0 \in A_i$, if $\beta \geq \alpha_0$, then $\beta \in A_j$ for some $j \geq i$, and by definition of the ordering on B_n we have $f_{ji}(\beta) \geq \alpha_0$, so $g_{ni}(\beta) \geq \alpha_0$.)

Since y^n is a subnet of x^i , the net $y_\alpha^n(i)$ converges for all indices $i < n$. Let (x^n, A_n) be a subnet of (y^n, B_n) that converges in position n .

By induction we obtain, for the ordinal $N+1$, a subnet $y_\alpha = y_\alpha^{N+1}$ that converges in all coordinates. This means y_α converges in X . ■

Axiom of Choice. The use of the Axiom of Choice in the preceding proof cannot be dispensed with, in the strong sense that Tychonoff's theorem *implies* the Axiom of Choice. Note that this is stronger than the commonly-heard statement 'you need the Axiom of Choice to construct a non-measurable set'.

Partitions of unity.

Theorem 9.18 *Let X be a compact Hausdorff space, and let \mathcal{U} be an open cover of X . Then there is a finite subcover U_i and functions $0 \leq f_i(x) \leq 1$ supported on U_i such that $\sum_1^n f_i(x) = 1$.*

Proof. For each $x \in X$ there is an open set $U \in \mathcal{U}$ and a continuous function $f \geq 0$ supported in U , such that $f(x) = 1$. By compactness there is a finite set of such functions such that the open sets $\{x : f_i(x) > 0\}$ cover X . Then $g(x) = \sum f_i(x) > 0$ at every point; replacing $f_i(x)$ by $f_i(x)/g(x)$ gives the desired result. ■

Lebesgue number. Corollary. Given an open covering \mathcal{U} of a compact metric space X , there is an $r > 0$ such that for every $x \in X$, there is a $U \in \mathcal{U}$ with $B(x, r) \subset U$. The number r is called the *Lebesgue number* of \mathcal{U} .

Proof. Construct a partition of unity subordinate to $U_1, \dots, U_n \in \mathcal{U}$; then for every x there is an i such that $f_i(x) \geq 1/n$, and by uniform continuity of the functions f_i there is an $r > 0$ such that $f_i(x) > 0$ on $B(x, r)$; then $B(x, r) \subset \{f_i > 0\} \subset U_i \in \mathcal{U}$. ■

Local constructions.

Theorem 9.19 *Any compact manifold X admits a metric.*

Proof. Take a finite collection of charts $\phi : U_i \rightarrow \mathbb{R}^n$, a partition of unity f_i subordinate to U_i , and let $g(v) = \sum f_i |D\phi_i(v)|^2$. ■

Maximal ideals in $C(X)$.

Theorem 9.20 *Let X be a compact Hausdorff space; then the maximal ideals in the algebra $C(X)$ correspond to the point evaluations.*

Proof. Let $I \subset C(X)$ be a (proper) ideal. Suppose for all $x \in C(X)$, I is not contained in the maximal ideal M_x of functions vanishing at x . Then we can find for each x a function $f \in I$ not vanishing on a neighborhood of x . By compactness, we obtain $g = f_1^2 + \cdots + f_n^2$ vanishing nowhere. Then $(1/g)g \in I$ so $I = C(X)$, contradiction. So I is contained in some M_x . ■

Spectrum. Given an algebra A over \mathbb{R} , let

$$\sigma(f) = \{\lambda \in \mathbb{R} : \lambda + f \text{ has no inverse in } A\}.$$

Then for $A = C(X)$, we have $\sigma(f) = f(X)$, and thus we can reconstruct $\|f\|_\infty$ from the algebraic structure on A .

Also for $A = C(X)$ we can let Y be the set of multiplicative linear functionals, and embed Y into \mathbb{R}^A by sending ϕ to the sequence $(\phi(f)) : f \in A$. Then in fact $\phi(f) \in [-\|f\|, \|f\|]$, so Y is compact, and Y is homeomorphic to X .

Local compactness. A topological space X is **locally compact** if the open sets U such that \bar{U} is compact form a base for the topology.

For example, \mathbb{R}^n is locally compact.

Alexandroff compactification. Let X be a locally compact Hausdorff space, and let $X^* = X \cup \{\infty\}$, and define a neighborhood base at infinity by taking the complements $X^* - K$ of all compact sets $K \subset X$.

Theorem 9.21 *X^* is a compact Hausdorff space, and the inclusion of X into X^* is a homeomorphism.*

Proof. Compact: if you cover X^* , once you've covered the point at infinity, only a compact set is left. Hausdorff: because of local compactness, every $x \in X$ is contained in a U such that \bar{U} is compact, and hence $V = X^* - \bar{U}$ is a disjoint neighborhood of infinity. ■

This space is called the *one-point compactification* of X .

Examples: \mathbb{N}^* ; $S^n = (\mathbb{R}^n)^*$.

Proper maps. A useful counterpart to local compactness is the notion of a *proper map*. A map $f : X \rightarrow Y$ is **proper** if $f^{-1}(K)$ is compact whenever K is compact. Intuitively, if x_α leaves compact sets in X , then $f(x_\alpha)$ leaves compact sets in $f(X)$. Thus $x_\alpha \rightarrow \infty$ implies $f(x_\alpha) \rightarrow \infty$, and so f extends to a continuous map from X^* to Y^* . This shows:

Theorem 9.22 *A continuous bijection between locally compact Hausdorff spaces is a homeomorphism iff it is proper.*

Example: There is a bijective continuous map $f : \mathbb{R} \rightarrow S^1 \cup [1, \infty) \subset \mathbb{C}$.

The Stone-Čech compactification.

Theorem 9.23 *Let X be a normal space. Then there is a unique compact Hausdorff space $\beta(X)$ such that:*

1. X is dense in $\beta(X)$;
2. Every bounded continuous $f : X \rightarrow \mathbb{R}$ extends to a continuous function on $\beta(X)$;
3. If X is compactified by another Hausdorff space Y , in the sense that the inclusion $X \subset Y$ is dense, then $\beta(X)$ is bigger than Y : there is a continuous map $\phi : \beta(X) \rightarrow Y$.

Proof. Let \mathcal{F} be the family of all continuous $f : X \rightarrow [0, 1]$, let Z be the compact product $[0, 1]^{\mathcal{F}}$, and let $\beta(X) \subset Z$ be the closure of X under the embedding $x \mapsto (x_f)$ where $x_f = f(x)$. The first two properties are now evident.

Finally let Y be another compactification of X , and let \mathcal{G} be the family of all continuous maps $g : Y \rightarrow [0, 1]$. Then there is an embedding $Y \subset [0, 1]^{\mathcal{G}}$, and the inclusion $\mathcal{G} \subset \mathcal{F}$ gives a natural projection map $[0, 1]^{\mathcal{F}} \rightarrow [0, 1]^{\mathcal{G}}$. This projection sends $\beta(X)$ into Y . ■

Example: $X = \beta(\mathbb{N})$. In this space, a sequence $x_n \in \mathbb{N}$ converges iff it is eventually constant. Thus X is compact but the sequence $x_n = n$ has no convergent subsequence! (However it does have convergent subnets; for such a net, $f(n_\alpha)$ converges for every $f \in \ell^\infty(\mathbb{N})$!)

Stone-Čech and dual spaces. Another way to look at $\beta(\mathbb{N})$ is that each $n \in \mathbb{N}$ provides a map $n : \ell^\infty(\mathbb{N}) \rightarrow \mathbb{R}$ by $a \mapsto a_n$, and that $\beta(\mathbb{N})$ is the closure

of these maps. Note that the maps in the closure are **bounded, linear functionals**. A typical example is provided by the ultrafilter limit we constructed before. In general the closure consists of those finitely-additive measures on \mathbb{N} such that $\mu(\mathbb{N}) = 1$ and $\mu(E) = 0$ or 1 for all $E \subset \mathbb{N}$.

Algebras of continuous functions. We now come to the issue of *approximations* of continuous functions on X . The following central result shows that any reasonably large *algebra* inside of $C(X)$ is dense. Key applications of this result include the completeness of Fourier series.

Theorem 9.24 (Stone–Weierstrass) *Let X be a compact Hausdorff space, and let $A \subset C(X)$ be a subalgebra that contains the constants and separates points. Then A is dense in $C(X)$.*

Examples. In $C[0, 1]$, the functions of bounded variation, or Lipschitz, or C^k , or Hölder continuous, or polynomials, or those that are real-analytic, all form subalgebras that separate points and contain the constant.

In any product of compact Hausdorff spaces $X \times Y$, the linear span of the functions of the form $f(x)g(y)$ forms an algebra that separates points (by Urysohn’s lemma), and hence is dense in $C(X \times Y)$.

We now proceed to the proof.

Lemma 9.25 *The closure of A is a lattice.*

Proof. We must show that $f, g \in A \implies f \vee g \in A$, where $(f \vee g)(x) = \max(f(x), g(x))$. Note that $f \vee g$ is the average of $f + g$ and $|f - g|$. So it suffices to show $f \in A \implies |f| \in A$. Now if $\epsilon < f < 1$, then $\sqrt{f} = \sqrt{1 - (1 - f)} \in A$, because $\sqrt{1 - x} = \sum a_n x^n$ can be expanded in a power series convergent in $B(0, 1)$, and hence uniformly convergent in $B(0, 1 - \epsilon)$. Then

$$|f| = \lim_{\epsilon \rightarrow 0} \sqrt{f^2 + \epsilon},$$

so $|f|$ is in A , and hence $f \vee g$ is in A . ■

Proof of Stone–Weierstrass. As above we replace A by its closure; then A is an algebra as well as a lattice.

Given $g \in C(X)$, let $\mathcal{F} = \{f \in A : f \geq g\}$. To show $g \in A$, it suffices to show for each x that $g(x) = \inf_{\mathcal{F}} f(x)$. Indeed, if that is the case, then for any $\epsilon > 0$ and $x \in X$, there is a neighborhood U of x and an $f \in \mathcal{F}$ such that $g \leq f \leq g + \epsilon$ on U . Taking a finite sub-cover, we obtain a finite

number of functions such that $g \leq f_1 \wedge \cdots \wedge f_n \leq g + \epsilon$ on all of X . Since A is a lattice, we are done.

It remains to construct, for $\epsilon > 0$ and $x \in X$, and function $f \in A$ such that $f \geq g$ and $g(x) \leq f(x) \leq g(x) + \epsilon$. By replacing g with $ag + b$, we may assume $g(x) = 0$ and $\sup |g| \leq 1$.

Pick a neighborhood U of x on which $|g| < \epsilon$. Since A separates points, for each $y \notin U$ there is a function $h \in A$ with $h(x) = 0$, $h(y) = 2$. Taking a finite subcover of $X - U$ by balls on which $h > 1$, we obtain a function $f = h_1^2 + \cdots + h_n^2 + \epsilon$ with $f(x) = \epsilon = g(x) + \epsilon$, with $f \geq \epsilon > g$ on U , and with $f \geq 1 > g$ on $X - U$. Then $f \in \mathcal{F}$, and so $g(x) = \inf_{\mathcal{F}} f(x)$ as desired. ■

Fourier applications. The Stone-Weierstrass theorem has a generalization to algebras over \mathbb{C} : again we have density, provide A is closed under *complex conjugation*. The proof is to apply the original theorem to real and imaginary parts. Thus we have:

Theorem 9.26 *The functions $e_n(z) = z^n$, $n \in \mathbb{Z}$, span a dense subset of $C(S^1)$.*

Corollary 9.27 *These functions give an orthonormal basis for $L^2(S^1)$ with the inner product*

$$\langle f, g \rangle = (1/2\pi) \int_{S^1} f(z) \overline{g(z)} |dz|.$$

The Hardy space. The condition on complex conjugation is necessary, since the holomorphic functions form a closed subalgebra of $C(\overline{\Delta})$. Similarly the function $z^n, n \geq 0$ form an orthonormal basis for a closed subspace $H^2(S^1) \subset L^2(S^1)$, called the *Hardy space*.

Corollary 9.28 *The functions $1, \cos(nx)$ and $\sin(nx)$, $n = 1, 2, \dots$ form an orthogonal basis for $L^2[0, 2\pi]$.*

Proof. Orthogonality is readily verified; we must check completeness. By the results above, the complex functions of the form $f(z) = \sum_{-n}^n a_n z^n$ with $z = e^{inx}$ span a dense subspace of $C[0, 2\pi]$, and $\operatorname{Re} f(z)$ can be expressed in terms of the basis above. ■

Stone–Weierstrass for polynomials. Using the important idea of approximations to the identity, to be discussed in detail later (see Theorem 11.9), we can give an alternative proof of a special case of the Stone–Weierstrass theorem:

Theorem 9.29 *The polynomials $\mathbb{R}[x]$ are dense in $C[a, b]$ for any interval $[a, b]$.*

Approximate identities. We say a sequence of functions $K_n \in L^1(\mathbb{R})$ give an *approximate identity* if:

1. $K_n \geq 0$;
2. $\int K_n = 1$; and
3. $\int_U K_n \rightarrow 1$ for any neighborhood U of $x = 0$.

The *convolution* of $f \in L^1(\mathbb{R})$ with $g \in L^\infty(\mathbb{R})$ is defined by

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

Theorem 9.30 *Let $\langle K_n \rangle$ be an approximate identity and let $f \in C(\mathbb{R})$ have compact support. Then*

$$K_n * f \rightarrow f$$

uniformly on \mathbb{R} .

Proof. We will check convergence at $x = 0$, but the bound we obtain will only depend on K_n and f . This will prove uniform convergence.

