



Advice for the final exam

Practical information

The available dates and times for the final exam are:

- (a) Wednesday, February 25, 12:00–15:00 in the Erwin-Schrödinger-Zentrum (Rudower Chaussee 26), rooms 0'110 and 0'310
- (b) Monday, March 30 from 9:00–12:00 in the Erwin-Schrödinger-Zentrum (Rudower Chaussee 26), room 0'115

If you do not pass it on the first date, you can register to try again on the second date. (You can also register for the second date even if you have not registered for the first one, as long as you do so before the deadline set by the Prüfungsbüro.)

Extra office hours

My usual Friday office hour is cancelled between semesters, but if you'd like to come by my office with questions in preparation for the exam, I will be available between 14:00 and 16:00 on Monday 23.02, Tuesday 24.02, Thursday 26.03 and Friday 27.03.

You can also feel free to post questions to the moodle forum or send them to me directly by e-mail.

What is or is not allowed during the exam?

As indicated in the original course syllabus, this will be an **open-book** exam, so you may have your class notes, textbooks, past problem sets, the typed lecture notes from the course etc. on hand for the exam and should feel free to use them. Electronic devices are **not** allowed.

A word of caution: The fact that the exam is open-book means that you don't need to stress about *memorizing* things, but it doesn't mean that you don't need to *study*. See below for more advice on that.

Format

The format of the exam will be similar to that of a problem set: 4 or 5 problems with 2 or 3 parts each. The bulk of the problems will be designed to be doable within a total of two hours. Some parts will be doable by anyone who has learned anything in the course, and a few parts may only be doable by the top 10% of students—do not despair or panic if one or two of the “part (c)”s leave you completely stumped.

Since the exam is open-book, there will be no questions asking you to reproduce essential definitions or proofs of standard theorems. The problems will instead be designed to test how well you have understood the main ideas behind those theorems, whether you can adapt them to different contexts and apply them in examples. One piece of good advice I can give about preparation is to review the problem sets and make sure you understand the solutions that were discussed in the problem sessions—including the unstarred problems. It is not out of the question that some of those problems may reappear in nearly identical forms on the exam. Most things that were stated in lecture as “exercises” but not assigned for homework are also fair game for exam problems.

Examinable vs. non-examinable material

As a rule, I will not expect you on the exam to understand anything about any proof that was not explained in either a lecture or a homework problem. There are a few cases of such results that were only briefly sketched in lecture but are explained in detail in the lecture notes, and are important enough that I will assume you at least understand the *statements* and how to use them in applications. A notable example of this would be Theorem 13.28 in the notes, which gives the formula $\partial^\alpha(\varphi*\Lambda) = (\partial^\alpha\varphi)*\Lambda = \varphi*\partial^\alpha\Lambda$ for differentiating the convolution of a distribution with a test function. Another is Theorem 4.5, a consequence of the dominated convergence theorem that we have used many times (occasionally without mentioning it) to justify differentiation under the integral sign. You don’t need to know how to prove such results, but you should know how to use them.

On the other hand, I will *not* assume you know anything about the following topics, despite some of them being covered extensively in the lecture notes:

- Maximal and weakly integrable functions, Vitali’s covering lemma and the Hardy-Littlewood maximal inequality (§9.3 in the notes)
- Functions of bounded variation and the proof that non-increasing absolutely continuous functions are integrals (§9.5.2 in the notes)
- Alternative approaches to proving the FTC without Radon-Nikodým (§9.5.3 in the notes)
- Why certain functions are nowhere differentiable (§11.8 in the notes)
- Fubini’s theorem for distributions (§13.4 in the notes)
- The topology of the space $\mathcal{D}(\Omega)$ of test functions (§13.8 in the notes): to clarify, I will definitely expect you to understand what convergence of a sequence in $\mathcal{D}(\Omega)$ and continuity of a linear functional $\Lambda : \mathcal{D}(\Omega) \rightarrow V$ mean, but you need not worry about what the open sets in $\mathcal{D}(\Omega)$ are.
- The Sobolev spaces $H^s(\mathbb{R}^n)$ for $s < 0$, $H^s(\Omega)$, $\tilde{H}^s(\Omega)$ and $H_0^s(\Omega)$ for open domains $\Omega \subset \mathbb{R}^n$, and the properties of the operator $\Delta : H_0^1(\Omega) \rightarrow H^{-1}(\Omega)$ (§17.2 and §17.7 in the notes)
- The functional calculus for unbounded self-adjoint operators and Stone’s theorem (contents of the final lecture)

How to study

Aside from reviewing homework, I can give the following advice about reviewing the course material. When reviewing an important theorem, ask yourself and try to answer the following questions:

- (a) *What is it good for?*
Try to find a few examples of other theorems in the course or applications we discussed in which this theorem shows up as an essential ingredient.
- (b) *Why, in a nutshell, is it true?*
In other words, don't try to memorize the proof, but see if you can discern a main idea or trick that summarizes why the proof works.
- (c) *Why are the hypotheses what they are?*
Think about how each individual hypothesis is used in the proof, especially those hypotheses that involve essential definitions like the completeness of a Banach space, or a subspace being closed. Consider how far you can alter the hypotheses before the theorem becomes false. (There have sometimes been homework problems exploring such questions.)

