Humboldt-Universität zu Berlin
Institut für Mathematik
C. Wendl, S. Dwivedi,
L. Upmeier zu Belzen

Funktionalanalysis
WiSe 2020–21

Advice for the final exam

Practical information

The exam will take place **online** during a precise **3-hour** time window (with a half-hour buffer at the end for uploading solutions). **Important**: Everyone taking the exam must join a **new moodle course** that has been created specifically for this exam:

- Moodle "Funktionalanalysis Klausur 9.04.2021":
  `https://moodle.hu-berlin.de/course/view.php?id=103195`
- Enrolment key: `Banach`

Essential practical details about the exam, including a trial run planned for the Wednesday beforehand (in order to sort out technical issues), will be announced only via that moodle page, and not via the usual moodle page for this course. Please enrol yourself in the new moodle course as soon as possible if you are taking the exam. (**Students who already took the March 3 exam, take note:** This is yet another separate moodle course, *not* the same one that was used for the March 3 exam.)

Also important (but irrelevant to most of you): if you are a HU student, then you are required to use your HU-account for enrolling in the new moodle course, even if you normally use a different e-mail address for logging into moodle. You may in that case need to contact Moodle-Support to change the e-mail address that you use for logging in—see the passage beginning with "Grundsätzlich müssen sich die Studierenden der HU. . . " at `https://www.cms.hu-berlin.de/de/dl/e-assessment/guide/vorbereitung`

Extra office hours

I will be available for questions during my usual office hour (Tuesday 14:00-15:00) in the week of the exam, and also during the usual Thursday lecture time (13:15-14:45) on the day before. The Zoom link for my virtual office hour is
`https://hu-berlin.zoom.us/j/98760164105`

There will be another opportunity to ask mathematical questions during or right after the "trial run" planned for the Wednesday before the exam (see the new moodle for details).

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. In theory you can also use the internet, but I would advise against it, as the information you find there may be unreliable, and searching the web will eat up time that you should instead spend thinking about the problems. We will be available the entire time in a continuously running Zoom meeting (and also by phone in case that doesn't work) to answer any questions that may arise. **You may not communicate with other people besides the teaching staff during the exam.**

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. The best 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 definitely would have been proved in lecture if the semester had been two weeks longer, and these are important enough that I will assume you at least understand the *statements* and how to use them in applications. The notable examples are the results singled out at the top of Problem Set 9, namely:

- For any distribution Λ and test function φ, $\varphi * \Lambda$ is a smooth function satisfying $\partial^\alpha(\varphi * \Lambda) = (\partial^\alpha \varphi) * \Lambda = \varphi * \partial^\alpha \Lambda$ for all multi-indices α. (Theorem 10.27 in the notes)
- If $\Lambda \in \mathscr{D}'(\Omega)$ has first derivatives $\partial_1 \Lambda, \ldots, \partial_n \Lambda \in \mathscr{D}'(\Omega)$ that are all representable by continuous functions on $\Omega \subset \mathbb{R}^n$, then Λ is representable by a C^1-function on Ω. (Theorem 10.33 in the notes)

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 (§6.3 in the notes)
- Functions of bounded variation and the proof that non-increasing absolutely continuous functions are integrals (§6.5.2 in the notes)
- Alternative approaches to proving the FTC without Radon-Nikodým (§6.5.3 in the notes)
- Why certain functions are nowhere differentiable (§8.8 in the notes)

- Fubini's theorem for distributions (§10.4 in the notes)
- Distributions with compact support (§10.7 in the notes)
- The topology of the space $\mathscr{D}(\Omega)$ of test functions (§10.8 in the notes): to clarify, I will definitely expect you to understand what convergence of a sequence in $\mathscr{D}(\Omega)$ and continuity of a linear functional $\Lambda : \mathscr{D}(\Omega) \to V$ mean, but you need not worry about what the open sets in $\mathscr{D}(\Omega)$ are.
- The spaces $H^s(\mathbb{R}^n)$ for $s < 0$, $H^s(\Omega)$, $\widetilde{H}^s(\Omega)$ and $H_0^s(\Omega)$ for open domains $\Omega \subset \mathbb{R}^n$, and the properties of the operator $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ (§2 and the last four exercises in the separate notes on Fredholm operators)
- 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 $x_n \to x$ under the extra condition that $\|x_n\| \to \|x\|$ (Theorem 4.16 in the lecture notes).

Example 2. Theorem (1.8 in the lecture notes): If X is a uniformly convex Banach space, $K \subset X$ is a closed convex subset and $x \in X \backslash K$, 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 $x \in \mathcal{H} \backslash V$ 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 $H^s(\mathbb{T}^n) \hookrightarrow H^t(\mathbb{R}^n)$ 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, leading to the so-called *Fredholm alternative* for operators of the form $\mathbb{1}_X - K$ with $K : X \to X$ compact. 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 case, there is an orthonormal basis of eigenvectors.

Example 4. The open mapping theorem: surjective bounded linear operators $T : X \to Y$ 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 Midterm #4(a)), and that the resolvent $R_\lambda(T)$ of a closed operator T for $\lambda \in \mathbb{C}\backslash\sigma(T)$ is a bounded operator.

(b) *Why, in a nutshell, is it true?*
Because if $T \in \mathscr{L}(X, Y)$ 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 $T \in \mathscr{L}(\mathcal{H})$ 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 $T : L^2(\mathbb{R}) \to L^2(\mathbb{R}) : u \mapsto Fu$ for any continuous bounded function $F : \mathbb{R} \to (0, \infty)$ whose image has 0 in its closure.)