

Real Analysis: Part II

William G. Faris

June 3, 2004

Contents

1	Function spaces	1
1.1	Spaces of continuous functions	1
1.2	Pseudometrics and seminorms	2
1.3	\mathcal{L}^p spaces	2
1.4	Dense subspaces of \mathcal{L}^p	5
1.5	The quotient space L^p	6
1.6	Duality of L^p spaces	7
1.7	Orlicz spaces	8
1.8	Problems	10
2	Hilbert space	13
2.1	Inner products	13
2.2	Closed subspaces	16
2.3	The projection theorem	17
2.4	The Riesz-Fréchet theorem	18
2.5	Bases	19
2.6	Separable Hilbert spaces	21
2.7	Problems	23
3	Differentiation	27
3.1	The Lebesgue decomposition	27
3.2	The Radon-Nikodym theorem	28
3.3	Absolutely continuous increasing functions	29
3.4	Problems	31
4	Conditional Expectation	33
4.1	Hilbert space ideas in probability	33
4.2	Elementary notions of conditional expectation	35
4.3	The L^2 theory of conditional expectation	35
4.4	The L^1 theory of conditional expectation	37
4.5	Problems	39

5	Fourier series	41
5.1	Periodic functions	41
5.2	Convolution	42
5.3	Approximate delta functions	42
5.4	Summability	43
5.5	L^2 convergence	44
5.6	$C(T)$ convergence	46
5.7	Pointwise convergence	48
5.8	Problems	49
6	Fourier transforms	51
6.1	Fourier analysis	51
6.2	L^1 theory	52
6.3	L^2 theory	53
6.4	Absolute convergence	54
6.5	Fourier transform pairs	55
6.6	Poisson summation formula	56
6.7	Problems	57
7	Topology	59
7.1	Topological spaces	59
7.2	Measurable spaces	61
7.3	The Borel σ -algebra	62
7.4	Comparison of topologies	63
7.5	Bases and subbases	64
7.6	Compact spaces	65
7.7	The one point compactification	66
7.8	Nets	68
7.9	Problems	69
8	Product and weak* topologies	71
8.1	Introduction	71
8.2	The Tychonoff product theorem	72
8.3	Banach spaces and dual Banach spaces	73
8.4	Weak* topologies on dual Banach spaces	75
8.5	The Alaoglu theorem	76
8.6	Appendix: The Bourbaki fixed point theorem	79
8.7	Appendix: Zorn's lemma	80
9	Radon measures	83
9.1	Topology and measure	83
9.2	Borel isomorphisms	84
9.3	Metric spaces	85
9.4	Riesz representation	86
9.5	Lower semicontinuous functions	87
9.6	Weak* convergence	88

9.7	Central limit theorem for coin tossing	90
9.8	Wiener measure	91
9.9	Appendix: Measure theory on non-metric spaces	92
9.10	Problems	94
10	Distributions	95
10.1	Fréchet spaces	95
10.2	Distributions	96
10.3	Operations on distributions	97
10.4	Mapping distributions	98
10.5	The support of a distribution	99
10.6	Radon measures	99
10.7	Approximate delta functions	100
10.8	Tempered distributions	101
10.9	Poisson equation	105
10.10	Problems	107
11	Appendix: Standard Borel measurable spaces	109
11.1	Complete separable metric spaces	109
11.2	Embedding a Cantor set	110
11.3	Placement in a cube	111
11.4	A cube inside a Cantor set	111
11.5	A dynamical system	112
11.6	The unique Borel structure	112
11.7	The unique measure structure with given mass	113

Chapter 1

Function spaces

1.1 Spaces of continuous functions

This section records notations for spaces of real functions. In some contexts it is convenient to deal instead with complex functions; usually the changes that are necessary to deal with this case are minor.

Let X be a topological space. The space $C(X)$ consists of all continuous functions. The space $B(X)$ consists of all bounded functions. It is a Banach space in a natural way. The space $BC(X)$ consists of all bounded continuous functions. It is a somewhat smaller Banach space.

Consider now the special case when X is a locally compact Hausdorff space. Thus each point has a compact neighborhood. For example X could be \mathbf{R}^n . The space $C_c(X)$ consists of all continuous functions, each one of which has compact support. The space $C_0(X)$ is the closure of $C_c(X)$ in $BC(X)$. It is itself a Banach space. It is the space of continuous functions that vanish at infinity.

The relation between these spaces is that $C_c(X) \subset C_0(X) \subset BC(X)$. They are all equal when X compact. When X is locally compact, then $C_0(X)$ is the best behaved.

Recall that a Banach space is a normed vector space that is complete in the metric associated with the norm. In the following we shall need the concept of the *dual space* of a Banach space E . The dual space E^* consists of all continuous linear functions from the Banach space to the real numbers. (If the Banach space has complex scalars, then we take continuous linear function from the Banach space to the complex numbers.) The dual space E^* is itself a Banach space, where the norm is the Lipschitz norm.

For certain Banach spaces E of functions the linear functionals in the dual space E^* may be realized in a more concrete way. For example, If $E = C_0(X)$, then its dual space $E^* = M(X)$ is a Banach space consisting of signed Radon measures of finite variation. (A signed measure σ of finite variation is the difference $\sigma = \sigma_+ - \sigma_-$ of two positive measures σ_{\pm} that each assign finite measure

to the space X . Radon measures will be discussed later.) If σ is in $M(X)$, then it defines the linear functional $f \mapsto \int f(x) d\sigma(x)$, and all elements of the dual space E^* arise in this way.

1.2 Pseudometrics and seminorms

A pseudometric is a function $d : P \times P \rightarrow [0, +\infty)$ that satisfies $d(f, f) \geq 0$ and $d(f, g) \leq d(f, h) + d(h, g)$ and such that $d(f, f) = 0$. If in addition $d(f, g) = 0$ implies $f = g$, then d is a metric.

The theory of pseudometric spaces is much the same as the theory of metric spaces. The main difference is that a sequence can converge to more than one limit. However each two limits of the sequence have distance zero from each other, so this does not matter too much.

Given a pseudometric space P , there is an associated metric space M . This is defined to be the set of equivalence classes of P under the equivalence relation fEg if and only if $d(f, g) = 0$. In other words, one simply defines two points r, s in P that are at zero distance from each other to define the same point $r' = s'$ in M . The distance $d_M(a, b)$ between two points a, b in M is defined by taking representative points p, q in P with $p' = a$ and $q' = b$. Then $d_M(a, b)$ is defined to be $d(p, q)$.

A seminorm is a function $f \mapsto \|f\| \geq 0$ on a vector space E that satisfies $\|cf\| = |c|\|f\|$ and the triangle inequality $\|f + g\| \leq \|f\| + \|g\|$ and such that $\|0\| = 0$. If in addition $\|f\| = 0$ implies $f = 0$, then it is a norm. Each seminorm on E defines a pseudometric for E by $d(f, g) = \|f - g\|$. Similarly, a norm on E defines a metric for E .

Suppose that we have a seminorm on E . Then the set of h in E with $\|h\| = 0$ is a vector subspace of E . The set of equivalence classes in the construction of the metric space is itself a vector space in a natural way. So for each vector space with a seminorm we can associate a new quotient vector space with a norm.

1.3 \mathcal{L}^p spaces

In this and the next sections we introduce the spaces $\mathcal{L}^p(X, \mathcal{F}, \mu)$ and the corresponding quotient spaces $L^p(X, \mathcal{F}, \mu)$.

Fix a set X and a σ -algebra \mathcal{F} of measurable functions. Let $0 < p < \infty$. Define

$$\|f\|_p = \mu(|f|^p)^{\frac{1}{p}}. \quad (1.1)$$

Define $\mathcal{L}^p(X, \mathcal{F}, \mu)$ to be the set of all f in \mathcal{F} such that $\|f\|_p < \infty$.

Theorem 1.1 *For $0 < p < \infty$, the space \mathcal{L}^p is a vector space.*

Proof: It is obvious that L^p is closed under scalar multiplication. The problem is to prove that it is closed under addition. However if f, g are each in

L^p , then

$$|f + g|^p \leq [2(|f| \vee |g|)]^p \leq 2^p(|f|^p + |g|^p). \quad (1.2)$$

Thus $f + g$ is also in L^p . \square

The function x^p is increasing for every $p > 0$. In fact, if $\phi(p) = x^p$ for $x \geq 0$, then $\phi'(p) = px^{p-1} \geq 0$. However it is convex only for $p \geq 1$. This is because in that case $\phi''(x) = p(p-1)x^{p-2} \geq 0$.

Let $a \geq 0$ and $b \geq 0$ be weights with $a + b = 1$. For a convex function we have the inequality $\phi(au + bv) \leq a\phi(u) + b\phi(v)$. This is the key to the following fundamental inequality.

Theorem 1.2 (*Minkowski's inequality*) *If $1 \leq p < \infty$, then*

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p. \quad (1.3)$$

Proof: Let $c = \|f\|_p$ and $d = \|g\|_p$. Then by the fact that x^p is increasing and convex

$$\left| \frac{f+g}{c+d} \right|^p \leq \left(\frac{c}{c+d} \left| \frac{f}{c} \right| + \frac{d}{c+d} \left| \frac{g}{d} \right| \right)^p \leq \frac{c}{c+d} \left| \frac{f}{c} \right|^p + \frac{d}{c+d} \left| \frac{g}{d} \right|^p. \quad (1.4)$$

Integrate. This gives

$$\mu \left(\left| \frac{f+g}{c+d} \right|^p \right) \leq 1. \quad (1.5)$$

Thus $\|f + g\|_p \leq c + d$. \square

The preceding facts show that \mathcal{L}^p is a vector space with a seminorm. It is a fact that $\mu(|f|^p) = 0$ if and only if $f = 0$ almost everywhere. Thus for f in \mathcal{L}^p we have $\|f\|_p = 0$ if and only if $f = 0$ almost everywhere.

Theorem 1.3 (*dominated convergence for \mathcal{L}^p*) *Let $0 < p < \infty$. Let $f_n \rightarrow f$ pointwise. Suppose that there is a $g \geq 0$ in \mathcal{L}^p such that each $|f_n| \leq g$. Then $f_n \rightarrow f$ in \mathcal{L}^p , that is, $\|f_n - f\|_p \rightarrow 0$.*

Proof: If each $|f_n| \leq g$ and $f_n \rightarrow f$ pointwise, then $|f| \leq g$. Thus $|f_n - f| \leq 2g$ and $|f_n - f|^p \leq 2^p g^p$. Since g^p has finite integral, the integral of $|f_n - f|^p$ approaches zero, by the usual dominated convergence theorem. \square

It would be an error to think that just because $g_n \rightarrow g$ in the \mathcal{L}^p sense it would follow that $g_n \rightarrow g$ almost everywhere. Being close on the average does not imply being close at a particular point. Consider the following example. For each $n = 1, 2, 3, \dots$, write $n = 2^k + j$, where $k = 0, 1, 2, 3, \dots$ and $0 \leq j < 2^k$. Consider a sequence of functions defined on the unit interval $[0, 1]$ with the usual Lebesgue measure. Let $g_n = 1$ on the interval $[j/2^k, (j+1)/2^k]$ and $g_n = 0$ elsewhere in the unit interval. Then the \mathcal{L}^1 seminorm of g_n is $1/2^k$, so the $g_n \rightarrow 0$ in the \mathcal{L}^1 sense. On the other hand, given x in $[0, 1]$, there are infinitely many n for which $g_n(x) = 0$ and there are infinitely many n for which $g_n(x) = 1$. So pointwise convergence fails at each point.

Say that a seminormed vector space is sum complete if every absolutely convergent series is convergent to some limit. Recall that it is complete (as a pseudometric space) if every Cauchy sequence converges to some limit.

Lemma 1.4 *Consider a seminormed vector space. If the space is sum complete, then it is complete.*

Proof: Suppose that E is a seminormed vector space that is sum complete. Suppose that g_n is a Cauchy sequence. This means that for every $\epsilon > 0$ there is an N such that $m, n \geq N$ implies $\|g_m - g_n\| < \epsilon$. The idea is to show that g_n has a subsequence that converges very rapidly. Let ϵ_k be a sequence such that $\sum_{k=1}^{\infty} \epsilon_k < \infty$. In particular, for each k there is an N_k such that $m, n \geq N_k$ implies $\|g_m - g_n\| < \epsilon_k$. The desired subsequence is the g_{N_k} . Define a sequence $f_1 = g_{N_1}$ and $f_j = g_{N_j} - g_{N_{j-1}}$ for $j \geq 2$. Then

$$g_{N_k} = \sum_{j=1}^k f_j. \quad (1.6)$$

Furthermore,

$$\sum_{j=1}^{\infty} \|f_j\| \leq \|f_1\| + \sum_{j=2}^{\infty} \epsilon_{j-1} < \infty. \quad (1.7)$$

This says that the series is absolutely convergence. Since E is sum complete, the series converges to some limit, that is, there exists a g such that the subsequence g_{N_k} converges to g . Since the sequence g_n is Cauchy, it also must converge to the same g . Thus E is complete. Thus the theorem follows. \square

Theorem 1.5 *For $1 \leq p < \infty$ the space \mathcal{L}^p is complete.*

Proof: Suppose that $\sum_{j=1}^{\infty} f_j$ is absolutely convergent in \mathcal{L}^p , that is,

$$\sum_{j=1}^{\infty} \|f_j\|_p = B < \infty. \quad (1.8)$$

Then by using Minkowski's inequality

$$\left\| \sum_{j=1}^k f_j \right\|_p \leq \sum_{j=1}^k \|f_j\|_p \leq B. \quad (1.9)$$

By the monotone convergence theorem $h = \sum_{j=1}^{\infty} |f_j|$ is in \mathcal{L}^p with \mathcal{L}^p seminorm bounded by B . In particular, it is convergent almost everywhere. It follows that the series $\sum_{j=1}^{\infty} f_j$ converges almost everywhere to some limit g . The sequence $\sum_{j=1}^k f_j$ is dominated by h in \mathcal{L}^p and converges pointwise to $\sum_{j=1}^{\infty} f_j$. Therefore, by the dominated convergence theorem, it converges to the same limit g in the \mathcal{L}^p seminorm. \square

Corollary 1.6 *If $1 \leq p < \infty$ and if $g_n \rightarrow g$ in the \mathcal{L}^p seminorm sense, then there is a subsequence g_{N_k} such that g_{N_k} converges to g almost everywhere.*

Proof: Let $g_n \rightarrow g$ as $n \rightarrow \infty$ in the \mathcal{L}^p sense. Then g_n is a Cauchy sequence in the \mathcal{L}^p sense. Let ϵ_k be a sequence such that $\sum_{k=1}^{\infty} \epsilon_k < \infty$. Let N_k be a subsequence such that $n \geq N_k$ implies $\|g_n - g_{N_k}\|_p < \epsilon_k$. Define a sequence f_k such that

$$g_{N_k} = \sum_{j=1}^k f_j. \quad (1.10)$$

Then $\|f_j\|_p = \|g_{N_j} - g_{N_{j-1}}\|_p \leq \epsilon_{j-1}$ for $j \geq 2$. By the monotone convergence theorem

$$h = \sum_{j=1}^{\infty} |f_j| \quad (1.11)$$

converges in \mathcal{L}^p and is finite almost everywhere. It follows that

$$g = \sum_{j=1}^{\infty} f_j \quad (1.12)$$

converges in \mathcal{L}^p and also converges almost everywhere. In particular, $g_{N_k} \rightarrow g$ as $k \rightarrow \infty$ almost everywhere. \square

In order to complete the picture, define

$$\|f\|_{\infty} = \inf\{M \geq 0 \mid |f| \leq M \text{ almost everywhere}\}. \quad (1.13)$$

This says that $\|f\|_{\infty} \leq M$ if and only if $\mu(|f| > M) = 0$. In other words, $M < \|f\|_{\infty}$ if and only if $\mu(|f| > M) > 0$. The space $\mathcal{L}^{\infty}(X, \mathcal{F}, \mu)$ consists of all functions f in \mathcal{F} such that $\|f\|_{\infty} < \infty$. It is a vector space with a seminorm. The following theorem is also simple.

Theorem 1.7 *The space $\mathcal{L}^{\infty}(X, \mathcal{F}, \mu)$ is complete.*

1.4 Dense subspaces of \mathcal{L}^p

For a function to be in $\mathcal{L}^p(X, \mathcal{F}, \mu)$ it is not only required that $\mu(|f|^p) < \infty$, but also that f is measurable, that is, that f is in \mathcal{F} . This requirement has important consequences for approximation.

Theorem 1.8 *Let X be a set, \mathcal{F} a σ -algebra of real measurable functions on X , and μ an integral. Consider the space $\mathcal{L}^p(X, \mathcal{F}, \mu)$ for $1 \leq p < \infty$. Let L be a vector lattice of functions with $L \subset \mathcal{L}^p(X, \mathcal{F}, \mu)$. Suppose that the smallest monotone class including L is \mathcal{F} . Then L is dense in $\mathcal{L}^p(X, \mathcal{F}, \mu)$. That is, if f is in $\mathcal{L}^p(X, \mathcal{F}, \mu)$ and if $\epsilon > 0$, then there exists h in L with $\|h - f\|_p < \epsilon$.*

Proof: We know that the smallest monotone class including L^+ is \mathcal{F}^+ . Let g be in L^+ . Let S_g be the set of all $f \geq 0$ such that $f \wedge g$ is in the \mathcal{L}^p closure of L . Clearly $L^+ \subset S_g$, since if f is in L^+ then so is $f \wedge g$. Furthermore, S_g is closed under increasing and decreasing limits. Here is the proof for increasing

limits. Say that f_n is in S_g and $f_n \uparrow f$. Then $f_n \wedge g \uparrow f \wedge g \leq g$. By the \mathcal{L}^p monotone convergence theorem, $\|f_n \wedge g - f \wedge g\|_p \rightarrow 0$. Since $f_n \wedge g$ is in the \mathcal{L}^p closure of L , it follows that $f \wedge g$ is also in the \mathcal{L}^p closure of L . However this says that f is in S_g . It follows from this discussion that $\mathcal{F}^+ \subset S_g$. Since g is arbitrary, this proves that f in \mathcal{F}^+ and g in L^+ implies $f \wedge g$ is in the \mathcal{L}^p closure of L .

Let $f \geq 0$ be in \mathcal{L}^p . Let S'_f be the set of all $h \geq 0$ such that $f \wedge h$ is in the \mathcal{L}^p closure of L . From the preceding argument, we see that $L \subset S'_f$. Furthermore, S'_f is closed under increasing and decreasing limits. Here is the proof for increasing limits. Say that h_n is in S'_f and $h_n \uparrow h$. Then $f \wedge h_n \uparrow f \wedge h \leq f$. By the \mathcal{L}^p monotone convergence theorem, $\|f \wedge h_n - f \wedge h\|_p \rightarrow 0$. Since $f \wedge h_n$ is in the \mathcal{L}^p closure of L , it follows that $f \wedge h$ is also in the \mathcal{L}^p closure of L . However this says that h is in S'_f . It follows from this discussion that $\mathcal{F}^+ \subset S'_f$. Since f is arbitrary, this proves that $f \geq 0$ in \mathcal{L}^p and h in \mathcal{F}^+ implies $f \wedge h$ is in the \mathcal{L}^p closure of L . Take $h = f$. Thus $f \geq 0$ in \mathcal{L}^p implies f is in the \mathcal{L}^p closure of L . \square

Corollary 1.9 *Take $1 \leq p < \infty$. Consider the space $\mathcal{L}^p(\mathbf{R}, \mathcal{B}, \mu)$, where μ is a measure that is finite on compact subsets. Let L be the space of step functions, or let L be the space of continuous functions with compact support. Then L is dense in $\mathcal{L}^p(\mathbf{R}, \mathcal{B}, \mu)$.*

This result applies in particular to the case $\mu = \lambda$ of Lebesgue measure. Notice that nothing like this is true for $\mathcal{L}^\infty(\mathbf{R}, \mathcal{B}, \lambda)$. The uniform limit of a sequence of continuous functions is continuous, and so if we start with continuous functions and take uniform limits, we stay in the class of continuous functions. But functions in $\mathcal{L}^\infty(\mathbf{R}, \mathcal{B}, \lambda)$ can be discontinuous in such a way that cannot be fixed by changing the function on a set of measure zero. Even a step function has this property.

Remark. So people might argue that the so-called delta function $\delta(x)$ is in \mathcal{L}^1 , since it has integral $\int_{-\infty}^{\infty} \delta(x) = 1$. Actually the delta function is a measure, not a function, so this is not correct. But there is a stronger sense in which this is not correct. Let h be an arbitrary continuous function with compact support. Look at the distance from $\delta(x)$ to $h(x)$. This is the integral $\int_{-\infty}^{\infty} |\delta(x) - h(x)| dx$ which always has a value one or bigger. The delta function is thus not even close to being in \mathcal{L}^1 .

1.5 The quotient space L^p

The space $\mathcal{L}^p(X, \mathcal{F}, \mu)$ is defined for $1 \leq p \leq \infty$. It is a vector space with a seminorm, and it is complete. One can associate with this the space $L^p(X, \mathcal{F}, \mu)$, where two elements f, g of $\mathcal{L}^p(X, \mathcal{F}, \mu)$ define the same element of $L^p(X, \mathcal{F}, \mu)$ provided that $f = g$ almost everywhere with respect to μ . Then this is a vector space with a norm, and it is complete. In other words, it is a Banach space.

This passage from a space of functions \mathcal{L} to the corresponding quotient space L^p is highly convenient, but also confusing. The elements of L^p are not

functions, and so they do not have values defined at particular points of X . Nevertheless they are come from functions.

It is convenient to work with the spaces L^p abstractly, but to perform all calculations with the corresponding functions in \mathcal{L}^p . For this reason people often use the notation L^p to refer to either space, and we shall follow this practice in most of the following, unless there is a special point to be made.

However be warned, these spaces can be very different. As an example, take the space $\mathcal{L}^1(\mathbf{R}, \mathcal{B}, \delta_a)$, where δ_a is the measure that assigns mass one to the point a . Thus the corresponding integral is $\delta_a(f) = f(a)$. This space consists of all Borel functions, so it is infinite dimensional. However two such functions are equal almost everywhere with respect to δ_a precisely when they have the same values at the point a . Thus the quotient space $L^1(\mathbf{R}, \mathcal{B}, \delta)$ is one dimensional. This is a much smaller space. But it captures the notion that from the perspective of the measure δ_a the points other than a are more or less irrelevant.

1.6 Duality of L^p spaces

In this section we describe the duality theory for the Banach spaces L^p .

Lemma 1.10 (*arithmetic-geometric mean inequality*) *Let $a \geq 0$ and $b \geq 0$ with $a + b = 1$. Let $z > 0$ and $w > 0$. Then $z^a w^b \leq az + bw$.*

Proof: Since the exponential function is convex, we have $e^{au+bv} \leq ae^u + be^v$. Set $z = e^u$ and $w = e^v$. \square

Lemma 1.11 *Let $p > 1$ and $q > 1$ with $1/p + 1/q = 1$. If $x > 0$ and $y > 0$, then*

$$xy \leq \frac{1}{p}x^p + \frac{1}{q}y^q. \quad (1.14)$$

Proof: Take $a = 1/p$ and $b = 1/q$, and substitute $e^a = x$ and $e^b = y$. \square

Theorem 1.12 (*Hölder's inequality*) *Suppose that $1 < p < \infty$ and that $1/p + 1/q = 1$. Then*

$$|\mu(fg)| \leq \|f\|_p \|g\|_q. \quad (1.15)$$

Proof: It is sufficient to prove this when $\|f\|_p = 1$ and $\|g\|_p = 1$. However by the lemma

$$|f(x)||g(x)| \leq \frac{1}{p}|f(x)|^p + \frac{1}{q}|g(x)|^q. \quad (1.16)$$

Integrate. \square

This lemma shows that if g is in $L^q(\mu)$, with $1 < q < \infty$, then the linear functional defined by $f \mapsto \mu(fg)$ is continuous on $L^p(\mu)$, where $1 < p < \infty$ with $1/p + 1/q = 1$. This shows that each element of $L^q(\mu)$ defines an element of the dual space of $L^p(\mu)$. It may be shown that every element of the dual space arises in this way. Thus the dual space of $L^p(\mu)$ is $L^q(\mu)$, for $1 < p < \infty$.

Notice that we also have a Hölder inequality in the limiting case:

$$|\mu(fg)| \leq \|f\|_1 \|g\|_\infty. \quad (1.17)$$

This shows that every element g of $L^\infty(\mu)$ defines an element of the dual space of $L^1(\mu)$. It may be shown that if μ is σ -finite, then $L^\infty(\mu)$ is the dual space of $L^1(\mu)$.

On the other hand, each element f of $L^1(\mu)$ defines an element of the dual space of $L^\infty(\mu)$. However in general this does not give all elements of the dual space of $L^\infty(\mu)$.

The most important spaces are L^1 , L^2 , and L^∞ . The nicest by far is L^2 , since it is a Hilbert space. The space L^1 is also common, since it measures the total amount of something. The space L^∞ goes together rather naturally with L^1 . Unfortunately, the theory of the spaces L^1 and L^∞ is more delicate than the theory of the spaces L^p with $1 < p < \infty$. Ultimately this is because the spaces L^p with $1 < p < \infty$ have better convexity properties.

Here is a brief summary of the facts about duality. The dual space of a Banach space is the space of continuous linear scalar functions on the Banach space. The dual space of a Banach space is a Banach space. Let $1/p + 1/q = 1$, with $1 \leq p < \infty$ and $1 < q \leq \infty$. (Require that μ be σ -finite when $p = 1$.) Then the dual of the space $L^p(X, \mathcal{F}, \mu)$ is the space $L^q(X, \mathcal{F}, \mu)$. The dual of $L^\infty(X, \mathcal{F}, \mu)$ is not in general equal to $L^1(X, \mathcal{F}, \mu)$. Typically $L^1(X, \mathcal{F}, \mu)$ is not the dual space of anything. The fact that is often used instead is that $M(X)$ is the dual of $C_0(X)$.

There is an advantage to identifying a Banach space E^* as the dual space of another Banach space E . This can be done for $E^* = M(X)$ and for $E^* = L^q(X, \mathcal{F}, \mu)$ for $1 < q \leq \infty$ (with σ -finiteness in the case $q = \infty$). Then E^* is the space of all continuous linear functionals on the original space E . There is a corresponding notion of pointwise convergence in E^* , called weak * convergence, and this turns out to have useful properties that make it a convenient technical tool.

1.7 Orlicz spaces

It is helpful to place the theory of L^p spaces in a general context. Clearly, the theory depends heavily on the use of the functions x^p for $p \geq 1$. This is a convex function. The generalization is to use a more or less arbitrary convex function.

Let $H(x)$ be a continuous function defined for all $x \geq 0$ such that $H(0) = 0$ and such that $H'(x) > 0$ for $x > 0$. Then H is an increasing function. Suppose that $H(x)$ increases to infinity as x increases to infinity. Finally, suppose that $H''(x) \geq 0$. This implies convexity.

Example: $H(x) = x^p$ for $p > 1$.

Example: $H(x) = e^x - 1$.

Example: $H(x) = (x + 1) \log(x + 1)$.

Define the size of f by $\mu(H(|f|))$. This is a natural notion, but it does not have good scaling properties. So we replace f by f/c and see if we can make

the size of this equal to one. The c that accomplishes this will be the norm of f .

This leads to the official definition of the *Orlicz norm*

$$\|f\|_H = \inf\{c > 0 \mid \mu(H(|f/c|)) \leq 1\}. \quad (1.18)$$

It is not difficult to show that if this norm is finite, then we can find a c such that

$$\mu(H(|f/c|)) = 1. \quad (1.19)$$

Then the definition takes the simple form

$$\|f\|_H = c, \quad (1.20)$$

where c is defined by the previous equation.

It is not too difficult to show that this norm defines a Banach space $L_H(\mu)$. The key point is that the convexity of H makes the norm satisfy the triangle inequality.

Theorem 1.13 *The Orlicz norm satisfies the triangle inequality*

$$\|f + g\|_H \leq \|f\|_H + \|g\|_H. \quad (1.21)$$

Proof: Let $c = \|f\|_H$ and $d = \|g\|_H$. Then by the fact that H is increasing and convex

$$H\left(\left|\frac{f+g}{c+d}\right|\right) \leq H\left(\frac{c}{c+d}\left|\frac{f}{c}\right| + \frac{d}{c+d}\left|\frac{g}{d}\right|\right) \leq \frac{c}{c+d}H\left(\left|\frac{f}{c}\right|\right) + \frac{d}{c+d}H\left(\left|\frac{g}{d}\right|\right). \quad (1.22)$$

Integrate. This gives

$$\mu\left(H\left(\left|\frac{f+g}{c+d}\right|\right)\right) \leq 1. \quad (1.23)$$

Thus $\|f + g\|_H \leq c + d$. \square

Notice that this result is a generalization of Minkowski's inequality. So we see that the idea behind L^p spaces is convexity. The convexity is best for $1 < p < \infty$, since then the function x^p has second derivative $p(p-1)x^{p-2} > 0$. (For $p = 1$ the function x is still convex, but the second derivative is zero, so it not strictly convex.)

One can also try to create a duality theory for Orlicz spaces. For this it is convenient to make the additional assumptions that $H'(0) = 0$ and $H''(x) > 0$ and $H'(x)$ increases to infinity.

The dual function to $H(x)$ is a function $K(y)$ called the *Legendre transform*. The definition of $K(y)$ is

$$K(y) = xy - H(x), \quad (1.24)$$

where x is defined implicitly in terms of y by $y = H'(x)$.

This definition is somewhat mysterious until one computes that $K'(y) = x$. Then the secret is revealed: The functions H' and K' are inverse to each other. Furthermore, the Legendre transform of $K(y)$ is $H(x)$.

Examples:

1. Let $H(x) = x^p/p$. Then $K(y) = y^q/q$, where $1/p + 1/q = 1$.
2. Let $H(x) = e^x - 1 - x$. Then $K(y) = (y + 1) \log(y + 1) - y$.

Lemma 1.14 *Let $H(x)$ have Legendre transform $K(y)$. Then for all $x \geq 0$ and $y \geq 0$*

$$xy \leq H(x) + K(y). \quad (1.25)$$

Proof: Fix y and consider the function $xy - H(x)$. Since $H'(x)$ is increasing to infinity, the function rises and then dips below zero. It has its maximum where the derivative is equal to zero, that is, where $y - H'(x) = 0$. However by the definition of Legendre transform, the value of $xy - H(x)$ at this point is $K(y)$. \square

Theorem 1.15 (Hölder's inequality) *Suppose that H and K are Legendre transforms of each other. Then*

$$|\mu(fg)| \leq 2\|f\|_H\|g\|_K. \quad (1.26)$$

Proof: It is sufficient to prove this when $\|f\|_H = 1$ and $\|g\|_K = 1$. However by the lemma

$$|f(x)||g(x)| \leq H(|f(x)|) + K(|g(x)|). \quad (1.27)$$

Integrate. \square

This is just the usual derivation of Hölder's inequality. However if we take $H(x) = x^p/p$, $K(y) = y^q/q$, then the H and K norms are not quite the usual L^p and L^q , but instead multiples of them. This explains the extra factor of 2. In any case we see that the natural context for Hölder's inequality is the Legendre transform for convex functions. For more on this general subject, see Appendix H (Young-Orlicz spaces) in R. M. Dudley, *Uniform Central Limit Theorems*, Cambridge University Press, 1999.

1.8 Problems

1. Consider Lebesgue measure λ defined for Borel functions \mathcal{B} defined on \mathbf{R} . Let $1 \leq q < r < \infty$. (a) Give an example of a function in \mathcal{L}^q that is not in \mathcal{L}^r . (b) Give an example of a function in \mathcal{L}^r that is not in \mathcal{L}^q .
2. Consider a finite measure μ (the measure of the entire set or the integral of the constant function 1 is finite). Show that if $1 \leq q \leq r \leq \infty$, then $\mathcal{L}^r \subset \mathcal{L}^q$.
3. Let ϕ be a smooth convex function, so that that for each a and t we have $\phi(t) \geq \phi(a) + \phi'(a)(t - a)$. Let μ be a probability measure. Let f be a real function in \mathcal{L}^1 . Show that $\phi(\mu(f)) \leq \mu(\phi(f))$. (This is Jensen's inequality.) Hint: Let $a = \mu(f)$ and $t = f$. Where do you use the fact that μ is a probability measure?

4. Let ϕ be a smooth convex function as above. Deduce from the preceding problem the simple fact that if $0 \leq a$ and $0 \leq b$ with $a + b = 1$, then $\phi(au + bv) \leq a\phi(u) + b\phi(v)$. Describe explicitly the probability measure μ and the random variable f that you use.
5. Suppose that f is in \mathcal{L}^r for some r with $1 \leq r < \infty$. (a) Show that the limit as $p \rightarrow \infty$ of $\|f\|_p$ is equal to $\|f\|_\infty$. Hint: Obtain an upper bound on $\|f\|_p$ in terms of $\|f\|_r$ and $\|f\|_\infty$. Obtain a lower bound on $\|f\|_p$ by using Chebyshev's inequality applied to the set $|f| > a$ for some a with $0 < a < \infty$. Show that this set must have finite measure. For which a does this set have strictly positive measure? (b) Show that the result is not true without the assumption that f belongs to some \mathcal{L}^r .
6. Consider $1 \leq p < \infty$. Let \mathcal{B} denote Borel measurable functions on the line. Consider Lebesgue measure λ and the corresponding space $L^p(\mathbf{R}, \mathcal{B}, \lambda)$. If f is in this L^p space, the translate f_a is defined by $f_a(x) = f(x - a)$. (a) Show that for each f in L^p the function $a \mapsto f_a$ is continuous from the real line to L^p . (b) Show that the corresponding result for L^∞ is false. Hint: Take it as known that the space of step functions is dense in the space L^p for $1 \leq p < \infty$.
7. Define the Fourier transform for f in $L^1(\mathbf{R}, \mathcal{B}, \lambda)$ by

$$\hat{f}(k) = \int_{-\infty}^{\infty} e^{-ikx} f(x) dx. \quad (1.28)$$

Show that if f is in L^1 , then the Fourier transform is in L^∞ and is continuous. Hint: Use the dominated convergence theorem.

