MATH 205B: Real Analysis II

Lecturer: Professor Rafe Mazzeo

Notes by: Andrew Lin

Winter 2023

Introduction

The syllabus, as well as any additional class postings, can be found on the Canvas page for the course. Office hours will be set in the next day or so, but we can email if we need to meet with the course staff individually to chat about anything.

Problem sets will be assigned weekly and are important for our preparation for the qualifying exam – we should work with each other but write up solutions on our own. Communication is important, and we should make sure to write in complete sentences and be clear with what we mean. (They should be handed in on Thursdays on class.) We'll have a take-home final exam (around week 8) which is basically a glorified individual problem set, but no in-class exams.

Math 205A is fairly disjoint from 205B – the two classes cover two different aspects of real analysis, and what we'll need most is useful properties of L^p spaces and some basic facts about the Fourier transform. We'll be most closely following Bühler and Salamon's book, which is fairly new and with a freely accessible PDF online. But there's supplementary material as well (with Fourier analysis and distribution theory) for which other references will be useful too.

1 January 9, 2023

Functional analysis is an important subject that has applications to lots of areas of mathematics, and so this course will give examples (mostly from other parts of analysis, like PDE) that are useful in other fields. We can either work from the most general objects (**topological vector spaces**) and work towards the special ones (Hilbert spaces) or go in the reverse direction, and we'll kind of do something in between. So this might seem like a glorified version of linear algebra, but that's a naive way of thinking about it because we're looking at particular types of linear maps useful for applications.

Remark 1. *From a real analysis point of view, one motivation for what we'll be doing in this class is the study of fine regularity properties of functions. We may ask about how continuous or differentiable a function is, as well as other ways that functions are "better" than just being measurable. There are many gradations of regularity, and some we'll discuss include the Holder and Sobolev spaces. And when we discuss generalized functions (distributions), we'll see that again measures of regularity come up, and function spaces coming from norms are a useful consideration.*

Remark 2. *In the study of partial differential equations, consider the Laplacian $\Delta : u \mapsto \Delta u = \sum_{j=1}^n \frac{\partial^2 u}{\partial x_j^2}$ is a linear mapping. The idea is to think of $\Delta u = f$ as an infinite-dimensional matrix equation, and then the questions of existence* *and regularity are relevant (for example, if f lives in some function space, we may want to know about the properties of u). Functional analysis then allows us to separate out the steps, for example showing that a solution is first in some class of functions and then that it is in fact satisfying even nicer properties.*

Quantum mechanics provides some motivations for functional analysis as well – it turns out that quantum mechanics was being developed around the same time and was exactly the right tool needed for formalism. And these concepts also come up in representation theory (if we think about infinite-dimensional representations), probability, and many other fields.

Definition 3

A **topological vector space** $(X, \mathcal{T})$ is a vector space with a compatible topology. In other words, X is a vector space over $\mathbb{R}$ or $\mathbb{C}$, and $\mathcal{T}$ is a topology on X which is invariant under vector space operations, meaning that for any open set $U \in \mathcal{T}$, we have $U + x, \lambda U \in \mathcal{T}$, and such that the operations of addition $X \times X \to X$ and multiplication $\mathbb{C} \times X \to X$ are continuous.

Notice that by translation-invariance, we can specify $\mathcal{T}$ by giving just the collection of open sets containing the origin, which we'll call $\mathcal{T}_0$. This is the most general object we will be considering, but we may impose various other conditions:

- Suppose $\mathcal{T}$ is generated by open sets in $\mathcal{T}_0$ which are **convex**. (Recall that convex sets U are those where for any $p, q \in U$, the line segment $\{sp + (1-s)q : 0 \le s \le 1\}$ is contained in U. Then we want the convex sets to form a base for the topology.) Convexity is a very helpful condition, so it is often useful to assume that we have a **locally convex** topological vector space.
- Next, suppose the base $\mathcal{T}_0'$ contains only convex, **balanced** sets. (In other words, for any $x \in U$, $\{\lambda x : |\lambda| \le 1\}$ is all contained in U – this is a "symmetric" condition.)
- For any convex, balanced open set U, we may then want to associate a **seminorm** (a way of measuring distance) by taking, for any $x \in X$,
$$||x||_U = \inf\{\lambda \in \mathbb{R}^+ : x \in \lambda U\}.$$
(In other words, we dilate our set U as little as possible to contain x.) **But** in order to ensure that this infimum is nonempty, we must place an additional assumption, which is that our set is **absorbing**.

Definition 4

A **seminorm** $p : X \to \mathbb{R}_{\ge 0}$ is a function which is positively homogeneous (so $p(\lambda x) = |\lambda| p(x)$ for all $\lambda \in \mathbb{R}$) and satisfies the triangle inequality $p(x + y)) \le p(x) + p(y)$.

The idea is that unlike a norm, we often have large infinite-dimensional spaces on which the seminorms take value zero. As an exercise, we can check that the definition of $||x||_U$ above does satisfy the triangle inequality, so it is a seminorm, and the point is that **our topology is generated by a large set of seminorms** (as many as there are open sets in the base). So we've now arrived at studying the locally convex topological vector spaces (with the additional conditions on the base), and now we can place further additional conditions:

- **Frechet spaces**: the topology is generated by a **countable** family of seminorms p_j for $j \in \mathbb{Z}_{\ge 0}$, and then we get a family of progressively stronger seminorms
$$|| \cdot ||_j = p_0 + \cdots + p_j.$$

So our countable base is $\{x : p_j(x) < \frac{1}{k}\}$ over the set of all j, k. For example, $C^\infty(S^1)$ is a Frechet space with seminorms given by $p_j(f) = \sup_{x \in S^1} |f^{(j)}(x)|$, and we get $||f||_j = \sup_{i \le j, x \in S^1} |f^{(i)}(x)|$ as a comparable norm to the $|| \cdot ||_j$ above (if one norm is bounded above and below by the other, we should really think of them as equivalent).

- **Banach spaces**: the topology is generated by a single seminorm, which is actually a norm, such that the norm only vanishes on the zero vector. The space $C^0(S^1)$ with the sup norm is an example of this. (More precisely, we're actually quotienting out by $\{x \in X : ||x|| = 0\}$, which is a subspace by the triangle inequality, and we can actually do this in the Frechet setting as well. We'll often do this without comment.)
- **Hilbert spaces**: the norm actually comes from an inner product $\langle , \rangle : X \times X \to \mathbb{R}$ or $\mathbb{C}$, with $||f||^2 = \langle f, f \rangle$. For example, $L^2(S^1)$ with the inner product is a Hilbert space.

And in fact, in all three of these cases we actually need the space to be **complete**, so that we can easily talk about Cauchy sequences. But completeness is tricky with respect to a countable family of seminorms: since the intersection of only finitely many open sets is open, we can require $||f_j|| < C_j$ for only $0 \le j \le J$, but there are still infinitely many directions in which things can fail to converge. Instead, the point is that if we have an increasing collection of seminorms $\{|| \cdot ||_j\}_{j=0}^\infty$, we can define a translation-invariant metric

$$d(f, 0) = \sum 2^{-j} \frac{||f||_j}{1 + ||f||_j}.$$

As an exercise, we can check that this generates a topology, and we can check if it is the same as the one from the countable family of seminorms.

Example 5

Consider the function space $C^\infty(\mathbb{R})$ (no requirement that these functions are bounded, just smooth). For any k, ℓ, we will define

$$||f||_{k,\ell} = \sup_{j \le k, |x| \le \ell} ||f^{(j)}(x)||.$$

In other words, we take an expanding sequence of compact sets and take the norms from before, and we can do this because $\mathbb{R}$ can be written as a countable union of such sets. This is a Frechet space.

Example 6

On the other hand, the compactly supported functions $C_c^\infty(\mathbb{R})$ do not form a Frechet space –instead this is what's called a nuclear space. We then say that $f_j \to f$ converges if and only if there is some compact set $K \in \mathbb{R}$ such that all f_j have support in K and $f_j \to f$ in $C^\infty(K)$. In other words, this space is really a union of $C^\infty(K)$ over compact K, with topology given by an **inductive limit**. But this example is important and will haunt us a lot this quarter.

Banach spaces are a nice middle ground, so we'll start with analyzing some of their properties. One fundamental difficulty of infinite dimensional spaces is that **infinite-dimensional spaces are not locally compact**, so we often want conditions under which a sequence $f_j \in X$ of vectors will converge or have a convergent subsequence. (In a finite-dimensional space, we can just assume that $||f_j|| \le C$ for all j, but that isn't enough in infinite dimensions.) We won't describe the most abstract version of this next result, just enough for it to be useful. First, we'll look at just metric spaces:

Proposition 7 (Arzela-Ascoli)
Let X and Y be metric spaces with metrics d_X, d_Y. Then we may define the space $C(X, Y) = \{A : X \to Y \text{ continuous}\}$. Suppose X is compact (otherwise we need to do this with the compact open topology), and define for any $f, g : X \to Y$ the distance

$$d(f, g) = \sup_{x \in X} d_Y(f(x), g(x)).$$

We say that a family $\mathcal{F} \in C(X, Y)$ is **equicontinuous** if for any $\varepsilon > 0$, there exists some $\delta > 0$ such that $d_X(x, x') < \delta$ implies that $d_Y(f(x), f(x')) < \varepsilon$ for all f uniformly. Then if X is compact, Y is complete, and $\mathcal{F} \subset C(X, Y)$, then $\mathcal{F}$ is precompact (closure is compact) if and only if it is equicontinuous and "pointwise precompact" – in other words, $\{f(x)\}_{f \in \mathcal{F}} \subset Y$ is precompact for all x.

(To connect this back to the more familiar version, equicontinuity is the same as the usual notion for functions, and pointwise precompact is a fancy way of saying bounded.)

Proof. For the forward direction, for any $x \in X$ we can define an evaluation map $E_x : C(X, Y) \to Y$ sending f to $f(x)$. By definition this is a continuous mapping. If $\mathcal{F}$ is precompact, then $\overline{\mathcal{F}}$ is compact, so $E_x(\overline{\mathcal{F}})$ is also compact (the image of a compact map under a continuous mapping is compact), which is the same as saying that $\overline{E_x(\mathcal{F})}$ is compact and thus that we have pointwise precompactness. For equicontinuity, if we fix $\varepsilon > 0$, we can choose functions $f_1, \cdots, f_N \in C(X, Y)$ such that $\overline{\mathcal{F}} \subset \bigcup_{j=1}^N B_\varepsilon(f_j)$ (this is saying that we're "totally bounded," and it follows by precompactness), so because each f_i is uniformly continuous and we only have finitely many f_is, there exists some δ such that $d_X(x, x') < \delta$ implies that $d_Y(f_i(x), f_i(x')) < \varepsilon$ for all $x, x' \in X$ and all $1 \le i \le N$. So now taking any f, we know that

$$d_Y(f(x), f(x')) \le d_Y(f(x), f_i(x)) + d_Y(f_i(x), f_i(x')) + d_Y(f_i(x'), f(x')) < 3\varepsilon,$$

which is the desired equicontinuity.

For the reverse direction, suppose we have a sequence inside $\mathcal{F}$ satisfying equicontinuity and pointwise precompactness. Then there exists a collection $\{x_k\} \subset X$ which is dense (again total boundedness – for all N, find a finite collection of $\frac{1}{N}$ balls covering X and put the centers in our collection), so we can choose a subsequence $\tilde{f}_n$ of the f_ns so that $f_n(x_k)$ converges as $n \to \infty$ for each k by repeated diagonalization. Relabel these back to f_n again. Then $\{f_n\}$ is actually Cauchy in $C(X, Y)$, because for all $\varepsilon > 0$ there is some δ such that $d_Y(f(x), f(x')) < \varepsilon$ if $d_X(x, x') < \delta$ for all $f \in \mathcal{F}$. Now $\{B_\delta(x_k)\}$ is an open cover of X, and by the same chain of inequalities

$$d_Y(f_n(x), f_m(x)) < d_Y(f_n(x), f_n(x_k)) + d_Y(f_n(x_k), f_m(x_k)) + d_Y(f_m(x_k), f_m(x))$$

so for any large enough m and n all three terms on the right side are at most ε. So this is indeed a Cauchy sequence. We can then define $f(x)$ by taking the limit of the sequence because Y is complete, and we just need to check that f is continuous and that's the same proof as before. $\square$

The simple picture to keep in mind for this is the following: suppose $X = [0, 1]$, and $\mathcal{F} = \{f \in C^0(I)\}$ satisfy equicontinuity and $||f|| \le M$ for f. That means all functions live in a rectangle, which we can break up into a finer grid of vertical gradation ε and former grid δ. Then if we consider the set of piecewise linear functions whose slope is at most $\frac{\varepsilon}{\delta}$, we're basically saying that our functions are limited to moving by one gridline vertically for every pool.

We also have Arzela-Ascoli for noncompact spaces – our proof here used compactness crucially – and we'll see one such one on our homework.

Example 8

The set of functions $\mathcal{F} = \{f \in C^1(S^1) : |f| \le C_0, |f'| \le C_1\}$ is not equicontinuous in S^1, but if we think of f as sitting inside C^0 then $\mathcal{F}$ is precompact because $|f(x) - f(x')| \le C_1|x - x'|$. This principle of equicontinuity following from derivatives will be seen many times throughout the course.

Next time, we'll start to get into some basic constructions of Banach spaces (dual spaces, spaces of bounded operators, quotient spaces) and then look at three important structural theorems between Banach spaces.

2 January 12, 2023

We'll usually begin class by explaining any questions from the previous lecture.

Remark 9. *A typical use of Arzela-Ascoli is as follows: suppose we have a class of functions, and we're interested in whether there is a convergent subsequence (this is a common use case in PDE). Then if we have a sequence of continuous functions with a bound on the first derivative, then they're equicontinuous, or if we have a bound on the second d The point is that in any particular circumstance, we need uniform control with some higher regularity.*

Remark 10. *A key example to keep in mind for a Frechet space which is not a Banach space is C^∞. And remember that we had a metric on C^∞ (with the 2^{-j} factors from the countably many norms), but it is not a norm because it doesn't satisfy homogeneity.*

Definition 11

A **Banach space** is a vector space X along with a norm $||\cdot||$, such that X is complete with respect to the norm. (Here, the norm is inducing a metric $d(x, y) = ||x - y||$ on X.) A vector space in which the norm is not necessarily complete is just called a **normed space**.

Banach spaces are a nice middle ground between very specialized spaces (Hilbert spaces) and more general spaces (Frechet spaces), in that we can often phrase problems nicely in terms of them. One nice fact is that the unit ball $B = \{x : ||x|| \le 1\}$ is always convex in a Banach space.

Example 12

The sequence spaces $\ell^p = \{x = (x_1, x_2, \cdots) : x_j \in \mathbb{C}, ||x||_p = (\sum |x_j|^p)^{1/p} < \infty\}$ are Banach spaces for all $1 \le p \le \infty$. (For $p = \infty$, we define the norm to be $||x||_\infty = \sup_j |x_j|$.)

For some motivation, if we're only looking at finite-length sequences, it turns out that all norms $(\mathbb{R}^n, ||x||_p)$ are **equivalent** – in other words, for any p, p' and any fixed n, there are absolute constants C_1, C_2 (depending on p, p', n) such that $C_1||x||_p \le ||x||_{p'} \le C_2||x||_p$. In particular, sequences that are Cauchy with respect to one norm are also Cauchy with respect to the other. (In $\mathbb{R}^2$, the unit balls look like a diamond for $p = 1$, a circle for $p = 2$, and a square for $p = \infty$. The point is that the shape of the unit ball doesn't have to look very "uniform" for a Banach space). And a pictorial way to see that these norms are equivalent is that when we look at the ℓ^p and $\ell^{p'}$ unit balls in $\mathbb{R}^n$, we can scale each of them so that it completely contains the other.

On the other hand, this equivalence does not hold for the ℓ^p spaces, and in fact we have $\ell^p \subseteq \ell^{p'}$ for all $p \le p'$. (An easy way to remember this is that $\ell^1 \subseteq \ell^\infty$, since any sequence whose sum of entries is bounded must have bounded entries overall.)

Example 13

For any measure space $(M, d\mu)$, we have the function spaces $L^p(M, d\mu) = \{f : M \to \mathbb{C} : ||f||_p = \left(\int |f|^p d\mu\right)^{1/p} < \infty\}$ for $1 \le p \le \infty$ (and where the L^∞ norm is the essential supremum).

It turns out that if we look at function spaces on a finite measure space like $B_1(0) \subset \mathbb{R}^n$, then the $L^p(B_1(0))$ spaces actually have the **reverse** inclusion. This is true by Holder's inequality – since $|\int fgd\mu| \le ||f||_p||g||_q$ whenever $\frac{1}{p} + \frac{1}{q} = 1$, we can get bounds that show that one norm is finite if the other is. For example ,$L^\infty \subset L^1$ because $\int |f| \le \sup |f| \int_B 1 = |B|||f||_\infty$. For the more general case, suppose $p' < p$. Letting $s = \frac{p}{p'}$, we find that

$$\int |f|^{p'} \le \left(\int (|f|^p)^s\right)^{1/s} \left(\int 1^t\right)^{1/t}$$

where t is chosen so that $\frac{1}{s} + \frac{1}{t} = 1$. Then if f is in L^p, the first term on the right-hand side is finite, and the second term is finite if we're integrating over finite volume, so f is also in $L^{p'}$.

The reason the inclusions go in reverse directions is then sort of that ℓ^p is a measure space in which we have an atomic mass on each of the integers of measure 1, and so we have a total finite mass – it turns out there's a theorem that being atomic is what gives the inclusion in that direction.

There are many other spaces we'll also be studying throughout this class, and we'll mention them now:

Example 14

Let $\Omega \subset \mathbb{R}^n$ be a nice domain (convex and bounded for simplicity). Recall that C^k is the space of functions $\{f : \Omega \to \mathbb{C} : \sup_{x\in\Omega} \max_{|\beta|\le k} |\partial^\beta f| < \infty\}$ (where β is a multi-index $(\beta_1, \cdots, \beta_n)$, $\partial^\beta f$ is taking β_i derivatives in the ith coordinate, and $|\beta| = \beta_1 + \cdots + \beta_n$). Also recall that each C^k is a complete space (by the uniform limit theorem). The **Holder spaces** are defined for any $0 < \alpha < 1$ to be

$$C^{0,\alpha}(\Omega) = \left\{f \in C^0 : \sup_{x \ne x'} \frac{|f(x) - f(x')|}{|x - x'|^\alpha} < \infty\right\}$$

with Holder norm $||f||_{0,\alpha} = ||f||_0 + \sup_{x\ne x'} \frac{|f(x)-f(x')|}{|x-x'|^\alpha}$. Then $C^{k,\alpha}$ is the space of Holder continuous functions whose derivatives up to order k are in $C^{0,\alpha}$.

In other words, being α-continuous for $0 < \alpha < 1$ is a little weaker than being differentiable, and this is a way of giving us a modulus of continuity (if we know that $||f||_{0,\alpha} \le 1$, then we have a way of seeing how quickly $f(x')$ converges to $f(x)$ as $x' \to x$). Essentially we should think of being in $C^{k,\alpha}$ as having "k derivatives" and then an additional "α fractional derivative."

Fact 15

Subspaces of Banach spaces need not be closed. For example, the subspace of continuous functions $C^0([0,1])$ in $L^2([0,1])$ is not closed, because there can be a Cauchy sequence of continuous functions in L^2 which does not converge to a continuous function. The point is that the L^2 norm is much weaker than the C^0, so demanding convergence in L^2 does not give any reason for us to expect continuity. (And remember that C^0 is in fact dense in L^2, because we can approximate simple functions with step functions with continuous functions. Alternatively, we can use the Stone-Weierstrass theorem, which helps prove these kinds of things in more generality.)

Fact 16

Also, even closed subspaces of Banach spaces need not be complemented – in other words, if Y is a closed subspace of a Banach space X, there does not necessarily exist a closed subspace W such that $Y \oplus W = X$ (where this means that we can write the decomposition $x = y + w$ uniquely, such that the norms $||x||$ and $||y|| + ||w||$ are equivalent – we'll talk about this more later).

Theorem 17

Let X be a Banach space. Then the following are equivalent:

- $\dim(X) < \infty$,
- The unit ball $B = \{x : ||x|| \le 1\}$ is compact,
- The unit sphere $S = \{x : ||x|| = 1\}$ is compact.

Proof. The implications that (1) imply (2) imply (3) are clear (since we can take our usual norm on $\mathbb{R}^n$ and use Heine-Borel, and S is a closed subspace of B). For the final implication that (3) implies (1), first suppose $\dim(X) = \infty$. We'll produce a sequence of points $x_1, x_2, \cdots \in S$ as follows: pick some $x_1 \in S$, then pick some x_2 such that $||x_2 - x_1|| > \frac{1}{2}$, and then more generally choose x_{j+1} such that $\inf_{y \in \text{span}(x_1, x_2, \cdots, x_j)} ||x_{j+1} - y|| > \frac{1}{2}$. To do this general step, we can choose some $z \in S$ that is not in V_j (which we can do because we're assuming $\dim(X) = \infty$). We claim that $d = \text{dist}(z, V_j) = \inf_{y \in V_j} ||z - y|| > 0$ – this is true because V_j is closed (it's a finite-dimensional space) so z is a positive distance away from it. Choose some y_j that achieves almost this infimum, so that we have some $||z - y_j|| \le 2d$. (The idea is that $z - y_j$ is "kind of perpendicular" to V_j.) We may then choose

$$x_{j+1} = \frac{z - y_j}{||z - y_j||}.$$

For any $w \in V_j$, we then have

$$\begin{aligned} ||x_{j+1} - w|| &= \left|\left| \frac{z - y_j}{||z - y_j||} - w \right|\right| \\ &= \frac{1}{||z - y_j||} ||z - (y_j + ||z - y_j||w)|| \\ &\ge \frac{d}{||z - y_j||} \end{aligned}$$

because $(y_j + ||z - y_j||w)$ is an element of V_j, and this last expression is indeed greater than $\frac{1}{2}$. (We can get this arbitrarily close to 1 but not exactly equal to 1 – this has to do with the lack of complementability.) So S is not compact because there is no Cauchy subsequence. □

Definition 18

A space X is a **Hilbert space** if it is a Banach space and there is a sesquilinear form $\langle , \rangle : X \times X \to \mathbb{C}$ (linear in the first argument and conjugate linear in the second), such that $||x||^2 = \langle x, x \rangle$ (so in particular the form is nondegenerate and positive definite).

Hilbert spaces include $\mathbb{R}^n, L^2, \ell^2$, but there are also nonseparable Hilbert spaces that are "big." And in general, it turns out that Hilbert spaces are a little bit nicer in terms of complementability and the "closest distance" point from above:

Proposition 19

Let X be a Hilbert space, and let $Y \subsetneq X$ be a proper closed subspace. Then there exists some closed subspace W such that $X = Y \oplus W$ and $W = Y^\perp$, and there is some $x \in X$ with $||x|| = 1$ and $d(x, Y) = 1$.

Before we discuss this further, though, we'll backtrack to Banach spaces again:

Definition 20

Let X and Y be Banach spaces. The **space of bounded linear mappings from X to Y** is defined as

$$\mathcal{B}(X, Y) = \{A : X \to Y : \exists C > 0 \text{ such that } \sup_{x \neq 0} \frac{||Ax||}{||x||} \le C\}.$$

The point is that there are unbounded linear mappings (that are very useful) when our spaces are not finite-dimensional anymore, but the bounded linear mappings are the ones where $||Ax||_Y \le C||x||_X$ for some absolute constant C. And $\mathcal{B}(X, Y)$ is a normed vector space in which we can define the **uniform operator norm**

$$||A|| = \sup_{x \neq 0} \frac{||Ax||}{||x||}.$$

This is a related norm to some other ones we might place on linear transformations in $\mathbb{R}^n$, but the point is that as n gets larger the norms will become not comparable.

Proposition 21

Let X and Y be normed spaces, and suppose Y is complete (that is, Banach). Then $\mathcal{B}(X, Y)$ is again a Banach space.

Proof. Let $\{A_j\}$ be a Cauchy sequence of operators. In other words, for all $\varepsilon > 0$ there is some i_0 such that for all $i, j \ge i_0$, we have $||A_i - A_j|| < \varepsilon$. Our first step is to fix some $x \in X$ and look at the sequence of elements $A_j x \in Y$ – we claim that this is Cauchy in Y. Indeed, we know that for all $i, j > i_0$,

$$||A_j x - A_i x|| \le ||A_i - A_j|| ||x|| < \varepsilon ||x||,$$

so we can define for any x the limiting object $Ax = \lim_{i \to \infty} A_i x$. This operator A is linear (because we can interchange limits and finite sums), and we claim that A is bounded. Indeed, by the triangle inequality we have for each fixed i that

$$||A|| = \sup_{x \neq 0} \frac{||Ax||}{||x||} \le \sup_{x \neq 0} \frac{||(A - A_i)x||}{||x||} + \frac{||A_i x||}{||x||}.$$

But we know that

$$\frac{||(A - A_i)x||}{||x||} = \lim_{j \to \infty} \frac{||(A_j - A_i)x||}{||x||}$$

so as long as $i > i_0$, the first term is at most ε, the second term is bounded for each fixed i and thus the whole right-hand side is finite. So A is bounded, and now we must show that $||A - A_i|| \to 0$ to finish the proof but that's exactly what the calculation above does by taking $\varepsilon \to 0$. $\square$

Corollary 22

If X is Banach and Y is the base field (say $\mathbb{C}$) of our vector space, the space of bounded mappings $\mathcal{B}(X, \mathbb{C})$ is called the **dual space** X^* of X. (In other words, X^* is the space of bounded linear functionals ℓ on X, meaning that $|\ell(x)| \le C||x||$ for all x.) Then X^* is a Banach space if Y is complete.

It's often useful to identify X^* more explicitly if we know X:

Example 23

For any $1 \le p < \infty$, the dual space of ℓ^p is $(\ell^p)^* = \ell^q$, where $\frac{1}{p} + \frac{1}{q} = 1$. (When $p = 1$ we take $q = \infty$, and this is false when $p = \infty$.)

To see why this holds, notice that for any $x \in \ell^p, y \in \ell^q$, Holder's inequality yields

$$\langle x, y \rangle = \sum_{j=1}^{\infty} x_j \overline{y}_j.$$

So $\ell^q \subseteq (\ell^p)^*$ (any element $y \in \ell^q$ yields a linear functional $\ell(x) = \langle x, y \rangle$), and for the reverse inclusion, if we have a linear map $\ell : \ell^p \to \mathbb{C}$ we can take $y_j = \overline{\ell(e_j)}$ and show that that does indeed give us something in ℓ^q. (The details are in the book for us.)

Example 24

Similarly, the dual space of $L^p(M, d\mu) = X$ for any $1 \le p < \infty$ is $X^* = L^q$.

Holder's inequality again proves that $L^q \subseteq (L^p)^*$ (and in fact the L^q norm of g is equal to the norm of the operator sending f to fg, so we have an **isometric inclusion**). The reverse inclusion was done in 205A: the idea is that if we have some linear functional $f : L^p \to \mathbb{C}$, we can define a signed measure ν on M via $\nu(A) = \int_A d\mu$. (We're basically testing against characteristic functions instead of basis vectors.) Taking the Radon-Nikodym derivative, we get $g = \frac{d\nu}{d\mu}$, and then we do some work to show that $g \in L^q$. And in general, the dual of L^∞ is usually very hard to identify and pretty ugly.

Example 25

By the Riesz representation theorem, the dual of C^0 is the space of signed measures. On the other hand, the duals of C^1 or $C^{0,\alpha}$ are much more complicated – trying to characterize the dual of C^1 already requires us to talk about distributions.

In the case where we have a Hilbert space, the characterization is much nicer:

Theorem 26 (Riesz)

Let X be a (not necessarily separable) Hilbert space. For any $\ell \in X^*$, there is some $v \in X$ such that $\ell(x) = \langle x, v \rangle$.

Such a linear functional is always bounded because the Cauchy-Schwarz inequality tells us that $|\langle x, v \rangle| \le ||x|| \, ||v||$. And the Riesz representation theorem says that any linear functional is indeed of this form (taking the inner product with some v). The idea of the proof here is that if we have some linear functional ℓ, we get some hyperplane $H = \{y : \ell(y) = 0\}$ (it's closed because ℓ is continuous and the zeroset of a continuous function is a closed subset).

Then we can find a vector z of norm 1 perpendicular to H; the idea is that the linear functional ℓ is basically projecting along z. And that's how all of this will tie back to Proposition 19, but we're going to prove something a bit more general:

Theorem 27

Let X be a Hilbert space, and let $K \subset X$ be a closed convex **subset** (not necessarily subspace). Then there exists a unique element $\overline{x} \in K$ with $||\overline{x}|| = \inf_{x \in K} ||x||$.

Proof. If $0 \in K$, then we can and must take $\overline{x} = 0$. Otherwise, define $d = \text{dist}(K, 0)$ (which is strictly positive because K is closed – any sequence of elements with norm going to zero must converge to zero). Choose $x_i \in K$ such that $||x_i|| \to d$. It turns out that this sequence will **automatically be Cauchy** – indeed, the midpoint of x_i and x_j is also in K by convexity, so $||x_i + x_j|| > 2d$ by definition of d. Now by the parallelogram rule,

$$||x_i - x_j||^2 = 2||x_i||^2 + 2||x_j||^2 - ||x_i + y_j||^2,$$

and now if we choose i, j large enough so that $||x_i|| < d + \varepsilon$, the right-hand side is at most $4(d+\varepsilon)^2 - 4d^2$, which goes to zero. So now because we're in a complete space, we can define $\overline{x} = \lim x_i$, and this is also in K because the subset is closed. Finally, to check uniqueness, if we have $\overline{x}$ and $\tilde{x}$ both with minimum norm, then again by the parallelogram law

$$||\overline{x} - \tilde{x}||^2 = 2||\overline{x}||^2 + 2||\tilde{x}||^2 - ||\overline{x} + \tilde{x}||^2 \le 0$$

because the first two terms on the left-hand side are each exactly $2d^2$ and the last term is at least $(2d)^2$. Thus $||\overline{x} - \tilde{x}|| = 0$. $\square$

So with that, we can prove the Riesz representation theorem (Theorem 26): let $K = \{x \in X, \ell(x) = 1\}$, which is a closed subset not containing zero. By our previous result, there is some $\overline{x} \in K$ such that $||\overline{x}|| = d(K, 0)$. Then take some $y \in \ell^{-1}(0)$ and notice that we have $||\overline{x} + ty||^2 \ge d^2$ for all real numbers t, with minimum achieved at $t = 0$. Differentiating this expression, we see that $\langle \overline{x}, y \rangle = 0$, so $\overline{x}$ is in fact perpendicular to all $y \in \ell^{-1}(0)$. And the rest of the proof is like how it looks in $\mathbb{R}^m$ – the point is to prove that the orthogonal complement is in fact one-dimensional.

3 January 17, 2023

Last time, we claimed that if X is a Hilbert space, then there is a natural isomorphism $X \to X^*$ which is actually an isometry. In other words, we need to show that for any $v \in X$ we can find an $\ell_v \in X^*$, and vice versa, so that $\ell(y) = \langle y, x \rangle$. To do this, we look at $\ell^{-1}(0)$ and its "perpendicular subspace" $\ell^{-1}(0)^\perp = \{w \in X : \langle w, y \rangle = 0 \ \forall y \in \ell^{-1}(0)\}$. This can be quickly checked to be a closed subspace, and it is in fact one-dimensional, because for any $u_1, u_2 \in \ell^{-1}(0)^\perp$ we have $\ell(u_1), \ell(u_2) \neq 0$, so $\ell(au_1 + bu_2) = 0$ for some a, b, meaning $au_1 + bu_2 \in \ell^{-1}(0) \cap \ell^{-1}(0)^\perp = \{0\}$ and thus there is a linear dependence between any u_1, u_2. Then our next step is to take any unit vector v_0 in this one-dimensional subspace $\ell^{-1}(0)^\perp$ and define $\ell(y) = a\langle y, v_0 \rangle$, and we can show that this does work. (Alternatively, the point is that $\ell^{-1}(1)$ is a closed convex set, so we can find a unique shortest element in it. So a lot of these arguments in Hilbert spaces are rather Euclidean in nature.) So this is a useful fact, and we can check that $||v|| = ||\ell|$ (remember that the norm of a functional ℓ is $\sup_{x \neq 0} \frac{|\ell(x)|}{||x||}$) because $\sup \frac{|\langle x, v \rangle|}{||x||}$ is at most $||v||$ by Cauchy-Schwarz, and the supremum is attained at v itself.

We'll come back to Hilbert spaces again when we come back to Fourier spaces and need the setting of L^2. But for now, we'll turn our attention to some structural theorems that are important in more generality.

Definition 28

Let (X, d) be a metric space. A subset $A \subset X$ is **nowhere dense** if the closure $\overline{A}$ has empty interior. A is **meagre** (also **first category**) if it is a countable union of nowhere dense sets. A is **nonmeagre** (also **second category**) if it is not meager and **residual** if $X \setminus A$ is meagre.

Being meagre is a different notion of being a "small set" than having small Lebesgue measure – it's possible to have meager sets that are full in measure. But the point is that there is still "a lot of stuff left over" if we're working in a complete metric space.

Example 29

The integers $\mathbb{Z}$ form a meagre set in $\mathbb{R}$, and so do Cantor sets (including fat Cantor sets with positive Lebesgue measure).

It's useful to note that A **is nowhere dense if and only if** $X \setminus A$ **contains a dense open set** – if $\overline{A}$ has empty interior, then $X \setminus \overline{A}$ is open, smaller than $X \setminus A$, and dense because for any ball around a point in X cannot be entirely contained inside $\overline{A}$.

Theorem 30 (Baire category theorem)

Let (X, d) be a nonempty complete metric space. Then the following are true:

1. Every residual set is dense,
2. If a set U is open and nonempty, then it is nonmeagre,
3. If A_i is a (countable) sequence of closed sets with no interior, then $\bigcup_i A_i$ has no interior.
4. If U_is are open and dense, then $\bigcap U_i$ is dense,
5. Every residual set is nonmeagre.

It turns out that condition (3) will be the most useful for us – for example, suppose we have a linear mapping between two Banach spaces. If we look at the image of the unit ball of the first space, we want to show that it has an interior in the second space; otherwise we can let A_ns be the images of the balls of radius n and apply this result. And once we have an interior, we can show that maps between these spaces are invertible.

Proof. For (1) implies (2), if U is an open set, then $X \setminus U$ is not dense and thus not residual. For (2) implies (3), notice that $\bigcup A_i$ is meagre by definition because each A_i is nowhere dense, so the interior of $\bigcup A_i$ is also meagre (since it's a smaller subset). But that means the interior must be empty if it is nonmeagre. To get from (3) to (4), we just take the complements of the A_is. Then to show that (4) implies (1), let R be a residual set. That means that there are dense open sets U_i such that $\bigcap U_i \subset R$ (take the complement of R, which is meagre and thus contained in the union of closed sets with no interior). But because $\bigcap U_i$ is dense, property (4) tells us that R is also dense, which is what we want. Thus, all of (1), (2), (3), (4) are equivalent. Next, we show that (2) implies (5). If a residual set R were also meagre, then $X = (X \setminus R) \cup R$ must be meagre as well, which is not true because X is open and nonempty. Thus R is nonmeagre. (So all of this is basically a bunch of set theory equivalences.)

It now suffices to show that (4) is true, which will prove that all of these claims are true. Suppose we have open sets U_i that are open and dense (so they all "spread almost everywhere"). We must show that for any $x_0 \in X$ and $\varepsilon_0 > 0$, we have $B_{\varepsilon_0}(x_0) \cap \bigcap U_i \neq \varnothing$. We'll construct points $x_k \in U_k$ such that $0 < \varepsilon_k < 2^{-k}\varepsilon_0$, such that

$\overline{B_{\varepsilon_k}(x_k)} \subset U_k \cap B_{\varepsilon_{k-1}}(x_{k-1})$ for all $k \geq 1$ (so we basically have nested closed balls). We'll do this inductively: for $k = 1$, since U_1 is an open **dense** set, we know that $U_1 \cap B_{\varepsilon_0}(x_0)$ is nonempty, so we can take any x_1 inside this set and choose $\varepsilon_1 < \frac{\varepsilon_0}{2}$ so that $\overline{B_{\varepsilon_1}(x_1)} \subset U_1 \cap B_{\varepsilon_0}(x_0)$. The point is that x_1 is in both open sets U_1 and $B_{\varepsilon_0}(x_0)$ so we can find a small enough ball so that even the closure is contained within both open sets. Going from $(k-1)$ to k is exactly the same procedure. But now because $d(x_{k-1}, x_k) < \varepsilon_{k-1} < \varepsilon_0 2^{-k}$, we see that

$$d(x_0, x_k) \leq d(x_0, x_1) + \cdots + d(x_{n-1}, x_n) < \varepsilon_0$$

for all k, and in fact the x_ns are Cauchy and thus converging to some $\overline{x}$. But now x_ℓ is in $B_{\varepsilon_k}(x_k)$ for all $\ell \geq k$ by the same triangle inequality argument as before, so the limit point $\overline{x}$ must be in $\overline{B_{\varepsilon_k}(x_k)}$ (here is where we see the advantage of closed balls), which lies inside U_k for all k. So this limit point is in all of the open sets, and $\overline{x}$ is also inside $\overline{B_{\varepsilon_1}(x)} \subset B_{\varepsilon_0}(x_0)$. So $B_{\varepsilon_0}(x_0) \cap \bigcap_i U_i$ is indeed nonempty. □

The key point is that a intersecting (countably many) dense sets still gives us something dense. And even if this proof is relatively unintuitive, it already has some applications:

Proposition 31

There is no infinite-dimensional Banach space with a countable basis.

