Notes for MATH 313 — Functional Analysis
(Winter 2018)

SIYAN DANIEL LI

These are live-TEX'd notes for a course taught at the University of Chicago in Winter 2018 by Professor Charlie Smart. Any errors are attributed to the note-taker. If you find any such errors or have comments at large, don't hesitate to contact said note-taker at `lidansiyan@gmail.com`.

1 January 3, 2018

Well, nobody is talking, so it must be time to start. This is Analysis II (functional analysis), and I'm Charlie Smart, one of the faculty here. There's some administrative stuff, but it's mostly just to read the syllabus on Canvas. The grading will be 40% homework, 30% midterm, and 40% final. The midterm is... sometime in February, probably? I should *also* read the syllabus.

There are no official textbooks, but there are at least three good options you could follow for the course:

- The main notes that I will roughly follow are notes by Bühler–Salamon, two fellows at ETH Zürich,
- The traditional textbook for Analysis II is by Brezis,
- Another good one is by Lax. It's only not our course text because I think it's out of print.

Now let's get to the course itself. What *is* functional analysis? It's the use of topology to tame infinite-dimensional linear algebra. We all know finite-dimensional linear algebra, point-set topology, and metric spaces, but there isn't much we can prove for infinite-dimensional linear algebra a priori. We would like our bread-and-butter linear algebra (over $\mathbb{R}$ or $\mathbb{C}$ or something like that) to work in infinite dimensions, for various applications, so that's the motivation for the course.

Let's consider a certain problem in which this comes up: the *vibrating membrane*. I got this impression from action movies that everything has *resonant frequencies*, like the Tacoma Narrows Bridge or like pieces of metal (which we hear by hitting them with hammers). But what *are* resonant frequencies? Consider an open bounded subset $\Omega \subset \mathbb{R}^2$, and consider a map $u : \mathbb{R} \times \overline{\Omega} \longrightarrow \mathbb{R}$ which is supposed to send time and space to displacement. Say it satisfies a differential equation $\partial_t^2 u = (\partial_x^2 + \partial_y^2)u$ on $\mathbb{R} \times \Omega$, which is saying that the membrane is accelerating at every point proportionally to the curvature there. This seems intuitive, but let's not say more about why this is the case—we're not in a physics class, after all. Let's also impose $u = 0$ on $\mathbb{R} \times \partial\Omega$, which corresponds to fixing the boundary of our membrane.

It'll take us several weeks to solve this in general, but we can already perform some special cases. Say that $\Omega = (0, 1)^2$ is a square. Then it's an undergraduate exercise to find a separation of variables solution. Suppose that $u(t, x, y)$ is of the form $v(t)w(x, y)$. Then our equations become

$$\begin{aligned} \partial_t^2 v &= \lambda v \text{ in } \mathbb{R}, \\ (\partial_x^2 + \partial_y^2)w &= \lambda w \text{ in } \Omega, \\ w &= 0 \text{ in } \partial\Omega, \\ \lambda &\in \mathbb{R}. \end{aligned}$$

We know how to find some solutions to these equations:

$$u_{n,m}(t,x,y) \coloneqq \sin(\pi(m+n)t)\sin(\pi nx)\sin(\pi my) \text{ or}$$
$$\widetilde{u}_{n,m}(t,x,y) \coloneqq \cos(\pi(m+n)t)\underbrace{\sin(\pi nx)\sin(\pi my)}_{w_{n,m}},$$

and Fourier analysis indicates that *all* solutions are linear combinations of the above.

The next question is: does this method work if Ω isn't a square? In general the answer is no, though it *does* work in special cases. For example, if $\Omega = B$ is a disk, we can use polar coordinates to give solutions in the form of Bessel functions of the first kind. If we go to three dimensions are take Ω to be a sphere, then we can use spherical harmonics. But in general…? These special results were known classically, even to Renaissance mathematicians.

In our above work, we started with a PDE from physics, and then we proceeded to separate variables and solve the separated equations. Abstractly, what are we doing? What's really happening is that we have an operator $L : C_c^\infty(\Omega) \longrightarrow C_c^\infty(\Omega)$ given by, say, $Lw \coloneqq (\partial_x^2 + \partial_y^2)w$, and the $w_{n,m}$ form a basis of eigenvectors. (Nevermind for now that the $w_{n,m}$ aren't actually compactly supported—we'll deal with that later.)

In the same way that the Greeks didn't know about e or arbitrary real numbers (as opposed to special examples like π or $\sqrt{2}$), which are obtained from rational numbers via a completion process, Renaissance mathematicians didn't know about larger function spaces, which are also obtained from completing with respect to some metric. I'm pretty sure much of what I'm saying is wrong—I only thought about this over the summer, when I was trying to think of history to present in my first lecture, and I haven't thought about it since! However, you can consult an article entitled "The Establishment of Functional Analysis" by Birkhoff–Kreyszig.

That was all history—what is functional analysis itself? I'll assume Hahn–Banach, Hilbert spaces, and Banach spaces. We'll start with a possibly new perspective on Lebesgue integration.

1.1 Definition. A *metric space* is a set X and a function $d : X \times S \longrightarrow \mathbb{R}$ such that

(1) $d(x,y) = 0$ if and only if $x = y$,

(2) $d(x,y) = d(y,x)$,

(3) $d(x,z) \leq d(x,y) + d(y,z)$.

This endows X with a topology, generated by the *metric balls*

$$B(x,r) \coloneqq \{y \in X : d(x,y) < r\}.$$

We say that X is *complete* if

(4) every Cauchy sequence is convergent.[1]

1.2 Theorem. *For every metric space X, there is a complete metric space $\overline{X}$ and isometric $i : X \longrightarrow \overline{X}$ such that, for any complete metric space $\overline{Y}$ and continuous map $f : X \longrightarrow \overline{Y}$, there is a unique continuous $\overline{f} : \overline{X} \longrightarrow \overline{Y}$ for which $f = \overline{f} \circ i$.*

We're only recalling this in metric space generality to remind ourselves that completions are purely metric space theoretic. Now let's add more algebraic structure.

[1] The originally stated criterion was that if $B(x_1,r_1) \supseteq B(x_2,r_2) \supseteq \cdots$, then $\bigcap_n B(x_n,r_n)$ is nonempty. However, this original criterion is actually stronger than spherical completeness, and there are certainly complete metric spaces that are not spherically complete.

1.3 Definition. A *normed vector space* is a (real) vector space V together with a function $\|\cdot\| : V \longrightarrow \mathbb{R}$ such that

(1) $\|x\| = 0$ if and only if $x = 0$,

(2) $\|tx\| = |t|\|x\|$,

(3) $\|x + y\| \leq \|x\| + \|y\|$.

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

For the most part, I want to stick to using real vector spaces in this course. When we discuss the spectral theorem for compact operators, we'll want to use complex vector spaces instead. This V forms a metric space via setting $d(x, y) := \|x - y\|$ and therefore has a topology and notion of completeness. Recall the following result, which is on your homework.

1.4 Theorem. *If $f : V \longrightarrow W$ is linear, then f is continuous if and only if f is bounded.*

The completion also satisfies a universal property with respect to normed vector spaces, but it doesn't follow directly from the corresponding metric space result. The following is also on your homework.

1.5 Theorem. *If V is a normed vector space, then there is a Banach space $\overline{V}$ and an isometric linear map $i : V \longrightarrow \overline{V}$ such that, if $\overline{W}$ is a Banach space and $f : V \longrightarrow \overline{W}$ is a bounded linear map, then $f = \overline{F} \circ i$ for a unique bounded linear map $\overline{f} : \overline{V} \longrightarrow \overline{W}$.*

1.6 Remark. Almost (but not all!) of our examples of Banach spaces will come from Theorem 1.5. We will usually take a function space, define some norm, and then take its completion with respect to said norm.

1.7 Example. Take the Lebesgue measure μ on $\mathbb{R}^d$, let p be in $[1, \infty)$, and consider $L^p(\mathbb{R}^d, \mu)$. This is isometric to the completion of $C_c^\infty(\mathbb{R}^d)$ with respect to the L^p-norm. I think you know this from Analysis I, but I've put it as a homework problem.

This has applications to integration. Taking $p = 1$ and using Theorem 1.5 shows that Lebesgue integration is the completion of Riemann integration from $C_c^\infty(\mathbb{R}^d)$ to $L^1(\mathbb{R}^d, \mu)$. This indicates that Lebesgue integration is quite natural—we constructed it originally to solve problems, but this shows that it's perhaps inevitable. Despite the fact that $L^1(\mathbb{R}^d, \mu)$ arises as a completion, its explicit description of equivalence classes of integrable functions remains useful, since it's far more concrete and simple than the Cauchy sequence completion description.

Next time, we'll start *really* talking about functional analysis, rather than just review and points of view.

2 January 5, 2018

Last time, we discussed completions of normed vector spaces. Recall from the toy example of completing $\mathbb{Q}$ to $\mathbb{R}$ that we take completions to obtain theorems which required "filling in holes" (like the intermediate value theorem), and a similar heuristic applies to spaces of functions. For example, we need these properties to solve extremal problems in function spaces, such as perform the calculus of variations. In order to ensure the existence of "minimal" points given certain conditions, you need completeness.

Today, we'll discuss dual spaces and alternative topologies. For the latter, our motivation is that we often want certain subspaces to be compact, which is tricky and in general not true for the usual topology on infinite-dimensional spaces. Our strategy is to rectify this by cheating and taking a coarser topology in which our desired subspaces *are* compact. For the problems we face, we can often pass back to our original topologies via some argument. Much of today will likely be review.

2.1 Definition. Let X and Y be normed vector spaces. Then $\mathscr{L}(X,Y)$ denotes the space of bounded linear maps $f : X \longrightarrow Y$ equipped with the *operator norm*

$$\|f\|_{\mathscr{L}(X,Y)} \coloneqq \sup\{\|f(x)\|_Y : \|x\|_X \leq 1\}.$$

This satisfies the triangle inequality due to the properties of sup.

2.2 Theorem. *If Y is Banach, then $\mathscr{L}(X,Y)$ is Banach.*

Proof. Take a Cauchy sequence of operators $\{f_n\}_n$ in $\mathscr{L}(X,Y)$. For any x in X, we have

$$\|f_n(x) - f_m(y)\|_Y \leq \|f_n - f_m\|_{\mathscr{L}(X,Y)}\|x\|_X,$$

so $\{f_n(x)\}_n$ is a Cauchy sequence in Y. The completeness of Y allows us to define the limit

$$f(x) \coloneqq \lim_{n\to\infty} f_n(x),$$

and we can show that $f_n \to f$ in the operator norm. □

In particular, our operator spaces allow us to define dual spaces.

2.3 Definition. The dual space of X is defined to be $X^* \coloneqq \mathscr{L}(X,\mathbb{R})$.

We're going to spend a lot of time thinking about duals (and duals of duals), and in fact we'll use them to obtain our aforementioned alternative topologies. Here's a useful fact about dual spaces.

2.4 Theorem. *For any x in X, we have*

$$\|x\|_X = \max_{\substack{f\in X^*\\ \|f\|_{X^*}\leq 1}} f(x).$$

Proof. This follows from the Hahn–Banach theorem. First, we immediately get

$$\|x\|_X \geq \sup_{\substack{f\in X^*\\ \|f\|_{X^*}\leq 1}} f(x)$$

from the definition of the operator norm. Next, set $Y \coloneqq \{tx : t \in \mathbb{R}\} \subseteq X$. By scaling, we may assume that $\|x\|_X = 1$. Let $\ell : Y \longrightarrow \mathbb{R}$ be defined by sending $tx \mapsto t$. Because this linear functional satisfies $\ell(y) \leq \|y\|_X$ for y in Y, the Hahn–Banach theorem yields an extension $\ell : X \longrightarrow \mathbb{R}$ satisfying $\ell(y) \leq \|y\|_X$ for all y in Y. Therefore $\|\ell\|_{X^*} \leq 1$, and the fact that $\ell(x) = 1$ shows that we attain equality in the above inequality (and that the supremum is actually attained, so we have a maximum). □

2.5 Definition. Write $J : X \longrightarrow X^{**}$ for the canonical evaluation map.

2.6 Exercise. Check that J is an isometric injection.[2]

The next concept will appear over and over again, so we give it a name:

2.7 Definition. We say that X is *reflexive* if J is bijection (equivalently, by the above exercise, if J is a surjection).

[2] This follows immediately from Theorem 2.4.

2.8 Example. If H is a Hilbert space, that is, if it is Banach and if the norm comes from some inner product $\langle\cdot,\cdot\rangle_H$ via $\|x\|_H^2 = \langle x, x\rangle_H$, then the Riesz representation theorem says that H is isometric to H^* via the map $x \mapsto \Lambda_x$, where $\Lambda_x(y) := \langle x, y\rangle_H$.

Hilbert spaces are the prototypical example of reflexive Banach spaces, and they form a motivating example for reflexivity in general. Throughout the course, we'll see that many proofs are much easier for Hilbert spaces. We often end up proving these theorems in greater generality than the Hilbert setting by retooling the proofs to only use reflexivity rather than $H = H^*$.

Anyways, let's continue with examples of normed vector spaces.

2.9 Example. If (X, μ) is a finite measure space, and we have $1 \le p < \infty$ and $1 < q \le \infty$ satisfying $\frac{1}{p} + \frac{1}{q} = 1$, then $L^q(X,\mu)$ is isometric to $L^p(X,\mu)^*$ via the map $g \mapsto \Lambda_g$, where $\Lambda_g(f) := \int_X \mathrm{d}\mu\, fg$. This is a consequence of Hölder (which shows that $\Lambda.$ is injective) and Radon–Nikodym (which shows that $\Lambda.$ is surjective).

In general, we have $L^1(X,\mu) \subsetneq (L^\infty(X,\mu))^* = L^1(X,\mu)^{**}$. It's an exercise to show that the latter isn't quite a space of countably additive measures but rather of finitely additive ones, which aren't too useful—it's a disaster. That's why this isn't usually covered in texts.

2.10 Example. If K is a compact metric space, then let $C(K)$ be the space of continuous functions on K equipped with the sup norm. We can show that $C(K)^*$ is the space of signed Borel measures on K with the total variation norm (which was *not* covered last quarter, so I'll probably prove it sometime later).

The following reflexivity statement follows from our discussion in Example 2.9.

2.11 Theorem. *For $1 < p < \infty$ and a finite measure space (X, μ), the space $L^p(X,\mu)$ is reflexive.*

Next, let's turn towards building the *weak* and *weak-$*$* topologies. Our usual topology on X comes from its norm. We can build a new topology from a map to another topological space Y by taking the coarsest topology for which all the inverse images of open sets of Y are now open in X. This is the coarsest topology for which said map is continuous.[3]

More explicitly, if S is a set, and $\mathscr{F}$ is a collection of maps from S to a topological space Y, then the coarsest topology for which all the maps in $\mathscr{F}$ are continuous has a sub-basis given by

$$\{f^{-1}(U) : f \in \mathscr{F},\, U \subseteq Y \text{ open}\}.$$

Denote this topology using $\sigma(S, \mathscr{F})$.

2.12 Definition. The *weak topology* on X is $\sigma(X, X^*)$, that is, the coarsest topology for which all *strongly continuous* linear functionals are continuous. The *weak-$*$ topology* on X^* is $\sigma(X^*, X)$, where we view $X \hookrightarrow X^{**}$ via J.

2.13 Remark. Because the maps in X^* (respectively X) were already continuous in the strong topology, the weak (respectively weak-$*$) topologies are coarser than the norm topology. Similarly, because $X \subseteq X^{**}$, we see that the weak-$*$ topology is coarser than the weak topology on X^*.

In the coming days, we'll prove that the unit ball is compact in the weak-$*$ topology. However, what we really care about is the weak topology. Fortunately, we'll also prove that X is reflexive in some situations, so in those settings we can upgrade this compactness to the weak topology.

2.14 Theorem. *The weak and weak-$*$ topologies are Hausdorff.*

[3] I use map in the singular here and maps in the plural below, but by taking the product of many maps with domain X, the distinction between them evaporates.

Proof. For the weak-$*$ topology $\sigma(X^*, X)$, take f and g in X^*, and suppose that $f \neq g$. Reorder f and g such that there exists an x in X with $f(x) < g(x)$. Take two real numbers a and b such that $f(x) < a < b < g(x)$. Then f and g lie in the disjoint open subsets $\{h \in X^* : h(x) < a\}$ and $\{h \in X^* : h(x) > b\}$, respectively.

As for the weak topology $\sigma(X, X^*)$, this is harder and requires a result from last quarter. Take $x \neq y$ in X. Because $\{x\}$ and $\{y\}$ are closed and disjoint, the geometric Hahn–Banach theorem yields a bounded linear map f in X^* and $a < b$ in $\mathbb{R}$ such that $f(x) < a < b < f(y)$. At this point, the same construction as above yields the desired disjoint neighborhoods. □

We conclude today by pointing out that we can put the weak, weak-$*$, and strong topologies in the same category by using *locally convex topological vector spaces*. Their topologies are generated by families of seminorms rather than norms, which makes sense because linear functionals really feel like seminorms, as they seem to measure "distance" (albeit with a sign, which is lost after taking absolute values) but can also have nontrivial kernel.

3 January 8, 2018

Let's continue talking about the weak and weak-$*$ topologies today. We'll be discussing compactness this time, via the *Banach–Alaoglu theorem*, which I think was *not* covered last quarter. Recall that the weak topology $\sigma(X, X^*)$ is the coarsest topology for which every f in X^* is continuous, whereas the strong topology is the usual metric space topology (which we can view as the coarsest topology for which all translates of the norm are continuous).

3.1 Remark. If X is finite-dimensional, I claim that the weak and strong topologies coincide. As we already have $\sigma(X, X^*) \subseteq \sigma(X, x \mapsto \|x + s\|)$, it suffices to show containment in the other direction. It's enough to show that every open ball centered at a point contains a $\sigma(X, X^*)$-open neighborhood of said point.

Proving this isn't too hard. Choose a basis $e_1, \dots, e_n$ for X, and choose f_k in X^* such that $x = \sum_{k=1}^n f_k(x) e_k$ for all x in X. Then the triangle inequality yields

$$\|x\| \leq \sum_{k=1}^n \|f_k(x) e_k\| \leq \underbrace{\sum_{k=1}^n \|f_k\|_{X^*} \cdot \|e_k\|_X}_{M} \cdot \|x\|.$$

Applying the above to $x - y$ shows that

$$x \in \bigcap_{k=1}^n \{y : |f_k(y - x)| < M^{-1}\varepsilon\} \subseteq B(x, \varepsilon),$$

giving us our desired neighborhood. Therefore the weak topology is only interesting in infinite dimensions.

Let us now recall the Riesz lemma, which was used to prove Hahn–Banach last quarter.

3.2 Lemma (Riesz)**.** *If X is a normed vector space, $Y \subsetneq X$ is a closed linear subspace, and δ is a real number in $(0, 1)$, then there exists x in X such that $\|x\| = 1$ and*

$$\inf_{y \in Y} \|x - y\| > \delta.$$

If you haven't seen the Riesz lemma before, you should prove it as a fun exercise.

3.3 Corollary. *If X is a normed vector space, then the set $\overline{B}_1 := \{x \in X : \|x\| \leq 1\}$ is compact if and only if $\dim X < \infty$.*

Proof. This is compact in finite dimensions by, say, Heine–Borel. In infinite dimensions, one can use the Riesz lemma to find a sequence in $\overline{B}_1$ that has no convergent subsequence. $\square$

Let's continue by exploring the difference between the strong and weak topologies.

3.4 Example. If X is a normed vector space with $\dim X = \infty$, then $S_1 := \{x : \|x\| = 1\}$ is not weakly closed (in contrast, it's immediately seen to be strongly closed). In fact, we can prove a *stronger*[4] statement: the weak closure of S_1 is equal to $\overline{B}_1$. This shall follow from the these two results:

(1) $\overline{B}_1$ is weakly closed,

(2) B_1 is contained in the weak closure $\overline{S}_1^{\sigma(X,X^*)}$ of S_1.

For (2), it is enough to show that for all x in B_1 and weakly open neighborhoods $V \in \sigma(X, X^*)$ of x, the intersection $V \cap S_1$ is nonempty. To see this, it suffices to check it for $\|x\| < 1$ and for V in a base of $\sigma(X, X^*)$, so suppose that

$$V = \bigcap_{k=1}^{n} \{y : |f_k(y - x)| < \varepsilon\}$$

for some $f_1, \ldots, f_n$ in X^* and positive ε. Note that $\ker f_k$ has codimension at most 1 (it could have codimension 0, as f_k could be 0), so $\bigcap_{k=1}^n \ker f_k$ has codimension at most n. Because $\dim X = \infty$, there exists a nonzero z such that $f_k(z) = 0$ for all $1 \leq k \leq n$. Therefore $V = tz + V$ for any t in $\mathbb{R}$. If we choose t such that $\|x + tz\| = 1$, then we see that $x + tz$ lies in $V \cap S_1$.

For (1), it follows from Theorem 3.5, as $\overline{B}_1$ is convex and strongly closed.

3.5 Theorem. *If X is a normed vector space and $C \subseteq X$ is convex, then C is strongly closed if and only if C is weakly closed.*

Proof. As the weak topology is coarser, being weakly closed certainly implies being strongly closed. Conversely, suppose that C is strongly closed. It shall suffice to show that $X \smallsetminus C$ is weakly open, that is, every x in $X \smallsetminus C$ has a weak neighborhood in $X \smallsetminus C$. Since C is convex and the singleton x is compact, Hahn–Banach yields an f in X^* and a and b in $\mathbb{R}$ such that

$$f(x) > b > a > f(y)$$

for all y in C. Then $V := \{f : f(y) > b\}$ is a weakly open subset for which $V \cap C = \varnothing$, as desired. $\square$

Theorem 3.5 shows that weak and strong closedness coincide for convex subsets, while Example 3.4 shows that they can differ wildly for non-convex subsets. However, Theorem 3.6 shall show that, in some sense, this is the only discrepancy.

3.6 Theorem (Mazur)**.** *If $x_n \to x$ weakly, then there exists a sequence $\{y_n\}_n$ in the convex hull of the x_n such that $y_n \to x$ strongly.*

3.7 Remark. In the course, I shall usually denote the topology in which a sequence $x_n \to x$ converges using words. However, the traditional notation of the field is to write $x_n \rightharpoonup x$ for weak convergence and $x_n \overset{*}{\rightharpoonup} x$ for weak-$*$ convergence.

I shall now prove the *Banach–Alaoglu* theorem.

[4] Get it?

3.8 Theorem (Banach–Alaoglu)**.** *The subset*

$$\overline{B}_1^* \coloneqq \{f \in X^* : \|f\|_{X^*} \leq 1\}$$

is compact in the weak-$*$ topology $\sigma(X^*, X)$.

Banach–Alaoglu was the original motivation for Tychonoff's theorem—the former is just a killer app of the latter. While the proof of Tychonoff's theorem requires some real work and depends on the axiom of choice, the proof of Banach–Alaoglu is extremely soft and is essentially a direct consequence of Tychonoff. We just have to be careful with which category we're working in.

Proof. We break it down into three abstract steps:

Step 1. Identify $X^* = \{f : X \longrightarrow \mathbb{R} \text{ bounded linear}\}$ as a subset of $\mathbb{R}^X$.

Step 2. I claim that $\sigma(X^*, X)$ is the subspace topology on X^* inherited from the product topology on $\mathbb{R}^X$. Now the product topology is the coarsest topology for which the projection maps are continuous, which in the case of X^* corresponds to evaluation maps.

Step 3. Write $\overline{B}_1^* = K_1 \cap K_2 \cap K_3$, where

$$\begin{aligned} K_1 &\coloneqq \{f \in \mathbb{R}^X : \text{ for all } x, y \in X,\ f(x+y) = f(x) + f(y)\}, \\ K_2 &\coloneqq \{f \in \mathbb{R}^X : \text{ for all } x \in X \text{ and } t \in \mathbb{R},\ f(tx) = tf(x)\}, \\ K_3 &\coloneqq \{f \in \mathbb{R}^X : \text{ for all } x \in X,\ |f(x)| \leq \|x\|\}. \end{aligned}$$

Note that the K_i are cut out by closed conditions and hence are closed subsets of $\mathbb{R}^X$. Furthermore, note that K_3 equals

$$\prod_{x \in X} [-\|x\|, \|x\|] \subseteq \mathbb{R}^X,$$

which is compact by Tychonoff's theorem.

Now we're done, because we've exhibited $\overline{B}_1^*$ as a closed subspace of a compact space. □

This proof is incredibly abstract and non-constructive, as it must be since its result is incredibly general. There's also an interesting alternative proof for separable X. We begin by recalling what separability means.

3.9 Definition. A topological space X is *separable* if it contains a countable dense subset.

Many proofs are easier in the separable case, and in this quarter many of our results will begin with a treatment of the separable special case.

3.10 Remark. When X is separable, there is a more constructive proof of (a slightly stronger version of) Banach–Alaoglu. As a sketch, in this setting you only use your countable dense subset of X for your evaluation maps, and then you dovetail to get up to all of X, by density. The whole proof can be found in any of the recommended course textbooks (as well as in Theorem 4.1).

4 January 10, 2018

I realized that I want to do the proof of separable Banach–Alaoglu, because I want to apply this stronger result in the context of ergodic theory. We're doing this just to see *some* sort of application of the abstract machinery we've been developing.