8. Show that if f is in L^1 , then the Fourier transform of f vanishes at infinity. Hint: Take it as known that the space of step functions is dense in the space L^1 . Compute the Fourier transform of a step function.
9. (a) Let $1 \leq p < \infty$. Show that the L^p norm of the integral is bounded by the integral of the L^p norm. More specifically, let μ be a measure defined for functions on X , and let ν be a measure defined for functions on Y . Suppose that μ and ν are each σ -finite. Let f be a product measurable function on $X \times Y$. Then $\nu(f | 1)$ denotes the ν partial integral of f keeping the first variable fixed, and $\|f | 2\|_p$ is the \mathcal{L}^p norm with respect to μ keeping the second variable fixed. The assertion is that

$$\|\nu(f | 1)\|_p \leq \nu(\|f | 2\|_p). \quad (1.29)$$

That is,

$$\left(\int \left| \int f(x, y) d\nu(y) \right|^p d\mu(x) \right)^{\frac{1}{p}} \leq \int \left(\int |f(x, y)|^p d\mu(x) \right)^{\frac{1}{p}} d\nu(y). \quad (1.30)$$

(b) What is the special case of this result when ν is a counting measure on two points? (c) What is the special case of this result when μ is a counting measure on two points? Hint: For the general inequality it is enough to give the proof when f is a positive function. Write $\alpha(y) = (\int f(x, y)^p d\mu(x))^{\frac{1}{p}}$ and set $\alpha = \int \alpha(y) d\nu(y)$. Then

$$\left(\frac{1}{\alpha} \int f(x, y) d\nu(y)\right)^p = \left(\int \frac{f(x, y)}{\alpha(y)} \frac{\alpha(y)}{\alpha} d\nu(y)\right)^p \leq \int \left(\frac{f(x, y)}{\alpha(y)}\right)^p \frac{\alpha(y)}{\alpha} d\nu(y). \quad (1.31)$$

Apply the μ integral and interchange the order of integration.

10. Let K be the Legendre transform of H . Thus $K(y) = xy - H(x)$, where x is the solution of $y = H'(x)$. (a) Show that $K'(y) = x$, in other words, K' is the inverse function to H' . (b) Show that if $H''(x) > 0$, then also $K''(y) > 0$. What is the relation between these two functions?

Chapter 2

Hilbert space

2.1 Inner products

A *Hilbert space* H is a vector space with an inner product that is complete. The vector space can have real scalars, in which case the Hilbert space is a real Hilbert space. Or the vector space can have complex scalars; this is the case of a complex Hilbert space. Both cases are useful. Real Hilbert spaces have a geometry that is easy to visualize, and they arise in applications. However complex Hilbert spaces are better in some contexts. In the following most of the attention will be given to complex Hilbert spaces.

An *inner product* is defined so that it is linear in one variable and conjugate linear in the other variable. The convention adopted here is that for vectors u, v, w and complex scalars a, b we have

$$\begin{aligned}\langle u, av + bw \rangle &= a\langle u, v \rangle + b\langle u, w \rangle \\ \langle au + bw, v \rangle &= \bar{a}\langle u, v \rangle + \bar{b}\langle w, v \rangle.\end{aligned}\tag{2.1}$$

Thus the inner product is conjugate linear in the left variable and linear in the right variable. This is the convention in physics, and it is also the convention in some treatments of elementary matrix algebra. However in advanced mathematics the opposite convention is common.

The inner product also satisfies the condition

$$\langle u, v \rangle = \overline{\langle v, u \rangle}.\tag{2.2}$$

Thus $\langle u, u \rangle$ is real. In fact, we require for an inner product that

$$\langle u, u \rangle \geq 0.\tag{2.3}$$

The final requirement is that

$$\langle u, u \rangle = 0 \Rightarrow u = 0.\tag{2.4}$$

The inner product defines a norm $\|u\| = \sqrt{\langle u, u \rangle}$. It has the basic homogeneity property that $\|au\| = |a|\|u\|$. The most fundamental norm identity is

$$\|u + v\|^2 = \|u\|^2 + \langle u, v \rangle + \langle v, u \rangle + \|v\|^2. \quad (2.5)$$

Notice that the cross terms are real, and in fact $\langle u, v \rangle + \langle v, u \rangle = 2\Re\langle u, v \rangle$. This leads to the following basic result.

Theorem 2.1 (Schwarz inequality)

$$|\langle u, v \rangle| \leq \|u\|\|v\|. \quad (2.6)$$

Proof: If either u or v is the zero vector, then the inequality is trivial. Otherwise let $u_1 = u/\|u\|$ and $v_1 = e^{i\theta}v/\|v\|$. Then these are unit vectors. By the fundamental identity

$$-2\Re(e^{i\theta}\langle u_1, v_1 \rangle) \leq 2. \quad (2.7)$$

Pick θ so that $-\Re(e^{i\theta}\langle u_1, v_1 \rangle) = |\langle u_1, v_1 \rangle|$. Then the equation gives

$$2|\langle u_1, v_1 \rangle| \leq 2. \quad (2.8)$$

This leads immediately to the Schwarz inequality. \square

Theorem 2.2 (Triangle inequality)

$$\|u + v\| \leq \|u\| + \|v\|. \quad (2.9)$$

Proof: From the fundamental identity and the Schwarz inequality

$$\|u + v\|^2 \leq \|u\|^2 + 2\|u\|\|v\| + \|v\|^2 = (\|u\| + \|v\|)^2. \quad (2.10)$$

\square

Each element u of H defines an element u^* of the dual space H^* of continuous linear functions from H to \mathbf{C} . The definition of u is that it is the linear function $v \mapsto \langle u, v \rangle$. That is, the value

$$u^*(v) = \langle u, v \rangle. \quad (2.11)$$

From the Schwarz inequality we have

$$\|u^*(v)\| \leq \|u\|\|v\|, \quad (2.12)$$

which proves that u^* is Lipschitz, hence continuous.

This notation shows the advantage of the convention that inner products are continuous in the right variable. In fact, this notation is rather close to one that is common in linear algebra. If u is a column vector, then u^* is given by the corresponding row vector with complex conjugate entries. The algebraic properties of this correspondence are just what one would want: $(u + w)^* = u^* + w^*$ and $(au)^* = \bar{a}u^*$.

The notation is also compatible with the standard notation for the adjoint of a linear transformation from one Hilbert space to another. We can identify the vector u with the linear transformation $a \mapsto au$ from \mathbf{C} to H . The adjoint of this transformation is the transformation $w \mapsto \langle u, w \rangle$ from H to \mathbf{C} . This is just the identity $\overline{\langle u, w \rangle} a = \langle w, au \rangle$. So it is reasonable to denote the transformation $w \mapsto \langle u, w \rangle$ by the usual notation u^* for adjoint.

One particularly convenient aspect of this notation is that one may also form the outer product vu^* . This is a linear transformation from H to itself given by $w \mapsto v\langle u, w \rangle$. This following result is well known in the context of integral equations.

Proposition 2.3 *The linear transformation vu^* has eigenvalues $\langle u, v \rangle$ and 0. If μ is not equal to either of these two values, then the inverse of $\mu I - vu^*$ is*

$$(\mu I - vu^*)^{-1} = \frac{1}{\mu} \left[I + \frac{1}{\mu - \langle u, v \rangle} vu^* \right]. \quad (2.13)$$

If $\langle u, v \rangle = 1$, then vu^* is a slant projection onto v along the directions perpendicular to u . In particular, if $u^*u = 1$, then uu^* is an orthogonal projection onto u .

In the physics literature a vector v is called a ket. A dual vector u^* is called a bra. The complex number that results from the pairing of a bra and a ket is $\langle u, v \rangle$, hence a bracket. This rather silly terminology is due to Dirac.

If u and v are vectors, we write $u \perp v$ if $\langle u, v \rangle = 0$. They are said to be perpendicular or orthogonal.

Theorem 2.4 (Pythagorus) *If $u \perp v$, then*

$$\|u + v\|^2 = \|u\|^2 + \|v\|^2. \quad (2.14)$$

The theorem of Pythagorus says that for a right triangle the sums of the squares of the sides is the square of the hypotenuse. It follows immediately from the fundamental identity. There is a better result that does not need the hypothesis of orthogonality.

Theorem 2.5 (parallelogram law) *In an inner product space*

$$\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2. \quad (2.15)$$

It says that for an arbitrary triangle there is an associated parallelogram consisting of $0, u, v, u + v$, and the sum of the squares of the diagonals is equal to the sum of the square of the four sides.

This parallelogram law is a fundamental convexity result. Write it in the form

$$\left\| \frac{u - v}{2} \right\|^2 + \left\| \frac{u + v}{2} \right\|^2 = \frac{1}{2}\|u\|^2 + \frac{1}{2}\|v\|^2. \quad (2.16)$$

This immediately shows that the function $u \rightarrow \|u\|^2$ has a convexity property, at least with weights $1/2$ and $1/2$. Furthermore, there is a kind of strict convexity, in that there is an extra term on the left that is strictly positive whenever $u \neq v$.

2.2 Closed subspaces

The classic examples of Hilbert spaces are sequence spaces \mathbf{C}^n (finite dimensional) and ℓ^2 (infinite dimensional). A more general class of infinite dimensional examples is given by $L^2(X, \mathcal{F}, \mu)$, the quotient space formed from \mathcal{F} measurable functions that are square integrable with respect to μ . In this case the inner product comes from

$$\langle f, g \rangle = \mu(\bar{f}g) = \int \overline{f(x)}g(x) d\mu(x). \quad (2.17)$$

The norm is

$$\|f\| = \sqrt{\mu(|f|^2)} = \sqrt{\int |f(x)|^2 d\mu(x)}. \quad (2.18)$$

A subspace (or vector subspace or linear subspace) M of a Hilbert space H is a subset with the zero vector in it and that is closed under vector addition and scalar multiplication.

Theorem 2.6 *A subspace M of a Hilbert space is a closed subspace if and only if it is itself a Hilbert space (with the same inner product).*

Proof: Suppose that M is a closed subspace. Since it is a closed subset of a complete metric space, it is complete. Therefore M is a Hilbert space.

Suppose on the other hand that M is a Hilbert space. Since it is a complete subset of a metric space, it follows that M is a closed subset. So M is a closed subspace. \square

Examples:

1. Consider the space ℓ^2 of square summable sequences. Fix n . The subspace consisting of all sequences x in this space with $x_k = 0$ for $k \leq n$ is a closed infinite dimensional subspace.
2. Consider the space ℓ^2 . The subspace consisting of all sequences x such that there exists an n such that $x_k = 0$ for $k > n$ is an infinite dimensional subspace that is not closed. In fact, its closure is all of ℓ^2 .
3. Consider the space ℓ^2 . Fix n . The subspace consisting of all sequences x in this space with $x_k = 0$ for $k > n$ is a closed finite dimensional subspace.
4. Consider the space $L^2(\mathbf{R}, \mathcal{B}, \lambda)$ of Borel functions on the line that are square summable with respect to Lebesgue measure λ . The subspace consisting of all step functions is not closed. In fact, its closure is the entire Hilbert space.
5. Consider the space $L^2(\mathbf{R}, \mathcal{B}, \lambda)$ of Borel functions on the line that are square summable with respect to Lebesgue measure λ . The subspace consisting of all continuous functions with compact support is not closed. In fact, its closure is the entire Hilbert space. Notice that the subspaces in the last two examples have only the zero vector in common.

If M is a subspace of a Hilbert space H , then we write $w \perp M$ if for all v in M we have $w \perp v$. Then M^\perp consists of all vectors w in H such that $w \perp M$.

Theorem 2.7 *If M is a subspace, then M^\perp is a closed subspace. Furthermore, $M \subset M^{\perp\perp}$.*

Notice that in general, it is not the case that M is equal to $M^{\perp\perp}$. As an example, let M be the subspace that is the intersection of the subspaces in examples 1 and 2 above. Then M^\perp is the subspace in example 3. However $M^{\perp\perp}$ is the subspace in example 1.

2.3 The projection theorem

Lemma 2.8 *Let M be a subspace of H . Let u be a vector in H . Then v is a vector in M that is closest to u if and only if v is in M and $u - v$ is perpendicular to M .*

Proof: Suppose that v is in M and $u - v$ is perpendicular to M . Let v' be another vector in M . Then $u - v' = u - v + v - v'$. By the theorem of Pythagoras $\|u - v'\|^2 = \|u - v\|^2 + \|v - v'\|^2$. Hence $\|u - v\| \leq \|u - v'\|$. So v is the vector in M closest to u .

Suppose on the other hand that v is the vector in M closest to u . Let $w \neq 0$ be another vector in M . Let cw be the projection of $u - v$ onto the subspace generated by w given by taking $c = \langle u - v, w \rangle / \langle w, w \rangle$. Let $v' = v + cw$. Then the difference $(u - v') - (v' - v) = u - v'$ is orthogonal to $v' - v$. By the theorem of Pythagoras $\|u - v\|^2 = \|u - v'\|^2 + \|v' - v\|^2$. Since v is closest to u , the left hand side must be no larger than the first term on the right hand side. Hence the second term on the right hand side is zero. Thus $v' - v = cw$ is zero, and hence $c = 0$. We conclude that $u - v$ is perpendicular to w . Since w is an arbitrary non-zero vector in M , this proves that $u - v$ is perpendicular to M . \square

Theorem 2.9 (Projection theorem) *Let M be a closed subspace of the Hilbert space H . Let u be a vector in H . Then there exists a unique vector v in M that is closest to u . In particular, $u - v$ is perpendicular to M .*

Proof: Let a be the infimum of the numbers $\|u - v\|^2$ for v in M . Let v_n be a sequence of vectors in M such that $\|u - v_n\|^2 \rightarrow a$ as $n \rightarrow \infty$. Apply the parallelogram identity to two vectors $u - v_m$ and $u - v_n$. This gives

$$\|\frac{1}{2}(v_m - v_n)\|^2 + \|u - \frac{1}{2}(v_m + v_n)\|^2 = \frac{1}{2}\|u - v_m\|^2 + \frac{1}{2}\|u - v_n\|^2. \quad (2.19)$$

Since each $\frac{1}{2}(v_m + v_n)$ is in M , it follows that

$$\|\frac{1}{2}(v_m - v_n)\|^2 + a \leq \frac{1}{2}\|u - v_m\|^2 + \frac{1}{2}\|u - v_n\|^2. \quad (2.20)$$

As m, n get large, the right hand side tends to a . This proves that the v_n form a Cauchy sequence. Since M is complete, there exists a v in M such that $v_n \rightarrow v$ as $n \rightarrow \infty$. Since the $\|u - v_n\|^2 \rightarrow a$ and the norm is a continuous function on H , it follows that $\|u - v\|^2 = a$. This proves that v is the vector in M closest to u . \square

The vector v is the orthogonal projection (or projection) of u on M . It is clear that the vector $u - v$ is the orthogonal projection of u on M^\perp . Thus every vector in H may be written as the sum of a vector in M and a vector in M^\perp .

Corollary 2.10 *If M is a closed subspace, then $M^{\perp\perp} = M$.*

Proof: It is evident that $M \subset M^{\perp\perp}$. The hard part is to show that $M^{\perp\perp} \subset M$. Let u be in $M^{\perp\perp}$. By the projection theorem we can write $u = w + v$, where v is in M and w is in M^\perp . Then $0 = \langle u, w \rangle = \langle w, w \rangle + \langle v, w \rangle = \langle w, w \rangle$. So $w = 0$ and $u = v$ is in M . \square

2.4 The Riesz-Fréchet theorem

Theorem 2.11 (Riesz-Fréchet representation theorem) *Let H be a Hilbert space. Let L be a continuous linear function $L : H \rightarrow \mathbf{C}$. Then there exists a unique vector u in H such that $L = u^*$. That is, the vector u satisfies $L(v) = \langle u, v \rangle$ for all v in H .*

Proof: Let

$$E[u] = \frac{1}{2}\|u\|^2 - \Re L(u). \quad (2.21)$$

Since L is continuous, it is Lipschitz. That is, there is a constant c such that $|Lw| \leq c\|w\|$. It follows that $E[w] \geq \frac{1}{2}\|w\|^2 - c\|w\| \geq -\frac{1}{2}c^2$. Thus the $E[w]$ are bounded below. Let a be the infimum of the $E[w]$. Let w_n be a sequence such that $E[w_n] \rightarrow a$ as $n \rightarrow \infty$. By the parallelogram identity and the fact that L is linear we have

$$\|\frac{1}{2}(w_m - w_n)\|^2 + 2E[\frac{1}{2}(w_m + w_n)] = E[w_m] + E[w_n]. \quad (2.22)$$

Hence

$$\|\frac{1}{2}(w_m - w_n)\|^2 + 2a \leq E[w_m] + E[w_n]. \quad (2.23)$$

However, the right hand side tends to $2a$. This proves that the w_n form a Cauchy sequence. Since H is complete, they converge to a vector u in H . By continuity $E[u] = a$.

The rest of the proof amounts to taking the derivative of the function $E[w]$ at $w = u$ in the direction v . The condition that this derivative is zero should give the result. However it is perhaps worth writing it out explicitly. Let v be a vector in H . Then for each real number $t > 0$ we have

$$a = E[u] \leq E[u + tv]. \quad (2.24)$$

This says that

$$0 \leq t\Re\langle u, v \rangle - t\Re L(v) + \frac{1}{2}t^2\|v\|^2. \quad (2.25)$$

Divide both sides by $t > 0$. Since the resulting inequality is true for each $t > 0$, we have $0 \leq \Re\langle u, v \rangle - \Re L(v)$. Thus we have proved that $\Re L(v) \leq \Re\langle u, v \rangle$ for all v . The same argument applied to $-v$ shows that $-\Re L(v) \leq -\Re\langle u, v \rangle$. Hence $\Re L(v) = \Re\langle u, v \rangle$. The above reasoning applied to $-iv$ shows that $\Im L(v) = \Im\langle u, v \rangle$. We conclude that $L(v) = \langle u, v \rangle$. \square

The Riesz-Fréchet representation theorem says that every element of the dual space H^* comes from a vector u in H . However it does not quite say that a Hilbert space is naturally isomorphic to its dual space, at least not in the case of complex scalars. In fact, the correspondence $u \mapsto u^*$ from H to H^* is conjugate linear. However it does preserve the norm. (In the real case it is completely accurate to say that a Hilbert space is naturally self-dual.)

2.5 Bases

In the following we need the notion of isomorphism of Hilbert spaces. If H and K are each Hilbert spaces, then an *isomorphism* (or *unitary transformation*) from H to K is a bijection $U : H \rightarrow K$ that is linear and preserves the inner product. The latter means that $\langle Uv, Uw \rangle = \langle v, w \rangle$, where the inner product on the left is that of K and the inner product on the right is that of H . An isomorphism of Hilbert spaces obviously also preserves the norm. As a consequence, an isomorphism of Hilbert spaces is simultaneously an isomorphism of vector spaces and an isomorphism of metric spaces. Clearly the inverse of an isomorphism is also an isomorphism.

Let H be a Hilbert space. Let J be an index set. An orthonormal family is a function $j \mapsto u_j$ from J to H such that $\langle u_j, u_k \rangle = \delta_{jk}$.

Proposition 2.12 (Bessel's inequality) *Let f be in H and define coefficients c by*

$$c_j = \langle u_j, f \rangle. \quad (2.26)$$

Then c is in $\ell^2(J)$ and

$$\|c\|^2 \leq \|f\|^2. \quad (2.27)$$

Proof: Let J_0 be a finite subset of J . Let $g = \sum_{j \in J_0} c_j u_j$. Then $f - g \perp u_k$ for each k in J_0 . By the theorem of Pythagorus $\|f\|^2 = \|f - g\|^2 + \|g\|^2$. Hence $\|g\|^2 \leq \|f\|^2$. This says that $\sum_{j \in J_0} |c_j|^2 \leq \|f\|^2$. Since this is true for arbitrary finite $J_0 \subset J$, we have $\sum_{j \in J} |c_j|^2 \leq \|f\|^2$. \square

Proposition 2.13 (Riesz-Fischer) *Let $j \mapsto u_j$ be an orthonormal family. Let c be in $\ell^2(J)$. Then the series*

$$g = \sum_j c_j u_j \quad (2.28)$$

converges in the Hilbert space sense to a vector g in H . Furthermore,

$$\|g\|^2 = \|c\|^2. \quad (2.29)$$

Proof: Since c is in $\ell^2(J)$, there are only countably many values of j such that $c_j \neq 0$. So we may consider this countable sum indexed by natural numbers. Let g_n be the n th partial sum. Then for $m > n$ we have $\|g_m - g_n\|^2 = \sum_{j=n+1}^m |c_j|^2$. This approaches zero as $m, n \rightarrow \infty$. So the g_n form a Cauchy sequence. Since H is complete, the g_n converge to some g in H . That is, $\|g - \sum_{j=0}^n c_j e_j\| \rightarrow 0$ as $n \rightarrow \infty$. This is the Hilbert space convergence indicated in the statement of the proposition. \square

Theorem 2.14 *Let $j \mapsto u_j$ be an orthonormal family of vectors in the Hilbert space H . Then there exists a closed subspace M of H such that the map $c \rightarrow \sum_j c_j u_j$ is an isomorphism from ℓ^2 to M . The map that sends f in H to $g = \sum_j \langle u_j, f \rangle u_j$ in H is the orthogonal projection of f onto M .*

Sometimes physicists like to write the orthogonal projection onto M with a notation that somewhat resembles

$$E = \sum_j u_j u_j^*. \quad (2.30)$$

This just means that the orthogonal projection of a vector f onto the span of the e_j is $Ef = \sum_j u_j \langle u_j, f \rangle$.

An orthonormal family $j \mapsto u_j$ is a *basis* for the Hilbert space H if every vector f in H has the representation

$$f = \sum_j \langle u_j, f \rangle u_j. \quad (2.31)$$

In this case the correspondence between H and $\ell^2(J)$ given by the basis is an isomorphism of Hilbert spaces. That is, the coefficient vectors c with $c_j = \langle u_j, f \rangle$ give an alternative description of the Hilbert space.

Proposition 2.15 *Suppose that $j \mapsto e_j$ is a maximal orthonormal family in H . Then it is a basis for H .*

Proof: Let M be the collection of all linear combinations $\sum_j c_j u_j$ where c is in ℓ^2 . Then M is a closed subspace. If the family does not form a basis, then M is a proper subset of H . Let f be in H and not in M . Let g be the projection f on M . Then $f - g$ is orthogonal to M and is non-zero, and so can be normalized to be a unit vector. This gives a strictly larger orthonormal family, so the original family is not maximal. \square

If H is an arbitrary Hilbert space, then it follows from the axiom of choice via Zorn's lemma that there is a maximal orthonormal family $j \mapsto u_j$ defined on some index set J . This is a basis for H . In other words, every Hilbert space has a basis. It follows that for every Hilbert space there is a set J such that H is

isomorphic to $\ell^2(J)$. This method of constructing bases involves lots of arbitrary choices and is not particularly practical. However it is of considerable theoretical interest: it says that Hilbert spaces of the same dimension (cardinality of index set) are isomorphic.

Proposition 2.16 . *Suppose that $j \mapsto u_j$ is an orthonormal family such that for every vector f in H we have*

$$\|f\|^2 = \sum_j |\langle u_j, f \rangle|^2. \quad (2.32)$$

Then $u \mapsto u_j$ is a basis for H .

Proof: Let M be the collection of all linear combinations $\sum_j c_j u_j$ where c is in ℓ^2 . Then M is a closed subspace. If the family does not form a basis, then M is a proper subset of H . Let f be in H and not in M . Let g be the projection of f on M . By the theorem of Pythagorus, $\|f\|^2 = \|g\|^2 + \|f - g\|^2 > \|g\|^2$. This violates the equality. \square

Proposition 2.17 *Suppose that $j \mapsto e_j$ is an orthonormal family, and suppose that the set of all finite linear combinations of these vectors is dense in H . Then it is a basis.*

Proof: Let M be the collection of all linear combinations $\sum_j c_j u_j$ where c is in ℓ^2 . Then M is a closed subspace. Let f be in H . Consider $\epsilon > 0$. Then there exists a finite linear combination h such that $\|f - h\| < \epsilon$. Furthermore, h is in M . Let g be the projection of f onto M . Since g is the element of M that is closest to f , it follows that $\|f - g\| \leq \|f - h\|$. Hence $\|f - g\| < \epsilon$. Since $\epsilon > 0$ is arbitrary, we have $f = g$. \square

2.6 Separable Hilbert spaces

Theorem 2.18 (Gram-Schmidt orthonormalization) *Let $k \mapsto v_k$ be a linearly independent sequence of vectors in the Hilbert space H . Then there exists an orthonormal sequence $k \mapsto u_k$ such that for each m the span of u_1, \dots, u_m is equal to the span of v_1, \dots, v_m .*

Proof: The proof is by induction on m . For $m = 0$ the sequences are empty, and the corresponding spans are both just the subspace with only the zero vector. Given u_1, \dots, u_m orthonormal with the same span as v_1, \dots, v_m , let q_{m+1} be the orthogonal projection of v_{m+1} on this subspace. Let $p_{m+1} = v_{m+1} - q_{m+1}$. Then p_{m+1} is orthogonal to u_1, \dots, u_m . This vector is non-zero, since otherwise v_{m+1} would depend linearly on the vectors in the subspace. Thus it is possible to define the unit vector $u_{m+1} = p_{m+1}/\|p_{m+1}\|$. \square

Theorem 2.19 *Let H be a separable Hilbert space. Then H has a countable basis.*

Proof: Since H is separable, there is a sequence $s : \mathbf{N} \rightarrow H$ such that $n \mapsto s_n$ has dense range. In particular, the set of all finite linear combinations of the vectors s_n is dense in H . Define a new subsequence v_k (finite or infinite) by going through the s_n in order and throwing out each element that is a linear combination of the preceding elements. Then the v_k are linearly independent, and the linear span of the v_k is the same as the linear span of the s_n . Thus we have a linearly independent sequence t_k whose linear span is dense in H . Apply the Gram-Schmidt orthonormalization procedure. This gives an orthonormal sequence u_k with the same linear span. Since the linear span of the u_k is dense in H , the u_k form a basis for H . \square

Corollary 2.20 *Every separable Hilbert space is isomorphic to some space $\ell^2(J)$, where J is a countable set.*

For most applications separable Hilbert spaces are sufficient. In fact, a separable Hilbert space is either finite dimensional or has countable infinite dimension. All countable infinite dimensional separable Hilbert spaces are isomorphic. In fact, they are all isomorphic to $\ell^2(\mathbf{N})$.

The space $L^2([0, 1], \mathcal{B}, \lambda)$ of Borel functions on the unit interval with

$$\lambda(|f|^2) = \int_0^1 |f(x)|^2 dx < +\infty \quad (2.33)$$

is a separable infinite dimensional Hilbert space. An example of an orthonormal basis is given by the Walsh functions.

Consider a natural number $n \geq 1$. Divide the interval from 0 to 1 in 2^n equally spaced parts, numbered from 0 to $2^n - 1$. The Rademacher function r_n is the function that is 1 on the even numbered intervals and -1 on the odd numbered intervals. A Walsh function is a product of Rademacher functions. Let $S \subset \{1, 2, 3, \dots\}$ be a finite set of strictly positive natural numbers. Let the Walsh function be defined by

$$w_S = \prod_{j \in S} r_j. \quad (2.34)$$

Notice that when S is empty the product is 1.

The Walsh functions may be generated from the Rademacher functions in a systematic way. At stage zero start with the function 1. At stage one take also r_1 . At stage two take r_2 times each of the functions from the previous stages. This gives also r_2 and $r_1 r_2$. At stage three take r_3 times each of the functions from the previous stages. This gives also r_3 and $r_1 r_3$ and $r_2 r_3$ and $r_1 r_2 r_3$. It is clear how to continue. The Walsh functions generated in this way oscillate more and more.

Theorem 2.21 *The Walsh functions form an orthonormal basis of $L^2([0, 1], \mathcal{B}, \lambda)$ with respect to the inner product*

$$\langle f, g \rangle = \lambda(\bar{f}g) = \int_0^1 \overline{f(x)}g(x) dx. \quad (2.35)$$

Thus for an arbitrary function f in $L^2([0,1])$ there is an L^2 convergent Walsh expansion

$$f(x) = \sum_S \langle w_S, f \rangle w_S(x). \quad (2.36)$$

Proof: The 2^n Walsh functions w_S with $S \subset \{1, \dots, n\}$ are linearly independent. It follows that they span the 2^n dimensional space of binary step functions with step width $1/2^n$. Thus the linear span of the Walsh functions is the space of all binary step functions. These are dense in L^2 . In fact, they are even uniformly dense in the space of continuous functions. \square

Consider a natural number $n \geq 0$. Divide the interval from 0 to 1 in 2^n equally spaced parts, numbered from 0 to $2^n - 1$. The binary rectangular function $f_{n;k}$ is the function that is the indicator function of the k th interval.

A Haar function is a multiple of a product of a binary rectangular function with a Rademacher function. For $n \geq 0$ and $0 \leq k < 2^n$ define the Haar function to be

$$h_{n;k} = c_n f_{n;k} r_{n+1}, \quad (2.37)$$

and define $h_{-1;0} = 1$. For $n \geq 0$ the coefficient $c_n > 0$ is determined by $c_n^2 = 1/2^n$. The function $h_{-1;0}$ together with the other Haar functions $h_{j;k}$ for $j = 0$ to $n - 1$ and $0 \leq k < 2^j$ form a basis for the binary step functions with width $1/2^n$. Note that the number of such functions is $1 + \sum_{j=0}^{n-1} 2^j = 2^n$.

The Haar functions may be generated in a systematic way. At stage zero start with the function 1. At stage one take also r_1 . At stage two take also $f_{1;0}r_2$ and $f_{1;1}r_2$. At stage three take also $f_{2;0}r_3$ and $f_{2;1}r_3$ and $f_{2;2}r_3$ and $f_{2;3}r_3$. The Haar functions generated in this way become more and more concentrated in width.

Theorem 2.22 *The Haar functions form an orthonormal family of vectors with respect to the inner product*

$$\langle f, g \rangle = \lambda(\bar{f}g) = \int_0^1 \overline{f(x)}g(x) dx. \quad (2.38)$$

For an arbitrary function f in $L^2([0,1])$ there is an L^2 convergent Haar expansion

$$f(x) = \sum_{n=-1}^{\infty} \sum_{0 \leq k < 2^n} \langle h_{n;k}, f \rangle h_{n;k}(x). \quad (2.39)$$

Proof: The 2^n partial sum of the Haar series is the same as the 2^n partial sum of the Walsh series. Each of these is the projection onto the space of rectangular functions of width $1/2^n$. \square

2.7 Problems

1. Consider the Hilbert space $\mathcal{H} = L^2(\mathbf{R}, \mathcal{B}, \lambda)$, where $\lambda(f) = \int_{-\infty}^{\infty} f(x) dx$ is Lebesgue measure. Consider the closed subspace M consisting of all

functions g in \mathcal{H} satisfying $g(-x) = -g(x)$. Find the orthogonal projection of the function $f(x) = (1+x)e^{-x}/(1+e^{-2x})$ onto this subspace.

2. Let

$$E = x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} + z \frac{\partial}{\partial z} \quad (2.40)$$

be the Euler operator. Let

$$\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \quad (2.41)$$

be the Laplace operator. The space P_ℓ of solid spherical harmonics of degree ℓ consists of all polynomials p in variables x, y, z with $Ep = \ell p$ and $\Delta p = 0$. It may be shown that P_ℓ has dimension $2\ell + 1$. (a) Prove that P_ℓ is invariant under rotations about the origin. (b) Find the spaces P_ℓ for $\ell = 0, 1, 2, 3$.

3. Let S be the unit sphere in \mathbf{R}^3 defined by the equation $x^2 + y^2 + z^2 = 1$. The space H_ℓ of surface spherical harmonics consists of the restrictions of the solid spherical harmonics to S . Show that P_ℓ and H_ℓ are isomorphic. Give an explicit formula that expresses a solid spherical harmonic in terms of the corresponding surface spherical harmonic.
4. Let $x = r \sin \theta \cos \phi$, $y = r \sin \theta \sin \phi$ and $z = r \cos \theta$ be expressed in spherical polar coordinates. Then $E = r\partial/\partial r$. Furthermore,

$$\Delta = \frac{1}{r^2} (E(E+1) + \Delta_S), \quad (2.42)$$

where

$$\Delta_S = \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \sin \theta \frac{\partial}{\partial \theta} + \frac{1}{\sin^2 \theta} \frac{\partial^2}{\partial \phi^2}. \quad (2.43)$$

Consider the space $L^2(S)$ with the rotation invariant measure $\sin \theta d\theta d\phi$, where the co-latitude θ goes from 0 to π and the longitude ϕ goes from 0 to 2π . (a) Show that for smooth functions in $L^2(S)$ we have the identity $\langle \Delta_S u, v \rangle = \langle u, \Delta_S v \rangle$. When you integrate by parts, be explicit about boundary terms. (b) Show that for u in H_ℓ we have $\Delta_S u = -\ell(\ell+1)u$.