The details here are left to us as an exercise. Here, we're using the linear algebraic definition where elements must be written as finite linear combinations of basis elements – something like Fourier series in the Hilbert space basis sense allows countable linear combinations, so it's not a counterexample.

Proposition 32

Let $f : \mathbb{R}^+ \to \mathbb{R}$ be continuous, and suppose $\lim_{n\to\infty} f(nt) = 0$ for all $t > 0$ (so behavior is good along subsequences). Then $\lim_{t\to\infty} f(t) = 0$ as well.

This problem may look rather artifically constructed, but there are versions of this useful for example in convergence of Fourier series.

Proof. For each positive integer n, define the closed set

$$A_n = \{t : |f(mt)| \leq \varepsilon \text{ for all } m \geq n\}.$$

We know that the union of the A_ns is all of $\mathbb{R}^+$ by assumption, so one of them must have an interior point. Thus some A_n contains a nontrivial interval $[a, b]$ of $\mathbb{R}^+$. But now the intervals $[a, b], [2a, 2b], [3a, 3b]$ eventually overlap, because there is some large enough N such that $Na < (N-1)b$, specifically because $\frac{b}{a} > 1$ and thus $\frac{N}{N-1}$ is eventually less than it. So once t gets into this overlapping range, $|f(t)|$ will be at most ε everywhere. □

Remark 33 (Aside)**.** *Theorems like this can be very delicate – we used the structure $f(nt)$ explicitly in this proof. Here's an unrelated but similar-looking theorem that might be interesting to us: if f is a smooth decreasing function on $\mathbb{R}$, and $\hat{f}(\pm\sqrt{n})$ for all n, then $f = 0$. And this is related to the densest sphere packing in certain dimensions.*

Proposition 34

Let $p \in C^\infty(\mathbb{R})$, and suppose that for all x there exists some n such that $p^{(n)}(x) = 0$. Then all derivatives vanish past some point, meaning that p is a polynomial.

(This one is left as an exercise to us, and it's similar to the setup of the previous example.) This next result is less artificial:

Theorem 35

The set of functions $\{f \in C^0(I) : f \text{ nowhere differentiable}\}$ is dense in C^0.

Partial sketch. Define the sets

$$E_{m,n} = \{f \in C^0(I) : \exists x : \text{ for some } y \text{ with } 0 < |y - x| \le \frac{1}{m}, \text{ we have } |f(x) - f(y)| \le n|x - y|\}$$

In other words, it's the set where some difference quotient close enough to x is at most n. Then $\{f : f \text{ differentiable somewhere}\}$ is contained within $\cup E_{m,n}$, but each $E_{m,n}$ is closed and has empty interior in C^0 (we just need to perturb our function in uniform norm slightly). Then by Baire the union of the sets must be meagre. □

We haven't applied any properties of functional analysis yet – this is the "bare" theorem, so to speak – and now we'll see how this relates to Banach spaces.

Theorem 36 (Uniform boundedness principle / Banach-Steinhaus)

Let X be a Banach space, and let Y_α be a family of normed spaces (not necessarily complete), indexed by some possibly uncountable set. Suppose $A_\alpha : X \to Y_\alpha$ are linear maps, and for all $x \in X$, there is some C_x such that $\sup_{\alpha \in A} ||A_\alpha x|| \le C_x$. Then there is some absolute constant such that $||A_\alpha|| \le C$ (in other words, C_x is just linearly bounded in x).

One thing it's important to know here is that there are many different topologies we can place on the set of bounded linear operators. Specifically, suppose $A_i : X \to Y$ is a sequence of such linear operators. We have the following three notions:

- Uniform operator topology: $A_i \to A$ if $||A_i - A|| \to 0$ in operator norm.
- Strong topology: $A_i \to A$ converges strongly if $A_i x \to Ax$ for all x. (So we're testing our operators against particular vectors.)
- Weak topology: for all $x \in X$ and all $\ell \in Y^*$, $\ell(A_i x)$ converges. (So here we're testing against particular vectors **and** pairing against linear functionals on Y.)

The point is that the weaker topologies are much weaker, but it is easier to prove convergence in them.

Proof. Define the functions

$$f_\alpha(x) = ||A_\alpha x||_{Y_\alpha}$$

for every α. We know that $f_\alpha(x) \le C_x$ for all α, and now we can define the sets

$$\mathcal{F}_{\alpha,n} = \{x : |f_\alpha(x)| \le n\}, \quad \mathcal{F}_n = \bigcap_\alpha \mathcal{F}_{\alpha,n}.$$

These are all closed sets because arbitrary intersections of closed sets are closed. Then $X = \bigcup_n \mathcal{F}_n$ by assumption, so by Baire there is some $\mathcal{F}_n$ with nonempty interior. Thus there is some n, some $x_0 \in X$, and some $\varepsilon > 0$, such that $B_\varepsilon(x_0) \subset \mathcal{F}_n$. In other words,

$$\sup_\alpha \sup_{x \in B_\varepsilon(x_0)} |f_\alpha(x)| \le n,$$

so we already have the uniformity over all points $x \in B_\varepsilon(x_0)$: $\sup_\alpha \sup_{x \in B_\varepsilon(x_0)} ||A_\alpha x||_{Y_\alpha} \le C$ for some finite C.

Now if we take any $x \in X$ with $||x|| = 1$, the points $x_0 \pm \frac{\varepsilon}{2}x$ are in the ball and thus $||A_\alpha(x_0 \pm \frac{\varepsilon}{2}x)|| \le C$. But now by linearity we know that

$$\varepsilon A_\alpha x = A_\alpha \left(x_0 + \frac{\varepsilon}{2}x\right) - A_\alpha \left(x_0 - \frac{\varepsilon}{2}x\right),$$

so $||A_\alpha x|| \le \frac{2}{\varepsilon}C$ by the triangle inequality. Since this is true for any x of norm 1, by homogeneity we see that all operator norms are indeed uniformly bounded. $\square$

The useful trick here is that we go from a ball far from the origin to a ball near the origin using properties of the norm. And this has a big application in Fourier analysis:

Theorem 37

There is a continuous function $f \in C^0(S^1)$, such that if its Fourier series is $\sum_{-\infty}^{\infty} a_n e^{inx}$ (here defined by $a_n = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-inx} f(x)dx$, and by the L^2 theory we know that the finite partial sums $S_N f = \sum_{-N}^{N} a_n e^{inx}$ converge in L^2 to f), then $S_N f(0)$ is unbounded.

In other words, we can have a continuous function which does not converge pointwise. (If we choose something more subtle than the finite partial sums or if f has better regularity, we do get much better convergence theorems, though. We'll talk more about this next month.)

Proof. We can think of the partial sums as

$$S_N f(x) = \sum_{-N}^{N} \frac{1}{2\pi} e^{inx} \int_0^{2\pi} f(y) e^{-iny} dy,$$

and now interchanging the order of finite sum and integral yields

$$= \int_0^{2\pi} \left(\sum_{-N}^{N} \frac{1}{2\pi} e^{in(x-y)} \right) f(y) dy.$$

This inner parenthetical term $D_N(x-y)$ is easy to compute because it is a finite geometric series – it turns out to be $\frac{1}{2\pi}\frac{\sin(N+\frac{1}{2})x}{\sin\left(\frac{1}{2}x\right)}$. But then the map $E_N : f \to S_N f(0)$ can then be represented as integration against $D_N(x)$. Since L^1 is dual to L^∞, the operator norm of E_N is bounded by the L^1 norm of D_N, which is $\int_0^{2\pi} |D_N(x)|dx$. (Here it's important that we're using the C^0 norm and not the L^2 one for f.) So now there are two possibilities: if $S_N f(0)$ is bounded for any fixed f, then $||E_N||$ must be uniformly bounded. And that can be proven false by explicit integral estimation: D_N is basically a damped oscillation, and we can estimate its L^1 norm by estimating each trig part by a triangle. It then turns out we get an estimate of the form $\int |D_N| \ge C_1 + C_2 \sum_{k=1}^{N} \frac{1}{k}$, and the right-hand side is unbounded. $\square$

4 January 19, 2023

The first problem set's solutions will be posted this afternoon – we should look over them and make sure we understand how everything works.

Last time, we proved the Baire category theorem, which is a useful result for proving some other notable facts. These abstract theorems are nice structurally, but it will take us a while before the applications all become clear.

Theorem 38 (Alternate statement of Banach-Steinhaus)
Let $A_i : X \to Y$ be a sequence of mappings between two fixed Banach spaces. Then the following are equivalent:

1. A_i "converges strongly," meaning that $\{A_i x\}$ converges for each $x \in X$,
2. $\sup_i ||A_i||$ is finite, and there exists a dense set $D \subset X$ on which $\{A_i x\}$ is Cauchy for all $x \in D$,
3. $\sup_i ||A_i||$ is finite, there is some A such that $A_i \to A$ strongly (this is the actual statement of strong convergence), and $||A|| \le \liminf ||A_i||$.

Proof. It is clear that (3) implies both (1) and (2) (we can take the dense set to be everything). To show that (1) implies (3), the uniform boundedness principle implies that $\sup ||A_i|| \le C$, and then we can define a linear map by setting $Ax = \lim_{i\to\infty} A_i x$ for all x. This map is linear, and for any x we have

$$||Ax|| = \lim_{i\to\infty} ||A_i x|| \le \liminf_{i\to\infty} ||A_i x|| \le (\liminf ||A_i||)||x||$$

(it's useful to make sure we understand why we need liminf instead of lim in the last inequality here), and this implies that $||A|| \le \liminf ||A_i||$.

Finally, to show that (2) implies (3), we just need to produce A. Define $C = \sup ||A_i||$ and notice that for any $x \in X$, we may choose some $\overline{x} \in D$ such that $||x - \overline{x}|| < \varepsilon$. Then

$$||A_i x - A_j x|| \le ||A_i(x - \overline{x})|| + ||(A_i - A_j)\overline{x}|| + ||A_j(\overline{x} - x)||.$$

Since $\overline{x} \in D$, the middle term is small, and the first and last terms are small because the A_is have bounded norm. Thus we have a Cauchy sequence and can in fact define Ax. □

The idea is that we often want to define a family of operators on some dense subspace (like step functions or C^∞) where analysis is easier to perform. And the Fourier series case gives us a sense that we often care about whether operators are bounded – here's an example of a question we'll be answering later on:

Problem 39
Suppose $f \in L^2(S^1)$, so that we have a sequence of Fourier coefficient $a_n = \frac{1}{2\pi}\int_0^{2\pi} e^{-inx} f(x)dx$ with $\sum |a_n|^2 < \infty$ (so $\{a_n\} \in \ell^2$). If $f \in L^1(S^1)$ instead, then we know that $\{a_n\} \in \ell^\infty$, and in fact by Riemann-Lebesgue we have $\{a_n\} \in c_0$ (prove this first for indicators of intervals, then for indicators of step functions, and linear combinations of step functions are dense in L^1). So the Fourier operator is a mapping $L^2 \to \ell^2$ and also a mapping $L^1 \to c_0$, so we may ask what happens to the mapping on L^p for some $1 < p < 2$ – we may want to ask to prove that it's a bounded linear operator on some space. (It turns out to map to ℓ^q, where $\frac{1}{p} + \frac{1}{q} = 1$, using interpolation theorems.) And the point is that we usually need to verify the conditions for one of our abstract theorems instead of doing things bare-handed.

Theorem 40 (Open mapping theorem)
Let $A : X \to Y$ be a bounded linear map between Banach spaces, and suppose that A is surjective. Then A is an open map, meaning that the image of any open set in X is open in Y.

Proof. Let B be the unit ball in X, and let C be its image under A. The first step is to show that there is some δ such that $B_\delta(0)$ in Y is contained in $\overline{A(B)}$. (In other words, the image of the unit ball is not squished too much – its closure has 0 as an interior point.) Indeed, X is the union of dilates of the unit ball, so

$$X = \bigcup_{n=1}^{\infty} nB \implies Y = \bigcup_{n=1}^{\infty} nC$$

by surjectivity of A. Then Baire tells us that the nCs cannot all be nowhere dense, so there is some n such that $\overline{nC}$ has interior – in particular, this means that $\overline{\frac{1}{2}C}$ also has interior. If we now pick any interior point, namely a point $y_0 \in Y$ such that $B_\delta(y_0) \subset \overline{\frac{1}{2}C}$, then for any y with $||y|| < \delta$ we have $y + y_0 \in B_\delta(y_0) \subset \overline{\frac{1}{2}C}$ and $y_0 \in \overline{\frac{1}{2}C}$, and thus there are some $x_i, x_i' \in \frac{1}{2}B$ such that $Ax_i \to y + y_0$ and $Ax_i' \to y_0$. Thus $y = \lim_{i\to\infty} A(x_i - x_i')$. Since x_i, x_i' are in $\frac{1}{2}B$, $x_i - x_i'$ is in B and thus $y \in \overline{A}(B)$.

The next step is to show that $B_\delta(0)$ is actually inside $CC = A(B)$, not just its closure. Fix any y with $||y|| < \delta$ and let $\varepsilon = \delta - ||y||$, and construct a sequence x_n as follows. First choose x_0 so that $||x_0|| < \frac{1}{\delta}||y||$, and then pick subsequent terms so that

$$||x_n|| < \frac{\varepsilon}{\delta 2^k}||y - A(x_0 + \cdots + x_n)|| < \frac{\varepsilon}{\delta 2^{k+1}}$$

(we want these to be successively better approximations as solutions to $y = Ax$, and the point is to stay within the ball of radius δ). In the first step, we showed basically that for any $y \in Y$, y lies in the set $\overline{\{Ax : ||x|| < \frac{1}{\delta}||y||\}}$. Indeed, if we take $\frac{\delta}{1+\eta}\frac{y}{||y||}$ (which has norm strictly less than δ), this lies in $B_\delta(0)$ and thus lies in $\frac{\delta}{1+\eta}\frac{y}{||y||} \in \overline{A(B)}$. Multiplying through, this means that for every $\eta > 0$ we get

$$y \in \overline{A\left(\frac{1+\eta}{\delta}||y||B\right)}$$

so taking $\eta \to 0$ yields the result. So then our goal is to pick some x_0 so that $||y - Ax_0|| < \frac{\varepsilon}{2}$ and $||x_0|| < \frac{1}{\delta}||y||$, and now for the undictve step replace y with $y - Ax_0$ and do the same thing. (For example, there exists an x_1 so that $||x_1|| < \frac{1}{\delta}||y - Ax_0||$, and we can choose it so that $||y - A(x_0 + x_1)|| < \frac{\varepsilon}{4}$, and so on.) So now we can define $x = x_0 + x_1 + \cdots$, and by the triangle inequality we have $||x|| \le \frac{1}{\delta}||y|| + \frac{\varepsilon}{\delta}$. But because we chose $\varepsilon = \delta - ||y_0||$, this is exactly telling us that $||x|| < 1$, so we have found a point inside the unit ball such that $y = Ax$.

Finally, suppose $U \subset X$ is an open set in X, and suppose we have some arbitrary $y_0 \in A(U)$. There is some $x_0 \in U$ such that $Ax_0 = y_0$, and there exists some $B_\varepsilon(x_0) \subset U$; we claim that $B_{\delta\varepsilon}(y_0) \subset A(U)$. The point is that (from step 2)

$$||y - y_0|| < \delta\varepsilon \implies \left|\left|\frac{y - y_0}{\varepsilon}\right|\right| < \delta,$$

so there is some $\hat{x} \in B$ such that $A\hat{x} = \frac{y-y_0}{\varepsilon} = \frac{y-A(x_0)}{\varepsilon}$. In particular, this means $A(x_0 + \varepsilon\hat{x}) = y$, so being within a small ball of an arbitrary y_0 means we're in the image of A. □

Intuitively, the point is that between Banach spaces, either the image is a compressed subspace (for example the inclusion mapping $\iota : C^1 \to C^0$) with no interior, or it fills out everything and is "really continuous.'

Corollary 41

If $A : X \to Y$ is a bounded bijective linear map (meaning the inverse map exists), then A^{-1} is also a **bounded** linear map.

(Remember that being bounded and being continuous are the same thing for Banach spaces.)

Proposition 42

Suppose Y, Z are two closed subspaces of a Banach space X. Then we can look at their sum $Y + Z = \{y + z : y \in Y, z \in Z\}$. If we have that $Y + Z = X$ and $Y \cap Z = \{0\}$ (purely algebraic statements about decomposing as a sum), then $X \cong Y \oplus Z$; that is, there are constants c_1, c_2 such that $c_1||x|| \le ||y|| + ||z|| \le c_2||x||$ for all $x \in X$.

Proof. We know that $||x|| \le ||y|| + ||z||$ by the triangle inequality, so we can always take $c_1 = 1$. Now define the topological Banach space

$$W = Y \oplus Z = \{(y, z) : y \in Y, z \in Z, ||(y, z)||_W = ||y|| + ||z||\}.$$

Then the map $W \to X$ sending (y, z) to $y + z$ is bijective by hypothesis, so the open mapping theorem tells us that the inverse map is bounded and thus $||(y, z)||_W \le C||x||$ as well. □

There's also the case where Y and Z can have nontrivial intersection:

Corollary 43

Suppose that $X = Y + Z$ and $V = Y \cap Z$. Then for all $x \in X$, there is **some** decomposition $y \in Y$ and $z \in \mathbb{Z}$ such that $||x|| \le ||y|| + ||z||$.

We'll now mention the final big structural theorem that we should know, but first we should keep the following example in mind:

Example 44

We may think of the derivative map $\frac{d}{dx}$ as a map $C^1 \to C^0$, which makes it a bounded operator. or we may think of it as a map $C^0 \to C^0$ with domain restricted only to $D(A) = C^1$.

Definition 45

If $A : X \to Y$ is only defined on some dense subset $D(A) \subset X$, then we can define the graph $\text{gr}(A) = \{(x, Ax) : x \in D(A)\} \subset X \oplus Y$. An operator is a **closed operator** if $\text{gr}(A)$ is a closed subspace, and it is **closable** (that is, it doesn't have any "vertical lines" in the usual function graph sense) if $\overline{\text{Gr}(A)} \cap (\{0\} \times Y) = \{(0, 0)\}$ (this is sufficient by linearity).

The point, though, is that this notion of "closedness" is not an interesting generalization if $D(A)$ is the whole space:

Theorem 46 (Closed graph theorem)

Let $A : X \to Y$ be a linear operator with $D(A) = X$. Then A is bounded if and only if it is closed.

Proof. If A is bounded, then it is indeed closed. To check this, if (x_i, Ax_i) is a sequence of points in $\text{gr}(A)$, then we must check that $x_i \to x$ in X and $Ax_i \to y$ in Y, then $y = Ax$ (the limit is also in the graph). But this is true by continuity of the operator (since it is bounded, thus continuous).

The opposite direction requires the open mapping theorem: if $\text{Gr}(A)$ is closed, then it is a closed subspace of a Banach space, so it is also a Banach space (with norm $||(x, Ax)|| = ||x|| + ||Ax||$). Let π_X be the projection $\text{Gr}(A) \to X$; this is a bijective map because $D(A) = X$ and A is a function, and it is bounded because $||x|| \le ||x|| + ||Ax||$ (so it has norm at most 1). Thus its inverse is also a bounded mapping, so the map $x \mapsto (x, Ax)$ is bounded. Composing this with the bounded map π_Y shows that A is bounded. □

Example 47

As before, take $\frac{d}{dx}$ as an operator from $C^0 \to C^0$, with $D(\frac{d}{dx}) = C^1$ (by definition C^1 is the largest space on which the derivative lands in C^0). We claim $\frac{d}{dx}$ is a closed operator (but not bounded).

Indeed, if we start with some sequences $u_i, u_i' = f_i \in C^0$, and we know that $u_i \to u$ and $f_i \to f$ in C^0, then $u' = f$ (because u_i, u_i' both converge uniformly, and thus u_i converges to some $v \in C^1$, but u_i converges to u in C^0, so $u = v$ because converging in C^1 implies converging in C^0; thus $u \in C^1$ and $u' = f$). So the graph is indeed closed.

Example 48

Now consider the operator $\frac{d}{dx} : C^0 \to C^0$ but now with the domain $D(\frac{\partial}{\partial x}) = C^\infty$. We claim this operator is closable but not closed.

Here, the closure of the graph of $\text{Gr}((\frac{d}{dx}, C^\infty))$ is exactly $\text{Gr}((\frac{d}{dx}, C^1))$ (exercise). And we can extend this example slightly: take an arbitrary open set Ω in $\mathbb{R}^n$, and consider the linear differential operator

$$P = \sum_{|\alpha| \le m} a_\alpha(x) \partial_x^\alpha,$$

where each α is a multi-index $(\alpha_1, \cdots, \alpha_n)$ and each a_α is a $C^\infty(\Omega)$ function. We claim that $(P, C_0^\infty(\Omega))$ is **always closable** in any $L^p(\Omega)$. In other words, if we have a sequence of functions $u_i \in C_0^\infty(\Omega)$ such that $u_i \to u$ in L^p, and $Pu_i = f_i \to f$ in L^p, then in fact we have $Pu = f$. And in fact, remembering the definition of closability, we can just check that if the limit of the u_is is **zero**, then the limit of the fs also must be zero. For this, we consider the adjoint operator

$$P^t = \sum_{|\alpha| \le m} (-1)^{|\alpha|} \partial_x^\alpha (a_\alpha \cdot).$$

Then for any $u, v \in C_0^\infty(\Omega)$ (this space chosen so we can do integration by parts), we have $\int_\Omega (Pu)v = \int_\Omega u(P^t v)$ because we pick up a negative sign for each derivative we must take. So now fix some $v \in C_0^\infty$; since $Pu_i \to f$ in L^p, we have

$$\int (Pu_i)v \to \int fv, \int u_i(P^t v) \to \int uP^v = 0.$$

But because $\int (Pu_i)v = \int u_i(P^t v)$ by definition, this means $\int fv = 0$ for any v. Therefore $f = 0$ and we do have closability. (So more explicitly, the thing we checked – **closability** – is that $(u_i, Pu_i) \to (0, f)$ implies $f = 0$, so even when we take the closure we won't get multiple points lying over 0 in the first coordinate.

Example 49

On the other hand, an example of a nonclosed operator is the following linear functional: consider the map $\ell : L^2(\mathbb{R}) \to \mathbb{R}$ with domain $D(\ell) = C_0^\infty$, such that $\ell(f) = f(0)$. But then there is a sequence $f_n \in C_0^\infty$ with $f_n \to 0$ in L^2 and $\ell(f_n) \to b$ in $\mathbb{R}$ but $b \ne 0$. And this is true if we take bump functions which get shrunk in the horizontal direction.

The difficulty here is that L^2 topology is too weak to capture the C_0^∞ behavior, and that's why closability fails.

5 January 24, 2023

We'll be discussing the last big structural theorem of Banach spaces today, but first we'll take a general view of "why we do functional analysis." One question that we often like to answer is whether we can solve an operator equation $Ax = y$, which breaks down into a few parts: understanding the correct spaces on which we're trying to solve the equation, knowing what the spaces look like (such as the underlying geometry), deciding whether the operator is bounded, and so on. This last question can come up in a lot of ways, and we saw this last time with our discussion of the closed graph theorem. To review, we started entertaining the possibility of only defining an operator $A : X \to Y$ on a dense subspace $D(A) \subset X$ (for example, have the Laplacian act on continuous functions, but really only having it act on the sufficiently differentiable ones). Then the closed graph theorem says that when $D(A)$ is actually everything and the graph is closed, then A is actually bounded.

Example 50

Let $A : L^2 \to L^2$ be a bounded operator, and suppose we actually know that for any $u \in C^0$, Au is also continuous (so the smaller space is also preserved). So we've defined a map $C^0 \to C^0$, and we may want to know whether it is bounded. It turns out the answer is yes, and that's a nice exercise for us.

So the "best" operators are the bounded operators, and the next best ones are the closed operators (meaning that $\text{gr}(A)$ is a closed subspace of $X \times Y$). After that are the closable operators, which is where A is defined on some domain $D(A)$ (such as C^∞) such that the subspace $\overline{\text{gr}(A)}$ is still a "graph," meaning that its intersection with $\{0\} \times Y$ is just $\{(0, 0)\}$. Beyond that is the world of nonclosable operators, and we won't think about those much at all.

We'll first talk about the last big structural theorem of Banach spaces, though:

Definition 51

A function p from a topological vector space to $\mathbb{R}$ is a **seminorm** if $p(\lambda x) = |\lambda| p(x)$ (homogeneity) and $p(x+y) \leq p(x) + p(y)$ (triangle inequality).

Notice that this is in fact a function from a topological vector space to $\mathbb{R}_{\geq 0}$, since

$$0 = p(0) + p(x + (-x)) \leq p(x) + p(-x) = 2p(x)$$

for any x.

Example 52

Any norm is a seminorm, and for any linear functional $\ell : X \to \mathbb{R}$, $p(x) = |\ell(x)|$ is a seminorm. Notice that the set of points where $p(x) \leq 1$ is a unit ball with respect to the former seminorm is a unit ball but an "infinite slab" with respect to the latter.

For our next example, let X be any topological vector space, and let $S \subset X$ be a subset. A point x_0 is an **interior point** if for all $y \in X$, there is some ε such that $\{x_0 + ty : |t| < \varepsilon\}$ is contained within S. (So this ε distance that we can travel may depend on the direction, so it's a very weak sense of being interior.)

Example 53
Let K be a convex set containing 0 as an interior point, and define

$$p_k(x) = \inf\left\{a : \frac{x}{a} \in K\right\}.$$

The point is that for large enough a, we will have $\frac{1}{a}$ small enough that we can contract x to be close to 0, so $p_k(x)$ is finite for all x.

We claim that this is in fact a seminorm. Homogeneity is clear, and for triangle inequality we will use convexity. For any $x, y \in X$, we may choose a, b so that $\frac{x}{a} + \frac{y}{b} \in K$. By convexity,

$$\frac{a}{a+b} \cdot \frac{x}{a} + \frac{b}{a+b} \cdot \frac{y}{b} = \frac{x+y}{a+b}$$

is thus also in K, so $p_k(x+y) \le a+b$. Taking the infimum over all a and b shows that this is indeed the case. The point is then that the $p(x)$s over all linear functionals ℓ gives us the weak topology, and over any topological space a system of convex sets like K will also give us a topology. But we'll now turn to our big structural theorem:

Theorem 54 (Hahn-Banach)
Let X be a topological vector space (with base field eal), and let $p : X \to \mathbb{R}$ be a seminorm. Let $Y \subset X$ be any linear subspace, and suppose $\wp : Y \to \mathbb{R}$ is a function with $\phi(x) \le p(x)$ (in particular this means $|\phi(x)| \le p(x)$). Then there is a linear functional $\Phi : X \to \mathbb{R}$ such that $\Phi = \wp$ on Y and $|\Phi(x)| \le p(x)$.

So weak control on the magnitude allows us to extend our function to the whole space with control on the magnitude everywhere.

Proof. First we check that we can extend the space by one dimension. Choose any $x_0 \notin Y$ and consider the space $Y' = Y + \mathbb{R}x_0$. We wish to define $\wp' : Y' \to \mathbb{R}$ bounded by $p(x)$ for all $x \in Y'$. Any element of Y' can be written as $x = y + \lambda x_0$ for $y \in Y$, and thus we define $\wp'(x) = \wp(y) + \lambda\wp'(x_0)$. So we just need to determine whether there's a way to choose this single real number $a = \wp'(x_0)$.

Notice that it suffices by homogeneity to check $\lambda = \pm 1$ (since we can always divide by λ and get something in Y, plus or minus x_0). The point is that we want to make sure $\wp(y) \pm a \le p(y \pm x_0)$: in other words, we require

$$\wp(y) - p(y - x_0) \le a, \quad a \le p(y' + x_0) - \wp(y')$$

for all $y, y' \in Y$. But we know that

$$\wp(y) + \wp(y') = \wp(y + y') \le p(y + y') \le p(y + x_0) + p(y' - x_0)$$

(middle step by the seminorm estimate), and thus we always have

$$\wp(y') - p(y' - x_0) \le p(y + x_0) - \wp(y).$$

but that's exactly what we wanted to show since $\sup(\wp(y') - p(y' - x_0) \le \inf(p(y + x_0) - \wp(y))$, so we can choose some a which is somewhere in between those two values.

For the next step, we use Zorn's lemma: consider the set of extensions $(\tilde{Y}, \tilde{\phi})$ (with partial order given by containment of the Ys and agreement of the ϕs on the smaller one), such that $Y \subseteq \tilde{Y} \subseteq X$, $\tilde{\phi}$ restricts to ϕ on Y, and

$|\tilde{\phi}(\tilde{y})| \le p(\tilde{y})$. There is some maximal element $(\overline{Y}, \overline{\phi})$, but if this were anything smaller than X then we could use the first step to get a larger extension. Thus there is some linear functional on all of X, as desired. □

Corollary 55

Let X be a Banach space. Then the dual space X^* of continuous linear functionals is always nontrivial (it contains a nonzero element).

Proof. Start by taking any $x_0 \in X$ (say of unit norm) and define $\ell(x_0) = 5$. Then we can extend this to a linear functional on all of X satisfying the bound $\ell(x) \le 5||x||$. □

There's a more geometric version of Hahn-Banach as well, which will help us enhance the dual space significantly. Here's the setup: let X be a topological vector space, and let K be a convex set with 0 in its interior. We thus get a function $p_k(x)$ from Example 53. Notice that $p_k(x) \le 1$ if $x \in K$, with $p_k(x) < 1$ for any interior point.

Proposition 56

If p is a seminorm, $\{x : p(x) \le 1\}$ is convex, and $\{x : p(x) < 1\}$ is convex with 0 in the interior.

Theorem 57

Let K be nonempty and convex, and suppose that all points of K are interior. Then there is some y not in K, some linear functional $\ell : X \to \mathbb{R}$, and some $c \in \mathbb{R}$ such that $\ell(y) = c$ and $\ell(x) < c$ for all $x \in K$

Proof. Assume that $0 \in K$ and let p be the seminorm associated to K. Remember that $p_k(x) < 1$ if and only if $x \in K$. Then $p(x) < 1$ if and only if $x \in K$. We can set $\ell(y) = 1$, so that $\ell(ay) = a$ by homogeneity. Then $\ell(y) \le p(y)$ on the subspace spanned by y, since $\ell(y) = 1 \le p(y)$ (because $y \notin K$). Now extend this linear functional ℓ to everything and it will satisfy the same bound by Hahn-Banach. □

The geometric Hahn-Banach theorem is basically an upgraded version of this:

Theorem 58 (Geometric Hahn-Banach)

Let K_1, K_2 be two disjoint convex sets. Then there is some linear functional ℓ such that $\ell(y_1) \le c \le \ell(y_2)$ for all $y_1 \in K_1, y_2 \in K_2$.

Proof. Consider the set $\tilde{K} = K_1 - K_2$. This set is convex and does not contain 0. Choose ℓ so that $\ell(x) \le 0$ for all $\tilde{K}$ by the previous result and such that $\ell(0) = 0$. Then for any x that we can write in the form $y_1 - y_2$, we have $\ell(y_1) \le \ell(y_2)$ by linearity, so $\sup_{y_1} \ell(y_1) \le \inf_{y_2} \ell(y_2)$ and thus ℓ does separate the two convex sets. □

This next corollary of Hahn-Banach also tells us something useful about dual spaces:

Proposition 59

Let X be a Banach space. Then for all $x \in X$, there is some linear functional $\ell \in X^*$ with $||\ell||_{X^*} = 1$ and setting $\ell(x) = ||x||$.

In the case where X is a Hilbert space, this is easy because we can choose ℓ to be the inner product with the unit-length vector in the x-direction.

Proof. Define the linear functional by letting $\ell(x) = ||x||$ and then extend ℓ with that control everywhere. $\square$

The point is that there's always a linear functional which "sees things" in the direction of x. And we'll now see this in a context we'll be using going forward:

Definition 60

For a Banach space X, the **weak topology** is the smallest topology such that all $\ell \in X^*$ are continuous.

Basically, we can define the dual space (bounded linear functionals on X), and we want to only keep the open sets that make each of those functionals continuous. For a fixed ℓ, this means we want the preimage $\ell^{-1}(U)$ for any open set in U, so the slabs $\ell^{-1}(-\varepsilon, \varepsilon)$ will be open sets in the weak topology. Then the basic open sets that are neighborhoods of the origin in this topology are of the form

$$\left\{ \bigcap_{j=1}^{N} \ell_j^{-1}((-\varepsilon_j, \varepsilon_j)) \right\}.$$

The reason this is a weak topology is that each of these open sets has an "infinite dimension not caught by functionals," so a lot of stuff can escape to infinity. And in particular it's not metrizable because it's so weak. But it is Hausdorff because of Hahn-Banach (we can always separate points), and it's nice because sequences can often converge and we can still extract some information about this.

We'll do one final general construction before jumping back into information with Fourier transforms and various function spaces:

Definition 61

For any Banach space X, we have a dual space X^*, and then we get the dual of that $(X^*)^* = X^{**}$, which we call the **double dual**.

An element $\mu \in X^{**}$ of the dual takes in linear functionals $\ell \in X^*$ and yields real numbers. And for any $x \in X$, we get an evaluation map $\mu_x : X^* \to \mathbb{R}$ which sends ℓ to $\ell(x)$. We must check that μ_x is actually a bounded linear operator with $||\mu_x||_{X^{**}} = ||x||_X$, and this map $x \mapsto \mu_x$ is injective. So there is an isometric embedding $X \to X^{**}$; if this is a bijection then we say that X is **reflexive**.

Example 62

Hilbert spaces are always reflexive, and L^p is reflexive for any $1 < p < \infty$ because we go from L^p to L^q to L^p. A similar thing happens for the ℓ_p spaces. On the other hand, $L^1, L^\infty, C^{0,\alpha}$ and so on are far from being reflexive.

It'll turn out that this will be a topology even weaker than the weak topology, and reflexivity tells us something about the unit ball too. But that's also left for later in the class – for now we're going to enrich our class of examples and start talking about Fourier analysis, looking at the Fourier transform, Schwartz space, distributions, and Sobolev spaces. (References can be found from the Stanford library and are posted on Canvas.)

Definition 63

Let $u \in C_0^\infty(\mathbb{R}^n)$ (compactly supported smooth functions). The **Fourier transform** of u is the function

$$\hat{u}(\xi) = (2\pi)^{-n} \int e^{-ix\xi} u(x) dx.$$

The Fourier transform is basically meant to capture how oscillatory u is – oscillatory integrals are ubiquitous in certain areas of math, and we should think that if u is smooth, then the regular oscillations should make the integral quite well-behaved for large ξ because of cancellations, and as $\xi \to \infty$ we have $\hat{u}(\xi) \to 0$. But in fact we can say something more specific – local singularities exactly correspond to failure to decay at infinity.

The problem is that C_0^∞ has a slightly unpleasant topology: we say that u_j converge to u if **(1)** there is some compact set K such that all u_js are supported on K and **(2)** we have uniform convergence of kth derivatives for any k. The reason that we need this is to avoid the "traveling bump" which keeps moving to the right and should not converge to zero. So we're going to take a slightly larger space:

Definition 64

The **Schwartz space** on $\mathbb{R}^n$ is the set of functions

$$\{u \text{ smooth} : \left|x^\alpha \partial_x^\beta u\right| \le C_{\alpha\beta} \forall \alpha, \beta \text{ multi-indices}\}$$

In other words, we don't require that we're compactly suppored, but we do require that any derivative of u decays faster than any polynomial. And we in fact get a Frechet space out of the countable family of seminorms: often this is all phrased slightly differently as the space of rapidly decreasing functions

$$\mathcal{S} = \left\{ u : \sum_{|\beta| \le \ell} (1 + |x|)^k |\partial_x^\beta u| \le C_{k,\ell} \right\}.$$

(A good example to keep in mind is $e^{-|x|^2}$.) What we care about is the Fourier transform as a map from $\mathcal{S}$ to $\mathcal{S}$, and the point is that we can do basically anything to $\mathcal{S}$ and it'll stay Schwartz (multiplying by polynomials or taking derivatives).

The Fourier transform has a few formal properties: for example, integration by parts tells us that

$$\int e^{-ix\cdot\xi} \partial_{x_j} u dx = i\xi_j \int e^{-ix\xi} u(x) dx$$

(legitimate calculation here because of fast decay) and thus $\frac{1}{i}\partial_{x_j}$ **can be interchanged with** ξ_j. Similarly, we have

$$\int e^{-ix\xi} x_j u(x) dx = i\partial_{\xi_j} \int e^{-ix\xi} u(x) dx,$$

and thus x_j is exchanged with $i\partial_{\xi_j}$. So indeed growth at infinity (how big of a polynomial we can multiply by) is related to differentiability. And the corollary of those computations is the following:

Corollary 65

The Fourier transform $\mathcal{F}$ maps $\mathcal{S} \to \mathcal{S}$; that is, if $u \in \mathcal{S}$, then $\hat{u} \in \mathcal{S}$.

Proof. We know that up to some factors of i,

$$\xi^\alpha \partial_\xi^\beta \hat{u} = \mathcal{F}(i^n \partial_x^\alpha x^\beta u)$$

and because the $\partial_x^\alpha x^\beta u$ is still rapidly decreasing (for example faster than $\frac{1}{1+|x|^{2n+1}}$) since $u \in \mathcal{S}$, the left-hand side will be bounded. $\square$

Next time we'll show that this mapping is actually an isomorphism, so in particular it is a continuous invertible mapping. That'll allow us to look at $\mathcal{S}^*$ and then look at a nice subspace of Sobolev spaces.

6 January 26, 2023

Last time, we discussed the (continuous) Fourier transform, which sends a function $u(x)$ to $\hat{u}(\xi) = (2\pi)^{-n} \int u(x)e^{-ix\xi}dx$. We discussed the class of Schwartz functions and how smoothness relates to decay – in fact lack of decay will indicate irregularity, and that's because the Fourier transform and the original function are "basically the same." Today, we'll prove the Fourier inversion theorem in a slightly different way from last quarter, doing the "correct proof."

