SPECTRAL THEORY

EVAN JENKINS

Abstract. These are notes from two lectures given in MATH 27200, Basic Functional Analysis, at the University of Chicago in March 2010. The proof of the spectral theorem for compact operators comes from [Zim90, Chapter 3].

1. The Spectral Theorem for Compact Operators

The idea of the proof of the spectral theorem for compact self-adjoint operators on a Hilbert space is very similar to the finite-dimensional case. Namely, we first show that an eigenvector exists, and then we show that there is an orthonormal basis of eigenvectors by an inductive argument (in the form of an application of Zorn's lemma).

We first recall a few facts about self-adjoint operators.

Proposition 1.1. *Let V be a Hilbert space, and $T : V \to V$ a bounded, self-adjoint operator.*

1. *If $W \subset V$ is a T-invariant subspace, then $W^\perp$ is T-invariant.*
2. *For all $v \in V$, $\langle Tv, v\rangle \in \mathbb{R}$.*
3. *Let $V_\lambda = \{v \in V \mid Tv = \lambda v\}$. Then $V_\lambda \perp V_\mu$ whenever $\lambda \neq \mu$.*

We will need one further technical fact that characterizes the norm of a self-adjoint operator.

Lemma 1.2. *Let V be a Hilbert space, and $T : V \to V$ a bounded, self-adjoint operator. Then $\|T\| = \sup\{|\langle Tv, v\rangle| \mid \|v\| = 1\}$.*

Proof. Let $\alpha = \sup\{|\langle Tv, v\rangle| \mid \|v\| = 1\}$. Evidently, $\alpha \leq \|T\|$. We need to show the other direction.

Given $v \in V$ with $Tv \neq 0$, setting $w_0 = Tv/\|Tv\|$ gives $|\langle Tv, w_0\rangle| = \|Tv\|$. Thus,

$$\begin{aligned}\|T\| &= \sup_{\substack{v\in V,\\ \|v\|=1}} |\langle Tv, w_0\rangle| \\ &\leq \sup_{\substack{v,w\in V,\\ \|v\|=\|w\|=1}} |\langle Tv, w\rangle|.\end{aligned}$$

Thus, it suffices to show that $|\langle Tv, w\rangle| \leq \alpha\|v\|\|w\|$ for all $v, w \in V$. Without loss of generality, we may multiply w by a unit complex number to make $\langle Tv, w\rangle$ is real and positive.

We now compute

$$\begin{aligned}\langle T(v + w), v + w\rangle &= \langle Tv, v\rangle + \langle Tv, w\rangle + \langle Tw, v\rangle + \langle Tw, w\rangle \\ &= \langle Tv, v\rangle + 2\langle Tv, w\rangle + \langle Tw, w\rangle,\end{aligned}$$

and

$$\langle T(v - w), v - w\rangle = \langle Tv, v\rangle - \langle Tv, w\rangle - \langle Tw, v\rangle + \langle Tw, w\rangle$$

$$= \langle Tv, v\rangle - 2\langle Tv, w\rangle + \langle Tw, w\rangle .$$

Subtracting these two identities yields

$$4\langle Tv, w\rangle = \langle T(v+w), v+w\rangle - \langle T(v-w), v-w\rangle .$$

Taking absolute values and applying the triangle and parallelogram inequalities yields

$$\begin{aligned}\langle Tv, w\rangle &\leq \frac{\alpha}{4}(\|v+w\|^2 + \|v-w\|^2)\\ &\leq \frac{\alpha}{2}(\|v\|^2 + \|w\|^2).\end{aligned}$$

Without loss of generality, we assume $v, w \neq 0$. If we replace v by $\sqrt{a}v$ and w by $\sqrt{a}^{-1}w$, where $a = \|w\|/\|v\|$, the left-hand side remains unchanged, while the right-hand side becomes

$$\frac{\alpha}{2}((\|w\|/\|v\|)\|v\|^2 + (\|v\|/\|w\|)\|w\|^2) = \alpha\|v\|\|w\|$$

as desired. □

We are now ready to prove the spectral theorem.