5. Show that the subspaces H_ℓ of surface spherical harmonics are orthogonal in $L^2(S)$ for different values of ℓ .
6. Consider the Hilbert space $L^2(\mathbf{R}, \mathcal{B}, \gamma)$, where γ is the measure $\gamma(h) = \int_{-\infty}^{\infty} h(x) \exp(-x^2) dx$. Thus the inner product is

$$\langle f, g \rangle = \gamma(\bar{f}g) = \int_{-\infty}^{\infty} \overline{f(x)} g(x) e^{-x^2} dx. \quad (2.44)$$

Let the polynomial $h_n(x)$ of degree n be defined by

$$h_n(x) = \left(2x - \frac{d}{dx} \right)^n 1. \quad (2.45)$$

Thus $h_0(x) = 1$, $h_1(x) = 2x$, $h_2(x) = 4x^2 - 2$, and so on. Find $h_3(x)$, $h_4(x)$, $h_5(x)$. Show that the h_n form an orthogonal family of vectors in the Hilbert space. Hint: Integrate by parts.

7. Show that the inner product in a Hilbert space is determined by its norm. Prove that a Banach space isomorphism between Hilbert spaces is actually a Hilbert space isomorphism. (That is, prove that if there is a linear bijective correspondence that preserves the norm, then it preserves the inner product.) Hint: For a real Hilbert space this follows from $4\langle u, v \rangle = \|u + v\|^2 - \|u - v\|^2$. How about for a complex Hilbert space?
8. Let $j \rightarrow u_j$ be an orthogonal family of non-zero vectors in H indexed by $j \in J$. Suppose that finite linear combinations of these vectors are dense in H . Let $w_j = 1/\langle u_j, u_j \rangle$ be numerical weights associated with these vectors. Let $\ell^2(J, w)$ be the Hilbert space of all sequences c with $\sum_j |c_j|^2 w_j < \infty$. Show that H is isomorphic to $\ell^2(J, w)$ by the isomorphism that sends f to c given by $c_j = \langle u_j, f \rangle$. Find the inverse isomorphism. Check that both isomorphisms preserve the norm.
9. This is a continuation. Let $T = (-\frac{L}{2}, \frac{L}{2}]$ be the circle of length L , and consider the Hilbert space $H = L^2(T, \mathcal{B}, \lambda)$ with norm squared equal to

$$\|f\|^2 = \int_T |f(x)|^2 dx. \quad (2.46)$$

Let k be an integer multiple of $\frac{2\pi}{L}$. Let u_k be the element of H defined by $u_k(x) = \exp(-ikx)$. According to the theory of Fourier series the finite linear combinations of these vectors are dense in H . These vectors are orthogonal but not orthonormal. Do not normalize them! Instead, find the weights $w_j = 1/\langle u_j, u_j \rangle$. (a) Write explicitly the formula for the coefficients c in the space $\ell^2(\frac{2\pi}{L}\mathbf{Z}, w)$ of a function f in H . (b) Write explicitly the formula for f in terms of the coefficients c . (c) Write explicitly the equation that expresses the equality of norms (squared) in the two Hilbert spaces.

10. This is a continuation. Fix f smooth with compact support. Let $L \rightarrow \infty$. What are the limiting formulas corresponding to (a),(b),(c) above. Why was it important not to normalize the vectors?

Chapter 3

Differentiation

3.1 The Lebesgue decomposition

Consider a measurable space X, \mathcal{F} . Here \mathcal{F} can stand for a σ -algebra of subsets, or for the corresponding σ -algebra of real measurable functions. Let ν be a measure. As usual we write the integral of a measurable function $f \geq 0$ as $\nu(f) \geq 0$. We also write the measure of a measurable subset E as $\nu(E) = \nu(1_E)$.

Consider two such measures ν and μ . Then ν is said to be *absolutely continuous* with respect to μ if every measurable subset E with $\mu(E) = 0$ also has $\nu(E) = 0$. In that case we write $\nu \prec \mu$.

Consider two such measures ν and μ . Then ν and μ are said to be *singular* if there exists a measurable set A with complement $A^c = X \setminus A$ such that $\mu(A) = 0$ and $\nu(A^c) = 0$. In that case we write $\nu \perp \mu$.

Theorem 3.1 (Lebesgue decomposition) *Let X, \mathcal{F} be a measurable space. Let μ and ν be finite measures. Then $\nu = \nu_{ac} + \nu_s$, where $\nu_{ac} \prec \mu$ and $\nu_s \perp \mu$.*

Proof: This Hilbert space proof is due to von Neumann. The trick is to compare ν not directly with μ but with $\mu + \nu$ instead. In fact, it is also possible to compare μ to $\mu + \nu$ in the same way.

The main technical device is to look at ν as a linear functional on the real Hilbert space $L^2(X, \mathcal{F}, \mu + \nu)$. We have

$$|\nu(f)| \leq \nu(|f|) \leq (\mu + \nu)(|f|) = \langle 1, |f| \rangle. \quad (3.1)$$

By this Schwarz inequality and the fact that $\mu + \nu$ is a finite measure we have the continuity

$$|\nu(f)| \leq \|1\|_2 \|f\|_2. \quad (3.2)$$

It follows from the Riesz-Fréchet theorem that there exists a function g in $L^2(X, \mathcal{F}, \mu + \nu)$ so that ν is given by the inner product by $\nu(f) = \langle g, f \rangle$. In other words

$$\nu(f) = (\mu + \nu)(gf). \quad (3.3)$$

A little algebra gives

$$\mu(f) = (\mu + \nu)((1 - g)f). \quad (3.4)$$

Let E be the set where $g < 0$. If $(\mu + \nu)(E) > 0$, then $\nu(E) = (\mu + \nu)(g1_E) < 0$, which is a contradiction. Thus $(\mu + \nu)(E) = 0$ and so $0 \leq g$ almost everywhere with respect to $\mu + \nu$. Let F be the set where $g > 1$. If $(\mu + \nu)(F) > 0$, then $\mu(F) = (\mu + \nu)((1 - g)1_F) < 0$, which is a contradiction. So also $(\mu + \nu)(F) = 0$. So we may as well assume that $0 \leq g \leq 1$.

One consequence of this result is that the last two displayed equations hold for all measurable functions $f \geq 0$, by the monotone convergence theorem.

Let A be the set where $g = 1$. Let $A^c = X \setminus A$ be its complement, so $0 \leq g < 1$ on A^c . Let $\nu_s(f) = \nu(f1_A)$ and let $\nu_{ac}(f) = \nu(f1_{A^c})$. Since $\mu(A) = 0$ and $\nu_s(A^c) = 0$, we have $\nu_s \perp \mu$. If $\mu(E) = 0$, let $E' = E \cap A^c$. Then $\mu(E') = 0$ and so $(\mu + \nu)((1 - g)1_{E'}) = 0$. Since $1 - g > 0$ on E' , we have $(\mu + \nu)(E') = 0$ and so $\nu(E') = 0$. Thus $\nu_{ac}(E) = 0$. This proves that $\nu_{ac} \prec \mu$. \square

The Lebesgue decomposition also holds true when ν and μ are σ -finite measures. It also holds true in the context when ν is a finite signed measure (a difference of two finite measures) and μ is a σ -finite measure.

3.2 The Radon-Nikodym theorem

Theorem 3.2 (Radon-Nikodym) *Let X, \mathcal{F} be a measurable space. Let μ and ν be finite measures. Suppose $\nu \prec \mu$. Then there exists a relative density $h \geq 0$ with $\mu(h) < \infty$ and with $\nu(f) = \mu(hf)$ for all $f \geq 0$.*

Proof: This is the von Neumann method of proof again. Suppose that $\nu \prec \mu$. Then from the Hilbert space argument above we find g with $0 \leq g \leq 1$ such that $\nu(u) = (\mu + \nu)(gu)$ and $\mu(u) = (\mu + \nu)((1 - g)u)$. These identities are valid for all measurable functions $u \geq 0$. Let A be the set where $g = 1$. Then $\mu(A) = (\mu + \nu)((1 - g)1_A) = 0$. Hence by absolute continuity $\nu(A) = 0$. Let $h = g/(1 - g)$ on A^c and $h = 0$ on A . Thus

$$(1 - g)h = g1_{A^c}. \quad (3.5)$$

This equation is crucial: it shows that h is related to g in a quite non-linear way. All that we know about h is that $0 \leq h < +\infty$.

Let $f \geq 0$. Then we use the first identity with $u = f1_{A^c}$ and the second identity with $u = hf$ to get

$$\nu(f) = \nu(f1_{A^c}) = (\mu + \nu)(g1_{A^c}f) = (\mu + \nu)((1 - g)hf) = \mu(hf). \quad (3.6)$$

It follows that $\mu(h) = \nu(1) < +\infty$, so h is integrable. \square

The Radon-Nikodym theorem is also true when ν is a finite measure and μ is a σ -finite measure. It also holds true in the more general context when ν is a finite signed measure and μ is a σ -finite measure.

The function h is called the Radon-Nikodym derivative of ν with respect to the reference measure μ . Some justification for this terminology may be found in the problems.

3.3 Absolutely continuous increasing functions

Proposition 3.3 *Let ν be a finite measure and let μ be another measure. Then $\nu \prec \mu$ is equivalent to the condition that for every $\epsilon > 0$ there is a $\delta > 0$ such that for every measurable subset E we have $\mu(E) < \delta \Rightarrow \nu(E) < \epsilon$.*

Proof: Suppose that $\nu \prec \mu$. Suppose there exists $\epsilon > 0$ such that for every $\delta > 0$ there is a measurable subset E such that $\mu(E) < \delta$ and $\nu(E) \geq \epsilon$. Consider such an ϵ . For each n choose a measurable subset E_n such that $\mu(E_n) < 1/2^{n+1}$ and $\nu(E_n) \geq \epsilon$. Let $F_k = \bigcup_{n=k}^{\infty} E_n$. Then $\mu(F_k) \leq 1/2^k$. Let $F = \bigcap_{k=1}^{\infty} F_k$. Since for each k we have $F \subset F_k$, we have $\mu(F) \leq \mu(F_k) \leq 1/2^k$. Thus $\mu(F) = 0$. On the other hand, $E_k \subset F_k$, so $\nu(F_k) \geq \nu(E_k) \geq \epsilon$. Since $F_k \downarrow F$ and ν is a finite measure, it must be that $\nu(F_k) \downarrow \nu(F)$. Hence $\nu(F) \geq \epsilon$. The existence of F with $\mu(F) = 0$ and $\nu(F) \geq \epsilon$ implies that $\nu \prec \mu$ is false. This is a contradiction. Thus the $\epsilon - \delta$ condition holds. Thus the implication follows.

The converse is considerably easier. Suppose that the $\epsilon - \delta$ condition is satisfied. Suppose $\mu(E) = 0$. Let $\epsilon > 0$. It follows from the condition that $\nu(E) < \epsilon$. Since $\epsilon > 0$ is arbitrary, we have $\nu(E) = 0$. This is enough to show that $\nu \prec \mu$. \square

Consider the real line \mathbf{R} with the notion \mathcal{B} of Borel set and Borel function. Let F be an increasing right continuous real function on \mathbf{R} . Then there is a unique measure ν_F with the property that $\nu_F((a, b]) = F(b) - F(a)$ for all $a < b$. This measure is finite on compact sets. The measure determines the function up to an additive constant. Clearly ν_F is a finite measure precisely when F is bounded.

One other fact that we need is that the measure of a Borel measurable set is determined from the function F by a two-stage process. The first stage is to extend the measure from intervals $(a, b]$ to countable unions of such intervals. The second stage is to approximate an arbitrary measurable subset from outside by such countable unions.

It is not very difficult to show that the same result may be obtained by using open intervals (a, b) instead of half-open intervals $(a, b]$. The most general open subset U of the line is a countable union of open intervals. So the second stage of the approximation process gives the condition of outer regularity:

$$\nu_F(E) = \inf\{\nu_F(U) \mid U \text{ open, } E \subset U\}. \quad (3.7)$$

An increasing function F is said to be absolutely continuous increasing if for every $\epsilon > 0$ there is $\delta > 0$ such that whenever V is a finite union of disjoint open intervals with total length $\lambda(V) < \delta$ the corresponding sum of increments of F is $< \epsilon$.

A Lipschitz increasing function from the real line to itself is an absolutely continuous increasing function. The converse is false.

An absolutely continuous increasing function is a uniformly continuous function from the real line to itself. However again the converse is false: not every uniformly continuous increasing function from the line to itself is absolutely continuous. The Cantor function provides an example.

Theorem 3.4 *The measure $\nu_F \prec \lambda$ if and only if F is an absolutely continuous increasing function.*

Proof: Suppose that $\nu_F \prec \lambda$. Then the fact that F is absolutely continuous increasing follows from the proposition above.

Suppose on the other hand that F is an absolutely continuous increasing function. Consider $\epsilon > 0$. Choose $\epsilon' < \epsilon$ with $\epsilon' > 0$. Then there exists a $\delta > 0$ such that whenever U is a finite union of disjoint open sets, then the sum of the corresponding increases of F is $< \epsilon'$. Suppose that E is a Borel measurable subset with $\lambda(E) < \delta$. Since λ is outer regular, there exists an open set U with $E \subset U$ and $\lambda(U) < \delta$. There is a sequence U_k of finite disjoint unions of k open intervals such that $U_k \uparrow U$ as $k \rightarrow \infty$. Since $\lambda(U_k) < \delta$, it follows that the sum of the increases $\nu_F(U_k) < \epsilon'$. However the sequence $\nu_F(U_k) \uparrow \nu_F(U)$ as $k \rightarrow \infty$. Hence $\nu_F(U) \leq \epsilon'$. Hence $\nu_F(E) \leq \epsilon' < \epsilon$. This establishes the $\epsilon - \delta$ condition that is equivalent to absolute continuity. \square

Suppose that F is an absolutely continuous increasing function that is bounded. Then the corresponding finite measure ν_F is absolutely continuous with respect to Lebesgue measure, and hence there is a measurable function $h \geq 0$ with finite integral such that

$$\nu_F(f) = \lambda(hf). \quad (3.8)$$

Explicitly, this says that

$$\int_{-\infty}^{\infty} f(x) dF(x) = \int_{-\infty}^{\infty} f(x)h(x) dx. \quad (3.9)$$

Take $f(x)$ to be the indicator function of the interval from a to b . Then we obtain

$$F(b) - F(a) = \int_{-a}^b h(x) dx. \quad (3.10)$$

So the absolutely continuous increasing functions are precisely those functions that can be written as indefinite integrals of positive functions.

There is a more general concept of absolutely continuous function that corresponds to a signed measure that is absolutely continuous with respect to Lebesgue measure. These absolutely continuous functions are the indefinite integrals of integrable functions.

It is not true that the derivative of an absolutely continuous function exists at every point. However a famous theorem of Lebesgue says that it exists at almost every point and that the function can be recovered from its derivative by integration.

3.4 Problems

1. Say that μ is a finite measure and $h \geq 0$ is a measurable function. Find the function g that minimizes the quantity $\frac{1}{2}\mu((1+h)g^2) - \mu(hg)$.
2. Show that if $\nu \prec \mu$, then the derivative of ν with respect to μ is h . That is, show that if there is no division by zero, then

$$\lim_{\epsilon \downarrow 0} \frac{\nu(t \leq h < t + \epsilon)}{\mu(t \leq h < t + \epsilon)} = t. \quad (3.11)$$

Hint: Prove in fact the bounds

$$t \leq \frac{\nu(t \leq h < t + \epsilon)}{\mu(t \leq h < t + \epsilon)} \leq t + \epsilon. \quad (3.12)$$

Chapter 4

Conditional Expectation

4.1 Hilbert space ideas in probability

Consider a probability space Ω, \mathcal{S}, μ . Here μ will denote the expectation or mean defined for \mathcal{S} measurable real functions. In particular $\mu(1) = 1$.

Recall that the set Ω is the set of outcomes of an experiment. A real measurable function f on Ω is called a random variable, since it is a real number that depends on the outcome of the experiment. If $\omega \in \Omega$ is an outcome, then $f(\omega)$ is the corresponding experimental number.

A measurable subset $A \subset \Omega$ is called an event. The probability of the event A is written $\mu(A)$. If $\omega \in \Omega$ is an outcome, then the event A happens when $\omega \in A$.

A real measurable function f is in $L^1(\Omega, \mathcal{S}, \mu)$ if $\mu(|f|) < +\infty$. In this case the function is called a random variable with finite first moment. The expectation $\mu(f)$ is a well-defined real number. The random variable $f - \mu(f)$ is called the centered version of f .

A real measurable function f is in $L^2(\Omega, \mathcal{S}, \mu)$ if $\mu(f^2) < +\infty$. In this case it is called a random variable with finite second moment, or with finite variance. The second moment is $\mu(f^2)$. The variance is the second moment of the centered version. In the Hilbert space language this is

$$\mu((f - \mu(f))^2) = \|f - \mu(f)\|^2. \quad (4.1)$$

This equation may be thought of in terms of projections. The projection of f onto the constant functions is $\mu(f)$. Thus the variance is the square of the length of the projection of f onto the orthogonal complement of the constant functions. It is a quantity that tells how non-constant the function is.

In probability a common notion for variance is

$$\text{Var}(f) = \mu((f - \mu(f))^2). \quad (4.2)$$

As mentioned before, this is squared length of the component orthogonal to the

constant functions. There is a corresponding notion of covariance

$$\text{Cov}(f, g) = \mu((f - \mu(f))(g - \mu(g))). \quad (4.3)$$

This is the inner product of the components orthogonal to the constant functions. Clearly $\text{Var}(f) = \text{Cov}(f, f)$.

Another quantity encountered in probability is the correlation

$$\rho(f, g) = \frac{\text{Cov}(f, g)}{\sqrt{\text{Var}(f)}\sqrt{\text{Var}(g)}}. \quad (4.4)$$

The Hilbert space interpretation of this is the cosine of the angle between the vectors (in the subspace orthogonal to constants). This explains why $-1 \leq \rho(f, g) \leq 1$.

In statistics there are similar formulas for quantities like mean, variance, covariance, and correlation. Consider, for instance, a sample vector f of n experimental numbers. Construct a probability model where each index has probability $1/n$. This is called the empirical distribution. Then f is a random variable, and so its mean and variance can be computed in the usual way. These are called the sample mean and sample variance. Or consider instead a sample of n ordered pairs. This can be regarded as an ordered pair f, g , where f and g are each a vector of n experimental numbers. Then f and g are each random variables with respect to the empirical distribution on the n index points, and the covariance and correlation is computed as before. These are called the sample covariance and sample correlation. (Warning: Statisticians often use a slightly different definition for the sample variance or sample covariance, in which they divide by $n - 1$ instead of n . This does not matter for the sample correlation.)

The simplest (and perhaps most useful) case of the weak law of large numbers is pure Hilbert space theory. It says that averaging n uncorrelated random variables makes the variance get small at the rate $1/n$.

Proposition 4.1 (Weak law of large numbers) *Let f_1, \dots, f_n be random variables with means $\mu(f_j) = m$ and covariances $\text{Cov}(f_j, f_k) = \sigma^2 \delta_{jk}$. Then their average (sample mean) satisfies*

$$\mu\left(\frac{f_1 + \dots + f_n}{n}\right) = m \quad (4.5)$$

and

$$\text{Var}\left(\frac{f_1 + \dots + f_n}{n}\right) = \frac{\sigma^2}{n}. \quad (4.6)$$

4.2 Elementary notions of conditional expectation

In probability there is an elementary notion of conditional expectation given an event B with probability $\mu(B) > 0$. It is

$$\mu(f | B) = \frac{\mu(f1_B)}{\mu(B)}. \quad (4.7)$$

This defines a new expectation, corresponding to a world in which it is known that the event B has happened. There is also the special case of conditional probability

$$\mu(A | B) = \frac{\mu(A \cap B)}{\mu(B)}. \quad (4.8)$$

Even these elementary notions can be confusing. Here is a famous problem, a variant of the shell game.

Suppose you're on a game show and you're given the choice of three doors. Behind one is a car, behind each of the others is a goat. You pick a door, say door a , and the host, who knows what's behind the other doors, opens another door, say b , which has a goat. He then says : "Do you want to switch to door c ?" Is it to your advantage to take the switch?

Here is a simple probability model for the game show. Let X be the door with the car. Then $P[X = a] = P[X = b] = P[X = c] = 1/3$. Suppose the contestant always initially chooses door a .

Solution 1: The host always opens door b . Then we are looking at conditional probabilities given $X \neq b$. Then $P[X = a | X \neq b] = (1/3)/(2/3) = 1/2$ and $P[X = c | X \neq b] = (1/3)/(2/3) = 1/2$. There is no advantage to switching. However this careless reading of the problem overlooks the hint that the host knows X .

Solution 2. The host always opens a door without a car. The door he opens is $g(X)$, where $g(b) = c$ and $g(c) = b$ and where for definiteness $g(a) = b$. Let f be defined by $f(a) = b$ and $f(b) = c$. Then the contestant can choose a or can switch and choose $f(g(X))$. There is no need to condition on $g(X) \neq X$, since it is automatically satisfied. The probabilities are then $P[X = a] = 1/3$ and $P[X = f(g(X))] = 2/3$. It pays to switch. This is the solution that surprised so many people.

4.3 The L^2 theory of conditional expectation

The idea of conditional expectation is that there is a smaller σ -algebra of measurable functions \mathcal{F} with random variables that convey partial information about the result of the experiment.

For instance, suppose that g is a random variable that may be regarded as already measured. Then every function $\phi(g)$ is computable from g , so one may think of $\phi(g)$ as measured. The σ -algebra of functions $\sigma(g)$ generated by g consists of all $\phi(g)$, where ϕ is a Borel function.

Given a σ -algebra $\mathcal{F} \subset \mathcal{S}$, we have the closed subspace

$$L^2(\Omega, \mathcal{F}, \mu) \subset L^2(\Omega, \mathcal{S}, \mu). \quad (4.9)$$

Suppose f is in $L^2(\Omega, \mathcal{S}, \mu)$. The *conditional expectation* $\mu(f | \mathcal{F})$ is defined to be the orthogonal projection of f onto the closed subspace $L^2(\Omega, \mathcal{F}, \mu)$.

The conditional expectation satisfies the usual properties of orthogonal projection. Thus $\mu(f | \mathcal{F})$ is a random variable in \mathcal{F} , and $f - \mu(f | \mathcal{F})$ is orthogonal to $L^2(\Omega, \mathcal{F}, \mu)$. This says that for all g in $L^2(\Omega, \mathcal{F}, \mu)$ we have $\langle \mu(f | \mathcal{F}), g \rangle = \langle f, g \rangle$, that is,

$$\mu(\mu(f | \mathcal{F})g) = \mu(fg). \quad (4.10)$$

If we take $g = 1$, then we get the important equation

$$\mu(\mu(f | \mathcal{F})) = \mu(f). \quad (4.11)$$

This says that we can compute the expectation $\mu(f)$ in two stages: first compute the conditional expectation random variable $\mu(f | \mathcal{F})$, then compute its expectation. In other words, work out the prediction based on the first stage of the experiment, then use these results to compute the prediction for the total experiment.

Proposition 4.2 *The conditional expectation is order-preserving. If $f \leq g$, then $\mu(f | \mathcal{F}) \leq \mu(g | \mathcal{F})$.*

Proof: First we prove that if $h \geq 0$, then $\mu(h | \mathcal{F}) \geq 0$. Consider $h \geq 0$. Let E be the set where $\mu(h | \mathcal{F}) < 0$. Then 1_E is in \mathcal{F} , so $\mu(\mu(f | \mathcal{F})1_E) = \mu(f1_E) \geq 0$. This can only happen if $\mu(f | \mathcal{F}) = 0$ almost everywhere on E . We can then apply this to $h = g - f$. \square

Corollary 4.3 *The random variables $\mu(f | \mathcal{F})$ and $\mu(|f| | \mathcal{F})$ satisfy $|\mu(f | \mathcal{F})| \leq \mu(|f| | \mathcal{F})$.*

Proof: Since $\pm f \leq |f|$, we have $\pm\mu(f | \mathcal{F}) \leq \mu(|f| | \mathcal{F})$. \square

Corollary 4.4 *The expectations satisfy $\mu(|\mu(f | \mathcal{F})|) \leq \mu(|f|)$.*

Here is the easiest example of a conditional expectation. Suppose that there is a partition of Ω into a countable family of disjoint measurable sets B_j with union Ω . Suppose that the probability of each B_j is strictly positive, that is, $\mu(B_j) > 0$. Let \mathcal{B} be the σ -algebra of measurable functions generated by the indicator functions 1_{B_j} . The functions in \mathcal{B} are constant on each set B_j . Then the conditional expectation of f with finite variance is the projection

$$\mu(f | \mathcal{B}) = \sum_j \frac{\langle 1_{B_j}, f \rangle}{\langle 1_{B_j}, 1_{B_j} \rangle} 1_{B_j}. \quad (4.12)$$

Explicitly, this is

$$\mu(f | \mathcal{B}) = \sum_j \mu(f | B_j) 1_{B_j}. \quad (4.13)$$

Now specialize to the case when $f = 1_A$. Then this is the conditional probability random variable

$$\mu(A | \mathcal{B}) = \sum_j \mu(A | B_j) 1_{B_j}. \quad (4.14)$$

This is the usual formula for conditional probability. It says that the conditional probability of A given which of the events B_j happened depends on the outcome of the experiment. If the outcome is such that a particular B_j happened, then the value of the conditional probability is $\mu(A | B_j)$.

In this example the formula that the expectation of the conditional expectation is the expectation takes the form

$$\mu(f) = \sum_j \mu(f | B_j) \mu(B_j). \quad (4.15)$$

The corresponding formula for probability is

$$\mu(A) = \sum_j \mu(A | B_j) \mu(B_j). \quad (4.16)$$

Sometimes the notation $\mu(f | g)$ is used to mean $\mu(f | \sigma(g))$, where $\sigma(g)$ is the σ -algebra of measurable random variables generated by the random variable g . Since $\mu(f | g)$ belongs to this σ -algebra of functions, we have $\mu(f | g) = \phi(g)$ for some Borel function ϕ . Thus the conditional expectation of f given g consists of the function $\phi(g)$ of g that best predicts f based on the knowledge of the value of g . Notice that the special feature of probability is not the projection operation, which is pure Hilbert space, but the nonlinear way of generating the closed subspace on which one projects.

4.4 The L^1 theory of conditional expectation

Consider again a probability space Ω, \mathcal{S}, μ . The conditional expectation may be defined for f in $L^1(\Omega, \mathcal{S}, \mu)$. Let $\mathcal{F} \subset \mathcal{S}$ be a σ -algebra of functions. Let $f_n = f$ where $|f| \leq n$ and let $f_n = 0$ elsewhere. Then $f_n \rightarrow f$ in $L^1(\Omega, \mathcal{S}, \mu)$, by the dominated convergence theorem. So $L^2(\Omega, \mathcal{S}, \mu)$ is dense in $L^1(\Omega, \mathcal{S}, \mu)$. Furthermore, by a previous corollary the map $f \mapsto \mu(f | \mathcal{F})$ (defined as a projection in Hilbert space) is uniformly continuous with respect to the $L^1(\Omega, \mathcal{S}, \mu)$ norm. Therefore it extends by continuity to all of $L^1(\Omega, \mathcal{S}, \mu)$. In other words, for each f in $L^1(\Omega, \mathcal{S}, \mu)$ the conditional expectation is defined and is an element of $L^1(\Omega, \mathcal{S}, \mu)$.

It is not hard to see that for f in $L^1(\Omega, \mathcal{S}, \mu)$ the conditional expectation $\mu(f | \mathcal{F})$ is the element of $L^1(\Omega, \mathcal{S}, \mu)$ characterized by the following two properties. The first is that f is in \mathcal{F} , or equivalently, that f is in $L^1(\Omega, \mathcal{F}, \mu)$. The second is that for all g in $L^\infty(\Omega, \mathcal{F}, \mu)$ we have $\mu(\mu(f | \mathcal{F})g) = \mu(fg)$.

Here is a technical remark. There is another way to construct the L^1 conditional expectation by means of the Radon-Nikodym theorem. Say that $f \geq 0$ is in $L^1(\Omega, \mathcal{S}, \mu)$. The idea is to look at the finite measure ν defined on \mathcal{F} measurable functions $g \geq 0$ by $\nu(g) = \mu(fg)$. Suppose that $\mu(g) = 0$. Then the set where $g > 0$ is in \mathcal{F} with μ measure zero, and so the set where $fg > 0$ is in \mathcal{S} with μ measure zero. So $\nu(g) = 0$. This shows that $\nu \ll \mu$ as measures defined for \mathcal{F} measurable functions. By the Radon-Nikodym theorem there exists an $h \geq 0$ in $L^1(\Omega, \mathcal{F}, \mu)$ such that $\nu(g) = \mu(hg)$ for all $g \geq 0$ that are \mathcal{F} measurable. This h is the desired conditional expectation $h = \mu(f | \mathcal{F})$.

Here is an example where conditional expectation calculations are simple. Say that $\Omega = \Omega_1 \times \Omega_2$ is a product space. The σ -algebra $\mathcal{S} = \mathcal{S}_1 \otimes \mathcal{S}_2$. There is a product reference measure $\nu = \nu_1 \times \nu_2$. The actual probability measure μ has a density w with respect to this product measure:

$$\mu(f) = (\nu_1 \times \nu_2)(fw) = \int \int f(x, y)w(x, y) d\nu_1(x) d\nu_2(y). \quad (4.17)$$

Thus the experiment is carried on in two stages. What prediction can we make if we know the result for the first stage? Let $\mathcal{F}_1 = \mathcal{S}_1 \otimes \mathbf{R}$ consist of the functions $g(x, y) = h(x)$ where h is in \mathcal{S}_1 . This is the information given by the first stage. It is easy to compute the prediction $\mu(f | \mathcal{F}_1)$ for the second stage. The answer is

$$\mu(f | \mathcal{F}_1)(x, y) = \frac{\int f(x, y')w(x, y') d\nu_2(y')}{\int w(x, y') d\nu_2(y')}. \quad (4.18)$$

This is easy to check from the definition. Notice that the conditional expectation only depends on the first variable, so it is in \mathcal{F}_1 . For those who like to express such results without the use of bound variables, the answer may also be written as

$$\mu(f | \mathcal{F}_1) = \frac{\nu_2 \circ (fw)^{|1}}{\nu_2 \circ w^{|1}}. \quad (4.19)$$

A notation such as $w^{|1}$ means the function that assigns to each x the function $y \mapsto w(x, y)$ of the second variable. Thus $\nu_2 \circ w^{|1}$ means the composite function that assigns to each x the integral $\nu_2(w^{|1}(x)) = \int w(x, y) d\nu_2(y)$ of this function of the second variable.

4.5 Problems

1. Deduce the weak law of large numbers as a consequence of Hilbert space theory.
2. Consider the game show problem with the three doors a, b, c and prize $X = a, b$, or c with probability $1/3$ for each. Recall that the host chooses $g(X)$, where $g(c) = b$ and $g(b) = c$ and also $g(a) = b$, though this is not known to the contestant. (i) Find $P[X = a \mid g(X) = b]$ and $P[X = f(g(X)) \mid g(X) = b]$. If the game show host chooses b , does the contestant gain by switching? (ii) Find $P[X = a \mid g(X) = c]$ and $P[X = f(g(X)) \mid g(X) = c]$. If the game show host chooses c , does the contestant gain by switching? (iii) Find the probabilities $P[g(X) = b]$ and $P[g(X) = c]$. (iv) Consider the random variable with value $P[X = f(g(X)) \mid g(X) = b]$ provided that $g(X) = b$ and with value $P[X = f(g(X)) \mid g(X) = c]$ provided that $g(X) = c$. Find the expectation of this random variable.
3. Say that f is a random variable with finite variance, and g is another random variable. How can one choose the function ϕ to make the expectation $\mu((f - \phi(g))^2)$ as small as possible?
4. Let $\lambda > 0$ be a parameter describing the rate at which accidents occur. Let W_1 be the time to wait for the first accident, and let W_2 be the time to wait from then until the second accident. These are each exponentially distributed random variables, and their joint distribution is given by a product measure. Thus

$$\mu(f(W_1, W_2)) = \int_0^\infty \int_0^\infty f(w_1, w_2) \lambda \exp(-\lambda w_1) \lambda \exp(-\lambda w_2) dw_1 dw_2. \quad (4.20)$$

Let $T_1 = W_1$ be the time of the first accident, and let $T_2 = W_1 + W_2$ be the time of the second accident. Show that

$$\mu(h(T_1) \mid T_2) = \frac{1}{T_2} \int_0^{T_2} h(u) du. \quad (4.21)$$

That is, show that given the time T_2 of the second accident, the time T_1 of the first accident is uniformly distributed over the interval $[0, T_2]$. Hint: Make the change of variable $t_1 = w_1$ and $t_2 = w_1 + w_2$ and integrate with respect to $dt_1 dt_2$. Be careful about the limits of integration.

5. Can a σ -algebra of measurable functions (closed under pointwise operations of addition, multiplication, sup, inf, limits) be a finite dimensional vector space? Describe all such examples.

Chapter 5

Fourier series

5.1 Periodic functions

Let T be the circle parameterized by $[0, 2\pi)$ or by $[-\pi, \pi)$. Let f be a complex function in $L^2(T)$. The n th Fourier coefficient is

$$c_n = \frac{1}{2\pi} \int_0^{2\pi} e^{-inx} f(x) dx. \quad (5.1)$$

The goal is to show that f has a representation as a Fourier series

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx}. \quad (5.2)$$

Another goal is to establish the equality

$$\frac{1}{2\pi} \int_0^{2\pi} |f(x)|^2 dx = \sum_{n=-\infty}^{\infty} |c_n|^2. \quad (5.3)$$

There are two problems with the Fourier series representation. One is to interpret the sense in which the series converges. The second is to show that it actually converges to f .

