1 Prerequisites

The official prerequisites are real analysis and linear algebra. I am expecting you to have a passing familiarity with the Lebesgue integral, know the dominated convergence theorem and similar facts, and know what a measure is and that you can integrate with respect to it. I am hoping you have seen L^p spaces; if not, you may want to read up a bit on the basic properties. I am also expecting you to know various facts from point-set topology about metric and topological spaces.

None of the above is especially essential, however. It's fair game to ask me to prove things that I claim you should have seen and you haven't. I reserve the right to not actually do this, but I'll at least send you a reference.

2 Structure

We will be following Haim Brezis, "Functional Analysis, Sobolev Spaces, and Partial Differential Equations" reasonably closely; you will need access to this book. You can access an online copy for free through the Rutgers library.

This book has many, many exercises of varying difficulty levels. I will assign some of them for you to do, typically on Thursdays. You will then turn them in to me the following Thursday. I will grade some subset of them. This will constitute about 15% of your grade.I would encourage you to at least think through and discuss other exercises from the text. This is the best (only?) way to really learn the material. Yes, I am aware that some of these exercises have solutions at the back of the book. You are also free to work together as you see fit on them (but turn in your own work).

I will sometimes ask one of you to present a solution to a homework problem, to help guide the class through a proof, or to work through one of the problems at the end of the book and present your findings. This will generally be on a volunteer basis, unless you abstain from volunteering. If the task is difficult I will assign it a bit in advance. Actively participating in this will account for 15% of your grade; if you participate often, this will be weighed heavily towards your best showings. I have never tried this in graduate classes before, so if it's a disaster I reserve the right to modify this.

There will be a take-home midterm exam worth 30% of your grade. There will be a final exam of some sort worth 40% of your grade; it will almost surely contain a take-home component, and might also contain either a written or oral component if I deem it warranted.

I expect to cover the first six chapters of Brezis (with some omissions), and hopefully Chapter 8. This may be adjusted somewhat to accommodate student interest.

3 What is functional analysis and why study it?

Functional analysis is an approach to linear algebra on infinite-dimensional vector spaces.

The most obvious reason one might encounter an infinite-dimensional vector space is because they are studying spaces of functions. Given a set X and a vector space V, V^X (the set of functions from X to V) is also a vector space. More usefully, so are various subspaces of this.

Thinking of collections of functions like this is a somewhat modern innovation, diffusing from algebra to geometry and then to analysis at the end of the nineteenth and beginning of the twentieth century. However, once you make this observation, many questions take on a distinctly linear-algebra interpretation. To give some examples:

- Consider any linear ordinary differential equation. This means searching for a function $u : \mathbb{R} \to \mathbb{R}^n$

satisfying a relation of the form

$$\sum_{k=0}^{K} a_k(t)u^{(k)}(t) = f(t),$$

where a_k are matrix-valued functions, the superscripts are derivatives in t and f is a function into $\mathbb{R}^n$. Assume a_k, f are all smooth. Possibly some boundary or initial conditions are imposed on u. This is a typical problem in analysis, and one may attempt to solve in various ways, including via power series, via transformations, via direct integration, via fundamental solutions, etc. One may, however, interpret it as follows: define a function from $C^K(\mathbb{R};\mathbb{R}^n)$ (the vector space of functions from $\mathbb{R}$ to $\mathbb{R}^n$ which admit K continuous derivatives) to $C^0(\mathbb{R};\mathbb{R}^n)$ given by

$$L(v) = \sum_{k=0}^{K} a_k v^{(k)}.$$

This is a linear map (differentiation is linear; that's the main point here) between these two spaces of functions. The question of whether the differential equation is solvable is the question of whether f lies in the image of L. The question of how to solve it is the most basic question addressed by linear algebra. The study of the properties of the solution is the study of the structure of the map L. The desire to diagonalize L lies at the heart of every transformation method designed to study these equations.

- Ordinary differential equations were actually not the (only) problem which drove the early development of the subject. Here is another: let g_i be a collection of (infinitely many) square-integrable functions (say from $\mathbb{R} \to \mathbb{R}$ with Lebesgue measure; this is unimportant), and α_i be a sequence of real numbers. Can one find a single square-integrable function f which satisfies all of the constraints

$$\int f g_i = \alpha_i$$

for each i? Square-integrability ensures that the above makes sense (why?). To understand why this is a "linear algebra" problem, notice that we may, given any square-integrable f, define a linear map $L_f : L^2(\mathbb{R}) \to \mathbb{R}$ ($L^2(X)$ is the space of square-integrable functions) via

$$L_f(v) = \int fv.$$

Therefore, the goal here is to produce a linear map of this type passing through a collection of points (g_i, α_i).