Since f has compact support, it is bounded and uniformly continuous. Suppose $|f| \leq M$ and given $\epsilon > 0$, there is an $r > 0$ such that $|f(x) - f(y)| \leq \epsilon$ when $|x - y| < r$. Choose N such that $\int_{-r}^r K_n(x) dx \geq 1 - \epsilon$ for all $n \geq N$. Note that the convolution of K_n with the constant function $f(0)$ is $f(0)$ again. Thus if we set $g(x) = f(x) - f(0)$, we have $|g(x)| \leq 2M$ and

$$|(K_n * g)(0)| = |(K_n * f)(0) - f(0)|.$$

Now we have

$$\begin{aligned} |(K_n * g)(0)| &\leq \int_{-r}^r K_n(-x)|g(x)| dx + \int_{\mathbb{R} - (-r, r)} K_n(-x)|g(x)| \\ &\leq \epsilon + 2M\epsilon \end{aligned}$$

for all $n \geq N$. Since ϵ was arbitrary, this shows uniform convergence. ■

A polynomial kernel. The kernel we will use is:

$$K_n(x) = \frac{1}{I_n}(1-x^2)^n \chi_{[-1,1]},$$

where I_n is chosen so that $\int K_n = 1$. Clearly $K_n \geq 0$, and $K_n|_{[-1,1]}$ is a polynomial in x .

To evaluate I_n , we first review some definite integrals. Given integers $a, b \geq 0$, let

$$F(a, b) = \int_0^1 x^a(1-x)^b dx.$$

(The function $B(a-1, b-1)$ is usually called the beta function.) Clearly $F(a, b) = F(b, a)$ and $F(a, 0) = 1/(a+1)$. By integration by parts, one also finds that

$$F(a, b) = \frac{b}{a+1} F(a+1, b-1).$$

This allows one to inductively prove that

$$F(a, b) = \frac{a!b!}{(a+b+1)!}.$$

We now observe that

$$I_n = \int_{-1}^1 (1-x^2)^n dx = \int_0^2 x^n(2-x)^n dx = 2^{2n+1} \int_0^1 u^n(1-u)^n du,$$

using the change of variables $u = x/2$. Since $F(n, n)$ has already been evaluated above, this shows:

$$I_n = 2^{2n+1} \frac{(n!)^2}{(2n+1)!} = 2 \frac{2 \cdot 4 \cdots 2n}{1 \cdot 3 \cdots (2n+1)}.$$

What is critical for us is the fact that

$$\lim I_n^{1/n} = 1, \tag{9.1}$$

as is easily verified. This implies that $K_n(x) \rightarrow 0$ uniformly outside $(-r, r)$, for any $r > 0$, and hence K_n is an approximate identity.

Proof of Stone–Weierstrass for polynomials. It suffices to treat the case of $C[-a, a]$ where a is small (say $a < 1/2$). By the results above, we have $K_n * f \rightarrow f$ uniformly for any compactly supported $f \in C(\mathbb{R})$.

Now consider $f \in C[-a, a]$. After adding to f a linear function (which is a polynomial), we can assume $f(a) = f(-a) = 0$, and thus f extends

to a compactly supported function on all of \mathbb{R} , vanishing outside $[-a, a]$. Moreover, for $x \in [-a, a]$ we have

$$(K_n * f)(x) = \int_{-a}^a K_n(x-t)f(t) dt = I_n^{-1} \int_{-a}^a c_n(1-(x-t)^2)^{2n} f(t) dt$$

since $|x-t| \leq 2a < 1$. But the right-hand side is clearly a polynomial function of x (of degree at most $2n$), so polynomials are dense. ■

This special case can substitute for Lemma 9.25 above, since it shows:

Corollary 9.31 *Let $A \subset C(X)$ be a subalgebra with identity. Then for any $f \in \overline{A}$ and any continuous function g defined on the range of f , we have $g \circ f \in \overline{A}$.*

Proof. Approximate g by polynomials p_n , and observe that if $f \in A$ then $p_n(f) \in A$. ■

Paracompactness. For local constructions like making a metric, what's needed is not so much compactness (finiteness of coverings) as paracompactness (local finiteness). This says that any open covering has a locally finite *refinement*. Using this property one can show, for example, that any *paracompact manifold* admits a metric.

All metric spaces are paracompact (a hard theorem). However there exists a manifold which is *not* paracompact, namely the *long line*. It is obtained from the first uncountable ordinal Ω by inserting an interval between any two adjacent points, and introducing the order topology.

This space X has the amazing property that **every sequence has a convergent subsequence**. Indeed, since a sequence is countable, it is bounded above by some countable ordinal α , and (by induction) the segment $[0, \alpha]$ is homeomorphic to $[0, 1]$, hence compact.

On the other hand, X is not compact, since the open covering by all intervals of the form $[0, \alpha)$ has no finite subcover. Thus X is not metrizable. Therefore X is not paracompact.

10 Banach Spaces

The theory of Banach spaces is a combination of infinite-dimensional linear algebra and general topology. The main themes are *duality*, *convexity* and *completeness*.

The first two themes lead into the Hahn-Banach theorem, separation theorems for convex sets, weak topologies, Alaoglu's theorem, and the Krein-Milman theorem on extreme points. The third theme leads to the '3 principles of functional analysis', namely the open mapping theorem, the closed graph theorem and the uniform boundedness principle. These three results all rest on the Baire category theorem and hence make crucial use of completeness.

Continuous linear maps. Let $T : X \rightarrow Y$ be a linear map between Banach spaces. The norm of T is defined by

$$\|T\| = \sup_{x \neq 0} \frac{\|Tx\|}{\|x\|}.$$

Note: if $Y = \mathbb{R}$ we use the usual absolute value on \mathbb{R} as a norm.

A linear map is *bounded* if its norm is finite.

Theorem 10.1 *A linear map is bounded iff it is continuous.*

Proof. Clearly boundedness implies (Lipschitz) continuity. Conversely, if ϕ is continuous, then $\phi^{-1}B(0, 1)$ contains $B(0, r)$ for some $r > 0$ and then $\|\phi\| \leq 1/r$. ■

The space of operators. We let $\mathcal{B}(X, Y)$ denote the space of all bounded linear operators $T : X \rightarrow Y$ with the norm above. With this norm, $\mathcal{B}(X, Y)$ is itself a Banach space.

Construction of elements of the dual. Note that X^* is just another notation for $\mathcal{B}(X, \mathbb{R})$. The next theorem shows that X^* is quite large. Without a result of this type, it is difficult to even show that X^* is *nontrivial* for an abstract Banach space X .

Theorem 10.2 (Hahn-Banach) *Let $\phi : S \rightarrow \mathbb{R}$ be a linear map defined on a subspace $S \subset X$ in a Banach space such that $|\phi(x)| \leq M\|x\|$ for all $x \in S$. Then ϕ can be extended to a linear map on all of X with the same inequality holding.*

Proof. Using Zorn's lemma, we just need to show that any maximal such extension of ϕ is defined on all of X . So it suffices to consider the case $S \neq X$ and show that ϕ can be extended to the span of S and y where $y \in X - S$.

We may assume $M = 1$. The extension will be determined by its value $\phi(y) = z$, and the extension will continue to be bounded by $M = 1$ so long as we can insure that z is chosen so for all $s \in S$ we have:

$$-\|y + s\| \leq \phi(y + s) = z + \phi(s) \leq \|y + s\|.$$

To show such a z exists amounts to showing that for any $s, s' \in S$ we have

$$-\phi(s) - \|y + s\| \leq -\phi(s') + \|y + s'\|,$$

so that there is a number z between the sup and inf. Now this inequality is equivalent to:

$$\phi(s') - \phi(s) \leq \|y + s\| + \|y + s'\|,$$

and this one is in fact true, since

$$\phi(s' - s) \leq \|s' + y - y - s\| \leq \|s + y\| + \|s' + y\|.$$

■

Linear functionals on $L^\infty[0, 1]$. We can now show more rigorously that $L^1[0, 1]$ is not reflexive: namely take point evaluation on $C[0, 1]$, and extend it by Hahn-Banach to a linear functional ϕ on $L^\infty[0, 1]$. It is clear then $\phi|_{C[0, 1]}$ is not given by an element in $L^1[0, 1]$.

Embedding into X^{} .**

Theorem 10.3 *For any $x \in X$ there is a $\phi \in X^*$ such that $\|\phi\| = 1$ and $\phi(x) = \|x\|$.*

Proof. Define ϕ first on the line through x , then extend it by Hahn-Banach.

Corollary 10.4 *The embedding of X into X^{**} is isometric.*

Application: the dual of $L^\infty[0, 1]$. The Hahn-Banach theorem can be used to show ‘concretely’ that the dual of $L^\infty[0, 1]$ is larger than $L^1[0, 1]$. Namely, given $p \in [0, 1]$ define $\phi_p : C[0, 1] \rightarrow \mathbb{R}$ by $\phi_p(f) = f(p)$. Then we have

$$|\phi_p(f)| \leq \|f\|_\infty$$

on the subspace $C[0, 1]$. Thus ϕ_p extends (in many ways) to a linear functional $\psi : L^\infty[0, 1] \rightarrow \mathbb{R}$ with $\|\psi\| \leq 1$. But there is no $g \in L^1[0, 1]$ such that $f(p) = \int fg$ for all continuous f , so $\psi \in X^{**}$ but $\psi \notin X = L^1[0, 1]$.

L^p examples. Let $f \in L^p(\mathbb{R})$ with $\|f\|_p = 1$ and $1 < p < \infty$, there is a *unique* ϕ of norm 1 in the dual space such that $\phi(f) = 1$: namely $\phi = \text{sign}(f)|f|^{p/q}$, which satisfies

$$\phi(f) = \int f\phi = \int |f|^p = 1.$$

This reflects the ‘smoothness’ of the unit ball in L^p : there is a unique supporting hyperplane at each point.

For L^1 things are different: for example, if $\text{supp } f = [0, 1]$ there is a huge space of $\phi \in L^\infty$ such that $\|\phi\|_\infty = 1$ and $\phi(f) = 1$.

Non-example: L^∞ . Now let $f(x) = x$ in $L^\infty[0, 1]$, and suppose $\phi \in L^1[0, 1]$ has norm 1. Choose $a < 1$ such that $t = \int_0^a |\phi| > 0$. Then we have:

$$\phi(f) \leq \int_0^a x|\phi| + \int_a^1 |\phi| \leq at + (1-t) = 1 - (1-a)t < 1.$$

Thus $\phi(f)$ can never be 1!

This gives a (rather indirect) proof that $X = L^1[0, 1]$ is a proper subset of its double dual $X^{**} = L^\infty[0, 1]^*$. Indeed, there *is* a $\phi \in L^\infty[0, 1]^*$ such that $\|\phi\| = 1$ and $\phi(f) = 1$.

More non-reflexive spaces. For the little ℓ^p spaces we have the following, rather rich non-reflexive example:

$$c_0^* = \ell^1, (\ell^1)^* = \ell^\infty, (\ell^\infty)^* = m(\mathbb{Z}).$$

It turns out the last space can be identified with the space of finitely-additive measures on \mathbb{Z} .

Weak closure. The Hahn-Banach theorem implies:

Theorem 10.5 *Let $S \subset X$ be a linear subspace of a Banach space. Then S is weakly closed iff S is norm-closed.*

Proof. Any weakly closed space is norm closed. Conversely, if S is norm closed, for any $y \notin S$ we can find a linear functional $\phi : X \rightarrow \mathbb{R}$ that vanishes on S and sends y to 1, so y is not in the weak closure of S . ■

(More generally, as we will see later, any norm-closed convex set is weakly closed.)

The weak* topology. We say $\phi_\alpha \rightarrow \phi$ in the weak* topology on X^* if $\phi_\alpha(x) \rightarrow \phi(x)$ for every $x \in X$.

Example: weak closures of continuous functions. The space $C[0, 1]$ is dense in $L^\infty[0, 1]$ in the weak* topology. Indeed, if $g \in L^\infty$ then there are continuous $f_n \rightarrow g$ pointwise a.e. with $\|f_n\|_\infty \leq \|g\|_\infty$. Now for any $h \in L^1[0, 1]$ the dominated convergence theorem implies

$$\langle h, f_n \rangle = \int h f_n \rightarrow \int h g = \langle h, g \rangle.$$

On the other hand, $C[0, 1]$ is already *closed* in the weak topology, since it is norm closed.

Theorem 10.6 (Alaoglu) *The unit ball $B^* \subset X^*$ is compact in the weak* topology.*

Proof. Let B be the unit ball in X . Then there is a tautological embedding of B^* into $[-1, 1]^B$. Since linearity and boundedness are preserved under pointwise limits, the image is closed. By Tychonoff, it is compact! ■

Metrizability. Theorem. If X is separable, then the unit ball B in X^* is a compact metrizable space in the weak* topology.

Proof. Let x_n be a dense sequence in X ; then the balls

$$B = \{\phi : |\phi(x_n) - p/q| < 1/r\}$$

form a countable base. By Urysohn's metrization theorem, B is metrizable.

Example: the space of measures. Naturally $C[0, 1]$ is separable. Thus $P[0, 1]$, the space of probability measures with the weak* topology, is a compact metric space. It can be thought of as a sort of infinite-dimensional simplex; indeed the measures supported on $\leq n$ points form an $(n + 1)$ -simplex.

Banach limits.

Theorem 10.7 *There is a linear map $\text{Lim} : \ell^\infty(\mathbb{N}) \rightarrow \mathbb{R}$ such that*

$$\begin{aligned} \text{Lim}(a_n) &\geq 0 \text{ if } a_n \geq 0 \\ \text{Lim}(1) &= 1, \text{ and} \\ \text{Lim}(a_{n+1}) &= \text{Lim}(a_n). \end{aligned}$$

Note that $|\text{Lim}(a_n)| \leq \|a_n\|$ and that Lim extends the usual limit on c and agrees with the Césaro limit when that exists.

Proof. Let $\phi_N(a) = N^{-1} \sum_1^N a_n$ and let Lim be the limit point of a convergent subnet. Note that $\phi_N(a_{n+1}) - \phi_N(a_n) = O(1/N)$. ■

Stone-Čech compactification of \mathbb{N} . The unit ball B in $\ell^\infty(\mathbb{N})^*$, while compact, is not metrizable! Indeed, the integers embed via $\phi_n(a) = z_n$, but $\langle \phi_n \rangle$ has no convergent subsequence! (If ϕ_{n_k} is a subsequence, then we can choose $a \in \ell^\infty$ such that $a_{n_k} = (-1)^k$; then $\phi_{n_k}(a)$ does not converge, so ϕ_{n_k} does not converge in the weak* topology.)