Recall that the Schwartz class is defined as the infinitely differentiable functions which decay faster than any polynomial:

$$\mathcal{S} = \{u \in \mathbb{R}^n : |(1+|x|)^k \partial_x^\alpha| \le C_{k,\alpha} \forall k, \alpha\}.$$

We could define basically exactly the same spaces where we require that $(1+|x|)^k \partial_x^\alpha u \in L^2$ instead of just being bounded – we often just need shift k and α around. And we saw that because multiplying by x_j is the same as $i\partial_{\xi_j}$, and differentiating also corresponds to multiplication, the Fourier transform does send $\mathcal{S}$ to $\mathcal{S}$.

Theorem 66

Define the "inverse Fourier transform" on $\mathcal{S}$ via $\mathcal{F}^{-1}(\hat{u})(x) = \int e^{ix\xi}\hat{u}(\xi)d\xi$. Then we actually have $\mathcal{F}^{-1}\hat{u} = u$.

Proof. We'll make life easier by not worrying about factors of 2π here. We must show that up to some constant factor, taking the Fourier transform and then the inverse Fourier transform gives us back the original function. Indeed,

$$\int\int e^{ix\xi}u(y)e^{-iy\xi}dyd\xi = \int\int e^{i(x-y)\xi}u(y)dyd\xi.$$

The problem is that this is not a decaying integral in ξ until we've done the y-integral, so we'll introduce decay in the ξ direction. Our goal is then to compute

$$\lim_{\varepsilon\to 0}\int\int e^{i(x-y)\xi - \varepsilon|x|^2}u(y)dyd\xi,$$

and this is a nicely convergent integral so we can swap the order of integration and do other manipulations. But now by completing the square,

$$i(x-y)\xi - \varepsilon|x|^2 = \left(\frac{1}{2\sqrt{\varepsilon}}(x-y) + i\sqrt{\varepsilon}\xi\right)\cdot\left(\frac{1}{2\sqrt{\varepsilon}}(x-y) + i\sqrt{\varepsilon}\xi\right) - \frac{1}{4\varepsilon^2}|x-y|^2$$

(remembering this is a vector identity) and substituting that in gives us

$$= \int\int e^{\left(\frac{1}{2\sqrt{\varepsilon}}(x-y)+i\sqrt{\varepsilon}\xi\right)\cdot\left(\frac{1}{2\sqrt{\varepsilon}}(x-y)+i\sqrt{\varepsilon}\xi\right) - \frac{1}{4\varepsilon^2}|x-y|^2}u(y)dyd\xi.$$

Making a change-of-variable $z = \frac{x-y}{\sqrt{\varepsilon}}$, $\eta = \sqrt{\varepsilon}\xi$, the Jacobian factor ends up being exactly 1 here because there's n copies of z and n copies of η:

$$= \int\int e^{\left(\frac{z}{2}+i\eta\right)\cdot\left(\frac{z}{2}+i\eta\right) - \frac{|z|^2}{4}}u(x-\sqrt{\varepsilon}z)dzd\eta.$$

We now shift in the complex plane, replacing η_j by $\eta_j + \frac{iz_j}{2}$. (The point is that we integrate around a big rectangular box with vertices at $-R, R, -R+\frac{iz_j}{2}, R+\frac{iz_j}{2}$, and the contributions from the vertical parts are small as $R \to \infty$ but the total integral is 0 by Cauchy's integral theorem, so the integral along the shift doesn't change.) We then end up with

$$= \int\int e^{i\eta\cdot i\eta - \frac{|z|^2}{4}}u(x-\sqrt{\varepsilon}z)d\eta dz = \int\int e^{-|\eta|^2 - \frac{|z|^2}{4}}u(x-\sqrt{\varepsilon}z).$$

The η-integral is now going to give us some factors of 2π, and by the dominated convergence theorem we can replace

$u(x-\sqrt{\varepsilon}z)$ with $u(z)$ and basically $e^{-|z|^2/4}$ becomes a constant times approximation of identity. (This will make more sense when we talk about distributions next week.) Thus overall we do get a universal constant $cu(x)$, (mostly) proving Fourier inversion. $\square$

Remark 67. *Depending on convention, the factors of 2π may appear in other places: some choices use $\frac{1}{(2\pi)^{n/2}}$ for both transforms or $\frac{1}{(2\pi)^n}$ for the inverse transform, and others put the 2πs in the exponent. So we should be very careful about that.*

So we now know that $\mathcal{F}:\mathcal{S}\to\mathcal{S}$ is an isomorphism, and moreover $\mathcal{F}\circ\mathcal{F}$ isn't quite the identity but it's very close – we get $(2\pi)^{-n}R$, where R is the reflection operator sending $u(x)$ to $u(-x)$. In particular, that means $\mathcal{F}^4$ is the identity up to powers of 2π. We should think of the Fourier transform as a unitary operator which is basically a "big rotation," and its eigenvalues can only be the fourth roots of unity. (And we can actually find all the eigenfunctions.)

As mentioned last time, there is a dual space

$$\mathcal{S}' = \{u : \mathcal{S}\to\mathbb{C} \text{ continuous linear functional}\}.$$

Since $\mathcal{S}$ is a Frechet space, being a continuous linear functional means that there **exists some** k,ℓ such that for all $\phi\in\mathcal{S}$, $\langle u,\phi\rangle$ (which we should think of as "the integral of u with ϕ") is bounded as

$$|\langle u,\phi\rangle| \le C\sum_{|\alpha|\le\ell}\sup(1+|x|)^k|\partial_x^\alpha\phi|.$$

In other words, we just need to be continuous with respect to one of the seminorms – we don't require ϕ to vanish rapidly or decay beyond just the first k powers of x and the ℓ derivatives.

Example 68

For any function $u\in L^1$, $\langle u,\phi\rangle=\int u\phi$ is a valid linear functional, since $|\int u\phi|\le ||u||\sup|\phi|$ is a bound in terms of the bound with $k=\ell=0$. And if u is a polynomial, then $\langle u,\phi\rangle=\int u\phi$ still makes sense if we just bound with a higher value of k.

But we can also do more complicated things like

$$\langle u,\phi\rangle = \sum_{|\alpha|\le N}\int (P_\alpha(x)\partial_x^\alpha\phi)dx$$

for polynomials P_α, or (letting u be the "delta function")

$$\langle u,\phi\rangle = \phi(0).$$

These are all valid examples of things in $\mathcal{S}'$, which is called the space of **tempered distributions** – what's not allowed is anything growing faster than polynomial. And it's nice that $L^p\subset\mathcal{S}'$ for any p.

Definition 69

Let $u\in\mathcal{S}'$. Then the Fourier transform $\hat{u}\in\mathcal{S}'$ is defined in the following indirect way: on any element $\phi\in\mathcal{S}$, we have $\langle\hat{u},\phi\rangle=\langle u,\hat{\phi}\rangle$.

Once we've done this, notice that this in particular tells us how to take the Fourier transform of L^p functions for any p, though the answer may not necessarily be nice. To understand why this is a valid definition, notice that for any

$u, v \in \mathcal{S}$, we have

$$\int \hat{u}(\xi)v(\xi)d\xi = \int e^{-ix\xi}u(x)v(\xi)d\xi = \int u(x)\int e^{-ix\xi}v(\xi)d\xi = \int u(x)\hat{v}(x),$$

so swapping the Fourier transform does make sense. And it will turn out that $\hat{S}$ is dense in $\hat{S}'$ for some appropriate topology, so we can approximate (in the appropriate very weak topology) any of these potentially-wildly-behaved linear functionals by Schwartz functions.

Proposition 70 (Plancherel)
For any $u \in \mathcal{S}$, we have $\int |u|^2 = \int |\hat{u}|^2$.

This is basically a specialization of the previous calculation: notice that if we define $v \in \mathcal{S}$ so that $\overline{u} = \hat{v}$, then

$$\int |u|^2 = \int u\overline{u} = \int u\hat{v} = \int \hat{u}v.$$

But now with this definition, we can manipulate conjugates and exponentials a bit and see that $v = \overline{\hat{u}}$, so this last expression is indeed $\int \hat{u}\overline{\hat{u}} = \int |\hat{u}|^2$.

The idea now is the following old-fashioned idea (before distributions were studied): for any $u \in L^2$, we can define a cutoff function $u_R = \chi_R u$, where χ_R is just 1 in $B_R(0)$ and 0 elsewhere. Then u_R is in L^1, so by classical integration we can define $\hat{u}_R$ for any R. If we now look at a sequence of radii $R_j \to \infty$, then by Plancherel we have

$$||u_{R_j} - u_{R_k}||_2 = ||\hat{u}_{R_j} - \hat{u}_{R_k}||_2,$$

and now the left-hand side can be made arbitrarily small if j, k are big enough, and thus the $\hat{u}_R$s are Cauchy as well – thus $\hat{u} = \lim \hat{u}_n$ in L^2. So we can basically interpret the Fourier transform concretely in terms of the L^1 integral or abstractly in terms of the $\mathcal{S}'$ definition, and this is a case where we can make a closed graph theorem argument – $\mathcal{F}$ happens to map L^2 to L^2 even though it was originally defined on $\mathcal{S}' \to \mathcal{S}'$, and so we can show it's a bounded linear operator on L^2 as well.

As a sidenote, we can also do analogous arguments for Fourier **series**: for any function $u(\theta) \in S^1$, we may define $\hat{u}(n) = \frac{1}{2\pi}\int e^{-in\theta}u(\theta)d\theta$, and similarly we can do the same thing on $S^1 \times \cdots \times S^1 = T^k$ and get a multidimensional Fourier series

$$\hat{u}_{n_1,\cdots,n_k} = \int e^{-i(n_1\theta_1+\cdots+n_k\theta_k)}u(\theta)d\theta(2\pi)^{-n}.$$

We then get an analogous version of Plancherel's theorem $\int |u|^2 = \sum |\hat{u}_n|^2$, which is basically saying that $\mathcal{F}$ takes $L^2(S^1)$ to ℓ_2. We can also check that it takes $L^1(S^1) \to \ell_\infty$ (in fact we land in c_0 by Riemann-Lebesgue), and there is an inverse transform which takes a sequence to its corresponding power series $\sum a_n e^{in\theta}$ which sends $\ell_2 \to L^2$.

One basic question we may want to ask at this point is to ask if L^p maps to a particular space. We know that $1 \in L^\infty$, and we will justify later that $\mathcal{F}(1)$ is in fact the delta function δ_0. So it is very possible for an L^p function to turn into a generalized distribution – thus this question is really only answerable for $1 \le p \le 2$. It turns out that $\mathcal{F}$ will map from L^p to L^q (where q is the usual dual exponent) – if we try to prove this directly, it is very challenging, but we can instead use an abstract method to get this done (and we can do the same thing for Fourier series). And there are other questions like this too: for example, for a periodic function $f \in L^p(S^1)$, we have Fourier coefficients $\hat{f}_n$, and we may ask if we can say anything about the convergence of the partial sums $(S_N f)(\theta) = \sum_{-N}^{N} \hat{f}_n e^{in\theta}$ in L^p. This is easy when $p = 2$, but again boundedness of operators will help us do it for the other values of p.

The tool we'll use is called the Riesz-Thorin interpolation theorem, and we'll need some preparation for it:

Lemma 71

If $f \in L^{p_0} \cap L^\infty$, then $f \in L^{p_1}$ for all $p_1 > p_0$.

Proof. Let $M = \sup |f|$. Then $\int |f|^{p_1} \le M^{p_1 - p_0} \int |f|^{p_0}$ is indeed finite. ☐

Lemma 72

Now suppose $f \in L^{p_1}$, and suppose the support K of f has finite measure. Then $f \in L^{p_0}$ for all $1 \le p_0 \le p_1$.

Proof. Apply Holder's inequality to $|f|^{p_1}$ and 1 to show that $\int |f|^{p_0}$ is bounded by $\left(\int |f|^{p_1}\right)^{p_0/p_1} \left(\int_K 1\right)^{\text{some fixed power}}$. ☐

Lemma 73

Suppose $p_0 < p < p_1$. Then if $f \in L^{p_0} \cap L^{p_1}$, then $f \in L^p$.

Proof. Define $E = \{x : |f| \le 1\}$. We claim that $|\mathbb{R}^n \setminus E| < \infty$; indeed, $\int_{\mathbb{R}^n \setminus E} 1 \le \int_{\mathbb{R}^2 \setminus E} |f|^{p_0} < \infty$ by assumption. So now we can write $f = f\chi_E + f(1 - \chi_E)$, the latter of which is in $L^{p_0} \cap L^\infty$ and the former of which is in $L^{p_1} \cap$ (support has finite measure). Thus combining the two previous lemmas yields the result (since both terms lie in L^p). ☐

The idea is that on an infinite domain, a function $f \in L^p$ can be decomposed as $f_0 + f_1$ with $f_0 \in L^{p_0}$ and $f_1 \in L^{p_1}$, where f_0 is going to be of the form $f\chi_E$ and f_1 is $f(1 - \chi_E)$. (As a sidenote, we can prove that the function $p \mapsto ||f||_p^p$ is convex in p, which is sometimes useful.)

Theorem 74 (Riesz-Thorin)

Let $(M, \mu), (N, \nu)$ be two measure spaces (we can think of these as $\mathbb{R}^n$). Suppose we have numbers $(p_0, q_0), (p_1, q_1)$, where $1 \le p_j, q_j \le \infty$ (not necessarily Holder conjugates), and suppose there are operators $A_0 : L^{p_0}(M) \to L^{q_0}(N)$ and $A_1 : L^{p_1}(M) \to L^{q_1}(N)$ with finite operator norms k_0 and k_1, such that $A_0 = A_1$ on the intersection $L^{p_0} \cap L^{p_1}$. (In other words, an operator extends in two different ways to A_0 and A_1.) Then A is actually a bounded operator for all L^p with p between p_0 and p_1. Specifically, for any $0 \le t \le 1$ we may define

$$\frac{1}{p_t} = \frac{1-t}{p_0} + \frac{t}{p_1}, \quad \frac{1}{q_t} = \frac{1-t}{q_0} + \frac{t}{q_1}.$$

Then there is some $A_t : L^{p_t} \to L^{q_t}$ with $||A_t|| \le k_0^{1-t} k_1^t$ (and A_ts agreeing).

This is a classical interpolation theorem – there are many more abstract results that just involve abstract Banach spaces, but there is some more setup required to explain how to interpolate those spaces. (The L^p spaces are indeed doing that though – that's what the lemmas were showing.) And recall that the reason we're doing it here is that we know the Fourier transform is bounded from L^2 to L^2 and from L^1 to L^∞, so this result tells us that it also maps from L^p to L^q (where $\frac{1}{p} + \frac{1}{q} = 1$). And something useful can be said for Fourier series too: the exact same application of interpolation tells us that

$$\left(\sum |\hat{f}(n)|^q\right)^{1/q} \le ||f||_{L^p}$$

for any $1 < p < 2$, and the inverse Fourier series says that

$$\left(\int_{S^1} \left|\sum a_n e^{in\theta}\right|^q d\theta\right)^{1/q} \le C\left(\sum |a_n|^p\right)^{1/p}.$$

So these are all nice results that are hard to prove directly but very useful. This proof uses complex analysis, and we'll need two main facts:

Lemma 75
Suppose $\Omega \subset \mathbb{C}$ is a domain and u is holomorphic in Ω. Then $|u(z)|$ has no local interior maxima or minima unless u is constant.

What's important about holomorphic functions is that there are power series at any point: we can write power series expressions $u(z) = a_0 + a_1(z - z_0) + a_2(z - z_0)^2 + \cdots$ (let's take $z_0 = 0$ for simplicity), and if the first nonconstant term in the power series is $a_n z^n$, we now have $|u|^2 - |a_0|^2 = 2|a_0 a_n| r^n \cos(n\theta - \theta_0) + O(r^{n+1})$, but the cosine oscillates so we can't achieve a maximum at any particular point.

Theorem 76 (Three lines theorem)
Let S be the rectangular strip where $\mathrm{Re}(z) \in [0, 1]$. Suppose $f : S \to \mathbb{C}$ is holomorphic, and suppose it is bounded by m_0 on $\mathrm{Re}(z) = 0$ and by m_1 on $\mathrm{Re}(z) = 1$. Then if $|f(z)| \le Ce^{c_2|z|^{2-\delta}}$ for some $\delta > 0$ (meaning that we're growing slower than e^{z^2}), then f is bounded on $\mathrm{Re}(z) = t$ by $m_0^{1-t} m_1^t$.

Proof. This is similar to the previous result, except S is not compact so we can't just use the same proof. But in fact we get a more numerical result here, getting an explicit bound in terms of the bounds on the boundary. First define the function

$$F(z) = \frac{f(z)}{m_0^{1-z} m_1^z},$$

so that we're bounded by 1 on $\mathrm{Re}(z) = 0$ and $\mathrm{Re}(z) = 1$ and we wish to show that we're bounded by 1 in the interior as well. First assume $|F| \to 0$ as $y \to \pm\infty$. Then we can make F at most $\frac{1}{2}$ in the subregion of S where $\mathrm{Im}(z) < C$, and then the maximum modulus principle immediately gives us the result.

Otherwise, consider the function (for $z = x + iy$)

$$F_n(z) = F(z)e^{(z^2-1)/n} = F(z)e^{\frac{x^2-y^2-1}{n}} e^{2ixy/n}$$

Notice that $x^2 - 1$ is negative and goes to zero as $n \to \infty$, and the $e^{-y^2/n}$ beats out the growth of $|f(z)|$ as $y \to \infty$. So for each n there is some large enough C so that the "bounding rectangle" argument above works. Then taking $n \to \infty$ yields the result. $\square$

To see how Riesz-Thorin is actually set up in this framework, instead of thinking of having $A : L^{p_t} \to L^{q_t}$, we instead will imagine that we have $L^{p(z)} \to L^{q(z)}$, and to check that an operator is bounded we need to check that Au is bounded, meaning that for any $g \in L^{q'}$ we have $\int (Au)g$ to be finite for all u and all g. And now we can think about "holomorphic functions" by first looking at simple functions (with $a_j, b_k \ge 0$)

$$u = \sum a_j e^{i\theta_j} \chi_{A_j}, \quad g = \sum b_k e^{i\theta_k} \chi_{B_k},$$

and what we'll do is create a holomorphic family of functions $u(z), g(z)$ by varying those coefficients a_j, b_k. We'll think more about this next time!

7 January 31, 2023

We'll finish the Riesz-Thorin theorem's proof today – the main idea relies on the three lines theorem from complex analysis, and the rigidity there comes up in lots of places. Recall that we're trying to prove that if we have a bounded operators $A_0 : L^{p_0} \to L^{q_0}$ and $A_1 : L^{p_1} \to L^{q_1}$ which agree on the intersection, then we get $A_t : L^{p_t} \to L^{q_t}$ for all $t \in [0,1]$ (where $\frac{1}{p_t} = \frac{1-t}{p_0} + \frac{t}{p_1}$ and $\frac{1}{q_t} = \frac{1-t}{q_0} + \frac{t}{q_1}$) with A_t agreeing with A_0 and A_1 on the corresponding restrictions L^{p_0}, L^{p_1}. Furthermore, we have the explicit bound on the operator norm $||A_t|| \le ||A_0||^{1-t}||A_1||^t$. (And remember that any function in an intermediate space can be written as a sum of a part in L^{p_0} and a part in L^{p_1}.) Here's the proof:

Proof. First of all, we know how we must define A_t: for any $f \in L^{p_t}$, we write $f = f_0 + f_1$ with $f_0 \in L^{p_0}$ and $f_1 \in L^{p_1}$, so then we must set $A_t f = A_0 f_0 + A_1 f_1$. We could probably use the closed graph theorem to show that this is bounded, but the size of the norm is the key fact.

Since the norm of an operator $A : L^p \to L^q$ is the best constant such that $||Au||_q \le C||u||_p$, this is the same as saying that

$$||A|| = \sup_{v \in L^{q'}, ||u||_p = 1, ||v||_{q'} = 1} |\langle Au, v\rangle|.$$

So now we can define, for any **complex** z,

$$\frac{1}{p(z)} = \frac{1-z}{p_0} + \frac{z}{p_1}, \quad \frac{1}{q'(z)} = \frac{1-z}{q_0'} + \frac{z}{q_1'},$$

where the primes denote dual exponents as usual. It suffices to study $\langle Au, v\rangle$ for **simple** functions of norm 1, and for u and v we will use (respectively)

$$f = \sum a_j e^{i\alpha_j}\chi_{A_j}, \quad g = \sum b_k e^{i\beta_k}\chi_{B_k},$$

where $a_j, b_k \in \mathbb{R}^+$ and the A_js all have disjoint support and so do the B_ks. Then define

$$\phi(\cdot, z) = \sum a_j^{p/p(z)} G_j, \quad \psi(\cdot z) = \sum b_j^{q'/q(z)} H_k,$$

where $G_j = e^{i\alpha_j}\chi_{A_j}$ and $H_k = e^{i\beta_k}\chi_{B_k}$. (The point is that ϕ agrees with f at zero and ψ agrrees with g at 1.) We can check that $\phi(\cdot, z) \in L^{p_j}$ and $\psi(\cdot, z) \in L^{q_j'}$ when $z = iy$ or $z = 1 + iy$ because $f \in L^p$ and $g \in L^{q'}$ from the start by assumption. So now we can define

$$F(z) = \int A\phi(\cdot, z)\psi(\cdot, z) d\nu;$$

this is a holomorphic function of z because we just have finite sums of certain phase factors times characteristic functions. Furthermore, because the A_js have disjoint support,

$$||\phi(\cdot, iy)||_{p_0} = \int \sum |a_j^{p/p(iy))}|^{p_0},$$

and the idea is that when $z = iy$, then any imaginary part of z becomes an imaginary exponent so we just get

$$= \int \sum (|a_j|^{p/p_0})^{p_0} = \sum |a_j|^p = 1.$$

Similarly, $||\phi(\cdot, 1 + iy)||_{p_1}^{p_1} = 1$. Thus by ordinary Holder, we have

$$|F(iy)| \le ||A\phi(iy)||_{q_0} ||\psi(iy)||_{q_0'} \le k_0,$$

where recall that k_0 is the norm of $A\phi$ on the left strip. Similarly, we have $|F(1+iy)| \le ||A\phi(\cdot, 1+iy)||_{q_1} ||\psi(1+iy)||_{q_1'} \le k_1$. Thus $|F(t)| \le k_0^{1-t}k_1^t$ by the three-lines lemma, and that's what we wanted to show. $\square$

The point is that holomorphic interpolation is difficult in general, but for simple functions we can write down an extension explicitly. And this is the typical use of Riesz-Thorin – we look at a class of functions where extending makes sense. And the fact that we already get boundedness properties of the Fourier transform (specifically, that it is a bounded operator $L^p \to L^{p'}$) shows how powerful this technique can be, and there are lots of situations where studying the Fourier transform in various L^p spaces is useful (particularly with PDEs). We'll now see another application which will be useful as we move into a closer study of the Fourier transform:

Proposition 77

For any functions $f, g \in \mathbb{R}^n$, we have the convolution $(f * g)(x) = \int f(x-y)g(y)dy = \int f(y)g(x-y)dy$. Then if $\frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r}$, then

$$||f * g||_r \le ||f||_p ||g||_q,$$

Proof. First fix $g \in L^1$, and consider the operator $C_g(f) = f * g$. Then we wish to show that $||f * g||_1 \le ||f||_1 ||g||_1$, but this is true because

$$\int \left| \int f(x_y)g(y)dy \right| dx \le \int \int |f(x-y)||g(y)|dydx$$

and now we can split apart the right-hand side into two integrals once we change variables $x - y \mapsto z$ and we split apart exactly into $||f||_1$ and $||g||_1$. Thus C_g maps $L^1 \to L^1$. But also C_g maps $L^\infty \to L^\infty$, since we can pull out $\sup |f|$ from the integral and that gives us the L^1 norm of g. So by Riesz-Thorin, this means C_g maps $L^p \to L^p$ for all p.

For the general case, fix any $f \in L^p$ and consider $C_f(g) = f * g$. There are two cases to first sort out. First of all, C_f takes $L^1 \to L^p$ from the previous step (but inverting the roles of f and g). But also C_f takes $L^{p'} \to L^\infty$, since by Holder's inequality

$$\sup_x \left| \int f(x-y)g(y)dy \right| \le ||f||_p ||g||_{p'}.$$

So C_f interpolates domain space from L^1 to $L^{p'}$, and we want to show that C_f takes $L^q \to L^r$. So we want

$$\frac{1}{q} = \frac{1-t}{1} + \frac{t}{q'}, \quad \frac{1}{r} = \frac{1-t}{p} + \frac{t}{\infty}.$$

Indeed, q does lie between 1 and p', since we want $\frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r}$ and $\frac{1}{p} + \frac{1}{p'} = 1$ (so having q be larger than p' would not work). And in fact we must have

$$\frac{1}{q} = 1 + t\left(\frac{1}{p'} - 1\right) = 1 - \frac{t}{p},$$

and thus

$$\frac{1}{r} = \frac{1}{p} - \frac{t}{p} = \frac{1}{p} + \frac{1}{q} - 1.$$

But this is exactly what we wanted – we have a bound on the r-norm of $||f * g||$ in terms of the p-norm of f and q-norm of g. $\square$

We can now return to the world of Fourier series:

Definition 78

Let $f = \sum_{-\infty}^{\infty} a_n e^{in\theta}$ be a Fourier series. Define the operator H, called the **Hilbert transform**, via

$$Hf = \sum_{n=1}^{\infty} ia_n e^{in\theta} - \sum_{n=-\infty}^{-1} ia_n e^{in\theta}.$$

So we basically multiply the positive-Fourier parts by i and the negative-Fourier parts by $-i$. The motivation for this is the following: we want to know which functions $u(\theta)$ are **boundary values of holomorphic functions** in D. Being holomorphic in D just means we have a radius of convergence at least 1, meaning that we want a function $U(z) = \sum_{n=0}^{\infty} b_n z^n$ converging absolutely on $|z| < 1$, and then setting $z = e^{i\theta}$ yields $\sum_{n=0}^{\infty} b_N e^{in\theta}$. So we do sometimes care about "only having the positive-Fourier parts." (There is a notion of a distributional boundary value, but we won't focus too much on that right now.)

The point now is that if $u(\theta) = \sum c_n e^{in\theta}$ is **real-valued**, then it is equal to its own conjugate, so the condition we need is

$$c_n e^{in\theta} + c_{-n} e^{-in\theta} = \overline{c_n} e^{-in\theta} + \overline{c_{-n}} e^{in\theta} \implies \overline{c_n} = c_{-n}$$

for all n. (The reason the different ns don't mix beyond this is basically because of orthogonality.) So now we have a relation to a classic problem: if we're given a real-valued function u on the boundary, we want to find another real-valued function v so that $(u + iv)(\theta)$ is a boundary value of some holomorphic function. The answer turns out to be $v = Hu$ – the Hilbert transform finds the harmonic extension, takes its harmonic conjugate, and then takes the boundary value. And all of this also has to do with pointwise convergence of Fourier series: we can check that the partial sums satisfy

$$\boxed{S_N f(\theta) = \sum_{n=-N}^{N} a_n e^{in\theta} = \frac{1}{2i}\left(e^{iN\theta} H(e^{-iN\theta} f) - e^{-iN\theta} H(e^{iN\theta} f)\right) + \frac{1}{2}\hat{f}(N) e^{iN\theta} + \frac{1}{2}\hat{f}(-N) e^{-iN\theta}}.$$

(This basically just comes from assembling geometric series.) So now we want to generalize an easy result for L^2 and ask whether we have $||S_N f - f||_p \to 0$, so that the partial sums converge in L^p. The answer turns out to lie in the formula above – we want to show that the Hilbert transform is bounded in L^p. It's true for $p = 2$ because $||f||_{L^2}^2 = \sum_{-\infty}^{\infty} |a_n|^2$ and $||S_N f||_{L^2}^2 = \sum_{n=-N}^{N} |a_n|^2$, and the tail sums coming from the difference must go to zero. And also we can check that

$$|\hat{f}(\pm N)| = |\int f(\theta) e^{\mp iN\theta} \theta| \le ||f||_p$$

by Holder's inequality and the fact that we're on the circle (which has finite measure); in fact they go to zero. All of this gives us the necessary background for the big result:

Theorem 79

The Hilbert transform is bounded as an operator $L^p \to L^p$ for all $1 < p < \infty$.

We'll just do the case $p \ge 2$ for now:

Part of the proof. As mentioned above, H maps $L^2 \to L^2$ by Hilbert space theory. On the other hand, we can see that H is not a bounded linear operator $L^1 \to L^1$, meaning $\frac{||Hf||_1}{||f||_1}$ is not bounded. Indeed, if it were, then we could bound $||S_N f||_1 \le C ||f||_1$ for all f, since the right-hand side of the boxed expression has each term controlled by $||f||_1$. But then by uniform boundedness, that would mean that the norm of the operators S_N are uniformly bounded, which is not true.

Our trick instead is now to show that H is bounded $L^{2k} \to L^{2k}$ for all integers $k \ge 2$ by a direct computation, so in particular H is bounded from $L^p \to L^p$ for all $p \ge 2$ by interpolation. Indeed, we can assume that f is real-valued (we can prove this separately for the real and complex parts), meaning that $\hat{f}(-n) = \overline{\hat{f}(n)}$. We can check that this means Hf is also real-valued, since Hf satisfies this same conjugation condition. So now consider the operator P defined via

$$Pf = \sum_{n \ge 1} \hat{f}(n) e^{in\theta}$$

(that is, projecting onto the positive Fourier coefficients), and we can notice that $f + iHf = \hat{f}(0) + 2Pf$. Assume that $\hat{f}(0) = 0$. Since P has no constant term, we have

$$0 = \int (2Pf)^{2k} = \sum_{j=0}^{2k} \binom{2k}{j} \int f^j (iHf)^{2k-j},$$

and taking the real parts, which basically makes the odd j terms go away because f is real-valued, we get

$$0 = \sum_{r=0}^{k} \int \binom{2k}{2r} f^{2r} (-1)^{k-r} (Hf)^{2k-2r}.$$

Thus moving the highest-power H term to one side, and then using the triangle inequality, we have

$$\int (Hf)^{2k} \le \sum_{r=1}^{k} c_r f^{2r} (Hf)^{2k-2r},$$

where we take all c_rs to be positive. By Holder's inequality with $p = \frac{k}{r}$ and $p' = \frac{k}{k-r}$ on each term, we then find that

$$\int (Hf)^{2k} \le \sum_{r=1}^{k} c_r \left(\int f^{2k} \right)^{r/k} \left(\int (Hf)^{2k} \right)^{1-r/k}$$

So now if we define $X = \frac{||Hf||_{2k}}{||f||_{2k}}$, dividing both sides of our equation above by $||f||^{2k}$ yields

$$X^{2k} \le \sum_{r=0}^{k-1} c_r X^{2r}.$$

But now the left-hand side grows faster than the right-hand side, so the only way this can happen is if X is bounded by some constant in terms of $c_1, \cdots, c_k$. Thus, we've proven that the Hilbert transform is bounded for all $2 \le p < \infty$ – we'll discuss the other case $1 < p < 2$ later. □

Theorem 80 (Riesz)

If $f \in L^p$ for some $1 < p < \infty$, then $||S_n f - f||_{L^p} \to 0$

Proof. This result is clear if f is some trigonometric polynomial $\sum_{n=-M}^{M} a_n e^{in\theta}$, since $S_N f$ is eventually just equal to f. The set of trigonometric polynomials is dense in any L^p, and we know that $\sup_N \frac{||S_N f||_p}{||f||_p}$ is finite, since again the Hilbert transform is uniformly bounded and we have the boxed expression for S_N in terms of the Hilbert transform and the $\hat{f}(N)$s, which are all properly controlled. Convergence on a dense set plus uniform bound on the norm everywhere means that Banach-Steinhaus (uniform boundedness) implies the desired result. □

We're now going to turn back to the "lighter" general topic of Fourier transforms and distributions. Recall that $\mathcal{F}$ was defined as an isomorphism from $\mathcal{S}$ (the space of smooth, rapidly decaying functions) to itself, and $\mathcal{S}$ has countably many seminorms that make it into a Frechet space. We then defined the dual space $\mathcal{S}'$ of continuous linear functionals $\mathcal{S} \to \mathbb{C}$ – the important thing here is that a linear functional v is continuous if there exists just a **single** k, ℓ such that we're bounded with respect to that seminorm:

$$|\langle u, v \rangle| \le C \sup_x \sup_{|\alpha| \le \ell} (1 + |x|)^k |D_x^\alpha u(x)|.$$

Remember that by definition the Fourier transform of the functional v is then defined via $\langle \hat{u}, v \rangle = \langle u, \hat{v} \rangle$, and we can check that this definition is invertible because the Fourier transform is invertible on $\mathcal{S}$. So we have a map $\mathcal{F} : \mathcal{S}' \to \mathcal{S}'$, and thus functions that are L^1_{loc} or L^p for any p can have a Fourier transform defined in a nice way.

So remember that L^p is mapped to $L^{p'}$ for $1 \le p \le 2$, but if we start with something in L^p for $p > 2$ we might end up with something messy. For example, we have the functional $1 \in \mathcal{S}'$, and then we want

$$\langle \hat{1}, u \rangle = \langle 1, \hat{u} \rangle = \int 1 \int e^{-ix\xi} u(x) dx d\xi,$$

and this is a problem because we will not have a convergent ξ-integral. But $v_\varepsilon = e^{-\varepsilon |x|^2}$ is Schwartz for any $\varepsilon > 0$, and we can check that $\hat{v_\varepsilon} = \varepsilon^{-n/2} e^{-|\xi|^2/(4\varepsilon)}$ is basically itself but rescaled a bit. Then v_ε converges to 1 in $\mathcal{S}'$ (that is, $\langle v_\varepsilon, u \rangle \to \langle v, u \rangle$ for all u – this is just dominated convergence theorem), so $\hat{v_\varepsilon}$ will also converge to $\hat{v}$ in the same weak sense. On the other hand, $\hat{v_\varepsilon}$ is extremely strongly peaked at zero and goes to zero rapidly for any nonzero ξ, and $\int \hat{v_\varepsilon}$ stays a constant as $\varepsilon \to 0$. So really $\hat{v_\varepsilon} = \hat{1}_\varepsilon$ is the **delta function** at the origin, which is defined to be the functional $\langle \delta_0, u \rangle = u(0)$. And indeed the Fourier inversion formula does tell us that $u(0) = \int \int e^{-ix\xi} u(x) dx d\xi$ (by taking the Fourier transform, then evaluating the inverse Fourier transform at 0).

We'll end with a definition of an object in Fourier analysis we'll be studying a lot in the coming lectures:

Definition 81

For any $s \in \mathbb{R}$, define the **L^2-based Sobolev spaces**

$$H^s(\mathbb{R}^n) = \{u \in \mathcal{S}' : \hat{u} \in L^1_{\text{loc}}, \int |\hat{u}(\xi)|^2 (1 + |\xi|^2)^s d\xi < \infty\}.$$

In other words, we have L^2 functions with an additional weighting factor but on the other side of the Fourier transform. And if s is very big, then that means $\hat{u}$ must decay pretty rapidly, but that's corresponding to smoothness of u. So the larger s is, the smoother our function is, and the more negative s is, the more singular our function can be. It turns out we'll basically get everything in $\mathcal{S}'$ as s varies.

8 February 2, 2023

We'll first quickly finish our discussion of the Hilbert transform – last time we showed that it's not bounded on L^1 but that it is bounded on L^p for $p \ge 2$. We'll do the case $1 < p < 2$ this time, and it's based on the following fact: up to a sign, the Hilbert transform is skew-adjoint:

$$\int (Hf)\overline{g} = -\int f\overline{Hg}.$$

This is just a matter of checking the definitions. And what this means is that if we have $f \in L^p$ and $g \in L^{p'}$, then

$$||Hf||_{L^p} = \sup_{||g||_{p'}=1} \left| \int (Hf)\overline{g} \right| = \left| \sup_{||g||_{p'}=1} f\overline{Hg} \right| \le ||f||_p ||g||_{p'}$$

by Cauchy-Schwarz and the boundedness of the Hilbert transform on the larger L^p spaces. So the real fact we're using is that boundedness corresponds to boundedness of the adjoint in the dual space.