Before turning to these issues, it is worth looking at the intuitive significance of these formulas. Write $e^{inx} = \cos(nx) + i \sin(nx)$. Then

$$f(x) = \frac{1}{2} a_0 + \sum_{n=1}^{\infty} [a_n \cos(nx) + b_n \sin(nx)], \quad (5.4)$$

where $a_n = c_n + c_{-n}$ and $b_n = i(c_n - c_{-n})$ for $n \geq 0$. Note that $b_0 = 0$. Also $2c_n = a_n - ib_n$ and $2c_{-n} = a_n + ib_n$ for $n \geq 0$. Furthermore, $|a_n|^2 + |b_n|^2 = 2(|c_n|^2 + |c_{-n}|^2)$.

In some applications $f(x)$ is real and the coefficients a_n and b_n are real. This is equivalent to $c_{-n} = \bar{c}_n$. In this case for $n \geq 0$ we can write $a_n = r_n \cos(\phi_n)$

and $b_n = r_n \sin(\phi_n)$, where $r_n = \sqrt{a_n^2 + b_n^2} \geq 0$. Thus ϕ_0 is an integer multiple of π . Then the series becomes for real $f(x)$

$$f(x) = \frac{1}{2}r_0 \cos(-\phi_0) + \sum_{n=1}^{\infty} r_n \cos(nx - \phi_n). \quad (5.5)$$

We see that the r_n determines the amplitude of the wave at angular frequency n , while the ϕ_n is a phase. Notice that the complex form coefficients are then $2c_n = r_n e^{-i\phi_n}$ and $c_{-n} = r_n e^{i\phi_n}$. Thus the coefficients in the complex expansion carry both the amplitude and phase information.

5.2 Convolution

It is possible to define Fourier coefficients for f in $L^1(T)$. The formula is

$$c_n = \frac{1}{2\pi} \int_0^{2\pi} e^{-inx} f(x) dx. \quad (5.6)$$

It is clear that the sequence of coefficients is in ℓ^∞ .

If f and g are in $L^1(T)$, then we may define their *convolution* $f * g$ by

$$(f * g)(x) = \frac{1}{2\pi} \int_T f(x-y)g(y) dy = \frac{1}{2\pi} \int_T f(z)g(x-z) dz. \quad (5.7)$$

All integrals are over the circle continued periodically.

Proposition 5.1 *If f has Fourier coefficients c_n and g has Fourier coefficients d_n , then $f * g$ has Fourier coefficients $c_n d_n$.*

Proof: This is an elementary calculation. \square

Another useful operation is the *adjoint* of a function in $L^1(T)$. The adjoint function f^* is defined by $f^*(x) = \overline{f(-x)}$.

Proposition 5.2 *If f has Fourier coefficients c_n , then its adjoint f^* has Fourier coefficients \bar{c}_n .*

5.3 Approximate delta functions

An approximate delta function is a sequence of functions δ_a for $a > 0$ with the following properties.

1. For each $a > 0$ the integral $\int_{-\infty}^{\infty} \delta_a(x) dx = 1$.
2. The function $\delta_a(x) \geq 0$ is positive.
3. For each $c > 0$ the integrals satisfy $\lim_{a \rightarrow 0} \int_{|x| \geq c} \delta_a(x) dx = 0$.

Theorem 5.3 Let δ_a for $a > 0$ be an approximate δ function. Then for each bounded continuous function f we have

$$\lim_{a \rightarrow 0} \int_{-\infty}^{\infty} f(x-y)\delta_a(y) dy = f(x). \quad (5.8)$$

Proof: By the first property

$$\int_{-\infty}^{\infty} f(x-y)\delta_a(y) dy - f(x) = \int_{-\infty}^{\infty} [f(x-y) - f(x)]\delta_a(y) dy. \quad (5.9)$$

By the second property

$$\left| \int_{-\infty}^{\infty} f(x-y)\delta_a(y) dy - f(x) \right| \leq \int_{-\infty}^{\infty} |f(x-y) - f(x)|\delta_a(y) dy. \quad (5.10)$$

Consider $\epsilon > 0$. Then by the continuity of f at x there exists $c > 0$ such that $|y| < c$ implies $|f(x-y) - f(x)| < \epsilon/2$. Suppose $|f(x)| \leq M$ for all x . Break up the integral into the parts with $|y| \geq c$ and $|y| < c$. Then using the first property on the second term we get

$$\left| \int_{-\infty}^{\infty} f(x-y)\delta_a(y) dy - f(x) \right| \leq 2M \int_{|y|>c} \delta_a(y) dy + \epsilon/2 \quad (5.11)$$

The third property says that for sufficiently small $a > 0$ we can get the first term also bounded by $\epsilon/2$. \square

There is also a concept of approximate delta function for functions on the circle T . This is what we need for the application to Fourier series. In fact, here is the example. For $0 \leq r < 1$ let

$$P_r(x) = \sum_{n=-\infty}^{\infty} r^{|n|} e^{inx} = \frac{1-r^2}{1-2r \cos(x) + r^2}. \quad (5.12)$$

The identity is proved by summing a geometric series. Then the functions $\frac{1}{2\pi} P_r(x)$ have the properties of an approximate delta function as r approaches 1. Each such function is positive and has integral 1 over the periodic interval. Furthermore,

$$P_r(x) \leq \frac{1-r^2}{2r(1-\cos(x))}, \quad (5.13)$$

which approaches zero as $r \rightarrow 1$ away from points where $\cos(x) = 1$.

5.4 Summability

Theorem 5.4 Let f in $C(T)$ be a continuous function on the circle. Then

$$f(x) = \lim_{r \uparrow 1} \sum_{n=-\infty}^{\infty} r^{|n|} c_n e^{inx}. \quad (5.14)$$

Proof: Proof: It is easy to compute that

$$\frac{1}{2\pi} \int_0^{2\pi} P_r(y) f(x-y) dy = \sum_{n=-\infty}^{\infty} r^{|n|} c_n e^{inx}. \quad (5.15)$$

Let $r \uparrow 1$. Then by the theorem on approximate delta functions

$$f(x) = \lim_{r \uparrow 1} \sum_{n=-\infty}^{\infty} r^{|n|} c_n e^{inx}. \quad (5.16)$$

□

Corollary 5.5 *Let f in $C(T)$ be a continuous function on the circle. Suppose that the Fourier coefficients c of f are in ℓ^1 . Then*

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx}. \quad (5.17)$$

The convergence is uniform.

Proof: If in addition c is in ℓ^1 , then the dominated convergence theorem for sums says it is possible to interchange the limit and the sum. □

5.5 L^2 convergence

The simplest and most useful theory is in the context of Hilbert space. Let $L^2(T)$ be the space of all (Borel measurable) functions such that

$$\|f\|_2^2 = \frac{1}{2\pi} \int_0^{2\pi} |f(x)|^2 dx < \infty. \quad (5.18)$$

Then $L^2(T)$ is a Hilbert space with inner product

$$\langle f, g \rangle = \frac{1}{2\pi} \int_0^{2\pi} \overline{f(x)} g(x) dx. \quad (5.19)$$

Here T is the circle, regarded as parameterized by an angle that goes from 0 to 2π .

Let

$$\phi_n(x) = \exp(inx). \quad (5.20)$$

Then the ϕ_n form an orthonormal family in $L^2(T)$. It follows from general Hilbert space theory (theorem of Pythagoras) that

$$\|f\|_2^2 = \sum_{|n| \leq N} |c_n|^2 + \|f - \sum_{|n| \leq N} c_n \phi_n\|_2^2. \quad (5.21)$$

In particular, Bessel's inequality says that

$$\|f\|_2^2 \geq \sum_{|n| \leq N} |c_n|^2. \quad (5.22)$$

This shows that

$$\sum_{n=-\infty}^{\infty} |c_n|^2 < \infty. \quad (5.23)$$

The space of sequences satisfying this identity is ℓ^2 . Thus we have proved the following proposition.

Proposition 5.6 *If f is in $L^2(T)$, then its sequence of Fourier coefficients is in ℓ^2 .*

Theorem 5.7 *If f is in $L^2(T)$, then*

$$\|f\|_2^2 = \sum_n |c_n|^2. \quad (5.24)$$

Proof: The function \overline{f} has Fourier coefficients c_n . The adjoint function f^* defined by $f^*(x) = \overline{f(-x)}$ has complex conjugate Fourier coefficients $\overline{c_n}$. The coefficients of a convolution are the product of the coefficients. Hence $g = f^* * f$ has coefficients $\overline{c_n}c_n = |c_n|^2$.

Suppose that f is in $L^2(T)$. Then $g = f^* * f$ is in $C(T)$. In fact,

$$g(x) = \frac{1}{2\pi} \int_T \overline{f(y-x)} f(y) dy = \langle f_x, f \rangle, \quad (5.25)$$

where f_x is f translated by x . Since translation is continuous in $L^2(T)$, it follows that g is a continuous function. Furthermore, since f is in $L^2(T)$, it follows that c is in ℓ^2 , and so $|c|^2$ is in ℓ^1 . Thus the theorem applies, and

$$g(x) = \frac{1}{2\pi} \int_T \overline{f(y-x)} f(y) dy = \sum_n |c_n|^2 e^{inx}. \quad (5.26)$$

The conclusion follows by taking $x = 0$. \square

Theorem 5.8 *If f is in $L^2(T)$, then*

$$f = \sum_{n=-\infty}^{\infty} c_n \phi_n \quad (5.27)$$

in the sense that

$$\lim_{N \rightarrow \infty} \|f - \sum_{|n| \leq N} c_n \phi_n\|_2^2 = 0. \quad (5.28)$$

Proof: Use the identity

$$\|f\|_2^2 = \sum_{|n| \leq N} |c_n|^2 + \|f - \sum_{|n| \leq N} c_n \phi_n\|_2^2. \quad (5.29)$$

The first term on the right hand side converges to the left hand side, so the second term on the right hand side must converge to zero. \square

5.6 $C(T)$ convergence

Define the function spaces

$$C(T) \subset L^\infty(T) \subset L^2(T) \subset L^1(T). \quad (5.30)$$

The norms $\|f\|_\infty$ on the first two spaces are the same, the smallest number M such that $|f(x)| \leq M$ (with the possible exception of a set of x of measure zero). The space $C(T)$ consists of continuous functions; the space $L^\infty(T)$ consists of all bounded functions. The norm on $L^2(T)$ is given by $\|f\|_2^2 = \frac{1}{2\pi} \int_0^{2\pi} |f(x)|^2 dx$. The norm on $L^1(T)$ is given by $\|f\|_1 = \frac{1}{2\pi} \int_0^{2\pi} |f(x)| dx$. Since the integral is a probability average, their relation is

$$\|f\|_1 \leq \|f\|_2 \leq \|f\|_\infty. \quad (5.31)$$

Also define the sequence spaces

$$\ell^1 \subset \ell^2 \subset c_0 \subset \ell^\infty. \quad (5.32)$$

The norm on ℓ^1 is $\|c\|_1 = \sum_n |c_n|$. Then norm on ℓ^2 is given by $\|c\|_2^2 = \sum_n |c_n|^2$. The norms on the last two spaces are the same, that is, $\|c\|_\infty$ is the smallest M such that $|c_n| \leq M$. The space c_0 consists of all sequences with limit 0 at infinity. The relation between these norms is

$$\|c\|_\infty \leq \|c\|_2 \leq \|c\|_1. \quad (5.33)$$

We have seen that the Fourier series theorem gives a perfect correspondence between $L^2(T)$ and ℓ^2 . For the other spaces the situation is more complex.

Lemma 5.9 (Riemann-Lebesgue) *If f is in $L^1(T)$, then the Fourier coefficients of f are in c_0 , that is, they approach 0 at infinity.*

Proof: Each function in $L^2(T)$ has Fourier coefficients in ℓ^2 , so each function in $L^2(T)$ has Fourier coefficients that vanish at infinity. The map from a function to its Fourier coefficients gives a continuous map from $L^1(T)$ to ℓ^∞ . However every function in $L^1(T)$ may be approximated arbitrarily closely in $L^1(T)$ norm by a function in $L^2(T)$. Hence its coefficients may be approximated arbitrarily well in ℓ^∞ norm by coefficients that vanish at infinity. Therefore the coefficients vanish at infinity. \square

In summary, the map from a function to its Fourier coefficients gives a continuous map from $L^1(T)$ to c_0 . That is, the Fourier coefficients of an integrable function are bounded (this is obvious) and approach zero (Riemann-Lebesgue lemma). Furthermore, it may be shown that the Fourier coefficients determine the function uniquely.

The map from Fourier coefficients to functions gives a continuous map from ℓ^1 to $C(T)$. An sequence that is absolutely summable defines a Fourier series that converges absolutely and uniformly to a continuous function.

For the next result the following lemma will be useful.

Lemma 5.10 *Say that*

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx} \quad (5.34)$$

with $L^2(T)$ convergence. Then the identity

$$f'(x) = \sum_{n=-\infty}^{\infty} in c_n e^{inx} \quad (5.35)$$

obtained by differentiating holds, again in the sense of $L^2(T)$ convergence. Here the relation between f and f' is that f is an indefinite integral of f' . Furthermore f' has integral zero and f is periodic.

Proof: Let $h(x) = \sum_{n \neq 0} b_n e^{inx}$ be in $L^2(T)$. Define the integral Vh by $(Vh)(x) = \int_0^x h(y) dy$. Then

$$\|Vh\|_2 \leq \|Vh\|_{\infty} \leq (2\pi)\|h\|_1 \leq (2\pi)\|h\|_2. \quad (5.36)$$

This shows that V is continuous from $L^2(T)$ to $L^2(T)$. Thus we can apply V to the series term by term. This gives

$$(Vh)(x) = \sum_{n \neq 0} b_n \frac{e^{inx} - 1}{in} = C + \sum_{n \neq 0} \frac{b_n}{in} e^{inx}. \quad (5.37)$$

Thus the effect of integrating is to divide the coefficient by in . Since differentiation has been defined in this context to be the inverse of integration, the effect of differentiation is to multiply the coefficient by in . \square

Theorem 5.11 *If f is in $L^2(T)$ and if f' exists (in the sense that f is an integral of f') and if f' is also in $L^2(T)$, then the Fourier coefficients are in ℓ^1 . Therefore the Fourier series converges in the $C(T)$ norm.*

Proof: The hypothesis of the theorem means that there is a function f' in $L^2(T)$ with $\int_0^{2\pi} f'(y) dy = 0$. Then f is a function defined by

$$f(x) = c_0 + \int_0^x f'(y) dy \quad (5.38)$$

with an arbitrary constant of integration. This f is an absolutely continuous function. It is also periodic, because of the condition on the integral of f' .

The proof is completed by noting that

$$\sum_{n \neq 0} |c_n| = \sum_{n \neq 0} \frac{1}{|n|} |nc_n| \leq \sqrt{\sum_{n \neq 0} \frac{1}{n^2}} \sqrt{\sum_{n \neq 0} n^2 |c_n|^2}. \quad (5.39)$$

In other words,

$$\sum_{n \neq 0} |c_n| \leq \sqrt{\frac{\pi^2}{3}} \|f'\|_2. \quad (5.40)$$

\square

5.7 Pointwise convergence

There remains one slightly unsatisfying point. The convergence in the L^2 sense does not imply convergence at a particular point. Of course, if the derivative is in L^2 then we have uniform convergence, and in particular convergence at each point. But what if the function is differentiable at one point but has discontinuities at other points? What can we say about convergence at that one point? Fortunately, we can find something about that case by a closer examination of the partial sums.

One looks at the partial sum

$$\sum_{|n| \leq N} c_n e^{inx} = \frac{1}{2\pi} \int_0^{2\pi} D_N(x-y) f(y) dy. \quad (5.41)$$

Here

$$D_N(x) = \sum_{|n| \leq N} e^{inx} = \frac{\sin((N + \frac{1}{2})x)}{\sin(\frac{1}{2}x)}. \quad (5.42)$$

This Dirichlet kernel $D_N(x)$ has at least some of the properties of an approximate delta function. Unfortunately, it is not positive; instead it oscillates wildly for large N at points away from where $\sin(x/2) = 0$. However the function $1/(2\pi)D_N(x)$ does have integral 1.

Theorem 5.12 *If for some x the function*

$$d_x(z) = \frac{f(x+z) - f(x)}{2 \sin(z/2)} \quad (5.43)$$

is in $L^1(T)$, then at that point

$$f(x) = \sum_{n=-\infty}^{\infty} c_n \phi_n(x). \quad (5.44)$$

Note that if $d_x(z)$ is continuous at $z = 0$, then its value at $z = 0$ is $d_x(0) = f'(x)$. So the hypothesis of the theorem is a condition related to differentiability of f at the point x . The conclusion of the theorem is pointwise convergence of the Fourier series at that point. Since f may be discontinuous at other points, it is possible that this Fourier series is not absolutely convergent. Thus the series must be interpreted as the limit of the partial sums over $|n| \leq N$, taken as $N \rightarrow \infty$.

Proof: We have

$$f(x) - \sum_{|n| \leq N} c_n e^{inx} = \frac{1}{2\pi} \int_0^{2\pi} D_N(z) (f(x) - f(x-z)) dz. \quad (5.45)$$

We can write this as

$$f(x) - \sum_{|n| \leq N} c_n e^{inx} = \frac{1}{2\pi} \int_0^{2\pi} 2 \sin((N + \frac{1}{2})z) d_x(-z) dz. \quad (5.46)$$

This goes to zero as $N \rightarrow \infty$, by the Riemann-Lebesgue lemma. \square

5.8 Problems

1. Let $f(x) = x$ defined for $-\pi \leq x < \pi$. Find the $L^1(T)$, $L^2(T)$, and $L^\infty(T)$ norms of f , and compare them.
2. Find the Fourier coefficients c_n of f for all n in Z .
3. Find the ℓ^∞ , ℓ^2 , and ℓ^1 norms of these Fourier coefficients, and compare them.
4. Use the equality of L^2 and ℓ^2 norms to compute

$$\zeta(2) = \sum_{n=1}^{\infty} \frac{1}{n^2}.$$

5. Compare the ℓ^∞ and L^1 norms for this problem. Compare the L^∞ and ℓ^1 norms for this problem.
6. Use the pointwise convergence at $x = \pi/2$ to evaluate the infinite sum

$$\sum_{k=0}^{\infty} (-1)^k \frac{1}{2k+1},$$

regarded as a limit of partial sums. Does this sum converge absolutely?

7. Let $F(x) = \frac{1}{2}x^2$ defined for $-\pi \leq x < \pi$. Find the Fourier coefficients of this function.
8. Use the equality of L^2 and ℓ^2 norms to compute

$$\zeta(4) = \sum_{n=1}^{\infty} \frac{1}{n^4}.$$

9. Compare the ℓ^∞ and L^1 norms for this problem. Compare the L^∞ and ℓ^1 norms for this problem.
10. At which points x of T is $F(x)$ continuous? Differentiable? At which points x of T is $f(x)$ continuous? Differentiable? At which x does $F'(x) = f(x)$? Can the Fourier series of $f(x)$ be obtained by differentiating the Fourier series of $F(x)$ pointwise? (This last question can be answered by inspecting the explicit form of the Fourier series for the two problems.)

Chapter 6

Fourier transforms

6.1 Fourier analysis

The general context of Fourier analysis is an abelian group and its dual group. The elements x of the abelian group are thought of as space (or time) variables, while the elements of the dual group are thought of as wave number (or angular frequency) variables.

Examples:

1. Let $\Delta x > 0$. The finite group consists of all $x = j\Delta x$ for $j = 0, \dots, N - 1$ with addition mod $L = N\Delta x$. The dual group is the finite group $k = \ell\Delta k$ with $\ell = 0, \dots, N - 1$ with addition mod $N\Delta k$. Here $\Delta x\Delta k = 2\pi/N$.
2. Let $L > 0$. The compact group consists of all x in the circle of circumference L with addition mod L . The dual group is the discrete group $k = \ell\Delta k$ with $\ell \in \mathbf{Z}$. Here $\Delta k = 2\pi/L$.
3. The discrete group consists of all $x = j\Delta x$ with $j \in \mathbf{Z}$. The dual group is the compact group of k in the circle of circumference $2\pi/\Delta x$.
4. The group consists of all x in the real line. The dual group is all k in the (dual) line.

In each case the formula are the essentially the same. Let $\lambda > 0$ be an arbitrary constant. We have dual measures dx/λ and $\lambda dk/(2\pi)$. Their product is $dx dk/(2\pi)$. The Fourier transform is

$$\hat{f}(k) = \int e^{-ikx} f(x) \frac{1}{\lambda} dx. \quad (6.1)$$

The integral is over the group. For a finite or discrete group it is a sum, and the dx is replaced by Δx . The Fourier representation is then given by the inversion formula

$$f(x) = \int e^{ikx} \hat{f}(k) \lambda \frac{dk}{2\pi}. \quad (6.2)$$

The integral is over the dual group. For a finite or discrete dual group it is a sum, and the dk is replaced by Δk .

The constant $\lambda > 0$ is chosen for convenience. There is a lot to be said for standardizing on $\lambda = 1$. In the case of the circle one variant is $\lambda = L$. This choice makes dx/λ a probability measure and $\lambda\Delta k/(2\pi) = 1$. Similarly, in the case of the discrete group $\lambda = \Delta x$ makes $\Delta x/\lambda = 1$ and $\lambda dk/(2\pi)$ a probability measure. In the case of the line $\lambda = 1$ is most common, though some people prefer the ugly choice $\lambda = \sqrt{2\pi}$ in a misguided attempt at symmetry.

6.2 L^1 theory

Let f be a complex function on the line that is in L^1 . The Fourier transform \hat{f} is the function defined by

$$\hat{f}(k) = \int_{-\infty}^{\infty} e^{-ikx} f(x) dx. \quad (6.3)$$

Note that if f is in L^1 , then its Fourier transform \hat{f} is in L^∞ and satisfies $\|\hat{f}\|_\infty \leq \|f\|_1$. Furthermore, it is a continuous function.

Similarly, let g be a function on the (dual) line that is in L^1 . Then the inverse Fourier transform \check{g} is defined by

$$\check{g}(x) = \int_{-\infty}^{\infty} e^{ikx} g(k) \frac{dk}{2\pi}. \quad (6.4)$$

We can look at the Fourier transform from a more abstract point of view. The space L^1 is a Banach space. Its dual space is L^∞ , the space of essentially bounded functions. An example of a function in the dual space is the exponential function $\phi_k(x) = e^{ikx}$. The Fourier transform is then

$$\hat{f}(k) = \langle \phi_k, f \rangle = \int_{-\infty}^{\infty} \overline{\phi_k(x)} f(x) dx, \quad (6.5)$$

where ϕ_k is in L^∞ and f is in L^1 .

Proposition 6.1 *f, g are in $L^1(\mathbf{R}, dx)$, then the convolution $f * g$ is another function in $L^1(\mathbf{R}, dx)$ defined by*

$$(f * g)(x) = \int_{-\infty}^{\infty} f(x - y)g(y) dy. \quad (6.6)$$

Proposition 6.2 *If f, g are in $L^1(\mathbf{R}, dx)$, then the Fourier transform of the convolution is the product of the Fourier transforms:*

$$(\widehat{f * g})(k) = \hat{f}(k)\hat{g}(k). \quad (6.7)$$

Proposition 6.3 *Let $f^*(x) = \overline{f(-x)}$. Then the Fourier transform of f^* is the complex conjugate of \hat{f} .*

Theorem 6.4 *If f is in L^1 and is also continuous and bounded, we have the inversion formula in the form*

$$f(x) = \lim_{\epsilon \downarrow 0} \int_{-\infty}^{\infty} e^{ikx} \hat{\delta}_\epsilon(k) \hat{f}(k) \frac{dk}{2\pi}, \quad (6.8)$$

where

$$\hat{\delta}_\epsilon(k) = \exp(-\epsilon|k|). \quad (6.9)$$

Proof: The inverse Fourier transform of this is

$$\delta_\epsilon(x) = \frac{1}{\pi} \frac{\epsilon}{x^2 + \epsilon^2}. \quad (6.10)$$

It is easy to calculate that

$$\int_{-\infty}^{\infty} e^{ikx} \hat{\delta}_\epsilon(k) \hat{f}(k) \frac{dk}{2\pi} = (\delta_\epsilon * f)(x). \quad (6.11)$$

However δ_ϵ is an approximate delta function. The result follows by taking $\epsilon \rightarrow 0$. \square

6.3 L^2 theory

The space L^2 is its own dual space, and it is a Hilbert space. It is the setting for the most elegant and simple theory of the Fourier transform.

Lemma 6.5 *If f is in $L^1(\mathbf{R}, dx)$ and in $L^2(\mathbf{R}, dx)$, then \hat{f} is in $L^2(\mathbf{R}, dk/(2\pi))$, and $\|f\|_2^2 = \|\hat{f}\|_2^2$.*

Proof: Let $g = f * f$. Then g is in L^1 , since it is the convolution of two L^1 functions. Furthermore, it is continuous and bounded. This follows from the representation $g(x) = \langle f_x, f \rangle$, where f_x is translation by x . Since $x \rightarrow f_x$ is continuous from \mathbf{R} to L^2 , the result follows from the Hilbert space continuity of the inner product. Finally, the Fourier transform of g is $|\hat{f}(k)|^2$. Thus

$$\|f\|_2^2 = g(0) = \lim_{\epsilon \downarrow 0} \int_{-\infty}^{\infty} \hat{\delta}_\epsilon(k) |\hat{f}(k)|^2 \frac{dk}{2\pi} = \int_{-\infty}^{\infty} |\hat{f}(k)|^2 \frac{dk}{2\pi} \quad (6.12)$$

by the monotone convergence theorem. \square

Theorem 6.6 *The Fourier transform F initially defined on $L^1(\mathbf{R}, dx) \cap L^2(\mathbf{R}, dx)$ extends by uniform continuity to $F : L^2(\mathbf{R}, dx) \rightarrow L^2(\mathbf{R}, dk/(2\pi))$. The inverse Fourier transform F^* initially defined on $L^1(\mathbf{R}, dk/(2\pi)) \cap L^2(\mathbf{R}, dk/(2\pi))$ extends by uniform continuity to $F^* : L^2(\mathbf{R}, dk/(2\pi)) \rightarrow L^2(\mathbf{R}, dx)$. These are linear transformations that preserve L^2 norm and preserve inner product. Furthermore, F^* is the inverse of F .*

Proof: It is easy to see that $L^1 \cap L^2$ is dense in L^2 . Here is the proof. Take f in L^2 and let A_n be a sequence of sets of finite measure that increase to all of \mathbf{R} . Then $1_{A_n}f$ is in L^1 for each n , by the Schwarz inequality. Furthermore, $1_{A_n}f \rightarrow f$ in L^2 , by the L^2 dominated convergence theorem.

The lemma shows that F is an isometry, hence uniformly continuous. Furthermore, the target space L^2 is a complete metric space. Thus F extends by uniform continuity to the entire domain space L^2 . It is easy to see that this extension is also an isometry.

The same reasoning shows that the inverse Fourier transform F^* also maps L^2 onto L^2 and preserves norm.

Now it is easy to check that $(F^*h, f) = (h, Ff)$ for f and h in $L^1 \cap L^2$. This identity extends to all of L^2 . Take $h = Fg$. Then $\langle F^*Fg, f \rangle = \langle Fg, Ff \rangle = \langle g, f \rangle$. That is $F^*Fg = g$. Similarly, one may show that $FF^*u = u$. These equations show that $F^* = F^{-1}$ is the inverse of F . \square

Corollary 6.7 *Let f be in L^2 . Let A_n be a sequences subsets of finite measure that increase to all of \mathbf{R} . Then $1_{A_n}f$ is in $L^1 \cap L^2$ and $F(1_{A_n}f) \rightarrow F(f)$ in L^2 as $n \rightarrow \infty$. That is, for fixed n the function with values given by*

$$F(1_{A_n}f)(k) = \int_{A_n} e^{-ikx} f(x) dx \quad (6.13)$$

is well defined for each k , and the sequence of such functions converges in the L^2 sense to the Fourier transform Ff , where the function $(Ff)(k)$ is defined for almost every k . Explicitly, this says that the Fourier transform of an L^2 function f is the L^2 function $\hat{f} = Ff$ characterized by

$$\int_{-\infty}^{\infty} |\hat{f}(k) - \int_{A_n} e^{-ikx} f(x) dx|^2 \frac{dk}{2\pi} \rightarrow 0 \quad (6.14)$$

as $n \rightarrow \infty$.

6.4 Absolute convergence

We have seen that the Fourier transform gives a perfect correspondence between $L^2(\mathbf{R}, dx)$ and $L^2(\mathbf{R}, dk/(2\pi))$. For the other spaces the situation is more complicated.

Theorem 6.8 (Riemann-Lebesgue lemma) *The map from a function to its Fourier transform gives a continuous map from $L^1(\mathbf{R}, dx)$ to $C_0(\mathbf{R})$. That is, the Fourier transform of an integrable function is continuous and bounded and approaches zero at infinity.*

Proof: We have seen that the Fourier transform of an L^1 function is bounded and continuous. The main content of the Riemann-Lebesgue lemma is that it also goes to zero at infinity. This can be proved by checking it on a dense subset, such as the space of step functions. \square

One other useful fact is that if f is in $L^1(\mathbf{R}, dx)$ and g is in $L^2(\mathbf{R}, dx)$, then the convolution $f * g$ is in $L^2(\mathbf{R}, dx)$. Furthermore, $\widehat{f * g}(k) = \hat{f}(k)\hat{g}(k)$ is the product of a bounded function with an $L^2(\mathbf{R}, dk/(2\pi))$ function and therefore is in $L^2(\mathbf{R}, dk/(2\pi))$.

However the same pattern of the product of a bounded function with an $L^2(\mathbf{R}, dk/(2\pi))$ function can arise in other ways. For instance, consider the translate f_a of a function f in $L^2(\mathbf{R}, dx)$ defined by $f_a(x) = f(x - a)$. Then $\hat{f}_a(k) = \exp(-ika)\hat{f}(k)$. This is also the product of a bounded function with an $L^2(\mathbf{R}, dk/(2\pi))$ function.

One can think of this last example as a limiting case of a convolution. Let δ_ϵ be an approximate δ function. Then $(\delta_\epsilon)_a * f$ has Fourier transform $\exp(-ika)\hat{\delta}_\epsilon(k)\hat{f}(k)$. Now let $\epsilon \rightarrow 0$. Then $(\delta_\epsilon)_a * f \rightarrow f_a$, while $\exp(-ika)\hat{\delta}_\epsilon(k)\hat{f}(k) \rightarrow \exp(-ika)\hat{f}(k)$.

Theorem 6.9 *If f is in $L^2(\mathbf{R}, dx)$ and if f' exists (in the sense that f is an integral of f') and if f' is also in $L^2(\mathbf{R}, dx)$, then the Fourier transform is in $L^1(\mathbf{R}, dk/(2\pi))$. As a consequence f is in $C_0(\mathbf{R})$.*

Proof: $\hat{f}(k) = (1/\sqrt{1+k^2}) \cdot \sqrt{1+k^2}\hat{f}(k)$. Since f is in L^2 , it follows that $\hat{f}(k)$ is in L^2 . Since f' is in L^2 , it follows that $k\hat{f}(k)$ is in L^2 . Hence $\sqrt{1+k^2}\hat{f}(k)$ is in L^2 . Since $1/\sqrt{1+k^2}$ is also in L^2 , it follows from the Schwarz inequality that $\hat{f}(k)$ is in L^1 . \square

6.5 Fourier transform pairs

There are some famous Fourier transforms. Fix $\sigma > 0$, and consider first the Gaussian

$$g_\sigma(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{x^2}{2\sigma^2}\right). \quad (6.15)$$

Its Fourier transform is

$$\hat{g}_\sigma(k) = \exp\left(-\frac{\sigma^2 k^2}{2}\right). \quad (6.16)$$

Here is a proof of this Gaussian formula. Define the Fourier transform $\hat{g}_\sigma(k)$ by the usual formula. Check that

$$\left(\frac{d}{dk} + \sigma^2 k\right)\hat{g}_\sigma(k) = 0. \quad (6.17)$$

This proves that

$$\hat{g}_\sigma(k) = C \exp\left(-\frac{\sigma^2 k^2}{2}\right). \quad (6.18)$$

Now apply the equality of L^2 norms. This implies that $C^2 = 1$. By looking at the case $k = 0$ it becomes obvious that $C = 1$.

Let $\epsilon > 0$. Introduce the Heaviside function $H(k)$ that is 1 for $k > 0$ and 0 for $k < 0$. The two basic Fourier transform pairs are

$$f_\epsilon(x) = \frac{1}{x - i\epsilon} \quad (6.19)$$

with Fourier transform

$$\hat{f}_\epsilon(k) = 2\pi i H(-k) e^{\epsilon k}. \quad (6.20)$$

and its complex conjugate

$$\overline{f_\epsilon(x)} = \frac{1}{x + i\epsilon} \quad (6.21)$$

with Fourier transform

$$\overline{\hat{f}_\epsilon(-k)} = -2\pi i H(k) e^{-\epsilon k}. \quad (6.22)$$

These may be checked by computing the inverse Fourier transform. Notice that f_ϵ and its conjugate are not in $L^1(\mathbf{R})$.