4.1 Theorem (Separable Banach–Alaoglu)**.** *Let X be a separable normed vector space. Then the subset $\overline{B}_1^*$ is weak-$*$ sequentially compact.*

Proof. Given sequence $\{f_n\}_n$ in X^* for which $\|f_n\|_{X^*}$ is bounded by $K < \infty$, we want to find a weak-$*$ convergent subsequence $\{f'_n\}_n$. For this, begin by choosing a dense countable subset $\{x_m\}_m$ of X. Now

$$|f_n(x_m)| \leq \|f_n\|_{X^*}\|x_m\|_X \leq K\|x_m\|_X.$$

By diagonalization and the sequential compactness of closed intervals in $\mathbb{R}$, we can choose a subsequence $\{f'_n\}_n$ of $\{f_n\}_n$ and real numbers a_m such that $f'_n(x_m) \to a_m$ as $n \to \infty$ for all m. The density of $\{x_m\}_m$ shows that $\{f'_n\}_n$ converges pointwise everywhere to a function $f : X \longrightarrow \mathbb{R}$, and continuity ensures that f lies in X^*. The fact that $f(x) = \lim_{n\to\infty} f'_n(x)$ for all x in X shows that $\{f'_n\}_n$ converges to f in the weak-$*$ topology. □

Let's talk through how Banach–Alaoglu and separable Banach–Alaoglu differ.

4.2 Example. Let $X = \ell^\infty$ be the set of bounded sequences of real numbers with the supremum norm. Consider the f_n in $(\ell^\infty)^*$ given by

$$(x_1, x_2, \dots) \mapsto x_n.$$

Then the f_n clearly lie in $\overline{B}_1^*$, but I claim that $\{f_n\}_n$ has no convergent subsequence. To see this, let $\{f_{n_k}\}_k$ be any subsequence of $\{f_n\}_n$, where $\{n_k\}_k$ is some increasing sequence of positive integers. Define

$$y \in \ell^\infty \text{ via } y_m := \begin{cases} 1 & \text{if } m = n_{2j}, \\ -1 & \text{if } m = n_{2j+1}, \\ 0 & \text{otherwise.} \end{cases}$$

Then $f_{n_k}(y) = (-1)^k$ does not converge, so $\{f_{n_k}\}_k$ can't converge. So we see that $\overline{B}_1^*$ is not weak-$*$ sequentially compact. This shows that our separable and non-separable Banach–Alaoglu are different statements, and it also gives an alternative proof that ℓ^∞ is not separable (which packages the diagonalization argument of the usual proof of this into the proof of separable Banach–Alaoglu). It also shows that the weak-$*$ topology on ℓ^∞ is not metrizable, because if it were then sequential compactness would agree with compactness.

Let's now move into our ergodic-theoretic application, which is about *invariant measures*. Our setup is a (non-empty) compact metric space (K, d), along with a homeomorphism $\phi : K \longrightarrow K$.

4.3 Example. Our prototypical example will be $K = \mathbb{R}/\mathbb{Z}$, where we shall take ϕ to be some translation.

To even say what invariance means, we need to expand our setup. Write $C(K)$ for the space of continuous functions $K \longrightarrow \mathbb{R}$ equipped with the supremum norm, and write $\mathcal{B}$ for the Borel σ-algebra of K. Let $\mathscr{M}(K)$ denote the set of signed Borel measures $\mu : \mathcal{B} \longrightarrow \mathbb{R}$ topologized by the *total variation* norm

$$\|\mu\|_{\mathscr{M}(K)} := \sup_{A\in\mathcal{B}} \mu(A) - \mu(K \smallsetminus A).$$

Recall from Example 2.10 that $\mathscr{M}(K)$ is isometric to $C(K)^*$ via the map induced from the pairing

$$\langle \mu, f\rangle := \int_K \mathrm{d}\mu\, f.$$

4.4 Definition. Let μ be in $\mathscr{M}(K)$. We say μ is *probability* if $\mu(K) = 1$ and $\mu(A) \geq 0$ for all A in $\mathcal{B}$. We say μ is (ϕ-)*invariant* if it is probability and satisfies

$$\int_K \mathrm{d}\mu\, f \circ \phi = \int_K \mathrm{d}\mu\, f.$$

In the language of our pairing, this is the same as saying that

$$\langle \mu, f \circ \phi \rangle = \langle \mu, f \rangle,$$

which is equivalent to saying that μ equals its pushforward $\phi_\sharp \mu$ under ϕ, as $\phi_\sharp \mu$ is defined by this property. We write $\mathscr{M}(\phi)$ for the space of ϕ-invariant probability measures on K. This lecture is taking so long because this is our first discussion of ergodic theory, so I have to set up all the notation—next time, we'll just start where we left off and pretend that you remember everything I wrote down previously.

Our first observation is that $\mathscr{M}(\phi)$ is bounded and convex. Boundedness follows from the fact that the μ in $\mathscr{M}(\phi)$ are probability, and convexity follows from the fact that invariance and being probability are preserved under weighted averaging. Our next observation will be less obvious:

4.5 Lemma. *The set $\mathscr{M}(\phi)$ is non-empty.*

The idea will be to take any point in K, move it around using ϕ, and construct a measure from there.

Proof. Fix x in K, and consider the (weighted) *empirical measure* μ_n for the points $x, \phi(x), \ldots, \phi^{n-1}(x)$:

$$\mu_n \coloneqq \frac{1}{n} \sum_{j=0}^{n-1} \delta_{\phi^j(x)}.$$

Equivalently, we could define μ_n via our pairing as follows:

$$\langle \mu_n, f \rangle \coloneqq \frac{1}{n} \sum_{j=0}^{n-1} f(\phi^j(x)).$$

Then $\{\mu_n\}_n$ is a perfectly good sequence of probability measures lying in the closed unit ball. Since $C(K)$ is separable, separable Banach–Alaoglu shows that $\{\mu_n\}_n$ has some subsequence $\{\mu_{n_k}\}_k$ that converges to another measure μ in the weak-$*$ topology as $k \to \infty$, where μ lies in $\overline{B}_1^*$. We have that

$$\int_K \mathrm{d}\mu_{n_k} f + \frac{f(\phi^n(x)) - f(x)}{n_k} = \frac{1}{n_k} \sum_{j=1}^{n_k} f(\phi^j(x))) = \frac{1}{n_k} \sum_{j=0}^{n_k - 1} f(\phi(\phi^j(x))) = \int_K \mathrm{d}\mu_{n_k} f \circ \phi,$$

and taking the limit $k \to \infty$ (using the fact that f is bounded) shows that μ is invariant.

How do we check that μ is a probability measure? We can read it off our identification of $\mathscr{M}(K)$ with $C(K)^*$—being non-negative is equivalent to being non-negative on all non-negative functions, and having total volume 1 is equivalent to the integral of 1 equaling 1, both of which are true via taking the limit. $\square$

Now that we have the existence of an invariant measure, we can ask: is it unique? On the flip side, if we change the point x used in the above process, do we get different measures? (The answer is pretty easy for the circle.) Note that the set of invariant measures is convex, and later we shall see that the extremal points of this convex subset are *ergodic*.

5 January 12, 2018

We talked about separable Hahn–Banach last time, and then we proceeded to apply it to construct invariant measures. Today, we'll talk about things that *won't* see applications for a while, but they fit into the theory at this stage: the *Kakutani* and *Eberlein–Šmulian theorems*.

5.1 Theorem (Kakutani–Eberlain–Šmulian). *If X is a Banach space, then the following are equivalent:*

(1) X is reflexive,

(2) $\overline{B}_1$ is weakly compact (that is, compact in the topology $\sigma(X, X^*)$),

(3) $\overline{B}_1$ is weakly sequentially compact.

The equivalence of (1) and (2) in Theorem 5.1 is Kakutani's theorem, which is the relatively easy part. The proof of (1) $\implies$ (2) is rather simple, whereas the proof of (2) $\implies$ (1) will take up most of today's lecture. The equivalence of (2) and (3) in Theorem 5.1 is the Eberlain–Šmulian theorem, which we probably won't get to today.

Proof of (1) $\implies$ (2). Suppose that X is reflexive. Then $J : X \xrightarrow{\sim} X^{**}$ is a homeomorphism from $\sigma(X, X^*)$ to $\sigma(X^{**}, X^*)$, and $J(\overline{B}_1) = \overline{B}_1^{**}$ since J is isometric. Now Banach–Alaoglu tells you that $\overline{B}_1^{**}$ is compact in $\sigma(X^{**}, X^*)$, and we see that $\overline{B}_1$ is compact via pullback through J. □

For the proof of (2) $\implies$ (1) in Theorem 5.1, let me first introduce the following theorem of Helly.

5.2 Theorem (Helly). *Let X be a normed vector space. Let $f = (f_1, \dots, f_n)$ be an n-tuple of elements in $(X^*)^n$, and let $c = (c_1, \dots, c_n)$ be in $\mathbb{R}^n$. Then the following are equivalent:*

(1) For every positive ε, there exists an x in $\overline{B}_1$ such that $|f_i(x) - c_i| < \varepsilon$ for all integers $1 \leq i \leq n$,

(2) For all $\lambda = (\lambda_1, \dots, \lambda_n)$ in $\mathbb{R}^n$, we have

$$\left| \sum_{i=1}^n \lambda_i c_i \right| \leq \left\| \sum_{i=1}^n \lambda_i f_i \right\|_{X^*}.$$

Proof. For (1) $\implies$ (2), let ε be any positive number, and take x as in the statement of (1). Then

$$\begin{aligned} \left| \sum_{i=1}^n \lambda_i c_i \right| &\leq \left| \sum_{i=1}^n \lambda_i f_i(x) \right| + \left(\sum_{i=1}^n |\lambda_i| \right) \varepsilon = \left| \sum_{i=1}^n (\lambda_i f_i)(x) \right| + \left(\sum_{i=1}^n |\lambda_i| \right) \varepsilon \\ &\leq \|x\|_X \left\| \sum_{i=1}^n \lambda_i f_i \right\|_{X^*} + \left(\sum_{i=1}^n |\lambda_i| \right) \varepsilon \leq \left\| \sum_{i=1}^n \lambda_i f_i \right\|_{X^*} + \left(\sum_{i=1}^n |\lambda_i| \right) \varepsilon, \end{aligned}$$

and taking $\varepsilon \to 0$ yields the desired statement.

For (2) $\implies$ (1), begin by defining $\varphi : X \longrightarrow \mathbb{R}^n$ by sending $x \mapsto (f_1(x), \dots, f_n(x))$. Note that (1) is equivalent to saying that c lies in $\overline{\varphi(\overline{B}_1)}$ (via Remark 3.1). Thus if (1) fails, then Hahn–Banach for the finite-dimensional vector space $\mathbb{R}^n$ allows us to choose $\lambda = (\lambda_1, \dots, \lambda_n)$ in $\mathbb{R}^n$ and numbers $a < b$ such that

$$\sum_{i=1}^n \lambda_i c_i = \lambda \cdot c > b > a > \lambda \cdot \varphi(y) = \left(\sum_{i=1}^n \lambda_i f_i \right)(y)$$

for all y in $\overline{B}_1$. Taking the supremum over such y shows that (2) also fails. □

Next, we'll present the following consequence of Helly's theorem.

5.3 Theorem (Goldstine)**.** *Let X be a Banach space. Then $J(\overline{B}_1)$ is dense in $\overline{B}_1^{**}$ with respect to the weak-$*$ topology $\sigma(X^{**}, X^*)$.*

Proof. We shall show that, for any x^{**} in $\overline{B}_1^{**}$ and any neighborhood V of x^{**} in $\sigma(X^{**}, X^*)$, the intersection $J(\overline{B}_1) \cap V$ is nonempty. For this, it suffices to take V as an element of a basis for $\sigma(X^{**}, X)$, so we assume that

$$V = \{y^{**} \in X^{**} \mid |y^{**} - x^{**}(f_i)| < \varepsilon \text{ for all integers } 1 \leq i \leq n\},$$

where the f_i are some points in X^*, and ε is a positive number. We want to find an x in $\overline{B}_1$ for which Jx lies in V, which is equivalent to asking that

$$|f_i(x) - x^{**}(f_i)| = |(Jx - x^{**})(f_i)| < \varepsilon.$$

Set $c_i = x^{**}(f_i)$. Because x^{**} has norm at most 1 in X^{**}, we see that

$$\left|\sum_{i=1}^{n} \lambda_i c_i\right| = \left|\sum_{i=1}^{n} \lambda_i x^{**}(f_i)\right| = \left|x^{**}\left(\sum_{i=1}^{n} \lambda_i f_i\right)\right| \leq \left\|\sum_{i=1}^{n} \lambda_i f_i\right\|_{X^*}$$

Therefore applying Helly's theorem yields the desired x. □

We can now show that (2) $\Longrightarrow$ (1) in Theorem 5.1, completing our proof of Kakutani's theorem.

Proof of (2) $\Longrightarrow$ (1) in Theorem 5.1. Suppose that $\overline{B}_1$ is weakly compact. In any case, $J : X \longrightarrow X^{**}$ is always a continuous map from $\sigma(X, X^*)$ to $\sigma(X^{**}, X^*)$. Therefore the image $J(\overline{B}_1)$ of $\overline{B}_1$ is compact in $\sigma(X^{**}, X^*)$, and Goldstine's theorem tells us that $J(\overline{B}_1)$ is dense in $\overline{B}_1^{**}$ with respect to $\sigma(X^{**}, X^*)$. But this topology is Hausdorff, so $J(\overline{B}_1)$ is closed and hence equal to $\overline{B}_1^{**}$. By scaling, this proves that J is surjective, and the injectivity of J is always present, by Exercise 2.6. □

As for the Eberlain–Šmulian theorem, we'll first need a couple of lemmas on Banach spaces. I didn't manage to assign these lemmas as homework problems, so I suppose we'll have to discuss them in class.

5.4 Lemma. *Let X be a normed vector space. If X^* is separable, then X is separable.*

This is one of those things where the proof is so easy, why would I write it in my notes? Why, to embarrass myself in front of all of you, of course! I think there's only one possible proof, so let's find it.

Proof. Let's use our first hypothesis: choose a countable sequence $\{x_n^*\}_n$ in X^* that is dense, and assume all the x_n^* are nonzero. Next, let's use the only other thing we know: choose $\{x_n\}_n$ in X with $\|x_n\| = 1$ and $|x_n^*(x_n)| \geq \frac{1}{2}\|x_n^*\|_{X^*}$, which is possible by the definition of the supremum norm on X^*. Note that the $\mathbb{Q}$-span of $\{x_n\}_n$ in X^* is countable, and the density of $\mathbb{Q}$ in $\mathbb{R}$ shows that it is dense in

$$\operatorname{span}_{\mathbb{R}}\{x_n\}_n.$$

For a contradiction, suppose that X is not the closure of $\operatorname{span}_{\mathbb{R}}\{x_n\}_n$. Then Hahn–Banach gives us a nonzero y^* in X^* for which $y^*(x) \neq 0$ for all x in $\operatorname{span}_{\mathbb{R}}\{x_n\}_n$. The density of $\{x_m^*\}_m$ (and the fact that y is nonzero) lets us find a subsequence $\{x_m^*\}_m$ such that $\|y^* - x_m^*\| \leq \frac{1}{4}\|x_m^*\|$ for all m. Thus we get

$$0 = \|y^*(x_m)\| \geq \|x_m^*(x_m)\| - \tfrac{1}{4}\|x_m^*\| \geq \tfrac{1}{2}\|x_m^*\| - \tfrac{1}{4}\|x_m\| > 0,$$

which is a contradiction. □

6 January 17, 2018

Recall that we proved Kakutani's theorem last time. Today, we'll proceed to Eberlain–Šmulian theorem. In addition to the lemma I had trouble with last time, we'll need another lemma:

6.1 Lemma. *Let X be a Banach space, and let Y be a closed subspace. If X is reflexive, then Y is reflexive.*

There's a direct proof of Lemma 6.1 in Bühler–Salamon that's very hands on and takes two pages, but given that we've already proven Kakutani's theorem, we can use it to get an easy proof of this.[5]

Proof. The space Y has two kinds of weak topologies:

- the usual weak topology $\sigma(Y, Y^*)$,
- the topology induced from Y being a subset of X with the weak topology $\sigma(X, X^*)$.

By Hahn–Banach, these two topologies coincide. Kakutani's theorem impies that Y is reflexive if and only if $\overline{B}_1^Y$ is compact in $\sigma(Y, Y^*)$, and the coincidence of our weak topologies implies that this this is equivalent to $\overline{B}_1^Y$ being compact in $\sigma(X, X^*)$. The reflexivity of X and Kakutani's theorem indicate that $\overline{B}_1^X$ is compact in $\sigma(X, X^*)$, and the fact that $\overline{B}_1^Y = \overline{B}_1^X \cap Y$ completes our proof. □

Lemma 6.1 is easy once you think about it in the right way, just like everything else in math.

6.2 Lemma. *Let X be a Banach space, and let M be a finite-dimensional subspace of X^*. Then there exist $x_1, \ldots, x_n$ in $\overline{B}_1$ such that, for any y^* in M, we have*

$$\max_{1\leq k\leq n} y^*(x_k) \geq \frac{1}{2}\|y^*\|.$$

Of course we can always find points in X whose values on y^* approximate $\|y^*\|$ arbitrarily well, but Lemma 6.2 says that we can do this roughly uniformly, to an extent, for any finite-dimensional subspace of X^*.

Outline of the proof. By Banach–Alaoglu, $\overline{B}_1^*$ is compact in the weak-$*$ star topology. Choose a $\frac{1}{4}$-net $y_1^*, \ldots, y_n^*$ of $\overline{B}_1^* \cap M$, and choose x_k in $\overline{B}_1$ such that $y_k^*(x_k) \geq \frac{3}{4}\|y_k^*\|$. □

With Lemma 6.2 in hand, we can now prove the Eberlain–Šmulian theorem.

Proof of (1) $\Longleftrightarrow$ (3) in Theorem 5.1. First, suppose that X is reflexive, and let $\{x_n\}_n$ be a sequence in $\overline{B}_1$. Write Y for $\overline{\operatorname{span}_{\mathbb{R}}\{x_n\}_n}$, which is separable, as remarked in the proof of Lemma 5.4. This Y is also reflexive by Lemma 6.1, which makes $Y^{**} = Y$ is separable. Applying Lemma 5.4 shows that Y^* is separable. From here, applying separable Banach–Alaoglu to Y^* indicates that $\overline{B}_1^{Y^{**}} = \overline{B}_1^Y$ is weakly sequentially compact, and hence $\{x_n\}_n$ has a weakly converging subsequence.

Conversely, suppose that $\overline{B}_1$ is weakly sequentially compact. We want to show that $J : \overline{B}_1 \hookrightarrow \overline{B}_1^{**}$ is surjective, and the proof shall be one of those slick things. Fix x^{**} in $\overline{B}_1^{**}$. We want to choose x_k in $\overline{B}_1$, x_k^* in $\overline{B}_1^*$, and n_k in $\mathbb{N} := \mathbb{Z}_{>0}$ such that

(1) $n_k < n_{k+1}$,

(2) $\max_{n_k<i\leq n_{k+1}} y^{**}(x_i^*) \geq \frac{1}{2}\|y^{**}\|$ for all y^{**} in $\operatorname{span}_{\mathbb{R}}\{x^{**}, Jx_1, \ldots, Jx_k\}$,

[5] Ishan Banerjee observes that Lemma 6.1 also follows immediately from the surjective half of the five lemma.

(3) $|x^{**}(x_i^*) - x_i^*(x_{k+1})| < \frac{1}{k+1}$ for all integers $1 \leq i \leq n_{k+1}$.

We do this as follows: inductively apply Lemma 6.2 to obtain n_{k+1} and $x^*_{n_k+!}, \dots, x^*_{n_{k+1}-1}$ that satisfy (1) and (2). As for (3), use Goldstine's theorem to choose the appropriate x_{k+1}.

Because $\overline{B}_1$ is weakly sequentially compact by assumption, we can assume that the $\{x_k\}_k$ converge to some x in $\overline{B}_1$ by replacing $\{x_k\}_k$ with a subsequence—we see that this preserves our desired conditions. Using (3) and taking $k \to \infty$, we obtain $x^{**}(x_i^*) = x_i^*(x)$ by weak convergence. Now (2) gives

$$\sup_{i \geq 1} y^{**}(x_i) \geq \frac{1}{2}\|y^{**}\|$$

for all y^{**} in $\overline{\operatorname{span}_{\mathbb{R}}\{x^{**}, Jx_1, Jx_2, \dots\}}$, where we can pass to strong convergence (which is necessary for the continuity of $\|\cdot\|$) by Mazur's theorem. Observe that $x^{**} - Jx$ lies in $\overline{\operatorname{span}_{\mathbb{R}}\{x^{**}, Jx_1, Jx_2, \dots\}}$, so

$$0 = \sup_{i \geq 1}(x^{**} - Jx)(x_i^*) \geq \frac{1}{2}\|x^{**} - Jx\|.$$

and thus $x^{**} = Jx$. □

This is a very slick proof—it works by trickily finding precisely the right sequence to prove reflexivity. Of course, it's been remixed over the years until it reached the textbook form you see now.

Finally, we'll do one more abstract result and then apply our results to ergodic theory. This abstract result shall be in convex analysis: we'll be discussing *extremal points*.

6.3 Definition. Let X be a normed vector space, and let K be a convex subset of X. A *face* of K is a subset F of K such that, if x lies in F and x_1 and x_2 are points in K satisfying $x = \lambda x_1 + (1-\lambda)x_2$ for some λ in $(0,1)$, then x_1 and x_2 lie in F.

In other words, F is closed under taking endpoints of segments through F.

6.4 Definition. Let K be a convex subset, and let x be a point of K. We say that x is *extremal* in K if $\{x\}$ is a face of K. Write $\mathscr{E}(K)$ for the set of extremal points of K.

I have but a few minutes left—let's see whether I can rush through the proof!

6.5 Theorem (Krein–Milman). *If X is a normed vector space and K is a convex subset of X, then*

$$K = \overline{\operatorname{hull}(\mathscr{E}(K))}.$$

Proof. Let's break it down into four steps:

Step 1. Let $\mathscr{K} := \{K \subseteq X : K \text{ is compact, convex, and non-empty}\}$. We put an ordering on $\mathscr{K}$ by saying that $K_1 \preceq K_2$ if and only if K_1 is a face of K_2. Note that descending chains have minimums because the intersection of nested compact spaces is nonempty, so we can run Zorn's lemma on $\mathscr{K}$.

Step 2. If K is in $\mathscr{K}$ and f is in X^*, then $K \cap f^{-1}(\max_K f) \preceq K$. This is true due to the linearity of f.

Step 3. Minimal elements of $\mathscr{K}$ are singletons, for if K in $\mathscr{K}$ were not a singleton, then by Hahn–Banach we could find f in X^* that is not constant on K. Applying Step 2 would then yield $K \cap f^{-1}(\max_K f)$ in $\mathscr{K}$ that is strictly less than K.

Step 4. I claim $K = \overline{\operatorname{hull}(\mathscr{E}(K))}$. Write Y for $\overline{\operatorname{hull}(\mathscr{E}(K))}$. Then $Y \subseteq K$, and if there exists a point x in $K \smallsetminus Y$, then we can choose f in X^* such that $f(x) > \max_Y f$. Applying Step 2 yields a $\widetilde{K} \preceq K$ in $\mathscr{K}$ that does not intersect Y. Any minimal element lying below $\widetilde{K}$ (which exists by Step 1) will yield an extremal point of K that isn't in Y, which contradicts the construction of Y.

□

Next time, we'll discuss some applications to ergodic theory.

7 January 19, 2018

Last time, we discussed Eberlain–Šmulian and Krein–Milman. I had promised that we'd get to applications of these theorems to ergodic theory, but I realized we need a slightly more general version of Krein–Milman for this. More precisely, we'll need a version that works for the weak and weak-$*$ topologies. For this, we establish the following framework that I alluded to in Lecture 2.

7.1 Definition. Let X be a vector space. A topology on X is *locally convex* if it is generated by a family of seminorms $p_a : X \longrightarrow \mathbb{R}$, where a varies over some indexing set A.

7.2 Example. Let X be a normed vector space. Then $\sigma(X, X^*)$ and $\sigma(X^*, X)$ are locally convex, where the former case the seminorms are of the form $p_{f_1,\dots,f_n}(x) = |f_1(x)| + \cdots + |f_n(x)|$ for any $f_1, \dots, f_n$ in X^*. The story for the latter is analogous.

7.3 Theorem. *The Krein–Milman theorem applies for locally convex vector spaces X.*

Proof. The proof of the Krein–Milman theorem from last time works, because its most important step, the Hahn–Banach theorem, works for arbitrary locally convex vector spaces. While Brezis does not prove Hahn–Banach in this generality, the proof of Hahn–Banach in this setting remains the same as well. □

Let's now return to ergodic theory! Let (K, d) be a compact metric space, and let $\phi : K \longrightarrow K$ be a homeomorphism.

7.4 Remark. Recall from Lecture 4 that the space $\mathscr{M}(\phi)$ of ϕ-invariant Borel probability measures on K is a convex, non-empty, weak-$*$ compact subset of $C(K)^*$.

7.5 Examples. Let $K = \mathbb{R}/\mathbb{Z}$, and let $\phi_a(x) = x + a$ for any a in $\mathbb{R}$, where a only depends on its image in $\mathbb{R}/\mathbb{Z}$. Here are some examples of ϕ-invariant Borel probability measures on K:

(1) The Lebesgue measure (though recall that most treatments of the Lebesgue measure define it on not just the Borel σ-algebra but rather its completion with respect to the Lebesgue measure),

(2) If $a = p/q$ is rational, where p and q are integers and q is positive, then for any b in $\mathbb{R}$, the measure

$$\mu_{b,q} := \frac{1}{q}\sum_{k=0}^{q-1} \delta_{b+k/q}$$

is $\phi_{p/q}$-invariant.

(3) In the situation of (2), the convexity of $\mathscr{M}(\phi_{p/q})$ implies that scaled linear combinations of the $\mu_{b,q}$ (for varying b) are invariant.