Remark 82. *Notice that in all of our reasoning so far, H is bounded as a linear operator with norm at most 1 for all $1 < p < \infty$, but then it suddenly becomes ∞ at the two endpoints $p = 1, \infty$, so there's no sense of "continuity" of the operator norms.*

We'll now return to Sobolev spaces: recall that $H^s(\mathbb{R}^n)$ contains the elements of $\mathcal{S}'$ which are in L^1_{loc} and such that

$$||u||_s^2 = \int |\hat{u}(\xi)|^2(1+|\xi|^2)^s d\xi < \infty.$$

We may have heard about Sobolev spaces before with a different definition using weak derivatives, in which $H^0 = L^2$, and then we have

$$H^1 = \{u \in L^2 : Du \in L^2 \text{ in the weak sense}\}.$$

Here, note that we're going to use the convention that $D_j = \frac{1}{i}\partial_{x_j}$ to make behavior with the Fourier transform nicer, and in particular we have the same adjoint property

$$\int (D_j u)\overline{v} = \int u\overline{D_j v}$$

So what we are basically asking is for the map $v \mapsto \int u\overline{D_j v}$ to be a continuous linear functional on L^2. And finally, we can define the space H^1 to be the space of L^2 functions whose difference quotients exist almost everywhere and so that the derivative is in L^2, and similarly define $H^k = \{u : D^\alpha u \in L^2 \forall |\alpha| \le k\}$ for any $k \in \mathbb{N}$. It turns out all of these definitions will be equivalent, but we'll work with weak derivatives (distributional derivatives) primarily here. Indeed, notice that $u \in L^2$ if and only if $\hat{u} \in L^2$ by Plancherel, and $D_j u \in L^2$ if and only if $\xi_j \hat{u} \in L^2$. And furthermore, those conditions together give us exactly the equivalent condition to requiring $\int(1+|\xi|^2)|\hat{u}(\xi)|^2 < \infty$. More generally, notice that the Fourier transform of $D^\alpha u$ is $\xi^\alpha \hat{u}$, so being in H^k with the weak derivative definition is equivalent to

$$\int \sum_{|\alpha|\le k} |\xi^\alpha \hat{u}|^2 = \int \left(\sum_{|\alpha|\le k} |\xi^\alpha|^2\right) |\hat{u}(\xi)|^2 < \infty,$$

and now the parenthetical term is comparable (up to constants) with $(1+|\xi|^2)^k$. So indeed one of the two expressions being finite is the same as the other one being finite, and thus the two definitions of H^k for nonnegative integers agree. Then we can interpolate between these spaces (from H^k to H^{k+1}) and take duals (to get from H^s to H^{-s}) and that'll give us the definition of H^s for all real numbers s. And that interpolation involves finding some function $F(z,\cdot)$ which agrees on the corresponding axes: we know that $\hat{u} = |\hat{u}(\xi)|e^{i\alpha(\xi)}$, so we can just put some exponents on $|\hat{u}(\xi)|$ similarly to how we proved Riesz-Thorin so that $F(s,\cdot) = \hat{u}$ and so that we interpolate between H^k and H^{k+1} on the imaginary axis and on $\text{Re}(z) = 1$. The point is that many families of function spaces do interpolate in this way – the notation is that $[H^0, H^1]_s = H^s$.

Fact 83

It turns out that $[C^{0,1/2}, C^{1,1/2}]_{1/2}$ is not actually C^1 – it's called the Zygmund space.

The last piece of the puzzle here is that $H^{-s} = (H^s)^*$ comes from the fact that the weighted $(1+|\xi|^2)^s$ factors cancel each other out and thus we get a corresponding pairing:

$$\left|\int u\overline{v}\right| = \left|\int \hat{u}\overline{\hat{v}}\right| \le \int |\hat{u}|^2(1+|\xi|^2)^s \int |\hat{v}|^2(1+|\xi|^2)^{-s}.$$

So now if we try to relate all of the H^s spaces, we need to think about how smooth the functions are and how they decay at infinity. It will turn out that $\bigcap_{s\in\mathbb{R}} H^s \subset C^\infty$ but there is some funny additional condition, and similarly $\bigcup_{s\in\mathbb{R}} H^s \subset \mathcal{S}'$ but again we don't get everything.

Proposition 84 (Sobolev embedding theorem)
We have $H^s \subset C^0$ if $s > \frac{n}{2}$ and $H^s \subset C^k$ if $s > \frac{n}{2} + k$.

It turns out that $H^s \subset C^{k,\alpha}$ if $s > \frac{n}{2} + k + \alpha$, but we won't prove that here. And notice that $u \in H^s$ implies that $D_j u \in H^{s-1}$, basically immediately from the definition, so we just need to prove the first inclusion. (The reason these are nice to work with is that even if we're in a very negative-s Sobolev space, we can still keep taking derivatives in the distributional sense.)

Proof. Suppose $s > \frac{n}{2}$ and $u \in H^s$. Then we have

$$u(x) = \int e^{ix\cdot\xi}\hat{u}(\xi)d\xi,$$

and we wish to show that this is a uniformly continuous integral so that we will have continuity in x by the dominated convergence theorem. We'll use Cauchy-Schwarz here:

$$\left|\int e^{ix\cdot\xi}\hat{u}(\xi)d\xi\right|^2 \le \int |\hat{u}(\xi)|^2(1+|\xi|^2)^s \int (1+|\xi|^2)^{-s},$$

and because s is large enough the last term decays fast enough to have a convergent integral, and the first term on the right-hand side is finite by assumption, we do indeed have continuity when $s > \frac{n}{2}$. □

Thinking about how this relates to distributions now, we'll do a few examples:

- Recall that δ_0 is the distribution such that $\hat{\delta_0} = 1$, and $\int |1|^2(1+|\xi|^2)^s < \infty$ whenever $s < -\frac{n}{2}$. So the delta function is in H^s for all $s < -\frac{n}{2}$, and the precise rate of growth of the Fourier transform (being exactly constant) will mean that we're in an open interval of weighted L^p spaces.
- We also have the distribution p.v.$\frac{1}{x}$, defined by setting

  $$\langle \text{p.v.}\frac{1}{x}, \phi\rangle = \lim_{\varepsilon\to 0}\int_{|x|\ge\varepsilon}\frac{\phi(x)}{x}dx$$

  The point is that we get a singularity at zero because $\frac{1}{x}$ gets large near the origin, so we take out a **symmetric** interval around the origin (and we don't have to worry about the decay at infinity because Schwartz functions decay rapidly there – some amount of extra regularity beyond continuity is important). It turns out that this distribution will again be in H^s for $s < -\frac{1}{2}$, and these two examples are "both homogeneous of degree -1."
- In higher dimensions, we may consider the sphere $S^{n-1} \subset \mathbb{R}^n$ and consider the delta function on S (also called the **delta measure** supported on the sphere):

  $$\langle \delta_{S^{n-1}}, \phi\rangle = \int_{S^{n-1}} \phi.$$

  To figure out what Sobolev space this is in, we must take the Fourier transform: we want to compute $\hat{\delta}_S = e^{-i\omega\xi}d\omega$, and the idea is to divide up into a coordinate $\omega' \in S^{n-2}$ (higher-dimensional longitude) and an angle $\omega_n = \cos\phi$ for $-\frac{\pi}{2} < \phi < \frac{\pi}{2}$ (measuring the latitude). Then the integral can be rewritten to $\int e^{-i|\xi|\cos\phi}(\cos\phi)^{n-2}d\phi d\omega$, and the ω-integral now just gives us the volume of S^{n-2}. But we're left with a one-dimensional integral of the form $\int_{-\pi/2}^{\pi/2}(\cos\phi)^{n-2}e^{-i|\xi|\cos\phi}d\phi$, and we want to know how quickly this decays as ξ goes to infinity. So for **geometric distributions** like this (badly singular but happening geometrically nicely) we can often understand what's going on by explicit contributions.

Example 85

More generally, the well-behavedness depends on the curvature of the surface – flat points mean slow decay rates. To see that, we'll consider the delta function $\delta_{\mathbb{R}^k}$ in $\mathbb{R}^n$, so we set $x = (x', x'')$ and let $\mathbb{R}^k = \{(x', 0)\}$.

Then we have (when we write down this integral, we really mean a distributional pairing)

$$\hat{\delta}_{\mathbb{R}^n} = \int_{x''=0} e^{-ix'\cdot\xi' - ix''\cdot\xi''} dx' = \int_{x''=0} e^{-ix'\cdot\xi'} dx'$$

which is the inverse Fourier transform of the function 1 and thus yields $\delta(\xi')$. But what we wrote down doesn't really make sense a priori, since $\langle e^{-ix\xi}, \delta_{\mathbb{R}^k}\rangle$ is pairing an element of $\mathcal{S}'$ with something that's not in Schwartz space. And because we're not pointwise defined, we can't be in a Sobolev space with this definition. Instead, the idea is to take some smooth cutoff function $\chi(x)\delta_{\mathbb{R}^k}$, so that

$$\langle e^{-ix\xi}, \delta_{\mathbb{R}^n}\chi(x)\rangle = \langle e^{-ix\xi}\chi(x), \delta_{\mathbb{R}^k}\rangle$$

and now this actually makes sense as a distributional pairing. As a function of ξ', then, we go to zero very rapidly, and in the ξ'' direction it's constant. But it's in a weighted L^2 space $\mathbb{R}^{n-k}$, so we just need to choose $\boxed{s < \frac{-(n-k)}{2}}$ so that the decay in the ξ'' direction is not a problem.

Definition 86

Let $\phi \in C_0^\infty$. Then define the mapping $M_\phi : \mathcal{S}' \to \mathcal{S}'$ given by $M_\phi(u) = \phi u$ (in other words, when we pair with any Schwartz function ψ, $\langle \phi u, \psi\rangle = \langle u, \phi\psi\rangle$, and this makes sense because $\phi\psi$ is still Schwartz).

We do not actually need ϕ to be in C_0^∞ – it's enough to be smooth and have derivatives only grow polynomially, since we just need $\phi\psi$ to be Schwartz.

Proposition 87

The map $M_\phi : H^s \to H^s$ is bounded for all s.

We'll see two different proofs of this:

Proof 1. Clearly M_ϕ maps $L^2 \to L^2$, and from the weak derivatives arguments above it also maps $H^k \to H^k$ (since we can differentiate ϕu up to k times by using the product rule). So then by Riesz-Thorin, we're bounded on any H^s for all $s \geq 0$. And then again the same duality trick as with the Hilbert transform shows boundedness for $s < 0$ as well. □

Proof 2. This time we'll use Fourier transforms and convolutions: by definition, we have

$$\widehat{\phi u}(\xi) = \int \hat{\phi}(\eta)\hat{u}(\xi - \eta) d\eta$$

(convolution turns into multiplication and vice versa); if u is in one of the weighted L^2 spaces then this will make sense for almost every ξ because ϕ is decaying so quickly. So now we just need to estimate

$$\int |\widehat{\phi u}|^2 (1 + |\xi|^2)^s = \int \left| \int \hat{\phi}(\xi - \eta)\hat{u}(\eta) d\eta \right|^2 (1 + |\xi|^2)^s,$$

and the idea is that we can bound (by rapid decay of ϕ)

$$|\hat{\phi}(\xi - \eta)| \le C(1 + |\xi - \eta|^2)^{-N}$$

for every N. So now by Cauchy-Schwarz, we can bound the previous expression as

$$\int |\widehat{\phi u}|^2 (1 + |\xi|^2)^s \le \int \left(\int |\hat{\phi}(\xi - \eta)|^2 (1 + |\eta|^2)^{-s} d\eta \right) \left(|\hat{u}(\eta)|^2 (1 + |\eta|^2)^s d\eta \right) (1 + |\xi|^2)^s d\xi,$$

and now we can bound this in terms of the H^s norm of u as

$$\le ||u||_{H^s}^2 = \int \int (1 + |\xi - \eta|^2)^{-N} (1 + |\eta|^2)^{-s} (1 + |\xi|^2)^s d\xi d\eta,$$

but this doesn't decay properly near $\xi - \eta$ so the Cauchy-Schwarz isn't working out correctly. Instead, we'll bound as

$$\le \int \left| \int (1 + |\xi|^2)^{s/2} (1 + |\xi - \eta|^2)^{-N} |\hat{u}(\eta)| d\eta \right|^2 d\xi,$$

and now we can use the inequality

$$(1 + |\xi|^2)^{s/2} \le C(1 + |\xi - \eta|^2)^{|s|/2} + (1 + |\eta|^2)^{s/2}.$$

To understand this inequality, substitute in $\xi \mapsto \xi - \eta$. Then for $s = 2$, we're basically saying that $1 + |\xi - \eta|^2 \le C(1 + |\xi|^2)(1 + |\eta|^2)$, which is true because we can estimate the cross-term by Cauchy-Schwarz. So then we can raise both sides to the $s/2$ power for any positive s, and when s is negative we need to rearrange the terms a little bit. Substituting in this inequality, we thus have

$$\int |\widehat{\phi u}|^2 (1 + |\xi|^2)^s \le \int \left| \int (1 + |\xi - \eta|^2)^{|s|/2 - N} |\hat{u}(\eta)| (1 + |\eta|^2)^{s/2} d\eta \right|^2 d\xi,$$

which we can now split apart as

$$\le \int \left(\int (1 + |\xi - \eta|^2)^{|s|/2 - N} d\eta \right) \left(\int (1 + |\xi - \eta|^2)^{|s|/2 - N} (1 + |\eta|^2)^s d\eta \right) d\xi$$

The first term is independent of ξ and is just a constant that we can pull out, and what remains is a convergent double integral because the ξ-integral has a rapidly-decaying factor and we are bounded by the H^s norm of u. So we can swap the order of integration and that does give us some constants times the H^s norm of u. (The whole point is that with these proofs we often have to do Cauchy-Schwarz and other manipulations in just the right order.) $\square$

Since we've been refering to more general distributions than $\mathcal{S}'$, we'll go ahead and make a formal definition now:

Definition 88

Let $\Omega \subset \mathbb{R}^n$ be open. Let $C_0^\infty(\Omega)$ denote the set of smooth functions C^∞ in Ω such that $\text{supp}\phi \subset \Omega$.

Recall that $\phi_j \to \phi$ converging in $C_0^\infty(\Omega)$ means that the support of all ϕ_js is contained in the same compact set K within Ω (without going to the boundary), and then the C^k norm of $\phi_j - \phi$ goes to zero for all k. So we get almost a Frechet space but with a slightly worse condition, and this is really writing Ω as a countable limit of compact sets and taking the inverse inductive limit over the Frechet spaces. But the point is that we have basically convergence in each C^k norm.

The point is that distributions are just the dual of this space:

Definition 89

The space $\mathcal{D}'(\Omega)$ of distributions on Ω is the dual of $C_0^\infty(\Omega)$.

This allows us to make a more localized definition of distributions – we often call elements of $C_0^\infty(\Omega)$ **test functions**, and $u \in \mathcal{D}'$ is basically a black box until we test it against some ϕ. The topology of this is pretty messy because it's the dual of a nuclear Frechet space, but what we really care about is convergence. Unpacking the definitions, $u \in C_0^\infty(\Omega)^*$ means that for each compact set K, and for all ϕ with support in K, there is some k such that we're bounded by that corresponding seminorm, so that $|\langle u, \phi\rangle| \le C_{K,k}||\phi||_{C^k}$. So on every K, u does something which uses up to k derivatives of ϕ.

Example 90

We can have $\langle u, \phi\rangle$ be $\phi(0)$, or it can be $\sum_{|\alpha|\le k} C_\alpha D^\alpha \phi(0)$, or we can take the values of ϕ at even infinitely many values as long as there is no accumulation point (since we will have finitely many on each compact set). For a more general example, we can take $\sum_{|\alpha|\le k} D_x^\alpha f_\alpha$, where $f_\alpha \in L^1_{\text{loc}}$.

This last case means that we have

$$\langle u, \phi\rangle = \sum_{|\alpha|\le k} \int (-1)^{|\alpha|}(D^\alpha \phi) f_\alpha,$$

and in fact this is basically the most general type of distribution – if we fix K, then any distribution is a sum of L^1_{loc} functions with derivatives. And then having $u_j \to u$ in $\mathcal{D}'$ is the weak definition that $\langle u_j, \phi\rangle \to \langle u, \phi\rangle$ for any test function ϕ. (So distributions look very abstract, but they're rather concrete at the end of the day and convergence isn't too bad.) We can see that this topology is rather weak though:

Proposition 91

$C^\infty(\Omega)$ is dense in $\mathcal{D}'(\Omega)$.

Proof. Given a distribution u, we define

$$u_\varepsilon(x) = \int \chi_\varepsilon(x-y)u(y)dy = \langle u, \chi_\varepsilon(x-\cdot)\rangle,$$

where we think of χ as being in C_0^∞, nonnegative, and with integral 1. Then we're basically convolving u and smearing it out (using the usual approximations to the identity), and we claim that we always get a smooth function. Each χ_ε is a smooth function of $\cdot$, so this pairing is well-defined. To prove that it is smooth, we must first check that it is continuous. Indeed, $|\langle u, \chi_\varepsilon(x-\cdot)\rangle| \le C||\chi_\varepsilon(x-\cdot)||_C^k$ for some k (and where everything's living inside some compact set K), thus as we vary x the $\chi_\varepsilon(x - \dot{)}$s will vary continuously as well, since

$$|\langle u, \chi_\varepsilon(x-\cdot) - \langle u, \chi_\varepsilon(x'-\cdot)| \le C||X_\varepsilon(x-\cdot) - \chi_\varepsilon(x'-\cdot)||_{C^k}$$

and the right-hand side goes to zero because the C^k norms behave well with translation. And the same works for higher derivatives too, just passing the derivatives to the χ_εs in the pairing. It now remains to check that u_ε converges to u in this distributional sense. But

$$\langle u_\varepsilon, \phi\rangle = \langle u * \chi_\varepsilon, \phi\rangle = \langle u, \chi_\varepsilon * \phi\rangle$$

(we can move the convolution over), and thus this converges to $\langle u, \phi\rangle$ because u is continuous and $\chi_\varepsilon * \phi$ converges to ϕ in C^∞. $\square$

For our last result, we'll go back to Sobolev spaces but in the language of distributions:

Theorem 92 (Rellich)
If $s > t$, then $\{u \in H^s : \text{supp}(u) \subset K\}$ is contained in H^t, and the inclusion is compact.

It's not true that H^s must sit in H^t, but if we demand that things are supported in a compact set, then we do get this inclusion. And having a compact inclusion means that if $u_j \in H^s$ with all supports contained within K, and all norms are bounded, then there is a subsequence u_j' which converges in H^t. This is basically a more doctored-up version of Arzela-Ascoli – regularity means continuity in a weaker space.

We'll do the proof next time, but the point is that **each $\hat{u}_j$ is C^∞**, since $\hat{u}_j = \int e^{-ix\cdot\xi} u_j(x)dx$ is actually smooth in ξ. We've only defined this integral in a weak sense, but because of the support condition, $u_j = \chi u_j$ where $\chi = 1$ on the support of u_j and $\chi \in C_0^\infty$. So this integral is actually the same as $\int e^{-ix\cdot\xi}\chi(x)u_j(x)$, and now this pairing doesn't need to be made in a distributional sense anymore. So then reinterpreting this as $\langle u_j, \chi(x)e^{-ix\cdot\xi}\rangle$, we're now pairing against a Schwartz function. As ξ varies, this is a smoothly varying family of functions, and using the continuity definition we see that the pairing will also be smooth in ξ. And this has nothing to do with being in a Sobolev space – it's just that we can multiply by a cutoff function without any harm – so being able to apply Arzela-Ascoli is not too shocking.

9 February 7, 2023

We'll back up a bit to go over a few more basics of distribution theory and then continue our discussion of last time. One thing to remember is that the smaller our space is, the larger the dual space will be and thus the less well-behaved the objects will look like.

Recall that the space of tempered distributions $\mathcal{S}'$ is a nice space to talk about globally on $\mathbb{R}^n$ (encapsulating how bad distributions can be locally, with some growth condition at infinity meaning that they can be paired with rapidly decaying functions). For comparison, if we take any $\Omega \subset \mathbb{R}^n$, we have the space $C_0^\infty(\Omega)$ (with the weird condition that we need to look at supports within compact subsets), which allows us to define the dual space $\mathcal{D}'(\Omega)$. The nice thing about $\mathcal{D}'(\Omega)$ is that it can be defined on any Ω, but it behaves a little less nicely than $\mathcal{S}'$ in a certain way. Remember that for any $u \in \mathcal{D}'$ we must have that for any compact set K there is some k, C such that $|\langle u, \phi\rangle| \le C||\phi||_{C^k}||$ for all ϕ supported in K. So ϕ sees a finite number of derivatives in any compact set, much like tempered distributions do, but it's possible for it to still have "infinite order" in the following way:

Example 93
Consider an infinite sequence of points in Ω, and let $u = \sum \delta_{p_j}^{(j)}$, where we define $\delta_{p_j}^{(j)}$ via

$$\langle \delta_{p_j}^{(j)}, \phi\rangle = (-1)^j \phi^{(j)}(p_j).$$

Then there's finitely many points p_j so we only use up badness to C^k on any compact set, but we still have infinite order and thus can be "arbitarily bad near the boundary."

Definition 94
For any distribution $u \in \mathcal{D}'$, the **support** of u is defined via

$$(\text{supp } u)^c = \{x : \langle u, \phi\rangle = 0 \quad \forall \phi \in C_0^\infty(B_\varepsilon(x)) \text{ for some } \varepsilon > 0\}.$$

So for any point x not in the support, there is some ε so that any function supported within that ball of radius ε is "not seen by u." **From here on**, if we say "distribution" we're talking about $\mathcal{D}'$.

Proposition 95

If u has support equal to a single point (which we may take to be the origin), then $u = \sum_{|\alpha| \le m} c_\alpha \delta_\theta^{(\alpha)}$.

Proof. Take any compact set K with 0 in its interior. By continuity in the sense described above, u must have finite order m in some fixed ball $\overline{B}_\varepsilon(0)$. (We can check that the order must be the same if we take a smaller ball, since what's happening between the balls must always be zero.)

Now take any general test function ϕ, and let $\chi \in C_0^\infty$ be a cutoff function so that $\chi = 1$ on $B_\varepsilon(0)$. Then $\langle u, \phi \rangle = \langle u, \chi\phi \rangle$, since $\langle u, (1-\chi)\phi \rangle = 0$ and then we can use linearity. So expand ϕ in a Taylor series up to order m, so

$$\phi = \sum_{|\alpha| \le |m|} \frac{1}{\alpha!} \phi^{(\alpha)}(0) x^\alpha + r_{m+1}(x), \quad \chi\phi = \sum_{|\alpha| \le m} \left(\frac{\chi(x)}{\alpha!} \phi^{(\alpha)}(0) x^\alpha \right) + \chi r_{m+1}.$$

Pairing against u, we see that

$$\langle u, \chi\phi \rangle = \sum_{|\alpha| \le m} \langle u, \frac{\chi(x)}{\alpha!} \phi^{(\alpha)}(0) x^\alpha \rangle + \langle u, \chi r_{m+1} \rangle.$$

But if χ is supported only in a ball of radius ε, then χr_{m+1} looks like ε^{m+1}, so its C^m norm is bounded by $C\varepsilon$. That's the only part that depends on ε, so in fact the remainder term contributes zero. But now in fact other than the partial derivative terms, we just have a bunch of numbers that don't depend on the ϕ function, so we indeed have $\langle u, \phi \rangle = \sum_{|\alpha| \le m} c_\alpha \phi^{(\alpha)}(0)$, as desired. $\square$

Definition 96

The space $\mathcal{E}'(\Omega)$ is the subset of $\mathcal{D}'(\Omega)$ with compact support.

It turns out that this space is the dual space of $C^\infty(\Omega)$ (which allows function to behave arbitrarily at the boundary other than being bounded). So on $\mathbb{R}^n$ we have $\mathcal{E}' \subset \mathcal{S}' \subset \mathcal{D}'$ (since $\mathcal{D}'$ has no growth restrictions at the boundary, $\mathcal{S}'$ has some conditions, and $\mathcal{E}'$ cannot grow at all).

A useful fact is that distributions are **localizable and coordinate-invariant**. For the former, we're saying that if $u \in \mathcal{D}'$ and $\chi \in C_0^\infty(\Omega)$, then $\chi u \in \mathcal{D}'$ as well, since we may define $\langle \chi u, \phi \rangle = \langle u, \chi\phi \rangle$. And for the latter, suppose we have a diffeomorphism $F : \Omega \to \Omega'$, given a function $\phi \in C_0^\infty(\Omega')$ we also have a pullback function $F^*\phi \in C_0^\infty(\Omega)$ (defined via $F^*\phi(x) = \phi(F(x))$. This is a reasonable thing to do because a diffeomorphism takes compact sets to compact sets. Then we can define "pushing forward" by duality: if u were a smooth function (which are dense in distributions, so combined with continuity this would give us everything), then

$$\langle u, F^*\phi \rangle = \int u(x)\phi(F(x))dx = \int u(F^{-1}(y))\phi(y)(JF^{-1})dy$$

by the change-of-variables formula $y = F(x)$, where JF^{-1} indicates the Jacobian of the inverse map. (Here note that $\phi(y)JF^{-1}$ is still C_0^∞.) So really, we're saying that we should define

$$\langle (F^{-1})^* u, \phi \rangle = \langle u, (F^*\phi)(JF^{-1})^{-1} \rangle.$$

This is continuous with respect to the topology of u, so this does define the pullback. The point is that if M is any manifold, then $\mathcal{D}'(M)$ will be well-defined (since we have an open cover of pieces of $\mathbb{R}^n$ up to diffeomorphism – use a partition of unity to transform any object to things contained within coordinate charts).

Notice that because $u \in \mathcal{E}' \subset \mathcal{S}'$, then $\hat{u} \in \mathcal{S}'$, and we claim that in fact $\hat{u} \in C^\infty(\mathbb{R}^n)$. Indeed, because u has compact support, we can say that

$$\hat{u}(\xi) = \langle u, e^{-ix\xi}\chi(x)\rangle$$

where $\chi \in C_0^\infty$ with $\chi = 1$ on the support of u and 0 far enough away, so this Fourier transform does actually make sense as a function of ξ. By continuity of the distribution, we have

$$\left|\langle u, e^{-ix\xi}\chi(x)\rangle\right| \le C||e^{-ix\xi}\chi(x)||_{C^m},$$

but continuity tells us more: If we want to take the derivative of the Fourier transform, we have

$$\partial_{x_j}\hat{u}(\xi) = \lim_{h\to 0}\frac{\langle u, e^{-i\xi\cdot(x+he_j)}\chi(x)\rangle - \langle u, e^{-i\xi\cdot x}\chi(x)\rangle}{h} = \lim_{h\to 0}\left\langle u, \chi(x)\frac{e^{-i\xi\cdot he_j - 1}e^{-ix\cdot\xi}}{h}\right\rangle.$$

Now the mth derivative of $\frac{e^{-i\xi\cdot he_j-1}e^{-ix\cdot\xi}}{h}$ goes to zero as $h \to 0$, so in fact any difference quotient has a limit that exists and thus we can differentiate as many times as we want. In other words, we can differentiate under the integral sign because of being compactly supported.

We can now prove Rellich's theorem from last time, which states that for any real numbers $s > t$, we have a compact inclusion $H^s_K(\mathbb{R}^n) \subset H^t(\mathbb{R}^n)$, where $H^s_K(\mathbb{R}^n)$ is the set of functions in $H^s(\mathbb{R}^n)$ only supported in the compact set K.

Proof. Take any sequence $u_j \in H^s_K$ with $||u_j||_s \le 1$. We wish to show that some subsequence converges in H^t. The Sobolev norms involve Fourier transforms, and because we have compact support in fact $\hat{u}_j(\xi) \in C^\infty$ for all j, with $\int|\hat{u}_j(\xi)|^2(1+|\xi|^2)^s \le 1$. Choose R large enough so that $(1+R^2)^{t-s} < \varepsilon$, we see that

$$\int_{|\xi\ge R}(1+|\xi|^2)^t|\hat{u}(\xi)|^2 d\xi \le (1+R^2)^{t-s}\int_{|\xi\ge R}(1+|\xi|^2)^s|\hat{u}(\xi)|^2 d\xi \le (1+R^2)^{t-s} < \varepsilon,$$

so now we have uniform smallness at infinity, and now we can write $\hat{u}_j = \chi_R\hat{u}_j + (1-\chi_R)\hat{u}_j$, where we can bound the second term in norm already. Now if we take the inverse Fourier transform $u_{j,R}$ of the first term $\chi_R\hat{u}_j$, which is compactly supported, we get something in C^∞, and we have

$$u_{j,R} = \int_{|\xi|\le R} e^{ix\xi}\hat{u}_j(\xi)d\xi.$$

But this is a convergent integral, since we can take as many derivatives as we we want and the ξ factors are all bounded: we have a constant A_m so that

$$||u_{j,R}||_{C^m} \le A_m \quad \forall j.$$

So we have a sequence of functions uniformly bounded, and C^{n+1}-norm boundedness means we can extract a subsequence where C^n norms converge by Arzela-Ascoli, and choose subsequences that converge in larger and larger open sets. So we get some $u_{j',R}$ converging to $u_{\infty,R}$ in C^∞. So now we can divide up the integral

$$\int(1+|\xi|^2)^t|\hat{u}_j(\xi) - \hat{u}_k(\xi)|^2 d\xi$$

into the two parts $|\xi| \le R$ and $|\xi| \ge R$, using our two different facts, to show that we do have convergence as desired. $\square$

So the point of this inclusion is that "supports could drift off to infinity," but if we prevent that from happening then putting any slightly weaker topology allows us to make the tail parts converge too.

We might ask now why we care about Sobolev spaces for real valued s, rather than just integers. Here's an example where fractional spaces do come into play:

Proposition 97

Consider the map $R : u(x) \to u(x', 0)$, where $x = (x', x_n) \in \mathbb{R}^n$ and with $x' \in \mathbb{R}^{n-1}$. (In other words, for example for any continuous function, we restrict our function to just a hyperplane.) Then for any $s > \frac{1}{2}$, this extends to a bounded map $R : H^s(\mathbb{R}^n) \to H^{s-1/2}(\mathbb{R}^{n-1})$.

So in particular, we at least need to care about half-integer Sobolev spaces. We can do this with Fourier transforms – there's a more subtle result for L^p-based Sobolev spaces which drops by $\frac{1}{p}$ instead of $\frac{1}{2}$, but that can be done by interpolation.

Proof. Assume that u is Schwarz so that we have a nice Fourier integral (we can get a bound and extend by density). Write out the usual Fourier inversion formula

$$u(x', x_n) = \int e^{-ix'\cdot\xi' + ix_n\xi_n}\hat{u}(\xi', \xi_n)d\xi' d\xi_n,$$

and now define $v(x') = u(x', 0)$. We wish to show that $||v||_{H^{s-1/2}} \le C||u||_{H^s}$. Since v basically plugs in 0 for x_n, we also have

$$v(x') = \int e^{ix'\cdot\xi'}\hat{u}(\xi', \xi_n)d\xi' d\xi_n \implies \hat{v}(\xi') = \int \hat{u}(\xi', \xi_n)d\xi_n,$$

so in other words we take the Fourier transform of the whole thing and integrate out the ξ_n variable. Our goal is then to evaluate

$$\begin{aligned}\int |\hat{v}(\xi')|^2(1 + |\xi'|^2)^{(s-1/2)}d\xi' &= \int \left|\int \hat{u}(\xi', \xi_n)d\xi_n\right|^2 \left(1 + |\xi'|^2\right)^{s-1/2} d\xi' \\ &\le \int \left[\left(\int |\hat{u}(\xi', \xi_n)|^2(1 + |\xi|^2)^s d\xi_n\right)\left(\int (1 + |\xi|^2)^{-s}d\xi_n\right)\right](1 + |\xi'|^2)^{s-1/2}d\xi'\end{aligned}$$

by Cauchy-Schwarz on the red part. Now if we look at the one-dimensional integral

$$\int (1 + |\xi|^2)^{-s}d\xi_n = \int (1 + |\xi'|^2 + \xi_n^2)^{-s}d\xi_n,$$

the point is that if $s > \frac{1}{2}$ then this integral decays faster than $\frac{1}{\xi}$ so things are convergent. If we now do a change-of-variables $\xi_n = \sqrt{1 + |\xi'|^2}\eta_n$, that integral becomes

$$= \int (1 + |\xi'|^2)^{-s}(1 + \eta_n^2)^{-s}d\eta_n \cdot (1 + |\xi'|^2)^{1/2},$$

so this whole thing is some constant C times $(1 + |\xi'|^2)^{1/2-s}$. Plugging that into the blue part above, it cancels out exactly with the last term in the integral, so

$$\int |\hat{v}(\xi')|^2(1 + |\xi'|^2)^{s-1/2}d\xi' \le C\int\int |\hat{u}(\xi', x_n)|^2(1 + |\xi|^2)^s d\xi = C||u||_{H^s},$$

as desired. $\square$

We can now see an application of how everything fits together in a concrete way, solving constant-coefficient elliptic equations. Recall that if $P = \sum_{|\alpha| \le m} c_\alpha D_x^\alpha$ is a differential operator with $c_\alpha \in \mathbb{C}$, we can take the associated symbol

$$P(\xi) = \sum_{|\alpha| \le m} c_\alpha \xi^\alpha,$$

so with such an operator we're equivalently multiplying by polynomials in the Fourier space:

$$\widehat{P(D)u} = P(\xi)\hat{u}.$$

So we just need to invert the polynomial to invert the operator, and for a certain class this is easy. We'll do an example:

Proposition 98

Let $P = \Delta + 1$, where $\Delta = \sum_{j=1}^n D_j^2 = -\sum_{j=1}^n \frac{\partial^2}{\partial x_j^2}$ (this is the "analyst's Laplacian"), so that $P(\xi) = |\xi|^2 + 1$. Then $P(D) : \mathcal{S} \to \mathcal{S}$ is an isomorphism, and for all $f \in \mathcal{S}$ there is some $u \in \mathcal{S}$ with $Pu = f$.

Proof. This is equivalent to the statement that $(1 + |\xi|^2) : \mathcal{S} \to \mathcal{S}$ is an isomorphism. Indeed, if $(1 + |\xi|^2)\hat{u} = \hat{f}$, the function $\hat{u} = \frac{\hat{f}}{1+|\xi|^2}$ clearly solves our equation because $1 + |\xi|^2$ vanishes nowhere. □

Proposition 99

Now consider the same P as a mapping $\mathcal{S}' \to \mathcal{S}'$. Then we still have an isomorphism, so we can still solve $Pu = f$ for $f \in \mathcal{S}'$.

Proof. Again we take $\hat{u} = \frac{\hat{f}}{1+|\xi|^2}$, and the point is that we are indeed in $\mathcal{S}'$ because by definition

$$\langle \frac{\hat{f}}{1+|\xi|^2}, \hat{\phi} \rangle = \langle \hat{f}, \frac{\hat{\phi}}{1+|\xi|^2} \rangle$$

□

So given any $f \in \mathcal{S}$, there is a unique solution $u \in \mathcal{S}$ solving $Pu = f$ and this is the unique solution **even** if we look at all tempered distributions. But interestingly, there could still be solutions not in $\mathcal{S}'$ that aren't the one we've found:

Example 100

For the function e^{x_1}, we have $Pe^{x_1} = -\partial_{x_1}^2 e^{x_1} + e^{x_1} = 0$, which is the same as $P0 = 0$. In fact, there are plenty of complex solutions to $Pu = 0$ but they are exponentially growing – the point is that for any polynomial $P(\xi)$, there's always a whole variety of zeros corresponding to $e^{x \cdot \xi}$ that will be exponentially growing solutions in $\mathbb{R}^n$. But within the world of $\mathcal{S}'$ we can still get uniqueness statements.

In general for PDE questions, we often care about existence and regularity questions, and here's an example of a result we may care about:

Proposition 101

We have an isomorphism $P : H^{s+2}(\mathbb{R}^n) \to H^s(\mathbb{R}^n)$.

Proof. Again we take the same $\hat{u}$ as before, but now we need to show that u is actually in H^{s+2}. Indeed,

$$\int |\hat{u}|^2 (1 + |\xi|^2)^{s+2} = \int \frac{|\hat{f}|^2}{(1 + |\xi|^2)^2} (1 + |\xi|^2)^{s+2} = \int |\hat{f}|^2 (1 + |\xi|^2)^s < \infty.$$

□

Definition 102

A distribution $u \in \mathcal{D}'(\Omega)$ is said to be in $H^s(U)$ for some open set $U \subset \Omega$ if $\chi u \in H^s(\mathbb{R}^n)$ for all $\chi \in C_0^\infty(U)$.

In other words, we can throw away everything outside of U. The point is that multiplying anything in H^s by a cutoff function stays in H^s, so this is actually a reasonable definition (we can take different χs and nothing will change.)

Theorem 103

Suppose $u, f \in \mathcal{D}'(\Omega)$, and suppose $Pu = f$ in the sense of distributional derivatives (that is, in $\mathcal{D}'(\Omega)$). Then if $f \in H^s(U)$, then $u \in H^{s+2}(U)$.

So this is a "local result" strengthening the previous ones: even if f behaves badly, as long as it is regular in U, we gain two orders of differentiability for u "locally."

Proof. Take a sequence of cutoff functions: let $\chi_0 \in C_0^\infty(\Omega)$ so that the support of χ_0 is contined in U. If we want to show that u is in $H^{s+2}(U)$, it suffices to show that for any smaller open set we're in H^{s+2} there. So we'll choose a cascading set $U_0 \supset U_1 \supset U_2 \supset \cdots \supset U_\ell$, and we'll show that u is H^{s+2} in U_ℓ. Notice that

$$P(\chi_0 u) = \chi_0 P u + [P, \chi_0]u,$$

and remembering that $[D_j, \chi_0] = D_j \chi_0$ as operators for any j, and $D_j \chi_0$ is a 0th order operator even though D_j is a first order operator. And the point is that $[P, \chi_0]$ only involve things up to order $2 - 1$ with all coefficients supported in U. So now $\chi_0 u$ is compactly supported, so it is in some H^t (since it's a distribution with some finite order, and the Fourier transform grows at some polynomial rate so we just need to choose t enough to make that integrable). Then

$$P(\chi_0 u) = \chi_0 f + [P, \chi_0]u,$$

and the first term is in H^s and the latter term is in H^{t-1} since the commutator is only first-order and the coefficients are all nicely supported in U. So the left-hand side is globally in some Sobolev space, and $\chi_0 u \in H^{\min(s+2, (t-1)+2)} = H^{\min(s+2, t+1)}$. So if we assume $\chi_0 u$ is in H^t, it's actually also in H^{t+1} if s is much bigger than t. Now we can keep applying this argument again and again, taking smaller and smaller cutoff functions, finitely often until we show that we're actually in H^{s+2} (when $s + 2$ becomes smaller than $t + 1$). $\square$

So constant-coefficient equations work nicely if we do things globally, and for local statements we need to mess aroudn with cutoff functions. The idea is always to look at the commutators and show that they're insignificant in the long run. We'll do one more argument of this sort before returning to more abstract functional analysis next time!