Now, given a collection of points of this form, there is an obvious necessary condition that must be satisfied: if a linear combination of the g_i gives 0, the corresponding linear combination of the α_i does too. It can be shown that in fact this is *sufficient* as well.

However, one quickly realizes that there is a problem here. Not every linear map from $L^2(X) \to \mathbb{R}$ is of the form L_f for some f; those that are satisfy an extra condition:

$$|L_f(v)| \le \|f\|_{L^2(X)} \|v\|_{L^2(x)},$$

and this in fact characterizes them (see the Reisz representation theorem for L^2; we will discuss this later in the course). This insight is the first inkling of how functional analysis departs from linear algebra: purely algebraic characteristics of a vector space either fail to provide useful understanding of our problems, or they answer different questions from those we had in mind. Notice the core role played by the norm in this example; this is not an algebraic property at all.

- As a final example, to illustrate some of the perils of functional analysis, consider another classic topic: the calculus of variations. Say that we wish to minimize the following quantity (here u is an arbitrary absolutely continuous function whose derivative is square-integrable; these form a vector space, say V):

$$E(u) = \int_0^1 a(x)(u')^2 + u^2$$

over $K = \{u \in V : u(0) = \alpha, u(1) = \beta\}$, where $0 < a_0 \leq a(x) \leq a_1 < \infty$ is fixed and measurable. Some of you may be able to immediately write down the solution, in fact, but hold back on doing this for a moment.

Is this even a functional analysis problem? Well, there is an obvious method to try and prove that it has a solution. Clearly we have $E \geq 0$ and there is at least one $u \in K$ so that $E(u) < \infty$. Set $m = \inf_{u \in K} E(u)$. Then there is a sequence of $u_k \in K$ so that $E(u_k) \to m$ (this is just from the definition of infimum, nothing has happened yet).

Now, there is a natural norm one may put on V: $\|v\|^2 = \int v^2 + (v')^2$. If defined suitably, one may even show that V is complete in the topology derived from this norm (we will consider this carefully near the end of the course, time permitting). We see that as $E(u_k) \leq m + 1$ for large k, and as

$$\|u_k\|^2 \leq \frac{1}{a_0} E(u_k) \leq \frac{m+1}{\alpha_0} < \infty,$$

all of the u_k live inside of a large ball in V. If V was a finite-dimensional Euclidean space, we could now argue that this large ball is closed and bounded, and hence compact; it would follow that there exists a subsequence $v_{k_j} \to v \in V$, in the norm, i.e. $\|v_{k_j} - v\| \to 0$ (compact and sequentially compact are equivalent in metric spaces). It is then not hard to show that E is continuous with respect to this convergence and that v must be in K (check; you need the fundamental theorem of calculus or similar for the second point). Hence, this $v \in K$ has $E(v) = m$; it is a minimizer.

This argument fails in infinite dimensional vector spaces. It fails because balls are not compact! (prove this). Understanding carefully why and how it fails and how to salvage it will take us most of the way through the course.

What we gather from the above is that if we want to apply linear algebra tools to spaces of functions in a productive manner, we must adapt them to care about the norm or topology that space comes with. Functional analysis is this study of topological vector spaces.

4 Basic Notions

You should likely have seen all of the concepts below, but if not, it would be best to review them.

Definition 4.1. *A* vector space V *over a field* F *is an abelian group (operation denoted by* $+$*) together with a mapping* $F \times V \to V$ *(denoted by adjacency) with the following properties* ($\forall \alpha, \beta \in F, v, w \in V$)*:*

1. $\beta(\alpha v) = (\beta\alpha)v$
2. $1_F v = v$
3. $(\alpha + \beta)v = \alpha v + \beta v$
4. $\alpha(v + w) = \alpha v + \alpha w.$

For the remainder of the course, F will be $\mathbb{R}$ or $\mathbb{C}$. A key point in functional analysis is not only the field, but also the *absolute value* on the field and the fact that it is topologically complete with respect to this absolute value. This means that if you are interested in fields like $\mathbb{Q}$, you will naturally end up studying their topological completions. There are other fields which are topologically complete besides these two: notably the p-adics and related things. These end up having extremely different properties due to the geometry of their absolute value, and are really a subject in their own right which we will not touch. I will note here that I am personally quite biased when writing the above comment, as I am a PDE specialist and PDE are studied over $\mathbb{R}$ or maybe $\mathbb{C}$; nonetheless, classically the subject has been constrained in this way.