The Banach-Tarski paradox . Using the same construction on \mathbb{Z} or \mathbb{Z}^n , we get finitely-additive measures by applying Lim to indicator functions. Because of these measures, you cannot cut \mathbb{Z} into a finite number of sets, move them by translation and re-assemble them to form 2 copies of \mathbb{Z} .

However, this type of re-construction *is* possible for a free group G on 2 generators!

Suppose μ is a finitely-additive invariant probability measure on G . Let $W_a, W_{a'}, W_b$ and $W_{b'}$ denote the partition of $G - \{e\}$ into reduced words beginning with a, a', b and b' . Then $a'W_a$ contains $W_a, W_b, W_{b'}$ and $\{e\}$. Since translation by a' preserves measures, we conclude that the extra sets $W_b, W_{b'}$ and $\{e\}$ have measure zero. By the same token, all the W 's have measure zero, which contradicts the assumption that $\mu(G) = 1$.

Cutting up the sun. Note that $G = a'W_a \cup W_{a'}$, and similar for W_b and $W_{b'}$. Thus we can cut G into 5 pieces, discard one of them (e), and re-assemble the other two into two copies of G .

Now embed G into $SO(3)$ by taking two random rotations. Then G acts on S^2 . Let $E \subset S^2$ be a transversal, consisting of one point from each G -orbit, so $S^2 = G \cdot E$. Now cut S^2 into pieces of the form $E_i = W_i \cdot E$, $i = 1, \dots, 4$. (There will be some S^2 left over.) Applying the left action of G to these pieces — that is, applying rotations — we can re-arrange W_1 and W_2 to form G , and so re-arrange E_1 and E_2 to form S^2 . Do the same thing with E_3 and E_4 , and we can make a second sphere!

Three basic principles of functional analysis. Let $T : X \rightarrow Y$ be a linear map between Banach spaces X and Y . Then we have:

1. *The open mapping theorem.* If T is continuous and onto, then it is open; that is, $Tx = y$ has a solution with $\|x\| \leq C\|y\|$.

Corollaries: If T is continuous and bijective, then it is an isomorphism. If X is complete with respect to two norms, and $\|x\|_1 \leq C\|x\|_2$, then a reverse inequality holds.

2. *The closed graph theorem.* If the graph of T is closed — meaning $x_n \rightarrow x, Tx_n \rightarrow y$ implies $Tx = y$ — then T is continuous.

3. *The uniform boundedness principle.* Let $\mathcal{F} \subset X^*$ satisfy that for each $x \in X$, $|f(x)| \leq M_x \|x\|$ for all $f \in \mathcal{F}$. Then there is an M such that $\|f\| \leq M$ for all $f \in \mathcal{F}$.

The same result holds if we replace X^* with $B(X, Y)$.

These principles should be compared to the following results that hold when X and Y are *compact*.

1. If $f : X \rightarrow Y$ is bijective and continuous, then f is a homeomorphism.
2. If $f : X \rightarrow Y$ has a closed graph, then f is continuous.
(Note that $f(x) = 1/x$ for $x \neq 0$, $f(0) = 0$, gives a map $f : \mathbb{R} \rightarrow \mathbb{R}$ with a closed graph that is not continuous.)
3. Let $\mathcal{F} \subset C(X)$ satisfy $|f(x)| \leq M_x$ for all $f \in \mathcal{F}$ and $x \in X$. Then there is a nonempty open set $U \subset X$ and a constant $M > 0$ such that $|f|_U \leq M$ for all $f \in \mathcal{F}$.

Open mapping theorem: Preparation. The technical heart of the open mapping theorem is captured by the following result.

Lemma 10.8 *Let $T : X \rightarrow Y$ be a bounded linear operator between Banach spaces, and let $B \subset X$ denote the open unit ball. If $\overline{T(B)}$ has nonempty interior, then T is an open mapping. (In particular, $T(B)$ is open and T is surjective.)*

Proof. Let U be a nonempty open subset of $\overline{T(B)}$. Since \overline{B} is convex and symmetric, we have $(U - U)/2 \subset \overline{T(B)}$ and hence $\overline{T(B)}$ contains a neighborhood of the origin, say $B(0, r)$.

Suppose $y \in B(0, r) \subset Y$. We will proceed to solve the equation $Tx = y$. First, we can find an x_0 with $\|x_0\| < 1$ such that $\|T(x_0) - y\|$ is as small as we like; say, less than $r/2$. Then we can find an x_1 with $\|x_1\| < 1/2$ that is nearly equal to $y - T(x_0)$; say $\|T(x_0) + T(x_1) - y\| \leq r/4$. Continuing in this way, we construct a sequence with $\|x_i\| \leq 1/2^i$ such that $T(\sum x_i) = y$ and $\|\sum x_i\| < 2$. This shows that $T(B(0, 2))$ contains $B(0, r)$. In other words, T is an open mapping at the origin. By linearity, T is an open mapping everywhere. ■

Open mapping theorem: Proof. We will prove a slightly stronger result.

Theorem 10.9 *Let $T : X \rightarrow Y$ be a bounded linear map between Banach spaces. Then either $T(X) = Y$ and T is an open map, or $T(X)$ is meager in Y .*

Proof. First suppose T is surjective. Then $T(X) = Y = \bigcup_{n=1}^{\infty} n \cdot T(B)$, where B is the unit ball in X . Since Y is complete, the Baire category theorem implies that $\overline{T(B)}$ has nonempty interior. Then by the preceding lemma, T is an open map. This is the open mapping theorem.

Now suppose T is not surjective. Then by the preceding lemma, $\overline{T(B)}$ is nowhere dense. Thus $T(X) = \bigcup_1^{\infty} nT(B)$ is contained in a countable union of closed, nowhere dense sets, so it is meager. ■

Open-mapping theorem: application. The open mapping theorem implies:

Corollary 10.10 *If X is complete in two norms, and $\|x\|_1 \leq C\|x\|_2$, then there is a C' such that $\|x\|_2 \leq C'\|x\|_1$.*

Here is a nice application due to Grothendieck.

Theorem 10.11 *Let $S \subset L^2[0, 1]$ be a closed subspace such that every $f \in S$ is continuous. Then S is finite-dimensional.*

Proof. We have $\|f\|_{\infty} \geq \|f\|_2$, so S is complete in both the L^2 and the L^{∞} norms. Thus there is an $M > 0$ such that $M\|f\|_2 \geq \|f\|_{\infty}$.

Now let f_1, \dots, f_n be an orthonormal set. Then for any $p \in [0, 1]$, we have

$$\left\| \sum f_i(p)f_i \right\|_2 = \left(\sum |f_i(p)|^2 \right)^{1/2},$$

and thus

$$M \left(\sum |f_i(p)|^2 \right)^{1/2} \geq \left\| \sum f_i(p)f_i \right\|_{\infty} \geq \sum f_i(p)^2,$$

which implies $\sum f_i(p)^2 \leq M^2$. Integrating from 0 to 1 gives $n \leq M^2$. ■

Closed graph theorem: proof. Let $|x| = \|x\| + \|Tx\|$. Now if $|x_n|$ is Cauchy, then $x_n \rightarrow x$ in X and $Tx_n \rightarrow y$ in Y ; since the graph of T is closed, we have $Tx = y$ and thus $|x_n - x| \rightarrow 0$. Thus X is complete in the $|\cdot|$ norm, so by the open mapping theorem we have $|x| \leq M\|x\|$ for some M ; thus $\|T\| \leq M$ and T is continuous. ■

Application of the closed graph theorem. A typical symmetric operator on $L^2(\mathbb{R})$ is given by $(Tf)(x) = \int K(x, y)f(y) dy$, where the kernel $K(x, y)$ is symmetric. The next result shows that if T is well-defined, it is automatically continuous.

Theorem 10.12 (Toeplitz) *Let $T : H \rightarrow H$ be a symmetric linear operator on Hilbert space, meaning $(Tx, y) = (x, Ty)$. Then T is continuous.*

Proof. Suppose $x_n \rightarrow x$ and $Tx_n \rightarrow z$. Then for all $y \in H$, we have

$$(y, z) = \lim(y, Tx_n) = \lim(Ty, x_n) = (Ty, x) = (y, Tx).$$

Thus $(y, Tx - z) = 0$ for all $y \in H$. Taking $y = Tx - z$, we find $Tx = z$. Thus T has a closed graph, and hence T is continuous. ■

Uniform boundedness theorem: proof. Let $F_M = \{x : |\phi(x)| \leq M \forall \phi \in \mathcal{F}\}$. By Baire category, some F_M contains an open ball; and by convexity and symmetric, it contains a ball $B(0, r)$ for some $r > 0$. Then for $\|x\| \leq r$ we have $|\phi(x)| \leq M$ for all $\phi \in \mathcal{F}$, and thus $\|\phi\| \leq M/r$ for all elements of \mathcal{F} . ■

This simple argument actually establishes a somewhat stronger result, which we will later use in applications to Fourier series.

Theorem 10.13 *Suppose we have a sequence of linear functionals $\phi_n \in X^*$ that satisfy $\|\phi_n\| \rightarrow \infty$. Then $\sup_n |\phi_n(x)| = \infty$ for a generic $x \in X$.*

Proof. Consider the closed set $X_M = \{x : \sup_n |\phi_n(x)| \leq M\}$. If it contained an open ball, then we would have a uniform bound on $\|\phi_n\|$. But $\|\phi_n\| \rightarrow \infty$, so X_M is also nowhere dense. Outside the meager set $\bigcup_{M=1}^{\infty} X_M$, we have $\sup_n |\phi_n(x)| = \infty$, and thus this property holds for a generic $x \in X$. ■

Example. Let $\phi_n \in X^*$ have the property that $\psi(x) = \lim \phi_n(x)$ exists for every $x \in X$. Then $\|\phi_n\| \leq M$ and hence $\psi \in X^*$. In other words:

You cannot construct an unbounded linear functional by taking a pointwise limit of bounded ones.

Unbounded linear functionals are something like non-measurable sets; they exist but all constructions of them involve the Axiom of Choice.

Remark. If a net satisfies $x_\alpha \rightarrow y$, is $\|x_\alpha\|$ necessarily bounded? No! Let α range over all finite subsets of \mathbb{N} , directed by inclusion, and let x_α be the minimum of α . Then $x_\alpha \rightarrow 0$ but $\sup x_\alpha = \infty$.

Convexity. A subset $K \subset X$ is *convex* if $x, y \in K \implies tx + (1-t)y \in K$ for all $t \in (0, 1)$.

Support.

Theorem 10.14 *Let $K \subset X$ be an open convex set not containing the origin. Then there is a $\phi \in X^*$ such that $\phi(K) > 0$.*

Proof. Geometrically, we need to find a closed, codimension-one hyperplane $H = \text{Ker } \phi \subset X$ disjoint from K . Consider all subspaces disjoint from K and let H be a maximal one (which exists by Zorn's lemma). If H does not have codimension one, then we can consider a subspace $S \supset H$ of two dimensions higher and all extensions $H' = H + \mathbb{R}v_\theta$ of H to S , $\theta \in S^\perp$.

Now consider the set $A \subset S^\perp$ of θ such that $H + \mathbb{R}_+v_\theta$ meets K . Then A is open, connected, and $A \cap A + \pi = \emptyset$; else K would meet H . It follows that A is an open interval of length at most π . Taking an endpoint of A , we obtain an extension of H to H' , a contradiction.

Thus H has codimension one. Since K is open, \overline{H} is also disjoint from K , and hence $H = \overline{H}$. Thus H is the kernel of the desired linear functional. ■

Separation.

Theorem 10.15 *Let $K, L \subset X$ be disjoint convex sets, with K open. Then there is a $\phi \in X^*$ separating K from L ; i.e. $\phi(K)$ and $\phi(L)$ are disjoint.*

Proof. Let $M = K - L$; this set is open, convex, and it does not contain the origin because K and L are disjoint. Thus by the support theorem, there is a linear functional with $\phi(M) \geq 0$. Then for all $k \in K$ and $\ell \in L$, we have $\phi(k - \ell) = \phi(k) - \phi(\ell) \geq 0$. It follows that $\inf \phi(K) \geq \sup \phi(L)$. Since $\phi(K)$ is open, these sets are disjoint. ■

Weak closure. Recall that $K \subset X$ is **weakly closed** if whenever a net $x_\alpha \in K$ satisfies $\phi(x_\alpha) \rightarrow \phi(x)$ for all $\phi \in X^*$, we have $x \in K$. The weak closure of a set is generally larger than the strong (or norm) closure. For

example, the sequence $f_n(x) = \sin(nx)$ in $L^1[0, 1]$ is closed in the norm topology (it is discrete), but its weak closure adds $f_0 = 0$.

Another good image to keep in mind is that K is weakly closed if for any $x \notin K$, there is a continuous linear map $\Phi : X \rightarrow \mathbb{R}^n$ such that $\overline{\Phi(K)}$ is disjoint from $\Phi(x)$. This is just because a base for the weak topology consists of finite intersections of sets of the form $\phi^{-1}(\alpha, \beta)$.

Theorem 10.16 *A convex set $K \subset X$ is weakly closed iff K is strongly (norm) closed.*

Proof. A weakly closed set is automatically strongly closed. Now suppose K is strongly closed, and $x \notin K$. Then there is an open ball B containing x and disjoint from K . By the separation theorem, there is a $\phi \in X^*$ such that $\phi(x) > \phi(K)$, and thus x is not in the weak closure of K . Thus K is weakly closed.

Linear combinations. By the preceding result, we see that the weak closure of a set $E \subset X$ is contained in $\text{hull}(E)$, the smallest norm-closed convex set containing E . Now $\text{hull}(E)$ can be described as the closure of finite convex combinations of points in E . So as an example we have:

Proposition. For any $\epsilon > 0$ there exist constants $a_n \geq 0$, $\sum a_n = 1$, such that

$$\left\| \sum_1^N a_n \sin(nx) \right\|_1 < \epsilon.$$

Problem. Prove this directly!