10 February 9, 2023

We briefly discussed Sobolev spaces on open subsets of $\mathbb{R}^n$, but we'll do it more carefully now. For arbitrary sets Ω it is a bit hard to understand, but $H^s_{\text{loc}}(\Omega)$ is easy to define: it is the set of $u \in \mathcal{D}'(\Omega)$ such that $\phi u \in H^s(\mathbb{R}^n)$ for all $\phi \in C_0^\infty(\Omega)$ (so whenever we localize a distribution it's in the appropriate Sobolev space). But because H^s also has some implicit growth conditions at infinity, there's complexity there – if Ω is smoothly bounded then we have an intelligible answer. For example, if we take $C_0^\infty(\Omega)$ and take the closure in the H^s topology (which makes sense, since for any such sequence $u_j - u_k$ are zero outside Ω and in fact vanishing on the boundary as long as that notion makes sense.

Separately, we can also define $H^1(\Omega), H^2(\Omega), \cdots$ in terms of weak derivatives for any domain Ω: we have $\{u : u \in L^2(\Omega), Du \in L^2(\Omega)\}$, where we're defining $\langle D_j u, \phi\rangle = -\langle u, \partial_{x_j}\phi\rangle$, and we have $\partial_{x_j} u \in L^2$ if there is some function $f_j \in L^2$ where we happen to agree with the usual pairing $\langle f_j, \phi\rangle = -\langle u, \partial_{x_j}\phi\rangle$. So this is well-defined regardless of Ω, and then for the fractional and negative-ordered spaces we can interpolate and dualize. But the point is that whenever Ω is complicated, this is not an easy task.

The problem with just trying to take our function $u \in L^2$ and taking its Fourier transform is that it may be artificial to extend it to being zero outside Ω. And for an example of a bad situation, something like $f = 1$ being extended by 0 outside of Ω would create an artificial singularity which isn't inherent to the original function. So this is connected to the question of "the right class of distributions" we've been asking – for example, if we want to define distributions on $\mathbb{R}^+$ we can just do $\mathcal{D}'(\mathbb{R}^+)$, but we can also do things with **supported distributions** (meaning that extending by zero still gives us a valid distribution) or **extended distributions** (restrictions of distributions defined everywhere). So H^s could contain the set of functions which extend to H^s on a larger set, and we could define norms in terms of the infimum H^s norm of any extension. One useful result of this form is that when s is positive we can indeed get a bounded map $E : H^s(\mathbb{R}^n_+) \to H^s(\mathbb{R}^n)$ of this type.

We can now continue our discussion from last time. We introduced the constant-coefficient differential operators last time, focusing on the operator $P = \Delta + 1 = \sum D_j^2 + 1$, and showing that the Fourier transform of Pu is $P(\xi)\hat{u}$, where $P(\xi) = |\xi|^2 + 1$ because derivatives become multiplication under the Fourier transform. We proved that P is an isomorphism $\mathcal{S} \to \mathcal{S}$, $\mathcal{S}' \to \mathcal{S}'$, and also $H^{s+2} \to H^s$ for all s, all of which basically involve dividing by $|\xi|^2 + 1$. We then showed that if $Pu = f$ and f restricted to some open set U gives us H^s_{loc}, then $u \in H^{s+2}_{\text{loc}}(U)$ – the proof here involved looking at $P(\chi u)$ for some cutoff function χ and using an inductive process to show that we can make the regularity better and better up until $s + 2$.

Remark 104. *This technique is very general – it works as long as $P = \sum_{|\alpha|\le m} a_\alpha D^\alpha$ for constants $a_\alpha \in \mathbb{C}$, or even if $a_\alpha \in M_k(\mathbb{C})$, as long as* **the corresponding symbol** *$P(\xi) = \sum_{|\alpha|\le m} a_\alpha \xi^\alpha$ has that $||P(\xi)^{-1}|| \le C(1 + |\xi|)^{-m}$ is bounded in matrix norm for all ξ. Then all of the arguments work exactly as stated, since $Pu = f$ is equivalent to $P(\xi)\hat{u} = \hat{f}$, which is equivalent to $\hat{u} = P(\xi)^{-1}\hat{f}$ (which always makes sense in $\mathcal{S}'$ because of the polynomial bound, and it's* **smooth** *in terms of ξ by Cramer's rule). But if we no longer have constant coefficients, things become significantly more messy.*

Definition 105

A differential operator $P = \sum_{|\alpha|\le m} D^\alpha$ is **elliptic of order m** if the top-order part $P_m(\xi) = \sum_{|\alpha|=m} a_\alpha \xi^\alpha$ is invertible with $||P_m(\xi)^{-1}|| \le C|\xi|^{-m}$ when $\xi \ne 0$, or **equivalently** we can define $P(\xi) = \sum_{|\alpha|\le m} a_\alpha \xi^\alpha$ (with not just the top order part) with $||P(\xi)^{-1}|| \le C(1 + |\xi|)^{-m}$.

The problem is that with this definition, we can't quite just divide by $P(\xi)$ near zero like we did with $\Delta + 1$, so we can't always find a $u \in \mathcal{S}$ such that $Pu = f$ for any $f \in \mathcal{S}$. Instead, we can "not quite, but almost" do this, and this idea will lead into **Fredholm theory** next time. Basically, there will only be a finite codimensional space where things are bad (that is, the intersection of zerosets of finitely many zero linear functionals), and it turns out we do have very clean checkable linear conditions for when we can solve the equation.

Proposition 106

Suppose $Pu = f$ with P elliptic of order m with constant coefficients, and suppose that $f|_U \in H^s_{\text{loc}}$. Then $u|_U \in H^{s+m}_{\text{loc}}$ (so the local regularity is correlated).

Proof. This time we can't directly write u in terms of f by dividing, but we can consider again

$$P(\chi_1 u) = \chi_1 f + [P, \chi_1]u$$

where $[P, \chi_1]$ is of order $(m-1)$. So it suffices to solve $Pv = h$ for $v, h \in \mathcal{E}'$ (recall that this denotes the compactly supported distributions), and we'll do so by defining the bounded operator $G : H^t \to H^{t+m}$

$$Gh = \mathcal{F}^{-1}\left(\frac{\hat{h}}{P(\xi)}\chi_2(\xi)\right)$$

(G maps between these spaces because the $P(\xi)$ in the denominator cancels out factors of $(1+|\xi|^2)$ in the Sobolev approximation), and χ_2 is chosen to be supported away from the compact set on which $P(\xi) = 0$. So now everywhere in the support of χ_2, P is invertible, and then

$$P(Gh) = \int e^{ix\xi}\frac{\hat{h}(\xi)}{P(\xi)}\chi(\xi)P(\xi)d\xi$$

since hitting the exponential factor gives us a factor of ξ for each derivative so we get a $P(\xi)$ everywhere. So $P(Gh) = \mathcal{F}^{-1}(\chi_2 \cdot h)$, and now if we write $\chi\hat{h} = \hat{h} - (1-\chi)\hat{h}$ we find that

$$(g\hat{h}) = h - \mathcal{F}^{-1}((1-\chi_2)\hat{h}) = h + Rh,$$

with $Rh \in C^\infty$. Then this R is a distribution because $(1-\lambda)\hat{h}$ is compactly supported. So now if we know that $\chi_1 u \in H^t$ and $\chi_1 f \in H^s$ for $t \ll s$ (so we don't know much about u yet), then $h = P(\chi u) = \chi_1 f + [P, \chi_1]u$, where this sum is in H^{t-m+1}. Dividing by $P(\xi)$ makes the decaying of h better by m orders, so in fact $G(h)$ is in H^{t+1} and thus $P(G(h))$ is in $h + C^\infty$. So if $w = G(h)$, then $P(xu - w) = Rh \in C^\infty$, and χu is compactly supported. But now have an identity of the time $P \circ G = I - R$, so up to this smooth error term R we know how to solve our equation – unfortunately we want information about $\chi u - w$ itself. So the trick now is to think about things algebraically – we really want P and G in the opposite order

$$G(P(\chi u)) = GP\mathcal{F}^{-1}(\widehat{\chi_1 u}) = G\mathcal{F}^{-1}(P(\xi)\widehat{\chi_1 u}) = \mathcal{F}^{-1}\left(P(\xi)\widehat{\chi_1 u}\frac{\chi_2(\xi)}{P(\xi)}\right),$$

and again we have the nice cancellation so that

$$G(P(\chi u)) = \mathcal{F}^{-1}(\hat{\chi_1} u \chi_2(\xi)).$$

In other words, we have the same observation as before where we can write the right-hand side as $v - Rv$, where $Rv \in C^\infty$. So overall we have the identities

$$P \circ G = I - R_1, \quad G \circ P = I - R_2, \quad R_1, R_2 : \mathcal{E}'(\mathbb{R}^n) \to C^\infty,$$

so if we take $P(\chi u) = \chi f + [P, \chi]u$, where the right-hand side is in H^{t-m+1}, and apply G to both sides, we get

$$\chi u = G(\hat{f}) + R_2(\chi u).$$

So again assuming that we're in H^t locally, we end up finding that we're actually in H^{t+1} for any cutoff function. Iterating this procedure again, we run into trouble when we get to $t > s$ again, and thus we can only conclude that $\chi u \in H^{s+m}$, which is what we wanted to prove. $\square$

So the point is that we've gotten information about regularity, even though this time we haven't actually found an exact formula for the solution. Instead, we found solutions up to these operators R_1, R_2, which are bounded operators but potentially with very large norm. So we're not trying to invert $P(\chi u) = f$ – we're just trying to get regularity statements and the R_is are completely invisible.

Corollary 107

If P is an elliptic operator and $Pu = 0$, then u is smooth.

In particular, we see that (letting $P = \Delta$) any harmonic function is smooth, and (letting $P = \frac{\partial}{\partial \overline{z}} = \frac{1}{2}(\partial_x + i\partial_y)$, with symbol $\frac{1}{2}(\xi + i\eta)$) if $\frac{\partial u}{\partial \overline{z}} = 0$, then $u \in C^\infty$ (so any holomorphic function is smooth). But on the other hand if we try looking at an operator like $\Box = \partial_t^2 - \partial_x^2$ in $\mathbb{R}^2$ (the wave operator), the symbol is now $-\tau^2 + \xi^2$ and this is not elliptic – in particular the zeros of this thing are not on a compact set, and then it's no longer true that $\Box u = 0$ implies that u is smooth.

Our next result is meant to be a "cautionary tale" (since Grothendieck had his theory of nuclear Frechet spaces long ago and took a lot of work to prove this, but the point is that the result is almost a tautology in certain aspects and can be proved quickly with the right mindset):

Theorem 108 (Schwartz kernel theorem)

Suppose $A : C_0^\infty(\Omega_1) \to \mathcal{D}'(\Omega_2)$ is a continuous operator (meaning that for any compact set K, and for every level of regularity m, the image is continuous when paired with a test function – that is, there is some C such that $|\langle A\phi, \psi \rangle| \le C||\phi||_{C^m}$). Then there is a distribution $\mathcal{K}_A \in \mathcal{D}'(\Omega_1 \times \Omega_2)$ such that $\langle A\phi, \psi \rangle = \langle \mathcal{K}_A, \phi \otimes \psi \rangle$.

Here, we define the tensor product by saying that if $\phi \in C_0^\infty(\Omega_1)$ and $\psi \in C_0^\infty(\Omega_2)$, then $\phi \otimes \psi \in C_0^\infty(\Omega_1 \times \Omega_2)$ evaluates to $\phi(x)\psi(y)$ at (x, y). And in fact we have $C_0^\infty(\Omega_1) \otimes C_0^\infty(\Omega_2) \subset C_0^\infty(\Omega_1 \times \Omega_2)$, where we're taking the tensor product in the usual algebraic sense. But then if we take the completion we end up getting all of $C_0^\infty(\Omega_1 \times \Omega_2)$.

So we can write this Schwartz kernel as basically wanting to find a $\mathcal{K}_A$ such that

$$\int \mathcal{K}_A(y, x)\phi(x)\psi(y)dxdy = \int (A\phi)(y)\psi(y)dy$$

if we have ordinary functions everywhere. So now if we send in two test functions ϕ, ψ and sending them to the complex number $\langle A\phi, \psi \rangle$, we've defined a distribution on the product space, and this is in fact bilinear – it's continuous because A is continuous. So the point is that we've actually defined our distribution on $C_0^\infty(\Omega_1) \otimes C_0^\infty(\Omega_2)$, and we're claiming it's a continuous linear functional on $C_0^\infty(\Omega_1 \times \Omega_2)$. So the two ingredients are that the former is dense in the latter, and we have a bilinear functional which is separately continuous that we want to prove is jointly continuous. And this last part is true by a corollary of uniform boundedness principle (so the Baire category theorem comes up again).

This result is useful because we now see that any operator A (a big infinite-dimensional matrix) can be thought of as a distribution $\mathcal{K}_A \in \mathcal{D}'$, so there are nice distributions that correspond to operators (or equivalently the Schwartz kernels of many operators turn out to be nice distributions, such as the delta function or the G in the proof we did earlier).

Our next topic will be the **weak and weak* topology**, and we'll now return to Buhler and Salamon. Throughout all of this, X will be a real normed vector space, meaning that we do not demand completeness but do have a norm $|| \cdot ||$. Recall that the **strong topology** corresponds to the set of "strongly open" sets $\mathcal{U}_S$ generated by the norm, while the **weak topology** corresponds to the much smaller set of open sets $\mathcal{U}_w$ enough to make any linear functional continuous and thus generated (as a subbase, meaning we're allowed finite intersections) by elements of the form

$\ell^{-1}((a, b))$ (infinite-dimensional slabs). In particular, we have weak convergence

$$x_j \xrightarrow{w} x \text{ if } \ell(x_j) \to \ell(x) \quad \forall \ell \in X^*,$$

though the rate of convergence may depend on ℓ. And finally, we have the **weak* topology** $\mathcal{U}_{W^*}$, which is a topology on X^* (which is always a Banach space and thus has a norm and dual space). As usual, we can define the strong and weak topology on X^*. But additionally, because X sits inside X^{**} (since the evaluation-at-x map $x(\ell) = \ell(x)$ gives us an element of X^{**} for any X), we can specify that we we don't need all linear functionals of X^* to be continuous, only the ones corresponding to these evaluation maps – that yields the weak* topology. And the following fact is why we care about this:

Theorem 109 (Banach-Alaoglu)
Let X be a real Banach space. The unit ball $B^* \subset X^*$ is weak*-compact.

We've talked about how infinite-dimensional balls are always badly noncompact in the strong topology, but in a weak enough topology we can get compactness. And we'll see that this is in fact quite useful.

Corollary 110
In any Hilbert space (which is isomorphic to its dual so that $X = X^* = X^{**}$), the unit ball is weakly compact.

Example 111
If we take the unit ball in ℓ^2 and consider the sequence $e_1, e_2, e_3, \cdots$, this sequence does not converge strongly but converges weakly to zero because $\langle e_i, x\rangle = x_j$ must go to zero because $\sum x_j^2 < \infty$. But the rate at which things go to zero does depend on the x that we choose.

(And remember that in Hilbert spaces and Banach spaces, sequential compactness is the same as compactness, even in the weak or weak* topology.)

11 February 14, 2023

Last lecture, we were discussing the weak*-topology, which is a topology on X^* given by $U_{w,*}$ (which is smaller than U_w, which is smaller than U_s – all of these topologies exist on any space which happens to be the dual of some other space). The idea is that we are only looking at the linear functionals coming from the evaluation maps coming from X, not all of X^{**}, and the $U_{w,*} = U_w$ if and only if $X^{**} = X$ (meaning that X is reflexive).

Proposition 112
Let X be a real normed space and $K \subset X$ be convex. Then K is closed (with respect to the strong topology) if and only if it is weakly closed.

Proof. For the forward direction, suppose K is closed. Then for any point $x_0 \notin K$, there is some $\delta > 0$ such that $B_\delta(x_0)$ is disjoint from K. Then by geometric Hahn-Banach, these two convex sets are separated by some $\ell \in X^*$, so $\ell(x_0) > c \geq \ell(y)$ for all $y \in K$, so x_0 is in some weak neighborhood disjoint from K, namely the set of points where $\ell(x) > c$. Conversely, suppose K is weakly closed. Then $X \setminus K$ is weakly open, but then any weakly open set is also in the strong topology, so it must also be strongly open. Thus K is indeed (strongly) closed. □

Proposition 113 (Mazur)

Suppose X is a real normed space, $x_i \in X$ is a sequence, and $x_i \xrightarrow{w} \overline{x}$ converges weakly. Let $\text{conv}(\{x_i\})$ be the **convex hull** of the x_is; that is, the set of all linear combinations $\sum_{j=1}^N \lambda_j x_j$ such that $\sum \lambda_j = 1$ and $\lambda_j \geq 0$. Then $\overline{x} \in \overline{\text{conv}(\{x_i\})}$.

In other words, $\overline{x}$ lies in the strong closure of the convex hull, so there is a convex combination of the x_is which converge strongly to x.

Proof. The convex hull is clearly convex, and the closure of a convex set is also convex. So $\overline{\text{conv}(\{x_i\})}$ is convex and closed, meaning that it is also **weakly closed** by Proposition 112. So if $x_i \in \overline{\text{conv}(\{x_i\})}$ converge to $\overline{x}$ weakly, the definition of being closed means that $\overline{x}$ is also in that set. □

(Note that the convex hull is not closed in general – for example consider the basis vectors $e_1, e_2, \cdots$.)

Lemma 114

Suppose our space is infinite-dimensional. Let $S = \{x : ||x|| = 1\}$ be the unit sphere and $B = \{x : ||x|| \leq 1\}$ be the unit ball. Then B is the weak closure of S, denoted $\overline{S}^w$.

Proof. Since B is strongly closed and convex, it is weakly closed, and it contains $\overline{S}^w$ by our previous result. For the other direction, take any $x_0 \in B$ (the case where $||x_0|| = 1$ is already clear, so we can assume $||x_0|| < 1$), and let U be some weak open set containing x_0. Then there must be some $\ell_1, \cdots, \ell_n \in X^*$ such that the set $\bigcap_{i=1}^N \{|\ell_i(x - x_0)| < \varepsilon\}$ is contained in U. But there is some $\xi \in X$ such that $\ell_i(\xi) = 0$ for all i (since there's only finitely many codimension-1 conditions to worry about). Then because $||x_0|| < 1$, we can find some t such that $x_0 + t\xi \in S$, but we have $\ell_i(x_0) + t\xi) = \ell_i(x_0)$ so we are still inside the open set and intersecting S. Thus any open set must intersect S and thus x_0 must be in $\overline{S}^w$. □

Corollary 115

Let X be infinite-dimensional. Passing to the dual space, we have $B^* = \overline{S^*}^{w^*}$.

(This is different, because we're taking the weak* closure instead of the corresponding weak closure for the dual space.)

Proof. For any $x \in S$, consider the set of functionals

$$F_x = \{x^* \in X^* : \langle x^*, x \rangle \leq 1\}.$$

This F_x is weak*-closed for all x, and now we can define $B^* = \bigcap_{x \in S} F_x$. (The idea is to look at the intersection of all half-spaces as we go around the unit sphere, if we're in something like a Hilbert space.) If we let $K = \overline{S^*}^{w^*}$, then $K \subset B^*$ (since S^* lies inside it and K is weak*-closed) and $B^* \subset K$ (by the same proof as before). □

Theorem 116

Let X be a separable Banach space, and let $\{\ell_j\}$ be a bounded sequence in X^*. Then there exists a weak*-convergent subsequence. (In other words, any closed bounded set is weak*-compact.)

Proof. Choose $\{x_\ell\}$ to be a countable dense subset of X (by separability), and do the following. First consider $\ell_1(x_1), \ell_2(x_1), \cdots$, which is a bounded sequence of numbers and thus has a convergent subsequence. We can then refine that subsequence to a smaller subsequence by considering $\ell_1(x_2), \ell_2(x_2), \cdots$ and so on. Then taking the diagonal elements yield a subsequence ℓ'_j such that for any i, $\ell_{j'}(x_i)$ converge, and by (a simple case of) uniform boundedness and density of x_is that means the $\ell_{j'}$s converge everywhere. □

We can now discuss the general Banach-Alaoglu theorem:

Proof of Theorem 109. For any $x \in X$, let $I_x = [-||x||, ||x||]$, and consider the space

$$\mathcal{I} = \prod_{x \in X} I_x.$$

An element f in here is precisely a choice of element $f(x) \in I_x$ for some $x \in X$ where $|f(x)| \leq ||x||$ (so somehow we are building in the norms this way). Now let $\mathcal{L} \cap \mathcal{I}$ be the subspace of **linear** functions; we see that in fact this is B^* itself. Since each I_x is compact, Tychonoff's theorem tells us that $\mathcal{I}$ is compact, so now we just need to show that B^* is closed in $\mathcal{I}$. A proof of Tychonoff's theorem can be found in Munkres or elsewhere, but in particular it uses the axiom of choice . (Here, the topology on $\mathcal{I}$ is generated in the following way: consider $\pi_X : \mathcal{I} \to \mathcal{I}_x$, let $x_1, \cdots, x_N$ be points, and let $U_j \subset I_{x_j}$ for each j. Then the topology is generated by the sets $\bigcap_{j=1}^N \pi_{x_j}^{-1}(U_j)$. So we trap finitely many of the coordinates and demand that the entries live in some given open set, and in fact this is exactly the weak* topology that we wanted.)

So to show that $\mathcal{L} \cap \mathcal{I}$ is weak*-closed, consider the maps

$$\phi_{x,y}(f) = f(x+y) - f(x) - f(y), \quad \psi_{x,x}(f) = f(\lambda x) - \lambda f(x),$$

which are all maps $\mathcal{L} \to \mathbb{R}$. Then $\mathcal{L}$ is exactly the intersection of the kernels of all $\phi_{x,y}$s and $\psi_{x,x}$s (since those are the conditions for being a linear operator); since each of those sets is closed so is their intersection. □

So in any reflexive space, the unit ball is weakly compact, and this is the form in which Banach-Alaoglu is often used. And the converse also turns out to hold.

Our next example is going to be a way to bring these big results together:

Example 117
Let $\Omega \subset \mathbb{R}^2$ be smoothly bounded, and consider the Sobolev completion $H_0^1(\Omega) = \overline{C_0^\infty}^{H^1}$.

We know that the set of functions $u \in H^1(\mathbb{R}^2)$ where $u = 0$ on the region $\mathbb{R}^2 \setminus \overline{\Omega}$ is a closed subspace, so $H_0^1(\Omega)$ is a closed subset of H^1 functions which vanish on the boundary of Ω – in fact, as long as the boundary is smooth, we will have $H_0^1(\Omega) = \{u \in H_0^1(\mathbb{R}^n) : u_{\mathbb{R}^n \setminus \overline{\Omega}} = 0\}$. Now define the **Rayleigh quotient**

$$Q(u) = \frac{\int |\nabla u|^2}{\int u^2}.$$

We want to show that this will lead us to a solution $\Delta u = \lambda u$ in Ω with $u|_{\partial\Omega} = 0$. Indeed, $Q(u) \geq 0$ for all u, and thus there is some infimum $\inf_{||u||=1} Q(u) = \mu_0 \geq 0$. We wish to show that the minimum is actually attained, meaning that there is some u_0 with $Q(u_0) = \mu_0$. (In particular, this will also show us that $\mu_0 > 0$.)

This is basically the direct method of the calculus of variations: pick any sequence u_j with $||u_j||_{L^2} = 1$ (this normalization is allowed because the definition of $Q(u)$ is homogeneous) and with $Q(u_j) \to \mu_0$. In particular, we see that by choosing j large enough so that the Rayleigh quotient is at most $\mu_0 + 1$, we have

$$||u_j||_{H_0^1}^2 = \int_\Omega u_j^2 + |\nabla u_j|^2 = 1 + (\mu_0 + 1) \le C,$$

so we have a bounded sequence in the H_0^1 norm. We can thus find a subsequence $\{v_j\}$ converging strongly in L^2 (by the Rellich theorem, since the inclusion into a lower Sobolev space is compact). Now choose a further subsequence $\{w_j\}$ using Banach-Aloaglu, so that w_j converges weakly to some u_0 on H_0^1. In particular, since we first found something converging strongly in L^2, the two convergence points must agree and u_0 is the limit in both cases.

Now we claim that $Q(u_0)$ achieves the infimum of the Rayleigh quotient. We cannot say that Q is continuous in H_0^1 norm, so instead we use Mazur's theorem to write a sequence of convex combinations $U_N = \sum_{j=1}^N \lambda_j^N u_j$ with $\sum \lambda_j = 1$ and all $\lambda_j \ge 0$, such that $U_N \to u_0$ converges strongly in H_0^1. In particular, we don't need to take convex combinations starting from the beginning of the sequence – we may choose to take them further and further out, specifically choosing U_N so that each U_N only involves U_j with $j \ge \frac{N}{2}$. But then because the convex combinations of our points u_j get closer to u_0 in L^2, the U_Ns must also converge in L^2 to u_0, and now the problem is to work with the $Q(u)$s. Now think about the quadratic functional $\hat{Q}(u) = \int |\nabla u|^2$, which we can think of as a graph over H_0^1. Since this is convex, for any $v_0, v_1 \in H_0^1$ and for any $0 < \lambda < 1$, we claim that $\hat{Q}(v) \le (1-\lambda)\hat{Q}(v_0) + \lambda \hat{Q}(v_1)$. Indeed,

$$\hat{Q}((1-\lambda)v_0 + \lambda v_1) = (1-\lambda)^2 \hat{Q}(v_0) + 2\lambda(1-\lambda) \int \nabla v_0 \nabla v_1 + \lambda^2 \hat{Q}(v_1),$$

and by Cauchy-Schwarz this middle term is bounded, since $2\int \nabla v_0 \nabla v_1 \le \int |\nabla v_0|^2 + \int |\nabla v_1|^2$:

$$\le ((1-\lambda)^2 + \lambda(1-\lambda))\hat{Q}(v_0) + (\lambda^2 + \lambda(1-\lambda))\hat{Q}(v_1) = (1-\lambda)Q(v_0) + \lambda Q(v_1).$$

By induction we can extend this to a convex combination of any finite number of terms. Then since **(1)** $||U_N||_{L^2} \to 1$, **(2)** each $\hat{Q}(U_N)$ is bounded from above by the corresponding $\sum \lambda_j Q(u_j)$, and **(3)** the values of $Q(u_j)$ converge to μ_0, the convex combinations also converge and we have $\hat{Q}(U_N) \to \mu_0$ (remember that because μ_0 is the infimum we can't get any lower).

(So summarizing what we did, we got compactness from the Rellich theorem, then we used Banach-Alaoglu to show weak convergence, and then we use Mazur's theorem to show that a convex combination converges strongly, and finally we used convexity to show that convex combinations behave well. So we did get strong convergence $U_N \to u_0$ in H_0^1, and we achieved the minimum $Q(u_0) = \mu_0$.)

So now we know that $Q(u) \ge Q(u_0)$ for all nonzero $u \in H_0^1$ – in particular, for any $v \in H_0^1$ and any $t \in \mathbb{R}$, we know that $Q(u_0 + tv)$ reaches a local minimum at $t = 0$. Expanding this out, we have

$$Q(u_0 + tv) = \frac{\int |\nabla u_0|^2 + 2t\nabla u_0 \cdot \nabla v + t^2 |\nabla v|^2}{\int u_0^2 + 2tu_0 v + t^2 v^2}.$$

Taking the derivative at $t = 0$, we find that this evaluates to $2\int \nabla u_0 \cdot \nabla v - \int |\nabla u_0|^2 \cdot 2\int u_0 v$, and that must be zero. So our minimizer must satisfy (for all v) $\int \nabla u_0 \cdot \nabla v = \int |\nabla u_0|^2 \int u_0 v$, and in particular this is true for all $v \in C_0^\infty(\Omega)$. Since u_0 can be thought of as a distribution (with the above identity holding for any test function v), and any distribution lives in some Sobolev space H^t, that means $\int \nabla u_0 \cdot \nabla v - \mu_0 u_0 v = 0$ for all $v \in H_0^1$. Thus, for all test functions v we have (by integration by parts, moving one derivative over on the first term)

$$\int (\Delta u_0 - \mu_0 u_0) v = 0 \implies -\Delta u_0 - \mu_0 u_0 = 0 \text{ in } \mathcal{D}'(\Omega).$$

But u_0 is supported in Ω, so in fact $\Delta u_0 = -\mu_0 u_0$ as a statement in H^{t-2}. And now we can go back to everything we did last week – localizing inside Ω, this is a constant-coefficient elliptic equation if we think about it in the form

$(\Delta + \mu_0)u_0 = 0$. So we see that in fact u_0 is smooth in the interior of Ω (if we localize by any cutoff function, we get a C^∞ function).

In other words, we've found a function satisfying the eigenvalue equation, and this is the **generalization of Fourier series** in an arbitrary domain. Applying this same argument for $u_0, u_1, u_2, \cdots \in H_0^1(\overline{\Omega}) \cap C^\infty(\Omega)$, in fact we get a minimizer equation where u_j achieves the infimum of $Q(u)$ over all u perpendicular to $u_0, u_1, \cdots, u_{j-1}$. Then $u_0 \perp u_1 \perp u_2 \perp \cdots$ gives us an orthonormal set of eigenfunctions with span dense in L^2, so for any $f \in L^2(\Omega)$ there are $a_j \in \mathbb{C}$ with $f = \sum a_j u_j$. This is relevant to spectral geometry and other fields, and we'll say a bit more about this next time!

12 February 16, 2023

Last time, we saw an important application of various big theorems we've proven in this class to constructing Fourier series on an arbitrary domain. There's one more application where we consider a PDE in a bounded domain which relies on the Lax-Milgram theorem – its statement is close to the Riesz representation theorem:

Theorem 118 (Lax-Milgram)
Let $B(x, y)$ be a bilinear form on a Hilbert space X, where $|B(x, y)| \le C||x||\,||y||$ and $c||x||^2 \le |B(x, x)|$ (this lower bound condition is called **coercivity**). Then for any $\ell \in X^*$ (which we can identify with X), there is some y_ℓ such that $\ell(x) = B(x, y_\ell)$.

In other words, fixing y gives us a bounded linear functional, and coercivity tells us that any such linear functional can be represented in terms of this bilinear form. (So this is Riesz representation if we have the standard bilinear form.)

Proof. The map $x \mapsto B(x, y)$ is in X^* (that is, it is a bounded linear functional), meaning that (by Riesz representation) there exists some $z_y \in X$ (depending on y) such that $B(x, y) = \langle x, z_y \rangle$. The map $T : y \mapsto z_y$ is linear, and we are claiming that T is invertible.

Indeed, notice that the range of T is closed, since

$$c_0||y||^2 \le B(y, y) = \langle y, T(y) \rangle \le ||y|| \cdot ||Ty||$$

for some constant c_0 by assumption, meaning that $\boxed{||Ty|| \ge c_0||y||}$. We claim that such an inequality automatically means the range is closed. Indeed, if we have a sequence of points (y_0, Ty_0) with $Ty_j \to \overline{z}$, we wish to show $\overline{z}$ is also in the range. But this is true because Ty_j is Cauchy and thus y_j must also be Cauchy by the boxed inequality, hence convergent. And the point is that this "bounding from below" inequality means T cannot have a nullspace, and thus $\ker(T) = 0$.

Now we can get the result we want by applying the open mapping theorem. To do that, we must show T is surjective – in fact, it suffices to show that it's dense, since we've just shown it's closed. Suppose not – then there would be some w perpendicular to the range of T by Hilbert space theory. But then $B(w, y) = \langle w, T(y) \rangle = 0$ for all y, so in particular $B(w, w) = 0$, but it is also at least $c||w||^2$ so we must have $w = 0$. So T^{-1} is also a bounded linear operator, so if ℓ corresponds to z under the Hilbert space dual identification, then we can plug in $y_\ell = T^{-1}(z)$ to finish the proof. $\square$

We now want to do something similar to the applications we've seen in previous classes, starting with a simple example:

Example 119

Let $\Omega \subseteq \mathbb{R}^n$ be a smoothly bounded domain, and define the functional $H_0^1 \times H_0^1 \to \mathbb{C}$ given by

$$B(u, v) = \int |\nabla u||\nabla v|.$$

Our goal is to use this to solve $\Delta u = f$.

We wish to check that this is coercive, and indeed we have $|B(u, v)| \le ||\nabla u|| \cdot ||\nabla v|| \le ||u||_{H_0^1} ||v||_{H_0^1}$, so we already have one side of the necessary inequalities. And now we want to show that $B(u, u) = \int |\nabla u|^2 \ge c \int |u|^2 + |\nabla u|^2$ – essentially the model case we want to avoid is the one where u is constant, but anything in $H_0^1 \times H_0^1$ is zero on the boundary so that case can't happen. To be more careful with this claim, we can make the following argument:

Theorem 120 (Crudest form of Poincaré inequality)

There is some $c > 0$ such that $c \int |u|^2 \le \int |\nabla u|^2$ for all $u \in M_0^1$.

Proof. Single out the x_1 direction, fix some point $x \in \Omega$, and draw a horizontal line in the $-x_1$ direction until we hit some first contact point $x_{1,d}$. Then because u is zero on the boundary,

$$u(x) = \int_{x_{1,d}}^{x_1} \partial_{x_1} u(t, x') dt \implies |u(x)|^2 \le \int_{x_{1,d}}^{x_1} \partial_{x_1} u(t, x') dt,$$

and now if we apply Cauchy-Schwarz and average over the whole domain, we get

$$\int_\Omega |u(x)|^2 \le \int_{x_{1,d}}^{x_1} |\partial_{x_1} u|^2 dt \le d \int_\Omega |\nabla u|^2,$$

where d is the diameter along the ray that we described. We can now recall from last time the definition of the Rayleigh quotient $Q(u) = \frac{\int |\nabla u|^2}{\int u^2}$ for a bounded domain; in particular we have $\inf(Q) = \mu_0 > 0$ and that infimum is attained. But that means that we get a bound on what the "best diameter" can be, and in particular taking $c = \frac{1}{\mu_0}$ will work. $\square$

The point now is that the bilinear form $H_0^1 \times H_0^1 \to \mathbb{C}$ is a valid bilinear form for applying Lax-Milgram – thus for all $\ell \in X^*$, there is some $y_\ell \in X$ such that $\ell(x) = B(x, y_\ell)$. So for any $f \in L^2$, we can consider the linear functional $\ell(v) = \int vf$, and there will be some u such that for all v we in fact have $\ell(v) = B(v, u)$. In other words, for any $f \in L^2$ there is some $u \in M_0^1$ such that $\int \nabla u \nabla v = \int fv$ for all $v \in M_0^1$. Letting $v \in C_0^\infty$ and integrating one more time, we find that $-\Delta u = f$. In "physics words," this means that we have managed to find a potential corresponding to a given charge density.

Furthermore, we can make some regularity statements: since $u \in H_0^1(\Omega) \subset H^1(\mathbb{R}^n)$, and $f \in L^2(\Omega) \subset L^2(\mathbb{R}^n)$, we see that $u \in H^2(\Omega) \cap H_0^1(\Omega)$ anywhere on the interior. But Lax-Milgram allowed us to take any ℓ in the dual of H_0^1, not H_0^1 itself, and with Sobolev spaces there's a duality between H^s and H^{-s}.

Fact 121

We have $H_0^1(\Omega)^* = H^{-1}(\Omega) = \{u_\Omega : u \in H^{-1}(\mathbb{R}^n)\}$.

So in particular we can take $f \in H^{-1}$ and take the pairing $\ell(v) = \int vf$, and $-\Delta u = f$ will be true in the distributional sense. Knowing that $f \in H^{-1}$ only gets us so much regularity for u (being $H^2(\Omega) \cap H_0^1(\Omega)$), but we also know that if f is smooth (or H^{25}, etc.), then u must be smooth (or H^{27}, and so on). The point is that this result is really following from Riesz representation rather abstractly.

Example 122
All of the PDEs we've chosen to solve have been of a simple form, but here's a method of approaching these things more generally: consider the operator

$$Lu = -\sum \frac{\partial}{\partial x_i} a_{ij}(x) \frac{\partial u}{\partial x_j} + b_i \frac{\partial u}{\partial x_i} + cu$$

for $u \in H_0^1$.

If we pair this against a function v, then

$$\langle Lu, v\rangle = \sum \int a_{ij}(x) \frac{\partial u}{\partial x_i}\frac{\partial v}{\partial x_j} + b_i \frac{\partial u}{\partial x_i} v + cuv,$$

meaning the upper bound for using Lax-Milgram is easy in these cases. For the lower bound, we need **positive definiteness** of the a_{ij}s, meaning that the lower eigenvalue is strictly bounded from below. But even that is not enough – if we try to define $B(u, v) = \sum \int_\Omega a_{ij}(x) \frac{\partial u}{\partial x_i}\frac{\partial v}{\partial x_j} + b_i \frac{\partial u}{\partial x_i} v + cuv$, the upper bound works, but we will not get the lower bound immediately. Instead (assuming all of our coefficients are smooth), the first term is bounded by $c||u||_{H_0^1}^2$ from below by positive definiteness. Then by using the fact that $\frac{\partial u}{\partial x_i} u \le \varepsilon \left(\frac{\partial u}{\partial x_i}\right)^2 + \frac{1}{4\varepsilon} u^2$, we can choose ε small enough to get an inequality of the form

$$B(u, u) \ge c_1 ||u||_{H_0^1}^2 - c_2 ||u||_{L^2}^2$$

for some constants $c_1 > 0$ and c_2, and thus we don't quite get coercivity in the form that we want. So instead we will modify our problem: we solve the equation $Lu + \lambda u = f$ (and $u|_{\partial\Omega} = 0$) for some sufficiently large λ canceling out the c_2 factor, so that for any $f \in H^{-1}$ there is some $u \in H_0^1$ such that

$$a_{ij}(x) \frac{\partial u}{\partial x_i}\frac{\partial \phi}{\partial x_i} + b_i(x) \frac{\partial u}{\partial x_i}\phi + (c(x) + \lambda) u\phi = \int f\phi$$

for any test function and thus every H_0^1 ϕ. Then once again we find that we can solve this general, slightly modified, equation using Lax-Milgram. (We'll see next week why we care about adding this factor λu and how to get rid of it.)