Take $1/\pi$ times the imaginary part. This gives the approximate delta function

$$\delta_\epsilon(x) = \frac{1}{\pi} \frac{\epsilon}{x^2 + \epsilon^2}. \quad (6.23)$$

with Fourier transform

$$\hat{\delta}_\epsilon(k) = e^{-\epsilon|k|}. \quad (6.24)$$

Take the real part. This gives the approximate principal value of $1/x$ function

$$p_\epsilon(x) = \frac{x}{x^2 + \epsilon^2} \quad (6.25)$$

with Fourier transform

$$\hat{p}_\epsilon(k) = -\pi i [H(k) e^{-\epsilon k} - H(-k) e^{\epsilon k}]. \quad (6.26)$$

6.6 Poisson summation formula

Theorem 6.10 *Let f be in $L^1(\mathbf{R}, dx)$ with \hat{f} in $L^1(\mathbf{R}, dk/(2\pi))$ and such that $\sum_k |\hat{f}(k)| < \infty$. Then*

$$2\pi \sum_{n \in \mathbf{Z}} f(2\pi n) = \sum_{k \in \mathbf{Z}} \hat{f}(k). \quad (6.27)$$

Proof: Let

$$S(t) = \sum_n f(2\pi n + t). \quad (6.28)$$

Since $S(t)$ is 2π periodic, we can expand

$$S(t) = \sum_k a_k e^{ikt}. \quad (6.29)$$

It is easy to compute that

$$a_k = \frac{1}{2\pi} \int_0^{2\pi} S(t)e^{-ikt} dt = \frac{1}{2\pi} \hat{f}(k). \quad (6.30)$$

So the Fourier series of $S(t)$ is absolutely summable. In particular

$$S(0) = \sum_k a_k. \quad (6.31)$$

□

6.7 Problems

1. Let $f(x) = 1/(2a)$ for $-a \leq x \leq a$ and be zero elsewhere. Find the $L^1(\mathbf{R}, dx)$, $L^2(\mathbf{R}, dx)$, and $L^\infty(\mathbf{R}, dx)$ norms of f , and compare them.
2. Find the Fourier transform of f .
3. Find the $L^\infty(\mathbf{R}, dk/(2\pi))$, $L^2(\mathbf{R}, dk/(2\pi))$, and $L^1(\mathbf{R}, dk/(2\pi))$ norms of the Fourier transform, and compare them.
4. Compare the $L^\infty(\mathbf{R}, dk/(2\pi))$ and $L^1(\mathbf{R}, dx)$ norms for this problem. Compare the $L^\infty(\mathbf{R}, dx)$ and $L^1(\mathbf{R}, dk/(2\pi))$ norms for this problem.
5. Use the pointwise convergence at $x = 0$ to evaluate an improper integral.
6. Calculate the convolution of f with itself.
7. Find the Fourier transform of the convolution of f with itself. Verify in this case that the Fourier transform of the convolution is the product of the Fourier transforms.
8. In this problem the Fourier transform is band-limited, that is, only waves with $|k| \leq a$ have non-zero amplitude. Make the assumption that $|k| > a$ implies $\hat{f}(k) = 0$. That is, the Fourier transform of f vanishes outside of the interval $[-a, a]$.

Let

$$g(x) = \frac{\sin(ax)}{ax}.$$

The problem is to prove that

$$f(x) = \sum_{m=-\infty}^{\infty} f\left(\frac{m\pi}{a}\right)g\left(x - \frac{m\pi}{a}\right).$$

This says that if you know f at multiples of π/a , then you know f at all points.

Hint: Let $g_m(x) = g(x - m\pi/a)$. The task is to prove that $f(x) = \sum_m c_m g_m(x)$ with $c_m = f(m\pi/a)$. It helps to use the Fourier transform

of these functions. First prove that the Fourier transform of $g(x)$ is given by $\hat{g}(k) = \pi/a$ for $|k| \leq a$ and $\hat{g}(k) = 0$ for $|k| > a$. (Actually, it may be easier to deal with the inverse Fourier transform.) Then prove that $\hat{g}_m(k) = \exp(-im\pi k/a)\hat{g}(k)$. Finally, note that the functions $\hat{g}_m(k)$ are orthogonal.

9. In the theory of neural networks one wants to synthesize an arbitrary function from linear combinations of translates of a fixed function. Let f be a function in L^2 . Suppose that the Fourier transform $\hat{f}(k) \neq 0$ for all k . Define the translate f_a by $f_a(x) = f(x - a)$. The task is to show that the set of all linear combinations of the functions f_a , where a ranges over all real numbers, is dense in L^2 .

Hint: It is sufficient to show that if g is in L^2 with $(g, f_a) = 0$ for all a , then $g = 0$. (Why is this sufficient?) This can be done using Fourier analysis.

Chapter 7

Topology

7.1 Topological spaces

Let X be a set. The power set $P(X)$ consists of all subsets of X . In the following we shall fix attention on a universe X of points and certain subsets $U \subset X$. Often we shall want to speak of sets of subsets. For clarity, we shall often speak instead of collections of subsets. Thus a collection is a subset of the power set $P(X)$.

Let $\Gamma \subset P(X)$ be a collection of sets. Recall the definitions of union and intersection:

$$\bigcup \Gamma = \{x \in X \mid \exists U (U \in \Gamma \wedge x \in U)\} \quad (7.1)$$

and

$$\bigcap \Gamma = \{x \in X \mid \forall U (U \in \Gamma \Rightarrow x \in U)\}. \quad (7.2)$$

Thus the union and intersection are each a subset of X .

A *topology* on X is a subcollection \mathcal{T} of $P(X)$ with the following two properties:

1. If $\Gamma \subset \mathcal{T}$, then $\bigcup \Gamma \in \mathcal{T}$.
2. If $\Gamma \subset \mathcal{T}$ is finite, then $\bigcap \Gamma \in \mathcal{T}$.

The set X with a given topology \mathcal{T} is called a *topological space*.

It follows from the first property that $\bigcup \emptyset = \emptyset \in \mathcal{T}$. It follows from the second property that $\bigcap \emptyset = X \in \mathcal{T}$. (The fact that $\bigcap \emptyset = X$ follows from the convention that for $\Gamma \subset P(X)$ the universe is X .) An *open set* is a subset that is in the topology. A *closed set* is a subset that is the complement of an open set.

The *interior* $\text{int } S$ of a set S is the union of all open subsets of it. It is the largest open subset of S . A point is in the interior of S iff it belongs to an open subset of S . The *closure* \bar{S} of a set S is the intersection of all closed supersets. It is the smallest closed superset of S . A point is in the closure of S iff every open set to which it belongs intersects S in at least one point.

Let X and Y each have a topology. A *continuous map* $f : X \rightarrow Y$ is a function such that for every open subset V of Y the inverse image $f^{-1}[V]$ is an open subset of X .

There is an alternate characterization of continuous maps that is often useful. A function $f : X \rightarrow Y$ is continuous if and only if for every closed subset F of Y the inverse image $f^{-1}[F]$ is a closed subset of X . This is a useful fact; it often used to show that the solutions of an equation form a closed set.

If $f : X \rightarrow Y$ is a continuous bijection with continuous inverse, then f is a topological isomorphism, or *homeomorphism*.

Examples:

1. The open unit interval $(0, 1)$ is homeomorphic to \mathbf{R} .
2. The open unit interval $(0, 1)$ is not homeomorphic to the circle S_1 . (One is compact; the other is not.)
3. The closed unit interval $[0, 1]$ is not homeomorphic to the circle S_1 . (One can be disconnected by removing a point; the other not.)

If $\Gamma \subset P(Y)$ is an arbitrary collection of subsets of X , then there is a least (smallest, coarsest) topology \mathcal{T} with $\Gamma \subset \mathcal{T}$. This is the topology generated by Γ .

For example, say that $\Gamma = \{U, V\}$, where $U \subset Y$ and $V \subset Y$. Then the topology generated by Γ is $\mathcal{T} = \{\emptyset, U \cap V, U, V, U \cup V, Y\}$ and can have up to 6 subsets in it.

Proposition 7.1 *Say that X has topology \mathcal{S} and Y has topology \mathcal{T} . Suppose also that Γ generates \mathcal{T} . Suppose that $f : X \rightarrow Y$ and that for every V in the generating set Γ the inverse image $f^{-1}[V]$ is an open set in \mathcal{S} . Then f is a continuous map.*

An example where this applies is when Y is a metric space. It says that in this case it is enough to check that the inverse images of open balls are open.

If Y is a topological space, and if Z is a subset of Y , then there is a *relative topology* on Z . It is defined as the collection of all $U \cap Z$ for $U \subset Y$ open. The most common way of defining a topology space is to take some well known space, such as $Y = \mathbf{R}^n$ and then indicate some subset $Z \subset Y$. Even though the topology of Z is the relative topology derived from Y , it is important that one can forget about this and think of Z as a topological space that is a universe with its own topology.

As an example, take the case when $Y = \mathbf{R}$ and $Z = [0, 1]$. Then a set like $[0, 1/2)$ is an open subset of Z , even though it is not an open subset of \mathbf{R} . The reason is that $[0, 1/2)$ is the intersection of $(-2, 1/2)$ with Z , and $(-2, 1/2)$ is an open subset of \mathbf{R} .

If $f : X \rightarrow Y$ is a continuous injection that gives a homeomorphism of f with $Z \subset Y$, where Z has the relative topology, then f is said to be an *embedding* of X into Y .

Examples:

1. There is an embedding of the open interval $(0, 1)$ into the circle S_1 . The range of the embedding is the circle with a single point removed.
2. There is no embedding of the circle into the open interval. A continuous image of the circle in the open interval is compact and connected, and thus is a closed subinterval. But a circle is not homeomorphic to a closed interval.

7.2 Measurable spaces

Let X be a set. A σ -algebra of subsets of X is a subcollection \mathcal{F} of $P(X)$ with the following three properties:

1. If $\Gamma \subset \mathcal{F}$ is countable, then $\bigcup \Gamma \in \mathcal{F}$.
2. If $\Gamma \subset \mathcal{F}$ is countable, then $\bigcap \Gamma \in \mathcal{F}$.
3. If $A \in \mathcal{F}$, then $X \setminus A \in \mathcal{F}$.

The set X with a given σ -algebra of measurable subsets is called a *measurable space*. An *measurable set* is a subset that is in the given σ -algebra.

Let X and Y each have a given σ -algebra of subsets. A *measurable map* $g : X \rightarrow Y$ is a function such that for every measurable subset B of Y the inverse image $g^{-1}[B]$ is an measurable subset of X .

If $\Gamma \subset P(Y)$ is an arbitrary collection of subsets of Y , then there is a least σ -algebra \mathcal{F} of subsets of Y such that $\Gamma \subset \mathcal{F}$. This is the σ -algebra of subsets generated by Γ .

For example, say that $\Gamma = \{U, V\}$, where $U \subset Y$ and $V \subset Y$. Then the σ -algebra of sets generated by Γ can have up to 16 subsets in it.

Proposition 7.2 *Say that X and Y are measurable spaces. Suppose also that Γ generates the σ -algebra of sets \mathcal{F} for Y . Suppose that $f : X \rightarrow Y$ and that for every set B in the generating set Γ the inverse image $f^{-1}[B]$ is a measurable set. Then f is a measurable map.*

An example where this applies is when $Y = \mathbf{R}$ and the generating set Γ consists of all intervals $(a, +\infty)$. This generates the Borel σ -algebra of subsets of \mathbf{R} . Thus if the inverse image of each of these intervals is measurable, then the inverse image of every Borel subset is measurable.

In the case when the set is the real line \mathbf{R} , the σ -algebra will always be the Borel σ -algebra, unless otherwise explicitly indicated. Let Y be a measurable space, equipped with a σ -algebra \mathcal{F} of measurable sets. Then the space of all measurable functions $f : Y \rightarrow \mathbf{R}$ is a σ -algebra of measurable functions. (That is, it is a vector lattice and a monotone class and it includes the constant function.) Conversely, given a σ -algebra of measurable functions on Y , then

there is a σ -algebra of measurable sets, those whose indicator functions are measurable.

Say that X and Y are measurable spaces. It is not difficult to show that $g : X \rightarrow Y$ is a measurable map (in the sense that g^{-1} maps measurable subsets of Y to measurable subsets of X) if and only if for all measurable functions $f : Y \rightarrow \mathbf{R}$ the function $f \circ g : X \rightarrow \mathbf{R}$ is measurable.

It is worth remarking that if X is a measurable space and $Z \subset X$ is a subset, then there is a natural structure of measurable space on Z . This is the relative σ -algebra consisting of all the intersections of measurable subsets of X with Z .

7.3 The Borel σ -algebra

Given a space Y with topology \mathcal{T} , then there is an associated σ -algebra of subsets generated by \mathcal{T} . This is the Borel σ -algebra of subsets of Y . That is, the Borel σ -algebra of subsets is the smallest σ -algebra of which each open subset is an element. There is a corresponding Borel σ -algebra of real measurable functions, consisting of all functions from Y to \mathbf{R} that are measurable (in the sense that the inverse image of each Borel set is a Borel set).

Theorem 7.3 *Suppose that X and Y are topological spaces. Then X and Y may also be regarded as measurable spaces with the Borel σ -algebras of subsets. If $f : X \rightarrow Y$ is continuous, then $f : X \rightarrow Y$ is measurable.*

Proof: This follows immediately from the fact that the topology of Y generates the Borel σ -algebra of Y . Furthermore, it is enough to check measurability on a generating set. \square

It turns out that the interaction between topology and measurability is uneasy, in large part because of the fact that topology allows uncountable unions. However for metric spaces the situation is somewhat better.

Theorem 7.4 *Let Y be a metrizable topological space. Then the vector lattice of continuous real functions on Y generates the Borel σ -algebra of measurable real functions on Y .*

Proof: Let U be an open subset of Y . Let F be its complement. Then F is a closed set. Let $h : Y \rightarrow \mathbf{R}$ be the distance to F . Then h is continuous, h is zero on F , and $h > 0$ on the complement of F , that is, on U . The sequence $nh \wedge 1$ increases to the indicator function of U . Thus the indicator function of the open subset U is in the σ -algebra of functions generated by the continuous functions. This is enough to imply that the continuous functions generate the entire Borel σ -algebra of continuous functions. \square

The conclusion of the theorem does not hold for all topological spaces. In general the vector lattice of continuous real functions may generate a σ -algebra of functions that does not include all Borel functions.

7.4 Comparison of topologies

There are situations in analysis when there is more than one topology. A standard example is an infinite dimensional Hilbert space H . The strong topology is the topology consisting of all open sets in the usual metric sense. The weak topology is the topology generated by all unions of sets of the form $\{u + w \mid u \in U, w \in M^\perp\}$, where U is an open subset of a finite dimensional subspace M . Such a set is restricted in finitely many dimensions. It is not hard to see that every open set in the weak topology is an open set in the strong topology. The weak topology is the coarser or smaller topology.

Let $n \mapsto s_n$ be a sequence in the Hilbert space H . Then $s_n \rightarrow w$ weakly as $n \rightarrow \infty$ if and only if for each finite dimensional subspace M with associated orthogonal projection P_M the function $P_M s_n \rightarrow P_M w$ as $n \rightarrow \infty$. Since finite dimensional projections are given by finite sums involving inner products, this is the same as saying that for each vector v in H the numerical sequence $\langle v, s_n \rangle \rightarrow \langle v, w \rangle$ as $n \rightarrow \infty$.

It is clear that strong convergence of a sequence implies weak convergence of the sequence. This is because $|\langle v, s_n \rangle - \langle v, w \rangle| = |\langle v, s_n - w \rangle| \leq \|v\| \|s_n - w\|$, by the Schwarz inequality.

The converse is not true. For example, let $n \mapsto e_n$ be a countable orthonormal family. Then $e_n \rightarrow 0$ weakly as $n \rightarrow \infty$. This is because $\sum_n |\langle v, e_n \rangle|^2 \leq \|v\|^2$. It follows from convergence of this sum that $\langle v, e_n \rangle \rightarrow 0$ as $n \rightarrow \infty$. However $\|e_n - 0\| = 1$ for all n , so there is certainly not strong convergence to zero.

If $\mathcal{T}' \subset \mathcal{T}$, then we say that the topology on \mathcal{T}' is *coarser* (or smaller), while the topology \mathcal{T} is relatively *finer* (or bigger). Sometime the terms weak and strong are used, but this takes some care, as is shown by the following two propositions.

Proposition 7.5 *Let X have topology \mathcal{S} . If $f : X \rightarrow Y$ and $\mathcal{T}' \subset \mathcal{T}$ are topologies on Y , then $f : (X, \mathcal{S}) \rightarrow (Y, \mathcal{T})$ continuous implies $f : (X, \mathcal{S}) \rightarrow (Y, \mathcal{T}')$ continuous.*

The above proposition justifies the use of the word weak to describe the coarser topology on Y . Thus strong continuity implies weak continuity for maps into a space.

Proposition 7.6 *Let Z have topology \mathcal{U} . If $f : Y \rightarrow Z$ and $\mathcal{T}' \subset \mathcal{T}$ are topologies on Y , then $f : (Y, \mathcal{T}') \rightarrow (Z, \mathcal{U})$ continuous implies $f : (Y, \mathcal{T}) \rightarrow (Z, \mathcal{U})$ continuous.*

The above proposition gives a context when the coarser topology imposes the stronger continuity condition. If we want to continue to use the word weak to describe a coarser topology, then we need to recognize that weak continuity for maps from a space is a more restrictive condition.

As an example, consider again infinite dimensional Hilbert space H with the weak topology. Let $f : H \rightarrow \mathbf{R}$ be given by $f(u) = \|u\|^2$. Then f is continuous

when H is given the strong topology. However f is not continuous when H is given the weak topology. This may be seen by looking at a sequence e_n that is an orthonormal family. Then $e_n \rightarrow 0$ weakly, but $f(e_n) = 1$ for all n .

The finest possible topology on a set X is the *discrete topology*, for which every subset is open. The coarsest possible topology on a set is the *indiscrete topology*, for which only the empty subset \emptyset and X itself are open subsets.

7.5 Bases and subbases

A *base* for a topology \mathcal{T} is a collection Γ of open sets such that every open set V in \mathcal{T} is the union of some subcollection of Γ .

Let X be a topological space and let Γ be a collection of open subsets. Then Γ is a *subbase* if the collection $\tilde{\Gamma}$ of all intersections of finite subcollections of Γ is a base. Notice that according to the convention $\bigcap \emptyset = X$, the set X automatically belongs to $\tilde{\Gamma}$.

Theorem 7.7 *Let X be a set. Let Γ be a collection of subsets. Then there is a coarsest topology \mathcal{T} including Γ , and Γ is a subbase for \mathcal{T} .*

Proof: Let Γ be a collection of subsets of X . Let $\tilde{\Gamma}$ be the collection consisting of intersections of finite subsets of Γ . Let \mathcal{T} be the collection consisting of unions of subsets of $\tilde{\Gamma}$. The task is to show that \mathcal{T} is a topology.

It is clear that the union of a subcollection of \mathcal{T} is in \mathcal{T} . The problem is to show that the intersection of a finite subcollection Δ of \mathcal{T} is in \mathcal{T} . Each W in Δ is a union of a collection of sets $\mathcal{A}_W \subset \tilde{\Gamma}$. By the distributive law the intersection is

$$\bigcap_{W \in \Delta} \bigcup \mathcal{A}_W = \bigcup_s \bigcap_{W \in \Delta} s(W).$$

Here s is summed over all possible selection functions with the property that $s(W)$ is in \mathcal{A}_W for each W . Since each $s(W)$ is a finite intersection of sets in Γ , it follows that each $\bigcap_{W \in \Delta} s(W)$ is a finite intersection of sets in Γ . Thus the finite intersection is a union of such finite intersections. Thus it is in \mathcal{T} . \square

A *neighborhood base* for a point x in a topological space is a family Γ_x of open sets V with $x \in V$ such that for every open set U with $x \in U$ there is a V in Γ_x with $V \subset U$.

A *neighborhood subbase* for a point x in a topological space is a family Γ_x of open sets V with $x \in V$ such that for every open set U with $x \in U$ there is a V that is a finite intersection of sets in Γ_x with $V \subset U$.

A topological space is *first countable* if every point has a countable neighborhood base. This is the same as having a countable neighborhood subbase. A metric space is first countable. A neighborhood base at x consists of the open balls centered at x of radius $1/n$, for $n = 1, 2, 3, \dots$. So being close to x is determined by countably many conditions.

A topological space is *second countable* provided that it has a countable base. This is the same as having a countable subbase.

If X is a topological space and S is a subset, then S is *dense* in X if its closure is X .

A topological space X is *separable* provided that there is a countable subset S with closure $\bar{S} = X$. In other words, X is separable if it has a countable dense subset.

Theorem 7.8 *If X is second countable, then X is separable.*

Proof: Let Γ be a countable base for X . Let $\Gamma' = \Gamma \setminus \{\emptyset\}$. Then Γ' consists of non-empty sets. For each U in Γ' choose x in U . Let S be the set of all such x . Let $V = X \setminus \bar{S}$. Since V is open, it is the union of those of its subsets that belong to Γ . Either there are no such subsets, or there is only the empty set. In either case, it must be that $V = \emptyset$. This proves that $\bar{S} = X$. \square

It is also not very difficult to prove that a separable metric space is second countable. It is not true in general that a separable topological space is second countable.

7.6 Compact spaces

A topological space K is *compact* if whenever Γ is a collection of open sets with $K = \bigcup \Gamma$, then there is a finite subcollection $\Gamma_0 \subset \Gamma$ with $K = \bigcup \Gamma_0$. This can be summarized in a slogan: Every open cover has a finite subcover.

Sometimes one wants to apply this definition to a subset K of a topological space X . Then it is customary to say that K is compact if and only if whenever Γ is a collection of open subsets of X with $K \subset \bigcup \Gamma$, then there is a finite subcollection $\Gamma_0 \subset \Gamma$ with $K \subset \bigcup \Gamma_0$. Again: Every open cover has a finite subcover. However this is just the same as saying that K itself is compact with the relative topology.

There is a dual formulation in terms of closed subsets. A topological space X is compact if whenever Γ is a collection of closed sets with $\bigcap \Gamma = \emptyset$, then there is a finite subcollection $\Gamma_0 \subset \Gamma$ with $\bigcap \Gamma_0 = \emptyset$.

A collection of sets Γ has the *finite intersection property* provided that for every finite subcollection $\Gamma_0 \subset \Gamma$ we have $\bigcap \Gamma_0 \neq \emptyset$.

Proposition 7.9 *A topological space is compact if and only if every collection Γ of closed subsets with the finite intersection property has $\bigcap \Gamma \neq \emptyset$.*

Again there could be a slogan: Every collection of closed sets with the finite intersection property has a common point.

Corollary 7.10 *A topological space is compact if and only if for every collection Γ of subsets with the finite intersection property there is a point x that is in the closure of each of the sets in Γ .*

Perhaps the slogan could be: For every collection of sets with the finite intersection property there exists a point near each set.

Proposition 7.11 *If K is compact, $F \subset K$, and F is closed, then F is compact.*

Proof: Let Γ be a collection of closed subsets of F with the finite intersection property. Since F is closed, each set in Γ is also a closed subset of K . Since K is compact, there is a point p in each set in Γ . This is enough to prove that F is compact. \square

Theorem 7.12 *Let $f : K \rightarrow L$ be a continuous surjection from K onto L . If K is compact, then L is compact.*

Proof: Let Δ be an open cover of L . Then the inverse images under f of the sets in Δ form an open cover of K . However K is compact. Therefore there exists a finite subset Δ_0 of Δ such that the inverse images of the sets in Δ_0 form an open cover of K . Since f is a surjection, every point y in L is the image of a point x in K . There is an open set V in Δ_0 such that that x is in the inverse image of V . It follows that y is in V . This proves that Δ_0 is an open cover of L . It follows that L is compact. \square

A topological space is *Hausdorff* provided that for each pair of points x, y in the space there are open subsets U, V with $x \in U$, $y \in V$, and $U \cap V = \emptyset$.

Proposition 7.13 *Each compact subset K of a Hausdorff space is closed.*

Proof: Let X be a Hausdorff space and $K \subset X$ a compact subset.

Fix $y \notin K$. For each $x \in K$ choose U_x and V_x with $x \in U_x$ and $y \in V_x$ and $U_x \cap V_x = \emptyset$. The union of the U_x for x in K includes K , so there is a finite subset S of K such that the union of the U_x for x in S included K . Let V be the intersection of the V_x for x in S . Then V is open and $y \in V$ and $V \cap K = \emptyset$.

Choose for every y in $X \setminus K$ an open set V_y with $y \in V_y$ and $V_y \cap K = \emptyset$. Then $X \setminus K$ is the union of the V_y with y in $X \setminus K$. This proves that $X \setminus K$ is open, and so K is closed. \square

7.7 The one point compactification

Theorem 7.14 (one point compactification) *Let X be a topological space that is not compact. Then there exists a topological space X^* with one extra point that is compact and such that X is a subspace of X^* with the induced topology.*

Proof: Let ∞ be a point that is not in X , and let $X^* = X \cup \{\infty\}$. The topology for X^* is defined as follows. There are two kinds of open sets of X^* . If $\infty \notin U$, then U is open if and only if U is an open set in the topology of X . If $\infty \in U$, then U is open if and only if U is the complement of a closed compact subset K of X . It is clear that the topology of X is the relative topology as a subspace of X^* .

Consider an open cover of X^* . This is a collection of open subsets of X^* whose union is X^* , so there must be at least one open subset that is the complement of a compact closed subset K of X . The union of the remaining open

subsets in the cover includes K . These open sets can be of two kinds. Some of them may be open subsets of X . The other are complements of closed subsets of X , so their intersections with X are open subsets of X . These subsets provide an open cover of K , so they have a finite subcover of K . This shows that the one open set whose complement is K together with the remaining finite collection of open sets that cover K form an open cover of X^* . This proves that X^* is compact. \square

A topological space X is *locally compact* if and only if for each point p in X there exists an open set U and a compact set K with $p \in U \subset K$.

Theorem 7.15 (one point Hausdorff compactification) *Let X be a topological space. Then its one point compactification X^* is Hausdorff if and only if X is both locally compact and Hausdorff.*

Proof: Here is a sketch of the fact that X locally compact Hausdorff implies X^* Hausdorff. The Hausdorff property says that each pair of distinct points are separated by open sets. The separation is clear for two points that are subsets of X . The interesting case is when one point is $p \in X$ and the other point is ∞ . Then by local compactness there exists an open subset U of X such that $p \in U \subset K$, where K is a compact subset of X . Since X is Hausdorff, K is also closed. Let V be the complement of K in X^* . Then $\infty \in V$ and $p \in U$, both U and V are open in X^* , and the two open sets are disjoint. Thus p and ∞ are separated by open sets. \square

Examples:

1. \mathbf{R}^n is locally compact and Hausdorff. Its one point compactification is homeomorphic to a sphere S_n .
2. Consider an infinite dimensional real Hilbert space H , for example ℓ^2 . It is a metric space and so is Hausdorff. However it is not locally compact. In fact, an open ball is not totally bounded, so it cannot be a subset of a compact set. The one point compactification of H is not even Hausdorff.
3. What if we give H the weak topology? As we shall see, then each closed ball is compact. But the space is still not locally compact, since there are no non-empty weakly open sets inside the closed unit ball. The non-empty weakly open sets are all unbounded.

The case of \mathbf{R} and \mathbf{C} are particularly interesting. From the point of analytic function theory, the one point compactification is quite natural. The one point compactification of \mathbf{R} is a circle, and the one point compactification of \mathbf{C} is the Riemann sphere.

On the other hand, \mathbf{R} has an order structure, and in this context it is more natural to look at a two point compactification $[-\infty, +\infty]$. This of course is not homeomorphic to a circle, but instead to an interval such as $[0, 1]$.

7.8 Nets

The following topic is optional; it is intended only to indicate that nets are a generalization of sequences. In a metric space the notion of sequence is important, because most topological properties may be characterized in terms of convergence of sequences. In more general spaces sequences are not enough to characterize convergence. However the more general notion of net does the job.

A *directed set* is an ordered set I with the property that every finite non-empty subset has an upper bound. For general topological spaces it is important that I is not required to be a countable set.

A *net* in X is a function $w : I \rightarrow X$. If X is a topological space, then a net w *converges* to x provided that for every open set U with $x \in U$ there is a j in I such that for all k with $j \leq k$ we have $w_k \in U$.

Examples:

1. A sequence in X is a net in X . This is the special case when the directed set is the set of natural numbers.
2. Let S be a set. Let I consist of all finite subsets of S . Notice that if S is uncountable, then the index set I is also uncountable. For H and H' in I we write $H \leq H'$ provided that $H \subset H'$. Fix a function $f : S \rightarrow \mathbf{R}$. Define the net $H \mapsto \sum_{s \in H} f(s)$ with values in \mathbf{R} . If this net converges to a limit, then this limit is a number $\sum f$ that deserves to be called the unordered sum of f . The set of f for which such an unordered sum exists is of course just $\ell^1(S)$. It turns out that this example is not so interesting after all, since each f in $\ell^1(S)$ vanishes outside of some countable subset.
3. Here is an example that shows how one can construct a directed set to describe convergence in a general topological setting. Let x be a point in X , and define the directed set I to consist of all open sets with $x \in U$. Let $U \leq U'$ mean that $U' \subset U$. Then I is a directed set, since the intersection of finitely many such open sets is open. We shall see in the proof of the next two theorems that this kind of directed set is a rather natural domain for a net.

Theorem 7.16 *Let E be a subset of X . A point x is in the closure \bar{E} if and only if there is a net w with values in E that converges to x .*

Proof: First note that the complement of \bar{E} is the largest open set disjoint from E . It follows that $x \notin \bar{E}$ is equivalent to $\exists U (x \in U \wedge U \cap E = \emptyset)$. Here U ranges over open subsets. As a consequence, $x \in \bar{E}$ is equivalent to $\forall U (x \in U \Rightarrow U \cap E \neq \emptyset)$.

Suppose w is a net in E that converges to x . Let U be an open set in X such that $x \in U$. Then there exists a j so that $w_j \in U$. Hence $U \cap E \neq \emptyset$. Since U is arbitrary, it follows that $x \in \bar{E}$.

Suppose on the other hand that x is in \bar{E} . Then for every open set U with $x \in U$ we have $U \cap E \neq \emptyset$. By the axiom of choice there is a point w_U with $w_U \in E$ and w_U in U .

Let I consist of the open sets U with $x \in U$. Let $U \leq U'$ provided that $U' \subset U$. Then I is a directed set, since the intersection of finitely many such open sets is an open set. Thus $U \mapsto w_U$ is a net in E that converges to x . \square

Theorem 7.17 *A function $f : X \rightarrow Y$ is continuous if and only if it maps convergent nets into convergent nets.*

Proof: Suppose f is continuous. Let w be a net in X that converges to x . Let V be an open set in Y such that $f(x) \in V$. Let $U = f^{-1}[V]$. Then $x \in U$, so there exists a j such that $j \leq k$ implies $w_k \in U$. Hence $f(w_k) \in V$. This shows that the net $j \mapsto f(w_j)$ converges to $f(x)$.

Suppose on the other hand that f maps convergent nets to convergent nets. Suppose that f is not continuous. Then there is an open set V in Y such that $f^{-1}[V]$ is not open.

Next, notice that a set G is open if and only if $\forall x x \in G \Rightarrow (\exists U(x \in U \wedge U \subset G))$. Here U ranges over open subsets. This is simply because the union of open sets is always open. Hence a set G is not open if and only if $\exists x x \in G \wedge (\forall U(x \in U \Rightarrow U \setminus G \neq \emptyset))$.

Apply this to the case $G = f^{-1}[V]$. Then there exists an x with $f(x) \in V$ but with the property that for every open set U with $x \in U$ there exists a point $w_U \in U$ with $f(w_U) \notin V$. The existence of this function $U \mapsto w_U$ is guaranteed by the axiom of choice.

Let I consist of the open sets U with $x \in U$. Let $U \leq U'$ provided that $U' \subset U$. Then I is a directed set, since the intersection of finitely many such open sets is an open set. Thus $U \mapsto w_U$ is a net in X that converges to x . However $U \mapsto f(w_U)$ does not converge to $f(x)$. This is a contradiction. Thus f must be continuous. \square

The net language also sheds light on the Hausdorff separation property. It may be shown that a topological space is Hausdorff if and only if every net converges to at most one point.

Nets are convenient for dealing with more general topological spaces. However, in an effort to keep the number of new concepts within bounds this account will not make much further use of nets.

7.9 Problems

1. Let $X = \mathbf{R}$. Show that the sets $(a, +\infty)$ for $a \in \mathbf{R}$, together with the empty set and the whole space, form a topology.
2. Give an example of a continuous function $f : \mathbf{R} \rightarrow X$, where \mathbf{R} has the usual metric topology and X has the topology of the preceding problem. Make the example such that f is not a continuous function in the usual sense.

3. What are the compact subsets of X in the example of the first problem?
4. Let $X = \mathbf{R}$. Show that the sets $(a, +\infty)$ for $a \in \mathbf{R}$ together with $(-\infty, b)$ for $b \in \mathbf{R}$ are not a base for a topology. They are a subbase for a topology. Describe this topology.
5. Let $s : \mathbf{N} \rightarrow X$ be a sequence in a metric space X . Let $T_n = \{s_k \mid k \geq n\}$. Show that the sets T_n for $n \in \mathbf{N}$ have the finite intersection property. Show that x is in the closure of every set T_n if and only if there is a subsequence that converges to x .

Chapter 8

Product and weak* topologies

8.1 Introduction

The following sections deal with important compactness theorems. The proofs of these theorems make use of Zorn's lemma. In the appendices to this chapter it is shown that the axiom of choice implies Zorn's lemma. It is quite easy to show that Zorn's lemma implies the axiom of choice.

Consider a non-empty partially ordered set. Suppose that every non-empty totally ordered subset has an upper bound. Zorn's lemma is the assertion that the set must have a maximal element.

In a sense, Zorn's lemma is an obvious result. Start at some element of the partially ordered set. Take a strictly larger element, then another, then another, and so on. Of course it may be impossible to go on, in which case one already has a maximal element. Otherwise one can go through an infinite sequence of elements. These are totally ordered, so there is an upper bound. Take a strictly larger element, then another, then another, and so on. Again this may generate a continuation of the totally ordered subset, so again there is an upper bound. Continue in this way infinitely many times, if necessary. Then there is again an upper bound. This process is continued as many times as necessary. Eventually one runs out of set. Either one has reached an element from a previous element and there is not a larger element after that. In that case the element that was reached is maximal. Or one runs at some stage through an infinite sequence, and this has an upper bound, and there is nothing larger than this upper bound. In this case the upper bound is maximal.