A further remark: functional analysis over $\mathbb{R}$ and $\mathbb{C}$ is very similar, with only occasional differences. We will work over $\mathbb{R}$. There is a section in Brezis which discusses the differences and how to deal with them, and I would encourage you to read it if analysis over $\mathbb{C}$ is what you need. If you continue to study functional analysis (especially the theory of operators) it is highly likely that you will need or prefer to work over $\mathbb{C}$ at some points, for the same reason that it is easier to discuss diagonalizability of matrices over $\mathbb{C}$ than over $\mathbb{R}$.

Definition 4.2. *A subset $E \subseteq V$ of a vector space is* convex *if $tx + (1-t)y \in E$ for any $x, y \in E$ and $t \in (0,1)$. A function $\phi : V \to \mathbb{R}$ is convex if $\phi(tx + (1-t)y) \leq t\phi(x) + (1-t)\phi(y)$ for all $x, y \in V$ and $t \in (0,1)$.*

This is an algebraic property; it makes sense for any vector space over $\mathbb{R}$. As such, it plays a critical role in functional analysis.

Definition 4.3. *A* normed vector space *is a vector space V equipped with a* norm*, a function $\|\cdot\| : V \to \mathbb{R}$ satisfying*

1. *$\|\alpha x\| = |\alpha|\|x\|$ for any $\alpha \in \mathbb{R}$ and $x \in V$.*
2. *$\|x\| = 0$ only if $x = 0$.*
3. *$\|x + y\| \leq \|x\| + \|y\|$ for any $x, y \in V$.*

It is typical to just write V for a normed vector space, even though it is really the pair of vector space and norm which is important (and one may put different norms on the same space, etc). When ambiguous one may use $\|\cdot\|_V$ for the norm, or other subscripts if still ambiguous.

Some obvious observations:

Proposition 4.1. *Let V be a normed vector space.*

1. *$\|\cdot\|$ is convex.*
2. *Setting $d(x, y) = \|x - y\|$, d is a metric on V. In particular, it induces a topology on V.*
3. *Let V be finite-dimensional. Then any isomorphism $V \to \mathbb{R}^n$ is a homeomorphism.*

It is important here to understand that (3) fails completely if V is not finite-dimensional, in the sense that there is absolutely no hope of salvaging it. Norms on Euclidean spaces are a geometric curiosity; interesting but at the end of the day revealing nothing new about the topology or structure of the set. On infinite-dimensional spaces, norms produce topologies which are completely different from one another, and are the key piece of structure we will be dealing with.

It is possible to obtain *topological vector spaces* that do not come supplied with a norm. This happens in e.g. the theory of distributions naturally; sometimes there is a metric which produces the topology, but the metric does not derive from a norm. You may see more of this in the second semester of this course, if it is offered and you take it. We will also put some other topologies on normed spaces when we study them; this is a bit different in spirit, as the topologies will be very different from the one induced by the norm.

Definition 4.4. *A* Banach space *is a normed vector space which is complete in the norm topology.*

There is an obvious incentive to working with Banach spaces when interested in topics in analysis, as naturally the completeness lends itself to existence results. It is also technically simpler, though a decent portion of the theory does not require it.

Some quick review of point-set topology:

A metric space V is *complete* if every Cauchy sequence (i.e. $d(u_k, u_j)- > 0$ as $k, j \to \infty$) converges (i.e. $d(u_k, u) \to 0$ for some $u \in V$). This means, roughly, that there are no "holes;" it says nothing of the size or quality of the space. There is a procedure where any metric space can be continuously included in a completion of it, and this can preserve the linear and norm structure, but we will not really be using it.

There is a hierarchy of properties which describe how well a topological space separates points. Let me remind you of it, as it will help frame some of what we will soon be doing.

T1 For any $x, y \in X$, there is an open set such that $x \in U$ and $y \notin U$.

T2 ("Hausdorff") For any $x, y \in X$, there are open disjoint U_x, U_y with $x \in U_x$ and $y \in U_y$.

T3 ("regular") T2 + For any closed $E \subseteq X$ and $x \in X \setminus E$, there are open disjoint U_E, U_x with $E \subseteq U_E$ and $x \in U_x$.

T4 ("normal") T2 + For any two closed sets $E, F \subseteq X$, there are open disjoint U_E, U_F with $E \subseteq U_E$ and $F \subseteq U_F$.