With that in mind, we'll now move onto a new chapter – this next part of our quarter will discuss the classes of **compact operators** and **Fredholm operators**. There is a nice structure theory behind them, where we should think of compact operators as being "morally zero" and Fredholm operators as "morally invertible." One example we've shown (in our homework) is that adding a small operator (in norm) to an invertible A still gives us something invertible, but being a compact operator (such as an operator whose range is finite-dimensional) is a much more robust way of quantifying being "morally negligible." So motivated by this, Fredholm operators will be of the form $I + B$ for some "negligible" B.

Remark 123. *Let X be Banach, $\mathcal{X}$ be linear functionals on X, and let $\mathcal{K}$ be the set of compact operators and $\mathcal{F}(X)$ be the set of Fredholm operators. It will turn out $\mathcal{F}(X)$ is open in $\mathcal{B}(X)$, and $\mathcal{K}(X)$ is closed in it, and we can show that compositions and combinations of these classes also work well (the terminology is that* **compact operators form an ideal** *in the set of operators). In particular, $\mathcal{C}(X) = \mathcal{B}(X)/\mathcal{K}(X)$ is then also a valid space, and $\mathcal{F}(X)$ is the set of invertible operators of that space.*

So we'll be proving all of these things now, but first we'll give some definitions:

Definition 124

A bounded linear operator K is **compact** if for any bounded sequence $\{x_i\} \in X$, the sequence $\{K(x_i)\}$ has a convergent subsequence.

In other words, an operator is compact if (letting B denote the unit ball) $\overline{K}(B)$ is compact – another phrasing is that K takes bounded sets to precompact sets. But note that if X is not reflexive, there may be some issues with sequential compactness versus compactness.

Definition 125

A bounded linear operator A is **Fredholm** if **(1)** the kernel of A is finite-dimensional, **(2)** the range of A is closed, and **(3)** $X/\text{ran}(A)$ is finite-dimensional.

The picture to have in mind is the following: we start with some linear operator $X \to X$, where the kernel and cokernel are both finite-dimensional. Then being Fredholm means that $\ker A$ and $\text{coker } A$ are finite-dimensional, and on the remaining spaces $(\ker A)^\perp \to \text{Ran}(A)$ we have an isomorphism (so once we excise a finite-dimensional piece from the domain and range, we get an isomorphism). And this secretly uses the fact that a bijective map from $(\ker A)^\perp \to \text{Ran}(A)$ must have a bounded inverse as well.

Definition 126

Let $A : X \to Y$ be a linear operator. Define the operator $A^* : Y^* \to X^*$ via the formula

$$(A^* y^*)(x) = y^* Ax.$$

There are a few important properties of this new operator we've constructed:

Lemma 127

If A is bounded, then so is A^*, and $||A|| = ||A^*||$. Also, $(AB)^* = B^* A^*$, and $(\text{id}_X)^* = \text{id}_{X^*}$.

Proof. By definition we have

$$||A^*|| = \sup_{||y^*|| \neq 0} \frac{||A^* y^*||}{||y^*||} = \sup_{x, y \neq 0} \frac{\langle A^* y^*, x\rangle}{||y^*|| \cdot ||x||}$$

(that is, we can check the norm against all unit vectors) and now by definition of A^* this is

$$= \sup_{x, y \neq 0} \frac{\langle y^*, Ax\rangle}{||y^*|| ||x||} = \sup_{x \neq 0} \frac{||Ax||}{||x||} = ||A||.$$

The other two statements follow more directly by manipulation. $\square$

Proposition 128

Let $A \in \mathcal{B}(X, Y)$. Then we have the following:

- We have $\{y^* \in Y^* : y^*|_{\text{Ran}(A)} = 0\} = \ker(A^*)$.
- A^* is injective if and only if the range of A is dense in Y.
- A is injective if and only if A^* has weak*-dense range.

Proof. These basically follow directly from the definition: if y^* annihilates the range, then $0 = y^*(Ax)$, meaning that we always have $(A^*y^*)(x) = 0$, so A^*y^* must be zero. The other two are similarly trivial. □

Theorem 129

The following are equivalent:

1. $\text{Ran}(A) = {}^\perp \ker(A^*)$,
2. The range of A is closed,
3. There is some $C > 0$ such that $\inf_{A\xi=0} ||x + \xi|| \le C||Ax||$,
4. $\text{Ran}(A^*) = (\ker A)^\perp$,
5. The range of A^* is weak*-closed in X^*,
6. The range of A^* is closed,
7. There is some $C > 0$ such that $\inf_{A^*\eta^*=0} ||y^* + \eta^*|| \le C||A^*y^*||$.

The main fact here is that closure of the range is equivalent for A and for A^*, and everything else is kind of alternate statements and properties of closed range.

First we explain what ${}^\perp \ker A^*$ is: recall that $\ker(A^*) = \{y^* \in^* : A^*y^* = 0\}$, so ${}^\perp \ker A^*$ is the orthogonal space in a "weak sense:" it's the set of all $y \in Y$ such that $\langle y^*, y\rangle = 0$ for all y^* in the nullspace of A^*.

Proof. (1) implies (2) because the common zeroset of linear functionals is always closed. (2) implies (3) comes from the fact that $X_0 = X/\ker(A)$ is a new Banach space in which $||[x]|| = \inf_{\xi:A\xi=0} ||x + \xi||$ for any equivalence class $[x]$. Then A maps X_0 to $Y_0 = \text{Range}(A)$ (since we're assuming the range is closed, so we can mod out the nullspace and restrict the target to just the range – these are both Banach spaces and bijective, so we have a boundedly invertible map by the open mapping theorem. Thus we have some estimate of the form $C||x|| \le ||Ax||$ by boundedness of A^{-1}.

Next, for (3) implies (4) (passing to the adjoint), we know that the range of A^* sits inside $(\ker A)^\perp$, since $\langle A^*y^*, x\rangle = \langle y^*, Ax\rangle$, and whenever $x \in \ker(A)$ this does pair to zero. For the other way around, suppose we have some x^* in the annihilator of the kernel. Then $x^*(\xi) = 0$ for all $\xi \in \ker(A)$, so by (3) we have $\langle x^*, x\rangle = \langle x^*, x+\xi\rangle \implies |\langle x^*, x\rangle| = \langle x^*, x + \xi\rangle$ is bounded by Cauchy-Schwarz by $C||x^*||||Ax||$. (So here we're using explicitly that we're taking an infimum over all ξs.) Next, define a mapping $\psi : \text{Ran}(A) \to \mathbb{C}$ such that whenever $y = Ax$, we have $\psi(y) = \langle x^*, x\rangle$ (this doesn't depend on the choice of x because x^* annihilates the differences). We know that $\psi(y)$ is bounded, since $|\psi(y)| \le ||x^*||||Ax|| = ||x^*||||y||$, so it extends continuously to Ran(A). And we know that $\psi \circ A = x^*$, so there exists some $y^* \in Y^*$ which extends ψ. Then

$$x^* = \psi \circ A = y^* \circ A = A^*y^*,$$

so if we have an element in $(\ker A)^\perp$ it is also in the range of A^*, which is what we wanted to show.

(4) implies (5) is the definition of the weak*-topology, and (5) implies (6) is the fact that weak*-closed sets are closed. (6) implies (7) is identical to (2) implies (3). And we'll do the last implication next time, basically using the open mapping theorem. □

13 February 21, 2023

Last lecture, we considered an operator between two Banach spaces $A : X \to Y$. We know that $\mathrm{Ran}(A)$ is a subspace of Y, but it may not be closed, and we were establishing a criterion for when it is closed. The point is that when we try to solve the equation $Ax = y$, Fredholm theory is meant to be a situation where we can "almost" always get a solution – specifically, if A is Fredholm, then y will lie in a finite codimensional subspace, meaning there are $\ell_1, \cdots, \ell_n \in Y^*$ such that $\ell_1(y) = \cdots = \ell_N(y) = 0$ if and only if we can solve for a solution y. And we can be in a situation of closed range with infinite codimension, in which the statement is more complicated but we still have some criterion for solvability.

Example 130

Consider $\Delta : L^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$, which is an unbounded operator (so it's not defined everywhere – it would map into H^{-2} if we wanted to define it everywhere). As discussed before, this is equivalent to the multiplication operator $M_{|\xi|^2} : L^2(\mathbb{R}^2) \to L^2(\mathbb{R}^2)$, and we may ask about the range of this multiplication operator.

In other words, we want to know, given $\hat{f} \in L^2$, whether there is some $\hat{u} \in L^2$ with $|\xi|^2\hat{u} = \hat{f}$. (We could ask this at the distribution level too, but here we're asking in the sense of functions.) Unfortunately, we can't really write this in any real effective criterion, and this is an exapmle where having non-closed range is not great.

Example 131 (Lewy's counterexample)

There is a differential operator which looks something like $L = \partial_x + i\partial_y + (x + iy)\partial_t$, such that the map $L : \mathcal{D}'(\Omega) \to \mathcal{D}'(\Omega)$ does not have closed range for all Ω.

Often in a large domain we might be able to shrink the domain smaller and get rid of problems (getting a perturbative result), since the differential operator will be almost constant-coefficient. But in this example that never happens, and there will always be smooth compactly supported functions which have no solution. (And this is all connected to ill-conditioning of very large matrices as well.)

Recall that a bound of the form $||x|| \le C||Ax||$ automatically gives us closed range (because we can see that the Ax_is being Cauchy means the x_is are Cauchy). Such an estimate is quite strong, because in particular it requires that A has no nullspace, and the appropriate generalization is that we pass to the quotient space and show that $\inf_{A\xi=0} ||x + \xi|| \le C||Ax||$. We'll now prove the last part of Theorem 129, showing that

$$\inf_{A^*\eta^*=0} ||y^* + \eta^*|| \le C||A^*y^*|| \implies \mathrm{Ran}(A) = {}^{\perp}\ker(A)^*,$$

where ${}^{\perp}\ker(A)^*$ is the predual (the set of $y \in Y$ such that $\ell(y) = 0$ for all $\ell \in \ker(A^*)$).

Proof. First we assume that A^* is injective. This automatically implies that the range is dense, and we want to show that it is the whole space. We know by assumption that $||y^*|| \le C||A^*y^*||$, and we can define the set $K = \{Ax : ||x|| < 1\}$. Our first step is to show that $\{y : ||y|| \le \frac{1}{c}\} \subset \overline{K}$ (where we're taking the strong closure). We know that K is convex and $\overline{K}$ is closed convex, and it suffices to show that if $y_0 \notin \overline{K}$ we have $|y_0| > \frac{1}{c}$. By Hahn-Banach, there exists some separating y_0^* such that $\langle y_0^*, y_0\rangle > \sup_{y\in\overline{K}}\langle y_0^*, y\rangle$. Then

$$||A^*y_0|| = \sup_{||x||<1} \langle A^*y_0^*, x\rangle = \sup_{||x||<1} \langle y_0^*, Ax\rangle = \sup_{y\in K}\langle\langle y_0^*, y\rangle < ||y_0^*|| \cdot ||y_0||,$$

so $||y_0|| \geq \frac{||A^* y_0^*||}{||y_0^*||} \geq \frac{1}{c}$ by our original hypothesis. So there is indeed a ball around any point in $\overline{K}$ contained inside within $\overline{K}$, meaning that $\{y : ||y|| < \frac{1}{c}\} \subset \overline{\{Ax : ||x|| < 1\}}$. From here, we use the same technique as in the proof of the open mapping theorem (taking a sequence of balls staying inside the closure), and specifically choosing a ball so that we instead have (now without the closure)

$$\left\{y : ||y|| < \frac{1}{c}\right\} \subset \{Ax : ||x|| < 1\} .$$

But now for any y, $\frac{y}{2c||y||}$ is inside the left-hand set, meaning that it is within the range of A; rescaling shows that the range of A is indeed everything.

Finally, for the general case where we don't assume A^* is injective, we can take $Y_0 = \overline{\text{Ran } A}$, consider $A : X \to Y_0$ and apply the same argument. □

Corollary 132

An operator A is surjective if and only if A has closed range and A^* is injective, and A^* is surjective if and only A^* is injective and has closed range. And then we can check (using ordinary linear algebra) that $(A^*)^{-1} = (A^{-1})^*$ whenever A is bijective (and thus A^* is also bijective).

Definition 133

An operator $K : X \to Y$ is **completely continuous** if whenever $\{x_j\}$ is a bounded sequence, then Kx_j has a convergent subsequence (or alternatively K sends weakly convergent sequences to strongly convergent sequences).

We previously defined that $K : X \to Y$ is compact if $\overline{K(B)}$ is compact, and it turns out this new notion is related:

Lemma 134

If K is compact, then it is completely continuous. And if X is reflexive, then X being completely continuous implies compactness.

Proof. We'll just prove this last part. If X is reflexive, the unit ball is weakly compact, so any bounded sequence is weakly convergent, and thus K applied to some subsequence will be strongly convergent. □

Definition 135

An operator K is **finite rank** if $\text{Ran}(K)$ is finite-dimensional.

Bounded sets in $\mathbb{R}^n$ have some convergent subsequence, so any finite rank operator is compact.

Lemma 136

If $K_j \to K$ and K_j are compact, then K is also compact. In particular, the set of compact operators $\mathcal{K}(X)$ is a closed subspace.

Proof. If we have a weakly convergent subsequence, a usual diagonal subsequence allows us to make $K_j(x_k)$ all converge. Rename this subsequence to x_k and notice that

$$K(x_k - x_\ell) = (K - K_J)(x_k - x_\ell) + K_J(x_K - x_\ell)$$

which can be made arbitrarily small because $K - K_J \to 0$ and and the x_ks are Cauchy in K_J. □

In particular, the operator norm limit of any sequence of finite rank operators is compact, but the closure of the finite rank operators is **not always** all compact operators. However, if X is a separable Hilbert space this will be true.

Lemma 137

If A is bounded and K is compact, then $A \circ K$ and $K \circ A$ are compact. Combined with the previous result, we see that $\mathcal{K}(X)$ is an ideal of the space of operator $\mathcal{B}(X)$.

(This follows pretty directly from the definitions.)

Lemma 138

K is compact if and only if K^* is compact.

FInite rank operators have closed range, but otherwise we never have closed range – we should be thinking of approximation properties as saying that we are dividing our overall space into a finite dimensional-bit and the rest. And outside of that finite-dimensional the norm is very small.

Proof. For one direction, suppose K is compact, and consider $M = \overline{K(B)}$ as a compact metric space with the norm metric. Then for any $y^* \in Y^*$, consider the collection of all $f_{y^*} : M \to \mathbb{R}$, which are functionals y^* restricted to M. Defining $\mathcal{F} = \{f_{y^*} : ||y^*|| \le 1\}$, we wish to apply Arzela-Ascoli. We know that

$$\sup_{y \in M} |f_{y^*}(y)| = ||f_{y^*}||_{C^0(M)} = \sup_{y \in M} |\langle y^*, y \rangle|$$

but this right-hand side is uniformly bounded because M is bounded. So $\mathcal{F}$ is a uniformly bounded family of functions, and in fact we can say a bit more: this is bounded by $\sup_{||x||<1} |\langle y^*, Kx \rangle| = \sup_{||x||<1} |\langle K^* y^*, x \rangle| = ||K^*|| = ||K||$m so we do have an explicit bound in terms of K. And uniform equicontinuity is clear because $|f_{y^*}(y) - f_{y^*}(y')| = |\langle y^*, y - y' \rangle| \le ||y^*|| \cdot ||y - y'|| \le ||y - y'||$. Now up to taking a subsequence we can assume that $f_{y_n}^* \to f$, so that $K^* y_n^*$ is Cauchy. But that's exactly what we need to show that the sequence has a convergent subsequence.

Now for the other direction, if K^* is compact, then K^{**} is compact by our previous result, and we want to show that K is compact. We have an isometric embedding $\iota_X : X \to X^{**}$. Then (x_n) being bounded means that $K^{**}(\iota_X(x_n))$ has a convergent subsequence, and since K^{**} is a mapping $X^{**} \to Y^{**}$, and we can "take double duals" or apply K/K^{**} and get the same result, this means we do have a convergent subsequence for the corresponding K maps. □

Example 139

Consider $H_0^s(B) \subset H_0^t(B)$ for $s > t$ for the unit ball B. By the Rellich theorem, this is a compact inclusion, and that's the same as saying that the inclusion map is a compact operator.

Recall that an operator is **Fredholm** if it has finite-dimensional kernel, closed range, and the cokernel of A is finite-dimensional. So the point is that if we want to solve $Ax = y$, then we have a situation where we have $\ker A$ and its complement X_0 in X, and then we have $\text{Ran}(A)$ and the complement $\text{coker}(A)$ in Y. Then A is a bijective map between the spaces except the kernel and cokernel. And since the range is closed we get a Banach subspace: $A : X_0 \to \text{Ran}(A)$ is then a bounded bijective map, hence boundedly invertible, and we have a map G which inverts on these good subspaces. (So we have a finite set of conditions to check whether we have a solution, and the solution set will also be finite dimensional.)

Definition 140

The **index** of a Fredholm operator A is $\dim \ker A - \dim \operatorname{coker} A$.

The reason this is an interesting quantity is that it is very stable under perturbations, meaning that it can actually be calculated. And A is Fredholm if and only if A^* is Fredholm – the numbers for $\dim \ker$ and $\dim \operatorname{coker}$ just get swapped, so that $\operatorname{Ind}(A^*) = -\operatorname{Ind}(A)$.

Theorem 141

The following are equivalent:

1. A is Fredholm,
2. A (as an element of the quotient algebra $\mathcal{C}(X) = \mathcal{B}(X)/\mathcal{K}(X)$) is invertible,
3. there is some $G : Y \to X$ such that $A \circ G = \operatorname{Id} - R_1$ and $G \circ A = \operatorname{Id} - R_2$, where R_1, R_2 are compact.

Here $\mathcal{B}(X)$ is an algebra where we can compose and take adjoints, and modding out by an ideal still gives an algebra (here this is called the Calkin algebra) – what we're saying is that the equivalence classes for G and A are inverses. So it is clear that (2) and (3) are really saying the same thing.

Proof. First we'll show (1) implies (3). Continue thinking about the decomposition of X and Y as we did above. Let $\tilde{G}$ be the inverse map $\operatorname{Ran}(A) \to X_0$ to A, and now define $G = \hat{G}$ on $\operatorname{Ran}(A)$ and 0 otherwise. So $G \circ A$ is the identity, minus the projection map onto $\ker(A)$. (When we're in a Banach space we must talk about what we're projecting onto, as well as what we're projecting off of, since we need to write things uniquely as a sum.) On the other hand, $A \circ G$ is the identity minus the projection onto the cokernel of A. Since both the kernel and cokernel are finite-dimensional, the projection operators are finite-rank, hence compact. (Notice that this means the "not-near-finite-rank compact operators" do not arise in this Fredholm argument.)

On the other hand, let's show that (3) implies (1). To see that $\ker(A)$ is finite-dimensional, suppose $\{x_j\} \in \ker(A)$ are all in the unit ball. Then $G \circ A x_j = (I - R_2)x_j$, so $x_j = R_2 x_j$ for all j. Then $R_2 x_j$ has a convergent subsequence by definition of compactness, meaning that x_j has a convergent subsequence, which means our ball must have been finite-dimensional to begin with. For closed range, suppose $Ax_j = y_j$ and $y_j \to \overline{y}$. Then $GAx_j = Gy_j = x_j - R_2 x_j$, and since G is a bounded operator the left-hand side is convergent. But the right-hand side may not have bounded x_js as stated – instead, we must first assume $x_j \in X_0$ to avoid large contributions from the kernel of A. (Here it's important that we have a topological complement so that we can do this projection.) Then we have $x_j = Gy_j + R_2 x_j$, and now we can rewrite this as

$$\frac{x_j}{||x_j||} = \frac{1}{||x_j||} G y_j + R_2 \frac{x_j}{||x_j||}.$$

Then the Gy_js converge, so if $||x_j|| \to \infty$ (this is the bad case) then the first term on the right-hand side converges to zero, and then repeating the argument from before shows that (because $\frac{x_j}{||x_j||}$s are now bounded) $R_2 \frac{x_j}{||x_j||}$ has a convergent subsequence, meaning that some subsequence of the $\frac{x_j}{||x_j||}$s converge to some $\overline{x}$ (again relabel this subsequence to be the main sequence). Now $A\left(\frac{x_j}{||x_j||}\right) = \frac{y_j}{||x_j||}$ by definition – if we take the limit on both sides, we get $A\overline{x}$ on the left and 0 on the right, since $\frac{\overline{y}}{||x_j||} \to 0$. But the x_js all live in X_0, and thus we have a unit vector $\overline{x}$ in the nullspace of A. We've chosen x to be complemented to the nullspace, so this case is impossible. So otherwise $||x_j||$ are bounded and the argument from before works. Finally, the cokernel is finite dimensional because we can take adjoints to get $G^* A^* = I - R_1^*$ and apply the same arguments as before, using that the adjoint of a compact operator is compact. $\square$

Proposition 142
An operator A has finite-dimensional kernel and closed range if and only if for some constant $C > 0$ and some compact operator $K \in \mathcal{K}$,

$$||x|| \le C(||Ax|| + ||Kx||).$$

The point here is that for any bounded sequence in the nullspace we can make the $K(x_j - x_k)$s converge, so again this implies that the unit ball is compact. And closed range will be the same argument as what we just did – next time, we'll continue this and discuss how the index behaves with respect to Fredholm operators. That'll lead into computations of index and some important Fredholm operators.

14 February 23, 2023

We'll continue our discussion of Fredholm operators and then look at properties of the index and the Toeplitz index theorem today. We have the class of Fredholm operators $\mathcal{F}(X, Y)$ and the class of compact operators $K(X, Y)$, and we showed that $A \in \mathcal{F}$ if and only if there is some $G \in \mathcal{B}(Y, X)$ with $A \circ G = I - K_1$ and $G \circ A = I - K_2$ for $K_1 \in \mathcal{K}(Y)$ and $K_2 \in \mathcal{K}(X)$. In other words, Fredholm operators are "almost invertible" in that way we can decompose $X = \ker A \oplus X_0$ and $Y = \text{Ran}(A) \oplus C_0$ (where C_0 is isomorphic to the cokernel)

Remark 143. *Here we use the fact that any finite-dimensional / finite-codimensional space has a complement. This is true because for a one-dimensional space spanned by v we can find a linear functional ℓ with $\ell(v)$ nonvanishing by Hahn-Banach and consider $\ell^{-1}(0)$ and then for higher dimensions intersect those spaces. And for the finite-codimension case, we can just lift the finite basis of the quotient $Y/Ran(A)$ to the whole space. The reason we can't work too well with algebraic bases is that they don't really give us any sense of whether subspaces will be closed.*

So now restricting G (this is called a **parametrix**) to the range of A, we have a map $A^{-1} : \text{Ran}(A) \to X_0$, and given any such decomposition of X and Y in this form we can define G to recover K_1, K_2 being the projections onto the kernel and cokernel. So for a general compact remainder term, this argument tells us that we can actually improve this to finite-rank operators – this is still consistent with the fact that not all compact operators are well-approximated by finite-rank operators because the parametrix is far from unique.

We'll next prove Proposition 142, which gives us an estimate looking like the closed-range estimate but with an extra $||Kx||$ term if we require a finite-dimensional kernel. (But notice that there's no requirement of finite-dimensional cokernel here.) In the study of PDEs we often do get bounds of this form, and unless we do more work it does not actually follow that the cokernel is finite-dimensional.

Proof. For the forward direction, choose a basis $\{x_1, \cdots, x_n\}$ for $\ker A$, and choose appropriate linear functionals $x_i^* \in X^*$ such that $x_i^*(x_j) = \delta_{ij}$. So the compact operator is essentially just a "projection onto the nullspace," and we define

$$Kx = \sum_{j=1}^{n} x_j^*(x)x_j.$$

Clearly K is the identity operator on the nullspace, and the common nullspace of the x_i^*s is a complement to $\ker(A)$. So now consider the map $x \mapsto (Ax, Kx) \in Y \oplus (\ker A)$; this is still a map between Banach spaces and this time it is injective. Thus by our closed-range estimate, $||x|| \le C||(Ax, Kx)|| = C(||Ax|| + ||Kx||)$, as desired.

For the other direction, for any bounded sequence $\{x_j\} \in \ker A$ we see that $||x_j|| \le C||Kx_j||$. But then Kx_j has a convergent subsequence, hence x_j does as well; thus $\ker A$ must indeed be finite-dimensional and we can write

$X = \ker A \oplus X_0$. So now suppose we have a sequence $Ax_j \to y_j$ with y_j converging to $\overline{y}$. We wish to show that the x_js also converge (at least have a convergent subsequence), but again it's possible that they get larger and larger along the $\ker A$ direction. So much like last proof, we can replace x_j with $x_j - w_j = \tilde{x}_j$, where $Aw_j = 0$ and $\tilde{x}_j \in X_0$ for all j. Now by our estimate, if the $\tilde{x}_j$s are bounded then we have

$$||\tilde{x}_j - \tilde{x}_k|| \le C||y_j - y_k|| + ||K(\tilde{x}_j - \tilde{x}_k)||;$$

we may assume that a subsequence of the $K\tilde{x}_j$s is Cauchy, so passing to that subsequence both terms on the right-hand side are small and thus the $\tilde{x}_j$s are also Cauchy and thus converge. So in this case we are fine; otherwise we must do the same argument as last time and take $\hat{x}_j = \frac{\tilde{x}_j}{||\tilde{x}_j||}$. Then $A\tilde{x}_j$ now actually converge to zero, and again by compactness a subsequence of the $K\hat{x}_j$ will converge, so now $\hat{x}_{j'}$ converges to $\overline{x}$. But now $\overline{x} \in X_0$ as well, and $A\overline{x} = 0$ even though $\overline{x}$ is of norm 1, which is again a contradiction. ☐

Example 144

To get a bit of motivation for how we might see this argument, suppose we have some general elliptic equation on a domain Ω of the form

$$L = \sum a_{ij}(x)\frac{\partial^2}{\partial x_i x_j} + b_i(x)\frac{\partial}{\partial x_i} + c(x),$$

and suppose we want to know the solvability of the equation $Lu = f$.

Then take $u = C_0^\infty(\Omega)$ (in particular vanishing on the boundary of our domain) and consider (different arguments for Sobolev or Holder or other spaces) the estimates that look something like

$$||u||_{H^2} \le C(||Lu||_{L^2} + ||u||_{L^2})$$

(this is called an **a priori estimate**, and these usually come from things like Cauchy-Schwarz or other real analysis trickery). If we know this is true for all $u \in C_0^\infty$, then taking the closure with respect to whatever norms are involved will give us something useful. Specifically, if $u_j \in C^\infty$ converge to u in L^2 and with Lu_j converging to f in L^2, we get the same kind of bound for those limiting functions, and then we can find that this bound works for every $u \in H^2 \cap H_0^1$. And in fact the inclusion H^2 into L^2 is a compact embedding by Rellich, so we've thus proved that we have finite-dimensional kernel and closed range for this operator L. But the cokernel may still be infinite-dimensional, and so the trick is to apply the same trick to the **adjoint** L^* instead. Then we find that L^* has closed range (which we already know), but also that the kernel of L^* is finite-dimensional. So L is then actually Fredholm. (And if we want to compute L^* explicitly we just integrate by parts to get $\int(Lu)v = \int u(L^*v)$, and it'll also be a differential operator of the same type. But then there's a whole theory of the case where a_{ij} is not necessarily differentiable, and there are reasons why we might be interested in that case as well.)

Proposition 145

Suppose $A \in \mathcal{F}(X, Y)$ and $B \in \mathcal{B}(X, Y)$ has sufficiently small operator norm with respect to A. Then $A + B \in \mathcal{F}$. (In other words, the space of Fredholm operators is open in the operator topology.) Also, if $A \in \mathcal{F}(X, Y)$ and $K \in \mathcal{K}(X, Y)$, then $A + K \in \mathcal{F}(X, Y)$ (so the space is invariant under translation by the space of compact operators).

Proof. The latter fact is easier to prove: if A is Fredholm, then we can choose a parametrix G. Then $G(A + K) = GA + GK = (I - R_1) + GK$ and $(A + K)G = AG + KG = (I - R_2) + KG$. But since GK and KG are both compact, we can combine the compact error terms and see that G will still be a parametrix for $A+K$, and thus $A+K$ is still Fredholm.

Remark 146. *This in fact closes the loop on the Lax-Milgram idea from a previous lecture (Example 122): recall that we showed that for some sufficiently large λ, $L+\lambda$ is an invertible map from $H_0^1 \to H^{-1}$. But now the inclusion $H_0^1 \to H^{-1}$ is compact, so $L = (L+\lambda) - \lambda$ is Fredholm because $L+\lambda$ is invertible, hence Fredholm, and λ is compact.*

For the former fact, we do the same trick: we know that

$$(A+B)G = I - R_1 + BG.$$

Now we can use the fact that if $||BG|| < 1$, then $I + BG$ is invertible, meaning that we can write

$$(A+B)(G(I+BG)^{-1}) = I - R_1(I+BG)^{-1},$$

and thus $G(I+BG)^{-1}$ can be the new parametrix and $R_1(I+BG)^{-1}$ can be the new compact error. The same thing works for modifying R_2, but now we have equations of the form

$$(A+B)G_1 = I - Q_1, \quad G_2(A+B) = I - Q_2.$$

But now if we run through the arguments to show that $A+B$ is Fredholm, **all of them still work**, and thus $A+B$ is Fredholm and there is actually some single good operator H such that $(A+B)H = I - \pi_1$ and $H(A+B) = I - \pi_2$ where π_1, π_2 are some projections onto a finite-dimensional space. And notice that

$$H(I-Q_1) = H(A+B)G_1 = (I-\pi_2)G_1, \quad (I-Q_2)H = G_2(A+B)H = G_2(I-\pi_1),$$

so equating the expressions for H we find that

$$H = G_1 - \pi_2 G_1 + HQ_1 = G_2 - G_2\pi_1 + Q_2 H.$$

So in fact G_1 and G_2 just differ by a compact operator since all the other terms are compact. $\square$

Recall that the **index** of a Fredholm operator A is $\dim\ker A - \dim\operatorname{coker} A$, and we'll now explore a few properties of it (it's a homomorphism, it's locally constant, and it can be computed for a few interesting operators). The intuition for why this object is good is that in spectral theory, we care about knowing when $A - \lambda I$ is invertible. In the case where A is self-adjoint (for example a finite-dimensional symmetric matrix), the eigenvalues are on the real line, and when we're away from an eigenvalue we are invertible, but when we hit an eigenvalue we suddenly gain a nullspace. But the cokernel will jump up by the same amount whenever the kernel jumps up.

Theorem 147

If $A : X \to Y$ and $B : Y \to Z$ are Fredholm operators, then $\operatorname{Ind}(BA) = \operatorname{Ind}(A) + \operatorname{Ind}(B)$.

Proof. We can write down a bunch of vector spaces and chase dimensions algebraically, but the picture to have in mind is that we decompose X into $\ker A$, another space $A^{-1}(\ker B) \cap X_0$ (which is the space in X_0 which maps to $\ker B$ but isn't immediately mapped to zero in Y). Then $A^{-1}(\ker B) \cap X_0$ is isomorphic to $\ker B \cap \operatorname{Ran}(A)$, and the other part of $\ker B$ is $\operatorname{coker}(A) \cap \ker B$. Finally, in Z we can think about the range of B, inside of which we have $\operatorname{Ran}(BA)$ – the complement of $\operatorname{Ran}(BA)$ is the set of things not hit by BA but hit by B, so they're not coming from the range of A and thus what we're left with is $\operatorname{coker}(A)$, except for the part $\operatorname{coker}(A) \cap \ker B$. So now with all of these relevant bits we can compute the index:

$$\begin{aligned}\dim\ker(BA) - \dim\operatorname{coker}(BA) = &\left[\dim(\ker A) + \dim(A^{-1}(\ker B)\cap X_0)\right] \\ &- \left[\dim\operatorname{coker}(B) + \dim\operatorname{coker}(A) + \dim(\operatorname{coker}(A)\cap\ker B)\right].\end{aligned}$$

The first and fourth terms then give us the index of A, the second and fifth terms give us the kernel of B, and thus combining that with the third term gives us the index of B. $\square$

Thus the index is actually a mapping from $\mathcal{F}\to\mathbb{Z}$ which is a homomorphism.

Theorem 148

Let A be Fredholm. Then for all Fredholm operators B with sufficiently small norm, $\operatorname{Ind}(A+B) = \operatorname{Ind}(A)$.

Proof. Again take $X = \ker(A)\oplus X_0$ and $Y = \operatorname{Ran}(A)\oplus C_0$, and define projectors onto the different pieces $P_{\ker(A)}, P_{X_0}$, $\pi_{\operatorname{Ran}(A)}$, and π_{C_0}. Now split up into four terms by writing identity operators as a sum of two projectors:

$$A + B = \pi_{C_0}(A+B)P_{\ker(A)} + \pi_{C_0}(A+B)P_{X_0} + \pi_{\operatorname{Ran}(A)}(A+B)P_{\ker(A)} + \pi_{\operatorname{Ran}(A)}(A+B)P_{X_0}.$$

Think of this as a block matrix $\begin{bmatrix}\pi_{C_0}(A+B)P_{\ker A} & \pi_{C_0}(A+B)P_{X_0}\\ \pi_{\operatorname{Ran}(A)}(A+B)P_{\ker A} & \pi_{\operatorname{Ran}(A)}(A+B)P_{X_0}\end{bmatrix}$. Then when $B = 0$, only the stuff in the bottom right corner remains, and by definition this is an invertible operator (it's $A : X_0 \to \operatorname{Ran}(A)$). So when B is small, we get a new operator $D = \begin{bmatrix}D_{11} & D_{12}\\ D_{21} & D_{22}\end{bmatrix}$ where D_{22} is invertible (since small perturbations of an invertible matrix are invertible) and the others are all finite-rank because we're either going into or out of a finite-dimensional space. And if we try to solve the system of equations

$$D_{11}x_1 + D_{12}x_2 = y_1, \quad D_{21}x_1 + D_{22}x_2 + y_2$$

for x, we have $x_2 = D_{22}^{-1}(y_2 - D_{21}x_1)$ and thus (substituting into the first equation) $D_{11}x_1 + D_{12}D_{22}^{-2}(y_2 - D_{21}x_1) = y_1$, which simplifies to $(D_{11} - D_{12}D_{22}^{-1}D_{21})x_1 = y_1 - D_{12}D_{22}^{-1}y_2$. So the point is that obstruction of invertibility is the failure of invertibility of $\mathcal{D} = D_{11} - D_{12}D_{22}^{-1}D_{21}$, which is a matrix from $\ker A \to \ker A$. Thinking of this as a function of our perturbation B and noting that $\mathcal{D}(0) = 0$, notice that for a family of matrices $\mathcal{D} : \mathbb{R}^N \to \mathbb{R}^N$ the index of $\mathcal{D}$ will always be constant by rank-nullity. So as long as D_{22} stays invertible we do have constant index. $\square$

Corollary 149

If A is Fredholm and K is compact, then $\operatorname{Ind}(A+K) = \operatorname{Ind}(A)$.

Proof. Do the same proof as before, but note that $A+K$ stays Fredholm even if it doesn't stay invertible. So if we look at the family of operators $A + tK$ for $t\in[0,1]$, they're all Fredholm and thus the index is locally constant, so by connectedness $A+K$ has the same index as A. $\square$

The reason for the name "index" is that the Fredholm operator splits as a disjoint union

$$\mathcal{F} = \bigsqcup_{k\in\mathbb{Z}} \mathcal{F}_k,$$

indexed by the index k, and $\mathcal{F}_k$ are unions of connected components. Furthermore, the index map is a homomorphism from $\mathcal{F}/\mathcal{F}_0 \to \mathbb{Z}$ and in fact $\pi_0\mathcal{F}\to\mathbb{Z}$ (here π_0 is the space of connected components), and we may ask what the $\mathcal{F}_k$s look like.

Fact 150 (Atiyah)

Suppose X is a separable Hilbert space. Then $\mathcal{F}_0$ is connected and the index map is onto, meaning that the space of Fredholm has infinitely many components indexed by the various integers which are basically translates of each other.

It turns out the unit sphere in a separable Hilbert space is contractible – we kill homotopy groups by adding more dimensions. And there's lots of results of that type – in this case we can find a continuous path from any Fredholm operator of index zero to any other one.

Example 151

A way to generate operators of any arbitrary index that's good to keep in mind is the following. Consider the sequence space ℓ^2, and define the left shift $S_L : (x_0, x_1, x_2, \cdots) \mapsto (x_1, x_2, x_3, \cdots)$ and right shift $S_R : (x_0, x_1, x_2, \cdots) \mapsto (0, x_0, x_1, \cdots)$. Then the index of S_L is 1 and the index of S_R is -1, and composing these gives us any arbitrary index that we want.

Related to this, we'll finish by using topology to do an interesting computation: identify $L^2(S^1)$ with ℓ^2 (with bi-infinite sequences coming from Fourier coefficients), and consider the space $\mathcal{H}_+ = \left\{u = \sum_{n=0}^{\infty} u_n e^{in\theta}\right\}$ (that is, the space with only positive Fourier coefficients). Let $\pi_+ : L^2(S^1) \to \mathcal{H}_+$ be the projector which removes the negative Fourier coefficient terms.