7.6 Definition. Let μ lie in $\mathscr{M}(\phi)$. We say μ is *ergodic* if $\mu(B) \in \{0, 1\}$ for any Borel subset B satisfying $\phi(B) = B$.

7.7 Example. In the setting of Examples 7.5, the Lebesgue measure is not ergodic when a is rational. We can see this by taking the union of a-translates of a small interval. However, in this situation, the $\mu_{b,a}$ *are* ergodic, as is Lebesgue measure itself when a is irrational instead.

7.8 Theorem. *Let μ lie in $\mathscr{M}(\phi)$. Then μ is ergodic if and only if μ is extremal in $\mathscr{M}(\phi)$.*

For the moment, we shall only prove this modulo a harder theorem.

7.9 Remark. Theorem 7.8 shows that ergodicity is a characterization of extremal that is intrinsic to $\mathscr{M}(\phi)$.

We begin with the easier direction of Theorem 7.8.

7.10 Lemma. *If μ is extremal, then it is ergodic.*

The idea is that, if μ is not ergodic, we can break it up into pieces.

Proof. Suppose μ is not ergodic. Then we can choose a Borel subset B such that $\phi(B) = B$ and $\mu(B)$ lies strictly between 0 and 1. Form new probability measures μ_0 and μ_1 via setting

$$\mu_0(A) = \frac{\mu(A \cap B)}{\mu(B)} \text{ and } \mu_1 = \frac{\mu(A \smallsetminus B)}{1 - \mu(B)}.$$

The ϕ-invariance of μ and B indicate that μ_0 and μ_1 remain ϕ-invariant. However, we have

$$\mu = \mu(B)\mu_0 + (1 - \mu(B))\mu_1,$$

and since $\mu_0(B) = 1$ and $\mu_1(B) = 0$, we see that the μ_0 and μ_1 are distinct. Thus μ is not extremal. $\square$

We won't be able to finish everything today, but we *will* first introduce von Neumann's *mean ergodic theorem*:

7.11 Theorem (Von Neumann)**.** *If μ in $\mathscr{M}(\phi)$ is ergodic and f lies in $L^2(\mu)$, then*

$$\frac{1}{n}\sum_{k=0}^{n-1} f \circ \phi^k \to \int_K \mathrm{d}\mu\, f$$

in $L^2(\mu)$.

Von Neumann's mean ergodic theorem shall be the consequence of a more general statement on Hilbert spaces. We won't prove it today—for now, we'll just prove how it finishes the proof of Theorem 7.8.

7.12 Lemma. *If μ_0 and μ_1 in $\mathscr{M}(\phi)$ are ergodic measures that satisfy*

$$\mu_1(B) = \mu_0(B)$$

for all Borel subsets B such that $\phi(B) = B$, then $\mu_0 = \mu_1$.

Basically, this says that ergodic measures are really characterized by their values on the σ-algebra of ϕ-invariant Borel subsets. This determination will use the fact that they're only valued in $\{0, 1\}$ on such subsets anyways.

Proof. Let be i in $\{0, 1\}$, and fix f in $C(K) \subseteq L^2(\mu_i)$. Von Neumann's mean ergodic theorem implies that

$$\frac{1}{n}\sum_{k=0}^{n-1} f \circ \phi^k \to \int_K \mathrm{d}\mu_i\, f$$

in $L^2(\mu_i)$. A result from last quarter shows that there exists a Borel subset B_i satisfying $\mu_i(B) = 1$ and a subsequence $n_k \to \infty$ such that

$$\frac{1}{n_k}\sum_{k=0}^{n_k-1} f \circ \phi^k(x) \to \int_K \mathrm{d}\mu_i\, f \tag{$\star$}$$

for all x in B_i, that is, we have a subsequence that converges almost everywhere. We can choose n_k independently of i by taking it for $i = 0$ and then refining it to satisfy the same condition for $i = 2$. Letting $A_i := \bigcap_{n\in\mathbb{Z}} \phi^n(B_i)$, we see that $\mu_i(A_i) = 1$ and A_i is ϕ-invariant. By hypothesis, we see that $\mu_{1-i}(A_i) = \mu_i(A_i)$. Equation $(\star)$ indicates that

$$\int_K \mathrm{d}\mu_0\, f = \int_K \mathrm{d}\mu_1\, f,$$

and letting f vary shows that $\mu_0 = \mu_1$. □

Let us now finish the proof of Theorem 7.8.

7.13 Lemma. *If μ in $\mathscr{M}(\phi)$ is ergodic, then it is extremal.*

Proof. By Krein–Milman, we can write μ as the scaled linear combination

$$\mu = \lambda\mu_0 + (1-\lambda)\mu_1$$

of two extremal μ_0 and μ_1, where λ lies in $[0, 1]$. Suppose that λ actually lies in $(0, 1)$. Lemma 7.10 shows that the μ_0 and μ_1 are ergodic. For any ϕ-invariant Borel subset B, if we have $\mu(B) = 0$, the above relation shows that $\mu_0(B) = \mu_1(B) = 0$. Similarly, if $\mu(B) = 1$, then $\mu_0(B) = \mu_1(B) = 1$, so Lemma 7.12 implies that $\mu_0 = \mu_1$. □

Next time, we'll go through the von Neumann mean ergodic theorem, as well as what ergodicity actually *means*.

8 January 22, 2018

I seem to have brought the notes I wanted to throw away and thrown away the ones I meant to bring. Oh well—we'll muddle through, as usual. Today, we'll complete the proof of Theorem 7.8 by proving von Neumann's mean ergodic theorem. We'll actually start with the *abstract ergodic theorem*, also due to von Neumann:

8.1 Theorem (Abstract ergodic theorem)**.** *Let H be a Hilbert space, and let $U : H \longrightarrow H$ be a bijective unitary (i.e. preserves the inner product) operator. Write H^U for the closed subspace $\{x \in H : Ux = x\}$, and write $\pi : H \longrightarrow H^U$ for the orthogonal projection map onto H^U. Then for any x in H, we have*

$$\frac{1}{n}\sum_{k=0}^{n-1} U^k x \to \pi(x) \tag{$\star$}$$

in the strong topology.

Von Neumann used the spectral theorem, but we haven't seen that yet, so we'll use a sneaky argument due to Riesz instead. When we prove the spectral theorem, we'll return to von Neumann's proof.

Proof. Because we're very smart[6], we'll guess exactly what H^U equals. Let $G := \{Uy - y : y \in H\}$. Our first observation is that G and H^U are orthogonal, since for all x in H^U, we have

$$\langle x, Uy - y\rangle = \langle x, Uy\rangle - \langle x, y\rangle = \langle U^{-1}x, y\rangle - \langle x, y\rangle = \langle x, y\rangle - \langle x, y\rangle = 0.$$

[6] Smart?

Our second observation is that the sum

$$\frac{1}{n}\sum_{k=0}^{n-1} U^k(Uy-y) = \frac{1}{n}(U^n y - y) \to 0$$

telescopes, and the limit goes to zero because U has operator norm 1. As ($\star$) is linear and immediately holds for x in H^U, we see that ($\star$) holds for $H^U + G$. Furthermore, as the operator

$$\frac{1}{n}\sum_{k=0}^{n-1} U^k$$

is uniformly bounded by 1, taking the limit yields a linear operator, so ($\star$) actually holds for $\overline{H^U + G}$.

Therefore it suffices to show that $\overline{H^U + G} = H$. For any z in H, write it as $w + z$ for some w in $\overline{G}$ and x in $\overline{G}^{\perp}$. It'll be enough to show that x lies in H^U. We know that $\langle x, Ux - x\rangle = 0$, so we obtain

$$\begin{aligned}\|x - Ux\|^2 &= \langle x - Ux, x - Ux\rangle = \langle x, x\rangle + \langle Ux, Ux\rangle - \langle x, Ux\rangle - \langle Ux, x\rangle \\ &= 2(\langle x, x\rangle - \langle x, Ux\rangle) = 2\langle x, x - Ux\rangle = 0\end{aligned}$$

as desired. □

The above proof is very nice, but it sheds absolutely no light as to why the abstract ergodic theorem is true. On your next homework, there'll be a quantitative version of this, which I stole from Terry Tao's blog. Hopefully this makes the theorem clearer. You *could* just read off the answer from Terry Tao's blog, but…try not to do that. Additionally, the proof by von Neumann is entirely clear, modulo the mystery of the spectral theorem.

This was for real Hilbert spaces—it also works for complex Hilbert spaces, where the use of symmetry should be replaced by the polarization identity. That's a general recipe for going between real and complex Hilbert spaces, and this trick only doesn't work if there's something truly strange going on.

Let's return to the proof of von Neumann's mean ergodic theorem. Let H be $L^2(K,\mu)$, and let U be the operator given by $Uf := f \circ \phi$. The unitariness of U follows from the ϕ-invariance of μ. We have the following immediate corollary of the abstract ergodic theorem.

8.2 Corollary. *For any function f in $L^2(K,\mu)$, we have*

$$\frac{1}{n}\sum_{k=0}^{n-1} f \circ \phi^k \to \pi(f),$$

where π denotes orthogonal projection onto $\{g \in L^2(K,\mu) : g \circ \phi = g\}$.

We remind ourselves that are equalities of elements in L^2 are taken almost everywhere. Note that the result of Corollary 8.2 only differs from von Neumann's mean ergodicity theorem by replacing $\pi(f)$ with $\mu(f)$. For this last step, we use the following tricky lemma regarding ergodicity.

8.3 Lemma. *Let μ be ergodic, let f be in $L^2(K,\mu)$, and suppose that $f \circ \phi = f$ almost everywhere. Then f is constant almost everywhere.*

Proof. Let $\widetilde{A}$ denote the set $\{x \in K : f(\phi(x)) \neq f(x)\}$. By assumption, we have $\mu(\widetilde{A}) = 0$. Now let $A := \bigcup_{n\in\mathbb{Z}} \phi^n(\widetilde{A})$, which also has measure zero and is ϕ-invariant. Next, form $c := \int_K \mathrm{d}\mu\, f$, and take the sets

$$\begin{aligned} B^+ &:= \{x \in K \smallsetminus A : f(x) > c\}, \\ B^0 &:= \{x \in K \smallsetminus A : f(x) = c\}, \\ B^- &:= \{x \in K \smallsetminus A : f(x) < c\}.\end{aligned}$$

These sets are invariant by the invariance of A and f, so ergodicity implies that their measures lie in $\{0, 1\}$. As K is a disjoint union of A, B^+, B^0, and B^-, exactly one of these subsets has measure one, and it can only be B^0. $\square$

Proof of von Neumann's mean ergodicity theorem. By Corollary 8.2, the desired sequence of averages converges to $\pi(f)$. By Lemma 8.3, we see that $\pi(f) = c$. $\square$

8.4 Theorem. *The following are equivalent:*

(1) μ is ergodic,

(2) μ is extremal in $\mathscr{M}(\phi)$,

(3) For all f in $L^2(K, \mu)$, we have $\frac{1}{n}\sum_{k=0}^{n-1} f \circ \phi^k \to \int_K \mathrm{d}\mu\, f$ in $L^2(K, \mu)$.

(4) For all f in $L^1(K, \mu)$, we have $\frac{1}{n}\sum_{k=0}^{n-1} f \circ \phi^k \to \int_K \mathrm{d}\mu\, f$ in $L^1(K, \mu)$.

Proof. The equivalence of (1) and (2) is precisely Theorem 7.8, and (1) $\Longrightarrow$ (3) is precisely the mean ergodic theorem. For (3) $\Longrightarrow$ (1), let B be a ϕ-invariant Borel subset. Consider the characteristic function $f = \mathbf{1}_B$, and note that $\frac{1}{n}\sum_{n=0}^{k-1} f \circ \phi^k = f$ by the invariance of B. Therefore $\mathbf{1}_B$ constantly equals $\int_K \mathrm{d}\mu\, f = \mu(B)$, and $\mathbf{1}_B$ only takes values in $\{0, 1\}$, so $\mu(B)$ lies in this range as well. Finally, the equivalence between (1) and (4) will be left for your homework because it's not functional analysis enough for lecture—it's too much of a callback to last quarter. $\square$

8.5 Example. Return to the setting of Examples 7.5. I claim that the Lebesgue measure μ is ergodic for ϕ_a when a is irrational. To see this, it suffices to show that μ satisfies condition (3) in Theorem 8.4, and in turn it suffices to use trigonometric sums for f. By linearity, this further reduces to powers of the exponential, which was then shown as a homework problem from last quarter.

9 January 24, 2018

I didn't quite get to everything I wanted to cover about ergodic theory last time, but that's okay—we'll move on. What I want to talk about next is *spectral theory*. This will take us a while. Our goal is to prove the analog of Jordan decomposition, except in the setting of infinite-dimensional vector spaces. I'll begin with two concepts that don't seem to have anything to do with spectral theory: *dual* and *compact* operators.

9.1 Definition. Let X and Y be normed vector spaces, and let $f : X \longrightarrow Y$ be a bounded linear map. Its *dual* $f^* : Y^* \longrightarrow X^*$ is the linear operator defined by $f^*(y^*) := y^* \circ f$.

Let's begin by thinking about what the norm of f^* would be.

9.2 Lemma. *The norm of the dual equals*

$$\|f^*\|_{\mathscr{L}(Y^*,X^*)} = \|f\|_{\mathscr{L}(X,Y)}.$$

Proof. We have

$$\begin{aligned}\|f^*\| &= \sup_{\|y^*\|\le 1} \|f^*(y^*)\| = \sup_{\|y^*\|\le 1}\sup_{\|x\|\le 1} \|(f^*(y^*))(x)\| = \sup_{\|y^*\|\le 1}\sup_{\|x\|\le 1} y^*(f(x)) \\ &= \sup_{\|x\|\le 1}\sup_{\|y^*\|\le 1} y^*(f(x)) = \sup_{\|x\|\le 1} \|f(x)\| = \|f\|.\end{aligned}$$

by Theorem 2.4. $\square$

There are all sorts of relationships between f and f^* that one can prove; Bühler–Salamon has a whole chapter on this entitled *Fredholm theory*, which is really precisely the study of relating these duals to the original operator. Let's continue discussing some immediate relationships.

9.3 Lemma. *Let $f : X \longrightarrow Y$ and $g : Y \longrightarrow Z$ be two bounded linear maps between normed vector spaces.*

(1) $(g \circ f)^* = g^* \circ f^*$,

(2) $J_Y \circ f = f^{**} \circ J_X$.

In other words, taking duals is a contravariant functor, and J is a natural transformation from the identity functor to the double-dual.

9.4 Example. My favorite concrete example of where duals appear in nature comes from ergodic theory. As usual, let (K, d) be a compact metric space, and let $\phi : K \longrightarrow K$ be a homeomorphism. Then ϕ induces an operator $T : C(K) \longrightarrow C(K)$ via $Tf := f \circ \phi$, and under our identification of $\mathscr{M}(K)$ with $C(K)^*$ from Example 2.10, the dual $T^* : \mathscr{M}(K) \longrightarrow \mathscr{M}(K)$ corresponds to the map sending $\mu \mapsto \phi_*\mu$, where $\phi_\sharp\mu$ is the measure on K with value $(\phi_\sharp\mu)(B) := \mu(\phi^{-1}(B))$.

9.5 Example. There's also the finite-dimensional case: for any m-by-n matrix A, write $L_A : \mathbb{R}^n \longrightarrow \mathbb{R}^m$ for the corresponding linear operator. Then $(L_A)^*$ equals L_{A^t}.

9.6 Example. Let H be a Hilbert space, and let $T : H \longrightarrow H$ be a bounded linear operator on H. The identification of H with H^* given by the inner product $\langle \cdot, \cdot \rangle$ also identifies $T^* : H \longrightarrow H$ with the operator characterized by

$$\langle x, Ty \rangle = \langle T^*x, y \rangle$$

for all x and y in H. In this situation, we call T^* the *adjoint* of T.

9.7 Theorem. *Let f lie in $\mathscr{L}(X, Y)$. Then*

(1) $\operatorname{im}(f)^\perp = \ker(f^*)$,

(2) $\operatorname{im}(f^*)^\perp = \ker(f)$,

(3) $\operatorname{im}(f)$ *is dense in Y if and only if f^* is injective,*

(4) $\operatorname{im}(f^*)$ *is weak-$*$ dense in X^* if and only if f is injective.*

Proof. I claim that these all immediately follow from unwrapping the definitions.

(1) We have

$$y^* \in \operatorname{im}(f)^\perp \iff \forall x \in X,\, y^*(f(x)) = 0 \iff \forall x,\, (f^*(y^*))(x) = 0 \iff y^* \in \ker f^*.$$

(2) By Hahn–Banach, we have

$$x \in \operatorname{im}(f^*)^\perp \iff \forall y^* \in Y^*,\, (f^*(y^*))(x) = 0 \iff \forall y^* \in Y^*,\, y^*(f((x)) = 0 \iff x \in \ker(f).$$

(3) By (1), we have

$$\operatorname{im}(f) \text{ is dense in } Y \iff \operatorname{im}(f)^\perp = 0 \iff \ker(f^*) = 0 \iff f^* \text{ is injective.}$$

(4) The proof is the same as that of (3), except we use the weak-$*$ topology and part (2) instead. $\square$

There are some more properties about dual operators that I've included on the homework, but let's stop here for now. At this point, we turn to compact operators.

9.8 Lemma. *Let f lie in $\mathscr{L}(X, Y)$. Then the following are equivalent:*

(1) if $\{x_n\}_n$ in X is bounded, then $\{f(x_n)\}_n$ in Y has a convergent subsequence,

(2) if S is a bounded subset of X, then $\overline{f(S)}$ in Y is compact,

(3) $\overline{f(B_1)}$ in Y is compact.

Someone offers a proof, but Charlie is too sleepy, so we'll take it on faith.

9.9 Definition. Let f lie in $\mathscr{L}(X, Y)$.

(1) We say that f is *compact* if it satisfies the equivalent conditions of Lemma 9.8.

(2) We say that f is *finite rank* if $\dim \operatorname{im} f$ is finite.

(3) We say that f is *completely continuous* if, for every weakly convergent sequence $x_n \rightharpoonup x$ in X, we have $f(x_n) \to f(x)$ in Y strongly.

Lemma 9.8 implies the following statement.

9.10 Lemma. *Any finite rank operator is compact.*

We also have the following lemma.

9.11 Lemma. *The space $\mathscr{K}(X, Y)$ of compact operators is a closed subset of $\mathscr{L}(X, Y)$.*

Proof. Use Lemma 9.8 and a diagonalization argument. □

9.12 Remark. Morally, the space of compact operators is the closure of the space $\mathscr{F}(X, Y)$ of finite rank operators. However, this is not true in general. What do we have instead?

9.13 Lemma. *The set $\mathscr{F}(X, Y)$ is dense in $\mathscr{K}(X, Y)$ if and only if, for any compact subset K of Y and positive number ε, there exists an f in $\mathscr{F}(X, Y)$ satisfying*

$$\sup_{y \in K} \inf_{x \in X} \|y - f(x)\| < \varepsilon.$$

9.14 Remark. The hypotheses of Lemma 9.13 are certainly true when Y is separable and X is infinite-dimensional, since we can take a finite ε-dense subset of Y and then build the desired f using (the higher-dimensional version of) Hahn–Banach.

We'll finish discussing Lemma 9.13 next time.

10 January 26, 2018

I mumbled a bit last time about compact operators—let's summarize what I was trying to do. We defined the notions of compact, completely continuous, and finite rank operators. We discussed the idea that finite rank operators should approximate compact ones in Lemma 9.13. Now, let's turn to completely continuous operators. Let X and Y be Banach spaces, and let $f : X \longrightarrow Y$ be a bounded linear operator.

10.1 Lemma. *If f is compact, then it is also completely continuous.*

Proof. Suppose that $x_n \to x$ weakly in X. By applying the uniform boundedness principle to the operators x_n on the Banach space X^* and using the fact that J_X is an isometry, we see that the x_n are bounded. Since f is compact, we can replace $\{x_n\}_n$ with a subsequence such that there exists y in Y satisfying $f(x_n) \to y$ strongly. It is enough to prove that $f(x) = y$. If we had $y \neq f(x)$, Hahn–Banach would give us a y^* in Y^* satisfying $y^*(y - f(x)) = 1$. We have

$$(y^* \circ f)(x_n) = y^*(f(x_n)) \to y^*(y),$$

but weak convergence would indicate that

$$(y^* \circ f)(x_n) \to (y^* \circ f)(x) = y^*(f(x)),$$

violating $y^*(y - f(x)) = 1$. □

In the presence of reflexivity, the converse of Lemma 10.1 also holds:

10.2 Lemma. *If X is reflexive, then f being completely continuous implies that f is compact.*

Proof. We use condition (1) of Lemma 9.8: Let x_n be a sequence in B_1^X. Eberlain–Šmulian implies that we may replace x_n with a subsequence such that that $x_n \to x$ weakly for some x in B_1^X. Complete continuity then implies that $f(x_n) \to f(x)$ strongly, as desired. □

The following theorem of Schauder is another one of those theorems with an abstract nonsense proof.

10.3 Theorem (Schauder)**.** *The operator f^* is compact if and only if f is compact.*

Proof. Suppose that f is compact, and write K for the subset $\overline{f(\overline{B}_1^X)}$ of Y. Then K is a compact metric space, as it's equipped with the norm metric from Y. Given a sequence y_k^* in $\overline{B}_1^{Y^*}$, we shall show that $f^*(y_k^*)$ has a convergent subsequence in X^*. Observe that the restrictions $y_k^*\big|_K$ satisfy

- $|y_k^*(f(x))| \leq \|f\|$ for all x in $\overline{B}_1^X$,
- $|y_k^*(f(x_1)) - y_k^*(f(x_2))| \leq \|f(x_1) - f(x_2)\|$ for all x_1 and x_2 in $\overline{B}_1^X$.

That is, the $y_k^*\big|_K$ are uniformly bounded and Lipschitz. Thus Arzelà–Ascoli implies that we can pass to a subsequence of $y_k^*\big|_K$ that converges to some g in $C(K)$. Now $g \circ f$ is linear on K and also obeys the above bounded and Lipschitz conditions, so, via linearity, we can extend $g \circ f$ to a bounded linear functional x^* on all of X that satisfies the analogous conditions. We see that $y_k^* \circ f \to x^*$.

Conversely, suppose that f^* is compact. By the above argument, we see that f^{**} is compact. Next, consider the commutative diagram

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "J_X"] & Y \arrow[d, "J_Y"] \\ X^{**} \arrow[r, "f^{**}"] & Y^{**}. \end{tikzcd}$$

For any bounded sequence x_n in X, the isometry of J_X implies that $J_X(x_n)$ is bounded. Thus $f^{**}(J_X(x_n)) = J_Y(f(x_n))$ converges after passing to some subsequence, so in particular it is Cauchy. The isometry of J_Y implies that $f(x_n)$ is Cauchy, and hence it converges by the completeness of Y. □

Speaking of abstract facts, let's prove another one!

10.4 Lemma. *We have* $\overline{\mathrm{im}(f)} = \ker(f^*)^\perp$.

Proof. We immediately obtain $\mathrm{im}(f) \subseteq \ker(f^*)^\perp$ from the adjointness of f^* and f with respect to the evaluation map. Therefore $\ker(f^*)^\perp$ also contains $\overline{\mathrm{im}(f)}$, as the former is closed. For the converse, suppose that there exists an element y in $\ker(f^*)^\perp \smallsetminus \overline{\mathrm{im}\, f}$. Hahn–Banach gives us a linear functional y^* in $(\overline{\mathrm{im}(f)})^\perp$ satisfying $y^*(y) = 1$. We see that y^* lies in $\ker(f^*)^\perp$, yet y lies in $\ker(f^*)$, so y^* has no right being nonzero on y. □

Let's now state the *Fredholm alternative* and see how far I get.

10.5 Theorem (Fredholm). *Let X be a Banach space, and suppose that $f : X \longrightarrow X$ is a compact operator.*

(1) $\ker(\mathrm{id} - f)$ *is finite-dimesional,*

(2) $\mathrm{im}(\mathrm{id} - f) = \ker(\mathrm{id} - f^*)^\perp$,

(3) $\ker(\mathrm{id} - f) = 0$ *if and only if* $\mathrm{im}(\mathrm{id} - f) = X$,

(4) $\dim \ker(\mathrm{id} - f) = \dim \ker(\mathrm{id} - f^*)$.

Part (3) of Theorem 10.5 is like the rank–nullity theorem. Why do we care about the Fredholm alternative? Because it allows us to solve eigenvalue equations.

10.6 Remark. The kernel $\ker(\mathrm{id} - f)$ is nontrivial if and only if $f(x) = x$ has a nonzero solution. Therefore part (1) of Theorem 10.5 says that any eigenvalue of a compact operator appears with finite multiplicity. We'll discuss and come to appreciate the other parts later, but for now let's just notice that (1) is like the spectral theorem.

We won't get to finish all the parts today, but let's at least tackle part (1).

Proof of (1). Let Y be the kernel of $\mathrm{id} - f$. Then Y is a closed subspace of X, and the equation $f(x) = x$ shows that it lies in $\mathrm{im}(f)$. Thus $\overline{B}_1^Y$ is a closed subset of the set $\overline{f(\overline{B}_1^X)}$, and the latter is compact by the compactness of f. So $\overline{B}_1^Y$ is compact, and Corollary 3.3 implies that Y is finite-dimensional. □

So we've found an actual use for the funhouse fact of Corollary 3.3.

11 January 29, 2018

Recall that we were in the middle of proving the Fredholm alternative, though we really only proved a tiny piece.

Proof of Theorem 10.5. Recall that we did part (1) last time.