Theorem 1.3 (Spectral Theorem for Compact Self-Adjoint Operators)**.** *Let V be a nonzero Hilbert space, and let $T : V \to V$ be a compact, self-adjoint operator. Then $V = \overline{\bigoplus_\lambda V_\lambda}$. For each $\lambda \neq 0$, $\dim V_\lambda < \infty$, and for each $\epsilon > 0$, $|\{\lambda \mid |\lambda| \geq \epsilon, \dim V_\lambda > 0\}| < \infty$.*

Proof. If $T = 0$, the conclusions are obvious, so we assume $T \neq 0$. We will first show the existence of a nonzero eigenvector with eigenvalue $\lambda = \|T\| \neq 0$. Since $\lambda = \sup\{|\langle Tv, v\rangle| \mid \|v\| = 1\}$ by Lemma 1.2, we can choose a sequence $v_n \in V$ with $\|v_n\| = 1$ such that $\langle Tv_n, v_n\rangle \to \lambda$. Since T is compact and the sequence (v_n) is bounded, we can pass to a subsequence and assume $Tv_n \to w$ for some $w \in V$. Since $\|Tv_n\| \geq \langle Tv_n, v_n\rangle \to \lambda \neq 0$, it follows that $\|w\| \geq \lambda \neq 0$, so that, in particular, $w \neq 0$. This is our candidate for an eigenvector.

We compute, using the fact that $\langle Tv_n, v_n\rangle \in \mathbb{R}$,

$$\begin{aligned}\|Tv_n - \lambda v_n\|^2 &= \langle Tv_n - \lambda v_n, Tv_n - \lambda v_n\rangle\\ &= \|Tv_n\|^2 - 2\lambda\langle Tv_n, v_n\rangle + \lambda^2\|v_n\|^2\\ &\leq 2\|T\|^2 - 2\lambda\langle Tv_n, v_n\rangle\\ &\to 0.\end{aligned}$$

Thus, $Tv_n - \lambda v_n \to 0$ as $n \to \infty$. But $Tv_n \to w$ as $\lambda \to \infty$, so we must have $\lambda v_n \to w$. Bounded operators take convergence sequences to convergent sequences, so applying T to both sides yields $\lambda Tv_n \to Tw$. But we know that $\lambda Tv_n \to \lambda w$ by definition, so we must have $Tw = \lambda w$ as desired.

By Zorn's lemma, we choose a maximal orthonormal set of eigenvectors. Let W be the closure of the span of these vectors. Suppose $W^\perp \neq 0$. Then $T|_{W^\perp}$ is self-adjoint and compact, so there exists an eigenvector for T in $W^\perp$, which contradicts our assumption of maximality. Thus, $W^\perp = 0$, and hence $V = W$.

To see the rest of the conclusions, we must show that $W_\epsilon = \overline{\bigoplus_{|\lambda|>\epsilon} V_\lambda}$ is finite dimensional. Suppose otherwise. Take an orthonormal eigenbasis $(e_i)_{i\in\mathbb{N}}$ of W_ϵ with respect to T, where e_i has eigenvalue λ_i. Then for $i \neq j$,

$$\|Te_i - Te_j\|^2 = \langle Te_i - Te_j, Te_i - Te_j\rangle$$

$$= |\lambda_i|^2 + |\lambda_j|^2$$
$$> 2\epsilon^2$$

Thus, $(Te_i)_{i\in\mathbb{N}}$ has no convergent subsequence, which contradicts compactness of T. ☐

An important extension of the spectral theorem is to commuting families of compact, self-adjoint operators.

Corollary 1.4. *Let V be a Hilbert space, and let $\mathcal{T}$ be a family of pairwise-commuting compact self-adjoint operators on V. Then V has an orthonormal basis $(e_i)_{i\in\mathbb{N}}$ for which each e_i is an eigenvector for every $T \in \mathcal{T}$.*

The spectral theorem for compact normal operators follows from this corollary, as we can write a normal operator as the sum of its self-adjoint and anti-self-adjoint parts, and these will commute with each other.