(Solution. Just take $a_n = 1/N$ for $n = 1, \dots, N$, and note that for orthogonal functions e_n the function $f = \sum a_n e_n$ satisfies

$$\|f\|_1^2 \leq \|f\|_2^2 \|1\|_2^2 = O(\sum |a_n|^2) = O(1/N).$$

Intuitively, $f(x)$ behaves like a random walk with N steps.

LCTVS. A *topological vector space* X is a vector space with a topology such that addition and scalar multiplication are continuous. By translation invariance, to specify the topology on X it suffices to give a basis at the origin.

A very useful construction comes from continuity of addition: for any open neighborhood U of the origin, there is a neighborhood V such that $V + V \subset U$.

Usually we assume X is Hausdorff (T_2). This is equivalent to assuming points are closed (T_1). Indeed, if points are closed and $x \neq y$, then we can

find a balanced open neighborhood U of the origin such that $y+U$ is disjoint from X . We can then find a balanced open V such that $V+V \subset U$, and then $x+V$ is disjoint from $y+V$.

Warning: Royden at times implicitly assumes X is Hausdorff. For example, if X is not Hausdorff, then an extreme point is not a supporting set, contrary to the implicit assumption in the proof of the Krein-Milman theorem.

Let X be a Banach space. Then the weak topology on X and the weak* topology on X^* are Hausdorff and locally convex. All the results like the Hahn-Banach theory, the separation theorem, etc. hold for locally convex topologies as well as the norm topology and weak topology.

Extreme points. Let K be convex. A point $x \in K$ is an *extreme point* if there is no open interval in K containing x . More generally, a *supporting set* $S \subset K$ is a closed, convex set with the property that, whenever an open interval $I \subset K$ meets S , then $I \subset S$. One should imagine a face of ∂K or a subset thereof.

Example: Let K be a convex compact set. Then the set of points where $\phi \in X^*$ assumes its maximum on $K \subset X$ is a supporting set. In particular, any compact convex set has nontrivial supporting sets.

Theorem 10.17 (Krein-Milman) *Let K be a compact convex set in a locally convex (Hausdorff) topological vector space X . Then K is the closed convex hull of its extreme points.*

Remark. The existence of *any* extreme points is already a nontrivial assertion.

Proof. We will show any supporting set contains an extreme point. Indeed, consider any minimal nonempty supporting set $S \subset K$; these exist by Zorn's lemma, using compactness to guarantee that the intersection of a nested family of nonempty supporting sets is nonempty. Now if S contains two distinct points x and y , we can find a $\phi \in X^*$ (continuous in the given topology) such that $\phi(x) \neq \phi(y)$. Then the set of points in S where ϕ assumes its maximum is nonempty (by compactness) and again a supporting set, contrary to minimality.

Now let $L \subset K$ be the closed convex hull of the extreme points. If there is a point $x \in K - L$, then we can separate x from L by a linear functional, say $\phi(x) > \phi(L)$. But then the set of points where ϕ assumes its maximum is a supporting set, and therefore it contains an extreme point, contrary to the assumption that ϕ does not assume its maximum on L .

Therefore $L = K$. ■

Prime example: The unit ball in X^* , in the weak* topology.

What are the extreme points of the unit ball B in $L^2[0, 1]$? Every point in ∂B is extreme! Because if $\|f\|_2 = 1$ then for ϵ and $g \neq 0$, we have

$$\|f \pm \epsilon g\|_2^2 = \|f\|_2^2 \pm 2\epsilon \langle f, g \rangle + \epsilon^2 \|g\|_2^2$$

and this is $> \|f\|_2^2 = 1$ if the sign is chosen properly.

What about the unit ball in $L^\infty[0, 1]$? Here the extreme points are functions with $|f| = 1$ a.e. Picture the finite-dimensional case — a cube.

What about in $L^1[0, 1]$? Here there are *no* extreme points! For example, if $f = 1$, then $f(x) + a \sin(2\pi x)$ has norm one for all small a , so f is not extreme. Similarly, for any $f \neq 0$ we can find a set A of positive measure on which $f > a > 0$ (or $0 > a > f$), and then a function g of mean zero supported on A such that $\|f \pm g\| = \|f\|$.

This fact is compatible with Krein-Milman only because $L^1[0, 1]$ is *not* a dual space. In fact the preceding remark *proves* that for any Banach space X , the dual X^* is not isomorphic to $L^1[0, 1]$.

For $X = C[0, 1]$, the dual X^* consists of signed measures of total mass one, and the extreme points are $\pm \delta_x$.

Stone-Weierstrass revisited. Let X be a compact Hausdorff space, and let $A \subset C(X)$ be an algebra of real-valued functions containing the constants and separating points. Then A is dense in $C(X)$.

Proof (de Brange). Let $A^\perp \subset M(X)$ be the set of measures that annihilate A . By the Hahn-Banach theorem, to show A is dense it suffices to show that A^\perp is trivial.

Let K be the intersection of A^\perp with the unit ball. Then K is a closed, compact, convex set in the weak* topology. Thus K is the closed convex hull of its extreme points.

Suppose $\mu \in K$ is a nonzero extreme point. We will deduce a contradiction.

First, let $E \subset X$ be the support of μ (the smallest closed set whose complement has measure zero). Suppose E is not a single point. Choose a function $f \in A$ such that $f|_E$ is not constant, and $|f| < 1$. Consider the two measures

$$\sigma = (1 + f)\mu/2, \quad \tau = (1 - f)\mu/2.$$

Since A is an algebra, both σ and τ are in A^\perp , and of course we have $\sigma + \tau = \mu$. Moreover, since $1 \pm f > 0$, we have

$$\|\mu\| = \|\sigma\| + \|\tau\| = 1.$$

Thus μ is a convex combination:

$$\mu = \|\sigma\| \frac{\sigma}{\|\sigma\|} + \|\tau\| \frac{\tau}{\|\tau\|}.$$

Since μ is an extreme point, it follows that $\mu = \sigma = \tau$. Therefore f is constant a.e. on E , a contradiction.

It follows that μ is a delta-mass supported on a single point. But the μ is not in A^\perp , since it pairs nontrivially with the constant function in A . ■

Haar measure. As a further application of convexity, we now develop the Kakutani fixed-point theorem and use it to prove the existence of Haar measure on a compact group. Our treatment follows Rudin, *Functional Analysis*, Chapter 5.

Theorem 10.18 (Milman) *Let $K \subset X$ be a compact subset of a Banach space and suppose $H = \overline{\text{hull}(K)}$ is compact. Then the extreme points of H are contained in K .*

Proof. Suppose x is an extreme point of H that does not lie in K , and let $r = d(x, K)$. Then by compactness we can cover K by a finite collection of balls $B(x_i, r)$, $i = 1, \dots, n$. Let H_i be the closed convex hull of $K \cap B(x_i, r)$. Since the ball is compact, we have $H_i \subset B(x_i, r)$.

Now $H = \text{hull}(\bigcup_1^n H_i)$, and thus $x = \sum t_i h_i$ is a convex combination of points $h_i \in H_i \subset H$. But x is an extreme point, so $x = h_i$ for some i . This implies $x \in B(x_i, r)$, contradicting the fact that $d(x, K) = 2r$. ■

Theorem 10.19 (Kakutani) *Let $K \subset X$ be a nonempty compact convex subset of a Banach space, and let G be a group of isometries of X leaving K invariant. Then there exists an $x \in K$ fixed by all $g \in G$.*

Proof. Let $L \subset K$ be a minimal, nonempty, compact convex G -invariant set; such a set exists by the Axiom of Choice. If L consists of a single point, we are done. Otherwise there are points $x \neq y$ in L . Let $z = (x+y)/2$. Then by minimality of L , we have $L = \overline{\text{hull}(G \cdot z)}$. Let z' be an extreme point of L . By Milman's theorem, z' is a limit of points in $G \cdot z$. By compactness of K , we can choose $g_n \in G$ such that $g_n z \rightarrow z'$, $g_n x \rightarrow x'$ and $g_n y \rightarrow y'$. But then $z' = (x' + y')/2$, so z' is not an extreme point. ■

Theorem 10.20 *Let G be a compact Hausdorff group. Then there is a unique left-invariant Borel probability measure μ on G , and μ is also right invariant.*

Proof. Let G be a compact topological group, and let $X = C(G)$. For each $g, h \in G$, the shift operators $L_g(f) = f(g^{-1}x)$ and $R_h(f) = f(xh)$ are isometries, and they commute. The only fixed-points for G are the constant functions.

Now fix $f \in C(G)$. Then f , and all its translates, are equicontinuous, and thus

$$L(f) = \overline{\text{hull}(G \cdot f)} \subset C(G)$$

is compact. Similarly, the closed convex hull of the right translates, $R(f) = \overline{\text{hull}(f \cdot G)}$, is also compact. By Kakutani's fixed-point theorem, each of these convex sets contains at least one constant function, $l(f)$ and $r(f)$.

The constant $l(f)$ can be approximated by averages of the form

$$T(f) = \sum \alpha_i L_{a_i}(f),$$

and similarly for $r(f)$. But the right and left averages commute, and leave the constants invariant, so $l(f) = r(f)$. Thus there is a unique constant function, $M(f)$, contained in both $L(f)$ and $R(f)$.

To show $M(f)$ corresponds to Haar measure, we must show $M(1) = 1$, $M(f) \geq 0$ when $f \geq 0$, and M is linear. The first two assertions are immediate. To show $M(f + h) = M(f) + M(h)$, choose a left-averaging operator T such that $M(f) \approx T(f)$. Then $T(h) \in L(h)$, so $M(T(h)) = M(h)$. Thus there is a second left-averaging operator S such that $S(T(h)) \approx M(h)$. But then $S(T(f + h)) \approx M(f) + M(h) \in L(f + h)$, so $M(f + h) = M(f) + M(h)$. ■

Examples of compact groups: Finite groups, products such as $(\mathbb{Z}/2)^\mathbb{N}$, inverse limits such as $\mathbb{Z}_p = \varprojlim (\mathbb{Z}/p^n)$; Lie groups such as $\text{SO}(n, \mathbb{R})$ and $\text{SU}(n, \mathbb{C})$; p -adic Lie groups such as $\text{SL}_n(\mathbb{Z}_p)$.

Here is a description of Haar measure on $G = \text{SO}(n, \mathbb{R})$. Consider the Lie algebra $\mathfrak{g} = \mathfrak{so}(n, \mathbb{R})$; it is the space of trace-zero matrices satisfying $A^t = -A$. There is a natural inner product on this space, given by $\langle A, B \rangle = \text{tr}(AB)$. This inner product is invariant under the adjoint action of G , so it gives rise to an invariant quadratic form on every tangent space $T_g G$. In the case of $\text{SO}(n)$, this metric is negative definite. Thus its negative determines a bi-invariant **metric** on $\text{SO}(n, \mathbb{R})$, and hence an invariant measure.

This measure can be described as follows: to choose a random frame in \mathbb{R}^n , one first pick a point at random on S^{n-1} , then a point at random on the orthogonal S^{n-2} , etc., using the rotation-invariant probability measures on each sphere. There is a unique choice for the final point on S^0 that makes the frame positively oriented.

Unimodularity. More generally, any locally compact group G carries right and left invariant measures, unique up to scale, but they need not agree. When they do, the group is *unimodular*. For example, the group $\mathrm{SL}_2(\mathbb{R})$ is unimodular, but its upper-triangular subgroup AN is not.

11 Fourier Series

In this section we turn to Fourier series for functions on S^1 , and study their convergence and their applications to differential equations.

Complex Hilbert spaces. The basics regarding Hilbert spaces were presented at the end of section §7.

Over the field \mathbb{C} , the natural form for a Hilbert space is a Banach space H with a *Hermitian form* $\langle x, y \rangle$ such that $\langle x, x \rangle = \|x\|^2$. In this case, $(x, y) = \mathrm{Re}\langle x, y \rangle$ makes H into a Hilbert space over \mathbb{R} . Examples:

$$\begin{aligned}\langle x, y \rangle &= \sum x_i \bar{y}_i \text{ on } \mathbb{C}^n \text{ or } \ell^2(\mathbb{N}) \otimes \mathbb{C}. \\ \langle f, g \rangle &= \int f \bar{g} \text{ on } L^2(\mathbb{R}^n).\end{aligned}$$

The circle. Our main focus will be on the space of *complex-valued* functions on the circle S^1 .

There are two natural coordinates on S^1 : the coordinate $\theta \in [0, 2\pi]$ or, more naturally, in $\mathbb{R}/2\pi\mathbb{Z}$, and the coordinate

$$z = \exp(i\theta) = \cos(\theta) + i \sin(\theta).$$

We will use both; when we write $f(z)$ or $f(\theta)$, these functions are meant to be related by the change of coordinates above. It is also sometimes useful to regard $f(\theta)$ as a function $f : \mathbb{R} \rightarrow \mathbb{C}$ such that

$$f(\theta + 2\pi) = f(\theta).$$

The Hilbert space of the circle. To define $L^2(S^1)$, we take the *normalized* inner product

$$\langle f, g \rangle = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) \overline{g(\theta)} d\theta = \frac{1}{2\pi} \int_{S^1} f(z) \overline{g(z)} |dz|.$$

This normalization simplifies our choice of an orthonormal basis: namely, the functions

$$e_n(z) = z^n = \cos(n\theta) + i \sin(n\theta),$$

$n \in \mathbb{Z}$, are obviously orthonormal, and by the complex form of the Stone–Weierstrass theorem, they span a dense subspace of $C(S^1)$ and hence form a basis for $L^2(S^1)$.