Theorem 152 (Toeplitz index theorem)

For any $f \in C^0(S^1)$, let $T_f = \pi_+ M_f \pi_+$ be an operator $\mathcal{H}_+ \to \mathcal{H}_+$ (this is called a **compression**). Then T_f is Fredholm if and only if $f \neq 0$ everywhere (meaning that we map from $S^1 \to \mathbb{C}^*$), and $\text{Ind}(T_f) = -W(f)$, where $W(f)$ is the winding number of f as that map $S^1 \to S^1$.

We'll look at the proof this result next time!

15 February 28, 2023

We'll do the proof of the Toeplitz index theorem and then move on to our next set of topics today. There is a general index theorem for elliptic operators (which we've proved are Fredholm) on manifolds, and a natural question is to ask whether we can compute the index of those operators. The Toeplitz index theorem is one such method, and there are various other index theorems culminating in the Atiyah-Singer index theorem as well. This one that we're proving is a predecessor to that.

Recall that we are trying to prove the following fact: let $\mathcal{H}_+$ be the space of functions in $L^2(S^1)$ whose negative Fourier coefficients are all zero. Then defining π_+ and π_- to be the projections onto $\mathcal{H}_+$ and its complement, we can define for any function f the operator $T_f = \pi_+ M_f \pi_+$. Then T_f is Fredholm if f is nowhere vanishing, and in that case we can view it as a map $S^1 \to \mathbb{C}^*$; the index of T_f is then $-W(f)$. Here note that if $f = \sum_k e^{ik\theta}$ and $u = \sum_{n=0}^{\infty} u_n e^{in\theta}$, then

$$fu = \sum_{k \in \mathbb{Z}} \sum_{\ell \in \mathbb{Z}_{\geq 0}} f_k u_\ell e^{i(k+\ell)\theta}.$$

Reindexing with $n = k + \ell$, this can be rewritten as

$$= \sum_{n=-\infty}^{\infty} \left(\sum_{\ell=0}^{\infty} f_{n-\ell} u_\ell \right) e^{in\theta},$$

and thus $T_f u$ just cuts off the negative-n terms and yields $\sum_{n=0}^{\infty} \sum_{\ell=0}^{\infty} f_{n-\ell} u_\ell e^{in\theta}$.

Proof. To prove that T_f is Fredholm, we can construct a parameterix, and we claim that $T_{1/f}$ does the job (here's where we use that f is nowhere vanishing). Indeed,

$$T_f \circ T_{1/f} = \pi_+ M_f \pi_+ M_{1/f} \pi_+ = \pi_+ M_f \left(\mathrm{Id} - \pi_-\right) M_{1/f} \pi_+ = \pi_+ M_f M_{1/f} \pi_+ - \pi_+ M_f \pi_- M_{1/f} \pi_+ = \mathrm{Id} - \pi_+ M_f \pi_- M_{1/f} \pi_+,$$

and thus our goal is to show that $\pi_+ M_f \pi_- M_{1/f} \pi_+$ is compact. Notice that $||Tf|| \le \sup |f|$, since $||T_f|| \le ||\pi_+|| \, ||M_f|| \, ||\pi_+|| \le M_f$, and the multiplication by f map has norm bounded by $\sup |f|$ and projection maps are of norm 1. So we can use that to show that this remainder term $\pi_+ M_f \pi_- M_{1/f} \pi_+$ is the limit of finite-rank operators. We split this up as $(\pi_+ M_f \pi_-)(\pi_- M_{1/f} \pi_+)$, and we focus on the $(\pi_- M_{1/f} \pi_+)$ part. We can write f as a limit in C^0 of trigonometric polynomials $f_N = \sum_{-N}^{N} a_n e^{ik\theta}$ (here we use Weierstrass approximation to see that we have denseness of trigonometric polynomials), and now the product $M_{f_N} u$ has Fourier components zero for all $k < -N$ (since we shift down each term by at most N when expanding out). But now that means $\pi_- M_{f_N} \pi_+$ has finite rank, and now $|\pi_- M_f \pi_+ - \pi M_{f_N} \pi_+| \to 0$ because this is at most $||M_{f-f_N}|| \le \sup |f - f_N| \to 0$ by the same argument as before. So the operator we want is a norm limit of finite-rank operators, meaning it is indeed compact.

Now recalling that the index doesn't change under small perturbations of our operator, we claim that there exists a continuous function $F : [0,1] \times S^1 \to \mathbb{C}^*$ with $F(0, \theta) = f(\theta)$ and $F(1, \theta) = e^{iN\theta}$, where $N = W(f)$. In other words, we can continuously deform a curve in the complex plane until it just goes N times around the unit circle. (We won't prove this – it's a property of the punctured plane.) When $T_{F(\varepsilon, \cdot)}$ is Fredholm for all $0 \le \varepsilon \le 1$ (because this is true locally), and thus $\mathrm{Ind}(T_f) = \mathrm{Ind}(T_{e^{iN\theta}})$. But now multiplication by $e^{iN\theta}$ is exactly the N-dimensional right-shift if we view our Fourier coefficients in a sequence, yielding no kernel and a dimension N cokernel, so the index is $-N$. Similarly, we get a left shift for negative N, meaning we again get the same form for the index. ☐

In the Atiyah-Singer index theorem the computation is more complicated – it's not about multiplication by functions, instead about differential operators on manifolds, but we do a similar kind of deformation.

These Toeplitz operators have their roots in various parts of mathematics, such as random matrix theory. This leads us to the Riemann-Hilbert problems (which is a core strategy involving "factoring" our function f):

Theorem 153

If $W(f) = 0$, then in fact T_f is invertible, and if the winding number is negative (resp. positive), then T_f is injective (resp. surjective).

So this is stronger than just computing the index, since we're saying that there is no kernel / cokernel.

Remark 154. *Also, there is a variant of this Toeplitz index theorem: if we look at functions $u : S^1 \to \mathbb{C}^k$ such that $u_n = 0$ for all negative n, and $f : S^1 \to Mat_k(\mathbb{C})$, then we can define similarly an operator $\pi_+ M_f \pi_+$. A similar proof to the one we just showed also works, and this time we have $Ind(T_f) = -W(\det f)$. But this vector-valued version does not exhibit the injectivity/surjectivity facts for the one-dimensional case.*

Proof. We claim that if $W(f) = 0$, then we can "take the log of f" in the following way: let $f(\theta) = e^{i\sigma(\theta)}$, and then we can break up σ into $\sigma_+ + \sigma_-$ with $\sigma_\pm \in \mathcal{H}_\pm$. Then $e^{\sigma_\pm} = f_\pm$, and then we have $f = f_+ f_-$. But f_+ has all nonzero

Fourier coefficients greater than zero, that's the same as saying that it extends to being a holomorphic function inside the disk, and f_- then extends outwards. (This turns out to be a general trick that is useful.) So if we want to solve the equation $T_f u = h$, then it suffices to solve the equation $fu = h + g_-$ with $h \in \mathcal{H}^+$ and $g_- \in \mathcal{H}^-$. But using the factorization of f above, we want to solve

$$f_+ u = \frac{h}{f^-} + \frac{g^-}{f^-}$$

(this is valid because $f_\pm$ are exponentials of something so never zero). But then $f_-^{-1} g_- \in \mathcal{H}_-$ (here we use that $f_-^{-1} = \exp(-\sigma_-)$, and the exponential of something with only negative Fourier coefficients still only has negative Fourier coefficients because we can take a power series expansion and see that each term in $\sum \frac{1}{k!}(e^{-in\theta})^k$ has negative Fourier coefficients), so

$$f_+ u = \pi_+(f_-^{-1} h) \implies u = f_+^{-1} \pi_+(f_-^{-1})h,$$

and we've now managed to isolate out the "negative stuff'.' So if f has winding number zero, then we can indeed invert our operator of multiplication by f within the relevant space.

For the other parts, replace f with $fe^{-iN\theta}$ with $N = W(f)$. And it's left as an exercise that making this transformation proves injectivity and surjectivity of the original operator. $\square$

This finishes our discussion of Fredholm theory for now – we'll now move on to one of the last major topics, the **spectral theorem**. We basically want to answer the following questions: given a bounded operator $A \in \mathcal{B}(X)$, what is the spectrum $\text{spec}(A)$ (the generalization of the set of eigenvalues of a finite-dimensional matrix)? And what happens if A is a closed but unbounded operator (for example, thinking of $-\frac{d^2}{d\theta^2} : H^2 \to L^2$ as an operator on $L^2(S^1)$ with domain $\mathcal{D}(A) = H^2$)?

The idea is that we'll want to look not just at $Ax = \lambda x$, which doesn't tell us the whole story, and look at functions of A instead and make sense of them. For example, suppose we have a finite-dimensional self-adjoint operator with eigenvalues on the real line, and we consider f to be a bump function which is 1 at some eigenvalue λ_j and 0 at the others. When we already have such a spectral decomposition, the operator $f(A)$ should have eigenvalues $f(\lambda_i)$ and the same eigenvectors v_i, and so in this case $f(A)$ would be taking an expansion $x = \sum a_i v_i$ and spitting out the "spectral projection' onto the jth eigenspace $a_j v_j$. Our goal is to generalize to infinite dimensions.

Definition 155

Let $A \in \mathcal{B}(X)$. The **spectrum** of A is

$$\sigma(A) = \{\lambda \in \mathbb{C} : (\lambda I - A) \text{ not (boundedly) invertible}\},$$

and $\rho(A) = \mathbb{C} \setminus \sigma(A)$ is the **resolvent set**.

There are three reasons why an operator would fail to be invertible, and that allows us to divide up

$$\sigma = \sigma_{\text{pp}} \cup \sigma_{\text{res}} \cup \sigma_c,$$

where we have the **pure-point spectrum**

$$\sigma_{\text{pp}} = \{\lambda \in \mathbb{C} : \lambda I - A \text{ not injective}\}$$

corresponding to the λs where there is actually an eigenvector $Ax = \lambda x$, the **residual spectrum**

$$\sigma_{\text{res}} = \{\lambda \in \mathbb{C} : \lambda I - A \text{ injective but has a range which is not dense}\},$$

and the **continuous spectrum**

$$\sigma_c = \{\lambda \in \mathbb{C} : \lambda I - A \text{ injective and dense range but not all of } X\}.$$

In this part of the class, we'll make frequent use of complex analysis – the point is that if we think of λ as a complex number, we can consider mapping $\lambda \mapsto (\lambda I - A)^{-1} = R_\lambda(A)$ for any λ in the resolvent set. So this is a set from some set $\rho(A)$ in $\mathbb{C}$ to operators, and it will be "holomorphic" in the appropriate sense. Specifically, we'll prove the "strongest possible sense" of holomorphicity: we'll show that $\rho(A)$ is open, so that for any $\lambda_0 \in \rho(A)$ there is some $\varepsilon > 0$ so that whenever $|\lambda - \lambda_0| < \varepsilon$ we have a Taylor series $\sum_{j=0}^{\infty} R_j(\lambda - \lambda_0)^j$. Alternatively, we can show a "weak sense" of holomorphicity: if we choose $x \in X$ and $\ell \in X^*$, we can apply the strong operator-valued function to a given vector, apply a linear functional to that, and consider the map $\lambda \mapsto \ell(R_\lambda(A)x)$ and show that that is holomorphic. It turns out these two forms are actually equivalent, and we'll talk about that more next week.

Example 156

Consider the left and right shifts S_L and S_R on ℓ^2. We can compute the spectrum manually for these operators.

Notice that

$$(\lambda I - S_L)x = (\lambda x_1 - x_2, \lambda x_2 - x_3, \cdots),$$

$$(\lambda I - S_R)x = (\lambda x_1, \lambda x_2 - x_1, \lambda x_3 - x_2, \cdots).$$

We see that $\lambda I - S_R$ always has no nullspace, since that sequence being zero means $x_1 = 0$, which means $x_2 = 0$, and so on (except when $\lambda = 0$, where again we get no nullspace). But in the former case, determining x_1 determines $x_2, x_3, \cdots$, and we need λ to have norm less than 1 so that the solution $(x_1, x_2, \cdots)$ is actually in ℓ^2. Since the adjoint of S_L is S_R and vice versa, we're either in the point spectrum or the residual spectrum for S_L and S_R, respectively, for any $|\lambda| < 1$. And finally, the boundary of the unit ball $|\lambda| = 1$ forms the continuous spectrum for both operators.

Lemma 157

$\sigma(A)$ is a closed and bounded set.

Proof. First we prove that the resolvent is open. Suppose that $(\lambda_0 I - A)^{-1}$ exists. If $|\lambda - \lambda_0| < \varepsilon$, then $(\lambda I - A) = (\lambda_0 I - A) + (\lambda - \lambda_0)I$; the first term on the right-hand side is invertible so for small enough norm the left-hand side is also invertible. While we're working with these facts, we will prove a **resolvent identity**: if λ, μ are both in the resolvent set, then we have (here we write $\lambda = (\lambda - \mu) + \mu$ and multiply the first two terms out)

$$(\lambda I - A)\left[(\lambda I - A)^{-1} - (\mu I - A)^{-1}\right](\mu I - A) = \left(I - (\lambda - \mu)(\mu I - A)^{-1} - I\right)(\mu I - A) = \mu - \lambda,$$

so in other words (rearranging and plugging in the definition of $R_\lambda(A)$)

$$\boxed{R_\lambda(A) - R_\mu(A) = (\mu - \lambda)R_\lambda(A)R_\mu(A)},$$

and this will turn out to be a useful identity later on. To show that $\sigma(A)$ is actually bounded, define the **spectral radius** of A to be $r_A = \lim_{n\to\infty} ||A^n||^{1/n}$. (This limit indeed exists because $||A^2|| \le ||A||^2 \implies ||A^n|| \le ||A||^n$ and then we can do a bit more work.) We claim that $r_A \ge \sup_{\lambda\in\sigma(A)} |\lambda|$, and we'll do this by showing that if $\lambda > r_A$ then λ is in the resolvent set. Indeed, then we have $r_{\lambda^{-1}A} < 1$ (the spectral radius scales), so in particular $\lambda^{-1}A$ has norm at most 1 and thus $I - \lambda^{-1}A$ is invertible. But then $\lambda I - A = \lambda(I - \lambda^{-1}A)$ is also invertible, hence in the resolvent. $\square$

We may now want to know what sets can actually show up as $\sigma(A)$ for some A – it turns out, for example, that we can always find Schrodinger operators that have a given spectrum $\sigma(A)$ within the restrictions we've found. We can specialize to **Hilbert spaces** to see this more clearly, and we'll consider **self-adjoint operators**.

Lemma 158

If X is Hilbert and A is self-adjoint, then $\sigma(A) \subset \mathbb{R}$.

Proof. Expand out the following expression using the inner product:

$$\boxed{||(\lambda I - A)x||^2} = |\lambda|^2||x||^2 - \langle \lambda x, Ax\rangle - \langle Ax, \lambda x\rangle + ||Ax||^2,$$

where the middle two terms are the same as $2\text{Re}(\lambda)\langle x, Ax\rangle$. We can thus rewrite this whole expression as

$$= (\text{Im}(\lambda)^2 + \text{Re}(\lambda)^2)||x||^2 - 2\text{Re}(\lambda)\langle x, Ax\rangle + ||Ax||^2 = \boxed{|\text{Im}(\lambda)|^2||x||^2 + ||\text{Re}(\lambda)I - A)x||^2}.$$

So now if the imaginary part of λ were nonzero, then we would have $||x|| \le C||(\lambda I - A)x||$, so the operator is closed range and injective. Doing the same thing for the adjoint, which is $\lambda^* I - A^*$ (which will also have a nonzero imaginary part), we see that the adjoint also has closed range and injective. So then $\lambda I - A$ is bijective, hence invertible by the open mapping theorem, and thus anything not in $\mathbb{R}$ is in the resolvent. □

So now if A is bounded and self-adjoint, $\sigma(A)$ is some bounded subset of $\mathbb{R}$, and we can now decompose the spectrum in a different way. We've mentioned wanting to define $f(A)$, and we'll do so in steps:

- If $p(x)$ is a polynomial, then we define $p(A)$ by expanding it out – in particular if the coefficients are real then $p(A)$ is another self-adjoint operator.
- Now for any p, q, we have $||p(A) - q(A)|| \le \sup_{x\in\sigma(A)} |p(x) - q(x)|$, so for any continuous function continuous on a neighborhood of $\sigma(A)$, we can take a sequence of real polynomials p_n with $p_n \to f$ in C^0 near $\sigma(A)$ by Weierstrass. So then we can define $f(A)$ to be the operator norm limit of the $p_n(A)$s.

The interesting way to think about the spectrum is now the following: fix $u, v \in X$ in our Hilbert space, and consider the number $\langle f(A)u, v\rangle$ which varies in f. Then $f \mapsto \langle f(A)u, v\rangle$ is a linear functional that depends continuously on f, so we can define a measure $\mu_{u,v}$ such that this mapping is integration $f \mapsto \int f d\mu_{u,v}$ (this is the Riesz representation on C^0). And now we can break up the measure into the atomic parts and the continuous parts, and thus the spectrum decomposes instead into

$$\sigma(A) = \sigma_{\text{pp}}(A) \cup \sigma_{\text{ac}}(A),$$

where the pure point spectrum corresponds to $d\mu$ being atomic, σ_{ac} is the part that's absolutely continuous with respect to the Lebesgue measure, and σ_{sc} is the singular part. And remember now that, for example, the atomic part of the measure can overlay on top of the continuous part. Our end goal will be to define a measure $d\mu_A$ on $\mathbb{R}$ so $d\mu_A(\lambda)$ meant to encode all of the spectral information below λ:

$$I = \int_{-\infty}^{\infty} d\mu_A(\lambda), \quad A = \int_{-\infty}^{\infty} \lambda d\mu_A(\lambda), \quad f(A) = \int f(\lambda) d\mu_A(\lambda).$$

So $\langle d\mu_A(\lambda)u, v\rangle$ can be thought of as a projection-valued measure for any λ, and the general spectral theorem says we can do such a thing. So in summary, there's a way of decomposing A as a sum over its "eigenspaces," but now we're continuously decomposing into $d\mu_A(\lambda)$s.

16 March 2, 2023

This lecture is being given by Romain Speciel. The topic for today is the **spectral theorem** – we've discussed properties of operators (self-adjointness, compactness, normality, etc.), and the point is that self-adjoint compact operators are basically scalings. (So we go from two nice, but not very strong, properties to something much more powerful when those assumptions are placed together.)

Theorem 159 (Spectral theorem for compact self-adjoint operators)
Let H be a compact Hilbert space, and let A be a self-adjoint, compact complex linear operator on $\mathcal{H}$. Then there exists an orthonormal basis of $\mathcal{H}$ consisting of eigenvectors of A, such that the eigenvalues converge to zero.

Recall that we've already proved last time that the eigenvalues of a self-adjoint operator are real.

Proof. We're going to once again use Rayleigh quotients here, and the point is that our assumptions will allow us to actually achieve the maximums and use that to compute our eigenvectors. We take the definition

$$R(x) = \frac{\langle x, Ax\rangle}{\langle x, x\rangle},$$

and our goal is to maximize this. We'll just consider the value of $|R(x)|$ on the unit sphere S (since we can always rescale a nonzero vector) and maximize $\langle x, Ax\rangle$. The best way to get a maximum is to get a continuous function on a compact space, but in this case S is not compact in the strong topology. On the other hand, S **is** weakly compact (at least look at the unit ball) by Banach-Alaoglu and reflexivity (here's where we use the Hilbert space structure), and R is continuous from $S_{\text{weak}} \to \mathbb{R}$. (Here the point is that a **compact** operator upgrades weak convergence to norm convergence – for more justification, A being compact is equivalent to $A : \mathcal{H}_{\text{weak}} \to \mathcal{H}_{\text{norm}}$ if we just unpack the definitions, which implies that $\langle x, Ax\rangle$ is continuous.)

So the point is that there is some $x_1 \in S$ which maximizes $|R(x)|$, and we claim that this value is an eigenvalue of A. Indeed, if $R(x_1) = 0$, then $A = 0$ (because if $\langle Ax, x\rangle = 0$, we can expand $0 = \langle A(x+y), x+y\rangle = \langle Ax, y\rangle + \langle x, Ay\rangle$ to find that $\langle Ax, y\rangle = 0$ for any x, y, which means $A = 0$). Otherwise, we can compute variations: take v such that $\langle v, x_1\rangle = 0$ (since we need to vary along the sphere) and compute

$$0 = \frac{\partial}{\partial\varepsilon}\bigg|_{\varepsilon=0} R(x_1 + \varepsilon v) = \frac{\partial}{\partial\varepsilon}\bigg|_{\varepsilon=0} \langle A(x_1 + \varepsilon v), x_1 + \varepsilon v\rangle = \frac{\partial}{\partial\varepsilon}\bigg|_{\varepsilon=0} (\varepsilon\langle Av, x_1\rangle + \varepsilon\langle Ax_1, v\rangle)$$

since the first term doesn't vary in ε and the last term has an ε^2 factor. But now because A is self-adjoint, we can move the A over to the other side to find that in fact

$$0 = \langle v, Ax_1\rangle + \langle Ax_1, v\rangle = 2\text{Re}\langle Ax_1, v\rangle.$$

But since we're doing this for arbitrary v and can scale by i as well, this means $\langle Ax_1, v\rangle = 0$ for any v perpendicular to x_1. Thus Ax_1 must be $\lambda_1 x_1$ for some λ, since the only part nonzero is the part along x_1. And then in fact we know that $R(x_1) = \lambda_1$ and $Ax_1 = \lambda_1 x_1$.

So now we can do this process inductively: let $\mathcal{H}_n$ be the orthogonal complement of $\text{span}(x_1, \cdots, x_n)$, and let A_n be the operator A restricted to $\mathcal{H}_n$. Then either A_n is eventually zero, which means we're done because on the orthogonal complement we can just take any basis and all of those will be eigenvectors of eigenvalue zero. But otherwise we get a sequence $(x_1, x_2, \cdots)$ of orthonormal eigenvectors, though we don't know that this is a basis yet. However, we claim that the λ_is go to zero by compactness of A – indeed, the sequence of norm-one vectors x_i converges to zero weakly, so $Ax_i = \lambda_i x_i$ must go to zero in norm, meaning that $\lambda_i \to 0$.

The idea now is to consider

$$\tilde{H} = \overline{\text{span}(x_i)}^{\perp},$$

the orthogonal complement of the closure of the span of all of the x_is in our sequence, and let $\tilde{A}$ be A restricted to this space. Our goal is to show that in fact $\tilde{H} = 0$. Now consider $\sup_{||x||=1, x \in \tilde{H}} \left|\langle x, \tilde{A}x\rangle\right|$; we know this is at most λ_i for all i (because we picked out the maximum at stage i), so because $\lambda_i \to 0$ we have $\tilde{A} = 0$. So on the complement of the span of the x_is, the operator A is just zero, meaning that we can indeed complete the basis with a bunch of eigenvectors of eigenvalue zero. □

Remark 160. *There are kind of two senses in which we can have a basis: a linearly independent set such that every vector is a finite linear combination of them, or one where every vector is the limit of such combinations. And these bases do not need to be countable, but this argument tells us that the orthogonal complement of the kernel of a compact operator is necessarily separable (since there are only countably many eigenvalues).*

Restricting ourselves to Hilbert spaces imposes a lot of structure, so we may ask the same question for just a general Banach space. That's what our next result will do:

Theorem 161

Let X be an infinite-dimensional Banach space, C a compact operator on X, and let $\sigma(C)$ be its spectrum. Then the following hold:

1. Every nonzero $\lambda \in \sigma(C)$ is an eigenvalue (so in fact the spectrum is just the set of eigenvalues, plus potentially zero, and we don't have the situation where $C - \lambda I$ is not injective but not surjective).
2. If we define $E_\lambda = \bigcup_{n=1}^{\infty} \ker((\lambda - C)^n)$, then there is some $m \in \mathbb{N}$ with $E_\lambda = \ker(\lambda - C)^m$.
3. The eigenvalues can accumulate only at zero – in particular, $\sigma(C)$ is countable.

The idea with (2) is that $\lambda - C$ is not injective, and if we exponentiate it the kernel can usually get larger and larger but in this case will always stabilize. And we don't have the decomposition of our Banach space in the same way as before, but that's expected because we have less structure. The proof has a different flavor, but we'll again use compactness:

Proof. For **(1)**, we may take $\lambda = 1$ without loss of generality by rescaling C. Suppose for the sake of contradiction that 1 is not an eigenvalue but it is in the spectrum, so $I - C$ is injective but not surjective. Define $Y_1 = \text{Ran}(I - C) \subsetneq X$ – this is a closed subspace, and here we're using that C being compact implies that $I - C$ has closed range. (We do this by taking a convergent sequence and then passing to a subsequence where Cx_n converges, then think about the cases where $||x_n|| \to \infty$ or $||x_n||$ is bounded.)

We can then define $Y_n = (I - C)Y_{n-1}$ for each n. The point is that $X \supsetneq Y_1 \supsetneq Y_2 \supsetneq \cdots$ (they are strict inclusions – $I - C$ is injective on each Y_{n-1}, and if it were surjective we could pass that back up to all of X). But now we can use the Riesz lemma (which says that in a proper closed subspace, we can get far-away points from that subspace) to produce unit-norm $y_n \in Y_n$ such that $\text{dist}(y_n, Y_{n+1}) > \frac{1}{2}$. Since C is compact, the sequence (Cy_n) has a convergent subsequence; passing to that subsequence, we see that if $n < m$, we have

$$||Cy_n - Cy_m|| = ||(C - I)y_n + y_n - (C - I)y_m - y_m||,$$

and now $(C - I)y_n - (C - I)y_m - y_m \in Y_{n+1}$ and thus the quantity above is at least $\frac{1}{2}$ (it's the distance between y_n and something in Y_{n+1}). So this is a contradiction with the Cy_ns being Cauchy. So 1 is indeed an eigenvalue. The point is that we cannot have a "infinite nullity," since sequences of nested spaces and compactness gives us easy ways to construct sequences of vectors.

Next, **(2)** will be a very similar proof: let $Z_m = \ker((\lambda - C)^m))$, and suppose that this sequence does not stabilize. Then we have $Z_1 \subsetneq Z_2 \subsetneq Z_3 \subsetneq$, a strict inclusion of closed subspaces (because these are all kernels of linear operators), and we can pick unit-norm $z_n \in Z_n$ so that $d(z_n, Z_{n-1}) > \frac{1}{2}$ for all n. Passing again to a subsequence, we know that (Cz_n) converges, and then we get the same contradiction as before. Finally, for **(3)**, if there were some accumulation point not at zero, then there would be some $\varepsilon > 0$ such that there are infinitely many λ_is with $|\lambda_i| > \varepsilon$. Letting $W_n = \text{span}(x_1, \cdots, x_n)$, where x_i is the eigenvector of λ_i, we again get a sequence of nested subspaces and are able to take the Riesz lemma step for $|\lambda_i| > \varepsilon$. Countability of eigenvalues then follows because there are finitely many of norm at least $\frac{1}{n}$ for any n. □

We'll spend the rest of the class stating some interesting properties of the spectrum:

Theorem 162 (Spectral radius formula)
Let X be a Banach space, and let A be a complex linear operator on X. Then $\lim_{n\to\infty} ||A^n||^{1/n} = \sup_{\lambda \in \sigma(A)} |\lambda|$, and in particular the limit exists.

We probably won't have time to prove this – one half of this has been shown already, but the other direction is much harder. The idea is that $||A^n|| \le ||A||^n$, but this equality case is often not attained, and the spectrum only gets big depending on how badly this inequality fails. (For example, consider the shift operator where we scale the first number by 2. This has norm 2, but any power of it still has norm 2, so that tells us that the spectral radius is 1.)

Theorem 163 (Spectrum of normal operators)
Let $\mathcal{H}$ be a Hilbert space, and let A be a complex linear operator which is **normal** (meaning $A^*A = AA^*$). Then $||A^n|| = ||A||^n$, and in particular $\sup_{\lambda\in\sigma(A)} |\lambda| = ||A||$.

Proof. We have for any unit-norm x that

$$||Ax||^2 = \langle Ax, Ax\rangle = \langle x, A^*Ax\rangle \le ||A^*Ax|| = ||A^2x||,$$

where the last step comes from taking $\langle A^*Ax, A^*Ax\rangle$ and moving the A^*s across the inner product. So we know that $||A^2|| \le ||A|^2 = \sup_{||x||=1} ||Ax||^2 = \sup_{||x||=1} ||A^2x|| = ||A^2||$, so we must have equality everywhere and $||A^2|| = ||A||^2$. This means that $||A^{2^m}|| = ||A||^{2^m}||$ for any integer m. But now choosing m large enough so that $n < 2^m$, we have

$$||A||^{2^m-n}||A||^n = ||A||^{2^m} = ||A^{2^m-n}A^n|| \le ||A^n||||A||^{2^m-n},$$

and then dividing by $||A||^{2^m-n}$ on both sides yields the result for general n. □

Theorem 164 (Spectrum of self-adjoint operators)
Let $\mathcal{H}$ be a complex Hilbert space, and let A be a self-adjoint operator on $\mathcal{H}$. We have already proved that $\sigma(A) \subset \mathbb{R}$. Then **(1)** $\sup \sigma(A) = \sup_{||x||=1}\langle x, Ax\rangle$ (with no absolute value on the right-hand side), **(2)** $\inf \sigma(A) = \inf_{||x||=1}\langle x, Ax\rangle$, and **(3)** $||A|| = \sup_{||x||=1} |\langle x, Ax\rangle|$.

In particular, we can figure out the full range in which the spectum is contained, not just a bound in terms of its absolute value.

Proof. For **(1)**, we may first shift our operator A to $A + \lambda I$ so that $\langle x, Ax\rangle \geq 0$ for all x (that is, A is positive). **We claim that** $\sigma(A) \subseteq [0, \infty)$ and that $||A|| = \sup_{||x||=1}\langle x, Ax\rangle$; in particular this is sufficient by the spectral radius formula because the norm is then equal to the supremum of $|\sigma(A)|$, which is actually $\sup \sigma(A)$ because $\sigma(A)$ only takes values in $\mathbb{R}^+$. From there, we can take A and replace it with $-A$ to get **(2)**, and then we can combine the statements of **(1)** and **(2)** to get **(3)**.

To prove the claim, let $\varepsilon > 0$, and notice that

$$\varepsilon||x||^2 = \langle x, \varepsilon x\rangle \leq \langle x, \varepsilon x + Ax\rangle \leq ||x||\,||(\varepsilon + A)x||.$$

Rearranging this means that $\varepsilon||x|| \leq ||(\varepsilon + A)x||$, so $\varepsilon + A$ is injective and has closed range. Then $\overline{\mathrm{im}((\varepsilon + A)x)} = \ker(\varepsilon + A)^{\perp} = \mathcal{H}$ (since $\varepsilon + A$ is injective), so in fact this means $\varepsilon + A$ is bijective and thus $-\varepsilon$ is not in the spectrum. Thus the spectrum is positive. The other part of the claim is more painful: we can define $a = \sup_{||x||=1}\langle x, Ax\rangle \geq 0$. Then for any $||x|| = 1$ we have

$$\langle x, Ax\rangle \leq ||x||\,||Ax|| \leq ||A||,$$

so $a \leq ||A||$. For the other inequality, we notice that for any $x, y \in \mathcal{H}$, we have

$$\mathrm{Re}(\langle x, Ay\rangle) = \frac{1}{4}\left(\langle x + y, A(x + y)\rangle - \langle x - y, A(x - y)\rangle\right),$$

but now both inner products are nonnegative and thus we have

$$-\frac{1}{4}\langle x - y, A(x - y)\rangle \leq \mathrm{Re}(\langle x, Ay\rangle) \leq \frac{1}{4}\langle x + y, A(x + y)\rangle.$$

Thus if $||x|| = ||y|| = 1$, we find that $||x - y|| \leq 2$ so

$$-a \leq -\frac{a}{4}||x - y||^2 \leq -\frac{1}{4}\langle x - y, A(x - y)\rangle.$$

This last step comes from the fact that if we divide the right-hand side by $||x - y||^2$, we get an expression bounded by a by definition. But then we can continue this chain of inequalities to get $a \leq \mathrm{Re}(\langle x, Ay\rangle)$, and similarly we can find that this is bounded from above by $\frac{1}{4}\langle x + y, A(x + y)\rangle \leq \frac{a}{4}||x + y||^2 \leq a$. So indeed $|\mathrm{Re}(\langle x, Ay\rangle)| \leq a$, and taking the supremum over all $||x|| = ||y|| = 1$ yields $||A|| \leq a$, so $a = ||A||$. $\square$

17 March 7, 2023

We'll continue discussing spectral theory today – we won't do the full generalized spectral theorem, but we'll build up to some interesting cases and see how everything plays out in situations where we don't just have a discrete set of eigenvalues. Last time, we discussed self-adjointness and compact operators in the Hilbert and Banach settings – specifically, to find the spectrum of a compact, self-adjoint operator A on X, we can consider the Rayleigh quotient $R(x) = \frac{||Ax||^2}{||x||^2}$ and look for a maximizer x_1 (compactness allows us to extract an L^2-convergent subsequence easily); we can then show that $\frac{d}{d\varepsilon}\big|_{\varepsilon=0} R(x_1 + \varepsilon y) = 0 \implies Ax_1 = \lambda_1 x_1$, and now we can repeat this process in the orthogonal complement. The only issues are that (1) we need to check that the λ_js go to zero, but this again comes from compactness of A (since our eigenvectors are all orthogonal), and (2) we need to check if the x_js with nonzero eigenvalues span the whole space, and in fact they form the orthogonal complement to the kernel of A. And the argument on Banach spaces is very different – the idea is that orthogonality helps a lot in Hilbert spaces, and so does the fact that $(A - \lambda I)^n x = 0 \implies (A - \lambda I)x = 0$ (which isn't true in general Banach spaces).

Our goal is to now use the theory of **holomorphic functions** in spectral theory – we will need to use some facts in complex analysis here. Let X be a Banach space, and suppose we have some domain $\Omega \subset \mathbb{C}$ where $z \in \Omega$ is sent to $A(z) \in \mathcal{B}(X)$. There are many different ways we can say that A is "holomorphic in z," which will in fact all be equivalent. The usual (**strongly holomorphic**) requirement is that **(1)** the derivative

$$\lim_{h \to 0} \frac{A(z+h) - A(z)}{h}$$

exists regardless of the direction in which h approaches zero. We can also have a **weakly holomorphic** requirement that **(2)** for any $x \in X$ and any $\ell \in X^*$, the map

$$z \mapsto \ell(A(z)x)$$

is a holomorphic function (in the ordinary sense). And then there's another potential requirement that **(3)** $A(z)$ has a power series expansion. The idea is that in ordinary complex analysis, the Cauchy integral formula makes this third assumption equivalent to the first one, and there's no difference between the first two because we are just dealing with numbers.

Remark 165. *Here, the Cauchy integral formula says that for a holomorphic function $f(z)$ and a curve $\gamma = \{w : |w - z_0| = r\}$, we have $f(z) = \frac{1}{2\pi i} \int \frac{f(w)}{w-z} dw$. We can evaluate this integral with a reparameterization $w = z_0 + re^{i\theta}$, upon which we just get a line integral. But the point from there is that we can use the power series expansion $\frac{1}{w-z} = \frac{1}{w\left(1 - \frac{z}{w}\right)} = \frac{1}{w} \sum_{j=0}^{\infty} \left(\frac{z}{w}\right)^j$, and plugging this in (and using that whenever $|z| \leq \frac{1}{2} r$ this is an absolutely convergent series, so we can interchange the integral and limit) gives us a power series expansion for $f(z)$ in a neighborhood of z.*

So if we look at our function $\ell(A(z)x)$, saying that this is holomorphic means that we have such a power series expansion of the form $\sum_{j=0}^{\infty} a_j (z - z_0)^j$, and these a_js must depend linearly on both ℓ and x so we can find the coefficients as well. But if $f(w)$ is taking on values in the **Banach** space X, we can still have this kind of integral representation, since all of the integrals we're doing are nice and are basically Riemann integrals (and addition and scaling makes sense in X).

Proposition 166

All of the holomorphicity requirements above are equivalent, and they are also equivalent to **(4)** $\ell(A(z)x)$ having a power series (in fact being represented by its Cauchy integral formula).

Proof. Based on everything we've already said, it suffices to prove that if (4) holds for all ℓ and all x, then (1) holds. If we define

$$B(w)x = \frac{1}{2\pi i} \int \frac{A(z)x}{(z-w)^2} dz$$

and are claiming that this should be our derivative, then we can choose some z_0 and some small $r > 0$ and consider the set $\{w : |w - z_0| \leq r\}$. Then letting $c = \sup_{|z - z_0| = r} ||A(z)||$, we can check that the difference quotient limit exists by showing that it is actually equal to B. Indeed, plugging in the integral representations,

$$\begin{aligned} \left|\left| \frac{A(w+h) - A(w)}{h} - B(w) \right|\right| &= \ell\left(\left(\frac{A(w+h) - A(w)}{h} - B(w) \right) x \right) \\ &= \frac{1}{2\pi i} \int \ell\left(\frac{A(z+h) - A(z)}{w - z - h} - \frac{A(z)}{w - z} - \frac{A(z)}{(z-w)^2} \right) x, \end{aligned}$$

and if we do the common denominator calculations we have $\frac{1}{h}\left(\frac{1}{z-w-h} - \frac{1}{z-w}\right) = \frac{1}{(z-w)^2} = \frac{h}{(z-w)^2(z-w-h)}$, and this is bounded by $\frac{Cr|h|}{(r-|w|^2(r-|w|-|h|)}$. So the point is that we get an upper bound for the difference between the difference quotient and $B(w)$ and thus the derivative does exist. $\square$

So the surprising thing is that taking a weak starting assumption (only requiring a derivative after pairing with both ℓ and x) actually gives us a strong derivative. That means we can indeed take power series of operators by first pairing with ℓ and x, considering the corresponding power series, and working with those.