(2) By Lemma 10.4, we have $\overline{\mathrm{im}(\mathrm{id} - f)} = \ker(\mathrm{id} - f^*)^\perp$. Therefore it suffices to prove that $\mathrm{im}(\mathrm{id} - f)$ is a closed subset of X. Suppose we have a sequence $x_n - f(x_n)$ that converges to u in X. Our idea will be to modify x_n by an element of $\ker(\mathrm{id} - f)$. Let d_n be the distance from x_n to $\ker(\mathrm{id} - f)$. Part (1) indicates that $\ker(\mathrm{id} - f)$ is finite-dimensional, so this distance d_n is actually achieved by some y_n in $\ker(\mathrm{id} - f)$.

I claim that d_n is bounded, for if it weren't, we could replace it with a subsequence such that $d_n \to \infty$. Let

$$z_n = \frac{x_n - y_n}{\|x_n - y_n\|}.$$

Then we have

$$(\mathrm{id}-f)z_n = d_n^{-1}(\mathrm{id}-f)x_n \to 0 \cdot u = 0.$$

By compactness, replace z_n with a subsequence such that $f(z_n) \to z$ for some z in X. Therefore $z_n \to z$ and $z - f(z) = 0$ by the above, indicating that z lies in $\ker(\mathrm{id}-f)$. But then we get

$$1 = d_n^{-1} \cdot d(x_n, \ker(\mathrm{id}-f)) = d(d_n^{-1}x_n, \ker(\mathrm{id}-f)) = d(z_n, \ker(\mathrm{id}-f)) \leq \|z_n - z\| \to 0,$$

which is a contradiction.

Now that we know that $x_n - y_n$ is bounded, the compactness of T allows us to replace $x_n - y_n$ with a subsequence such that $x_n - y_n \to v$ for some v in X. Then

$$x_n - f(x_n) = (x_n - y_n) - f(x_n - y_n) \implies u = v + u - T(v + u)$$

by taking $n \to \infty$, which shows that $\mathrm{im}(\mathrm{id}-f)$ is indeed closed.

(3) Suppose that $\ker(\mathrm{id}-f)$ is trivial, and for a contradiction suppose that $Y := \mathrm{im}(\mathrm{id}-f)$ is not all of X. Part (2) indicates that Y is closed and hence Banach. Furthermore, note that $f(Y) \subseteq Y$, so the restriction of f to Y is a compact operator on Y. The injectivity of $\mathrm{id}-f$ indicates that $(\mathrm{id}-f)(Y)$ does not equal Y, so we can iteratively form a strictly descending chain $Y_n := (\mathrm{id}-f)^n X$ of Banach spaces. The Riesz lemma allows us to choose y_n in Y_n satisfying $\|y_n\| = 1$ and $d(y_n, Y_{n+1}) \geq \frac{1}{2}$. If $n > m$, then the equation

$$f(y_n) - f(y_m) = (y_m - f(y_m)) - (y_n - f(y_n)) + y_n - y_m$$

shows that $\|f(y_n) - f(y_m)\| \geq \frac{1}{2}$. Therefore $f(y_n)$ has no Cauchy subsequence, and the completeness of X shows that f cannot be compact.

Conversely, suppose that $\mathrm{im}(\mathrm{id}-f) = X$. Then $\ker(\mathrm{id}-f^*) = \mathrm{im}(\mathrm{id}-f)^\perp = 0$, which implies that $\mathrm{im}(\mathrm{id}-f^*) = X^*$ by the above. Finally, we have $\ker(\mathrm{id}-f) = \mathrm{im}(\mathrm{id}-f^*)^\perp = 0$, as desired.

(4) We want to show that the dimensions $d := \dim\ker(\mathrm{id}-f)$ and $d^* := \ker(\mathrm{id}-f^*)$ are equal. For a contradiction, suppose that $d < d^*$. If we were doing more Fredholm theory, this proof would seem more natural, but as is it seems rather strange. Of course, Schauder's theorem implies that d^* is finite, via applying part (1). By using algebraic bases, we can obtain a bounded linear map $P : X \longrightarrow \ker(\mathrm{id}-f)$ that is a projector. Part (2) indicates that $\mathrm{im}(\mathrm{id}-f) = \ker(\mathrm{id}-f^*)^\perp$ has codimension d^* in X, so we may pick a subspace Z of X of dimension d^* that trivially intersects $\mathrm{im}(\mathrm{id}-f)$. Because $d < d^*$, we may choose a linear map $\Lambda : \ker(\mathrm{id}-f) \longrightarrow Z$ that is injective but not surjective. Consider the compact operator

$$S := f + \Lambda \circ P,$$

which is compact because $\Lambda \circ P$ has finite rank. I claim that $\ker(\mathrm{id}-S) = 0$. To see this, for any x in $\ker(\mathrm{id}-S)$, we have

$$0 = x - S(x) = \underbrace{(x - f(x))}_{\in \mathrm{im}(\mathrm{id}-f)} - \underbrace{\Lambda(P(x))}_{\in Z}.$$

But $\mathrm{im}(\mathrm{id}-f)$ and Z have trivial intersection, so $x - f(x) = 0$ and $\Lambda(P(x)) = 0$. The first equation indicates that x lies in $\ker(\mathrm{id}-f)$, and the second equation further indicates that $x = 0$.

From here, part (3) implies that $\operatorname{im}(\operatorname{id} - S) = X$. The non-surjectivity of Λ allows us to find y in Z not lying in its image, but it's not possible to have

$$y = x - S(x) = (x - f(x)) + \Lambda(P(x)).$$

Therefore $d \geq d^*$. Conversely, this inequality already shows that

$$\dim \ker(\operatorname{id} - f^{**}) \leq \dim \ker(\operatorname{id} - f^*) \leq \dim \ker(\operatorname{id} - f),$$

but we have $\ker(\operatorname{id} - f) \subseteq \ker(\operatorname{id} - f^{**})$ automatically. Therefore all the inequalities collapse to equalities. □

I wanted to prove the (real) *spectral theorem* today, but Fredholm alternatives just take too long. We'll do the real spectral theory first, even though we'll need to use and prove complex spectral theory later for applications to ergodic theory, because we'll use the real theory for that. Write $\mathscr{L}(X)$ for the space $\mathscr{L}(X, X)$ of bounded linear maps $X \longrightarrow X$.

11.1 Definition. Let f be in $\mathscr{L}(X)$.

- The *resolvent set* of f is $\rho(f) \coloneqq \{\lambda \in \mathbb{R} : f - \lambda \operatorname{id} \text{ is a bijection}\}$.
- The *spectrum* of f is $\sigma(f) \coloneqq \mathbb{R} \smallsetminus \rho(f)$.
- The *eigenvalues* of f is $\operatorname{ev}(f) \coloneqq \{\lambda \in \mathbb{R} : \ker(f - \lambda \operatorname{id}) \text{ is nontrivial}\}$.

Thus the eigenvalues are a subset of the spectrum, but this containment can be strict.

12 January 31, 2018

Let's begin with an example of the new notions we introduced at the end of last time.

12.1 Example. Let X be ℓ^2, and let $T : X \longrightarrow X$ be the operator sending $(x_1, x_2, \cdots) \mapsto (0, x_1, x_2, \ldots)$. Then T is not surjective, which shows that 0 lies in $\rho(T)$. However, T is injective, so 0 does not lie in $\operatorname{ev}(T)$.

12.2 Lemma. *The resolvent set $\sigma(T)$ lies in $[-\|T\|, \|T\|]$, and it is closed.*

Proof. Suppose that $|\lambda| > \|T\|$. We want to show that $T - \lambda \operatorname{id}$ is a bijection. Consider the equation

$$(T - \lambda \operatorname{id})x = y \iff x = \lambda^{-1}(Tx - y).$$

for any y in X. Because $\|\lambda^{-1}T\| < 1$, the right-hand side yields a map that is a contradiction in X. Thus the Banach fixed-point theorem indicates that there is a unique solution x in X. Thus $T - \lambda \operatorname{id}$ is bijective.

To show that $\sigma(T)$ is closed, we shall show that its complement $\rho(T)$ is open. Fix a λ_0 in $\rho(T)$, and try to solve the above equation again. By rewriting it as

$$(T - \lambda_0 \operatorname{id})x + (\lambda_0 - \lambda)x = y \iff x = (T - \lambda_0 \operatorname{id})^{-1}((\lambda - \lambda_0)x + y),$$

we note that the right-hand side is a contraction if $\|\lambda - \lambda_0\|\|(T - \lambda_0 \operatorname{id})^{-1}\| < 1$. Therefore $\rho(T)$ contains an open neighborhood of λ_0, as desired. □

Let us now turn to the spectral theorem.

12.3 Theorem. *Let X be an infinite-dimensional Banach space, and let $T : X \longrightarrow X$ be a compact operator.*

(1) 0 lies in $\sigma(T)$,

(2) $\sigma(T) \smallsetminus \{0\} = \mathrm{ev}(T) \smallsetminus \{0\}$,

(3) if $\sigma(T) \smallsetminus \{0\}$ is infinite, then $\sigma(T) \smallsetminus \{0\}$ consists of a countable set of numbers converging to zero.

Proof.

(1) If 0 lies in $\rho(T)$, then T is bijective. Thus the identity map $\mathrm{id} = T \circ T^{-1}$ is compact (as the inverse and composition of compact operators remains compact), so $\overline{B}_1^X$ is compact. Corollary 3.3 then concludes that X is finite-dimensional.

(2) If λ is nonzero and does not lie in $\mathrm{ev}(T)$, then $\ker(T - \lambda\,\mathrm{id})$ is trivial. Therefore $\ker(-\lambda^{-1}T + \mathrm{id})$ is trivial, so the Fredholm alternative implies that $\mathrm{im}(-\lambda^{-1}T + \mathrm{id})$ is all of X. Therefore $\mathrm{im}(T - \lambda\,\mathrm{id})$ is also all of X, which altogether shows that $T - \lambda\,\mathrm{id}$ is a bijection. Hence λ lies in $\rho(T)$.

(3) By Lemma 12.2, it suffices to show that 0 is the only accumulation point of $\sigma(T)$ is 0. Let λ_n be a sequence in $\sigma(T)$ of distinct elements that converges to some λ in $\sigma(T)$. By skipping ahead in the sequence and using distinctness, we may assume that all the λ_n are nonzero. By part (2), we may find x_n in X satisfying $\|x_n\|$ and $Tx_n = \lambda_n x_n$.

I claim that $\{x_1, \dots, x_n\}$ is linearly independent. To see this, we can induct on n and use the fact that the x_i are eigenvectors for T with distinct eigenvalues. At this point, the proof will follow a strategy similar to our proof of the Fredholm alternative. Let X_n be the span of $\{x_1, \dots, x_n\}$. Then $X_1 \subsetneq X_2 \subsetneq \cdots$ is a strictly ascending chain of finite-dimensional subspaces of X, so we may choose y_n in X_n satisfying $\|y_n\| = 1$ and $d(y_n, X_{n-1}) \geq \frac{1}{2}$ by the Riesz lemma. Observe that $(T - \lambda_n\,\mathrm{id})$ sends X_n to X_{n-1}, because this operator acts by scaling on $\{x_1, \dots, x_{n-1}\}$ and annihilates x_n. For $n > m$, we have

$$\left\| \frac{Ty_n}{\lambda_n} - \frac{Ty_m}{\lambda_m} \right\| = \left\| \underbrace{\frac{(T-\lambda_n)y_n}{\lambda_n}}_{\in X_{n-1}} + \underbrace{\frac{(T-\lambda_m)y_m}{\lambda_m}}_{\in X_{m-1}} - y_m + y_n \right\| \geq \frac{1}{2}.$$

Therefore Ty_n/λ_n has no Cauchy subsequence. As T is compact, the y_n/λ_n are unbounded. Since $\|y_n\| = 1$, we have $\liminf_{n\to\infty} |\lambda_n| = 0$, and because $\lambda_n \to \lambda$, we see that $\lambda = 0$. ☐

The following example will form all the examples of compact operators that we have at the moment.

12.4 Example. If λ_n is any sequence of real numbers that converges to zero, we can define an operator $T : \ell^2 \longrightarrow \ell^2$ by sending $(x_1, x_2, \dots) \mapsto (\lambda_1 x_1, \dots, \lambda_2 x_2, \dots)$. I claim that T is compact and that $\mathrm{ev}(T) = \{\lambda_n\}_n$. Now T is compact because it is the limit of its finite-rank truncations, and its resolvent set clearly contains $\{\lambda_n\}_n$. As $\mathrm{ev}(T)$ clearly contains $\{\lambda_n\}_n$, we see that everything becomes equal to what they should be.

When we solve the wave equation later, we'll obtain another collection of examples. Next, let's discuss self-adjoint operators.

12.5 Lemma. *Let H be a Hilbert space, and let $T : H \longrightarrow H$ be a self-adjoint operator. Form the numbers*

$$M^- := \inf_{\|x\|=1} \langle Tx, x\rangle \text{ and } M^+ := \sup_{\|x\|=1} \langle Tx, x\rangle.$$

Then $\{M^-, M^+\}$ lies in $\sigma(T)$, which in turn lies in $[M^-, M^+]$. Furthermore, we have $\|T\| = \max\{M^+, -M^-\}$.

The proof will resemble our proof of Lemma 12.2, except we'll be using variational methods, for which the Riesz representation theorem will replace our use of the Banach fixed point theorem.

Start of the proof of Lemma 12.5. Suppose that λ is greater than M^+, and define

$$\beta(x,y) \coloneqq \lambda\langle x,y\rangle - \langle Tx,y\rangle = \langle(\lambda - T)x,y\rangle.$$

for all x and Y in H. Then β is a perfectly good bilinear form on H, and we have

$$(\lambda - M^+)\|x\|^2 \leq \beta(x,x) \leq (\lambda - M^-)\|x\|^2,$$

which follows for the definitions of $M^\pm$ via renormalizing x to satisfy $\|x\| = 1$. Therefore the linear functionals induced by β are continuous. For any x in H, Riesz representation yields a unique y in H satisfying $\beta(x,\cdot) = \langle y,\cdot\rangle$. This exhibits the bijectivity of $T - \lambda\,\mathrm{id}$. □

The $\lambda < M^-$ case is entirely analogous, and we'll finish the entire proof of Lemma 12.5 next time.

13 February 2, 2018

Let's continue with the spectral theorem for compact operators—we'll then move to general operators.

Continuation of the proof of Lemma 12.5. Next, we want to show that the endpoints $\{M^-, M^+\}$ lie in $\sigma(T)$. First, define

$$\alpha(x,y) \coloneqq \langle(M^+ - T)x,y\rangle.$$

We see that $\alpha(x,x)$ is always non-negative by the definition of M^+, and we have

$$\alpha(x,y) \leq \alpha(x,x)^{1/2}\alpha(y,y)^{1/2}$$

by Cauchy–Schwartz. But this simply becomes

$$\langle(M^+ - T)x,y\rangle \leq \langle(M^+ - T)x,x\rangle^{1/2}\underbrace{\langle(M^+ - T)y,y\rangle^{1/2}}_{\leq C\|y\|},$$

and taking the supremum over y satisfying $\|y\| = 1$ implies that

$$\|(M^+ - T)x\| \leq C\langle(M^+ - T)x,x\rangle^{1/2}.$$

Now choose a sequence x_n with norm 1 satisfying $\langle x_n, Tx_n\rangle \to M^+ = \langle x_n, M^+x_n\rangle$. Then

$$\|(M^+ - T)x_n\| \leq C\langle(M^+ - T)x_n,x_n\rangle^{1/2} \to 0,$$

which implies that if $(M^+ - T)^{-1}$ existed, it could not possibly be bounded. Therefore $M^+ - T$ cannot exist, by the open mapping theorem. As before, the same argument works for M^- too.

Our next step will be in proving that $\|T\| \leq \max\{M^+, -M^-\}$, for which we need the self-adjointness of T. We have

$$\begin{aligned}
\langle x+y, T(x+y)\rangle &= \langle x,Tx\rangle + 2\langle x,Ty\rangle + \langle y,Ty\rangle,\\
\langle x-y, T(x-y)\rangle &= \langle x,Tx\rangle - 2\langle x,Ty\rangle + \langle y,Ty\rangle,\\
\implies 4\langle x,Ty\rangle &= \langle x+y, T(x+y)\rangle - \langle x-y, T(x-y)\rangle \leq M^+\|x+y\|^2 - M^-\|x-y\|^2\\
&\leq \max\{M^+, -M^-\}(\|x+y\|^2 + \|x-y\|^2).
\end{aligned}$$

The parallelogram rule then implies that

$$2\langle x, Ty\rangle \leq \max\{M^+, -M^-\}(\|x\|^2 + \|y\|^2).$$

Replacing x with εx, replacing y with $\varepsilon^{-1}y$, and optimizing ε implies that

$$\langle x, Ty\rangle \leq \max\{M^+, -M^-\}\|x\|\|y\|,$$

and setting $x = Ty$ yields the desired inequality. $\square$

With Lemma 12.5 in hand, we can proceed to the spectral theorem for self-adjoint compact operators on Hilbert spaces.

13.1 Theorem. *Let H be a Hilbert space, and let $T : H \longrightarrow H$ be a compact self-adjoint operator. Then there exist $\{x_n\}_n$ in H and $\{\lambda_n\}_n$ in $\mathbb{R}$ such that*

(1) $\{\lambda_n\}_n = \sigma(T) \smallsetminus \{0\}$,

(2) $Tx_n = \lambda_n x_n$, *and the x_n are orthonormal,*

(3) $H = \overline{\operatorname{span}\{x_n\}_n} + \ker T$.

Proof.

Step 1. If $\lambda \neq \lambda'$ are both in $\sigma(T)$, then $\ker(T - \lambda\,\mathrm{id})$ is orthogonal to $\ker(T - \lambda'\,\mathrm{id})$, which I hope you remember how to prove from linear algebra class.

Step 2. Because T is compact, Theorem 12.3 implies that $\sigma(T) \smallsetminus \{0\} = \mathrm{ev}(T) \smallsetminus \{0\} = \{\lambda_n\}_n$ for some $\{\lambda_n\}_n$. Furthermore, using Fredholm as in the proof of Theorem 12.3, we see that $\ker(T - \lambda\,\mathrm{id})$ is finite-dimensional for λ in $\sigma(T) \smallsetminus \{0\}$. We can choose orthonormal $\{x_n\}_n$ in X satisfying

$$\operatorname{span}\{x_n\}_n = Y := \operatorname{span} \bigcup_{\lambda \in \sigma(T) \smallsetminus \{0\}} \ker(T - \lambda\,\mathrm{id}).$$

To prove (3), it's enough to show that $Y^\perp$ equals $\ker T$. Observe that $T(Y)$ equals Y, so self-adjointness implies that $T(Y^\perp)$ lies in $Y^\perp$. That is,

$$\begin{aligned} T(Y^\perp) &= \{Tx : \langle x, y\rangle = 0 \text{ for all } y \in Y\} = \{Tx : \langle x, Ty\rangle = 0 \text{ for all } y \in Y\} \\ &= \{Tx : \langle Tx, y\rangle = 0 \text{ for all } y \in Y\} \subseteq \{x : \langle x, y\rangle = 0 \text{ for all } y \in Y\} = Y^\perp. \end{aligned}$$

Therefore $T|_{Y^\perp}$ is a compact operator on the Hilbert space $Y^\perp$, and by construction $T|_{Y^\perp}$ has no nonzero eigenvalues. By Lemma 12.5, we conclude that $T|_{Y^\perp} = 0$. $\square$

Let's now discuss spectral theory in the situation where we omit compactness hypotheses. We'll start for a couple of minutes, but then we'll have to wait a week to pick it back up, as you'll have an exam as well as a substitute lecture (which will *not* be on spectral theory) next week.

For the spectral theorem of general self-adjoint operators, we shall need to develop the *functional calculus*. You may recall from undergraduate ODEs class that matrix exponentials can be defined both in terms of a solution to a differential equation as well as a Taylor series. We'll adapt the latter strategy to applying functions to operators. The motivating idea is to use polynomial functions to approximate arbitrary operators using compact ones.

13.2 Definition. Let $p(\lambda) = a_0 + \cdots + a_n\lambda^n$ be a one-variable polynomial with real coefficients, let H be a Hilbert space, and let $T : H \longrightarrow H$ be a bounded linear operator. We define

$$p(T) \coloneqq a_0 \operatorname{id} + \cdots + a_n T^n,$$

where powers are given by composition. We see that $p(T)$ remains bounded linear.

For compact T, recall that $\sigma(T)$ is a compact subset of $[M^-, M^+]$.

13.3 Theorem. *If $p_n(\lambda)$ are polynomials that converge to a continuous function $f : C(\sigma(T)) \longrightarrow \mathbb{R}$, then the $p_n(T)$ converge in the operator norm. Denoting $\mathcal{P}$ for the subspace of polynomial functions in $C(\sigma(T))$, the map $\mathcal{P} \longrightarrow \mathscr{L}(H)$ given by $p \mapsto p(T)$ has a unique continuous extension $C(\sigma(T)) \longrightarrow \mathscr{L}(H)$.*

We'll come back to this in a week.

14 February 12, 2018

I haven't even *started* grading the midterms—last week was too crazy. Where were we? We were trying to talk about the spectral theory of self-adjoint operators, and we were trying develop the functional calculus. But why do we even need this? Recall that Theorem 13.1 gives us a spectral theorem for compact self-adjoint operators T on a Hilbert space H, part (3) of which implies that there exists a orthonormal set $\{e_n\}_n$ of H and real numbers $\{\lambda_n\}_n$ such that $Te_n = \lambda_n e_n$ and $H = \overline{\operatorname{span}\{e_n\}_n} + \ker T$. This implies that T is a (generally infinite) sum

$$T = \sum_{n=1}^{\infty} \lambda_n E_n,$$

where E_n denotes orthogonal projection $x \mapsto \langle x, e_n\rangle e_n$ onto the line generated by e_n. In the case of general, possibly non-compact T, we shall need to replace the above sum with an integral, and that's the reason for introducing the functional calculus.

Returning to the functional calculus itself, suppose that H is finite-dimensional, and let

$$p(\lambda) = c_0 + \cdots + c_n\lambda^n$$

be a polynomial in λ with entries in $\mathbb{R}$. One can form the polynomial

$$p(T) = c_0 + \cdots + c_n T^n$$

as usual, but one would like more than this: we want to form $f(T)$ for arbitrary continuous functions $f : \mathbb{R} \longrightarrow \mathbb{R}$. First, observe that if T is of the form

$$\begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_m \end{pmatrix},$$

then $p(T)$ is of the form

$$\begin{pmatrix} p(\lambda_1) & & \\ & \ddots & \\ & & p(\lambda_m) \end{pmatrix}.$$

Therefore, one way to form $f(T)$ for f in $C(\mathbb{R})$ is by

- taking polynomials p_n such that $p_n \to f$ uniformly on $\sigma(T) = \{\lambda_1, \ldots, \lambda_m\}$,
- forming the limit $f(T) \coloneqq \lim_{n\to\infty} p_n(T)$.

Note that this works well in the finite-dimensional case, because every self-adjoint operator is diagonalizable, i.e. conjugate to T of the above form.

14.1 Example. Furthermore, suppose that the eigenvalues of T are distinct, and order them via $\lambda_1 < \cdots < \lambda_m$. To compute the orthogonal projection E_k, just choose f in $C(\mathbb{R})$ such that

$$f(\lambda_j = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{otherwise.} \end{cases}$$

Then we see that $E_k = f(T)$.

For general Hilbert spaces, we won't be able to obtain E_k in this manner since the spectrum might not be discrete—we might not be able to choose an f that picks out exactly λ_k. This is where the measure theory comes in to patch our problems. Here, if we take f in $C(\mathbb{R})$ that approximates $\mathbf{1}_E$ sufficiently well for measurable $E \subseteq \mathbb{R}$, then $f(T)$ will approximate the integral of projections over the spectrum elements lying in E.

With all this motivation in hand, let's actually develop the functional calculus. We remind ourselves that throughout this discussion, H is a Hilbert space, and $T, S : H \longrightarrow H$ are self-adjoint operators on H. Recall the basic result Lemma 12.5, which allows us to further deduce the following useful fact:

14.2 Lemma. *We have* $\|T^n\| = \|T\|^n$.

Proof. For all x in H, Lemma 12.5 implies that

$$\|Tx\|^2 = \langle Tx, Tx\rangle = \langle x, T^2x\rangle \leq \|T^2\|\|x\|^2,$$

and taking square roots implies that $\|T\|^2 \leq \|T^2\|$. General principles imply that $\|T^2\| \leq \|T\|^2$ (because $T^2 = T \circ T$), which yields the result for $n = 2$. Iteratively applying this yields $\|T\|^{2^m} = \|T^{2^m}\|$, and applying the general principle $\|T \circ S\| \leq \|T\|\|S\|$ once more gives the general case. ☐

I might've not needed Lemma 14.2 just yet, but oh well.

14.3 Lemma. *We have* $\operatorname{dist}(\sigma(T), \sigma(S)) \leq \|T - S\|$.

Here, dist denotes the *Hausdorff metric*:

$$\operatorname{dist}(K, G) \coloneqq \max\left\{\max_{k\in K}\min_{g\in G}|k - g|, \max_{g\in G}\min_{k\in K}|g - k|\right\}$$

for any compact subsets K and G of $\mathbb{R}$. It turns out that dist yields a metric on

$$\{K \subseteq \mathbb{R} : K \text{ compact}\},$$

and this set is a complete metric space under the Hausdorff metric. Thus the content of Lemma 14.3 is that norm-convergent self-adjoint operators have Hausdorff-convergent spectra.