The Fourier transform. Since the basis elements e_n are indexed by \mathbb{Z} , the Fourier coefficients of $f \in L^2(S^1)$ naturally take values in $\ell^2(\mathbb{Z})$. More precisely, we define the (discrete) Fourier transform of f by

$$\widehat{f}(n) = \langle f, e_n \rangle = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) e^{-in\theta} d\theta.$$

Then for any $f \in L^2(S^1)$, we have

$$f = \sum_{n \in \mathbb{Z}} \widehat{f}(n) e_n,$$

in the sense of convergence in the L^2 norm.

We note that $\widehat{f}(n)$ also makes sense for any $f \in L^1(S^1)$. We will often be interested in f in spaces other than $L^2(S^1)$, especially $f \in C(S^1)$, the space of continuous functions, or $f \in C^k(S^1)$, the space of functions with k continuous derivatives.

Convergence of Fourier series. One of the main concerns of analysts for 150 years has been the following problem: given a function $f(\theta)$ on S^1 , in what sense do we have:

$$f(\theta) = \sum a_n e^{in\theta},$$

where $a_n = \widehat{f}(n) = \langle f, e_n \rangle$?

It is traditional to write

$$S_N(f) = \sum_{-N}^N a_n e_n.$$

We begin by briefly summarizing the main results.

1. If $f \in L^2(S^1)$, then $S_N(f) \rightarrow f$ in the L^2 -norm. (We have already seen this holds in a general Hilbert space.)
2. Consequently, if $f \in C(S^1)$, then (a_n) determine f and $a_n \rightarrow 0$ as $n \rightarrow \infty$.

3. If $f(x)$ is $C^2(S^1)$, then $a_n = O(1/n^2)$ and thus $S_N(f)$ converges to f uniformly. (Observe that $\widehat{f'}(n) = in\widehat{f}(n)$.)
4. (Dirichlet's theorem.) In fact, if $f \in C^1(S^1)$, then $S_N(f)$ converges to f pointwise.
5. (Fejér's theorem.) If $f \in C(S^1)$, then $(S_0(f) + \cdots + S_{N-1}(f))/N$ converges to $f(x)$ uniformly.
6. (Abel's theorem.) If $f \in C(S^1)$, then $f_r(z) \rightarrow f(z)$ uniformly as $r \rightarrow 1$, where

$$f_r(z) = \sum_{-\infty}^{\infty} a_n r^n z^n.$$

7. (DuBois-Reymond.) However, $S_N(f)$ does not even converge, for a generic $f \in C(S^1)$.

The theorems of Dirichlet, Fejér, Abel and DuBois-Reymond all depend on the idea of expressing $S_N(f)$ as a convolution. We begin by developing this background concept.

Convolution. A key fact in our analysis is that S^1 and \mathbb{Z} are groups. In fact they are dual topological groups, in the sense that the functor

$$\widehat{G} = \text{Hom}(G, S^1)$$

interchanges S^1 and \mathbb{Z} . That is, every continuous homomorphism from S^1 to S^1 is given by $z \mapsto z^n$ for some n , and every continuous homomorphism from \mathbb{Z} to S^1 is given by

$$n \mapsto e_n(\theta) = e^{in\theta}$$

for some $\theta \in S^1$.

The group structure allows us to define the operation of *convolution* on $L^1(S^1)$ and $\ell^1(\mathbb{Z})$. This is a continuous bilinear operation making each of these spaces into a Banach algebra. That is, they satisfy

$$\|f * g\|_1 \leq \|f\|_1 \cdot \|g\|_1.$$

On sequences, convolution is defined by

$$(a * b)_n = \sum_{i+j=n} a_i b_j;$$

while on functions, we set

$$(f * g)(\theta) = \frac{1}{2\pi} \int_{S^1} f(\alpha)g(\theta - \alpha) d\alpha.$$

We can also write this as

$$(f * g)(z) = \frac{1}{2\pi} \int_{S^1} f(w)g(z/w) |dw|.$$

Behavior of convolutions. In general, if f is in $L^1(G)$ and g is an a Banach space X of functions on G with a translation-invariant norm, then $f * g \in X$ and we have

$$\|f * g\|_X \leq \|f\|_1 \cdot \|g\|_X.$$

To see this, observe that $f * g$ is approximated by a linear combination of translates of g , and apply the triangle inequality for the norm on X .

This if $f, g \in L^2(S^1) \subset L^1(S^1)$, then $f * g \in L^2(S^1)$. In fact, $f * g \in C(S^1)$, since we can write

$$(f * g)(\theta) = \langle f_\theta, g \rangle,$$

where f_θ moves continuously in $L^2(S^1)$. So in this case convolution greatly improves the behavior of functions. As we will see below, this is ‘dual’ to the fact that $\ell^2(\mathbb{Z})$ is closed under multiplication.

On the other hand, $\ell^2(\mathbb{Z})$ is not contained in $\ell^1(\mathbb{Z})$, and $\ell^2(\mathbb{Z})$ is not closed under convolutions. As we will see below, this is dual to the fact that $L^2(S^1)$ is *not* closed under multiplication.

In our study of convergence of Fourier series, we will be convolving with smooth kernels K , and f will lie in $L^2(S^1)$ or better, so $K * f$ will always exist — in fact, it will be a smooth function.

Convolution and multiplication. The most important property of convolution is that it is the Fourier transform of multiplication. That is,

$$\widehat{fg} = \widehat{f} * \widehat{g},$$

and

$$\widehat{f * g} = \widehat{f}\widehat{g}. \tag{11.1}$$

The first is easily verified formally: we have

$$\left(\sum a_i z^i\right) \left(\sum b_j z^j\right) = \sum_n \left(\sum_{i+j=n} a_i b_j\right) z^n.$$

For the second, it is useful to observe that

$$e_n e_m = e_{n+m} \quad \text{and} \quad e_n * e_m = \delta_{nm} e_n.$$

This comes down to the computation:

$$(e_n * e_m)(z) = \frac{1}{2\pi} \int_{S^1} w^n (z/w)^m |dw| = \delta_{nm} z^n = \delta_{nm} e_n.$$

Note that $\widehat{e_n} = \delta_{in}$. Thus the computation above shows that

$$\widehat{e_n * e_m} = \widehat{e_n} \widehat{e_m},$$

and the general case, equation (11.1), follows from the fact that e_n is a basis for $L^2(S^1)$. Explicitly, we have:

$$\left(\sum a_i z^i \right) * \left(\sum b_j z^j \right) = \sum a_n b_n z^n.$$

The Dirichlet kernel. We now return to the subject of *convergence* of Fourier series. Given $f \in L^2(S^1)$ we let

$$a_n = \widehat{f}(n) = \langle f, e_n \rangle,$$

and we let

$$S_N(f) = \sum_{-N}^N a_n z^n.$$

This is nothing but the projection of f to the subspace spanned by e_n with $|n| \leq N$.

The *Dirichlet kernel* is defined by

$$D_N = \sum_{n=-N}^N e_n.$$

Since $e_i * e_j = \delta_{ij} e_i$, it satisfies

$$S_N(f) = D_N * f.$$

In particular the total mass of D_N is one: we have

$$\frac{1}{2\pi} \int_0^{2\pi} D_N(\theta) d\theta = 1,$$

since $S_N(e_0) = D_N * 1 = e_0 = 1$.

On the other hand, we can write an exact formula for this kernel by summing the geometric series:

$$D_N(z) = \sum_{-N}^N z^n = z^{-N}(1 + z + \cdots + z^{2N}) = \frac{1 - z^{2N+1}}{z^N(1 - z)}.$$

Dividing the top and bottom by $z^{(N+1)/2}$, this becomes

$$D_N(z) = \frac{z^{(N+1)/2} - z^{-(N+1)/2}}{z^{1/2} - z^{-1/2}} = \frac{\sin((N + 1/2)\theta)}{\sin(\theta/2)}.$$

(We have used the fact that $z - 1/z = 2i \sin(\theta)$.)

Note that the top and bottom of this fraction are only periodic with period 4π , not 2π . However they both change by a sign under $\theta \mapsto \theta + 2\pi$, so their ratio is in fact well-defined on S^1 .

Dirichlet's theorem on C^1 functions. We can now establish:

Theorem 11.1 (Dirichlet) *If $f \in C^1(S^1)$, then $S_N(f)$ converges pointwise to f .*

Proof. It is enough to prove convergence at $\theta = 0$. We write $S_N(f) = D_N * f$. Since $S_N(f) = f$ when f is a constant, we can also assume $f(0) = 0$. Our goal then is to show that

$$S_N(f)(0) = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) D_N(\theta) d\theta \rightarrow 0$$

as $N \rightarrow \infty$. (Here we have used the fact that the Dirichlet kernel is *even*.)

Here is the key point: since f is C^1 and $f(0) = 0$, the function

$$g(\theta) = \frac{f(\theta)}{\sin(\theta/2)}$$

is continuous at $\theta = 0$! Unfortunately it is not quite defined on S^1 , since the denominator is only periodic under $\theta \mapsto \theta + 4\pi$. But we can rewrite $S_N(f)(0)$ as $(1/2)$ of an integral from 0 to 4π ; then we have

$$S_N(f)(0) = \frac{1}{4\pi} \int_0^{4\pi} g(\theta) \sin((N + 1/2)\theta) d\theta.$$

By the Riemann–Lebesgue lemma, this integral tends to zero, and thus $S_N(f) \rightarrow f$ pointwise. ■

Localization. Dirichlet's proof of convergence of $S_N(f)$ at $\theta = 0$ only used the existence of $f'(\theta)$ at $\theta = 0$. Thus it can be generalized to show:

Theorem 11.2 *Suppose $f \in L^1(S^1)$ and $f'(\theta)$ exists. Then $S_N(f)(\theta) \rightarrow f(\theta)$.*

Corollary 11.3 *If f and g agree on a neighborhood of θ , then $S_N(f)(\theta) - S_N(g)(\theta) \rightarrow 0$ as $N \rightarrow \infty$.*

*Dirichlet's proof . . . left open the question as to whether the Fourier series of every Riemann integrable, or at least every continuous, function converged. At the end of his paper Dirichlet made it clear he thought that the answer was yes (and that he would soon be able to prove it). During the next 40 years Riemann, Weierstrass and Dedekind all expressed their belief that the answer was positive. —Körner, *Fourier Analysis*, §18.*

In fact this is false!

Theorem 11.4 (DuBois-Reymond) *There exists an $f \in C(S^1)$ such that $\sup_N |S_N(f)(0)| = \infty$.*

The proof will use the uniform boundedness principle, plus the following computation:

Lemma 11.5 *The Dirichlet kernel satisfies $\|D_N\|_1 \rightarrow \infty$ as $N \rightarrow \infty$.*

Proof. Recall that

$$D_N(\theta) = \frac{\sin((N + 1/2)\theta)}{\sin(\theta/2)}.$$

Clearly all the action occurs near $\theta = 0$; indeed, we have $|D_N(\theta)| = O(1/|\theta|)$ on $[-\pi, \pi]$. But near $\theta = 0$, there are periodic intervals on which $|\sin((N + 1/2)\theta)| > 1/2$. On these intervals, $|1/\sin(\theta/2)| \approx 2/|\theta|$. Since $\int d\theta/|\theta|$ diverges, we have $\|D_N\|_1 \rightarrow \infty$. In fact, the L^1 -norm behaves like $\int_{1/N}^1 d\theta/\theta \asymp \log(N)$. ■

Proof of the theorem of DuBois-Reymond. We can now state and prove a strong form of the existence of a continuous function with a divergent Fourier series.

Theorem 11.6 *The Fourier coefficients of a generic $f \in C(S^1)$ satisfy $\sup_N |\sum_{-N}^N a_n| = \infty$. Thus $\sum a_n e^{in\theta}$ does not converge at $\theta = 0$.*

Proof. The functions $D_N(\theta)$ are even; thus we have, for any $f \in C(S^1)$,

$$\phi_N(f) = S_N(f)(0) = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) D_N(\theta) d\theta.$$

This formula shows that $S_N(f)(0)$ is an element of $C(S^1)^*$ given by integration against D_N . Now the function $\text{sign } D_N$ is a pointwise limit a.e. of continuous functions with $|f| \leq 1$; thus we have

$$\|\phi_N\|_{C(S^1)^*} = \|D_N\|_1 \rightarrow \infty.$$

By the strong form of the uniform boundedness (Theorem 10.13), this implies that $\sup |\phi_N(f)| = \infty$ for generic $f \in C(S^1)$. For such an f , $\phi_N(f) = S_N(f)(0)$ cannot converge. ■

Post-modern theory. After this phenomenon was discovered, a common sentiment was that it was only a matter of time before a continuous function would be discovered whose Fourier series diverged everywhere. Thus it was even more remarkable when L. Carleson proved:

Theorem 11.7 *For any $f \in L^2(S^1)$, the Fourier series of f converges to f pointwise almost everywhere.*

The proof is very difficult.




Figure 3. The Dirichlet and Fejér kernels.

The Fejér kernel. However in the interim Fejér, at the age of 19, proved a very simple result that allows one to reconstruct the values of f from its Fourier series for any continuous function.

Theorem 11.8 For any $f \in C(S^1)$, we have

$$f(x) = \lim \frac{S_0(f) + \cdots + S_{N-1}(f)}{N}$$

uniformly on the circle.

This expression is a special case of *Césaro summation*, where one replaces the sequence of partial sums by their averages. This procedure can be iterated. In the case at hand, it amounts to computing $\sum_{-\infty}^{\infty} a_n$ as the limit of the sums

$$\frac{1}{N} \sum_{i=-N}^N (N - |i|)a_n.$$

Approximate identities. To explain Fejér's result, it is useful to first understand the idea of convolution and approximate identities.

Let K_n be a sequence in $L^1(S^1)$. As in §9, where we treated the case of \mathbb{R} instead of S^1 , we say K_n is an *approximate identity* if:

1. $K_n \geq 0$,
2. $\langle K_n, 1 \rangle = 1$, and
3. $\int_{S^1 - U} K_n(\theta) d\theta \rightarrow 0$ for any neighborhood U of $\theta = 0$.