Our next task is to prove the **spectral radius formula**: recall that this is given by the formula $r_A = \lim ||A^n||^{1/n}$.

Proof of spectral radius formula. We already did the easy case last time: if $|\lambda| > r_A$, then $r_{\lambda^{-1}A} = |\lambda|^{-1} r_A < 1$, so $\mathrm{Id} - \lambda^{-1}A$ is invertible and thus $\lambda I - A = \lambda(I - \lambda^{-1}\lambda)$. Thus $\sup |\lambda| \le r_A$.

The other direction takes more work – we use the variable $z = \frac{1}{\lambda}$

$$\Omega = \{z : z = 0 \text{ or } z^{-1} \in \rho(A)\},$$

and now we can define a function $R(z) : \Omega \to \mathcal{B}(X)$ by sending 0 to 0 and sending z to $(z^{-1}I - A)^{-1} = z(I - zA)^{-1}$, which is holomorphic on $\Omega \setminus \{0\}$. We know that $\sigma(A)$ contains $\{z : |z| < r_A^{-1}\}$ from the first part of our argument, and now we can expand

$$R(z) = z(I - zA)^{-1} = \sum_{k=0}^{\infty} z^{k+1} A^k,$$

telling us that any derivative $R^{(n)}$ is holomorphic in $\{|z| < r_A^{-1}\}$ for all n. So now suppose that $\sup |\lambda| < r_A$ (for the sake of contradiction). Then $\overline{B}_r(0) \subset \Omega$, so $\lambda(R(z)x) = \sum_{k=1}^{\infty} \ell(A^{k-1}x)2^k$ by reinterpreting our formula for $R(z)$. In particular, we have

$$\ell(A^{n-1}x) = \frac{1}{n!}\frac{d^n}{dz^n}\ell(R(z)x) = \frac{1}{2\pi i}\int \frac{\ell(R(z)x)}{z^{n+1}} dz,$$

since every differentiation gives us a factor of z when we use Cauchy's integral formula. And now if we have a disk D with sufficiently large radius r (which we can choose to be any value larger than $\sup |\lambda|$), then we get an integral representation

$$A^n = \frac{1}{2\pi i}\int_D \frac{R(z)}{z^{n+1}} dz = \int_0^1 \frac{R(\gamma(t))}{\gamma(t)^{n+1}} dt,$$

meaning that $||A^n|| \le r^{n+1} \sup_{|\lambda|=r} ||\lambda I - A||$. Letting this supremum be c, we have $||A^n|| \le r^n(rc)$, so $||A^n||^{1/n} \le r(rc)^{1/n}$, which does converge to r as $n \to \infty$. Thus $r_A \le r$ and therefore $r_A \le \sup |\lambda|$.

Finally, to show that the spectrum is not empty, the main point is that if it were, then $(\lambda I - A)^{-1}$ must be entire. For ordinary numbers we would then be able to use the maximum modulus principle, so we mirror that here: we can choose $\ell(A^{-1}x) = -1$, and now we can define an entire linear function

$$f(\lambda) = \ell((\lambda I - A)^{-1}x) \implies f(0) = 1$$

which is nonzero and thus nontrivial. Then for sufficiently large λ, we have $|f(\lambda)| \le \lim_{\lambda\to\infty} \frac{||\ell|| \, ||x||}{|\lambda| - ||A||}$, which goes to zero, so the maximum modulus principle tells us that it must be zero everywhere, a contradiction. □

Example 167

We may ask whether we can have an operator whose spectrum is just a single point, and we can model after the **left-shift operator** (sending e_{j+1} to e_j). Basically, if we send e_{j+1} to $a_j e_j$ with coefficients a_j growing, then the origin will be the only point in the spectrum.

We can now introduce the general **functional calculus** – we know that the spectrum is a nonempty closed set and can generally be arbitrarily complicated. So now if we take such a nonempty bounded closed set σ_A and we take some $\mathcal{U}$ containing σ_A, we should be able to extract some information out of it.

Definition 168

Let f be holomorphic on U. Then we have (here $\gamma \subset \mathcal{U}/\sigma_A$, which we should think of as being a collection of loops going once around counterclockwise)

$$f(A) = \frac{1}{2\pi i}\int \frac{f(z)}{zI - A}dz = \frac{1}{2\pi i}\int_\gamma f(z)R_A(z)dz.$$

where notice that we have an operator A instead of just w in the denominator, since we can represent A as a power series $R_A(z)$.

We may now ask what happens if $f(z) = 1$, or $f(z)$ is a polynomial (since in those cases we know what should happen – a polynomial in A should just be represented in the ordinary way using multiplication, but we've defined something similar and want to check that it's equivalent). **We claim that**

$$\frac{1}{2\pi}\int_\gamma R_A(z)dz = I,$$

which we can interpret in the weak sense by pairing against some x, ℓ. (We should think about why this holds and we'll talk about it more next time.) But then we see that

$$A = \frac{1}{2\pi i}\int_\gamma AR_A(z)dz = \frac{1}{2\pi i}\int (A - zI + zI)R_A(z)dz = \frac{1}{2\pi i}\int (A - zI + zI)R_A(z)dz.$$

Then $(A - zI)R_A(z) = -I$ is holomorphic everywhere and thus integrates to zero, so we see that $\boxed{A = \frac{1}{2\pi i}\int zR_A(z)dz}$, and that's exactly the special case of the definition for $f(A)$ when we have a polynomial. And similarly we also have $A^n = \frac{1}{2\pi i}\int A^n R_A(z)dz$, so all polynomials agree with this new definition. Now there's really a list of facts that are good for us to know:

- $\sigma_{f(A)} = f(\sigma_A)$,
- $g(f(A) = (g \circ f)(A)$ (this involves messing around with the integral expressions).
- If $\sigma_A = \sigma'_A \sqcup \sigma''_A$ with $\sigma'_A \subset U'$, σ_A'' $\subset U''$, and $U' \cap U'' = \varnothing$, then if we define

  $$P_1(A) = \frac{1}{2\pi i}\int_{\gamma_1} R_A(z)dz, \quad P_2(A) = \frac{1}{2\pi i}\int_{\gamma_2} R_A(z)dz$$

  (the picture here is that we decompose our space into two subspaces), then P_1, P_2 are projectors with $P_1^2 = P_1, P_2^2 = P_2$, $P_1P_2 = P_2P_1 = 0$, and if $X_j = P_j(X)$ then $X_1 \oplus X_2 = X$.

So the point is that functoinal calculus leads us to projectors – the idea is that the integral around a single point is just the projection onto the eigenspace. And the fact we use here is that we get back the identity operator if we integrate once around each point in winding number. We can check some of these facts: for example, if we have our two projectors, we then have

$$P_1P_2 = \frac{1}{(2\pi i)^2}\int_{\gamma_1} R_A(z_1)\int_{\gamma_2} R_A(z_2)dz_1dz_2,$$

and now we may use the resolvent identity:

$$(z_1 - z_2)R(z_1)R(z_2) = R(z_1) - R(z_2) \implies \frac{1}{(2\pi i)^2}\int_{\gamma_1}\int_{\gamma_2} \frac{R(z_1)}{z_1 - z_2} - \frac{R(z_2)}{z_1 - z_2}dz_1zd_2.$$

But now integrating the first term around γ_2 yields zero, and similarly integrating the second term will vanish. (We might worry about a situation where the two sets are repeatedly nested within each other, but then we have to check the winding number definition and see that everything works.) And checking that $P_1^2 = P_1$ is a similar manipulation. The point is that holomorphicity is a technical tool, but it lets us define functions of operators, and it agrees with our original definition if f is just a polynomial, and it lets us separate out our spectrum.

Our goal is then to do this in more specificity when we discuss self-adjoint operators in Hilbert spaces, and that's the last major topic of the course. The question we care about is how much information we can extract, and the only case where we actually get an eigenfunction decomposition is in the case where we have a self-adjoint operator. So from now on we'll assume X is Hilbert and that A is self-adjoint (recall that one of the ways to think about this is that $\langle x, Ax\rangle \in \mathbb{R}$ for all x).

Proposition 169

Let P be a polynomial. Defining $p(A)$ in the ordinary polynomial sense, we have $\sigma(p(A)) = p(\sigma(A)$. Then if A is normal (meaning $AA^* = A^*A$), then $||p(A)|| = \sup_{\lambda\in\sigma(A)} |p(\lambda)|$.

Proof. If $\lambda \in \sigma(A)$, then $p(z) - p(\lambda)$ vanishes when $z = \lambda$, so we can factor this out as $(z - \lambda)q(z)$ for some polynomial q. Then

$$p(\lambda)I - p(A) = (\lambda I - A)q(A) = q(A)(\lambda I - A),$$

so if λ is in the spectrum, we know the middle term is not bijective, so the left term cannot be bijective either, and thus $p(\lambda) \in \sigma(p(A))$. Thus $p(\sigma(A)) \subset \sigma(p(A))$. For the other inclusion, if $\lambda \notin p(\sigma(A))$, then defining $q_\lambda(z) = \frac{1}{\lambda - p(z)}$, we have $q_\lambda(A)(\lambda I - A) = I$ and thus $\lambda \notin \sigma(p(A))$, so we can indeed just invert the argument from above. $\square$

So now if we consider the algebra $\mathcal{A} = \{p(A), p \text{ polynomial}\}$ for A normal (the reason we immediately pass to normality is that we can write down an orthonormal basis of eigenvectors with $Ax = \sum \lambda_i \langle x, e_i\rangle e_i$), we may ask whether this is closed in any reasonable sense. But the estimate $||p(A)|| = \sup|\sigma(p(\lambda))|$ is exactly what we need: consider the map from the set of polynomials to $\mathcal{A}$ sending p to $p(A)$. Then in the topology of the space of continuous functions on $\sigma(A)$, if we restrict to polynomials this map $p \to p(A)$ is a continuous function. And so this implies that for a continuous function f on $\sigma(A)$, if f is the uniform limit of some polynomials p_n then we can just define $f(A)$ to be the limit of the $p_n(A)$s, since $||p_n(A) - p_m(A)|| \le \sup|p_n(\lambda) - p_m(\lambda)|$; the last step is then the Stone-Weierstrass theorem, which shows that every continuous function is a uniform limit of polynomials. So the point is that any function of A is defined in this way, and we'll see some applications of this later on.

18 March 9, 2023

Last time, we looked at the resolvent $R_A(\lambda) = (\lambda I - A)^{-1}$ of a bounded operator A on a Banach space X. We prove that $R_A(\lambda)$ is holomorphic in $\mathbb{C} \setminus \sigma(A)$, and we used this to define the holomorphic functional calculus – that is, if $U \subset \mathbb{C}$ is a set containing $\sigma(A)$, and γ is a path in U and f is a holomorphic function on U, then we can define via the Cauchy integral formula

$$f(A) = \frac{1}{2\pi i}\int_\gamma f(\lambda)R_A(\lambda)d\lambda.$$

(Making sense of what this means is just going through the usual integration theory, since everything we're integrating is nice.) The idea now is that if $\sigma(A)$ is some bounded set and γ is a loop that goes around the set once, then $\frac{1}{2\pi i}\int_\gamma R_A(\lambda)d\lambda$ should yield the identity operator. Indeed, we know that $\lambda I - A = \lambda(I - \lambda^{-1}A)$, and now $\lambda^{-1}A$ has small norm so we can expand $R_A(\lambda) = \lambda^{-1}(I - \lambda^{-1}A)^{-1}$ as a power series $\sum_{j=0}^\infty \lambda^{-j-1}A^j$ which is a norm-convergent sum. Thus we can swap the sum and integral, and when we integrate we just pick up the residue at $j = 0$ which is

1. This same consideration tells us that $\frac{1}{2\pi i}\int_\gamma \lambda R_A(\lambda) = A$ (since we just pick up the next lowest power of A), and repeating this tells us that $\frac{1}{2\pi i}\int_\gamma p(\lambda)R_A(\lambda)d\lambda = p(A)$.

But another interesting case is where the spectrum $\sigma(A)$ decomposes into two disjoint closed clumps $\sigma_1(A), \sigma_2(A)$ (that is, the two sets are contained in disjoint open sets U_1, U_2). Then we can take the holomorphic function f which is 1 on U_1 and 0 on U_2 and set $f(A) = P_1$ to be the projector onto U_1. The point is that the projectors separate our Banach space into two pieces, and we can check that we do actually have projectors because

$$P_1^2 = \frac{1}{(2\pi i)^2}\int_{\gamma_1}\int_{\tilde{\gamma}_1} R_A(\lambda_1)R_A(\tilde{\lambda}_1)d\lambda_1 d\tilde{\lambda}_1,$$

where $\gamma_1, \tilde{\gamma}_1$ are both loops of winding number 1 around U_1. We can then use the resolvent identity to get an integrand $\frac{R_A(\lambda_1)-R_A(\tilde{\lambda}_1)}{\lambda_1-\tilde{\lambda}_1}$; one of the two terms in the numerator disappears (if we think of one of $\gamma_1, \tilde{\gamma}_1$ as being enclosed in the other) and the other yields the expression for P_1.

However, notice that we've been requiring our functions f to be holomorphic, and so we know very little about the spectral content. The only thing we really know is that if λ_0 is an isolated point, then letting $P_0 = \frac{1}{2\pi i}\int_{\gamma_0} R_A(\lambda)d\lambda$ be the integral around a loop containing only that isolated point, then

$$AP_0 = \frac{1}{2\pi i}\int_{\gamma_0} AR_A(\lambda)d\lambda = \frac{1}{2\pi i}\int_{\gamma_0}(A-\lambda_0 I)R_A(\lambda)d\lambda + \frac{\lambda_0}{2\pi i}\int R_A(\lambda)d\lambda.$$

We want to say that this is $\lambda_0 P_0$, but the problem is that when we write $(A-\lambda I)R_A(\lambda) = (A-\lambda_0)R_A(\lambda) + (\lambda - \lambda_0)R_A(\lambda)$, the second term may not be holomorphic if $R_A(\lambda)$ has a higher-order pole:

Example 170

Let A be a finite-dimensional matrix with 1s on and directly above the diagonal. Then $(A-\lambda I)^{-1}$ has $\frac{1}{1-\lambda}$s on the diagonal, but $-\frac{1}{(1-\lambda)^2}$s off the diagonal.

So the inverse of the resolvent indeed has a higher-order pole, and that's the **defect** of this setup: the idea is that we will have something that looks like

$$R_A(\lambda) = \sum_{j=1}^{N}\frac{P_j}{(\lambda-\lambda_0)^j} + Q(\lambda),$$

where $Q(\lambda)$ is regular but the first few terms are bad singularities (and in fact we may even have essential singularities) – in particular, anything in the range of P_N will be eigenvectors, but the others will be in some generalized eigenspaces. So even near isolated singularities we don't have any control related to holomorphicity.

Corollary 171

Suppose $R_A(\lambda)$ has a finite-order pole at λ_0 (so the N is finite), and the P_js are finite rank. Then if $A(\varepsilon)$ is a family of operators which is continuous in the operator norm, we may define the generalized eigenspace

$$V_\varepsilon = \bigcup_{j=1}^{\infty}\ker(A(\varepsilon)-\lambda_0 I)^j,$$

(we will only need to go up to some finite power of j because the kernel will eventually stabilize). Then V_ε is continuous in ε.

So in the matrix example above, we can modify the entries very slightly and suddenly make everything diagonalizable, or we can suddenly develop higher-order poles. But the point is that even if the "flag structure" switches around, the overall space will stay the same.

Proof idea. We can check that the resolvent $R_{A(\varepsilon)}(\lambda)$ is continuous in ε, we have $\frac{1}{2\pi i}\int R_{A_\varepsilon}(\lambda)d\lambda = P_0(\varepsilon)$ is continuous in ε, so something about the expansion of $R_A(\lambda)$ is continuous. $\square$

So we'll return to the Hilbert space world now – we still have all of the functional calculus from before, but we'll now consider self-adjoint (bounded for now) operators A. We know that $\sigma(A) \subset \mathbb{R}$, and this time we'll define $f(A)$ very differently. We started by doing this when $f \in C(\sigma(A))$, and we start by defining this for polynomials – here we'll just think about real-valued polynomials to avoid some technical details. We can define $p(A)$ in the obvious way, and the big fact is that $||p(A)|| = \sup_{\lambda\in\sigma(A)} |p(\lambda)|$. (Intuitively, if A were diagonalizable, we can check by hand that $p(A)$ has the same eigenvectors as A, but with eigenvalues transformed by p.)

The way to interpret this is to think about the following situation: let $\mathcal{C} = C^0(\sigma(A))$ be the set of continuous functions on $\sigma(A)$, which contains the space of polynomials, and there's a map from polynomials into $B(X)$ sending p to $p(A)$. Our goal is to show that p is dense:

Theorem 172 (Weierstrass)

The space of polynomials $\mathcal{P}$ is dense in $C^0(\sigma(A))$, so the mapping $p \mapsto p(A)$ extends to a map $C \to B(X)$ (so we can define $f \mapsto f(A)$ in general).

Corollary 173

Applying this to the function $f(t) = \frac{1}{\lambda - t}$, we have

$$||(\lambda I - A)^{-1}|| = \frac{1}{\text{dist}(\lambda, \sigma(A))}$$

In particular, this means the resolvant cannot do any worse than a simple pole for an isolated point.

We can state Weierstrass's result more abstractly in the following way:

Theorem 174

Let M be a compact Hausdorff space, and let $\mathcal{A}$ be a subalgebra of $C(M)$, such that $1 \in \mathcal{A}$, $\mathcal{A}$ separates points (that is, there is some function which takes different values on any $x \neq y$), and $\mathcal{A}^* = \mathcal{A}$ (closed under complex conjugation). Then $\mathcal{A}$ is dense in $C(M)$.

(This result can basically be applied directly to our set of polynomials, except that we have to mod out by an ideal if σ is a finite set.)

Proof of Theorem 174 in the polynomial case. We'll assume that we work with real-valued functions here. We know that there is a set of polynomials $p_n : [-1, 1] \to [0, 1]$ that approximates $|s|$ (that is, such that $p_n(s) \to |s|$ uniformly for all $|s| \leq 1$). The idea is to define $p_0(s) = s$ and then set $p_n(s) = \frac{s + p_{n-1}(s)^2}{2}$; we can prove by induction that $-1 \leq p_n \leq p_{n+1} \leq 1$ for all n, so the sequence converges pointwise to some $r(s)$. But r satisfies $r = \frac{s+r^2}{2}$, so $r = \frac{2\pm\sqrt{4-4s}}{2} = 1 + \sqrt{1-s}$ (we do need the positive root), and then $r(1 - s^2) - 1 = |s|$. Because $p_n(s) \in \mathcal{A}$, we must have $|s| \in \overline{A}$. (Remember that we're trying to prove that $\overline{A}$ contains all continuous functions, so this is just an important one for us to compute first.)

Now we claim that if $f \in \overline{\mathcal{A}}$, then $|f| \in \overline{\mathcal{A}}$ as well. We know that the normalized function $\frac{f}{||f||}$ is in $\overline{\mathcal{A}}$, and now by our previous step we can choose a polynomial p_n so that $||s| - p_n(s)| < \frac{\varepsilon}{||f||}$; this means that

$$\left| |f| - ||f|| p_n \left(\frac{f}{||f||} \right) \right| < \varepsilon,$$

and thus the absolute value function is also in the closure. Next notice that if we have $f, g \in \overline{\mathcal{A}}$, then $\max\{f, g\} = \frac{1}{2}(f + g + |f - g|)$ and $\min\{f, g\} = \frac{1}{2}(f + g - |f - g|)$ are both also in $\overline{\mathcal{A}}$. From here we can use the separating points axiom – given any $f \in \overline{\mathcal{A}}$ and $x, y \in M$, we have

$$g(z) = \frac{h(z) - h(y)}{h(x) - h(y)} f(x) + \frac{h(x) - h(z)}{h(x) - h(y)} f(y),$$

where h is some function where $h(x) \neq h(y)$ and $h \in \mathcal{A}$. So we can choose a family of such functions $\{g_{x,y}\}$ over all $x, y \in M$ (here we use the axiom of choice) and define

$$U_{x,y} = \{z : g_{x,y}(z) > f(z) - \varepsilon\}, \quad V_{x,y} = \{z : g_{x,y}(z) < f(z) + \varepsilon\}.$$

Then we have $M = \bigcup_x U(x, y)$ (since $U_{x,y}$ in particular contains x), so we can extract a finite subcover $\{g_{x,y}^{(1)}, \cdots, g_{x,y}^{(n)}\}$ by compactness, and then we can define $g_y(z) = \max\{g_{x,y}^{(n)}\}$, which will also be in $\overline{\mathcal{A}}$ because finite maximums are still within this set. Then if $V_y = V_{x_1,y} \cap \cdots \cap V_{x_N,y}$, then by definition $z \in V_y$ means that z lies in some $V_{x_i,y}$, so $g_y(z) < f(z) + \varepsilon$. Since the V_y are an open cover, we can again extract a finite subcover of them and define $g(z)$ to be the minimum of those g_{y_j}s, and that satisfies $\sup |f - g| < \varepsilon$.

So the point is that once we've done the work of showing we're closed under certain approximations, we can approximate any arbitrary function f within ε by using our open covers and taking maximums and minimums. □

We can define a map $\rho : C(\sigma(A)) \to \mathcal{F}_A \subset B(X)$, where $\mathcal{F}(A)$ is the smallest closed "$*$-algebra" containing I and A (that is, the closure of the set of polynomials of A), and this is in bijective correspondence with continuous functions. It turns out that this is the first step in the spectral theorem, and that's because we're really after the projectors – we're going to enlarge this now from continuous functions to Borel measurable functions, and the question is what measure to take.

Corollary 175

If $A \geq 0$ (meaning that $\langle Ax, x \rangle \geq 0$ for all x), there there exists a **unique** $B \geq 0$ such that $B^* = B$ and $B^2 = A$.

Proof. Existence is easy because we can take $f(\lambda) = \sqrt{\lambda}$ and let $B = f(A)$ with our general definition. Furthermore, since f is a real-valued function, we have $B = B^*$. For uniqueness, suppose we have some $C = C^*$ with $C \geq 0$ and $C^2 = A$. Then C commutes with A (since $CA = C^3 = AC$), so C commutes with B as well by approximating B in that equation by limits of polynomials of A. Now

$$(B+C)(B-C) = B^2 - C^2 = 0 \implies 0 = \langle (B-C)x, (B+C)(B-C)x \rangle = \langle (B-C)x, B(B-C)x \rangle + \langle (B-C)x, C(B-C)x \rangle.$$

But now both terms here are nonnegative since $B \geq 0, C \geq 0$, so both of them must in fact be zero. And then taking differences, we also find that

$$0 = \langle (B - C)x, (B - C)^2 x \rangle = \langle x, (B - C)^3 x \rangle,$$

and $(B - C)$ is self-adjoint so $||(B - C)^3|| = ||B - C||^3$, which means we must have $B = C$ if the inner product there is always zero.

□

We're now at the point where we can build up to understand the spectral theorem in generalization – functional calculus doesn't seem to have much to do with diagonalization so far, but now we'll see how it fits together. Fix $x \in X$ and define the linear functional Λ_x via

$$\Lambda_x(f) = \langle x, f(A)x \rangle.$$

This is a continuous linear functional and preserves positivity (since $f \geq 0$ means $\langle x, f(A)x \rangle = 0$), so by Riesz representation this must be integration against some Borel measure:

$$\Lambda_x(f) = \int f \, d\mu_x,$$

with $\text{supp}(d\mu_x) \subseteq \sigma(A)$. If the spectrum of A only has isolated bits, we would just have μ_x being a sum of terms of the form $a_j \delta_{\lambda_j}$, and thus $\int f d\mu_x$ is then $\sum a_j f(\lambda_j)$. But because $f(A)$ takes eigenfunctions e_j to $f(\lambda_j)e_j$, we can think about starting with $x = \sum c_j e_j$, so that

$$f(A)x = \sum c_j f(\lambda_j) e_j \implies \Lambda_x(f) = \langle x, f(A)x \rangle = \sum c_j^2 f(\lambda_j).$$

So in this simpler case the spectrum just moves pointwise, but more generally the measure will spread out.

The final step now is to ask what the set of Borel measurable functions $\mathcal{B}$ looks like: the idea behind a **projection-valued measure** is that for every open set U, we can define an operator P_U such that $P_U^2 = P_U$ and $P_U^* = P_U$. We obviously want $P_\varnothing = 0$ and $P_{\sigma(A)} = I$, and then we want $P_{U \cap V} = P_U P_V$. (So the picture behind this is that we have the set of all Borel sets on $\sigma(A)$ with a nice property of intersections – projections correspond to subspaces, and we want intersections of open sets to correspond to compositions of projectors and thus intersections of subspaces. This is called a **lattice structure**, and then the lattice of subspaces will have the same algebraic structure as the open sets.) And finally, we are asking that if $U_i \cap U_j = \varnothing$, and $U = \bigcup U_i$ is a countable union, then we have $P_U = \lim_{n\to\infty} \sum_{i=1}^n P_{U_i}$, with limit taken in the Hilbert space limit.

We can also polarize from $d\mu_x$ to $d\mu_{x,y}$, given by the formula

$$\langle y, f(A)x \rangle = \int f \, d\mu_{x,y},$$

so that $d\mu_{x,y}$ is now a signed Borel measure. And this all now relates to our picture in the following way: the space of continuous functions sits inside the space of Borel measurable functions, and there is a way to extend the Borel measure from C to $\mathcal{B}$ so that functionals can be extended. And we care about the case where we have an interval I – we can define P_I by taking the limit of a continuous family of monotone functions $f_n \in C$ converging downward to the indicator function $1\{x \in I\}$. We'll explain this more next time – if we take the interval $(-\infty, \lambda]$, and we let P_λ be the projection onto everything to the left of λ, increasing in that we project onto a continuously increasing family of subspaces, until eventually we project onto the whole space. And that's the analogue of the diagonalizability statement in the simpler cases we've discussed – what we'll find is that every self-adjoint operator is multiplication on L^2 of some measure space, and thus A just looks like multiplication by λ_i on each summand of $\bigoplus_{i=I} L^2(\mathbb{R}, d\mu_i)$.

19 March 14, 2023

Last time, we proved the Stone-Weierstrass theorem – the idea is that a subalgebra $\mathcal{A}$ of the set of continuous functions $C(M)$ on a compact Hausdorff space M is dense if $1 \in \mathcal{A}$, $\mathcal{A}^* = \mathcal{A}$, and $\mathcal{A}$ separates points. The proof is rather "disjoint from the rest of mathematics" – the idea when we have polynomials is to reduce to the case of real-valued functions and prove that we have a sequence $p_n(s)$ converging in norm to $|s|$ on $|s| \leq 1$ and thus that $p_n(f) \to |f|$. Thus if $\mathcal{A}$ contains f, it also contains $|f|$, and thus it contains minimums and maximums of finitely many elements of

$\mathcal{A}$. We'll now go through the last part of the argument in more detail – the idea is that for any fixed $x \neq y$ and any continuous function f, we can find some $g_{x,y} \in \overline{\mathcal{A}}$ with $g(x) = f(x)$ and $g(y) = f(y)$ by modifying a function which separates points. We can then approximate f by putting together these $g_{x,y}$s – the sets $U_{x,y}$ and $V_{x,y}$ (telling us when $g_{x,y}$ is within ε of f) can be extracted to a finite cover, and then taking maximums and minimums will make the result stay between $f - \varepsilon$ and $f + \varepsilon$.

After that, we considered **self-adjoint** bounded operators on a Hilbert space X – we've now developed tools, such as the spectrum $\sigma(A)$, spectral radius formula r_A, and resolvent $R_A(\lambda)$. In particular, to any $f \in C(\sigma(A))$ we can assign to it $f(A)$ as follows: let $q_n(s)$ be polynomials such that $q_n(s) \to f(s)$ in $C(\sigma(A))$ (which we can do because $\sigma(A)$ is compact), meaning that $||q_n(A)|| = \sup_{\lambda \in \sigma(A)} |q_n(\lambda)|$ and thus the $q_n(A)$s are convergent in operator norm to $f(A)$. So then $f(A)$ will be self-adjoint if f is real-valued, and in general any two functions $f(A), g(A)$ commute with each other with $f(A)g(A) = (fg)(A)$ and satisfying $f(g(A)) = (f \circ g)(A)$, as we should expect from such a recipe for constructing operators. So in the self-adjoint case what's nice is that we can take functions **directly** on the spectrum, while in the non-self-adjoint case we don't know what's happening on $\sigma(A)$. So the moral is that being on the spectrum has some abstract characterizations, but it is often challenging to work with, but self-adjointness is meant to give us a slightly better (but still abstract) description.

The idea is that we fix $x \in X$ and consider the functional $\Lambda_x(f) = \langle x, f(A)x \rangle$. We know this is a bounded, sign-preserving operator with $|\Lambda_x(f)| \leq C \sup |f|$, and thus this linear functional is integration against a measure by Riesz representation (this is the characterization of $C(\sigma(A))^*$). Polarizing then gives us signed Borel measures $d\mu_{x,y}$ (these can also be viewed as correlation functions between states x, y in quantum mechanics). So now that we have a measure, we can look at the Borel σ-algebra of open sets generated by open sets on $\sigma(A)$ and define $\langle x, f(A)y \rangle$ for all **Borel** measurable f. (So we have $f(A)$ at least in the weak sense.)

Checking that all of this takes a lot of work, because there's lots of measure theory involved, but the point is that it does work and we do have a way to define $f(A)$. So now our last topic for the quarter is basically "what does all of this mean" – we'll try to interpret this. For any set $\Omega \in \sigma(A)$, we have the indicator function $1\{x \in \Omega\}$, which is Borel measurable. Thus we can define $P_\Omega = \chi_\Omega(A)$; because $\chi^2 = \chi$ we have $P_\Omega^2 = P_\Omega$, and similarly $P_\Omega^* = P_\Omega$. Calculations also show us that $P_\varnothing = 0, P_{\sigma(A)} = I$, $P_{\Omega_1 \cap \Omega_2} = P_{\Omega_1} P_{\Omega_2} = P_{\Omega_2} P_{\Omega_1}$, and that if Ω is a disjoint union $\bigcup_{j=1}^\infty \Omega_j$, with $\Omega_i \cap \Omega_j = \varnothing$ for all $i \neq j$, then $\sum_{i=1}^N P_{\Omega_i}$ converges in norm to P_Ω. All of this, as we said last time, encodes a **projection-valued measure** sending Ω to P_Ω satisfying all of these facts, with

$$\langle x, P_\Omega y \rangle = \int \chi_\Omega d\mu_{x,y}.$$

Example 176

Suppose we consider the sets $\Omega_\lambda = (-\infty, \lambda] \cap \sigma(A)$ and let $P_\lambda = P_{\Omega(\lambda)}$.

The idea is that we're projecting onto the eigenspaces of eigenvalues at most λ, but there may not be any eigenvectors at all – for example, even if λ_0 is an isolated point in the spectrum, P_{λ_0} may still be zero. But either way, we have subspaces

$$H_\lambda = P_\lambda(X),$$

which are all closed subspaces in X that are "continually increasing."

Example 177

Next, consider $A = -\Delta$ on $L^2(\mathbb{R}^n)$ (this is self-adjoint up to some details with unbounded operators, which we won't get into).

The Fourier transform transforms this operator to multiplication by $|\xi|^2$, so the spectrum of A is everything on the positive real axis. (Here we're using that the Fourier transform is unitary and thus preserves spectrum, since

$$U(A - \lambda I)U^{-1} = (UAU^{-1} - \lambda I).$$

Indeed, we have

$$(\lambda I + \Delta)^{-1} = (\lambda I - (-\Delta))^{-1},$$

so if we take $\hat{f} \mapsto \frac{\hat{f}}{\lambda - |\xi|^2}$ we get a bounded operator in L^2. But otherwise – if λ is nonnegative real – it's easy to find sequences of functions f_n such that the ratio of the norms goes to infinity. So we do know that the spectrum is $\mathbb{R}_+$.

So now to compute the projection-valued measure, we have Ω some subset of the real line, and we can consider the sequence of maps

$$f \mapsto \hat{f} \mapsto \chi_\Omega(\xi)\hat{f} \mapsto \mathcal{F}^{-1}(\chi_\Omega(\xi)\hat{f}(\xi)).$$

The result is a C^∞ function, and the right-hand side can be anything in L^2 supported on Ω. So the range of P_Ω is the set of $h \in L^2$ so that $\hat{h}$ is supported in Ω; in particular, P_λ lies in the set of **band-limited functions** $H_\lambda = \{h : \hat{h} = 0 \text{ if } |\xi| > \lambda\}$. So in general the spectral theorem tells us that the unitary operator does something like this, but we don't always know what that unitary operator is (like with the Laplacian).

Theorem 178

Let A be a self-adjoint operator on a Hilbert space X. Then we can write a decomposition into an orthogonal direct sum $U : X \to \oplus_i L^2(\Sigma_i, \mu_i)$ with $\Sigma_i \subset \sigma(A)$, not necessarily disjoint (that is, we can decompose X into a possibly large set of subspaces), such that $(UAU^{-1}\phi)_i(\lambda)$ acts as multiplication by the function λ: $(UAU^{-1}\phi)_i(\lambda) = \lambda\phi_i(\lambda)$. (So the λs are the variables on which the spaces $L^2(\Sigma_i, \mu_i)$ are defined.)

Example 179

Suppose A is compact. Then we have the spectrum from the Rayleigh quotient perspective, where eigenvalues may only accumulate at zero.

Then P_λ is like taking the sum of the eigenspaces for all $\lambda_i < \lambda$, and now it's important whether we're doing the projection onto $(-\infty, \lambda)$ or $(-\infty, \lambda]$. But the point is that we have

$$P_\lambda = \sum_{\lambda_j \le \lambda} P_{\lambda_j},$$

so in this case every time we cross any nonzero eigenvalue we just gain a small space, and we have a complete basis of eigenfunctions.

Example 180

On the other hand, if $A = -\Delta$ is the Laplacian, then P_λ corresponds to Fourier localization.

Note that $P_{\{\lambda\}} = 0$, because in the Fourier interpretation the function is supported on $|\xi|^2 = \lambda$, which is of measure zero. However, we can consider thin annuli $P_{(\lambda-\varepsilon,\lambda+\varepsilon)}$, and this will project onto a space of smooth functions who are supported near $|\xi|^2 = \lambda$. So there are no solutions to $-\Delta u = \lambda u$ for $\lambda > 0$, but infinitely many such that $||(-\Delta - \lambda)u||_{L^2} \le \varepsilon ||u||_{L^2}$. These images are called **wavepackets** – the idea is that we can find nontrivial elements arbitrarily close but not exactly.

Of course, $u = e^{ix\cdot\xi}$, but the problem is that none of these are in L^2. And if we take more general tempered distributions, they must be linear combinations or integrals of such functions – anything of the form

$$\int_{|\xi|=\sqrt{\lambda}} \phi(\xi) e^{ix\cdot\xi} d\xi$$

will solve the equation, but they will always be of the rate $r^{-n/2}$. This leads us to **generalized eigenfunctions** – we can talk about "things that lie a little bit outside' our Hilbert space which generate our spectrum. Somehow smearing by changing the angle of ξ doesn't do anything, but smearing in the magnitude does allow us to get necessary decay for being in L^2.

We'll finish with one more application of the functional calculus – this is closely related to how when we have an operator $-\Delta + V$, it is often easiest to study the resolvent, because that is related to try to solve equations like $(-\Delta + V - \lambda)u = f$. And the point is that we have yet another relation between all of our different tools relating spectra.

Theorem 181 (Stone)

Let A have spectrum $\sigma(A)$ and resolvent $R_A(\lambda)$. Then for any function f, we have

$$\frac{1}{2\pi i}\int_a^b (R(\lambda + i\varepsilon) - R(\lambda - i\varepsilon)) f \, d\lambda \to \frac{1}{2}(P_{[a,b]} + P_{(a,b)})f.$$

Proof sketch. We can get resolvents by integrating: we have

$$(\lambda + i\varepsilon - A)^{-1} - (\lambda - i\varepsilon - A)^{-1} = \frac{1}{2\pi i}\int \left(\frac{1}{\lambda + i\varepsilon - t} - \frac{1}{\lambda - i\varepsilon - t}\right) R_A(t) dt.$$

The difference of fractions becomes $\frac{2i\varepsilon}{(\lambda-t)^2+\varepsilon^2}$, which is the Poisson kernel in the upper half-plane (thus harmonic there if we think of ε as the vertical coordinate) and such that we are the delta function at λ. So then integrating that against $R_A(t)$, we'll end up getting the function 1 on the interval and 0 elsewhere. The reason we have to average the closed and open projections is that if we have a point spectrum on the boundary of (a, b), we only get half of the projection onto those eigenspaces. □

The idea is that trying to understand spectra of operators like this is a big story that goes beyond mathematics – understanding the spectrum of $-\Delta + V$, whether it's composed of eigenfunctions on a dense subset or just forming a continuous spectrum, can be useful in quantum chemistry. And one of the big conjectures permeating various fields is **Anderson localization** – the idea is to take an operator with nice spectrum coming from periodic V, where we end up getting an infinite collection of intervals each with a continuous spectrum. But in crystals, we often get an "impurity" which adds an additional term W to V, and the question is what happens to the spectrum. Anderson localization is then the conjecture that that the spectrum breaks up into pure-point spectra with localized eigenfunctions (fixed energy), and in fact continuous spectra corresponds to conductivity and oscillating functions.