Proof. Suppose that $d \coloneqq \operatorname{dist}(\sigma(T), \sigma(S)) > \|T - S\|$. Without loss of generality, choose μ in $\sigma(T)$ such that $|\mu - \lambda| \geq d$ for all λ in $\sigma(S)$. Therefore $(\lambda - d, \lambda + d)$ lies in the resolvent set $\rho(S)$ of S. In particular, this implies that $S - \lambda\,\mathrm{id}$ is invertible, and we have $\|(S - \lambda\,\mathrm{id})^{-1}\| \leq \frac{1}{d}$.

We'll use a fixed-point argument similar to those of last time, but mixed up a little for variety. Write

$$T - \lambda\,\mathrm{id} = (S - \lambda\,\mathrm{id}) + (T - S) = \underbrace{(S - \lambda\,\mathrm{id})}_{\text{invertible}}(\mathrm{id} + \underbrace{(S - \lambda\,\mathrm{id})^{-1}(T - S)}_{\text{norm}<1}).$$

Therefore we can construct the inverse of $T - \lambda\,\mathrm{id}$ via

$$(T - \lambda\,\mathrm{id})^{-1} = \left[\sum_{n=0}^{\infty}\Big((S - \lambda\,\mathrm{id})^{-1}(T - S)\Big)^n\right](S - \lambda\,\mathrm{id})^{-1},$$

which contradicts the fact that μ lies in $\sigma(T)$. □

The next lemma is secretly the main theorem, though I will cheat a little in its proof. We'll discuss the cheat afterwards, and then you'll patch it up on your homework.

14.4 Lemma. *Let p by a polynomial with real entries.*

(1) $\sigma(p(T))$ is equal to $p(\sigma(T))$,

(2) $\|p(T)\| = \max_{\lambda\in\sigma(T)} |p(\lambda)|$.

Proof. Let's begin by proving that (1) $\implies$ (2). Lemma 12.5 indicates that

$$\|T\| = \max_{\lambda\in\sigma(T)} |\lambda|,$$

and then (1) gives us

$$\|p(T)\| = \max_{\lambda\in\sigma(p(T))} |\lambda| = \max_{\lambda\in p(\sigma(T))} |\lambda| = \max_{\mu\in\sigma(T)} |p(\mu)|.$$

It doesn't seem that I've cheated yet. Turning to (1), let $\nu := p(\mu)$, where μ is some element of $\sigma(T)$. We want to see that ν is in $\sigma(p(T))$. Factorize

$$p(\lambda) - \nu = (\lambda - \mu)q(\lambda) \implies p(T) - \nu\,\mathrm{id} = (T - \mu\,\mathrm{id})q(T),$$

which works because $p(\lambda) - \nu = 0$. Because $T - \mu\,\mathrm{id}$ is not a bijection, we see that $p(T) - \nu\,\mathrm{id}$ could not possibly be a bijection, where we use the fact that $T - \mu\,\mathrm{id}$ commutes with $q(T)$. This shows that $p(\sigma(T))$ lies in $\sigma(p(T))$.

Conversely, suppose that ν is not of the form $p(\mu)$ for any μ in $\sigma(T)$. Then the factorization of the polynomial $p(\lambda) - \nu$ into linear factors contains no roots in $\sigma(T)$, and the same argument shows that $p(T) - \nu\,\mathrm{id}$ is indeed a bijection. □

The illegal step in the above theorem in showing that $\sigma(p(T))$ is contained in $p(\sigma(T))$, because we can't generally factor every real polynomial into linear factors. You'll work on a workaround for this in your homework, and we'll finish the spectral theorem next time.

15 February 14, 2018

Recall the setup of last time, where H is a Hilbert space, and T and S are self-adjoint (i.e. symmetric) bounded operators on H. Let's begin today with a corollary of Lemma 14.4.

15.1 Lemma. *Let f be a continuous function on $\sigma(T)$, let ε be positive, and suppose that*

$$\max_{\sigma(T)} |f - p| \leq \varepsilon \text{ and } \max_{\sigma(T)} |f - q| \leq \varepsilon$$

for some polynomials p and q. Then $\|p(T) - q(T)\| \leq 2\varepsilon$, and $\operatorname{dist}(\sigma(p(T)), \sigma(q(T))) \leq 2\varepsilon$.

Proof. This follows immediately from parts (2) and (1) of Lemma 14.4, via the triangle inequality. □

Lemma 15.1 indicates that the following definition is well-defined, i.e. independent of choices.

15.2 Definition. Let f be in $C(\sigma(T))$. Then $f(T) \coloneqq \lim_{n\to\infty} p_n(T)$, where $\{p_n\}_n$ is any sequence of polynomials that converges to f uniformly on $\sigma(T)$.

Write $\mathscr{S}(H)$ for the set of symmetric bounded linear operators on H. Our existing work allows us to deduce the following theorem, which we call the functional calculus.

15.3 Theorem. *The map $f \mapsto f(T)$ is an isometric isomorphism*

$$C(\sigma(T)) \xrightarrow{\sim} \overline{\operatorname{span}_{\mathbb{R}}\{T^n : n \geq 0\}} \subseteq \mathscr{S}(H)$$

that commutes with addition and composition of functions. Furthermore, we have $\|f(T)\| = \max_{\sigma(T)} |f|$ and $\sigma(f(T)) = f(\sigma(T))$.

Proof. Surjectivity follows from Stone–Weierstrass, and well-definedness follows from Lemma 15.1. The other properties follow from Lemma 15.1 and Lemma 14.4. □

We have the following application of the functional calculus.

15.4 Definition. We say that T is *non-negative* if $\langle Tx, x\rangle \geq 0$ for all x in H.

15.5 Theorem. *The operator T is non-negative if and only if $\sigma(T)$ lies in $[0, \infty)$.*

Proof. For one direction, suppose that T is non-negative. Then

$$\sigma(T) \subseteq \left[\inf_{\|x\|=1} \langle Tx, x\rangle, \sup_{\|x\|=1} \langle Tx, x\rangle\right]$$

by Lemma 12.5, and this lies in $[0, \infty)$ by our assumption. In the other direction, suppose that $\sigma(T)$ lies in $[0, \infty)$. The functional calculus allows us to take a square root $S = \sqrt{T}$ of T, so

$$\langle Tx, x\rangle = \langle S^2 x, x\rangle = \langle Sx, Sx\rangle \geq 0,$$

as desired.[7] □

Now that we have the functional calculus, let's turn to the *spectral decomposition*.

15.6 Example. Consider the case when $H = \mathbb{R}^n$ is finite-dimensional, and let T be of the form

$$\begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix}.$$

[7] Alternatively, we can use the result from Lemma 12.5 that $\inf_{\|x\|=1}\langle Tx, x\rangle$ lies in $\sigma(T)$ to deduce that $\langle Tx, x\rangle$ is always non-negative, which circumvents the use of functional calculus.

If we let $e_1, \ldots, e_n$ denote the standard basis, we see that

$$T = \sum_{k=1}^{n} \overbrace{\lambda_k}^{\text{weight}} \underbrace{e_k^* \otimes e_k}_{\text{projection}}.$$

We can view this sum as the integral of some measure on the space of projections. In this spirit, we can also treat $f(T)$ for any continuous function $f : \sigma(T) \longrightarrow \mathbb{R}$ similarly. Namely, we can write $f(T)$ as

$$\begin{pmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_n) \end{pmatrix} = \int_{\sigma(T)} \mathrm{d}\left(\sum_{k=1}^{n} (e_k \otimes e_k^*)\delta_{\lambda_k}\right) f,$$

where now the measure that we're integrating over is somehow projection-valued.

This motivates the following definition. We use $\mathcal{B}$ to denote the Borel σ-algebra.

15.7 Definition. Let K be a compact subset of $\mathbb{R}$. A *projection-valued measure on* K is a map $E : \mathcal{B}(K) \longrightarrow \mathscr{S}(H)$ such that

(1) $E(\varnothing) = 0$ and $E(K) = \mathrm{id}$,

(2) if A and B are disjoint elements of $\mathcal{B}(K)$, then $E(A \cup B) = E(A) + E(B)$,

(3) for all A and B in $\mathcal{B}(K)$, we have $E(A \cap B) = E(A)E(B)$.

By taking $A = B$, we see that condition (3) is what makes E valued in projections.

15.8 Exercise. This definition of projection-valued measures is enough to ensure that $\int_K \mathrm{d}E\, f$ is a well-defined element of $\mathscr{S}(H)$ for any f in $C(K)$.

Exercise 15.8 uses the idea that finite additivity is enough to integrate continuous functions on compact subsets. When it's all said and done, this is all really just a fancy version of Riesz representation.

15.9 Definition. Let x and y be in H. We write $\ell_{x,y} : C(\sigma(T)) \longrightarrow \mathbb{R}$ for the function defined by

$$\ell_{x,y}(f) \coloneqq \langle f(T)x, y\rangle.$$

The continuity of the bracket and the functional calculus indicate that $\ell_{x,y}$ is bounded, and it's immediately seen to be linear, so Riesz representation for $C(\sigma(T))$ indicates that there exists a measure $m_{x,y}$ on $\sigma(T)$ such that

$$\langle f(T)x, y\rangle = \int_{\sigma(T)} \mathrm{d}m_{x,y}\, f$$

for all f in $C(\sigma(T))$. This yields a function $m : H^2 \longrightarrow \mathscr{M}(\sigma(T))$ via sending $(x, y) \mapsto m_{x,y}$.

15.10 Lemma. *Our function* $m : H^2 \longrightarrow \mathscr{M}(\sigma(T))$

(1) is bilinear,

(2) is symmetric,

(3) satisfies $\|m_{x,y}\| \leq \|x\|\|y\|$*, where* $\|m_{x,y}\|$ *is the total variation norm of* $m_{x,y}$*, for all* x *and* y *in* H*,*

(4) satisfies $m_{x,x} \geq 0$ for all x in H.

Some of this should be immediate.

Proof. We obtain (1) and (2) via the uniqueness of $m_{x,y}$ as given by the Riesz representation theorem. That is, because the defining property $\langle f(T)x, y\rangle = \int_{\sigma(T)} \mathrm{d}m_{x,y}\, f$ of m is bilinear and symmetric, m itself retains these properties. The Riesz representation theorem also shows that $\|m_{x,y}\| = \|\ell_{x,y}\|$, and

$$|\ell_{x,y}(f)| = |\langle f(T)x, y\rangle| \leq \|f(T)\|\|x\|\|y\| = \max_{\sigma(T)} |f| \cdot \|x\|\|y\|,$$

and this maximum is the sup norm. Thus $\|\ell_{x,y}\| \leq \|x\|\|y\|$, which gives part (3). Finally, for non-negative f, the functional calculus shows us that $f(T)$ is non-negative. Therefore in this case we have $\int_{\sigma(T)} \mathrm{d}m_{x,x}\, f = \langle f(T)x, x\rangle \geq 0$, so $m_{x,x} \geq 0$. ☐

I always judge how much time I have left in class from the song that plays at the bell tower, but there seems to be variation in when it ends. I should make some sort of arrangement with the carillonneur or something.

15.11 Lemma. *If $b : H^2 \longrightarrow \mathbb{R}$ is a bilinear symmetric function that satisfies $|b(x,y)| \leq K\|x\|\|y\|$ for some positive K and for all x and y in H, then there exists a bounded symmetric operator $B : H \longrightarrow H$ satisfying $b(x,y) = \langle Bx, y\rangle$ and $\|B\| \leq K$.*

Proof. The Riesz representation theorem shows that, for all x in H, there exists a unique Bx in H such that $b(x,y) = \langle Bx, y\rangle$ for all y in H. The desired properties of B follow from the corresponding properties of b, where we use the uniqueness of Riesz representation once again. ☐

15.12 Example. For any Borel subset A of $\sigma(T)$, Lemma 15.10 indicates that the function $b_A(x,y) := m_{x,y}(A)$ satisfies the conditions of Lemma 15.11, so there exists an operator $E(A)$ in $\mathscr{S}(H)$ such that

$$\langle E(A)x, y\rangle = m_{x,y}(A).$$

We shall prove next time that the assignment $A \mapsto E(A)$ is indeed a projection-valued measure, and we'll also show that it satisfies the integration property $f(T) = \int_{\sigma(T)} \mathrm{d}E\, f$ as in Example 15.6.

16 February 16, 2018

Recall the setup of Example 15.12. We shall prove that the assignment $A \mapsto E(A)$ from Example 15.12 is indeed a projection-valued measure.

16.1 Theorem. *The following holds:*

(1) $E(\varnothing) = 0$,

(2) $E(\sigma(T)) = \mathrm{id}$,

(3) if A and B are disjoint Borel subsets of $\sigma(T)$, then $E(A \cup B) = E(A) + E(B)$,

(4) for all Borel subsets A and B of $\sigma(T)$, we have $E(A)E(B) = E(A \cap B)$.

16.2 Remark. Morally, the operator $E(A)$ should be $\mathbf{1}_A(T)$, even though the right-hand side is undefined since $\mathbf{1}_A$ is usually not continuous. If it *were* well-defined, however, and satisfied the functional calculus, then Theorem 16.1 would follow immediately.

Maybe we could try to just define $\mathbf{1}_A(T)$ as a limit of $f(T)$ for continuous f on $\sigma(T)$, but then you have to do work to show that the limit exists and is well-defined. Bühler–Salamon takes this approach, though it requires the axiom of choice. Our approach offers an alternative.

Proof. I claim that parts (1) through (3) are easy. We won't do all of them, but here's the proof of (2) as a sample: for any x and y in H, we have

$$\langle E(\sigma(T))x, y\rangle = m_{x,y}(\sigma(T)) = \int_{\sigma(T)} \mathrm{d}m_{x,y}\, 1 = \langle 1(T)x, y\rangle = \langle x, y\rangle.$$

The uniqueness of these conditions shows then that $E(\sigma(T)) = \mathrm{id}$.

The interesting part is proving (4). For this, let us recall a measure-theoretic result from last quarter—namely, if μ is a Borel measure and A is a Borel subset, then

$$\mu(A) = \sup\{\mu(F) : F \subseteq A \text{ is closed}\} = \inf\{\mu(G) : G \supseteq A \text{ is open}\}.$$

The Borel σ-algebra is the smallest σ-algebra containing the open subsets, and the above result is one of the few places where one uses this minimality property of the Borel σ-algebra.

Next, we say that a sequence of functions $\{f_n\}_n$ in $C(\mathbb{R})$ *approximates* $\mathbf{1}_A$ if for all closed F inside A and open G containing A, we have $\mathbf{1}_F \leq f_n \leq \mathbf{1}_G$ for sufficiently large n. For any such sequence, we have

$$\mu(A) = \lim_{n\to\infty} \int \mathrm{d}\mu\, f_n.$$

Furthermore, one can readily construct such sequences. In our situation, applying this to $m_{x,y}$ implies that

$$\langle E(A)x, y\rangle = m_{x,y}(A) = \lim_{n\to\infty} \int \mathrm{d}m_{x,y}\, f_n = \lim_{n\to\infty} \langle f_n(T)x, y\rangle.$$

This says that $f_n(T)x$ weakly converges to $E(A)x$. It'd be nice to convert this weak convergence into strong convergence. Recall that Mazur's theorem shows that, for any fixed x, we can choose some convex combination of these $f_n(T)x$ that converges strongly to $E(A)x$. The linearity of $f \mapsto f(T)$ (as well as the fact that the *approximation* property is stable under convex combinations) shows that we can replace $\{f_n\}_n$ with some convex combinations that make $f_n(T)x \to E(A)x$ strongly.

Performing the same actions for a Borel subset B and y in H yields $g_n(T)y \mapsto E(B)y$ strongly in a similar manner. Observe that $f_n g_n$ approximates $\mathbf{1}_{A\cap B}$. Symmetry and the functional calculus imply that

$$\langle f_n(T)x, g_n(T)y\rangle = \langle g_n(T)f_n(T)x, y\rangle = \langle (f_n g_n)(T)x, y\rangle,$$

and taking $n \to \infty$ shows that

$$\langle E(A)E(B)x, y\rangle = \langle E(A)x, E(B)y\rangle = \langle E(A\cap B)x, y\rangle,$$

where we once symmetry once again. As this holds for all x and y in H, we obtain the desired result. $\square$

Theorem 16.1 is a nice result, but a priori we've not sure whether it's good for anything. Let's change that by computing some corollaries.

16.3 Corollary. *The following holds:*

(1) for all Borel subsets A, we have $E(A)^2 = E(A)$,

(2) we have $\|E(A)\| \in \{0, 1\}$,

(3) if A and B are disjoint Borel subsets of $\sigma(T)$, then $\mathrm{im}(E(A)) \perp \mathrm{im}(E(B))$,

(4) for all Borel subsets A and B of $\sigma(T)$, we have $E(A)E(B) = E(B)E(A)$.

Proof. I claim that these are all easy consequences of part (4) of Theorem 16.1.

(1) This follows by taking $A = B$.

(2) Part (1) implies that $\|E(A)\| = \|E(A)^2\| \leq \|E(A)\|^2$. Therefore either $\|E(A)\| = 0$ and hence $E(A) = 0$, or $\|E(A)\| \geq 1$. In the latter case, we know that $E(A)$ is a nonzero projection operator, so it has a fixed point and hence attains $\|E(A)\| = 1$.

(3) We have $\langle E(A)x, E(B)y\rangle = \langle E(B)E(A)x, y\rangle = \langle E(B \cap A)x, y\rangle = \langle E(\varnothing)x, y\rangle = 0$.

(4) This follows from part (4) of Theorem 16.1 by noting that $A \cap B = B \cap A$. □

We summarize our spectral theory as follows.

16.4 Theorem. *Let H be a Hilbert space, and let $T : H \longrightarrow H$ be a bounded symmetric operator on H. Then there exists a projection-valued measure E on $\sigma(T)$ such that, for all f in $C(\sigma(T))$, we have*

$$f(T) = \int_{\sigma(T)} \mathrm{d}E\, f,$$

where our integral converges in the operator norm.

Proof. If g is a simple function on $\sigma(T)$, say of the form

$$g = \sum_{k=1}^{n} a_k \mathbf{1}_{A_k}$$

for some real numbers $a_1, \dots, a_k$ and disjoint Borel subsets $A_1, \dots, A_k$, recall that the integral

$$\int_{\sigma(T)} \mathrm{d}E\, g \coloneqq \sum_{k=1}^{n} a_k E(A_k).$$

Corollary 16.3 indicates that

$$\left\| \int_{\sigma(T)} \mathrm{d}E\, g \right\| \leq \max_k |a_k|.$$

Therefore if g are $\widetilde{g}$ are two simple approximations of f, then we get

$$\left\| \int_{\sigma(T)} \mathrm{d}E\, g - \int_{\sigma(T)} \mathrm{d}E\, \widetilde{g} \right\| = \left\| \int_{\sigma(T)} \mathrm{d}E\, (g - \widetilde{g}) \right\| \leq \|g - \widetilde{g}\|.$$

Thus our integrals converge, and the limiting processes on both sides precisely yield the desired equality. □

It's fun and all to use things like Borel subsets, Riesz representation, and Mazur's theorem to prove Theorem 16.4, but why is this even meaningful? Why should our suite of results on spectral theory matter?

16.5 Example. We'd like to at least recover the usual spectral theory for finite-dimensional Hilbert spaces $H = \mathbb{R}^d$. By rescaling, let's assume that $\sigma(T)$ lies in $[0,1)$. For any positive integer n, we have the decomposition

$$[0,1) = \bigcup_{k=0}^{2^{n-1}} [k2^{-n}, (k+1)2^{-n})$$

of the interval into 2^n pieces. Applying E to both sides and using Corollary 16.3 shows that

$$\mathrm{id} = \sum_{k=0}^{2^{n-1}} E\left([k2^{-n}, (k+1)2^{-n})\right).$$

The finite-dimensionality of H indicates that at most d of these mutually orthonormal projections in the above sum are nonzero, so taking $n \to \infty$ shows that there are $m \leq d$ points $\lambda_1, \dots, \lambda_m$ in $[0,1)$ such that

$$E = \sum_{k=1}^{m} E(\{\lambda_k\})\delta_{\lambda_k}.$$

In other words, E is supported at the finitely many points $\lambda_1, \dots, \lambda_m$, and these points form the spectrum of T in the traditional linear algebraic sense. The point λ_k is weighted by the factor $E(\{\lambda_k\})$.

The following example shows that in general, E need not be a countably additive sort of "measure."

16.6 Example. Let $H = \ell^2$, and let T be given by sending

$$(x_1, x_2, x_3 \dots,) \mapsto \left(\frac{x_1}{1}, \frac{x_2}{2}, \frac{x_3}{3}, \dots\right).$$

Let $P_k : H \longrightarrow H$ denote projection onto the k-th coordinate. Then unrolling the definitions shows that

$$E = \sum_{k=1}^{\infty} P_k \delta_{1/k},$$

where we take this sum in the topology on $\mathscr{L}(H)$ generated by the maps $T \mapsto Tx$ for every x in H, even though this sum does *not* converge in the operator norm.

Next time, I'll prove the *Sobolev inequalities* so that we can make more examples of compact operators.

17 February 19, 2018

Today I want to start about something new, in the sense that it's orthogonal to what we've been discussing. I'll return to material from last quarter and prove the Sobolev inequalities, which will not only provide interesting examples of compact operators but also are probably more important than most things we've done in class. We'll also discuss the *isoperimetric inequality*.

Let's start with a very simple idea. Suppose we have a function u in $C_c^\infty(\mathbb{R}^d)$, and write $e_1, \dots, e_d$ for the standard basis of $\mathbb{R}^d$. Then the fundamental theorem of calculus tells us we can recover u from its partial derivatives via integrating on rays, because it vanishes at infinity:

$$u(x) = -\int_0^\infty \mathrm{d}t\, (D_i u)(x + te_i)$$

for any integer $1 \leq i \leq d$, where D_i denotes the partial derivative with respect to the e_i-coordinate.

Given that we can recover compactly supported smooth functions u from their derivatives $D_i u$, one might hope that we can relate the L^p-norms of u and $D_i u$. Namely, we hope that there exists a constant $C_{d,p,q}$ such that, for all u in $C_c^\infty(\mathbb{R}^d)$, we have

$$\|u\|_{L^q} \leq C_{d,q,p}\|Du\|_{L^p}$$

independently of the u chosen, where D denotes the gradient. Assuming that such a statement holds, let's see what would happen if we just dilated u by a positive r. Writing $u_r(x) \coloneqq u(rx)$, we have

$$\|u_r\|_{L^q} = \left[\int_{\mathbb{R}^d} \mathrm{d}x\, |u_r(x)|^q\right]^{1/q} = \left[\int_{\mathbb{R}^d} \mathrm{d}x\, |u(rx)|^q\right]^{1/q} = \left[r^{-d}\int_{\mathbb{R}^d} \mathrm{d}y\, |u(y)|^q\right]^{1/q} = r^{-d/q}\|u\|_{L^q}$$

by the change-of-variables integral rule and the substitution $y = rx$. The chain rule indicates that $(Du_r)(x) = r(Du)(rx)$, and applying the above calculation to this again yields

$$\|Du_r\|_{L^p} = r^{1-d/p}\|Du\|_{L^p}.$$

Therefore if our desired inequality is to hold, we must have

$$r^{-d/q}\|u\|_{L^q} = \|u_r\| \leq C_{d,q,p}\|Du_r\|_{L^p} = C_{d,q,p}r^{1-d/p}\|Du\|_{L^p}.$$

Then we must have $1 - d/p = -d/q$. The Sobolev inequality says that, when we also impose that $p \geq 1$ and $q \geq 1$ (which is reasonable, in order to obtain honest L^p-norms), this is all we need.

17.1 Theorem (Sobolev). *Let p and q lie in $[1,\infty)$, and suppose that $1/q = 1/p - 1/d$. Then there exists a positive constant $C_{d,q,p}$ such that*

$$\|u\|_{L^q} \leq C_{d,q,p}\|Du\|_{L^p}$$

for all u in $C_c^\infty(\mathbb{R}^d)$.

We'll take a meandering path to prove the Sobolev inequality, because we want to cover other things on the way. Write p^* for $pd/(d-p)$, and note that in the situation of the Sobolev inequality, we have $q = p^*$. Consider the special case where $p = 1$ and thus $q = d/(d-1)$, which we call the *simple Sobolev inequality*.

17.2 Exercise. Prove that the simple Sobolev inequality implies the full version, by applying this case to an appropriate power of u and using Hölder's inequality.

We won't prove this special case today—rather, we'll just show its equivalence to another statement. Recall the *isoperimetric inequality*, which is the following statement.

17.3 Theorem. *There exists a positive C_d such that, for any compact subset K of $\mathbb{R}^d$, we have*

$$\mathrm{vol}(K)^{(d-1)/d} \leq C_d \cdot \mathrm{per}(K),$$

where one has to make sense of volumes *and* perimeters *first.*

You've probably heard of the isoperimetric inequality in the context of proving that spheres have the most volume for a given surface area. We can obtain that sort of result from the above statement finding an optimal constant C_d.

Making sense of volumes and perimeters is more straightforward when K is a smooth submanifold of $\mathbb{R}^d$, but in general one has to be careful. The isoperimetric inequality was first proved in this context. We've actually cheated a bit by going to perimeters—we should be going to the area of the *boundary* instead, but then we also have to figure out what those words mean.

17.4 Definition. For us, the volume will just be the Lebesgue measure, and the perimeter will be

$$\mathrm{per}(K) := \inf_{\varepsilon>0} \inf \left\{ \int \mathrm{d}x\, |Du| : u \in C_c^\infty(\mathbb{R}^d),\ u = 1 \text{ on } K,\ u = 0 \text{ on } \mathbb{R}^d \smallsetminus (K + B_\varepsilon) \right\}$$

for any compact subset K of $\mathbb{R}^d$, where B_ε is the open ball of radius ε at the origin.