In other words, almost all the mass of $K_n(\theta)$ is concentrated near $\theta = 0$. (This is usually verified by shown that $K_n \rightarrow 0$ uniformly outside U .)

Morally, this means that the *measures* $K_n(\theta)d\theta/2\pi$ converge to the δ -mass at $\theta = 0$.

Theorem 11.9 If K_n is an approximation to the identity, and $f \in C(S^1)$, then $K_n * f \rightarrow f$ uniformly on S^1 .

Proof. Think of $K_n * f(\theta)$ as a sum of the translates $f(\theta - \alpha)$ weighted by $K_n(-\alpha)$. The translates with $-\alpha$ small are uniformly close to $f(\theta)$ because f is uniformly continuous. The translates with α large make a small contribution, because their total weight with respect to K_n is small. Thus $K_n * f$ is uniformly close to f . (A more formal proof can easily be given along the same lines as the proof of Theorem 9.30.) ■

The argument shows more. Suppose $f \in L^\infty(S^1)$. Let $f_\pm(x)$ denote the right and left limits of f , if they exist. Then we have:

Theorem 11.10 *For any approximate identity $\langle K_n \rangle$:*

1. $K_n * f(x) \rightarrow f(x)$ at any point x where f is continuous.
2. $l(K_n * f)(x) \rightarrow (f_+(x) + f_-(x))/2$, at any point where f is piecewise-continuous.
3. If f is continuous on a neighborhood of $[a, b]$, then $K_n * f \rightarrow f$ uniformly on $[a, b]$.

These more refined results hold for any summation method that can be related to approximate identities; in particular, they will hold for Cesàro and Abel summation.

Proof of Fejér's theorem. If we let

$$T_N(f) = \frac{S_0(f) + \cdots + S_{N-1}(f)}{N},$$

then $T_N(f) = f * F_N$ where the *Fejér kernel* is given by

$$F_N = (D_0 + \cdots + D_{N-1})/N.$$

Of course $\langle F_N, 1 \rangle = 1$ since $\langle D_n, 1 \rangle = 1$. But in addition, F_N is *positive* and *concentrated near 0*, i.e. it is an approximation to the identity. Indeed, we have:

$$F_N(\theta) = \frac{\sin^2(N\theta/2)}{N \sin^2(\theta/2)},$$

and away from the zeros of $\sin(\theta/2)$ this expression converges uniformly to zero, because of the N in the denominator.

To see the positivity of F_N more directly, note for example that

$$\begin{aligned} (2N+1)F_{2N+1} &= z^{-2N} + 2z^{-2N+1} + \cdots + (2N+1) + \cdots + 2z^{2N-1} + z^{2N} \\ &= (z^{-N} + \cdots + z^N)^2 = D_N^2. \end{aligned}$$

■

Proof of Abel's theorem. The proof of Abel's theorem is very similar. It turns on the computation of the *Abel kernel*, given by

$$A_r(z) = \sum_{-\infty}^{\infty} r^{|n|} z^n.$$

This has the property that

$$U_r(f) = \sum_{-\infty}^{\infty} r^{|n|} a_n z^n = A_r * f.$$

So to complete the proof of Abel's theorem, it suffices to show that $A_r(z)$ is an approximation to the identity as $r \rightarrow 1$.

But as in the case of the Dirichlet kernel, the Abel kernel is simply a geometric series (or rather, the sum of two such series). Thus we can explicitly compute:

$$A_r(z) = -1 + \sum_0^{\infty} ((rz)^n + (r/z)^n) = -1 + \frac{1}{1-rz} + \frac{1}{1-r/z}.$$

Using the fact that $z + 1/z = 2 \cos \theta$, this gives:

$$A_r(z) = \frac{1-r^2}{1+r^2-2r \cos \theta}.$$

The key point now is that the denominator is greater than $(1-r)^2$, so we have $A_r(z) \geq 0$. Of course $\langle A_r, 1 \rangle = 1$, and if we avoid a neighborhood of $\theta = 0$, then the denominator is bounded below (it converges to $2(1 - \cos \theta)$ as $r \rightarrow 1$) while the numerator tends to zero uniformly. Thus A_r is indeed an approximate identity, and hence

$$U_r(f) = A_r * f \rightarrow f$$

uniformly, for any $f \in C(S^1)$. ■

Application of Fourier series to differential equations. We now discuss the application of Fourier series to several different differential equations:

1. *Laplace's equation*, $f_{xx} + f_{yy} = \Delta f = 0$;
2. *The heat equation*, $f_t = -f_{xx}$; and

3. *The wave equation, $f_{xx} - f_{yy} = 0$.*

Laplace's equation on the disk. Laplace's equation governs several physical phenomena, ranges from temperature distributions (in equilibrium) to electric fields (in the absence of charge).

One of the nicest applications is to the solution of the *boundary value* problem for Laplace's equation on the unit disk. We are given $f \in C(S^1)$ and we wish to find an extension of f to a function $F \in C(\overline{\Delta})$ with

$$\Delta F = 0$$

inside the unit disk. This F turns out to be unique (by the maximum principle), and it is given simply by

$$F(z) = a_0 + \sum_{n>0} a_n z^n + a_{-n} \bar{z}^n,$$

where $f(z) = \sum a_n z^n$. Since $a_n \rightarrow 0$, it is easy to see that $F(z)$ is the sum of a holomorphic and anti-holomorphic function on the disk; in particular, F is infinitely differentiable and $\Delta F = 0$. Moreover, by Abel's theorem, $F|_{S^1}(r)$ converges uniformly to $f|_{S^1}$, which implies that F gives a continuous extension of f to the disk.

Sine series. If we restrict to *odd* functions — where $f(-x) = -f(x)$ — then only sine terms appear, and we can identify this subspace with $L^2[0, \pi]$. Thus a function on $[0, \pi]$ has a natural Fourier series:

$$f(x) = \sum_{n=1}^{\infty} a_n \sin(nx).$$

Since $\int_0^\pi \sin^2(x) dx = \pi/2$ (its average value is $1/2$), we have

$$a_n = \frac{2}{\pi} \int_0^\pi f(x) \sin(nx) dx.$$

Example: if the graph of f is a triangle with vertex (p, x) , then

$$a_n = \frac{2h \sin(np)}{n^2 p(\pi - p)}.$$

The wave equation and the heat equation. A typical problem in PDE is to solve the *wave equation* with given initial data $f(x) = u(x, 0)$ on $[0, \pi]$.




Figure 4. Solutions to the wave equation (undamped and damped) and the heat equation.

This equation, which governs the motion $u(x, t)$ of a vibrating string, is given by

$$u_{tt} = u_{xx}$$

(where the subscripts denote differentiation). If we think of $u(x, t)$ as the motion of a string with fixed end points, it is natural to impose the boundary conditions $u(0, t) = u(\pi, t) = 0$. We will also assume $u_t(x, 0) = 0$, i.e. the string is initially stationary.

Since the wave equation is linear, it suffices to solve it for the Fourier basis functions $f(x) = \sin(nx)$. And for these we have simply

$$u(x, t) = \cos(nt) \sin(nx).$$

This solution can be discovered by separation of variables; the key is that $f(x)g(t)$ solves the wave equation if f and g are eigenfunctions with the same eigenvalues.

These basic solutions are ‘standing waves’ corresponding to the bass note and then the higher harmonics of the string.

The solution to the wave equation for ‘general’ $f(x)$ is then given by:

$$u(x, t) = \sum_{n=1}^{\infty} a_n \cos(nt) \sin(nx).$$

Note that $u(x, t + 2\pi) = u(x, t)$, i.e. the string has a natural frequency.

The *heat equation*

$$u_t = u_{xx}$$

governs the evolution of temperature with respect to time. In the case at hand the boundary conditions mean that the ends of the interval are kept at a constant temperature of zero. Now the basic solutions are given by

$$u(x, t) = e^{-n^2 t} \sin(nx).$$

and thus

$$u(x, t) = \sum_{n=1}^{\infty} a_n e^{-n^2 t} \sin(nx).$$

Note that the Fourier coefficients are severely damped for any positive time; $u(x, t)$ is in fact a real-analytic function of x for $t > 0$.

An actual plucked guitar string does not have a periodic motion but a motion that smooths and decays with time. It obeys a combination of the heat and wave equations:

$$u_{tt} + 2\delta u_t = u_{xx}.$$

Here the basic solutions are given by

$$u(x, t) = \exp(\alpha_n t) \sin(nx)$$

where $\alpha_n^2 + 2\delta\alpha_n + n^2 = 0$. So long as $0 \leq \delta \leq 1$ we get

$$\alpha_n = -\delta \pm i\sqrt{n^2 - \delta^2}$$

and thus the solution with $u_t(x, 0) = 0$ has the form

$$u(x, t) = \exp(-\delta t) \cos(\omega_n t + \sigma_n) \sin(nx),$$

where $\omega_n = \sqrt{n^2 - \delta^2}$ and $\tan(\sigma_n) = -\delta/\omega_n$. Note that the frequencies are now slowed and out of harmony — their ratios are no longer rational — and that $u(t, x)$ is damped but not smoothed out over time!

Weyl's equidistribution theorem. Here is another application of Fourier series that is different in spirit. It shows that the orbits of an irrational rotation are all very evenly distributed around the circle.

Theorem 11.11 (Weyl) *Let $\theta \in S^1 = \mathbb{R}/\mathbb{Z}$ be an irrational number. Then for any $f \in C(S^1)$, we have*

$$\frac{1}{N} \sum_{k=1}^N f(k\theta) \rightarrow \int_{S^1} f,$$

as $n \rightarrow \infty$.

Proof. Since the average is bounded by the sup-norm, $\|f\|$, it suffices to prove this theorem on a dense set of functions, such as the finite Fourier series. But for these it follows from the fact that if $n \neq 0$, then

$$\sum_{k=1}^N e^{ink\theta} = \frac{1 - e^{i(N+1)k\theta}}{1 - e^{ik\theta}}$$

is bounded, and hence its average goes to zero. ■

Example: the first-digit-1 phenomenon. You pick a country at random. I give you odds of 1 to 5 that the first digit of its area in meters is 1. I.e. if I'm right, you pay me a dollar, otherwise I pay you 5 dollars. Should you take the bet?

More concretely, consider the first digits of the sequence of numbers 2^n . How often should they be 1?

This is a question about the distribution of $x_n = n \log 2 \bmod \log 10$. Since $\log 2 / \log 10$ is irrational, x_n is uniformly distributed. Now $\log_{10}(2) = 0.3010300\dots$, so the first digit is 1 a full 30

If we assume the logarithm of the areas of countries (or weight of fish or Latin characters in a book...) is uniformly distributed, then we get the same answer. So my expected gain on this bet is $5 \times .3 - 1 \times .7 = 80$ cents, i.e. I can make almost a dollar per round.

12 Harmonic Analysis on \mathbb{R} and S^2 .

In this section we give a glimpse of harmonic analysis on two other spaces: \mathbb{R} and S^2 .

Fourier transform. One of the great ideas in analysis is the *Fourier transform* on $L^2(\mathbb{R})$. We define it on $f \in L^2(\mathbb{R})$ by

$$\widehat{f}(\xi) = \int_{\mathbb{R}} f(x) e^{-ix\xi} dx.$$

This integral at least makes sense when f is smooth and compactly supported.

We claim the inner product $(\widehat{f}, \widehat{f})$ is a constant multiple of (f, f) . Indeed, on any interval $[-\pi M, \pi M]$ large enough to contain the support of f , we have an orthonormal basis $g_i = e^{inx/M}/\sqrt{2\pi M}$; writing $f = \sum a_i g_i$ we find $a_i = \widehat{f}(n/M)/\sqrt{2\pi M}$, and thus

$$(f, f) = \sum |a_i|^2 = \frac{1}{2\pi M} \sum \widehat{f}(n/M) \rightarrow \frac{1}{2\pi} \int |\widehat{f}|^2$$

as $n \rightarrow \infty$. Thus f extends to all of L^2 as an isometry.

Fourier transform and differential equations. The Fourier transform reverse small-scale and large scale features of f . It turns differentiation d/dx_i into multiplication by x_i . Thus $\widehat{f}(0) = \int f$; if f is smooth then \widehat{f} decays rapidly at infinity; etc.

Since differentiation is turned into multiplication, it becomes easy to solve PDEs. For example, to solve $\Delta u = f$, you just pass to the transform side and divide \widehat{f} by $\sum x_i^2$. There is no difficulty near infinity for the result to be in L^2 ; this reflects that fact that Δ is a smoothing operator. There is difficulty near 0: both $\int f$ and the moments $\int f x_i$ should vanish for u to be in L^2 .

Spherical harmonics. We can also look from $L^2(S^1)$ in another direction — towards $L^2(S^{n-1})$, where the domain remains compact but its symmetry group becomes larger $G = \text{SO}(n)$. How do Fourier series generalize to the higher-dimensional spheres?

The case of a sphere is especially convenient because we can regard S^{n-1} as the unit ball $|x| = 1$ in \mathbb{R}^n . Let P_d denote the space of homogeneous polynomials of degree d on \mathbb{R}^n . We have

$$\dim P_d = \binom{d+n-1}{n-1}.$$

The Laplacian $\Delta = \sum d^2/dx_i^2$ maps P_d to P_{d-2} ; its kernel H_d is the space of *harmonic polynomials* of degree d . The key property of the Laplace operator is that it is $\text{SO}(n)$ -invariant.

Theorem 12.1 *We have $L^2(S^{n-1}) = \bigoplus_{d=0}^{\infty} H_d$.*

Generally a function $f \in H_d$ or its restriction to S^{n-1} is called a *spherical harmonic*. It can be shown that f , considered as a function the sphere, is actually an *eigenfunction* of the spherical Laplacian.

One can also study issues of pointwise convergence in this setting, for example one has:

Theorem 12.2 *If $f \in C^2(S^{n-1})$ then its Fourier series converges uniformly.*

To begin the proof that the spherical harmonics forms a ‘basis’ for $L^2(S^{n-1})$, we first show there is no relation between them.