Metric spaces satisfy all of these (why?). We will later look at topologies where the question of which of these are satisfied is more interesting.

Definition 4.5. *Let V, W be normed vector spaces and $T : V \to W$ be a linear map. This will often be called a* linear operator*. If $W = \mathbb{R}$, it will be called a* linear functional*. We say that T is* bounded *if there exists a constant M such that*

$$\|T(x)\|_W \leq M\|x\|_V$$

for all $x \in V$.

Proposition 4.2. *A linear operator is bounded if and only if it is a continuous function.*

(prove this.)

The problem with infinite-dimensional spaces is that linear maps are common, easy to construct, and typically pathological. To see why, recall that every vector space admits a basis (a Hamel basis, in the sense of a linearly independent spanning set with respect to finite linear combinations); this can be proved using the axiom of choice and is found in Exercise 1.5 in Brezis. A linear map is fully determined by how it acts on a basis, and one is free to choose where each basis element is mapped. The problem with this is that the resulting maps lose all relationship to the geometry and topology of our vector space. As such, we will often be more interested in bounded linear maps. These are much less common and do interact with the important structure fully.

5 List of Characters

This is a list of simple topological vector spaces you should have seen before and a quick summary of some of their properties.

- Euclidean spaces $\mathbb{R}^n$ with various norms. A norm can be constructed as follows: take K a symmetric ($x \in K$ means $-x \in K$) convex set for which 0 is in the interior. Let $\|x\| = \sup\{t > 0 : x/t \in K\}$. We will discuss this construction extensively later. Typical norms you should be aware of are

$$\|x\|_p = \begin{cases} (\sum_{i=1}^n |x_i|^p)^{\frac{1}{p}} & 1 \leq p < \infty \\ \max_i |x_i| & p = \infty. \end{cases}$$

  (You may check that these are indeed norms, i.e. they satisfy the triangle inequality. You probably know how to from your real analysis course.)

- Given a measure space (X, μ), take the space of measurable functions $u : X \to \mathbb{R}$ for which

$$\|x\|_{L^p(X)} = \begin{cases} (\int_X |u|^p d\mu)^{\frac{1}{p}} & 1 \leq p < \infty \\ \text{esssup}|u| & p = \infty. \end{cases}$$

  Actually this is not a normed vector space, as this "norm" (potentially) violates property (2) of being a norm. One can, however, consider equivalence classes of functions as above, with u v if $u = v$ μ-almost everywhere. The set of all such equivalence classes still forms a vector space, and the norm above is well-defined on it. These normed vector spaces are denoted by $L^p(X)$ (with μ appended if unclear). I am assuming that you have seen them before in real analysis, that you know that they are well-defined and that the norm satisfies the triangle inequality, and that you have shown that they are complete, i.e. Banach spaces.

- A common sub-example of the above is $X = \mathbb{N}$, μ the counting measure. This gives spaces of p-summable sequences, and is typically abbreviated l^p. A slight generalization is $l^p(E)$ for any set E, which is $L^p(E)$ equipped with the counting measure. Note that in these examples all of the equivalence classes up to a.e. contain one element, and so the issue may be ignored. We will consider a few other spaces of sequences, including c (the subspace of λ^∞ for which the limit of the sequence exists) and c_0 (the subspace of c for which the limit is 0).

- Given a topological space X, consider $C(X)$, the space of continuous functions from X to $\mathbb{R}$. This has the natural norm
$$\|u\|_{C(X)} = \sup_X |u|,$$
and is a Banach space. If $X = \mathbb{R}^n$, a common generalization is $C^k(X)$, the space of continuous functions with k continuous partial derivatives, equipped with the norm
$$\|u\|_{C^k(\mathbb{R}^n)} = \sum_{|\alpha|\le k} \|\partial^\alpha u\|_{C(\mathbb{R}^n)}.$$
This is using multi-index notation; the sum is over every possible combination of partial derivatives. This also makes sense on C^k-differentiable manifolds, for example by using charts.

  $C^\infty(\mathbb{R}^n)$ can not be made into a normed vector space (or at least, not in a way that gives the expected topology, where sequences converge if and only if all of the derivatives converge uniformly). The problem is that if you want to control infinitely many quantities (all of the sup-norms of all of the partial derivatives in this case), you are forced to make a choice about how you do so, similarly to why the l^p spaces are not all the same. There is a natural way to make it into a topological (indeed, metric) vector space (how?).