For those in the know, note that this is the $(d-1)$-th dimensional Hausdorff measure of K, which is infinite if the boundary of K doesn't have Hausdorff dimension $d-1$. To motivate this definition of perimeter, suppose that we're in the case when K has smooth boundary, and define

$$u_\varepsilon(x) := \max\left\{0, 1, 1 - \tfrac{1}{\varepsilon} d(x, K)\right\}$$

for any positive ε. Note that this u_ε lies in the subset of test functions considered in the definition of $\mathrm{per}(K)$.

17.5 Exercise. One can check that $\int \mathrm{d}x\, |Du_\varepsilon| \to \mathrm{per}(K)$ as $\varepsilon \to 0$.

I'm gonna cheat even more in our proofs.

17.6 Theorem. *The simple Sobolev inequality is equivalent to the isoperimetric inequality.*

Proof. Suppose that the simple Sobolev inequality holds. For any u in $C_c^\infty(\mathbb{R}^d)$ such that $u = 1$ on K, we have

$$C_d \int \mathrm{d}x\, |Du| \geq \left[\int \mathrm{d}x\, u^{d/(d-1)} \right]^{(d-1)/d} \geq \mathrm{vol}(K)^{(d-1)/d}.$$

Taking the infimum over all u in the definition of perimeter yields

$$C_d \mathrm{per}(K) \geq \mathrm{vol}(K)^{(d-1)/d},$$

as desired. The converse shall use two simple but important ideas from geometric measure theory: the *layer cake* representation and the *coarea formula*. Layer cake is the observation that

$$\int \mathrm{d}x\, |u|^{d/(d-1)} = \int_0^\infty \mathrm{d}t\ \mathrm{vol}\{|u|^{d/(d-1)} \geq t\},$$

which is a basic result in geometric measure theory as well as a real analysis exercise for students everywhere. This is essentially an infinitesimal version of the Lebesgue integral. Rewrite this as

$$\int_0^\infty \mathrm{d}t\ \mathrm{vol}\{|u|^{d/(d-1)} \geq t\} = \int_0^\infty \mathrm{d}t\ \mathrm{vol}\{|u| \geq t^{(d-1)/d}\} = \frac{d}{d-1} \int_0^\infty \mathrm{d}s\ \mathrm{vol}\{|u| \geq s\} s^{1/(d-1)},$$

where we have taken $s = t^{(d-1)/d}$. Next, the coarea formula says that

$$\int \mathrm{d}x\, |Du| = \int_0^\infty \mathrm{d}t\, \mathrm{per}\{|u| \geq t\}.$$

It is instructive to see what this means for the function $u(x) = \max\{1 - |x-1|, 0\}$ when $d = 1$. The isoperimetric inequality gives

$$\int_0^\infty \mathrm{d}t\, \mathrm{per}\{|u| \geq t\} \geq C_d^{-1} \int_0^\infty \mathrm{d}t\ \mathrm{vol}\{|u| \geq t\}^{(d-1)/d}.$$

Next, write $H(t) \coloneqq \mathrm{vol}\{|u| \geq t\}$. Our goal is to prove that

$$f(T) \coloneqq \left(\int_0^T \mathrm{d}t\, H(t)^{(d-1)/d}\right)^{d/(d-1)} \geq g(T) \coloneqq \frac{d}{d-1}\int_0^T \mathrm{d}t\, H(t) t^{1/(d-1)}$$

for all non-negative T, after which we'll be done, because it'd imply that

$$\int_0^\infty \mathrm{d}t \ \mathrm{vol}\{|u| \geq t\}^{(d-1)/d} \geq \left(\frac{d}{d-1}\int_0^\infty \mathrm{d}s \ \mathrm{vol}\{|u| \geq s\} s^{1/(d-1)}\right)^{(d-1)/d} = \left(\int \mathrm{d}x\, |u|^{d/(d-1)}\right)^{(d-1)/d}$$

by first taking $T \to \infty$ and then applying layer cake. The statement $f(0) \geq g(0)$ is immediate, and we shall prove that $f'(T) \geq g'(T)$ for all non-negative T. Showing this amounts to proving that

$$\frac{d}{d-1}\left(\int_0^T \mathrm{d}t\, H(t)^{(d-1)/d}\right)^{1/(d-1)} H(T)^{(d-1)/d} \geq \frac{d}{d-1} H(T) T^{1/(d-1)}$$

by the fundamental theorem of calculus, which reduces to showing that

$$\int_0^T \mathrm{d}t\, H(t)^{(d-1)/d} \geq H(T)^{(d-1)/d} T.$$

But this follows from the fact that $H(T)$ is a non-decreasing function, concluding the proof. □

Note that the proof of Theorem 17.6 shows that one can use the same constant C_d in both the Sobolev and isoperimetric inequalities. Next time, I won't just prove that two things are equivalent—I'll actually prove the isoperimetric inequality.

18 February 21, 2018

We're talking about the Sobolev inequalities, right? Recall from Theorem 17.6 (along with Exercise 17.2) that they're equivalent to the isoperimetric inequality. Also recall that we plan to take a scenic route to proving the Sobolev inequalities, though we shall also describe a simpler proof given in Brezis.

Proof of simple Sobolev. Recall from last time that we've already remarked that

$$u(x) = -\int_0^\infty \mathrm{d}t\, (D_i u)(x + te_i),$$

and taking the triangle inequality yields

$$|u(x)| \leq \int_0^\infty \mathrm{d}t\, |D_i u|(x + te_i) \leq \int_0^\infty \mathrm{d}t\, |Du|(x + te_i) \leq \int_{\mathbb{R}} \mathrm{d}y_i\, |Du|(x + y_i),$$

where y_i runs over $\mathbb{R}e_i$. Taking the product over integers $1 \leq i \leq d$ gives us

$$|u(x)|^{d/(d-1)} \leq \prod_{i=1}^d \left(\int_{\mathbb{R}} \mathrm{d}y_i\, |Du|(x + y_i)\right)^{1/(d-1)}.$$

Now, write $x_1, \ldots, x_d$ for the coordinates of x. Note that the y_1-factor in the above bound does not change as we vary x_1, so we can pull it out when integrating with respect to x_1:

$$\int_{\mathbb{R}} \mathrm{d}x_1\, |u(x)|^{d/(d-1)} \leq \left(\int_{\mathbb{R}} \mathrm{d}y_1\, |Du|(x + y_1)\right)^{1/(d-1)} \int_{\mathbb{R}} \mathrm{d}x_1 \prod_{i=2}^d \left(\int_{\mathbb{R}} \mathrm{d}y_i\, |Du|(x + y_i)\right)^{1/(d-1)}.$$

Applying Hölder's inequality to the $(d-1)$-fold product shows this is bounded above by

$$\left(\int_{\mathbb{R}} \mathrm{d}y_1\,|Du|(x+y_1)\right)^{1/(d-1)} \prod_{i=2}^{d}\left(\int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_i\,|Du|(x+y_i)\right)^{1/(d-1)}.$$

Using the same reasoning to pull out the y_2-factor when integrating over x_2, we obtain

$$\begin{aligned}\int_{\mathbb{R}} \mathrm{d}x_2 \int_{\mathbb{R}} \mathrm{d}x_1\,|u(x)|^{d/(d-1)} &\leq \left(\int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_2\,|Du|(x+y_2)\right)^{1/(d-1)} \int_{\mathbb{R}} \mathrm{d}x_2 \Bigg[\left(\int_{\mathbb{R}} \mathrm{d}y_1\,|Du|(x+y_1)\right)^{1/(d-1)} \\ &\qquad \cdot \prod_{i=3}^{d}\left(\int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_i\,|Du|(x+y_i)\right)^{1/(d-1)}\Bigg] \\ &\leq \left(\int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_2\,|Du|(x+y_2)\right)^{1/(d-1)} \left(\int_{\mathbb{R}} \mathrm{d}x_2 \int_{\mathbb{R}} \mathrm{d}y_1\,|Du|(x+y_1)\right)^{1/(d-1)} \\ &\qquad \cdot \prod_{i=3}^{d}\left(\int_{\mathbb{R}} \mathrm{d}x_2 \int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_i\,|Du|(x+y_i)\right)^{1/(d-1)}.\end{aligned}$$

By collapsing the integrals over x_j and y_j together for j in $\{1,2\}$, we see that this equals

$$\left(\int_{\mathbb{R}} \mathrm{d}x_2 \int_{\mathbb{R}} \mathrm{d}x_1\,|Du|(x)\right)^{2/(d-1)} \prod_{i=3}^{d}\left(\int_{\mathbb{R}} \mathrm{d}x_2 \int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_i\,|Du|(x+y_i)\right)^{1/(d-1)}.$$

Iteratively repeating this process shows that

$$\begin{aligned}\int_{\mathbb{R}} \mathrm{d}x_k \cdots \int_{\mathbb{R}} \mathrm{d}x_1\,|u(x)|^{d/(d-1)} &\leq \left(\int_{\mathbb{R}} \mathrm{d}x_k \cdots \int_{\mathbb{R}} \mathrm{d}x_1\,|Du|(x)\right)^{k/(d-1)} \\ &\qquad \cdot \prod_{i=k+1}^{d}\left(\int_{\mathbb{R}} \mathrm{d}x_k \cdots \int_{\mathbb{R}} \mathrm{d}x_1 \int_{\mathbb{R}} \mathrm{d}y_i\,|Du|(x+y_i)\right)^{1/(d-1)},\end{aligned}$$

and taking $k=d$ yields the desired result. □

The above proof is pretty quick, and it even tells you that you can take the constant $C_d = 1$ independently of the dimension d (though we remark that this is a very non-ideal constant—one can do much better). However, what I dislike about this proof is that it tells me nothing about the structure of our functions u.

Let's now work towards an alternative proof.

Another proof of simple Sobolev. In our first step, we could've actually used the fact that

$$u(x) = \int_0^\infty \mathrm{d}t\,\nu \cdot (Du)(x - t\nu)$$

for any unit vector ν in $\mathbb{R}^d$. Therefore we can integrate over all such ν to obtain

$$u(x) = \frac{1}{|\partial B_1|}\int_{\partial B_1} \mathrm{d}\nu \int_0^\infty \mathrm{d}t\,\nu \cdot (Du)(x - t\nu).$$

This expression basically amounts to an integral over polar coordinates. Setting $y = t\nu$ and changing back to rectangular coordinates then gives us

$$u(x) = \frac{1}{|\partial B_1|}\int_{\mathbb{R}^d} \mathrm{d}y\,\frac{y}{|y|^d} \cdot (Du)(x-y),$$

where this integral converges at infinity because u is compactly supported and near zero because $y/|y|^d$ is an L^1-function in any bounded region. Taking absolute values yields

$$|u(x)| \leq \frac{1}{|\partial B_1|} \int_{\mathbb{R}^d} \mathrm{d}y\, |x-y|^{1-d} |Du|(y),$$

where we made the change of variables $y \mapsto x - y$. Now Theorem 2.4, Example 2.9, and the density of compactly supported smooth functions in L^p indicates that

$$\|u\|_{L^{d/(d-1)}} = \sup_{\substack{v \in C_c^\infty(\mathbb{R}^d) \\ \|v\|_{L^d} \leq 1}} \int \mathrm{d}x\, uv.$$

Combining this with our previous inequality yields

$$\|u\|_{L^{d/(d-1)}} \leq \frac{1}{|\partial B_1|} \sup_{\substack{v \in C_c^\infty(\mathbb{R}^d) \\ \|v\|_{L^d} \leq 1}} \int_{\mathbb{R}^d} \mathrm{d}x \int_{\mathbb{R}^d} \mathrm{d}y\, v(x) |x-y|^{1-d} |Du|(y).$$

This leads us to the *Hardy–Littlewood–Sobolev inequality*.

18.1 Theorem (Hardy–Littlewood–Sobolev). *If p and q are greater than 1, λ lies in $(0, d)$, and $\frac{1}{p} + \frac{\lambda}{d} + \frac{1}{q} = 2$, then there exists a positive $C_{d,\lambda,p,q}$ such that, for all f and g in $C_c^\infty(\mathbb{R}^d)$, we have*

$$\int_{\mathbb{R}^d} \mathrm{d}x \int_{\mathbb{R}^d} \mathrm{d}y\, f(x) |x-y|^{-\lambda} g(y) \leq C_{d,p,q,\lambda} \|f\|_{L^p} \|g\|_{L^q}.$$

18.2 Exercise. Show that Hardy–Littlewood–Sobolev implies simple Sobolev.[8] □

I claim that Hardy–Littlewood–Sobolev is a more fundamental statement than just Sobolev, even though its statement is way more complex. The Hardy–Littlewood–Sobolev inequality relates the L^p- and L^q-norms of what's essentially the convolution of two functions with their original norms, and it just happens that our very singular process of taking derivatives (which yields the Sobolev inequality) fits the bill.

The proof of Hardy–Littlewood–Sobolev requires the *Vitali covering lemma*, which y'all unfortunately covered last quarter. I'm teaching measure theory next quarter, and I really wanted to give the proof!

18.3 Lemma (Vitali). *If $\mathcal{B}$ is a collection of balls, and $\sup_{B(x,r) \in \mathcal{B}} r < \infty$, then there exist a subcollection $\mathcal{B}' \subseteq \mathcal{B}$ consisting of disjoint balls such that*

$$\bigcup_{B \in \mathcal{B}} B \subseteq \bigcup_{B(x,r) \in \mathcal{B}'} B(x, 5r).$$

We'll use this to give some bounds on the *Hardy–Littlewood maximal function*, which you should also remember from last quarter.

18.4 Definition. Let u be a function in $C_c^\infty(\mathbb{R}^d)$. Then its *Hardy–Littlewood maximal function* is

$$(\mathcal{M}u)(x) \coloneqq \sup_{B(y,r) \ni x} \frac{1}{|B(y,r)|} \int_{B(y,r)} \mathrm{d}z\, |u|(z).$$

Therefore $\mathcal{M}u$ expresses how large the locally averaged values of u can be. Now recall some bounds.

[8] Plug in $\lambda = d - 1$, $f = v$, and $g = Du$.

18.5 Proposition. *The following bounds hold:*

- *for all positive λ, we have $|\{\mathcal{M}u > \lambda\}| \leq C_d\|u\|_{L^1}/\lambda$ for some C_d depending only on d,*
- *$\|\mathcal{M}u\|_{L^\infty} \leq \|u\|_{L^\infty}$, which is trivial,*
- *for any $1 < p \leq \infty$, we have $\|\mathcal{M}u\|_{L^p} \leq C_{d,p}\|u\|_{L^p}$, where $C_{d,p}$ only depends on d and p.*

We call the last point the *interpolated strong L^p bound*, as... that's exactly what it is.

Proof. It suffices to show this for non-negative u. Our first goal is to prove that

$$|\{\mathcal{M}u > \lambda\}| \leq C_d \frac{\|u \cdot \mathbf{1}_{\{u \geq \lambda/2\}}\|_{L^1}}{\lambda},$$

which would imply our first bullet point. Begin by covering $\{\mathcal{M}u > \lambda\}$ with balls $B(x,r)$ satisfying $\frac{1}{|B(x,r)|}\int_{B(x,r)} \mathrm{d}x\, u \geq \lambda$. Note that for such a ball $B(x,r)$, we have

$$\frac{1}{|B(x,r)|}\int_{B(x,r)} \mathrm{d}x\, u \cdot \mathbf{1}_{\{u \geq \lambda/2\}} \geq \lambda/2,$$

because taking the same integral with $\mathbf{1}_{\{u \leq \lambda/2\}}$ yields a number that is at most $\lambda/2$. Choose the $\mathcal{B}'$ corresponding to this covering $\mathcal{B}$ as in the Vitali covering lemma, which satisfies

$$|\{\mathcal{M}u > \lambda\}| \leq \left|\bigcup_{B \in \mathcal{B}} B\right| \leq \left|\bigcup_{B(x,r) \in \mathcal{B}'} B(x,5r)\right| \leq 5^d \left|\bigcup_{B' \in \mathcal{B}'} B'\right|.$$

From here, we get

$$\begin{aligned}\|u \cdot \mathbf{1}_{\{u \geq \lambda/2\}}\|_{L^1} &= \int_{\mathbb{R}^d} \mathrm{d}x\, u \cdot \mathbf{1}_{\{u \geq \lambda/2\}} \geq \sum_{B' \in \mathcal{B}'} \int_{B'} \mathrm{d}x\, u \cdot \mathbf{1}_{\{u \geq \lambda/2\}} \geq \sum_{B' \in \mathcal{B}'} |B'|\lambda/2 \\ &\geq (\lambda/2)\left|\bigcup_{B' \in \mathcal{B}'} B'\right| \geq (\lambda/2)5^{-d}|\{\mathcal{M}u > \lambda\}|.\end{aligned}$$

Setting $C_d = 2 \cdot 5^d$ yields the desired result. Skipping over the second bullet point (which is immediate) and heading to the third one, we have

$$\int_{\mathbb{R}^d} \mathrm{d}x\, (\mathcal{M}u)^p \leq 2^p \sum_{n \in \mathbb{Z}} |\{2^n < \mathcal{M}u \leq 2^{n+1}\}| 2^{np} \leq 2^p \sum_{n \in \mathbb{Z}} |\{2^n < \mathcal{M}u\}| 2^{np}.$$

We can then apply our intermediary result to bound the above by

$$2^p C_d \sum_{n \in \mathbb{Z}} \|u \cdot \mathbf{1}_{\{u \geq 2^{n-1}\}}\|_{L^1} 2^{n(p-1)} \sim \|u\|_{L^p},$$

where we use the fact that $p > 1$ to ensure that $p - 1 > 0$, and we collapse the overlapping $(p-1)$-th power terms together to form the necessary p-th power terms. $\square$

We'll talk next time more about applying Proposition 18.5 to the proof of Hardy–Littlewood–Sobolev.

19 February 23, 2018

Recall that while we proved the Sobolev inequality last time, we also wanted to use the Hardy–Littlewood–Sobolev inequality to give another proof. Of course, we had better actually prove the Hardy–Littlewood–Sobolev inequality first.

Proof of Hardy–Littlewood–Sobolev. This inequality is basically saying that if you convolve with a certain kernel, you can bound the result by the L^q-norm of something that's roughly the Hölder conjugate. It suffices to prove the inequality for non-negative f. Let's set up some notation: define the function

$$h(y) := \int \mathrm{d}x\, f(x)|x-y|^{-\lambda},$$

and take r such that $\frac{1}{r} + \frac{1}{q} = 1$. It suffices to show that $\|h\|_{L^r} \leq C\|f\|_{L^p}$ for some positive constant C independent of f, as we can then use Example 2.9 to obtain the final inequality. We'll use a number of ingredients to prove our desired result:

1. Layer cake for $|x-y|^{-\lambda}$: the layer cake here ends up showing that

$$h(y) = \int_0^\infty \mathrm{d}r\, r^{d-\lambda-1}\left[\frac{1}{|B(y,r)|}\int_{B(y,r)} \mathrm{d}x\, f(x)\right],$$

   which you can also show for any radially symmetric kernel in place of $|x-y|^{-\lambda}$.

2. Let's (immediately) bound our average from ingredient 1: we definitionally have

$$\frac{1}{|B(y,r)|}\int_{B(y,r)} \mathrm{d}x\, f(x) \leq \mathcal{M}f(y)$$

   for any positive r.

3. Now let's use a less trivial bound: applying Hölder's inequality to $f \cdot 1$ yields

$$\frac{1}{|B|}\int_B f \leq \frac{1}{|B|}\left(\int_B 1^{p/(p-1)}\right)^{(p-1)/p}\left(\int_B f^p\right)^{1/p} \leq |B|^{-1/p}\|f\|_{L^p} = Dr^{-d/p}\|f\|_{L^p},$$

   where D is some constant independent of f (coming from the formula for the volume of balls), and B is any ball of radius r.

This is a pretty classic trick in analysis: the bound in ingredient 3 decays very fast but is singular near 0, so near 0 we use the bound in ingredient 2. For any positive r^* (which serves as a cut-off between these two regimes), using ingredient 1 then nets us

$$\begin{aligned} h(y) &\leq \int_0^{r^*} \mathrm{d}r\, r^{d-\lambda-1}\mathcal{M}f(y) + D\int_{r^*}^\infty \mathrm{d}r\, r^{d-\lambda-1}r^{-d/p}\|f\|_{L^p} \\ &= \frac{1}{d-\lambda}r_*^{d-\lambda}\mathcal{M}f(y) + \frac{D}{d-\lambda-d/p}r_*^{d-\lambda-d/p}\|f\|_{L^p}, \end{aligned}$$

provided that $d-\lambda-1 > -1$ and $d-\lambda-1-d/p < -1$. The fact that λ lies in $(0,d)$ guarantees the first inequality, and

$$\frac{1}{p} + \frac{\lambda}{d} + \frac{1}{q} = 2 \implies \frac{d}{p} + \lambda + \frac{d}{q} = 2d \implies 0 > \frac{d}{q} - d = d - \lambda - \frac{d}{p}$$

guarantees the second inequality, so this calculation is valid. Setting $r_*^{d/p} = \|f\|_{L^p}/\mathcal{M}f(y)$ lets us combine the two above terms, and we obtain

$$h(y) \leq \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right) \|f\|_{L^p}^{(p/d)(d-\lambda)} \mathcal{M}f(y)^{1-(p/d)(d-\lambda)}.$$

Let's try to clean up the exponents a bit:

$$1 - \frac{p}{d}(d-\lambda) = 1 - p + \frac{\lambda p}{d} = 1 - p + \left(2p - 1 - \frac{p}{q}\right) = p - \frac{p}{q} = p \cdot \frac{q-1}{q} = \frac{p}{r},$$

where we used the fact that $1 + p\lambda/d + p/q = 2p$. Plugging this into our bound yields

$$\begin{aligned} h(y) &\leq \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right) \|f\|_{L^p}^{1-p/r} \mathcal{M}f(y)^{p/r} \\ \implies \int \mathrm{d}y\, h(y)^r &\leq \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right) \int \mathrm{d}y\, \|f\|_{L^p}^{r-p} \mathcal{M}f(y)^p \\ &= \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right) \|f\|_{L^p}^{r-p} \int \mathrm{d}y\, \mathcal{M}f(y)^p \\ &\leq \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right) C_{d,p}^p \|f\|_{L^p}^{r-p} \|f\|_{L^p}^p \end{aligned}$$

by the third bullet point of Proposition 18.5, so taking

$$C = \left(\frac{1}{d-\lambda} + \frac{D}{d-\lambda-d/p}\right)^{1/r} C_{d,p}^{p/r}$$

gives the desired result, concluding the proof of the Hardy–Littlewood–Sobolev inequality. □

Let's take a step back and see what we did. To prove our goal, we used our layer cake formulation to break our function down into two bounds, and we chose the way we broke it down (as well as our exponents) exactly to piece them together the right way. We're relating different L^p-norms, so it makes sense to try using Hölder's inequality. Altogether, this makes for a "natural" proof of the Sobolev inequality, which you could've come up with just by sitting down and trying it, say, if you were stuck on a desert island.

Next time, we'll use the Sobolev inequality to give examples of compact operators as well as finally answer our question of vibrating drums of arbitrary shape, as we have long promised.

20 February 26, 2018

I want to show today that the *inverse Laplacian* is compact. For this, I will need to introduce *Sobolev spaces*.

20.1 Definition. Let U be a bounded open subset of $\mathbb{R}^d$. For any positive integer k and real $1 \leq p < \infty$, we write $W^{k,p}(U)$ for the completion of $C^\infty(\overline{U})$ with respect to the norm

$$\|u\|_{W^{k,p}(U)} := \sum_{|\alpha| \leq k} \|D^\alpha u\|_{L^p(U)},$$

where α runs over multi-indices of size $|\alpha|$ at most k, and D^α denotes the corresponding partial derivative. Write $W_0^{k,p}(U)$ for the completion of $C_c^\infty(U)$ with respect to the same norm formula, where we view $C_c^\infty(U) \subseteq C^\infty(\overline{U})$ and hence $W_0^{k,p}(U) \subseteq W^{k,p}(U)$ via extension-by-zero. These are the *Sobolev spaces*. We also write $H^k(U)$ for $W^{k,2}(U)$, and we write $H_0^k(U)$ for $W_0^{k,2}(U)$.

You've seen the $d = 1$ and $U = (0, 1)$ version of these spaces already on your homework. Note that we have $H^0(U) = L^2(U)$ in particular. We think of $W_0^{k,p}(U)$ as the elements of $W^{k,p}(U)$ that are zero on ∂U.

20.2 Remark. I'm cheating a bit here—one usually first defines Sobolev spaces using *weak derivatives*, and then it's a theorem of Friedrichs that this definition matches the one we have given above. Because I don't want to discuss the measure-theoretic details of Friedrichs's theorem, I'll just define Sobolev spaces in terms of the result of Friedrichs's theorem.

Let's finally use the Sobolev inequalities to give examples of compact operators.

20.3 Theorem (Rellich). *If $1 \leq p < d$ and $1 \leq q < p^*$, then $W_0^{1,p}(U) \subseteq L^q(U)$, and the inclusion map is compact.*

Proof. For any u in $C_c^\infty(U)$, extension by zero allows us to interpret it as a function in $C_c^\infty(\mathbb{R}^d)$. The Sobolev inequality then yields

$$\|u\|_{L^{p^*}(U)} = \|u\|_{L^{p^*}(\mathbb{R}^d)} \leq \|Du\|_{L^p(\mathbb{R}^d)} = \|Du\|_{L^p(U)},$$

where we have taken the constant in the Sobolev inequality to be 1. Because U is bounded and $q < p^*$, we have $\|u\|_{L^q(U)} \leq |U|^{1/q - 1/p^*}\|u\|_{L^{p^*}(U)}$ by Hölder's inequality, and combining this with the above shows that our inclusion of interest is bounded.