Proposition 12.3 *The restriction map from \mathbb{R}^n to S^{n-1} is injective on $\oplus H_d$.*

Proof. A harmonic polynomial which vanishes on the sphere is everywhere zero, by the maximum principle. ■

Raising operator. Of course this result fails for general polynomials, because $r^2 = \sum x_i^2$ is constant on S^{n-1} . To take this into account, we introduce the *raising operator*

$$L : P_d \rightarrow P_{d+2}$$

defined by $L(f) = r^2 f$. Here are some of its key properties and their consequences.

1. If f is harmonic, then $\Delta L(f) = 2(n+d)f$. This is because

$$\Delta(r^2 f) = (\Delta r^2) f + (\nabla r^2) \cdot (\nabla f) + r^2 \Delta f$$

2. More generally, we have $\Delta L(f) = 2(n+d)f + L\Delta f$, i.e. $[\Delta, L] = 2(n+d)$. From this we find inductively:

$$\Delta L^{k+1} = C_k L^k + L^{k+1} \Delta,$$

where $C_k \neq 0$. This shows:

$$\Delta(r^{2k+2} H_d) = r^{2k} H_d$$

(and of course the map is an isomorphism because both sides have the same dimension).

3. We can now prove by induction:

$$P_d = H_d \oplus r^2 H_{d-2} \oplus r^4 H_{d-4} \oplus \cdots$$

Indeed, once this is known for P_d we simply consider $\Delta : P_{d+2} \rightarrow P_d$. This map has kernel H_{d+2} and maps $r^2 H_d$ bijectively to H_d , etc.

4. As a Corollary we immediately see that $\oplus H_d|S^{n-1}$ is the same space of functions as $\oplus P_d|S^{n-1}$, since $r = 1$ on S^{n-1} . In particular, $\oplus H_d$ is dense in $L^2(S^{n-1})$.
5. It remains to check that H_d and H_e are orthogonal for $d \neq e$. One way is to consider the spherical Laplacian and note that these are eigenspaces with different eigenvalues. Another way is to consider the character of $\text{SO}(2)$ acting on H_d .
6. The combination of these observations proves the spherical harmonics form a basis for $L^2(S^{n-1})$.

Low-dimensional examples. For example, when $n = 2$ we have $\dim H_0 = 1$ and $\dim H_d = 2$ for $d > 0$. A basis is given by $\text{Re } z^d$ and $\text{Im } z^d$.

For $n = 3$ we have $\dim h_d = 2d + 1 = 1, 3, 5, \dots$. It is traditional to form a complex basis Y_{md} for H_d where $-d \leq m \leq d$, and

$$Y_{md}(x, y, z) = (x \pm iy)^{|m|} P_d^m(z).$$

Here $P_d^m(z)$ is a Legendre polynomial.

The hydrogen atom. The simplest model for the hydrogen atom in quantum mechanics has as states of pure energy the functions f on \mathbb{R}^3 which satisfy

$$\Delta f + r^{-1}f = Ef.$$

It turns out a basis for such functions has the form of products of radial functions with spherical harmonics. The energy is proportional to $1/N^2$ where N is the principal quantum number. For a given N , the harmonics with $0 \leq d < N - 1$ all arise, each with multiplicity $2d + 1$, so there are N^2 independent states altogether. The states with $d = 0, 1, 2, 3, \dots$ are traditionally labeled s, p, d, f, g, h .

Irreducibility. Is there a finer Fourier series that is still natural with respect to rotations? The answer is no:

Theorem 12.4 *The action of $\text{SO}(n)$ on H_d is irreducible.*

Proof. There are many proofs of irreducibility; here is a rather intuitive, analytic one.

Suppose the action of $\text{SO}(n)$ on H_d splits nontrivially as $A \oplus B$. Then we can find in each subrepresentation a function such that $f(N) = 1$, where N is the ‘north pole’ stabilized by $\text{SO}(n - 1)$; and by averaging over $\text{SO}(n - 1)$,

we can assume f is constant on each sphere $S^{n-2} \subset S^{n-1}$ centered at N . In particular, if we consider a ball $B \subset S^{n-1}$ centered at N and of radius $\epsilon > 0$, we can find a nonzero $f \in H_d$ with $f|_{\partial B} = 0$ and $\max f|_B = 1$.

Consider the cone $U = [0, 2]B \subset \mathbb{R}^n$. Then f is a harmonic function which vanishes on all of the boundary of B except the cap $2B$. By homogeneity, $\max f|_{2B} = 2^d$. In addition, there is an $x \in B$ where $f(x) = 1$. By the mean value property of harmonic functions, $f(x)$ is the average of the values $f(y)$ over the points y where a random path initiated at x first exits U . But the probability that the path exits through the cap $2B$ is $p(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. Thus

$$1 = f(x) \leq 2^d p(\epsilon) \rightarrow 0,$$

a contradiction. ■

(Note: this argument gives *a priori* control over the diameter of a closed ‘nodal set’ for an eigenfunction of the Laplacian on S^{n-1} in terms of its eigenvalue.)

Spherical Laplacian. Here is a useful computation for understanding spherical harmonics intrinsically.

To compute the Laplacian of $f|_{S^{n-1}}$, we use the formula:

$$\Delta_s(f) = \nabla \cdot \pi_s(\nabla f),$$

where

$$\pi_s(\nabla f) = \nabla f - (\hat{r} \cdot \nabla f)\hat{r}$$

is the projection of ∇f to a vector field tangent to the sphere. Using the fact that $\nabla \cdot \hat{r} = n - 1$, this gives:

$$\Delta_s(f) = \Delta(f) - (n - 1)(df/dr) - d^2f/dr^2.$$

Now suppose f is a spherical harmonic of degree ℓ . Then $\Delta(f) = -\ell(\ell + n - 2)f$, $df/dr = \ell f$, and $d^2f/dr^2 = \ell(\ell - 1)f$, which yields:

Theorem 12.5 *If $f \in H_\ell(\mathbb{R}^n)$ then $f|_{S^{n-1}}$ is an eigenfunction of the spherical Laplacian, satisfying*

$$\Delta_s(f) = -\ell(\ell + n - 2)f.$$

13 General Measure Theory

Measures. A measure (X, \mathcal{B}, m) consists of a map $m : \mathcal{B} \rightarrow [0, \infty]$ defined on a σ -algebra of subsets of X , such that $m(\emptyset) = 0$ and such that $\sum m(B_i) = m(\bigcup B_i)$ for countable unions of disjoint $B_i \in \mathcal{B}$.

Countable/Co-countable measure. An example is the measure defined on any uncountable set X by taking \mathcal{B} to be the σ -algebra generated by singletons and $m(B) = 0$ or ∞ depending on whether B is countable or $X - B$ is countable.

Hausdorff measure. This is defined on the Borel subsets of \mathbb{R}^n by

$$m_\delta(E) = \lim_{r \rightarrow 0} \inf_{E = \bigcup E_i} \sum \text{diam}(E_i)^\delta,$$

where $\text{diam}(E_i) \leq r$. Appropriately scaled, m_n is equal to the usual volume measure on \mathbb{R}^n .

Dimension; the Cantor set. The Hausdorff *dimension* of $E \subset \mathbb{R}^n$ is the infimum of those δ such that $m_\delta(E) = 0$.

For example, the usual Cantor set E can be covered by 2^n intervals of length $1/3^n$, so its dimension is at most $\log 2 / \log 3$. On the other hand, there is an obvious measure on E such that $m(A) \leq C(\text{diam } E)^{\log 2 / \log 3}$ and from this it is easy to prove the dimension is equal to $\log 2 / \log 3$.

Linear maps and dimension. Clearly Hausdorff measure satisfies $m_\delta(\alpha E) = \alpha^\delta m_\delta(E)$. So for the Cantor set E built on disjoint subintervals of lengths a, b $a + b < 1$ in $[0, 1]$, one has $a^\delta + b^\delta = 1$ if $0 < m_\delta(E) < \infty$.

This makes it easy to guess the dimension of self-similar fractals. The self-affine case is much harder; cf. the M curve, of dimension $1 + 2^{\log 2 / \log 3}$.

Signed measures. To make the space of all measure into a linear space, we must allow measures to assume negative values.

A *finite signed measure* m on a σ -algebra \mathcal{B} is a map $m : \mathcal{B} \rightarrow [-M, M]$, such that for any sequence of disjoint B_i we have

$$\sum m(B_i) = m(\bigcup B_i).$$

Note that the sum above converges absolutely, since the sum of its positive terms individually is bounded above by M , and similar for the negative terms.

A general signed measure is allowed to assume at most one of the values $\pm\infty$, and the sum above is required to converge absolutely when $m(\bigcup B_i)$ is finite.

A *measure* is a signed measure assuming no negative values.

For simplicity we will restrict attention to **finite** signed measures.

Positive sets. Given a signed measure m , a set P is *positive* if $m(A) \geq 0$ for all $A \subset P$.

Theorem 13.1 *If $m(A) > 0$ then there is a positive set $P \subset A$ with $m(P) \geq m(A)$.*

Proof. Let $\lambda(A) = \inf\{m(B) : B \subset A\} \geq -M$. Pick a set of nonpositive measure, $B_1 \subset A$, with $m(B_1) < \lambda(A) + 1$. By induction construct a set of nonpositive measure $B_{n+1} \subset A_n = A - (B_1 \cup \dots \cup B_n)$ with $m(B_{n+1}) < \lambda(A_n) - 1/n$. Then $\sum |m(B_i)| < \infty$, so $m(B_i) \rightarrow 0$ and thus $\lambda(A_i) \rightarrow 0$.

Letting $P = \bigcap A_n$, we have $P \subset A_n$ so $\lambda(P) \geq \lim \lambda(A_n) = 0$. Thus P is a positive set, and $m(P) \geq m(A)$ since $m(B_i) \leq 0$ for each i . ■

The Hahn Decomposition.

Theorem 13.2 *Given a finite signed measure m on X , there is a partition of X into a pair A, B of disjoint sets, one positive and one negative.*

Proof. Let $p = \sup m(P)$ over all positive sets $P \subset X$. We claim p is achieved for some positive set A . Indeed, we can choose positive sets A_i with $m(A_i) \rightarrow p$ and just let $A = \bigcup A_i$.

Now let $B = X - A$. Then B contains a set of positive measure, then it contains a positive set P of positive measure; then $m(A \cup P) > m(A) = p$, contrary to the definition of p . Thus B is negative. ■

Jordan decomposition.

Theorem 13.3 *Let m be a signed measure on X . Then m can be uniquely expressed as $m = p - n$, where p and n are mutually singular (positive) measures.*

Here *mutually singular* means p and n are supported on disjoint sets.

Proof. Let $p = m|_A$ and $n = -m|_B$, where $A \cup B$ is the Hahn decomposition of X (unique up to null sets). This shows p and n of the required form exist.

Now assuming only that $m = p - n$, where p and n are mutually singular, we can assert that $p(A) = \sup\{m(B) : B \subset A\}$, and thus p is unique. Similarly n is unique. ■

Absolute continuity. Given a pair of measures μ and λ , we say $\mu \ll \lambda$, or μ is *absolutely continuous* with respect to λ , if $\lambda(E) = 0 \implies \mu(E) = 0$.

For example, $X = [0, 1]$ and $\mu(E) = \int_E f(x) dx$ for $f \in L^1[0, 1]$, then $\mu \ll \lambda$ if λ is Lebesgue measure on $[0, 1]$. In fact the converse holds.

The Radon-Nikodym theorem.

Theorem 13.4 *If $\mu \ll \lambda$ then there is an $f \geq 0$ such that*

$$\mu(E) = \int_E f(x) d\lambda.$$

Proof. If f has the form above, then the Hahn decomposition of μ is $\{f < 0\} \cup \{f > 0\}$. Similarly the Hahn decomposition of $\mu - \alpha\lambda$ is $\{f < \alpha\} \cup \{f > \alpha\}$.

So for each rational number α , let P_α be the positive set for the Hahn decomposition of $\mu - \alpha\lambda$. Then the P_α are nested (up to null sets). Define $f(x) = \sup\{\alpha : x \in P_\alpha\}$, and set $\nu(A) = \int_A f d\lambda$.

Now notice that for $\alpha < \beta$, for any A contained in

$$\{\alpha \leq f \leq \beta\} = P_\alpha - P_\beta,$$

we have $\nu(A)$ and $\mu(A)$ both contained in $[\alpha, \beta]\lambda(A)$. Chopping $[0, \infty]$ into intervals of length $1/n$, and pulling these intervals back to a decomposition E_i of a set E , we find that $\mu(E)$ is sandwiched between the upper and lower approximations to $\int_E f d\lambda$. Therefore equality holds. ■

Derivatives. The function f defined above is commonly written $f = d\mu/d\lambda$, so we have

$$\mu = \frac{d\mu}{d\lambda} d\lambda.$$

Absolutely continuous/singular decomposition. Given a pair of measures μ, ν on (X, \mathcal{B}) , we can naturally decompose $\nu = \nu_a + \nu_s$ with $\nu_a \ll \mu$ and $\nu_s \perp \mu$.

To do this, just let $\pi = \mu + \nu$, and write $d\mu = f d\pi$, $\nu = g d\pi$ (using Radon-Nikodym derivatives). Then we have $\nu = g/f d\mu$ on the set where $f > 0$, and $\nu \perp \mu$ on the set where $f = 0$. These two restrictions give the desired decomposition of ν .

Baire measures. We now pass to the consideration of measures μ on a *compact* Hausdorff space X compatible with the topology. The natural domain of such a measure is not the Borel sets but the **Baire** sets \mathcal{K} , the

smallest σ -algebra such that all $f \in C(X)$ are measurable.

A *Baire measure* is a measure m on (X, \mathcal{K}) .

What's the distinction? In \mathbb{R} , all closed sets are G'_δ 's, so their preimages under functions are also G_δ . Thus \mathcal{K} is generated by the closed G_δ 's in X , rather than all closed sets.

In a compact metric space, the Borel and Baire sets coincide.