Notice that this argument involves an incredible number of arbitrary choices. But the basic idea is simple: construct a generalized orbit that is totally ordered. Keep the construction going until a maximal element is reached, either as the result of a previous point in the orbit, or as the result of an previous sequence in the orbit.

8.2 The Tychonoff product theorem

Let $A = \prod_{t \in T} A_t$ be a product space. An element x of A is a function from the index set T with the property that for all $t \in T$ we have $x_t \in A_t$.

Suppose that each A_t is a topological space. For each t there is projection $p_t : A \rightarrow A_t$ defined by $p_t(x) = x_t$. The *product topology* is the coarsest topology on A such that each individual projection p_t is continuous. Thus if $U \subset A_t$ is open, then $p_t^{-1}[U]$ is an open set in A that has its t component restricted. Furthermore, a finite intersection of such sets is open. So there are open sets that are restricted in finitely many components.

Write the projection of x on the t coordinate as the value $x(t)$ of the function x . A net $j \mapsto w_j$ in $A = \prod_{t \in T} A_t$ converges to a point x in A if and only if for each t the net $j \mapsto w_j(t)$ converges to $x(t)$. For this reason the product topology is also called the topology of pointwise convergence.

Theorem 8.1 (Tychonoff) *The product of a family of compact spaces is compact.*

Before starting the proof, it is worth looking at an attempt at a proof that does not work. Suppose that for each $t \in T$ the space A_t is compact. Let Γ be a collection of closed subsets of A with the finite intersection property. We want to show that $\bigcap \Gamma \neq \emptyset$. This will prove that A is compact.

Fix t . Let Γ_t be the collection of all projected subsets $F_t = p_t[F]$ for $F \in \Gamma$. Then Γ_t has the finite intersection property. Since A_t is compact, there exists an element that belongs to the closure of each F_t in Γ_t . Choose such an element $x_t \in A_t$ arbitrarily.

Let $x \in A$ be the vector which has components x_t . If we could show that x is in each of the F in Γ , then this would complete the proof. However this where the attempt fails; there is no guarantee that this is so.

If fact, we could take a simple example in the unit square where this does not work. Let the set Γ consist of the single set $F = \{(0, 1), (1, 0)\}$. Then the projection on the first axis is the set $\{0, 1\}$ and the second projection is also $\{0, 1\}$. If we take $x_1 = 0$ and $x_2 = 0$, then $x = (0, 0)$ is far from belonging to F . The trouble is that the projections of a set do not do enough to specify the set. The solution is to specify the point in the product space more closely by taking a larger collection of sets with the finite intersection property. For instance, in the example one could take Γ' to consist of the set F together with the smaller set $\{(0, 1)\}$. Then the projected sets on the first axis have only 0 in their intersection, and the projected sets on the second axis have only 1 in their intersection. So from these one can reconstruct that point $(0, 1)$ in the product space that belongs to all the sets in Γ' and hence of Γ .

Notice that this enlarged collection of sets is somewhat arbitrary; one could have made another choice and gotten another point in the product space. However the goal is to single out a point, and the way to do this is to make a maximal specification of the point. One means to accomplish this is to take a maximal collection of sets with the finite intersection property. For instance, one could

take all subsets of which $(0, 1)$ is an element. This is an inefficient but sure way of specifying the point $(0, 1)$.

Proof: Suppose that for each $t \in T$ the space A_t is compact. Let Γ be a collection of closed subsets of A with the finite intersection property. We want to show that $\bigcap \Gamma \neq \emptyset$. This will prove that A is compact.

Consider all collections of sets with the finite intersection property that include Γ . By Zorn's lemma, there is a maximal such collection Γ' .

Fix t . Let Γ'_t be the set of all projected subsets $F_t = p_t[F]$ for $F \in \Gamma$. Then Γ'_t has the finite intersection property. Since A_t is compact, the intersection of the closures of the F_t in Γ'_t is non-empty. Choose an element x_t in the closure of each F_t for F in Γ' .

Let x be the vector which has components x_t . Let U be an open set with $x \in U$. Then there is a finite subset $T_0 \subset T$ and an open set $U_t \subset A_t$ for each t in T_0 such that the intersection of the sets $p_t^{-1}[U_t]$ for $t \in T_0$ is an open subset of U .

Consider t in T_0 . It is clear that $x_t \in U_t$. Since x_t is in the closure of each F_t for each F in Γ' , it follows that $U_t \cap F_t \neq \emptyset$ for each F in Γ' . Thus $p_t^{-1}[U_t] \cap F \neq \emptyset$ for each F in Γ' . Since Γ' is maximal with respect to the finite intersection property, it follows that $p_t^{-1}[U_t]$ is in Γ' .

Now use the fact that Γ' has the finite intersection property. Consider F in Γ' . Since each of the $p_t^{-1}[U_t]$ for t in the finite set T_0 is in Γ' , it follows that the intersection of the $p_t^{-1}[U_t]$ for t in T_0 with F is non-empty.

This shows that U has non-empty intersection with each element F of Γ' . Since U is arbitrary, this proves that x is in the closure of each element F of Γ' . In particular, x is in the closure of each element F of Γ . Since Γ consists of closed sets, x is in each element F of Γ . \square

8.3 Banach spaces and dual Banach spaces

This section is a quick review of the most commonly encountered Banach spaces of functions and of their dual spaces. Let E be a Banach space. Then its dual space E^* consists of the continuous linear functions from E to the field of scalars (real or complex). It is also a Banach space. There is a natural injection from E to E^{**} . The Banach space E is said to be *reflexive* if this is a bijection.

Let X be a set and \mathcal{F} be a σ -algebra, so that X is a measurable space. Fix a measure μ . The first examples consist of the Banach spaces $E = L^p(X, \mathcal{F}, \mu)$ for $1 \leq p < \infty$. (In the case $p = 1$ we require that μ be a σ -finite measure.) Then the dual space E^* may be identified with $L^q(X, \mathcal{F}, \mu)$, where $1 < q \leq \infty$. Here $1/p + 1/q = 1$. If u is in E and f is in E^* , the value of f on u is the integral $\mu(fu)$ of the product of the two functions. If $1 < p < \infty$, then the Banach space $L^p(X, \mathcal{F}, \mu)$ is reflexive. Thus if $1 < q \leq \infty$ the Banach space $L^q(X, \mathcal{F}, \mu)$ is a dual space. In general $L^1(X, \mathcal{F}, \mu)$ is not the dual of another Banach space. This is because the dual of $L^\infty(X, \mathcal{F}, \mu)$ is considerably larger than $L^1(X, \mathcal{F}, \mu)$. These facts are discussed in standard references, such as R. M. Dudley, *Real Analysis and Probability* (Cambridge University Press, Cambridge, 2002).

The following examples introduce topology in order to get a Banach space E with a dual space E^* that can play the role of an enlargement of L^1 . The space E will be a space of continuous functions, while the space E^* is identified with a space of finite signed measures. In order to avoid some measure theoretic technicalities, we shall deal only with continuous functions defined on metric spaces.

Let X be a compact metric space. Then $E = C(X)$ is a real Banach space. The norm of a function in $C(X)$ is the maximum value of its absolute value. Thus convergence in $C(X)$ is uniform convergence on X .

The space $C(X)$ of continuous real functions generates a σ -algebra of functions. These are called Borel functions, and there is a corresponding σ -algebra of Borel sets. The integrals or measures under consideration are defined on Borel functions or on Borel sets.

A *finite signed measure* is an object that is the difference of two finite measures that live on disjoint sets. That is, there are finite measures μ_+ and μ_- and measurable sets B_+ and B_- such that $\mu_+(B_-) = 0$ and $\mu_-(B_+) = 0$. The signed measure is then $\mu = \mu_+ - \mu_-$. Sometimes, in the context of signed measures, a measure of the usual kind is called a positive measure. Thus a finite signed measure is the difference of two finite positive measures.

There is a natural norm for finite signed measures. If μ is a finite signed measure, then $\|\mu\| = \mu_+(X) + \mu_-(X)$ is the norm.

A Riesz representation theorem is a theorem that identifies the dual of a Banach space of functions. For instance, the theorem that identifies the dual of L^p as L^q for $1 \leq p < \infty$ and $1/p + 1/q = 1$ is a Riesz representation theorem. An even more important Riesz representation theorem is the following.

Proposition 8.2 *Let X be a compact metric space. Let $E = C(X)$ be space of continuous real functions on X . Then the dual space E^* may be identified with the space of finite signed Borel measures on X . That is, each continuous real linear function on $C(X)$ is of the form $f \mapsto \mu(f)$ for a unique finite signed Borel measure on the compact space X .*

This result gives an example of a Banach space that is far from being reflexive. That is, the dual of the space E^* of finite signed measures is much larger than the original space $E = C(X)$. In fact, consider an arbitrary bounded Borel function f . Then the map $\mu \mapsto \mu(f)$ is continuous on E^* and hence is in E^{**} . However it is not necessarily given by an element of E .

The results have a useful generalization. Let X be a separable locally compact metric space. Then $C_0(X)$ is a real Banach space. Here f is in $C_0(X)$ provided that for every $\epsilon > 0$ there is a compact subset K of X such that $|f| < \epsilon$ outside of K . Such an f is said to vanish at infinity. The space $C_0(X)$ of continuous real functions that vanish at infinity generates the σ -algebra of Borel functions.

Proposition 8.3 *Let X be a σ -compact locally compact metric space. Let $E = C_0(X)$ be the space of continuous real functions on X that vanish at infinity.*

Then the dual space E^* may be identified with the space of finite signed Borel measures on X . That is, each continuous real linear function on $C_0(X)$ is of the form $f \mapsto \mu(f)$ for a unique finite signed Borel measure on the locally compact space X .

This proposition is only a slight variant on the preceding proposition. Let X^* be the one point compactification of X . Then the space $C_0(X)$ may be thought of as the functions in $C(X^*)$ that vanish at the point ∞ . Similarly, the finite signed measures μ on X may be identified with the measures on X^* that assign mass zero to the the set $\{\infty\}$.

8.4 Weak* topologies on dual Banach spaces

The *weak topology* on E is the coarsest topology such that every element of E^* is continuous. As a special case, the weak topology on E^* is the coarsest topology such that every element of E^{**} is continuous. The sets $W(f, V) = \{u \in E \mid f(u) \in V\}$, where f is in E^* and V is an open set of scalars, form a subbase for the weak topology of E .

The *weak* topology* on E^* is the coarsest topology such that every element of E defines a continuous function on E^* . This is the topology of pointwise convergence for the functions in E^* . The sets $W(u, V) = \{f \in E^* \mid f(u) \in V\}$, where u is in E and V is an open set of scalars, form a subbase for the weak topology* of E .

Proposition 8.4 *Let E be a Banach space, and let E^* be its dual space. The weak* topology on E^* is coarser than the weak topology on E^* . If E is reflexive, then the weak* topology on E^* is the same as the weak topology on E^* .*

Proof: Since each element of E defines an element of E^{**} , the weak* topology is the coarsest topology that makes all these elements of E^{**} that come from E continuous. The weak topology is defined by requiring that all the elements of E^{**} are continuous. Since more functions have to be continuous, the weak topology is a finer topology. \square

Examples:

1. Fix a σ -finite measure μ . Let $E = L^1$ be the corresponding space of real integrable functions. Then $E^* = L^\infty$. A sequence f_n in L^∞ converges weak* to f if for every u in L^1 the integrals $\mu(f_n u) \rightarrow \mu(fu)$.
2. Fix a measure μ . Let $E = L^p$ with $1 < p < \infty$. Then $E^* = L^q$ with $1 < q < \infty$. Here $1/p + 1/q = 1$. A sequence f_n in L^q converges weak* to f if for every u in L^p the integrals $\mu(f_n u) \rightarrow \mu(fu)$. Since L^p for $1 < p < \infty$ is reflexive, this is the same as weak convergence in L^q .
3. Fix a measure μ . Let $E = L^\infty$. Then E^* is an unpleasant space that includes L^1 but also has a huge number of unpleasant measure-like objects

in it. Notice that there is no weak* topology on L^1 , since it is not the dual of another Banach space.

4. Consider a compact metric space X . Let $E = C(X)$, the space of all continuous real functions on X . The norm on X is the supremum norm that describes uniform convergence. Then E^* consists of signed measures. These are of the form $\mu = \mu_+ - \mu_-$, where μ_+ and μ_- are finite measures. These are the Radon measures that will be described in more detail in a following chapter. The norm of μ is $\mu_+(X) + \mu_-(X)$. A sequence $\mu_n \rightarrow \mu$ in the weak* topology provided that $\mu_n(u) \rightarrow \mu(u)$ for each continuous function u .
5. Consider a locally compact metric space X . Let $E = C_0(X)$, the space of all continuous real functions on X that vanish at infinity. Then E^* again consists of signed measures. In fact, we can think of E as the space of all continuous functions on the one point compactification of X that vanish at the point ∞ . Then the measures in E^* are those measures on the compactification that assign measure zero to the set $\{\infty\}$. A sequence $\mu_n \rightarrow \mu$ in the weak* topology provided that $\mu_n(u) \rightarrow \mu(u)$ for each continuous function u that vanishes at infinity.

Such examples give an idea of the significance of the weak* topology. The idea is that for μ_n to be close to μ in this sense, it is enough that for each observable quantity u the numbers $\mu_n(u)$ get close to $\mu(u)$. The observation is a kind of blurred observation that does not make too many fine distinctions. In the case of measures it is the requirement of continuity that provides the blurring. This allows measures that are absolutely continuous with respect to Lebesgue measure to approach a discrete measure, and it also allows measures that are discrete to approach a measure that is absolutely continuous with respect to Lebesgue measure.

Examples:

1. The measures with density $n1_{[0,1/n]}$ approaches the point mass δ_0 . This is absolutely continuous to singular.
2. The singular measures $\frac{1}{n} \sum_{j=1}^n \delta_{j/n}$ approach the measure with density $1_{[0,1]}$. This is singular to absolutely continuous.

8.5 The Alaoglu theorem

Theorem 8.5 (Alaoglu) *Let E be a Banach space. Let B^* be the closed unit ball in the dual space E^* . Then B^* is compact with respect to the weak* topology.*

Proof: This theorem applies to either a real or a complex Banach space. Define for each u in E the set I_u of all scalars a such that $|a| \leq \|u\|$. This is an closed interval in the real case or a closed disk in the complex case. In either

case each I_u is a compact space. Let $P = \prod_{u \in E} I_u$. By the Tychonoff product theorem, this product space is compact. An element f of P is a scalar function on E with the property that $|f(u)| \leq \|u\|$ for all u in E . The product space topology on P is just the topology of pointwise convergence for such functions.

The unit ball B^* in the dual space E^* consists of all elements of P that are linear. The topology on B^* inherited from P is the topology of pointwise convergence. The topology on B inherited from the weak* topology on E^* is also the topology of pointwise convergence. So the task is to show that B^* is compact in this topology. For this, it suffices to show that B^* is a closed subset of P .

For each $u \in E$ the mapping $f \mapsto f(u)$ is continuous on P . Therefore, for each pair of scalars a, b and vectors u, v the mapping $f \mapsto f(au + bv) - af(u) - bf(v)$ is continuous on P . It follows that the set of all f with $f(au + bv) - af(u) - bf(v) = 0$ is a closed subset of P . The intersection of these closed sets for all a, b and all u, v is also a closed subset of P . However this intersection is just B^* . Since B^* is a closed subset of a compact space P , it must be compact. \square

If E is an infinite dimensional Banach space, then its dual space E^* with the weak* topology is not metrizable. This fact should be contrasted with the following important result.

Theorem 8.6 *If E is a separable Banach space, then the unit ball B^* in the dual space with the weak* topology is metrizable.*

Proof: Suppose E is separable. Let S be a countable dense subset of the unit ball B of E . Let I be the closed unit ball in the field of scalars. For each f in B^* there is a corresponding element $u \mapsto f(u)$ in I^S . Denote this element by $j(f)$. Thus $j(f)$ is just the restriction of f to S . From the fact that S is dense in B it is easy to see $j : B^* \rightarrow I^S$ is injective. Give I^S the product topology. Since for each u in S the map $f \mapsto f(u)$ is continuous, it follows that j is continuous.

The remaining task is to prove that the inverse $j(f) \mapsto f$ is continuous. To do this, consider a closed subset F of B^* . Since it is a closed subset of a compact space, it is compact. Since j is continuous, $j[F]$ is a compact subset of I^S . However a compact subset of Hausdorff space is closed. So $j[F]$ is closed. This says that the inverse image of each closed set under the inverse of j is a closed set. It follows that the inverse of j is continuous.

This proves that j is an embedding of B^* into I^S . However since I^S is a countable product of metric spaces, the product topology on this space is given by a metric. Such a metric induces a metric on B^* . This can be taken to have the explicit form $d(f, f') = \sum_{n=1}^{\infty} |f(s_n) - f'(s_n)|/2^n$. \square

Examples:

1. Let $E = L^1$, so $E^* = L^\infty$. The unit ball consists of all functions with absolute value essentially bounded by one. It is possible that a sequence of positive functions with essential bound one converges weak* to zero. For example, on the line the sequence of functions f_n that are the indicator

functions of intervals $[n, n + 1]$ converge to zero. This is because for each fixed u in L^1 we have $\mu(f_n u) \rightarrow 0$, by the dominated convergence theorem. Such an example is even possible when the measure space is finite. Here the example would be given by the indicator functions of the sets $[0, 1/n]$. Yet another example is convergence to zero by oscillation. Consider the functions $\cos(nx)$ on the interval $[0, 2\pi]$. These converge weakly to zero in the weak* topology of L^∞ , by the Riemann-Lebesgue lemma.

2. Let $E = L^p$ with $1 < p < \infty$, so $E^* = L^q$ with $1 < q < \infty$. Here $1/p + 1/q = 1$. It is possible that a sequence of positive functions with L^q norm equal to one converges weak* to zero. The sequence of indicator functions of the sets $[n, n + 1]$ provide the most obvious example. In the case when the measure space is finite, an example is where f_n is the $n^{\frac{1}{q}}$ times the indicator function of the set $[0, 1/n]$. This example is less obvious. It is clear that for u bounded we have $|\mu(f_n u)| \leq n^{\frac{1}{q}} \|u\|_\infty / n \rightarrow 0$. Since bounded functions are dense in L^p and we have a bound on the L^q norm of the f_n , it follows that we have $\mu(f_n u) \rightarrow 0$ for each u in L^p . There are yet more examples, such as convergence by oscillation.
3. Consider the space L^1 with the weak topology. The closed unit ball is not compact. In fact, let g_n be n times the indicator function of $[0, 1/n]$. If, for instance, w is a bounded continuous function, then $\mu(w g_n) \rightarrow w(0)$. This indicates that g_n is converging to something that acts like a point measure at the origin. This is no longer in the space L^1 . A sequence of densities with bounded total mass can converge to something that is not a density. In physical terms: conservation of mass is not enough to make something to remain a function. (This should be contrasted with the L^p with $p > 1$ case above. For L^2 this says that conservation of energy is enough to maintain the constraint of being a function.)
4. Consider a compact metric space X . Let $E = C(X)$. Then the signed measures in the unit ball of the dual space form a weak* compact set. Such a measure μ is an ordinary positive measure provided that for each positive continuous function $u \geq 0$ the value $\mu(u) \geq 0$. From this it is clear that the positive measures of total mass at most one form a weak* closed subset. (This is because $\mu \mapsto \mu(f)$ is continuous, so the inverse image of the closed set $[0, +\infty)$ is closed.) Therefore they are a compact subsets. Furthermore, the probability measures form a closed subset of these, since the requirement for a positive measure to be a probability measure is that $\mu(1) = 1$. (This is because $\mu \mapsto \mu(1)$ is continuous, so the image of the closed set $\{1\}$ is closed.) The conclusion is that the space of probability measures on a compact metric space is weak* compact. There is no way to lose probability from a compact space! Notice that this example explains what is going on in the preceding example. Consider the sequence μ_n of probability measures that have density with respect to Lebesgue measure that is n times the indicator function of $[0, 1/n]$. This sequence converges to the point measure δ_0 at the origin, which is still a probability measure.

5. Consider a separable locally compact metric space X . Let $E = C_0(X)$. Again the signed measures in the unit ball of E^* form a compact set. The positive measures of total mass at most one again form a compact subset. However the function 1 does not belong to the space $C_0(X)$. So we cannot conclude that the set of probability measures is closed or compact. In fact, we can take the measures with density given by the indicator function of $[n, n + 1]$. These probability measures converge weak* to zero. This seems mysterious until we choose to look instead at the one point compactification of X . Then it is seen that the probability has all gone to the point at infinity.

8.6 Appendix: The Bourbaki fixed point theorem

The following is an optional topic. It consists a result that does not depend on the axiom of choice. However it together with the axiom of choice will lead to a proof of Zorn's lemma.

Theorem 8.7 (Bourbaki) *Let A be a non-empty ordered set. Suppose that every non-empty totally ordered subset has a supremum. Let $f : A \rightarrow A$ be a function such that for all x in A we have $x \leq f(x)$. Then f has a fixed point.*

Proof: The function $f : A \rightarrow A$ is a dynamical system. Since A is non-empty, we can choose a in A as a starting point. Let $B \subset A$. We say that B is admissible if $a \in B$, $f[B] \subset B$, and whenever $T \subset B$ is totally ordered, then $\sup T \in B$. Thus f restricted to B is itself a dynamical system.

Let M be the intersection of all admissible subsets of A . It is not difficult to show that M is itself an admissible subset and that a is the least element of M . Thus f restricted to M is a dynamical system. We want to show that there is a sense in which M is a kind of generalized orbit of f starting at a . More precisely, we want to show that M is totally ordered. The rest of the proof is to establish that this is so. Then the fixed point will just be the supremum of this totally ordered set. In other word, the system starts at a and follows this generalized orbit until forced to stop.

Let $E \subset M$ be the set of points $c \in M$ such that for all x in M , the condition $x < c$ implies $f(x) \leq c$. Such a point c will be called a "choke point," for a reason that will be soon apparent.

Let $c \in E$. Let $M_c \subset M$ be the set of points x in M such that $x \leq c$ or $f(c) \leq x$. These are the points that can be compared unfavorably to c or favorably to $f(c)$.

First we check that M_c is admissible. First, it is clear that a is in M_c . Second, f maps the set of elements $x \leq c$ in M_c to M_c (since $x < c$ implies $f(x) \leq c$ and $x = c$ implies $f(c) \leq f(x)$), and f maps the set of elements x in M_c with $f(c) \leq x$ to itself. Third, if $T \subset M_c$ is totally ordered with supremum b , then either $x \leq c$ for all $x \in T$ implies $b \leq c$ (since b is the least upper bound),

or $f(c) \leq x$ for some x in T implies $f(c) \leq b$ (since b is an upper bound). Thus b is also in M_c .

So $M \subset M_c$, in fact they are equal. This works for arbitrary c in E . The conclusion is that c in E , x in M implies $x \leq c$ or $f(c) \leq x$. Thus each choke point c of M splits M into a part unfavorable to c or favorable to $f(c)$. This justifies the term “choke point.”

Next we check that the set of all choke points E is admissible. First, it is vacuously true that a is in E . Second, consider an arbitrary c in E , so that for x in M we have $x < c$ implies $f(x) \leq c$. Suppose that x is in M and $x < f(c)$. Since $M \subset M_c$, it follows that $x \leq c$ or $f(c) \leq x$. However the latter possibility is ruled out, so $x \leq c$. If $x < c$, then $f(x) \leq c \leq f(c)$, and if $x = c$ then again $f(x) \leq f(c)$. This is enough to imply that $f(x) \leq f(c)$. Thus for x in M we have that $x < f(c)$ implies $f(x) \leq f(c)$. This shows that $f(c)$ is in E . In other words, f leaves E invariant. Third, let T be a totally ordered subset of E with supremum b . Suppose x is in M with $x < b$. Since for all c in E we have $M \subset M_c$, either $f(c) \leq x$ for all c in T , or $x \leq c$ for some c in T . In the first case x is an upper bound for T , and so the least upper bound $b \leq x$. This is a contradiction. In the remaining second case $x \leq c$ for some c in T . If $x < c$, then $f(x) \leq c \leq b$, otherwise $x = c$ is in E and since $b \leq x$ is ruled out, again we have $f(x) \leq b$. Thus for all x in M we have that $x < b$ implies $f(x) \leq b$. Hence b is in E .

So $M \subset E$, in fact they are equal. Every point of M is an choke point.

Now we are done. Suppose that x and y are in M . Since $M \subset E$, it follows that x is in E . Since $M \subset M_x$, it follows that y is in M_x . Thus $y \leq x$ or $f(x) \leq y$. Hence $y \leq x$ or $x \leq y$. This proves that M is totally ordered. Therefore it has a supremum b . However $b \leq f(b) \leq b$. So b is a fixed point of f . \square

8.7 Appendix: Zorn’s lemma

The following is an optional topic. It is the proof that the axiom of choice applies Zorn’s lemma.

Theorem 8.8 (Hausdorff maximal principle) *Every partially ordered set has a maximal totally ordered subset.*

Proof: Let P be the ordered set. Consider the set A of all totally ordered subsets of P . Suppose that T is a non-empty totally ordered subset of A . Then $\bigcup T$ is a totally ordered subset of A , and it is the supremum of T . Suppose there is no maximal element of A . Then for each x in A the set of U_x of totally ordered subsets y of P with $x \subset y$ and $x \neq y$ is non-empty. By the axiom of choice there is a function $f : A \rightarrow A$ such that $f(x) \in U_x$. This f does not have a fixed point. This contradicts the Bourbaki fixed point theorem. \square

Theorem 8.9 (Zorn’s lemma) *Consider a non-empty partially ordered set such that every non-empty totally ordered subset has an upper bound. Then the set has a maximal element.*

Proof: Let P be the ordered set. By the Hausdorff maximal principle there is a maximal totally ordered subset X . Since P is not empty, X is not empty. Therefore there is a maximal element m in X . Suppose there were an element p with $m \neq p$ and $m < p$. Then we could adjoin p to X and get a strictly larger totally ordered subset. This is a contradiction. So m is maximal in P . \square

Chapter 9

Radon measures

9.1 Topology and measure

The interaction of topology and measure is complicated. A topological structure on a space X may somehow determine a measurable structure on X . The simplest example of this is that the topology itself generates the Borel σ -algebra. It turns out, however, that various technicalities arise. In particular, there are two directions that may be taken.

The first possible direction is to take X to be a locally compact Hausdorff space. This seems quite general, since such a space need not be a metrizable space. (A typical example where this is so is when X is an uncountable product of compact Hausdorff spaces, so that X is itself a compact Hausdorff space, not metrizable.) However in this generality there are technicalities due to the fact that a topology involves uncountable operations, and these interact uneasily with measure theory, which is primarily based on countable operations. This first direction is not emphasized in the present treatment. However there will be some discussion in an appendix.

The other direction is to take X to be a separable metric space. Often X is a complete separable metric space. This is general in a different way, since such a space need not be locally compact. (A typical example is when X is an infinite-dimensional Banach space; such a space is never locally compact.) In this setting compactness issues can be something of a struggle.

The best possible world is when the space X is a separable locally compact metric space. This is general enough for many applications, and most of the technicalities are gone.

In the next section we argue that the setting where X is a complete separable metric space the measurable structure is very close to being unique. The only possibilities for the Borel structure are that belonging to a countable set and that belonging to the unit interval $[0, 1]$.

9.2 Borel isomorphisms

Two measurable spaces are isomorphic if there is a bijection between the underlying sets that preserves the measurable subsets. If the spaces are topological spaces, and the measurable subsets are the Borel subsets, then the spaces are said to be Borel isomorphic.

Examples:

1. The measurable spaces $(0, 1)$ and $[0, 1)$ are Borel isomorphic. These spaces are homeomorphic to $(0, \infty)$ and $[0, +\infty)$. So to prove this, it is sufficient to show that $(0, +\infty)$ and $[0, +\infty)$ are Borel isomorphic. An isomorphism is given by $f(x) = n + 1 - (x - n)$ on $n < x \leq n + 1$. This is a Borel isomorphism, but it is far from being continuous.
2. The measurable space $[0, 1)$ and $[0, 1]$ are Borel isomorphic. It is obvious that $[0, 1)$ is homeomorphic to $(0, 1]$, so it is sufficient to show that $(0, 1]$ and $[0, 1]$ are Borel isomorphic. However $(0, 1)$ is a Borel subset of $(0, 1]$, and $[0, 1)$ is a Borel subset of $[0, 1]$. So we can take the isomorphism we got in the previous example between these subsets, and send 1 to 1.

Theorem 9.1 *Let X and Y be two uncountable separable metric spaces, each a Borel subset of a complete separable metric space. Then these spaces are Borel isomorphic.*

The theorem will be proved for complete separable metric spaces in the final chapter of these notes. Meanwhile, we shall remark on some consequences. Say that a measurable space is *standard* if it is isomorphic, as a measurable space, to a separable complete metric space with the Borel σ -algebra.

Examples:

1. Both $[0, 1]$ and $(-\infty, +\infty)$ are standard, since they are complete separable metric spaces.
2. $(0, 1)$ and $[0, 1)$ are standard, since they are Borel isomorphic to $[0, 1]$. It is true that these are not complete spaces, but they are Borel subsets of the complete space $[0, 1]$.

Corollary 9.2 *Every standard measurable space is isomorphic, as a measurable space, to a countable set with the discrete topology or to the unit interval $[0, 1]$ with the usual metric topology.*

Proof: If the space is uncountable, then it is isomorphic to $[0, 1]$, since $[0, 1]$ is a complete separable metric space. If the space is countable, then every subset is a Borel set. So it is Borel isomorphic to a finite set or to \mathbf{N} with the discrete topology. \square

Corollary 9.3 *Every standard measurable space is isomorphic, as a measurable space, to a compact metric space.*

Proof: The unit interval $[0, 1]$ is compact. Every finite set is compact. The remaining case is that of a countable infinite set. Take the space to be $\mathbf{N} \cup \{\infty\}$, the one point compactification of the natural numbers. This is a countable compact metric space. Every subset is a Borel set. \square

9.3 Metric spaces

Among topological spaces metric spaces are particularly nice. This section is an attempt to explain the special role of metric spaces.

A topological space is T_1 if for every pair of points there is an open set with the first point not in it and the second point in it. This is equivalent to the condition that single point sets are closed sets.

A topological space is *Hausdorff* or T_2 if every pair of points is separated by a pair of disjoint open sets.

A topological space is *regular* if it is T_1 and also satisfies condition T_3 : every pair consisting of a closed set and a point not in the set is separated by a pair of disjoint open sets.

A topological space is *normal* if it is T_1 and satisfies condition T_4 : every pair of closed sets is separated by a pair of disjoint open sets.

Theorem 9.4 *Consider a metric space. Given two disjoint closed sets, there is a continuous real function with values in $[0, 1]$ that is zero on one set and one on the other set. In particular, every metric space is normal.*

Proof: For every closed set E the function $x \mapsto d(x, E)$ is a continuous function that vanishes exactly on E . If E and F are disjoint closed sets, then

$$g(x) = \frac{d(x, E)}{d(x, E) + d(x, F)} \quad (9.1)$$

vanishes on E and is one on F . \square

Corollary 9.5 *Let X be a locally compact metric space. Let $K \subset X$ be a compact subset. Then there is a continuous function with values in $[0, 1]$ with compact support that has the value 1 on K .*

Proof: Since X is locally compact, it is not hard to prove that there is an open set U and a compact set L with $K \subset U \subset L$. Take f to be 1 on K and zero on the complement of U . Then f has support in L . \square

Notice that this result fails without the hypothesis of local compactness. For example, consider a point in infinite dimensional Hilbert space. Say that there is a real continuous function on the Hilbert space that is one at the point. Then it is non-zero on some non-empty open set. However a non-empty open set is never a subset of a compact set.

Corollary 9.6 *If f is a locally compact metric space, then the closure of the space $C_c(X)$ of continuous functions with compact support in $BC(X)$ is $C_0(X)$, the space of continuous functions that vanish at infinity.*

Proof: Since a function with compact support vanishes at infinity, it follows that the closure of $C_c(X)$ is a subset of $C_0(X)$. The converse is slightly more complicated. Suppose that f is in $C_0(X)$. Then there is a compact set K_n such that $|f| < 1/n$ on the complement of K_n . Let g_n be in $C_c(X)$ with $0 \leq g_n \leq 1$ and with $g_n = 1$ on K_n . Then $g_n f$ is in $C_c(X)$. Furthermore, $(1 - g_n)|f|$ is bounded by $1/n$. Hence $g_n f \rightarrow f$ uniformly. So f is in the closure of $C_c(X)$. \square

The following results are stated without proof. They may serve to clarify the relation between metric spaces and more general topological spaces.

Theorem 9.7 *A compact Hausdorff space is normal.*

Theorem 9.8 *A locally compact Hausdorff space is regular.*

Recall that in a topological space second countable implies separable. The converse is true for metric spaces.

Theorem 9.9 (Urysohn) *A second countable regular space is metrizable.*

Corollary 9.10 *A second countable locally compact Hausdorff space is metrizable.*

9.4 Riesz representation

A topological space is said to be σ -compact if it is a countable union of compact subsets. Suppose that X is a second countable locally compact Hausdorff space. Then it may be shown that X is σ -compact.

Let X be a separable locally compact metric space. Then X is second countable and hence σ -compact. It is then not hard to show that the σ -algebra generated by $C_c(X)$ is the same as the σ -algebra generated by $BC(X)$.

A separable (or second countable) locally compact metric space is an open subset of a compact space, its one point compactification. This space is also second countable, so it may be considered as a metric space. However a compact metric space is always a complete separable metric space, and hence gives a standard measure space. Thus the separable locally compact metric space also gives a standard measure space.