Similarly, when reviewing an important definition, ask yourself:

- (a) *What, in a nutshell, is it?*
Try to give an informal summary of the definition in only one sentence, even if only by an analogy (or a picture!).
- (b) *What are some examples and non-examples?*
If the definition is a condition, find some examples of things that satisfy the condition, and also some examples of things that don't.
- (c) *What is it good for?*
A common mathematical saying is that a good definition should always be a hypothesis for a good theorem. (If it is not, then the definition isn't good.) Find a few examples of theorems for which this definition is an essential hypothesis.

Here are some examples of this strategy in action.

Example 1. Definition: Uniform convexity.

- (a) *What, in a nutshell, is it?*
A normed vector space is uniformly convex if the unit ball is “round” in a quantitative sense described via the midpoint between two arbitrary points on its boundary. (This can be expressed better with a picture.)
- (b) *What are some examples and non-examples?*
All inner product spaces are uniformly convex, and so are most of our favorite Banach spaces, e.g. $L^p(X)$ for every $p \in (1, \infty)$, but $L^1(X)$ and $L^\infty(X)$ are not. (I don't think we ever talked about any other examples, so I wouldn't suggest trying to think up more examples now.)

(c) *What is it good for?*

The main immediate application is the theorem in Example 2 below, which is used in the proof of the Riesz representation theorem. One can also use uniform convexity to turn weakly convergent sequences $x_n \rightharpoonup x$ into strongly convergent sequences $x_n \rightarrow x$ under the extra condition that $\|x_n\| \rightarrow \|x\|$ (Theorem 7.16 in the lecture notes).

Example 2. Theorem (3.23 in the lecture notes): If X is a uniformly convex Banach space, $K \subset X$ is a closed convex subset and $x \in X \setminus K$, then K contains a unique point closest to x .

(a) *What is it good for?*

This was the main tool we used in proving that every Hilbert space \mathcal{H} is $V \oplus V^\perp$ for any closed subspace $V \subset \mathcal{H}$, because $x - v \in V^\perp$ for any $x \in \mathcal{H} \setminus V$ if v is the point in V closest to x . A similar trick provides the existence result needed for the Riesz representation theorem, both in Hilbert spaces and in $L^p(X)$ for $1 < p < \infty$.

(b) *Why, in a nutshell, is it true?*

Choose a sequence $v_k \in K$ whose distances to x converge to the infimum of all such distances and use the uniform convexity condition to prove that v_k is a Cauchy sequence.

(c) *Why are the hypotheses what they are?*

The theorem becomes false if X is uniformly convex but not complete, as the Cauchy sequence in the proof-sketch above need not converge. This is why, for instance, there can exist an (incomplete!) inner product space containing a closed codimension 1 subspace whose orthogonal complement is trivial (see Problem Set 2 #3).

Example 3. Definition: Compact operators.

(a) *What, in a nutshell, is it?*

An operator between Banach spaces is compact if it sends bounded sets to sets with compact closure.

(b) *What are some examples and non-examples?*

Finite-rank operators (in particular, every linear map to a finite-dimensional space) and limits of sequences of finite-rank operators are examples. More concrete examples include the Sobolev inclusions $H^s(\mathbb{T}^n) \hookrightarrow H^t(\mathbb{T}^n)$ for $s > t$, and certain operators defined via convolutions (e.g. Problem Set 11 #3). Non-examples include all Banach space isomorphisms in infinite dimensions.

(c) *What is it good for?*

Compact perturbations of Fredholm operators are always Fredholm and have the same index, so for instance, an operator of the form $\mathbb{1} + K$ with K compact is injective if and only if it is surjective. Relatedly, the spectral theory of compact operators is not so different from spectral theory in finite dimensions: the spectrum consists of isolated eigenvalues with finite multiplicity (with the possible exception of 0), and in the self-adjoint or normal case, there is an orthonormal basis of eigenvectors.

Example 4. The open mapping theorem: surjective bounded linear operators $T : X \rightarrow Y$ between Banach spaces X, Y map open sets to open sets.

(a) *What is it good for?*

Mainly for proving the *inverse* mapping theorem, which gives every bounded linear bijection a bounded inverse. The latter is used in proving e.g. that small perturbations of Fredholm operators are Fredholm with the same index, that an injective operator has closed image if and only if it is “bounded below” (Take-Home Mid-term #4(a)), and that the resolvent $R_\lambda(T)$ of a closed operator T for $\lambda \in \mathbb{C} \setminus \sigma(T)$ is a bounded operator.

(b) *Why, in a nutshell, is it true?*

Because if $T \in \mathcal{L}(X, Y)$ and $Y = \bigcup_{n \in \mathbb{N}} T(B_n(0))$, then the Baire category theorem implies that at least one of the $T(B_n(0))$ has closure containing a ball. (That’s not the entire proof, but the rest might be characterized as “nitpicky details”.)

(c) *Why are the hypotheses what they are?*

If Y is not complete then it does not satisfy the Baire category theorem, so all of the $T(B_n(0))$ might be nowhere dense, and it then becomes easy to find examples where the inverse mapping theorem fails: e.g. if $T \in \mathcal{L}(\mathcal{H})$ is injective but has 0 as an approximate eigenvalue, then it is a bijection to its necessarily non-closed image, and the inverse of this bijection is unbounded. (Easiest concrete example: a multiplication operator $T : L^2(\mathbb{R}) \rightarrow L^2(\mathbb{R}) : u \mapsto Fu$ for any continuous bounded function $F : \mathbb{R} \rightarrow (0, \infty)$ whose image has 0 in its closure.)