To prove that our inclusion of interest is also compact, we first want to smooth out u. Let ε be a positive number, and consider the function

$$u_\varepsilon \coloneqq u * \rho_\varepsilon,$$

where $\rho_\varepsilon(x) \coloneqq \varepsilon^{-d}\rho(\varepsilon^{-1}x)$ for some non-negative bump function ρ in $C_c^\infty(B_1)$ with integral 1. Our goal is to show that $\|u_\varepsilon\|_{L^q}$, $\|D(u_\varepsilon)\|_{L^\infty}$, and $\|u_\varepsilon - u\|_{L^q}$ are bounded by some nice combination of ε and $\|u\|_{W^{1,p}}$. Given any h in $\mathbb{R}^d$, write $(\tau_h u)(x) \coloneqq u(x+h)$ for the translation of u by h. We have

$$\begin{aligned}
&(\tau_h u - u)(x) = u(x+h) - u(x) = \int_0^1 \mathrm{d}t\, h \cdot (Du)(x+th) \\
\implies\ &\|\tau_h u - u\|_{L^1} \leq \int_{\mathbb{R}^d} \mathrm{d}x \int_{[0,1]} \mathrm{d}t\, |h \cdot (Du)(x+th)| \leq |h| \int_{[0,1]} \mathrm{d}t \int_{\mathbb{R}^d} \mathrm{d}x\, |Du|(x+th) = |h|\|Du\|_{L^1}
\end{aligned}$$

by our favorite representation of u in terms of Du. Recall that the definition of convolution gives us

$$\begin{aligned}
u_\varepsilon(x) = (u * \rho_\varepsilon)(x) = \int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h)(\tau_h u)(x) \implies \|u_\varepsilon - u\|_{L^1} &= \int_{\mathbb{R}^d} \mathrm{d}x \left| \int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h)(\tau_h u)(x) - u(x) \right| \\
= \int_{\mathbb{R}^d} \mathrm{d}x \left| \int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h) \left[(\tau_h u)(x) - u(x)\right] \right| &\leq \int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h) \int_{\mathbb{R}^d} \mathrm{d}x \left|(\tau_h u)(x) - u(x)\right| \leq |\varepsilon|\|Du\|_{L^1},
\end{aligned}$$

because ρ_ε has integral 1 and is supported on $\overline{B}_\varepsilon$. By the Sobolev inequality, the interaction between L^p-norms and convolution, and the fact that ρ_ε has integral 1, we have

$$\|u_\varepsilon - u\|_{L^{p^*}} \leq \|D(u_\varepsilon - u)\|_{L^p} \leq \|Du * \rho_\varepsilon\|_{L^p} + \|Du\|_{L^p} \leq (\|\rho_\varepsilon\|_{L^1} + 1)\|Du\|_{L^p} = 2\|Du\|_{L^p}.$$

Our goal is to have control over the L^q-norm, and because $1 \leq q < p^*$, we can interpolate our bounds for 1 and p^* to obtain one for q. Apparently you didn't do interpolation of L^p-spaces last quarter, so you'll want to learn about that and then do the following exercise.

20.4 Exercise. Show that there exist some positive C and θ that depend only on d, p, and q such that

$$\|u_\varepsilon - u\|_{L^q} \leq C\varepsilon^\theta \|Du\|_{L^p}.$$

Furthermore, we have

$$|u_\varepsilon(x)| = \left|\int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h)u(x+h)\right| \implies \|u_\varepsilon\|_{L^\infty} \le \|\rho_\varepsilon\|_{L^\infty}\|u\|_{L^1} \le \|\rho\|_{L^\infty}\varepsilon^{-d}\|u\|_{W^{1,p}}$$

by the definition of ρ_ε in terms of ρ. We also have

$$(D(u_\varepsilon))(x) = D\int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h)u(x+h) = D\int_{\mathbb{R}^d} \mathrm{d}h\, \rho_\varepsilon(h-x)u(h) = -\int_{\mathbb{R}^d} \mathrm{d}h\, (D\rho_\varepsilon)(h-x)u(h)$$

by making the change of variables $h \mapsto h - x$ and applying integration by parts. Taking the L^∞-norm and decomposing ρ_ε in terms of ε and ρ therefore yields

$$\|D(u_\varepsilon)\|_{L^\infty} \le \|D\rho_\varepsilon\|_{L^\infty}\|u\|_{L^1} \le \|D\rho\|_{L^\infty}\varepsilon^{-1-d}\|u\|_{W^{1,p}}.$$

With all these ingredients in hand, let $\{u_n\}_n$ be a sequence in $W_0^{1,p}(U)$ with norm at most 1. By fixing ε and applying our bounds

$$\|u_\varepsilon\|_{L^\infty} \le \|\rho\|_{L^\infty}\varepsilon^{-d}\|u\|_{W^{1,p}} \text{ and } \|D(u_\varepsilon)\|_{L^\infty} \le \|D\rho\|_{L^\infty}\varepsilon^{-1-d}\|u\|_{W^{1,p}}$$

to the sequence $\{u_n\}_n$, we see that the mollified $\{u_n^\varepsilon\}_n$ is pointwise pre-compact and equicontinuous, respectively. Therefore we may apply Arzelà–Ascoli to $\{u_n^\varepsilon\}_n$, so we can replace $\{u_n\}_n$ with a subsequence $\{\widetilde{u}_n\}_n$ such that $\widetilde{u}_n^\varepsilon$ converges uniformly in $C(U)$. By taking $\varepsilon = 1/m$ and letting m vary over all positive integers and diagonalizing, we can further refine $\{u_n\}_n$ such that the mollified functions $u_n * \rho_{1/m}$ converge uniformly as $n \to \infty$.

I claim that $\{u_n\}_n$ is Cauchy in the L^q-norm. To see this, use the triangle inequality to take

$$\begin{aligned}
\|u_n - u_k\|_{L^q} &\le \|u_n - u_n * \rho_{1/m}\|_{L^q} + \|u_k - u_k * \rho_{1/m}\|_{L^q} + \|u_n * \rho_{1/m} - u_k * \rho_{1/m}\|_{L^q} \\
&\le Cm^{-\theta}(\|Du_n\|_{L^p} + \|Du_k\|_{L^p}) + \|u_n * \rho_{1/m} - u_k * \rho_{1/m}\|_{L^q} \\
&\le 2Cm^{-\theta} + \|u_n * \rho_{1/m} - u_k * \rho_{1/m}\|_{L^q},
\end{aligned}$$

because the $W^{1,p}$-norms of the u_n are bounded by 1. In order to ensure that our result is less than some positive δ, choose m sufficiently large such that $2Cm^{-\theta} < \frac{1}{2}\delta$, and the uniform convergence of $\{u_n*\rho_{1/m}\}_n$ allows us to see that the other term is less than $\frac{1}{2}\delta$ for large n and k. This concludes the proof that $W_0^{1,p}(U) \subseteq L^q(U)$ is a compact linear map. □

The main point of the proof of Rellich's theorem (and power of the Sobolev inequalities) is that, if you have control over the $W^{1,p}$-norm, then you can have control of ε-mollifications of functions (up to a certain power of ε).

Next, let's cover the *Poincaré inequality*.

20.5 Theorem (Poincaré)**.** *There exists a positive number C_U such that, for all u in $C_c^\infty(U)$,*

$$\|u\|_{L^2(U)} \le C_U\|Du\|_{L^2(U)}.$$

Proof. Let's just prove it in the $d \ge 3$ case. Here, we have

$$2d > 2d - 4 \implies 2^* = \frac{2d}{d-2} > 2,$$

where the $d \ge 3$ hypothesis allows us to preserve order when dividing by $d - 2$. Therefore the Hölder and Sobolev inequalities yield

$$\|u\|_{L^2(U)} \le |U|^{1/2-1/2^*}\|u\|_{L^{2^*}(U)} \le |U|^{1/2-1/2^*}\|Du\|_{L^2},$$

so we can take $C_U = |U|^{1/2-1/2^*}$. □

21 February 28, 2018

Recall that we discussed the Sobolev spaces $W^{k,p}$ and H^k last time, as well as the Rellich's theorem and the Poincaré inequality. We can now discuss eigenvalues of the *Laplacian* and thus finally return to our motivating question from the beginning of the course.

What is the Laplacian? Let U be an open bounded subset of $\mathbb{R}^d$.

21.1 Definition. The *Laplacian* is the linear operator $\Delta : C_c^\infty(U) \longrightarrow C_c^\infty(U)$ given by

$$\Delta u := \sum_{i=1}^{d} \partial_i^2 u,$$

where ∂_i denotes the partial derivative with respect to the i-th coordinate.

It turns out that we can't interpret the Laplacian as a continuous operator on $C_c^\infty(U)$ in any sense, but we *will* be able to show that the inverse Laplacian is continuous and even compact when considered on $H_0^0(U)$. For this, we first turn to the *variational formulation* of the Laplacian.

For any f in $C_c^\infty(U)$, we want to find u in $C_c^\infty(U)$ that solves

$$-\Delta u = f, \tag{1}$$

which amounts to finding an inverse image for f under the Laplacian, by taking negations. This is a pointwise equality, so by using approximations to the identity at points in U, we see that it's equivalent to solving

$$\sum_{i=1}^{d} \int_U (\partial_i \varphi)(\partial_i u) = \sum_{i=1}^{d} \int_U -\varphi \partial_i^2 u = \int_U -\varphi \Delta u = \int_U \varphi f \tag{2}$$

for all φ in $C_c^\infty(U)$, where we have used integration by parts. We want to use a Hilbert space structure to solve our problem, so we reformulate this as

$$\langle \partial_i \varphi, \partial_i u \rangle_0 = \langle \varphi, f \rangle_0, \tag{3}$$

where we define our brackets as follows.

21.2 Definition. For any f and g in $H_0^0(U)$ (respectively $H_0^1(U)$), we write

$$\langle f, g \rangle_0 := \int_U fg \text{ (respectively } \langle f, g \rangle_1 := \sum_{i=1}^{d} \int_U (\partial_i f)(\partial_i g))$$

Note that $\langle \cdot, \cdot \rangle_0$ is precisely the inner product on $H_0^0(U) \subseteq H^0(U) = L^2(U)$.

21.3 Remark. And as for $\langle \cdot, \cdot \rangle_1$, recall that for any f in $C_c^\infty(U)$, we have

$$\|f\|_{H^1(U)} = \|f\|_{L^2(U)} + \sum_{i=1}^{d} \|\partial_i f\|_{L^2(U)}.$$

Now note that $\langle f, f \rangle_1 = \sum_{i=1}^{d} \|\partial_i f\|_{L^2(U)}^2 = \|Df\|_{L^2(U)}^2$, and by the Poincaré inequality, we have

$$\|f\|_{L^2(U)} \leq C_U \|Df\|_{L^2(U)}$$

for some positive C_U depending only on U. One can use this to show that the norm induced by $\langle \cdot, \cdot \rangle_1$ is equivalent to H^1-norm on $C_c^\infty(U)$, which allows us to equip $H_0^1(U)$ with a Hilbert space structure by taking the completion with respect to the norm induced by $\langle \cdot, \cdot \rangle_1$.

Our new norm $\langle \cdot, \cdot \rangle_1$ allows us to rewrite (3) as

$$\langle \varphi, u \rangle_1 = \langle \varphi, f \rangle_0. \tag{4}$$

We can finally invert the Laplacian. From now on, assume that ∂U has measure zero.[9]

21.4 Theorem. *There exists a bounded linear operator $A : H_0^0(U) \longrightarrow H_0^1(U)$ such that $\langle f, Ag \rangle_1 = \langle f, g \rangle_0$ for all f in $H_0^1(U)$ and g in $H_0^0(U)$.*

Proof. Fix g in $H_0^0(U) = L^2(U)$, where the values on ∂U don't matter because ∂U has measure zero. The map $f \mapsto \langle f, g \rangle_0$ is a bounded linear functional on the Hilbert space $H_0^1(U)$, so Riesz representation indicates that there exists a unique Ag in $H_0^1(U)$ such that $\langle f, g \rangle_0 = \langle f, Ag \rangle_1$ for all f in $H_0^1(U)$. We see that $\|Ag\|_{H_0^1(U)} \leq \|g\|_{H_0^0(U)}$, and uniqueness forces A to be linear. Thus A is also bounded, as desired. □

21.5 Remark. The equivalence of Equation (1) and Equation (4) implies that $u = Af$, so A is the (negation of) the inverse Laplacian, in the sense that $A(-\Delta)u = u$. To prove that A is also a right inverse of $-\Delta$, we need *regularity theory*, which we'll return to if we have time at the end. For now, we will indeed refer to A as $(-\Delta)^{-1}$ and call it the *inverse Laplacian*.

To prove more about A, we want to know more about our Sobolev spaces. And what do we know about these Sobolev spaces? Well, a good amount, actually!

21.6 Proposition.

- *The inverse Laplacian $(-\Delta)^{-1} : H_0^0(U) \xrightarrow{A} H_0^1(U) \hookrightarrow H_0^0(U)$ is a compact bounded operator,*
- *The inverse Laplacian is symmetric.*

Proof. Let's just prove it for $d \geq 3$. Here, Rellich's theorem tells us that the inclusion $H_0^1(U) \subseteq L^2(U) = H_0^0(U)$ is a compact map. Therefore the composition $H_0^0(U) \longrightarrow H_0^1(U) \longrightarrow H_0^0(U)$ is compact via Lemma 9.8, because the first map is bounded by Theorem 21.4. As for symmetry, note that

$$\langle f, Ag \rangle_1 = \langle f, g \rangle_0 = \langle g, f \rangle_0 = \langle g, Af \rangle_1 = \langle Af, g \rangle_1$$

for all f and g in $H_0^1(U)$. So the inverse Laplacian is symmetric on $H_0^1(U)$. As $H_0^1(U)$ is dense in $H_0^0(U)$, we obtain the same result in all of $H_0^0(U)$. □

Now that we know that the inverse Laplacian is a compact symmetric operator, we can apply our wealth of spectral theory to it. First off, what is its kernel? We have

$$f \in \ker(-\Delta)^{-1} \iff \langle g, (-\Delta)^{-1} f \rangle_0 = 0 \text{ for all } g \in H_0^0(U) \iff \langle g, (-\Delta)^{-1} f \rangle_1 = 0 \text{ for all } g \in H_0^1(U)$$

by using the density of $H_0^1(U)$ in $H_0^0(U)$ as well as the Poincaré inequality to see that the $\langle \cdot, \cdot \rangle_1$ and $\langle \cdot, \cdot \rangle_0$ norms are compatible. Equation (4) then indicates that this is equivalent to saying

$$\langle g, f \rangle_0 = 0 \text{ for all } g \in H_0^1(U) \iff \langle g, f \rangle_0 = 0 \text{ for all } g \in H_0^0(U) \iff f = 0$$

by another density argument. In other words, $\ker(-\Delta)^{-1}$ is trivial, and it's essentially due to the Poincaré inequality.

From here, applying the spectral theorem yields an orthonormal basis $\{\varphi_k\}_k$ of $H_0^1(U)$ and non-increasing positive sequence $\{\lambda_k\}_k$ in $\mathbb{R}$ such that $\lambda_k \to 0$ and $(-\Delta)^{-1}\varphi_k = \lambda_k \varphi_k$. We haven't yet proven that A is a right inverse to $-\Delta$, but if we had, we could then deduce that $-\Delta\varphi_k = \lambda_k^{-1}\varphi_k$, so $\{\varphi_k\}_k$ would form the

[9]This was not originally assumed, but it seems necessary to make these arguments work, and certainly this is not true for all U.

spectrum for our drum on U. Of course, we want more properties about our φ_k, like smoothness (which is the topic of *regularity theory*) and the distribution of λ_k, but this spectrum result is great already.

Well, let's begin asking these additional questions. For instance, how are the λ_k distributed? The answer is something called the *Weyl law*. For any positive T, write

$$N(T) \coloneqq \#\{k : \lambda_k^{-1} \leq T\}.$$

21.7 Theorem (Weyl)**.** *We have*

$$\lim_{T\to\infty} \frac{N(T)}{T^{d/2}} = (2\pi)^{-d} \cdot |B_1^{\mathbb{R}^d}| \cdot |U|.$$

The Weyl law is the consequence of the following two lemmas.

21.8 Lemma. *Let R be a positive number. Then the Weyl law holds true for $U = (0, R)^d$.*

Sketch of the proof. In this case, we have an explicit eigenbasis for Δ in terms of sines and cosines, was given in the first lecture and can be obtained from Fourier analysis. Therefore we can explicitly demonstrate the Weyl law in this case. □

21.9 Exercise. Explicitly do the computation necessary to prove Lemma 21.8.

Our next lemma is a variational reformulation of $N(T)$.

21.10 Lemma. *Returning to the setting of general U, we have*

$$N(T) = \sup\{\dim V : V \subseteq H_0^0(U) \text{ a subspace such that } \|(-\Delta)^{-1}u\|_0 \geq T^{-1}\|u\|_0 \text{ for all } u \in V\}.$$

Sketch of the proof. This follows from the *Courant–Friedrichs formula*—for more information on this formula, see any of your course texts. □

Proof of the Weyl law. If U contains two disjoint open subsets U_1 and U_2 such that $U \smallsetminus (U_1 \cup U_2)$ has measure zero, then Lemma 21.10 indicates that $N_{U_1}(T) + N_{U_2}(T) \leq N_U(T)$. Furthermore, we also have $H_0^1(U_1) \oplus H_0^1(U_2) \subseteq H_0^1(U)$ via extension by zero. Therefore, by sandwiching U between progressively finer cubes and using Lemma 21.8, the general case follows. □

22 March 2, 2018

Today, we move on to the discussion of the regularity of the spectrum $\{\varphi_k\}_k$ of $(-\Delta)^{-1}$. This spectrum lies in $H_0^1(U)$, but instead of having limits of functions, we want to obtain honest-to-goodness functions that we can evaluate at points. This is the subject of *regularity theory*, which is an incredibly deep topic.

22.1 Theorem. *Let U be an open bounded subset of $\mathbb{R}^d$ whose boundary is smooth. If u in $H_0^1(U)$ is an eigenvector for the negative Laplacian $-\Delta u = \lambda u$, then u is actually in $C^\infty(U)$.*

While we could prove Theorem 22.1 in its entirety if we used all our remaining class periods, we'll only prove parts of it because I want to cover other topics.

22.2 Theorem. *Let u lie in $H^1(B_2)$, let f lie in $H^k(B_2)$, and suppose that $-\Delta u = f$. Then the restriction of u to B_1 actually lies in $H^{k+2}(B_1)$.*

In other words, we can boost the regularity of values of the inverse Laplacian by going into a smaller sub-ball. There's actually a more sophisticated approach for general U in place of B_2, but that requires some more technical definitions, so I won't do it.

The idea of the proof is that, for all u in $C_c^\infty(\mathbb{R}^d)$, we have

$$\|D^2u\|_{L^2}^2 = \sum_{i,j=1}^d \int_{\mathbb{R}^d} (D_iD_ju)\cdot(D_iD_ju).$$

We can use integration by parts twice in order to move the D_i to one factor and the D_j to the other, and the two signs from these integrations by parts cancel. This then yields

$$\sum_{i,j=1}^d \int_{\mathbb{R}^d} (D_iD_iu)\cdot(D_jD_ju) = \int_{\mathbb{R}^d}(\Delta u)^2 = \|\Delta u\|_{L^2}^2.$$

From here, we can apply the Sobolev inequalities to conclude.

22.3 Remark. Everything here works for more general operators

$$Lu \coloneqq \sum_{i,j=1}^d D_i\cdot a^{ij}\cdot D_ju,$$

where $(a^{ij})_{ij}$ is a $(d\times d)$-symmetric matrix with entries in smooth functions on $\mathbb{R}^d$ such that there exists a $\Lambda \geq 1$ satisfying

$$\Lambda^{-1}\,\mathrm{id} \leq (a^{ij})_{ij} \leq \Lambda\,\mathrm{id},$$

where the inequalities mean that the difference is positive definite. In this context, $Lu = f$ becomes equivalent to saying

$$\sum_{i,j=1}^d \langle D_iu, a^{ij}D_j\varphi\rangle = \langle f,\varphi\rangle$$

for all φ in $C_c^\infty(U)$. And by gluing these local situations together, one can show that this discussion actually all works for operators of this form on arbitrary Riemannian manifolds. This is *elliptic regularity*.

Of course, we can't literally port these ideas over, since u is not differentiable, so we don't have objects like Du or D_iu. We rectify this problem by first introducing some finite difference derivatives.

22.4 Definition. Let g lie in $L^2(B_2)$, and let h be a number satisfying $0<|h|<1$. We define the function D_i^hg in $L^2(B_{2-|h|})$ via

$$(D_i^hg)(x) \coloneqq \frac{g(x+he_i)-g(x)}{h}.$$

22.5 Remark. Observe that summation by parts (which is the discretized version of integration by parts!) implies that

$$\langle D_i^hg,k\rangle_{L^2(B_1)} = -\langle g, D_i^{-h}k\rangle_{L^2(B_1+he_i)}$$

for all k in $C_c^\infty(B_1)$. This precisely analogous to integration by parts, except no actual derivatives are needed because we're just taking difference quotients.

Remark 22.5 allows us to obtain Lemma 22.7 below, but let's first prove an easier lemma.

22.6 Lemma. *Let g lie in $H^1(B_2)$. Then $\|D_i^h g\|_{L^2(B_1)} \leq \|D_i g\|_{L^2(B_2)}$.*

Proof. We've seen how to prove this before—we have

$$(D_i^h g)(x) = \frac{g(x+he_i) - g(x)}{h} = h^{-1} \int_0^1 \mathrm{d}t\, h(D_i g)(x+th).$$

Integrating over B_1 and using Jensen's inequality yields the desired result. □

A much more interesting result than Lemma 22.6 is the following partial converse, for which we shall actually need to use some of the functional analysis we learned.

22.7 Lemma. *If g lies in $L^2(B_2)$ and there exists a positive α such that $\|D_i^h g\|_{L^2(B_1)} \leq \alpha$ for all integers $1 \leq i \leq d$ and numbers $0 < |h| < 1$, then g lies in $H^1(B_1)$, and $\|D_i g\|_{L^2(B_1)} \leq \alpha$.*

Proof. By separable Banach–Alaoglu, we can find a sequence $h_k \to 0$ and some $\widetilde{g}_i$ in $H^1(B_1)$ such that $D_i^{h_k} g \to \widetilde{g}_i$ weakly in $L^2(B_1)$, because the reflexivity of $L^2(B_1)$ allows us to convert weak-$*$ convergence into weak convergence. We now want to show that $\widetilde{g}_i$ is the i-th derivative of g. Recall from Remark 22.5 that

$$\langle g, D_i^{-h_k} \varphi \rangle_{L^2(B_1 + h_k e_i)} = -\langle D_i^{h_k} g, \varphi \rangle_{L^2(B_1)}$$

for all φ in $C_c^\infty(B_1)$. The left-hand side equals $\langle g, \tau_{h_k e_i} D^{-h_k} \varphi \rangle_{L^2(B_1)}$, and taking $k \to \infty$ shows that this converges to the equation

$$\langle g, D_i \varphi \rangle_{L^2(B_1)} = -\langle \widetilde{g}_i, \varphi \rangle_{L^2(B_1)},$$

because the $h_k \to 0$ and the $D_i^{h_k} g \to \widetilde{g}_i$. Next, we mollify. Let ε be positive, and let η_ε be the corresponding standard mollifier (which we used in the proof of Rellich's theorem, under the name ρ_ε). Then

$$-\langle D_i(g * \eta_\varepsilon), \varphi \rangle_{L^2(B_1)} = \langle g * \eta_\varepsilon, D_i \varphi \rangle_{L^2(B_1)} = -\langle \widetilde{g}_i * \eta_\varepsilon, \varphi \rangle_{L^2(B_1)},$$

via integration by parts, and since both $g * \eta_\varepsilon$ and $\widetilde{g}_i * \eta_\varepsilon$ are smooth, we see that

$$\widetilde{g}_i * \eta_\varepsilon = D_i(g * \eta_\varepsilon) = (D_i g) * \eta_\varepsilon.$$

The fact that the g and $\widetilde{g}_i$ are L^2 implies that their mollified versions converge to the originals strongly in $L^2(B_1)$, concluding the proof. □

I now have ten minutes to prove elliptic regularity... let's do this.

Proof of Theorem 22.2. The idea shall be to prove a uniform bound on $\|D_i^h D_j u\|^2_{L^2(B_1)}$, for which we want to choose the right test function. We choose ψ in $C_c^\infty(B_2)$ such that $\psi = 1$ on B_1, and we form $\varphi \coloneqq D_j^{-h} \psi^2 D_j^h u$. Because we multiplied by ψ twice in the middle, this φ actually lies in $H_0^1(B_2)$. Now recall that we have

$$\langle D_i u, D_i \varphi \rangle = \langle f, \varphi \rangle$$

for all φ in $C_c^\infty(B_2)$ and hence $H_0^1(B_2)$ by density. Plugging in our chosen φ to the left-hand side yields

$$\begin{aligned}\langle D_i u, D_i D_j^{-h}\psi^2 D_j^h u\rangle &= \langle D_i u, D_j^{-h} D_i \psi^2 D_j^h u\rangle = -\langle D_j^h D_i u, D_i\psi^2 D_j^h u\rangle \\ &= -\langle D_j^h D_i u, \psi^2 D_j^h D_i u\rangle - \langle D_j^h D_i u, 2\psi(D_i\psi) D_j^h u\rangle \\ &\leq -\int \psi^2 |D_j^h D_i u|^2 - \int (D_j^h D_i u) 2\psi(D_i\psi) D_j^h u.\end{aligned}$$

Cauchy's inequality shows that this is bounded by

$$-\frac{1}{2}\int \psi^2 |D_j^h D_i u|^2 + C_\psi \int |D_j^h u|^2,$$

where C_ψ is a positive constant depending on ψ. We obtain a similar bound on the right-hand side in terms of $D_\psi \int |D_j^h u|^2$ and f^2, as desired, where D_ψ is another positive constant depending on ψ. □

23 February 5, 2018

The note-taker missed class today and thanks Hao Billy Lee for letting him consult his notes.