Theorem 9.11 (Dini) *Suppose that X is a compact space. Let $f_n \downarrow 0$ be a decreasing sequence of continuous functions that converges pointwise to zero. Then f_n converges uniformly to zero.*

Proof: Consider $\epsilon > 0$. The set where $f_n \geq \epsilon$ is a closed subset of a compact space and hence is compact. The pointwise convergence implies that the intersection of the collection of all these sets is zero. Hence there is a finite subcollection whose intersection is zero. However this is a decreasing sequence of sets. Therefore these sets are empty from some index on. That is, from some index on the set $f_n < \epsilon$ is the whole space. This implies uniform convergence. \square

A Radon measure on a space of real continuous functions is a linear function μ from the space to the reals that is order preserving. Thus in particular $f \geq 0$ implies $\mu(f) \geq 0$. A Radon measure might better be called a Radon integral, since it acts on functions rather than on sets, but both terms are used.

Theorem 9.12 (Riesz representation) *Let X be a separable locally compact metric space. Then there is a natural bijective correspondence between Radon measures on $C_c(X)$ and Borel measures that are finite on compact sets.*

Proof: Suppose each f_n is in $C_c(X)$ and that $f_n \downarrow 0$ pointwise. There is a fixed compact set K such that each f_n has support in K . According to Dini's theorem, $f_n \rightarrow 0$ uniformly. Let g be in $C_c(X)$ and have the value 1 on K . Then $0 \leq f_n \leq \|f_n\|_{sup}g$, so $0 \leq \mu(f_n) \leq \|f_n\|_{sup}\mu(g)$. Thus $\mu(f_n) \rightarrow 0$. That is, μ satisfies the monotone convergence theorem on $C_c(X)$. Thus μ is an elementary integral.

Since 1 is in the monotone class generated by $C_c(X)$, it follows that this monotone class coincides with the σ -algebra generated by $C_c(X)$. The elementary integral extends uniquely to this σ -algebra, which is the Borel σ -algebra. \square

Corollary 9.13 *Let X be a separable locally compact metric space. Then there is a natural bijective correspondence between Radon measures on $C_0(X)$ and finite Borel measures.*

Proof: It is not hard to see that if μ is not finite, then there exists a $f \geq 0$ in $C_0(X)$ such that $\mu(f) = +\infty$. \square

There is also a result where the Radon measure μ is not required to be order preserving but only to be continuous on $C_0(X)$. Then the conclusion is that μ is given by a finite signed Borel measure.

Corollary 9.14 *Let X be a compact metric space. Let $\mu : C(X) \rightarrow \mathbf{R}$ be linear and order preserving. Then there is a natural bijective correspondence between Radon measures on $C(X)$ and finite Borel measures.*

9.5 Lower semicontinuous functions

Let us look more closely at the extension process in the case of a Radon measure. We begin with the positive linear functional on the space $L = C_c(X)$ of continuous functions with compact support. The construction of the integral

associated with the Radon measure proceeds in the standard two stage process. The first stage is to consider the integral on the spaces $L \uparrow$ and $L \downarrow$. The second stage is to use this extended integral to define the integral of an arbitrary summable function.

A function f from X to $(-\infty, +\infty]$ is said to be *lower semicontinuous* (LSC) if for each real a the set $\{x \mid f(x) > a\}$ is an open set. A function f from X to $[-\infty, +\infty)$ is said to be *upper semicontinuous* (USC) if for each real a the set $\{x \mid f(x) < a\}$ is an open set. Clearly a continuous real function is both LSC and USC.

Theorem 9.15 *If each f_n is LSC and if $f_n \uparrow f$, then f is LSC. If each f_n is USC and if $f_n \downarrow f$, then f is USC.*

It follows from this theorem that space $L \uparrow$ consists of functions that are LSC. Similarly, the space $L \downarrow$ consists of functions that are USC. These functions can already be very complicated. The first stage of the construction of the integral is to use the monotone convergence theorem to define the integral on the spaces $L \uparrow$ and $L \downarrow$.

In order to define the integral for a measurable functions, we approximate such a function from above by a function in $L \uparrow$ and from below by a function in $L \downarrow$. This is the second stage of the construction. The details were presented in an earlier chapter.

The following is a useful result that we state without proof.

Theorem 9.16 *If μ is a Radon measure and if $1 \leq p < \infty$, then $C_c(X)$ is dense in $L^p(X, \mathcal{B}, \mu)$.*

Notice carefully that the corresponding result for $p = \infty$ is false. The uniform closure of $C_c(X)$ is $C_0(X)$, which in general is much smaller than $L^\infty(X, \mathcal{B}, \mu)$. A bounded function does not have to be continuous, nor does it have to vanish at infinity.

9.6 Weak* convergence

In order to emphasize the duality between the space of measures and the space of continuous functions, we sometimes write the value of the Radon measure μ on the continuous function f as

$$\mu(f) = \langle \mu, f \rangle. \tag{9.2}$$

As before, we consider only positive Radon measures, though there is a generalization to signed Radon measures. We consider finite Radon measures, that is, Radon measures for which $\langle \mu, 1 \rangle < \infty$. Such a measure extends by continuity to $C_0(X)$, the space of real continuous functions that vanish at infinity. In the case when $\langle \mu, 1 \rangle = 1$ we are in the realm of probability. Throughout we take X to be a separable locally compact metric space, though a more general setting is possible.

In this section we describe *weak* convergence* for Radon measures. In probability this is often called *vague convergence*. A sequence μ_n of finite Radon measures is said to weak* converge to a Radon measure μ if for each f in $C_0(X)$ the numbers $\langle \mu_n, f \rangle \rightarrow \langle \mu, f \rangle$.

The importance of weak* convergence is that it gives a sense in which two probability measures with very different qualitative properties can be close. For instance, consider the measure

$$\mu_n = \frac{1}{n} \sum_{k=1}^n \delta_{\frac{k}{n}}.$$

This is a Riemann sum measure. Also, consider the measure

$$\langle \lambda, f \rangle = \int_0^1 f(x) dx.$$

This is Lebesgue measure on the unit interval. Then $\mu_n \rightarrow \lambda$ in the weak* sense, even though each μ_n is discrete and λ is continuous.

A weak* convergent sequence can lose mass. For instance, a sequence of probability measures μ_n can converge in the weak* sense to zero. A simple example is the sequence δ_n . The following theory shows that a weak* convergent sequence cannot gain mass.

Theorem 9.17 *If $\mu_n \rightarrow \mu$ in the weak* sense, then $\langle \mu, f \rangle \leq \liminf_{n \rightarrow \infty} \langle \mu_n, f \rangle$ for all $f \geq 0$ in $BC(X)$.*

Proof: It is sufficient to show this for f in BC with $0 \leq f \leq 1$. Choose $\epsilon > 0$. Let $0 \leq g \leq 1$ be in C_0 so that $\langle \mu, (1-g) \rangle < \epsilon$. Notice that gf is in C_0 . Furthermore, $(1-g)f \leq (1-g)$ and $gf \leq f$. It follows that

$$\langle \mu, f \rangle \leq \langle \mu, gf \rangle + \langle \mu, (1-g) \rangle \leq \langle \mu, gf \rangle + \epsilon \leq \langle \mu_k, gf \rangle + 2\epsilon \leq \langle \mu_k, f \rangle + 2\epsilon \quad (9.3)$$

for k sufficiently large. \square

The following theorem shows that if a weak* convergent sequence does not lose mass, then the convergence extends to all bounded continuous functions.

Theorem 9.18 *If $\mu_n \rightarrow \mu$ in the weak* sense, and if $\langle \mu_n, 1 \rangle \rightarrow \langle \mu, 1 \rangle$, then $\langle \mu_n, f \rangle \rightarrow \langle \mu, f \rangle$ for all f in $BC(X)$.*

Proof: It is sufficient to prove the result for f in BC with $0 \leq f \leq 1$. The preceding result gives an inequality in one direction, so it is sufficient to prove the inequality in the other direction. Choose $\epsilon > 0$. Let $0 \leq g \leq 1$ be in C_0 so that $\langle \mu, (1-g) \rangle < \epsilon$. Notice that gf is in C_0 . Furthermore, $(1-g)f \leq (1-g)$ and $gf \leq f$. For this direction we note that the extra assumption implies that $\langle \mu_n, (1-g) \rangle \rightarrow \langle \mu, (1-g) \rangle$. We obtain

$$\langle \mu_n, f \rangle \leq \langle \mu_n, gf \rangle + \langle \mu_n, (1-g) \rangle \leq \langle \mu, gf \rangle + \langle \mu, (1-g) \rangle + 2\epsilon \leq \langle \mu, gf \rangle + 3\epsilon \leq \langle \mu, f \rangle + 3\epsilon \quad (9.4)$$

for n sufficiently large. \square

It is not true in general that the convergence works for discontinuous functions. Take the function $f(x) = 1$ for $x \leq 0$ and $f(x) = 0$ for $x > 0$. Then the measures $\delta_{\frac{1}{n}} \rightarrow \delta_0$ in the weak* sense. However $\langle \delta_{\frac{1}{n}}, f \rangle = 0$ for each n , while $\langle \delta_0, f \rangle = 1$.

We now want to argue that the convergence takes place also for certain discontinuous functions. A quick way to such a result is through the following concept. For present purposes, we say that a bounded measurable function g has μ -negligible discontinuities if for every $\epsilon > 0$ there are bounded continuous functions f and h with $f \leq g \leq h$ and such that $\mu(f)$ and $\mu(h)$ differ by less than ϵ .

Example: If λ is Lebesgue measure on the line, then every piecewise continuous function with jump discontinuities has λ -negligible discontinuities.

Example: If δ_0 is the Dirac mass at zero, then the indicator function of the interval $(-\infty, 0]$ does not have δ_0 -negligible discontinuities.

Theorem 9.19 *If $\mu_n \rightarrow \mu$ in the weak* sense, and if $\langle \mu_n, 1 \rangle \rightarrow \langle \mu, 1 \rangle$, then $\langle \mu_n, g \rangle \rightarrow \langle \mu, g \rangle$ for all bounded measurable g with μ -negligible discontinuities.*

Proof: Take $\epsilon > 0$. Take f and h in BC such that $f \leq g \leq h$ and $\mu(f)$ and $\mu(h)$ differ by at most ϵ . Then $\mu_n(f) \leq \mu_n(g) \leq \mu_n(h)$ for each n . It follows that $\mu(f) \leq \liminf_{n \rightarrow \infty} \mu_n(g) \leq \limsup_{n \rightarrow \infty} \mu_n(g) \leq \mu(h)$. But also $\mu(f) \leq \mu(g) \leq \mu(h)$. This says that $\liminf_{n \rightarrow \infty} \mu_n(g)$ and $\limsup_{n \rightarrow \infty} \mu_n(g)$ are each within ϵ of $\mu(g)$. Since $\epsilon > 0$ is arbitrary, this proves that $\lim_n \mu_n(g)$ is $\mu(g)$. \square

9.7 Central limit theorem for coin tossing

For each n consider the space $\{0, 1\}^n$ for the outcomes of n coin tosses. We shall think of 0 as heads and 1 as tails. The probability measure for fair tosses of a coin is the uniform measure that assigns measure $1/2^n$ to each of the 2^n one point sets.

For $j = 1, \dots, n$ let x_j be the function on $\{0, 1\}^n$ that has value $x_j(\omega) = 1$ if $\omega_j = 0$ and $x_j = -1$ if $\omega_j = 1$. Thus $x_1 + \dots + x_n$ is the number of heads minus the number of tails. It is a number between $-n$ and n .

Let μ_n be the probability measure on the line that is the image of the uniform probability measure on $\{0, 1\}^n$ under the map

$$\frac{1}{\sqrt{n}}(x_1 + \dots + x_n) : \{0, 1\}^n \rightarrow \mathbf{R}. \quad (9.5)$$

This measure assigns mass

$$p_k = \binom{n}{k} \frac{1}{2^n} \quad (9.6)$$

to the point $(2k - n)/\sqrt{n}$, for $k = 0, 1, \dots, n$. So it is very much a discrete measure.

Theorem 9.20 (Central limit theorem for coin tossing) *The measures μ_n converge in the weak* sense to the standard normal probability measure μ with density*

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \quad (9.7)$$

with respect to Lebesgue measure on the line.

9.8 Wiener measure

This section treats Wiener measure (otherwise known as the Einstein model for Brownian motion). Fix a time interval $[0, T]$. The space on which the measure lives is the space $X = C([0, T])$ of all real functions on $[0, T]$. These are considered as function from time to one dimensional space. According to Einstein there is a parameter $\sigma > 0$ that related time to space. It is defined so that the expectation of the square of the distance travelled is σ times the elapsed time.

The space X is a complete separable metric space. It is not locally compact. However there is an appropriate notion of convergence for probability measures. A sequence μ_n of probability measures is said to converge weak* to a probability measure μ if for every F in $BC(X)$ we have $\mu_n(F) \rightarrow \mu(F)$ as $n \rightarrow \infty$.

Consider a natural number N and define $\Delta t = T/N$. For each $j = 0, 1, 2, 3, \dots, N$ there are corresponding time instants $0, \Delta t, 2\Delta t, 3\Delta t, T$. For a coin toss sequence ω in $\{0, 1\}^N$ and a time t in $[0, T]$ with $j\Delta t \leq t \leq (j+1)\Delta t$ define

$$W(t, \omega) = \sigma\sqrt{\Delta t}(x_1(\omega) + \dots + x_j(\omega)) + \frac{t - j\Delta t}{\Delta t}x_{j+1}(\omega). \quad (9.8)$$

For each coin toss sequence ω the function $t \mapsto W(t, \omega)$ is a piecewise linear continuous path in the space $X = C([0, T])$. Define μ_N as the image of the coin tossing measure on $\{0, 1\}^N$ in the space $X = C([0, T])$.

Theorem 9.21 (Existence of Wiener measure) *Fix the total time T and the diffusion constant σ . For each N there is a probability measure μ_N in $X = C([0, T])$ defined as above by N tosses of a fair coin. The assertion is that there is a probability measure μ in $X = C([0, T])$ such that $\mu_n \rightarrow \mu$ in the weak* sense as $n \rightarrow \infty$.*

The Wiener measure has remarkable properties.

1. For each t in $[0, T]$ the expectation $\mu(W(t)) = 0$.
2. For each t in $[0, T]$ and $h \geq 0$ with $t+h$ in $[0, T]$ we have $\mu(W(t)W(t+h)) = \sigma^2 t$. In particular the variance is $\mu(W(t))^2 = \sigma^2 t$.
3. For each t the random variable $W(t)$ has the normal distribution with density

$$\phi_t(x) = \frac{1}{\sqrt{2\pi\sigma^2 t}} e^{-\frac{x^2}{2\sigma^2 t}}. \quad (9.9)$$

The second property implies that the next increment $W(t+h) - W(t)$ is uncorrelated with the present $W(t)$, that is, $\mu((W(t+h) - W(t))W(t)) = 0$.

As a consequence of these properties we see that for $h \geq 0$ we have

$$\mu((W(t+h) - W(t))^2) = \sigma^2 h \quad (9.10)$$

This is consistent with the fact that the paths are continuous. However for $h > 0$ this also says that

$$\mu\left(\left(\frac{W(t+h) - W(t)}{h}\right)^2\right) = \frac{\sigma^2}{h}. \quad (9.11)$$

This suggests that though the typical Wiener path is continuous, it is not differentiable.

9.9 Appendix: Measure theory on non-metric spaces

The following is an optional topic. It is a brief survey of the interaction of measure and topology when the topological space is not a separable metric space. In that case there are additional technical difficulties. Their origin is the following. A topological space is characterized by open sets allowing uncountable unions and finite intersections. A measurable space is characterized by measurable sets allowing countable unions, countable intersections, and complements. The tension arises from a situation when the uncountable operations for a topological space enter the measure theory.

Let X be a locally compact Hausdorff space. The *Baire* σ -algebra is the σ -algebra generated by the space $C_c(X)$ of real continuous functions on X , each of which has compact support. This is the same σ -algebra as that generated by the space $C_0(X)$ of real continuous functions on X that vanish at infinity.

The *extended Baire* σ -algebra is the σ -algebra generated by the space $BC(X)$ of bounded real continuous functions on X . This is the same σ -algebra as that generated by the space $C(X)$ of all real continuous functions on X . An example where the Baire σ -algebra is strictly smaller than the extended Baire σ -algebra is when X is an uncountable discrete space. This discrepancy is eliminated when X is a σ -compact locally compact Hausdorff space. In particular, it will be true for a second countable locally compact Hausdorff space.

The Borel σ -algebra is the σ -algebra generated by the open sets. It is possible that the Borel σ -algebra is strictly larger than the Baire σ -algebra.

An example is to take a compact Hausdorff space Y (such as the two point set $\{0,1\}$ or the unit interval $[0,1]$) and an uncountable index set I . Then $X = Y^I$ is a compact Hausdorff space, but it is not metrizable. Each set or function in the Baire σ -algebra depends only on countable many coordinate in I . On the other hand, each one point set in X is closed. Such a set is defined by restricting all of the coordinates in I to have fixed values.

If X is a compact metric space, then the Baire and Borel σ -algebras coincide. This is the case in many practical situations. In fact, if X is a compact metric space, then in particular it is a separable complete metric space. So the σ -algebra is a standard Borel σ -algebra.

Recall that a Radon measure on a space of real continuous functions is a linear functional μ from the space to the reals that is order preserving. Thus in particular $f \geq 0$ implies $\mu(f) \geq 0$.

Theorem 9.22 (Riesz representation) *Let X be a locally compact Hausdorff space. Suppose that it is also σ -compact. Then there is a natural bijective correspondence between Radon measures on $C_c(X)$ and Baire measures that are finite on compact sets.*

Corollary 9.23 *Let X be a locally compact Hausdorff space. Suppose that it is also σ -compact. Then there is a natural bijective correspondence between Radon measures on $C_0(X)$ and finite Baire measures.*

There is also a result where the Radon measure μ is not required to be order preserving but only to be continuous on $C_0(X)$. Then the conclusion is that μ is given by a finite signed Baire measure.

Corollary 9.24 *Let X be a compact Hausdorff space. Let $\mu : C(X) \rightarrow \mathbf{R}$ be linear and order preserving. Then there is a natural bijective correspondence between Radon measures on $C(X)$ and finite Baire measures.*

There are certain problems where it is desirable to have a correspondence between Radon measures and Borel measures. In order to make such a correspondence well defined, the Borel measure must be assumed to have regularity properties with respect to certain possibly uncountable supremum and infimum operations.

Here is an example where such issues arise. Consider the space $X = C([0, T])$ that was used in the example of Wiener measure. It is a nice complete separable metric space, and the notion of convergence, uniform convergence, seems quite natural. However it is not locally compact. It would be nice if it were possible to use a compact space in the construction of Wiener measure.

One approach is to use the product space $[-\infty, \infty]^{[0, T]}$. This has the topology of pointwise convergence. This space is much larger, but it is automatically compact, because of the way that the topology is defined. However it is not a metric space. Furthermore, the continuous functions are not a Baire subset, but only a Borel subset. So one has to deal with measure theory technicalities in order to get a measure on the space of continuous functions. This approach is so elegant, however, that it does lead to the temptation to study measures on compact Hausdorff spaces that are not metrizable. However using such a big space does not really save work, since the estimates that are needed to show that the measure lives in the space of continuous functions are rather similar to the estimates that are needed to establish the required compactness properties in the approach that works directly with $X = C([0, T])$.

Another approach is to take the Tychonoff-Čech compactification of X . Every metric space X has such a compactification K , but if the metric space is not compact, then K will not be metrizable. However K is a compact Hausdorff space. The functions in $BC(X)$ extend uniquely to functions in $C(K)$, so the dual of $BC(X)$ may be thought of as finite signed measures in K . The problem is that K is such a complicated space that does not give much of a handle on the probability measures in X , which is what one really wants to study. So here we again get compact Hausdorff spaces that are not metrizable, but this time they are so big that they do not give much insight.

9.10 Problems

1. Recall that X is locally compact if for every point x in X there is an open subset U and a compact subset K with $x \in U \subset K$. Suppose that X is locally compact. Prove that for every compact subset $M \subset X$ there is an open subset V and a compact subset N such that $M \subset V \subset N$.
2. Recall that $f : X \rightarrow \mathbf{R}$ is lower semicontinuous (LSC) if and only if the inverse image of each interval $(a, +\infty)$, where $-\infty \leq a \leq +\infty$, is open in X . Show that if $f_n \uparrow f$ pointwise and each f_n is LSC, then f is LSC. (This holds in particular if each f_n is continuous.) Challenge: Give an example of a function $f : \mathbf{R} \rightarrow \mathbf{R}$ that is LSC but is discontinuous at every point (except perhaps on a set of measure zero).
3. Show that if $f : X \rightarrow \mathbf{R}$ is LSC, and X is compact, then there is a point in X at which f has a minimum value. Show by example that if $f : X \rightarrow \mathbf{R}$ is LSC, and X is compact, then there need not be a point in X at which f has a maximum value.
4. Let H be a real Hilbert space. Let $L : H \rightarrow \mathbf{R}$ be continuous and linear. Thus in particular L is Lipschitz, that is, there is an M with $|L(u)| \leq M\|u\|$ for all u in H . Consider the problem of proving that there is a point in H at which the function $F : H \rightarrow \mathbf{R}$ defined by $F(u) = (1/2)\|u\|^2 - L(u)$ has a minimum value. One approach would be to look at a sufficiently large closed ball centered at the origin and argue that if there were a minimum, it would be in that closed ball. If H were finite dimensional, that ball would be compact, and the result is obvious. Show how to carry out the compactness proof for infinite dimensional H . Hint: Switch to the weak topology. Be explicit about which functions are continuous or lower semicontinuous.
5. The context is Borel functions on the real line. Let f be in L^2 and g be in L^1 . Let $T : L^2 \rightarrow L^2$ be defined by the convolution $Tf = g * f$. Then T is a continuous linear transformation. For each polynomial p define $\mu(p) = \langle f, p(T)f \rangle$. Then μ defines a Radon measure and hence is given by a Borel measure on the line. Find this measure explicitly, as the image of an absolutely continuous measure.

Chapter 10

Distributions

10.1 Fréchet spaces

In this chapter we consider complex vector spaces. This is mainly for convenience in the examples. Recall that a Banach space is a vector space with a norm that is a complete metric space. A Fréchet space is a vector space E with a countable family of seminorms that is a complete metric space.

Recall that a seminorm is a function $u \mapsto \|u\|$ that is positive and satisfies the homogeneity property and the triangle inequality. For a Fréchet there is a countable family of seminorms $u \mapsto \|u\|_m$, for $m = 0, 1, 2, 3, \dots$. The topology on the Fréchet space is defined as the coarsest topology including all of the open balls $B_m(v, \epsilon) = \{u \mid \|u - v\|_m < \epsilon\}$. In order that it be a Hausdorff space it is required that for each $w \neq 0$ there is an m such that $\|w\|_m > 0$. A metric that defines the topology may be defined by using the bounded function $b(t) = t/(1 + t)$ to define

$$d(u, v) = \sum_{m=0}^{\infty} \frac{b(\|u - v\|_m)}{2^{m+1}}. \quad (10.1)$$

The space is required to be complete with respect to this metric.

The example to keep in mind is $E = C_c^\infty(K)$, where K is a compact subset of the real numbers. This consists of all smooth complex functions ϕ on \mathbf{R} that have support in K . Consider the usual supremum norm $\|\phi\|_{\text{sup}}$. Then the family of seminorms is defined by

$$\|\phi\|_m = \left\| \left(\frac{d}{dx} \right)^m f \right\|_{\text{sup}}. \quad (10.2)$$

Proposition 10.1 *Let u_n be a sequence of elements of a Fréchet space. Then $u_n \rightarrow u$ if and only if for each m we have $\|u_n - u\|_m \rightarrow 0$.*

10.2 Distributions

A function ϕ is said to be smooth or $C^\infty(\mathbf{R})$ if it has derivatives of all orders. Consider the space $\mathcal{D} = C_c^\infty(\mathbf{R})$ of complex functions defined on the real line, smooth and with compact support. These will be called *distribution test functions*.

The first amazing thing is that there are such functions. Recall that a real analytic function is a function given at each point on the real line by a convergent power series. No real analytic function can belong to $C_c^\infty(\mathbf{R})$, since if a real analytic function is zero on an open interval, then it is zero at all points. However the function $\chi(t) = e^{-\frac{1}{t}}$ for $t > 0$ and $\chi(t) = 0$ for $t \leq 0$ is a smooth function with support the positive half-line. It is not analytic, since $t = 0$ is an essential singularity. Once we have this example, then $\phi(x) = \chi(1 - x^2)$ is an example of a smooth function with support the compact interval $[-1, +1]$.

A *distribution* is a linear functional F from $\mathcal{D} = C_c^\infty(\mathbf{R})$ to the complex numbers. It must also satisfy a certain continuity condition. This property is that for each compact subset K of \mathbf{R} , the function F restricted to $C_c^\infty(K)$ is continuous on this Fréchet space.

The value of the distribution F on the test function ϕ in $\mathcal{D} = C_c^\infty(\mathbf{R})$ is written $\langle F, \phi \rangle$. The space of all distributions is denoted \mathcal{D}' . Thus it is a dual space to the space of test functions \mathcal{D} .

It is possible to define a topology on \mathcal{D} in such a way that \mathcal{D}' is the space of continuous complex linear functions on \mathcal{D} . However the topology on \mathcal{D} is not a metric topology. For this reason it is often more convenient to use the equivalent definition of distribution as a complex linear function on $\mathcal{D} = C_c^\infty(\mathbf{R})$ whose restriction to each $C_c^\infty(K)$ is continuous.

If f is a locally integrable function, then it defines a distribution by

$$\langle F, \phi \rangle = \int_{-\infty}^{\infty} f(x)\phi(x) dx. \quad (10.3)$$

Thus many functions define distributions. This is why distributions are also called generalized functions.

A sequence of distributions F_n converges to a distribution F if for each test function ϕ the numbers $\langle F_n, \phi \rangle$ converge to the number $\langle F, \phi \rangle$. This is of course just convergence in the weak* topology of \mathcal{D}' . In other words, it is the topology of pointwise convergence on the space of test functions \mathcal{D} .

Example. The distribution δ_a is defined by

$$\langle \delta_a, \phi \rangle = \phi(a). \quad (10.4)$$

This is not given by a locally integrable function. However a distribution may be written as a limit of functions. For instance, let $\epsilon > 0$ and consider the function

$$\delta_\epsilon(x - a) = \frac{1}{\pi} \frac{\epsilon}{(x - a)^2 + \epsilon^2}. \quad (10.5)$$

The limit of the distributions defined by these locally integrable functions as $\epsilon \downarrow 0$ is δ_a . For this reason the distribution is often written in the incorrect but suggestive notation

$$\langle \delta_a, \phi \rangle = \int_{-\infty}^{\infty} \delta(x-a)\phi(x) dx. \quad (10.6)$$

The most important distributions are $\delta(x)$, PV $1/x$, $1/(x-i0)$, and $1/(x+i0)$. These are the limits of the functions $\delta_\epsilon(x) = (1/\pi)\epsilon(x^2 + \epsilon^2)$, $x/(x^2 + \epsilon^2)$, $1/(x - i\epsilon)$, $1/(x + i\epsilon)$ as $\epsilon \downarrow 0$. The relations between these functions are given by

$$\delta_\epsilon(x) = \frac{1}{\pi} \frac{\epsilon}{x^2 + \epsilon^2} = \frac{1}{2\pi i} \left(\frac{1}{x - i\epsilon} - \frac{1}{x + i\epsilon} \right). \quad (10.7)$$

and

$$\frac{x}{x^2 + \epsilon^2} = \frac{1}{2} \left(\frac{1}{x - i\epsilon} + \frac{1}{x + i\epsilon} \right). \quad (10.8)$$

It follows that

$$\delta(x) = \frac{1}{2\pi i} \left(\frac{1}{x - i0} - \frac{1}{x + i0} \right). \quad (10.9)$$

and

$$PV \frac{1}{x} = \frac{1}{2\pi i} \left(\frac{1}{x - i0} - \frac{1}{x + i0} \right). \quad (10.10)$$

10.3 Operations on distributions

Operations on distributions are defined by looking at the example of a distribution defined by a function and applying the integration by parts formula in that case. Here we consider two basic operations: differentiation and multiplication by a smooth function. These give a practical way of defining distributions: Start with locally integrable functions and apply these operations, perhaps repeatedly.

The derivative is defined by

$$\langle F', \phi \rangle = -\langle F, \phi' \rangle. \quad (10.11)$$

Example: Consider the locally integrable Heaviside function given by $H(x) = 1$ for $x > 0$, $H(x) = 0$ for $x < 0$. Then $H' = \delta$. Here δ is given by

$$\langle \delta, \phi \rangle = \phi(0). \quad (10.12)$$

Example: Consider the locally integrable log function $f(x) = \log|x|$. Then $f'(x) = PV 1/x$. Here the principal value PV $1/x$ is given by

$$\langle PV \frac{1}{x}, \phi \rangle = \lim_{\epsilon \downarrow 0} \int_{-\infty}^{\infty} \frac{x}{x^2 + \epsilon^2} \phi(x) dx. \quad (10.13)$$

This can be seen by writing $\log|x| = \lim_{\epsilon} \log \sqrt{x^2 + \epsilon^2}$.

A distribution can be approximated by more than one sequence of functions. For instance, for each $a > 0$ let $\log_a(|x|) = \log(|x|)$ for $|x| \geq a$ and $\log_a(|x|) =$

$\log(a)$ for $|x| < a$. Then $\log_a(|x|)$ also approaches $\log(|x|)$ as $a \rightarrow 0$. So its derivative is another approximating sequence for PV $1/x$. This says that

$$\langle \text{PV } \frac{1}{x}, \phi \rangle = \lim_{a \downarrow 0} \int_{|x| > a} \frac{1}{x} \phi(x) dx. \quad (10.14)$$

We can use this sequence to compute the derivative of PV $1/x$. We get

$$\int_{-\infty}^{\infty} \left(\frac{d}{dx} \text{PV } \frac{1}{x} \right) \phi(x) dx = - \lim_{a \rightarrow 0} \int_{|x| > a} \frac{1}{x} \phi'(x) dx = \lim_{a \rightarrow 0} \left[\int_{|x| > a} -\frac{1}{x^2} \phi(x) dx + \frac{\phi(a) + \phi(-a)}{a} \right]. \quad (10.15)$$

But $[\phi(a) + \phi(-a) - 2\phi(0)]/a$ converges to zero, so this is

$$\int_{-\infty}^{\infty} \left(\frac{d}{dx} \text{PV } \frac{1}{x} \right) \phi(x) dx = \lim_{a \rightarrow 0} \int_{|x| > a} -\frac{1}{x^2} [\phi(x) - \phi(0)] dx. \quad (10.16)$$

Another operator is multiplication of a distribution by a *smooth* function ψ in $C^\infty(\mathbf{R})$. This is defined in the obvious way by

$$\langle \psi \cdot F, \phi \rangle = \langle F, \psi \phi \rangle. \quad (10.17)$$

Distributions are not functions. They may not have values at points, and in general nonlinear operations are not defined. For instance, the square of a distribution is not always a well defined distribution.

Also, some algebraic operations involving distributions are quite tricky. Consider, for instance, the associative law. Apply this to the three distributions $\delta(x)$, x , and PV $1/x$. Clearly the product $\delta(x) \cdot x = 0$. On the other hand, the product $x \cdot \text{PV } 1/x = 1$ is one. So if the associate law were to hold, we would get

$$0 = 0 \cdot \text{PV } \frac{1}{x} = (\delta(x) \cdot x) \cdot \text{PV } \frac{1}{x} = \delta(x) \cdot (x \cdot \text{PV } \frac{1}{x}) = \delta(x) \cdot 1 = \delta(x). \quad (10.18)$$

10.4 Mapping distributions

Another operation on distributions is composition with smooth function, that is, the distribution F is to be composed with a smooth function to get $F \circ g$. (The composition in the other order is rarely well-defined.) It turns out that this is somewhat tricky, since this operation is in some sense unnatural.

A test function ϕ is naturally viewed as a covariant object, so the distribution F is contravariant. A proper function is a function such that the inverse image of each compact set is compact. It is natural to define the forward push of a distribution F by a smooth proper function g by $\langle g[F], \phi \rangle = \langle F, \phi \circ g \rangle$. Example: If F is the distribution $\delta(x - 3)$ and if $u = g(x) = x^2 - 4$, then the forward push is $\delta(u - 5)$. This is because $\int \delta(u - 5) \phi(u) du = \int \delta(x - 3) \phi(x^2 - 4) dx$.

However the usual operation of composition $F \circ g$ is quite different. The idea is to attempt to think of a distribution as a covariant object, since a distribution

is supposed to be a generalized function. The backward pull of the distribution by a smooth function g is defined, in at least some circumstances, by

$$\langle F \circ g, \phi \rangle = \langle F, g[\phi] \rangle. \quad (10.19)$$

Here

$$g[\phi](u) = \sum_{g(x)=u} \frac{\phi(x)}{|g'(x)|}. \quad (10.20)$$

Example. Let $u = g(x) = x^2 - 4$, with $a > 0$. Then the backward pull of $\delta(u)$ under g is $\delta(x^2 - 4) = (1/4)(\delta(x - 2) + \delta(x + 2))$. This is because

$$g[\phi](u) = \frac{\phi(\sqrt{u^2 + 4}) + \phi(-\sqrt{u^2 + 4})}{2\sqrt{u^2 + 4}}. \quad (10.21)$$

So if $F = \delta$, then

$$F \circ g = \frac{1}{4}(\delta_2 + \delta_{-2}) \quad (10.22)$$

Example: The backward pull is not always defined. To consider a distribution as a covariant object is a somewhat awkward act in general. Let $u = h(x) = x^2$. The backward pull of $\delta(u)$ by h is $\delta(x^2)$, which is not defined.