Regular contents. It is useful to have a characterization of those functions $\lambda : \mathcal{F} \rightarrow \mathbb{R}$ defined on the closed (hence compact) sets \mathcal{F} in X such that λ extends to a Baire measure. Here it is:

Theorem 13.5 *Let $\lambda(K) \geq 0$ be defined for all compact G_δ sets $K \subset X$ and satisfy:*

- (i) $\lambda(K_1) \leq \lambda(K_2)$ if $K_1 \subset K_2$;
- (ii) $\lambda(K_1 \cup K_2) = \lambda(K_1) + \lambda(K_2)$ if K_1 and K_2 are disjoint; and $\lambda(K) = \inf \lambda(\bar{U})$ over all open sets $U \supset K$.

Then there is a unique Baire measure μ such that $\mu(K) = \lambda(K)$ for all compact K .

Such a λ is called a *regular content* on X .

Sketch of the proof. Given λ , we can define a set-function (inner measure) by

$$\mu_*(E) = \sup_{K \subset E} \lambda(K),$$

define a set A to be measurable if $\mu_*(A \cap E) + \mu_*((X - A) \cap E) = \mu_*(E)$ for all E , show that the measurable sets contain the Baire sets and that $\mu = \mu_*$ is a Baire measure extending λ .

Positive functionals.

Theorem 13.6 *Let $\phi : C(X) \rightarrow \mathbb{R}$ be a linear map such that $f \geq 0 \implies \phi(f) \geq 0$. Then there is a unique Baire measure μ on X such that*

$$\phi(f) = \int_X f d\mu.$$

Proof. Let us say $f \in C(X)$ is *admissible* for a compact G_δ set K if $f \geq 0$ and $f \geq 1$ on K . Define $\lambda(K)$ as $\inf \phi(f)$ over all admissible f .

We claim λ is a regular content. (i) is clear; as for (ii), $\lambda(K_1 \cup K_2) \leq \lambda(K_1) + \lambda(K_2)$ is obvious. For the reverse inequality, use normality of X to get $g_1 + g_2 = 1$, $g_1 g_2 = 0$, $0 \leq g_i \leq 1$ with $g_i = 1$ on K_i . Then given f for

$K_1 \cup K_2$ we get competitors $f_i = g_i f$ for K_i with $\phi(f_1) + \phi(f_2) \leq \phi(f)$, so $\lambda(K_1 \cup K_2) \leq \lambda(K_1 \cup K_2)$.

Finally for (iii): choose f admissible for K with $\phi(f) \leq \lambda(K) + \epsilon$. Let $U = \{f > 1 - \epsilon\}$. Then $f/(1 - \epsilon)$ is admissible for \bar{U} , so $\lambda(\bar{U}) \leq (\lambda(K) + \epsilon)/(1 - \epsilon)$, and therefore $\lambda(K) = \inf \lambda(\bar{U})$.

Thus λ extends to a Baire measure μ . To show that integration against μ reproduces ϕ , first note that for any K there exist admissible f_n decreasing to χ_K pointwise, with f_n eventually vanishing on any compact set L disjoint from K (since K is a G_δ), and for which $\int f_n$ and $\phi(f_n)$ both converge to $\mu(K) = \lambda(K)$.

Thus we can approximate χ_K by a continuous function f with $\phi(f) \approx \int f$. Now approximate g from above by sums of indicator functions of compact sets, and approximate these from above by admissible functions f ; then we get $\phi(g) \leq \phi(f) \approx \int g d\mu$. Doing the same from below we find that $\phi(g) = \int g d\mu$. ■

Theorem 13.7 (Riesz) *Let X be a compact Hausdorff space. Then the dual of $C(X)$ is the space of Baire measures on X , with $\mu = |\mu|(X)$.*

Proof. One shows a linear functional can be decomposed into a positive and negative part, each of which is represented by a measure. ■

Corollary. The space of measures on a compact Hausdorff space is compact in the weak* topology.

Functions of bounded variation and signed measures on $[a, b]$. We can now address afresh the theory of differentiation of $f : [a, b] \rightarrow \mathbb{R}$. To each signed measure μ we can associate the function $f(x) = \mu[a, x]$. This function is continuous from above and of bounded variation. Conversely, to each such f one can attach a measure df . The weak topology is the one where $f_n \rightarrow f$ iff $f_n(x) \rightarrow f(x)$ for each x such that f is continuous at x .

Now signed measure correspond to functions in BV; absolutely continuous measures, to absolutely continuous functions; $f'(x)$ is $d\mu_a/d\lambda$; discontinuities correspond to atoms; singular measure correspond to f with $f' = 0$ a.e.

Compactness. As an alternative proof of compactness: consider a sequence of monotone increasing functions $f : [a, b] \rightarrow [0, 1]$ with $f(b) = 1$. (I.e. a sequence of probability measures.) Passing to a subsequence, we can get $f_n(x)$ to converge for all rational $x \in [a, b]$. Then there is a monotone

limit g , which can be arranged to be right-continuous, such that $f_n \rightarrow g$ away from its discontinuities.

Integration. Given a function f of bounded variation and $g \in C^\infty[a, b]$, we can define

$$I = \int_a^b g(x) df(x) = - \int_a^b f(x)g'(x) dx.$$

Now breaking $[a, b]$ up into intervals $[a_i, a_{i+1}]$ we get the approximation:

$$\begin{aligned} I &= - \sum f(a_i)(g(a_{i+1}) - g(a_i)) \\ &= + \sum (f(a_{i+1}) - f(a_i))g(a_i) = O(\|f\|_{BV}\|g\|_\infty). \end{aligned}$$

Thus integration against df gives a bounded linear functional on a dense subset of $C[a, b]$, so it extends uniquely to a measure.

This idea is the beginnings of the theory of distributions.

Sample application: Let $f : X \rightarrow X$ be a homeomorphism. Then there exists a probability measure μ on X such that $\mu(A) = \mu(f(A))$.

Proof. Take any probability measure — such as a point mass δ ; average it over the first n iterates of f ; and take a weak* limit. ■

Haar measure. If G is a compact Hausdorff topological group, for each open neighborhood U of the origin we define $\lambda_U(K) = [K : U]/[G : U]$ where $[E : U]$ is the minimal number of left translates gU needed to cover E . Then as U shrinks towards the identity, we can extract **some** (Banach) limit of λ_U , which turns out to be a content λ . In this way we obtain a left-invariant measure on G .

A Measurable A with $A - A$ nonmeasurable

A standard homework exercise in real analysis is to show that $A - A$ contains a nonempty interval whenever $A \subset \mathbb{R}$ has positive Lebesgue measure. (For the proof one can use Littlewood's first principle, which states that A is nearly a finite union of intervals.)

Of course A does not need to have positive measure for its difference set to contain an interval; in fact the standard Cantor middle-thirds set $K \subset [0, 1]$ has measure zero, and $K - K = [-1, 1]$.

What is less well-known is that the standard dictum that 'reasonable operations preserve measurability' fails here: even if A is measurable, its

difference set $A - A$ need not be. The Cantor set example already hints that such a pathology might occur: any subset of K is measurable, but $K - K$ contains many nonmeasurable sets, so one might guess that there is an $A \subset K$ such that $A - A$ is nonmeasurable. We will show this is indeed the case.

The proof will be by transfinite induction and appeal to the axiom of choice, as many constructions of nonmeasurable sets do. It will also rely on the following property of the map

$$\pi : K \times K \rightarrow [-1, 1]$$

given by $\pi(x, y) = x - y$.

Proposition A.1 *Almost all fibers of π contain perfect sets. In particular, $|\pi^{-1}(t)| = |\mathbb{R}|$ for almost every $t \in [-1, 1]$.*

Sketch of the proof. It is simpler to treat the averaging map $f(x, y) = (x + y)/2$, which sends $K \times K$ to $[0, 1]$ and has the same behavior as π . Suppose $t \in [0, 1]$, and for convenience, suppose t is not a rational of the form $p/3^n$. Then t has a unique ternary expansion $t = 0.t_1t_2t_3\dots_3$; let $N(t)$ be the number of times the digit $t_i = 1$ appears. Then it is readily verified that the fiber $F_t = |f^{-1}(t)| = 2^{N(t)}$ if $N(t)$ is finite, and that $f^{-1}(t)$ is a perfect set if $N(t)$ is infinite. In particular, F_t is perfect for almost every $t \in [0, 1]$, but F_t consists of a single point (namely t itself) for $t \in K$ (again assuming t is not a triadic rational). The proof is suggested in Figure 5; note that there are 2 squares above $[1/3, 2/3]$ and only 1 above $[0, 1/3]$ and $[2/3, 1]$. ■

Using this fact, we will show:

Theorem A.2 *There exists a set $A \subset K$ such that $A - A$ has positive outer measure, but $A - A$ contains no perfect set.*

Proof. Let \mathfrak{c} denote the smallest ordinal with $|\mathfrak{c}| = |\mathbb{R}|$. Then $|\alpha| < |\mathbb{R}|$ for all $\alpha \in \mathfrak{c}$. Since the open and closed subsets of \mathbb{R} themselves have the cardinality of the continuum, we can index them by \mathfrak{c} , and similarly for the particular types of open and closed sets we consider below.

Thus we let $(P_\alpha : \alpha \in \mathfrak{c})$ denote an enumeration of the perfect subsets of \mathbb{R} , and we let $(U_\alpha : \alpha \in \mathfrak{c})$ denote an enumeration of the open sets in \mathbb{R} with $m(U_\alpha) < 2$.

Our goal is to define, by transfinite induction on \mathfrak{c} , an increasing sequence of sets $K_\alpha \subset K$ and elements $p_\alpha \in P_\alpha$ such that for all $\alpha \in \mathfrak{c}$:




Figure 5. The averaging map from $K \times K$ onto $[0, 1]$.

1. $K_{\alpha+1} - K_{\alpha+1}$ is not contained in U_α ; and
2. $K_\alpha - K_\alpha$ does not meet $Z_\alpha = \{p_\beta : \beta < \alpha\}$.

We then set $A = \bigcup_{\alpha \in \mathfrak{c}} K_\alpha$. By the first property, the outer measure of $A - A$ satisfies $m^*(A - A) \geq 2$; otherwise, $A - A$ would be contained in an open set U with $m(U) < 2$, and we would have $U = U_\alpha$ for some α , contradicting the fact that $K_{\alpha+1} - K_{\alpha+1} \subset A - A$ is not contained in U_α . By the second property, $A - A = \bigcup (K_\alpha - K_\alpha)$ is disjoint from $\bigcup_{\beta \in \mathfrak{c}} p_\beta$, and hence $A - A$ does not contain any perfect set P_β .

It remains to construct K_α and p_α satisfying the two conditions above. We start the induction with $K_0 = \emptyset$. Whenever we construct K_α , we construct p_α immediately afterwards by choosing any point

$$p_\alpha \in P_\alpha \setminus (K_\alpha - K_\alpha).$$

Such a point exists because $|P_\alpha| = |\mathbb{R}|$, while $|K_\alpha - K_\alpha| \leq |\alpha|^2 < |\mathbb{R}|$.

To construct K_α at a limit ordinal, we let

$$K_\alpha = \bigcup_{\beta < \alpha} K_\beta.$$

It is then immediate by induction that K_α satisfies the two conditions above.

For a successor ordinal $\alpha = \gamma + 1$, we define

$$K_\alpha = K_\gamma \cup \{x, y\}$$

for a suitable pair of points $x, y \in K$.

We will choose (x, y) such that $x - y \notin U_\gamma$; then the first condition will hold. To be more precise, using the proposition above and the fact that $m(U_\alpha) < 2$, we will pick

$$t \in [-1, 1] \setminus U_\alpha$$

such that $F_t = \pi^{-1}(t)$ contains a perfect set; in particular, $|F_t| = |\mathbb{R}|$. The first condition then holds, for any $(x, y) \in F_t$.

We also need to choose (x, y) such that $K_\alpha - K_\alpha$ is disjoint from $Z_\alpha = Z_\gamma \cup \{p_\gamma\}$. Now by induction, $K_\gamma - K_\gamma$ is disjoint from Z_γ . We have also defined p_γ so it is not in $K_\gamma - K_\gamma$, and therefore $K_\gamma - K_\gamma$ is disjoint from Z_α . On the other hand, we have

$$K_\alpha - K_\alpha = (K_\gamma - K_\gamma) + \{0, x, y, x + y, -x, -y, -x - y\}.$$

Thus we need to choose (x, y) such that

$$x \notin E = (K_\gamma - K_\gamma) + Z_\alpha,$$

and similarly for $y, x + y, -x, -y, -x - y$.

Now note that the projection $(x, y) \mapsto x$ is one-to-one on F_t , and $|E| < |\mathbb{R}| = |F_t|$. Thus we can easily choose $(x, y) \in F_t$ so that $\pm x, \pm y \notin E$.

There remains the problem of insuring that $x + y$ and $-x - y$ are not in E , in other words, of insuring that $\pi(x, y) \notin (E \cup -E)$. But this simply requires that we refine our choice of t , so that we also have

$$t \notin E \cup (-E).$$

Since $|E| < |\mathbb{R}|$, most points in $[-1, 1] \setminus U_\alpha$ have this property, and we are done. ■

Since any set of positive measure contains a perfect set, we have:

Corollary A.3 *There exists a measurable set $A \subset \mathbb{R}$ such that $A - A$ is not measurable.*

(Note: Hugh Woodin provided hints on how to proceed.)

References

- [Con] J. H. Conway. *On Numbers and Games*. Academic Press, 1976.
 [Hal] P. Halmos. *Naive Set Theory*. Springer-Verlag, 1974.

- [Me] R. Mañé. *Ergodic Theory and Differentiable Dynamics*. Springer-Verlag, 1987.
- [Ox] J. C. Oxtoby. *Measure and Category*. Springer-Verlag, 1980.

MATHEMATICS DEPARTMENT
HARVARD UNIVERSITY
1 OXFORD ST
CAMBRIDGE, MA 02138-2901