We'll conclude the course with a discussion of ordinary differential equations with values in a Banach space! We begin by stating the following theorem, whose statement might be confusing, before explaining the objects involved.

23.1 Theorem. *Let X be a Banach space, let $f : X \longrightarrow X$ be a Lipschitz function, and let x be a point in X. Then there exists a unique continuous map $y : [0,\infty) \longrightarrow X$ such that*

$$y(t) = x + \int_0^t \mathrm{d}s\, f(y(s))$$

for all t in $[0,\infty)$.

When does Theorem 23.1 makes sense?

(1) For any Banach space X and continuous function $u : [0,T] \longrightarrow X$, the Riemann integral $\int_0^T \mathrm{d}s\, u(s)$ makes sense, where one emulates the usual construction of Riemann integration to define it.

(2) In a similar manner, one can construct derivatives of continuous functions $u : [0,T] \longrightarrow X$ (if they exist), and Theorem 23.1 is equivalent to saying that

$$y'(t) = f(y(t)) \text{ and } y(0) = x$$

for all t in $[0,\infty)$.

Proof of Theorem 23.1. Let K be the Lipschitz constant of f, and let Y be the space $C([0,1],X)$ equipped with the norm $\|y\|_Y \coloneqq \sup_{0\leq t\leq 1} e^{-\alpha t}\|y(t)\|_X$, where $\alpha \coloneqq 2K$. We define a map $I : Y \longrightarrow Y$ via

$$(Iy)(t) \coloneqq y(0) + \int_0^t \mathrm{d}s\, f(y(s)).$$

For any y and $\widetilde{y}$ in Y satisfying $y(0) = \widetilde{y}(0)$, we have

$$\begin{aligned}\|(I\widetilde{y} - Iy)(t)\|_X &= \left\|\int_0^t \mathrm{d}s\, f(\widetilde{y}(s)) - f(y(s))\right\|_X \leq \int_0^t \mathrm{d}s\, \|f(\widetilde{y}(s)) - f(y(s))\|_X \leq K\int_0^t \mathrm{d}s\, |\widetilde{y}(s) - y(s)| \\ &\leq K\int_0^t \mathrm{d}s\, e^{\alpha s}\|\widetilde{y} - y\|_Y = \frac{K}{\alpha}(e^{\alpha t} - 1)\|\widetilde{y} - y\|_Y \leq \frac{1}{2}e^{\alpha t}\|\widetilde{y} - y\|_Y.\end{aligned}$$

Therefore $\|I\widetilde{y} - Iy\|_Y \leq \frac{1}{2}\|\widetilde{y} - y\|_Y$, so I is a contraction mapping on the subspace of Y consisting of y satisfying $y(0) = x$. The contraction mapping theorem implies that I has a unique fixed point y on this subspace, and y satisfies the desired properties. By dilating our interval and passing to $[0, T]$ for $T \to \infty$, we obtain the desired result. $\square$

We would like to apply Theorem 23.1 to the heat equation

$$\partial_t v = \Delta v$$

by taking $f = \Delta$, but there are a few problems with this: Δ isn't defined on our entire function space X of interest, and it's also not bounded. We shall introduce the formalism of *Yosida approximation* to deal with both of these problems.

Let H be a Hilbert space, let $D(A)$ be a (linear) subspace of H, let $A : D(A) \longrightarrow H$ be a linear map, and let x lie in $D(A)$. Our goal is to find a $y : [0, \infty) \longrightarrow D(A)$ that solves

$$y'(t) + Ay(t) = 0 \text{ and } y(0) = x$$

for all t in $[0, \infty)$

23.2 Example. In our heat equation setup, take $H = L^2(\mathbb{R}^d)$, $D(A) = H^2(\mathbb{R}^d) \subset L^2(\mathbb{R}^d)$, and $A = -\Delta$.

The idea behind Yosida approximation is to build a bounded approximation to A, which we can then apply Theorem 23.1 to.

23.3 Definition. We say that the map A is *maximal monotone* if $\langle Ax, x\rangle_H \geq 0$ for all x in $D(A)$ and $\operatorname{im}(\operatorname{id} + A) = H$.

23.4 Example. It is known that our heat equation setup has A being maximal monotone.

The following lemma is crucial for our discussion of maximal monotonicity.

23.5 Lemma. *Suppose A is maximal monotone. Then*

(1) $D(A)$ is dense in H,

(2) A is a closed map,

(3) for all positive λ, the map $\operatorname{id} + \lambda A : D(A) \longrightarrow H$ is a bijection, and $\|\operatorname{id} + \lambda A\| \leq 1$.

Proof.

(1) Let x lie in H, and choose y in $D(A)$ such that $x = y + Ay$. If $x \perp D(A)$, then we have

$$0 = \langle x, y\rangle = \langle y + Ay, y\rangle = \|y\|^2 + \langle Ay, y\rangle \geq \|y\|^2.$$

Therefore y and hence x is equal to zero.

(2) Let λ be positive. I claim that if $\operatorname{id} + \lambda A$ is surjective, then it is in fact a bijection such that $\|(\operatorname{id} + \lambda A)^{-1}\| \leq 1$. To see this, suppose that $x \coloneqq y + \lambda Ay = z + \lambda Az$ for some y and z in $D(A)$. Then

$$0 = \langle (y - z) + \lambda A(y - z), y - z\rangle = \|y - z\|^2 + \lambda\langle A(y - z), y - z\rangle \geq \|y - z\|^2,$$

so $y = z$. In addition, we have

$$\|y\| \cdot \|x\| \geq \langle y, x\rangle = \langle y, y + Ay\rangle = \|y\|^2 + \lambda\langle Ay, y\rangle \geq \|y\|^2$$

by the Cauchy–Schwartz inequality. When $y \neq 0$, this indicates that $\|x\| \geq \|y\|$, so we see that $\|(\mathrm{id} + \lambda A)^{-1}\| \leq 1$.

Our hypotheses show that we can apply this for $\lambda = 1$. Given x_n in $D(A)$ such that Ax_n converges to y, we can choose the x_n such that x_n converges to some x. It suffices to show that x lies in $D(A)$, and that $Ax = y$. Because $x_n + Ax_n \to x + y$, applying the continuous operator $(\mathrm{id} + A)^{-1}$ yields

$$x_n = (\mathrm{id} + A)^{-1}(x_n + Ax_n) \to (\mathrm{id} + A)^{-1}(x + y) = x \in D(A).$$

Applying $\mathrm{id} + A$ to both sides yields $x + y = x + Ax$ and hence $y = Ax$, as desired.

(3) By our above claim, it remains to prove that $\mathrm{id} + \lambda A$ is surjective. We know this for $\lambda = 1$, and we want to prove it for all $\lambda > 0$ by "inducting" in both directions. Namely, suppose that it's true for λ, and let μ lie in $(\lambda/2, 2\lambda)$. We want to show that it's true for μ.

We do this via the contraction mapping theorem. We need to solve

$$(\mathrm{id} + \mu A)x = y$$

for all y in H. Multiplying by λ/μ shows that this is the same as

$$\frac{\lambda}{\mu}x + \lambda Ax = \frac{\lambda}{\mu}y \iff x + \lambda Ax = \left(1 - \frac{\lambda}{\mu}\right)x + \frac{\lambda}{\mu}y \iff x = (\mathrm{id} + \lambda A)^{-1}\left[\left(1 - \frac{\lambda}{\mu}\right)x + \frac{\lambda}{\mu}y\right].$$

By hypothesis, we have $\|(\mathrm{id} + \lambda A)^{-1}\| \leq 1$, and since $|1 - \lambda/\mu| < 1$, this equation yields a contraction mapping. This finishes our proof! ☐

With Lemma 23.5 in hand, we can make the following definitions.

23.6 Definition. Suppose A is maximal monotone, and let λ be positive. We write

$$J_\lambda := (\mathrm{id} + \lambda A)^{-1} \text{ and } A_\lambda := \frac{1}{\lambda}(\mathrm{id} - J_\lambda).$$

23.7 Remark. The operators J_λ approximate id as $\lambda \to 0$, and they also map H to the domain $D(A)$. Now the A_λ have bounded norm by Lemma 23.5.(3), and because we morally have

$$J_\alpha = (\mathrm{id} + \lambda A)^{-1} \approx \mathrm{id} - \lambda A + \lambda^2 A - \cdots \implies A_\lambda = \frac{1}{\lambda}(\mathrm{id} - J_\lambda) \approx A - \lambda A^2 + \lambda^2 A^3 - \cdots,$$

we see that A_λ is supposed to approximate A as $\lambda \to 0$. This is the idea behind Yosida approximation—we hope to solve our ODE for the A_λ and then take the limit to solve it for A.

We explicate Remark 23.7 via the following preparatory lemma.

23.8 Lemma. *Suppose that A is maximal monotone, let λ be positive, and let x be in $D(A)$. Then*

(1) $A_\lambda x = AJ_\lambda x = J_\lambda Ax$,

(2) $\|A_\lambda x\| \leq \|Ax\|$,

(3) $\lim_{\lambda \to 0} J_\lambda x = x$,

(4) $\lim_{\lambda \to 0} A_\lambda x = Ax$,

(5) $\langle A_\lambda x, x\rangle \geq 0$,

(6) $\|A_\lambda\| \leq \lambda^{-1}\|x\|$.

23.9 Remark. As Lemma 23.5.(1) indicates that $D(A)$ is dense in H, continuity indicates that most of these are true for all x in H.

Proof. These are all really consequences of (1), which is straightforward from the definitions:

(1) Suppose that $J_\lambda x = y$, so that $x = y + \lambda Ay$. Then $Ax = Ay + \lambda A \cdot Ay$, so $J_\alpha Ax = Ay = AJ_\lambda x$. We also have $A_\lambda x = (x-y)/\lambda = Ay$, as desired.

(2) Part (1) indicates that $\|A_\lambda x\| = \|J_\lambda Ax\| \leq \|J_\lambda\|\|Ax\| \leq \|Ax\|$, because Lemma 23.5.(3) indicates that $\|J_\alpha\| \leq 1$.

(3) Part (2) shows that, as we take $\lambda \to 0$,

$$\|x - J_\lambda x\| = \|(\mathrm{id} - J_\alpha)x\| = \|\lambda A_\lambda x\| \leq \lambda\|Ax\| \to 0.$$

(4) Part (1) and part (3) yield $\lim_{\lambda\to 0}\|A_\lambda x\| = \lim_{\lambda\to 0}\|J_\lambda Ax\| = \|Ax\|$.

(5) Part (1) indicates that

$$\langle A_\lambda x, x\rangle = \langle A_\lambda x, x - J_\lambda x\rangle + \langle A_\lambda x, J_\lambda x\rangle = \lambda\|A_\lambda x\|^2 + \langle AJ_\lambda x, J_\lambda x\rangle \geq \lambda\|A_\lambda x\|^2 \geq 0.$$

(6) Our calculation from part (5) shows that $\|A_\lambda x\|^2 \leq \lambda^{-1}\langle A_\lambda x, x\rangle$, which in turn is less than $\lambda^{-1}\|A_\lambda x\|\|x\|$ by Cauchy–Schwartz. Dividing by $\|A_\lambda\|$ yields the desired result. ☐

We can now move to the main result of *Hille–Yosida*.

23.10 Theorem (Hille–Yosida)**.** . *Suppose A is maximal monotone, and let x lie in $D(A)$. Then there exists a unique y in $C([0,\infty), D(A))$ such that*

$$y(t) = x - \int_0^t \mathrm{d}s\, A(y(t))$$

for all t in $[0,\infty)$. Moreover, y satisfies $\|y(t)\| \leq \|x\|$ and $\|A(y(t))\| \leq \|Ax\|$.

Proof. Uniqueness follows from the integral version of the following computation idea:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{2}\|\widetilde{y} - y\|^2\right) = \left\langle \frac{\mathrm{d}}{\mathrm{d}t}(\widetilde{y} - y), \widetilde{y} - y\right\rangle = \langle -A(\widetilde{y} - y), \widetilde{y} - y\rangle \leq 0.$$

On the other hand, existence will require lots of work. We'll chunk it up into many steps:

(1) Let λ be positive. We can readily solve

$$y_\lambda(t) = x - \int_0^t \mathrm{d}s\, A_\lambda y_\lambda(s)$$

via Theorem 23.1, because Lemma 23.5.(6) indicates that A_λ is bounded. In fact, the linearity of A_λ implies that this y_λ lies in $C^\infty([0,\infty), H)$, and we have

$$\frac{\mathrm{d}^{n+1}y_\lambda}{\mathrm{d}t^{n+1}} = -A_\lambda\frac{\mathrm{d}^n y_\lambda}{\mathrm{d}t^n}$$

for all non-negative integers n.

(2) We shall prove that $\|y_\lambda(t)\| \leq \|x\|$ and $\|A_\lambda y_\lambda(t)\| \leq \|A_\lambda x\|$. Our initial conditions for y_λ imply that

$$\|y_\lambda(0)\| = \|x\| \text{ and } \|y'_\lambda(0)\| = \|A_\lambda x\| \leq \|Ax\|.$$

From here, the analog of our uniqueness computation and Lemma 23.5.(5) yield

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{2}\|y_\lambda\|^2\right) = \langle -A_\lambda y_\lambda, y_\lambda\rangle \leq 0 \text{ and } \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{2}\|A_\lambda y_\lambda\|^2\right) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\left\|\frac{\mathrm{d}y_\lambda}{\mathrm{d}t}\right\|^2 = \langle -A_\lambda y'_\lambda, y'_\lambda\rangle \leq 0,$$

so our initial conditions imply the desired result.

(3) We want to piece these y_λ together and show that $\lim_{\lambda\to 0} y_\lambda(t)$ converges for all t in $[0,\infty)$, which is hard. Our goal will be to prove that y_λ is Cauchy for $\lambda \to 0$. As before, we begin by taking

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\|y_\lambda - y_\mu\|^2\right) = -\langle A_\lambda y_\lambda - A_\mu y_\mu, y_\lambda - y_\mu\rangle.$$

Next, observe that $\mathrm{id} = \lambda A_\lambda + J_\lambda$, so Lemma 23.5.(1) indicates that

$$-\langle A_\lambda y_\lambda - A_\mu y_\mu, y_\lambda - y_\mu\rangle = -\langle A_\lambda y_\lambda - A_\mu y_\mu, \lambda A_\lambda y_\lambda - \mu A_\mu y_\mu\rangle - \langle A(J_\lambda y_\lambda - J_\mu y_\mu), J_\lambda y_\lambda - J_\mu y_\mu\rangle.$$

Combining Lemma 23.5.(4) and Lemma 23.5.(5) shows that

$$\langle A(J_\lambda y_\lambda - J_\mu y_\mu), J_\lambda y_\lambda - J_\mu y_\mu\rangle \geq 0,$$

so our derivative is bounded above by

$$\begin{aligned}
-\langle A_\lambda y_\lambda - A_\mu y_\mu, \lambda A_\lambda y_\lambda - \mu A_\mu y_\mu\rangle &= -(\lambda\|A_\lambda y_\lambda\|^2 - \lambda\langle A_\mu y_\mu, A_\lambda y_\lambda\rangle - \mu\langle A_\lambda y_\lambda, A_\mu y_\mu\rangle + \mu\|A_\mu y_\mu\|^2) \\
&\leq -\lambda\|A_\lambda y_\lambda\|^2 - \mu\|A_\mu y_\mu\|^2 + (\lambda+\mu)\|A_\mu y_\mu\|\|A_\lambda y_\lambda\| \\
&= \mu\|A_\lambda y_\lambda\|^2 + \lambda\|A_\mu y_\mu\|^2 \\
&\quad - (\lambda+\mu)\left[(\|A_\lambda y_\lambda\| - \|A_\mu y_\mu\|)^2 + \|A_\lambda y_\lambda\|\|A_\mu y_\mu\|\right]
\end{aligned}$$

by Cauchy–Schwartz. Step (2) and Lemma 23.5.(2) indicate that this is bounded above by

$$\max\{\lambda,\mu\}\|Ax\|^2.$$

We integrate this estimate on our derivative to obtain the bound

$$\|y_\lambda(t) - y_\mu(t)\|^2 \leq \max\{\lambda,\mu\}\|Ax\|^2 t,$$

because $\|y_\lambda(0) - y_\mu(0)\|^2 = 0$. Finally, taking square roots shows that the y_λ converge uniformly on compact subsets $[0,T]$ of $[0,\infty)$, so they converge to some y in $C([0,\infty),H)$.

(4) Suppose first that x lies in $D(A^2) := A^{-1}(D(A))$. Then $z(t) := \lim_{\lambda\to 0} A_\lambda y_\lambda(t)$ exists for all non-negative t, by arguing as in step (3). This time, we use $\|A^2 x\|$ instead of $\|Ax\|$ for our bounds. Recall from step (1) that $A_\lambda y_\lambda = -\frac{\mathrm{d}y_\lambda}{\mathrm{d}t}$.

(5) We shall prove that, for x in $D(A^2)$, the theorem is true. Note that our existing results say that $y_\lambda \to y$ on every compact interval, and hence its negative derivative is given by $-\frac{\mathrm{d}y_\lambda}{\mathrm{d}t} = A_\lambda y_\lambda \to z$. Therefore the fundamental theorem of calculus indicates that

$$y(t) = x - \int_0^t \mathrm{d}s\, z(s),$$

so it suffices to prove that $z = Ay$. As Lemma 23.5.(1) gives $A_\lambda y_\lambda = AJ_\lambda y_\lambda$, and A is closed, it suffices to show that $J_\lambda y_\lambda(s) \to y(s)$ and then apply A. But we have

$$\|J_\lambda y_\lambda(s) - y(s)\| \leq \|J_\lambda(y_\lambda(s) - y(s))\| + \|J_\lambda y(s) - y(s)\| \to 0$$

as $\lambda \to 0$, because $y_\lambda(s)$ converges to $y(s)$, J_λ is continuous, and we have Lemma 23.5.(3).

(6) To finally pass to x in $D(A)$, we prove that $D(A^2)$ is dense in $D(A)$ under the *graph norm*

$$\|a\|_{D(A)} := \|a\| + \|Aa\|.$$

To see this, set $x_\lambda := J_\lambda x$, which is readily seen to lie in $D(A^2)$. Lemma 23.5 gives us the desired convergence $x_\lambda \to x$ in the graph norm. □

24 February 7, 2018

The note-taker missed class today and thanks Hao Billy Lee for letting him consult his notes.

24.1 Remark. In the setting of the Hille–Yosida theorem,

(1) Our result is equivalent to solving $y'(t) + Ay(t) = 0$ and $y(0) = x$ for all non-negative t.

(2) Our strategy for proving this was via taking the Yosida approximations

$$A_\lambda = \frac{1}{\lambda}\left(\mathrm{id} - (\mathrm{id} + \lambda A)^{-1}\right),$$

which had norms bounded by $1/\lambda$ and hence were amenable to standard ODE tactics. We then checked that the limit $\lim_{\lambda\to 0} y_\lambda$ of their solutions yielded a solution for the ODE with A.

(3) One can give an alternative proof of this fact by taking

$$y(t) = \lim_{n\to\infty}\left(\mathrm{id} + \frac{t}{n}A\right)^{-n} x$$

for the desired solution. This is called the *backwards Euler* method.

Let's now apply the Hille–Yosida theorem to a specific differential equation! Let $(a^{ij})_{ij}$ be a $(d \times d)$-matrix with entries in C^∞-functions on $\mathbb{R}^d$ such that there exists a $\Lambda \geq 1$ satisfying

$$\Lambda^{-1}\|x\|^2 \leq \sum_{i,j=1}^{d} x_i a^{ij} x_j \leq \Lambda\|x\|^2$$

for all $x = (x_1, \ldots, x_d)$ in $\mathbb{R}^d$. In other words, $(a^{ij})_{ij}$ is precisely the kind of matrix considered in Remark 22.3.

24.2 Example. As in our case of interest for elliptic regularity, we could take $(a^{ij})_{ij} = \mathrm{id}$.

Consider the *wave equation*

$$\begin{cases} D_t^2 u - Lu = 0 & \text{on } \mathbb{R}^d \times [0,\infty), \\ u = g & \text{on } \mathbb{R}^d \times \{0\}, \\ D_t u = h & \text{on } \mathbb{R}^d \times \{0\}, \end{cases}$$

where $Lu \coloneqq \sum_{i,j=1}^{d} D_i \cdot a^{ij} \cdot D_j u$, Our goal is to solve for u, given that g and h are fixed.

We begin by rewriting the wave equation as a first-order system in t. Setting $v \coloneqq D_i u$ yields

$$D_t u - v = 0 \text{ and } D_t v - Lu = 0$$

on $\mathbb{R}^d \times [0, \infty)$. By considering the pair (u, v), we can rewrite this as

$$D_t(u, v) + A(u, v) = 0,$$

where $A(u, v) \coloneqq -(v, Lu)$. If we focus on a bounded open subset U whose boundary has measure zero (instead of all of $\mathbb{R}^d$ for our space), then our solutions (u, v) of interest, as functions of space, lie in $H \coloneqq H_0^1(U) \times L^2(U)$. And because the operator A requires derivatives, we see that we must set

$$D(A) \coloneqq (H^2(U) \cap H_0^1(U)) \times H^1(U).$$

24.3 Proposition. *The operator A is maximal monotone.*

Therefore we can indeed solve the wave equation.

Finally, let us turn to the *semigroup* variant of Hille–Yosida. Suppose that A is maximal monotone, let x lie in $D(A)$, and let t be non-negative. If we set $S_t(x) \coloneqq y(t)$, where y is given by our existing Hille–Yosida theorem for x, we see that

(1) for all t, the operator $S_t : D(A) \longrightarrow D(A)$ is linear, and we have $\|S_t\| \leq 1$,

(2) $S_0 = \mathrm{id}$ and $S_{s+t} = S_s S_t$ for all non-negative s and t,

(3) for all x in $D(A)$, the map $t \mapsto S_t(x)$ is continuous.

Property (1) tells us that S_t is a bounded operator, and Lemma 23.5.(1) indicates that $D(A)$ is dense in H, so we may uniquely form the continuous linear extension $S_t : H \longrightarrow H$. This allows us to boost property (3) to all x in H.

24.4 Definition. Let $\{S_t\}_{t\geq 0}$ be a collection of linear operators on H. If properties (1)–(3) hold for $\{S_t\}_{t\geq 0}$, we say it is a *contraction semigroup*.

The above discussion shows that our first variant of Hille–Yosida provides a contraction semigroup for any maximal monotone A. Our second variant of Hille–Yosida gives a converse to this process.

24.5 Theorem (Hille–Yosida, revisited)**.** *Let $\{S_t\}_{t\geq 0}$ be a contraction semigroup. Then it comes from the maximal monotone operator defined via*

$$Ax \coloneqq \lim_{h\to 0} \frac{x - S_h x}{h},$$

where $D(A)$ is the subspace of H for which the above limit exists.

Proof. We first check that $\langle Ax, x\rangle \geq 0$ for all x in $D(A)$. To see this, we observe that

$$\left\langle \frac{x - S_h}{h}, x \right\rangle = h^{-1}\left(\|x\|^2 - \langle S_h x, x\rangle\right) \geq h^{-1}\left(\|x\|^2 - \|S_h x\| \cdot \|x\|\right) \geq h^{-1}\left(\|x\|^2 - \|x\|^2\right) = 0$$

by Cauchy–Schwartz and property (1). Next, set $R_\lambda x \coloneqq \int_0^\infty \mathrm{d}t\, e^{-\lambda t} S_t x$ for any positive λ. Then property (1) yields $\|R_\lambda\| \leq \frac{1}{\lambda}$, and I claim that

$$(\lambda \,\mathrm{id} + A) R_\lambda = \mathrm{id} \text{ and } R_\lambda(\lambda \,\mathrm{id} + A) = \mathrm{id}\,.$$

Then setting $\lambda = 1$ would show that $\operatorname{im}(\operatorname{id} + A) = H$. To verify the claim, we shall compute AR_λ. We have

$$\begin{aligned}
\frac{R_\lambda x - S_h R_\lambda x}{h} &= \frac{1}{h}\int_0^\infty \mathrm{d}t\,(e^{-\lambda t}S_t x - S_h e^{-\lambda t} S_t x) = \frac{1}{h}\int_0^\infty \mathrm{d}t\, e^{-\lambda t} S_t x - \frac{1}{h}\int_0^\infty \mathrm{d}t\, e^{-\lambda t} S_{t+h} x \\
&= \frac{1}{h}\int_0^\infty \mathrm{d}t\, e^{-\lambda t} S_t x - \frac{e^{\lambda h}}{h}\int_h^\infty \mathrm{d}t\, e^{-\lambda t} S_t x \\
&= \frac{1}{h}\int_0^h \mathrm{d}t\, e^{-\lambda t} S_t x - \frac{1 - e^{\lambda h}}{h}\int_h^\infty \mathrm{d}t\, e^{-\lambda t} S_t x
\end{aligned}$$

by properties (2) and (3). Taking $h \to 0$ yields $AR_\lambda x = x - \lambda R_\lambda x$, which yields $(\lambda \operatorname{id} + A)R_\lambda = \operatorname{id}$. We can prove $R_\lambda(\lambda \operatorname{id} + A) = \operatorname{id}$ in precisely the same manner. Finally, applying the fundamental theorem of calculus tells us that A indeed recovers the contraction semigroup $\{S_t\}_{t \geq 0}$. $\square$