Example: The general formula for the pull back of the delta function is

$$\delta(g(x)) = \sum_{g(a)=0} \frac{1}{|g'(a)|} \delta(x - a). \quad (10.23)$$

10.5 The support of a distribution

A distribution F is said to be zero on an open subset $U \subset \mathbf{R}$ if $\langle F, \phi \rangle = 0$ for all $C_c^\infty(U)$, the space of test functions with compact support in U . It may be shown that for each distribution F there is a largest open subset U on which F is zero. The complement of this subset is defined to be the support of F .

Examples:

1. The support of $\delta(x)$ is the origin.
2. The support of $PV1/x$ is the entire real line.

10.6 Radon measures

This section is a remark intended to clarify the relation between distributions and Radon measures.

Theorem 10.2 *Let F be a positive distribution defined on $\mathcal{D} = C_c^\infty(\mathbf{R})$. Thus $\phi \geq 0$ implies $F(\phi) \geq 0$. Then there exists a unique Radon measure μ defined on $C_c(\mathbf{R})$ so that the distribution F is the restriction of the Radon measure μ to $C_c^\infty(\mathbf{R})$.*

This theorem essentially says that a positive distribution is the same thing as a Radon measure. A positive distribution is a linear functional F on $C_c^\infty(\mathbf{R})$ such that for each test function ϕ the condition $\phi \geq 0$ implies that the value $\langle F, \phi \rangle \geq 0$. The theorem says that a positive distribution extends uniquely a linear functional μ on $C_c(\mathbf{R})$, the space of continuous functions with compact support satisfying the same positivity property.

It is common to write the value of a Radon measure μ in the form

$$\langle \mu, \phi \rangle = \int \phi(x) d\mu(x). \quad (10.24)$$

What is remarkable is that the theory of Lebesgue integration works for Radon measures. That is, given a real function $f \geq 0$ that is only required to be Borel measurable, there is a natural definition of the integral such that

$$0 \leq \int f(x) d\mu(x) \leq +\infty. \quad (10.25)$$

There is no such extension property for more general distributions.

10.7 Approximate delta functions

It might seem that one could replace the notion of distribution by the notion of a sequence of approximating functions. This is true in some sense, but the fact is that many different sequences may approximate the same distribution. Here is a result of that nature, for the special case of point masses.

Theorem 10.3 *Let $\delta_1(u) \geq 0$ be a positive function with integral 1. For each $\epsilon > 0$ define $\delta_\epsilon(x) = \delta_1(x/\epsilon)/\epsilon$. Then the functions δ_ϵ converge to the δ distribution as ϵ tends to zero.*

The convergence takes place in the sense of distributions (smooth test functions with compact support) or even in the sense of Radon measures (continuous test functions with compact support). Notice that there are no continuity or symmetry assumptions on the initial function.

Proof: Each δ_ϵ has integral one. Consider a bounded continuous function ϕ . Then

$$\int_{-\infty}^{\infty} \delta_\epsilon(x) \phi(x) dx = \int_{-\infty}^{\infty} \delta_1(u) \phi(\epsilon u) du. \quad (10.26)$$

The dominated convergence theorem shows that this approaches the integral

$$\int_{-\infty}^{\infty} \delta_1(u) \phi(0) du = \phi(0). \quad (10.27)$$

□

Here is an even stronger result.

Theorem 10.4 For each $\epsilon > 0$ let $\delta_\epsilon \geq 0$ be a positive function with integral 1. Suppose that for each $a > 0$

$$\int_{|x|>a} \delta_\epsilon(x) dx \rightarrow 0 \quad (10.28)$$

as $\epsilon \rightarrow 0$. Then the functions δ_ϵ converge to the δ distribution as ϵ tends to zero.

Proof: Let $H_\epsilon(a) = \int_0^a \delta_\epsilon(x) dx$. Then for each $a < 0$ we have $H_\epsilon(a) \rightarrow 0$, and for each $a > 0$ we have $1 - H_\epsilon(a) \rightarrow 0$. In other words, for each $a \neq 0$ we have $H_\epsilon(a) \rightarrow H(a)$ as $\epsilon \rightarrow 0$. Since the functions H_ϵ are uniformly bounded, it follows from the dominated convergence theorem that $H_\epsilon \rightarrow H$ in the sense of distributions. It follows by differentiation that $\delta_\epsilon \rightarrow \delta$ in the sense of distributions. \square

10.8 Tempered distributions

One problem with the general notion of distribution is that there is no satisfactory definition of Fourier transform. This is because it is impossible to have a function with compact support whose Fourier transform is a function with compact support. The only exception, of course, is the zero function. (The reason is that the Fourier transform of a function with compact support is analytic, and hence cannot have compact support unless it is the zero function.)

Tempered distributions are a slightly more restricted class for which there is a good definition of Fourier transform. The idea is to take a larger space of test functions, the Schwartz space. The dual space of tempered distributions is then a smaller space of distributions.

Let d/dx be the operator of differentiation, and let x be the operator of multiplication by the coordinate x . Let $\|\phi\|_{\text{sup}}$ be the usual supremum norm for functions on the line. The *Schwartz space* \mathcal{S} of rapidly decreasing smooth test functions is a Fréchet space. It consists of the smooth functions ϕ such that each seminorm $\|\phi\|_{m,p} = \|x^p(d/dx)^m \phi\|_{\text{sup}}$ is finite, for $m = 0, 1, 2, 3, \dots$ and $p = 0, 1, 2, 3, \dots$

A tempered distribution is a linear functional on \mathcal{S} satisfying the appropriate continuity property. The space of tempered distributions is denoted \mathcal{S}' .

A locally integrable function that defines a tempered distribution is called a *tempered locally integrable function*. An example of a locally integrable function that is not tempered is e^x . The integral of this function times a Schwartz function may not converge. On the other hand, there are functions that grow exponentially that are tempered. An example is $e^x \cos(e^x)$. In this case it is the oscillations that compensate for the exponential growth.

The derivative of a tempered distribution is a tempered distribution. For example, the function $\sin(e^x)$ is bounded and hence defines a tempered distribution. It follows that its derivative $e^x \cos(e^x)$ also defines a tempered distribution.

Furthermore, a tempered distribution may be multiplied by a *slowly increasing smooth function* to give another tempered distribution. A slowly increasing smooth function is a function ψ with the property that each derivative $(d/dx)^r \psi$ has absolute value bounded by some $C|x|^s$. Then if ϕ is in \mathcal{S} it follows that $\psi\phi$ is in \mathcal{S} . So we may define ψF by $\langle \psi F, \phi \rangle = \langle F, \psi\phi \rangle$ for an arbitrary tempered distribution F . The following theorem is an important special case that is key to the subsequent development of the Fourier transform for tempered distributions.

Theorem 10.5 *The operator of differentiation d/dx sends \mathcal{S}' into itself. The operator of multiplication by x sends \mathcal{S}' into itself.*

Every tempered distribution restricts to define a distribution. So tempered distributions are more special. The following lemma shows that the distribution uniquely determines the tempered distribution. So one can think of \mathcal{S}' as a subset of \mathcal{D}' .

Lemma 10.6 $\mathcal{D} = C_c^\infty(\mathbf{R})$ is dense in \mathcal{S} .

Proof: Let ϕ be in \mathcal{S} . Let ψ be in $C_c^\infty(\mathbf{R})$ with $0 \leq \psi \leq 1$ and $\psi = 1$ on some open ball centered at 0. For $a > 0$ let ψ^a be defined by $\psi^a(x) = \psi(ax)$. Then $\psi^a \phi$ is in $C_c^\infty(\mathbf{R})$. The idea is to show that $\psi^a \phi \rightarrow \phi$ in the topology of \mathcal{S} as $a \rightarrow 0$.

Calculate $x^p(d/dx)^m(\psi^a \phi - \phi)$. By the product rule, this gives a first term $x^p(\psi^a - 1)(d/dx)^m \phi$ plus additional terms involving derivatives of ψ^a .

Consider this first term. Consider $\epsilon > 0$. There is a large compact set K such that on the complement of this set $|x^p(d/dx)^m \phi|$ is bounded by $\epsilon/2$. Thus $|(\psi^a - 1)x^p(d/dx)^m \phi|$ is bounded by ϵ on the complement of K . On the other hand, for a sufficiently small the function $\psi^a = 1$ on K . For such a we have $(\psi^a - 1)x^p(d/dx)^m \phi = 0$ on K . So this first term becomes uniformly small as $a \rightarrow 0$.

Now consider the remaining terms. Each summand involves a derivative $(d/dx)^k \psi^a$ with $k \geq 1$. Such a derivative is bounded by a multiple of a^k . The remaining factors are bounded by constants independent of a . So these terms also approach zero uniformly as $a \rightarrow 0$. \square

The main lemma on tempered distributions is the following.

Lemma 10.7 *The Fourier transform sends \mathcal{S} into \mathcal{S} .*

Proof: Fix ϕ in \mathcal{S} . Consider arbitrary m and p . Then $(1+x^2)(d/dx)^p x^m \phi$ has bounded absolute value. It follows that $(d/dx)^p x^m \phi$ is in L^1 . Consequently, $k^p(d/dk)^m \hat{\phi}$ is in L^∞ . \square

Thus one can define the Fourier transform \hat{F} of an arbitrary tempered distribution F . The definition is

$$\langle \hat{F}, \phi \rangle = \langle F, \hat{\phi} \rangle. \quad (10.29)$$

Here $\hat{\phi}$ is the Fourier transform of ϕ . This gives the fundamental result.

Theorem 10.8 *The Fourier transform sends \mathcal{S}' into \mathcal{S}' .*

Practically everything that one says about the Fourier distribution also applies to the inverse Fourier. This is because the inverse Fourier transform $\check{\phi}$ is related to the Fourier transform $\hat{\phi}$ by $\check{\phi}(x) = \frac{1}{2\pi}\hat{\phi}(-x)$.

Proposition 10.9 *The Fourier transform of $\delta(x - a)$ is e^{-iak} . Thus, to synthesize a delta function one needs to weight all frequencies the same.*

Proof: This is just

$$\int_{-\infty}^{\infty} e^{-iak} \phi(k) dk = \hat{\phi}(a) = \int_{-\infty}^{\infty} \delta(x - a) \hat{\phi}(x) dx. \quad (10.30)$$

□

Proposition 10.10 *The Fourier transform of e^{ibx} is $2\pi\delta(k - b)$. Thus to synthesize a plane wave one puts all the mass on a single frequency.*

Proof: The 2π does not come for free. Indeed, the proof uses the inversion formula. Compute

$$2\pi \int_{-\infty}^{\infty} \delta(k - b) \phi(k) dk = 2\pi\phi(b) = \int_{-\infty}^{\infty} e^{ibx} \hat{\phi}(x) dx. \quad (10.31)$$

□

The preceding proposition illustrates why one only defines the Fourier transform for distributions that are tempered. Otherwise, one could take $b = -i$ imaginary in the above formula. Then the Fourier distribution of e^x would be $2\pi\delta(k + i)$. This is not a distribution, since smooth functions with compact support do not have analytic continuations to the complex plane.

Proposition 10.11 *The Fourier transform of d/dx is multiplication by ik . That is, for each tempered distribution F we have $\widehat{F'} = ik\hat{F}$.*

Proposition 10.12 *The Fourier transform of multiplication by x is id/dk . That is, for each tempered distribution F we have $\widehat{(xF)} = i(\hat{F})'$.*

Often one can compute the Fourier transform for a function and then pass to the limit to get a corresponding distribution formula. The following two propositions illustrate this situation.

Proposition 10.13 *The following Fourier transforms are valid for functions.*

$$\int_{-\infty}^{\infty} e^{-ikx} \frac{1}{x - i\epsilon} dx = 2\pi i e^{\epsilon k} H(-k) \quad (10.32)$$

and

$$\int_{-\infty}^{\infty} e^{-ikx} \frac{1}{x + i\epsilon} dx = -2\pi i e^{-\epsilon k} H(k). \quad (10.33)$$

Also

$$\int_{-\infty}^{\infty} e^{-ikx} \delta_{\epsilon}(x) dx = e^{-\epsilon|k|}. \quad (10.34)$$

and

$$\int_{-\infty}^{\infty} e^{-ikx} \frac{x}{x^2 + \epsilon^2} dx = -\pi i e^{-\epsilon|k|} \text{sign}(k). \quad (10.35)$$

Proposition 10.14 *The following Fourier transform formulas are valid for distributions.*

$$\mathcal{F}[1/(x - i0)] = 2\pi i H(-k) \quad (10.36)$$

and

$$\mathcal{F}[1/(x + i0)] = -2\pi i H(k). \quad (10.37)$$

Also,

$$\mathcal{F}[\delta(x)] = 1 \quad (10.38)$$

and

$$\mathcal{F}[\text{PV } \frac{1}{x}] = -\pi i \text{sign}(k). \quad (10.39)$$

For each of these formulas there is a corresponding inverse Fourier transform. For, instance, the inverse Fourier transform of 1 is

$$\delta(x) = \int_{-\infty}^{\infty} e^{ikx} \frac{dk}{2\pi} = \int_0^{\infty} \cos(kx) \frac{dk}{2\pi}. \quad (10.40)$$

Of course such an equation is interpreted by integrating both sides with a test function.

Another formula of the same type is gotten by taking the inverse Fourier transform of $-\pi i \text{sign}(k)$. This is

$$\text{PV } \frac{1}{x} = -\pi i \int_{-\infty}^{\infty} e^{ikx} \text{sign}(k) \frac{dk}{2\pi} = \int_0^{\infty} \sin(kx) dk. \quad (10.41)$$

Example: Here is a more complicated calculation. The derivative of PV $1/x$ is the distribution

$$\frac{d}{dx} \text{PV } \frac{1}{x} = -\frac{1}{x^2} + c\delta(x), \quad (10.42)$$

where c is the infinite constant

$$c = \int_{-\infty}^{\infty} \frac{1}{x^2} dx. \quad (10.43)$$

This makes rigorous sense if interprets it as

$$\int_{-\infty}^{\infty} \left(\frac{d}{dx} \text{PV } \frac{1}{x} \right) \phi(x) dx = \lim_{a \rightarrow 0} \int_{|x| > a} -\frac{1}{x^2} [\phi(x) - \phi(0)] dx. \quad (10.44)$$

One can get an intuitive picture of this result by graphing the approximating functions. The key formula is

$$\frac{d}{dx} \frac{x}{x^2 + \epsilon^2} = -\frac{x^2}{(x^2 + \epsilon^2)^2} + c_\epsilon \frac{2\epsilon^3}{\pi(x^2 + \epsilon^2)^2}, \quad (10.45)$$

where $c_\epsilon = \pi/(2\epsilon)$. This is an approximation to $-1/x^2$ plus a big constant times an approximation to the delta function.

The Fourier transform of the derivative is obtained by multiplying the Fourier transform of PV $1/x$ by ik . Thus the Fourier transform of $-1/x^2 + c\delta(x)$ is ik times $-\pi i \text{sign}(k)$ which is $\pi|k|$.

This example is interesting, because it looks at first glance that the derivative of PV $1/x$ should be $-1/x^2$, which is negative definite. But the correct answer for the derivative is $-1/x^2 + c\delta(x)$, which is actually positive definite. And in fact its Fourier transform is positive.

10.9 Poisson equation

We begin the study of fundamental solutions of differential equations. These are solutions of the equation $Lu = \delta$, where L is the differential operator, and δ is a point source.

Let us focus on the equation in one space dimension:

$$\left(-\frac{d^2}{dx^2} + m^2\right)u = \delta(x). \quad (10.46)$$

This is an equilibrium equation that balances a source with losses due to diffusion and to dissipation (when $m > 0$). Fourier transform. This gives

$$(k^2 + m^2)\hat{u}(k) = 1. \quad (10.47)$$

The solution is

$$\hat{u}(k) = \frac{1}{k^2 + m^2}. \quad (10.48)$$

There is no problem of division by zero. The inverse Fourier transform is

$$u(x) = \frac{1}{2m}e^{-m|x|}. \quad (10.49)$$

This is the only solution that is a tempered distribution. (The solutions of the homogeneous equation all grow exponentially.)

What happens when $m = 0$? This is more subtle. The equation is

$$-\frac{d^2}{dx^2}u = \delta(x). \quad (10.50)$$

Fourier transform. This gives

$$k^2\hat{u}(k) = 1. \quad (10.51)$$

Now there is a real question about division by zero. Furthermore, the homogeneous equation has solutions that are tempered distributions, namely linear combinations of $\delta(k)$ and $\delta'(k)$. The final result is that the inhomogeneous equation does have a tempered distribution solution, but it needs careful definition. The solution is

$$\hat{u}(k) = \frac{1}{k^2} + \infty\delta(k). \quad (10.52)$$

This may be thought of as the derivative of $-\text{PV } 1/k$. The inverse Fourier transform of $\text{PV } 1/k$ is $(1/2)i\text{sign}(x)$. So the inverse Fourier transform of $-d/dk \text{PV } 1/k$ is $-(-ix)(1/2)i\text{sign}(x) = -(1/2)|x|$. Thus

$$u(x) = -\frac{1}{2}|x| \quad (10.53)$$

is a solution of the inhomogeneous equation. The solutions of the homogeneous equation are linear combinations of 1 and x . None of these solutions are a good description of diffusive equilibrium. In fact, in one dimension there is no diffusive equilibrium.

Remark: The next case that is simple to compute and of practical importance is the equation in dimension 3. This is

$$(-\nabla^2 + m^2)u = \delta(x). \quad (10.54)$$

This is an equilibrium equation that balances a source with losses due to diffusion and to dissipation (when $m > 0$). Fourier transform. This gives

$$(k^2 + m^2)\hat{u}(k) = 1. \quad (10.55)$$

The solution is

$$\hat{u}(k) = \frac{1}{k^2 + m^2}. \quad (10.56)$$

The inverse Fourier transform in the three dimension case may be computed by going to polar coordinates. It is

$$u(x) = \frac{1}{4\pi|x|}e^{-m|x|}. \quad (10.57)$$

What happens when $m = 0$? The situation is very different in three dimensions. The equation is

$$\nabla^2 u = \delta(x). \quad (10.58)$$

Fourier transform. This gives

$$k^2\hat{u}(k) = 1. \quad (10.59)$$

The inhomogeneous equation has a solution

$$\hat{u}(k) = \frac{1}{k^2}. \quad (10.60)$$

But now this is a locally integrable function. It defines a tempered distribution without any regularization. Thus

$$u(x) = \frac{1}{4\pi|x|} \quad (10.61)$$

is a solution of the inhomogeneous equation. In three dimensions there is diffusive equilibrium. There is so much room that the effect of the source can be completely compensated by diffusion alone.

10.10 Problems

If F and G are distributions, and if at least one of them has compact support, then their convolution $F * G$ is defined by

$$\langle F * G, \phi \rangle = \langle F_x G_y, \phi(x + y) \rangle.$$

This product is commutative. It is also associative if at least two of the three factors have compact support.

1. If F and G are given by locally integrable functions f and g , and at least one has compact support, then $F * G$ is given by a locally integrable function

$$(f * g)(z) = \int_{-\infty}^{\infty} f(x)g(z - x) dx = \int_{-\infty}^{\infty} f(z - y)g(y) dy.$$

2. If G is given by a test function g , then $F * g$ is given by a smooth function

$$(F * g)(z) = \langle F_x, g(z - x) \rangle.$$

3. Calculate the convolution $1 * \delta'$.
4. Calculate the convolution $\delta' * H$, where H is the Heaviside function.
5. Calculate the convolution $(1 * \delta') * H$ and also calculate the convolution $1 * (\delta' * H)$. What does this say about the associative law for convolution?
6. Let L be a constant coefficient linear differential operator. Let u be a distribution that is a fundamental solution, that is, let $Lu = \delta$. Let G be a distribution with compact support. Show that the convolution $F = u * G$ satisfies the equation $LF = G$. Hint: Write $\langle LF, \phi \rangle = \langle F, L^\dagger \phi \rangle$, where L^\dagger is adjoint to L .
7. Take $L = -d^2/dx^2$. Is there a fundamental solution that has support in a bounded interval? Is there a fundamental solution that has support in a semi-infinite interval?

8. Show that $\frac{1}{x^{\frac{1}{3}}}$ is a locally integrable function and thus defines a distribution. Show that its distribution derivative is

$$\frac{d}{dx} \frac{1}{x^{\frac{1}{3}}} = -\frac{1}{3} \frac{1}{x^{\frac{4}{3}}} + c\delta(x), \quad (10.62)$$

where

$$c = \frac{1}{3} \int_{-\infty}^{\infty} \frac{1}{x^{\frac{4}{3}}} dx \quad (10.63)$$

Hint: To make this rigorous, consider $x/(x^2 + \epsilon^2)^{\frac{2}{3}}$.

9. Show that $\frac{1}{|x|^{\frac{1}{3}}}$ is a locally integrable function and thus defines a distribution. Show that its distribution derivative is

$$\frac{d}{dx} \frac{1}{|x|^{\frac{1}{3}}} = -\frac{1}{3} \frac{1}{x^{\frac{4}{3}}} \text{sign}(x). \quad (10.64)$$

The right hand side is not locally integrable. Explain the definition of the right hand side as a distribution. Hint: To make this rigorous, consider $1/(x^2 + \epsilon^2)^{\frac{1}{6}}$.

10. Discuss the contrast between the results in the last two problems. It may help to draw some graphs of the functions that approximate these distributions.

Chapter 11

Appendix: Standard Borel measurable spaces

11.1 Complete separable metric spaces

A measurable space is a set X together with a σ -algebra of measurable sets (or, equivalently, with a σ -algebra of measurable functions). A standard measurable space is one that is isomorphic to one of the following possibilities: 1) A countable set with the σ -algebra of all subsets; 2) The unit interval $[0, 1]$ with the Borel σ -algebra. This chapter treats standard measurable spaces. The emphasis is on the second possibility.

If X is a space with a topology, then X also is equipped with a σ -algebra of subsets, the Borel σ -algebra generated by the topology. This makes X into a measurable space. The purpose of this appendix is to sketch the proof of a remarkable theorem. It says that for every uncountable separable complete metric space the associated measurable space is isomorphic to the measurable space associated with the unit interval $[0, 1]$. In other words, for most practical purposes, there is just one measurable space of interest.

The proof will actually show that for every uncountable separable metric space the associated measurable space is isomorphic to the measurable space associated with the Cantor set. This is the same as the space of all infinite sequences of coin tosses.

The strategy is simple. The first part is to show that for every uncountable separable complete metric space has a Cantor set embedded in it. The second part is to show that every complete separable metric space may be placed inside an infinite dimensional cube. The third part is to place the cube inside a Cantor set. The final step is to argue that if the space has a Cantor set inside and is also inside a Cantor set, then it may be matched up with a Cantor set. This is done by a dynamical systems argument.

At the end of this chapter there is a remark on another striking result, which is an easy consequence of the theorem. This says that if X is an uncountable

complete separable metric space, and μ is a finite non-zero Borel measure with no point masses, then the measure space (X, \mathcal{B}, μ) is isomorphic to Lebesgue measure on some closed and bounded interval $[0, M]$ of real numbers, with $0 < M < +\infty$. There is a corresponding result for a σ -finite measure, where the interval $[0, +\infty)$ is also allowed. In other words, for most practical purposes there is just one class of continuous σ -finite measure spaces of interest, classified by total mass M .

11.2 Embedding a Cantor set

Proposition 11.1 *A Borel subset of a Borel set (with the induced topology) is a Borel set.*

It is certainly possible to have an uncountable complete metric space with no Cantor set inside. The obvious example is an uncountable discrete space. However with the additional hypothesis of separability this kind of example is ruled out.

Proposition 11.2 *Let B be a complete separable metric space. Suppose that B is uncountable. Then there exists a compact subset $A \subset B$ that is homeomorphic to $2^{\mathbb{N}^+}$.*

Proof: Let C be the set of all y in B such that there exists a countable open subset U with $y \in U$. Since B is separable, we can take the open sets to belong to a countable base. This implies that C is itself a countable set.

Let $A = B \setminus C$. Since B is uncountable, it follows that A is also uncountable. Each y in A has the property that each open set U with $y \in U$ is uncountable.

The next task is to construct for each sequence $\omega_1, \dots, \omega_m$ a corresponding closed subset $F_{\omega_1, \dots, \omega_m}$ of A . Each closed set has a non-empty interior and has diameter at most $1/m$. For different sequences the corresponding closed subsets are disjoint. The closed sets decrease as the sequence is extended. This is done inductively. Start with A . Since it is uncountable, it has at least two points. Construct the first two sets F_0 and F_1 to satisfy the desired properties. Say that the closed subset $F_{\omega_1, \dots, \omega_m}$ has been defined. Since it non-empty interior, there are uncountably many points in it. Take two points. About each of these points construct subsets $F_{\omega_1, \dots, \omega_m, 0}$ and $F_{\omega_1, \dots, \omega_m, 1}$ with the desired properties.

These closed subsets of A are also non-empty subsets of the complete metric space B with decreasing diameter. By the completeness of B , for each infinite sequence ω in $2^{\mathbb{N}^+}$ the intersection of the corresponding sequence of closed sets has a single element $g(\omega) \in B$. It is not hard to see that $g : 2^{\mathbb{N}^+} \rightarrow B$ is a continuous injection. Since $2^{\mathbb{N}^+}$ is compact, it is a homeomorphism onto a compact subset A of B . \square

11.3 Placement in a cube

Proposition 11.3 *Let B be a complete separable metric space. Then there exists a Borel subset $S \subset [0, 1]^{\mathbf{N}^+}$ such that B is homeomorphic to S . In particular, B is isomorphic to a Borel subset of a compact metric space.*

Proof: Let B be a complete separable metric space. There exists a metric on B with values bounded by one such that the map between the two metric spaces is uniformly continuous. So we may as well assume that B is a complete separable metric space with metric d bounded by one.

Since B is separable, there is a sequence $s : \mathbf{N}_+ \rightarrow B$ that is an injection with dense range. Let $I = [0, 1]$ be the unit interval, and consider the space $I^{\mathbf{N}^+}$ with the product metric d_p . Define a map $f : B \rightarrow I^{\mathbf{N}^+}$ by $f(x)_n = d(x, s_n)$. This is a homeomorphism f from B to a subset $S \subset I^{\mathbf{N}^+}$. It need not be uniformly continuous. The subset S is a Borel subset of $I^{\mathbf{N}^+}$. To prove this, consider S with metric d as a subset of its closure T in $I^{\mathbf{N}^+}$ with metric d_p . Then S is a complete metric space with metric d , and T is a complete metric space with metric d_p , and the two metrics define the same topology on S .

For each $n = 1, 2, 3, \dots$ define

$$U_n = \{x \in T \mid \exists \delta > 0 \forall y, z \in S (d_p(y, x) < \delta, d_p(z, x) < \delta \Rightarrow d(y, z) < 1/n)\}. \quad (11.1)$$

This is an open set in T . To see this, suppose that x is in U_n . Then there exists $\delta > 0$ such that $d_p(y, x) < \delta, d_p(z, x) < \delta \Rightarrow d(y, z) < 1/n$. Now consider nearby w with $d_p(w, x) < \delta/2$. Then $d_p(y, w) < \delta/2, d_p(z, w) < \delta/2 \Rightarrow d(y, z) < 1/n$. Thus w is in U_n . Furthermore, $S \subset U_n$. This follows from the fact that d and d_p define the same topology on S .

The last step is to show that $\bigcap_n U_n = S$. Suppose that x is in each U_n . Since S is dense in T , we may take a sequence x_m in S that converges in the d_p metric to x . Consider an arbitrary n . Since x is in U_n , there exists $\delta > 0$ such that y, z in S with $d_p(y, x) < \delta$ and $d_p(z, x) < \delta$ imply that $d(y, z) < \delta$. Pick m and k so large that $d_p(x_m, x) < \delta$ and $d_p(x_k, x) < \delta$. Then $d(x_m, x_k) < 1/n$. This is enough to prove that the sequence x_m is a Cauchy sequence with respect to d . Since S is complete, x_m must converge in the d metric to an element of S . So x must be in S .

The preceding arguments complete the proof that S is a Borel subset of $I^{\mathbf{N}^+}$. Recall that T is a closed subset of $I^{\mathbf{N}^+}$. Each open subset U_n of T is of the form $U_n = T \cap V_n$, where V_n is an open subset of $I^{\mathbf{N}^+}$. Thus $S = \bigcap_n U_n = T \cap \bigcap_n V_n$. This is enough to show that S is a Borel set. \square

11.4 A cube inside a Cantor set

Lemma 11.4 *There is a continuous bijective function from a Borel subset Y of $2^{\mathbf{N}^+}$ onto $[0, 1]^{\mathbf{N}^+}$ with a Borel measurable inverse. In particular, $[0, 1]^{\mathbf{N}^+}$ is Borel isomorphic with the Borel subset Y of $2^{\mathbf{N}^+}$.*

Proof: There is a continuous function from $2^{\mathbf{N}^+}$ onto $[0, 1]$ that is injective on a Borel subset W of $2^{\mathbf{N}^+}$. The inverse of this continuous function is a measurable function from $[0, 1]$ to the Borel subset W .

The continuous function from $2^{\mathbf{N}^+}$ to $[0, 1]$ defines a continuous function from $2^{\mathbf{N}^+ \times \mathbf{N}^+}$ to $[0, 1]^{\mathbf{N}^+}$. It maps $Y' = W^{\mathbf{N}^+}$ injectively onto $[0, 1]^{\mathbf{N}^+}$ and has a measurable inverse. Since there is a bijection of $\mathbf{N}_+ \times \mathbf{N}_+$ with \mathbf{N}_+ , there is a continuous bijection of $2^{\mathbf{N}^+ \times \mathbf{N}^+}$ with $2^{\mathbf{N}^+}$. This gives a continuous bijection of $Y' \subset 2^{\mathbf{N}^+ \times \mathbf{N}^+}$ with $Y \subset 2^{\mathbf{N}^+}$. \square

11.5 A dynamical system

Lemma 11.5 *Suppose that C is a complete separable metric space. Suppose that $A \subset B \subset C$ are Borel subsets. Suppose that C is Borel isomorphic to A . Then C is Borel isomorphic to B .*

Proof: Let $\phi : C \rightarrow C$ is a Borel transformation that gives a Borel isomorphism of C with A . Think of ϕ as a dynamical system on C . Since it is an injection, it decomposes the space into \mathbf{N} orbits and a part on which it is a bijection. Let $D = C \setminus A$. Then D consists of the starting points for the \mathbf{N} orbits of ϕ . The union of all the \mathbf{N} orbits then consists of the union of the Borel sets $\phi^n[D]$ for $n = 0, 1, 2, 3, \dots$. Let E be the set consisting of the part of C on which ϕ is a bijection. The points in the Borel set E consist of the intersection of all the $\phi^n[A]$ for $n = 0, 1, 2, 3, \dots$. That is, each point in E comes from an arbitrarily remote past.

Let $F = C \setminus B$ and $G = B \setminus A$. Then the union of the \mathbf{N} orbits consist of the union of all the $\phi^n[F]$ together with all the $\phi^n[G]$. Let ψ be a dynamical system that agrees with ϕ on the union of all the $\phi^n[F]$ and that is the identity on the union of all the $\phi^n[G]$. One can also make ψ be the identity on E . Then the \mathbf{N} orbits of ψ start on F , and the range of ψ is $B = C \setminus F$. Thus ψ gives a Borel isomorphism of C with B . \square

11.6 The unique Borel structure

Theorem 11.6 *Let B and B' be uncountable complete separable metric spaces. Then B is Borel isomorphic to B' .*

Proof: Let B be an uncountable complete separable metric space. Then there exists a compact subset $A \subset B$ that homeomorphic to $2^{\mathbf{N}^+}$. Furthermore, B is isomorphic to a Borel subset S of $[0, 1]^{\mathbf{N}^+}$. Since there is a Borel isomorphism of $[0, 1]^{\mathbf{N}^+}$ with a Borel subset Y of $2^{\mathbf{N}^+}$, there is a Borel isomorphism of S with a Borel subset Z of $2^{\mathbf{N}^+}$.

Thus we may consider $A \subset B \subset C$, where A and C are each Borel isomorphic to the coin tossing space $2^{\mathbf{N}^+}$. The last lemma shows that B must also be Borel isomorphic to the coin tossing space. The same reasoning shows

that B' is also Borel isomorphic to the coin tossing space. Thus B and B' are Borel isomorphic to each other. \square

11.7 The unique measure structure with given mass

Theorem 11.7 *Let μ be a finite Borel measure on an uncountable complete separable metric space X . Let $M = \mu(X)$ be the total mass and suppose that $M > 0$. Suppose also there are no one point sets with non-zero measure. Then the measure space (X, \mathcal{B}, μ) is isomorphic to the measure space $([0, M], \mathcal{B}, \lambda)$.*

This is not a difficult result, since one can reduce the problem to the analysis of a finite measure on $[0, 1]$. If the measure has no point masses, then it is given by an increasing continuous function from $[0, 1]$ to $[0, M]$. This maps the measure to Lebesgue measure on $[0, M]$. The only problem that there may be intervals where the function is constant, so it is not bijective. This can be fixed. See H. L. Royden, *Real Analysis*, 3rd edition, Macmillan, New York, 1988, Chapter 15, Section 5 for the detailed proof.

The conclusion is that for continuous finite measure spaces of this type the only invariant under isomorphism is the total mass. Otherwise all such measure spaces look the same. There is a corresponding theorem for σ -finite measures. In this situation there is the possibility of infinite total mass, corresponding to Lebesgue measure on the interval $[0, +\infty)$.

In the end, the measure spaces of practical interest are isomorphic to a countable set with point measures, an interval with Lebesgue measure, or a disjoint union of the two.