2. The Finite Dimensional Functional Calculus

We have taken the analogy between the spectral theory of operators on Hilbert spaces and that of operators on finite dimensional spaces about as far as it will go without requiring serious modification. If we drop the assumption that our operators are compact, it is easy to construct examples of bounded self-adjoint operators that have no eigenvectors whatsoever; for example, multiplication by x on $L^2([0,1])$ is certainly bounded, but any eigenvector would have to be supported at a single point, which is absurd.[1]

There is a modern reinterpretation of spectral theory, however, that is much more amenable to the general situation. This reinterpretation involves studying not just single operators, but entire **algebras** of operators, which are collections of operators closed under addition and composition. Let's motivate this shift by looking at the finite dimensional setting.

Let V be a finite dimensional Hilbert space over $\mathbb{C}$, and suppose that $T : V \to V$ is a self-adjoint operator.[2] Recall that any endomorphism of a finite dimensional Hilbert space has a polynomial associated to it called its **minimal polynomial**, which is the monic polynomial m_T of least degree such that $m_T(T) = 0$.

Proposition 2.1. *For any $T : V \to V$, m_T exists and is unique.*

Proof. Since $\mathrm{End}(V)$ is finite dimensional, the set $\{T^k\}_{k\in\mathbb{N}}$ is linearly dependent. Let n be minimal such that $\{T^k\}_{k\le n}$ is linearly dependent. Then we can write $T^n = \sum_{k=0}^{n-1} a_k T^k$ for some $a_k \in \mathbb{C}$. Since $\{T^k\}_{k\le n-1}$ is linearly independent by minimality of n, this decomposition is unique. Thus, $m_T(x) = x^n - \sum_{k=0}^{n-1} a_k x^k$ is the unique minimal polynomial for T. ☐

The spectral theorem allows us to describe the minimal polynomial of T very easily. Since T is self-adjoint, it is diagonalizable, so we may as well assume T is diagonal.

---

[1]It should be noted that physicists have a way around this dilemma by thinking of the Dirac δ-functions as eigenvectors for multiplication by x. These "functions" do not live in L^2 but rather in the larger space of distributions. Making this idea rigorous involves introducing something called a "rigged Hilbert space," which we will not do.

[2]The basic ideas here will also work with normal operators, but for simplicity, we restrict ourselves to self-adjoint operators.

Exercise 2.2. *Show that the minimal polynomial of a matrix is equal to that of any conjugate.*

If $\lambda_1, \ldots, \lambda_k$ are the eigenvalues of T (*not* counted with multiplicity), then the minimal polynomial of T will be $q(x) = \prod_{i=1}^k (x - \lambda_i)$. To see this, we note that $q(T) = 0$, so that m_T divides q. But it is clear that if we remove any linear factor from q, it will no longer satisfy $q(T) = 0$, so we must have $m_T = q$.

Given any polynomial $p \in \mathbb{C}[x]$, it makes sense to evaluate p on T. But since $m_T(T) = 0$, $p(T)$ will only depend on the class of p in $\mathbb{C}[x]$ modulo m_T. In other words, if $p_1 - p_2$ is a multiple of m_T, then $p_1(T) = p_2(T)$.

Let $\sigma(T) = \{\lambda_1, \ldots, \lambda_k\} \subset \mathbb{C}$ be the spectrum of T. Any polynomial $p \in \mathbb{C}[x]$ restricts to a function on $\sigma(T)$. But the value of the function at these points depends only on the class of p modulo m_T, as $m_T \equiv 0$ on $\sigma(T)$. Furthermore, a function $f : \sigma(T) \to \mathbb{C}$, corresponds to a *unique* class in $\mathbb{C}[x]$ modulo m_T, as two polynomials are equal on $\sigma(T)$ if and only if they differ by a multiple of m_T. (Here we are using our characterization of m_T as the product of $x - \lambda_i$.) The upshot is that it makes sense to take a function on $\sigma(T)$ and evaluate it on T by thinking of it as a class of polynomials!

Example 2.3. Let $V = \mathbb{C}^2$, $T = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}$. The minimal polynomial of T is $m_T(x) = (x+1)(x-1) = x^2 - 1$. Given a function f on $\sigma(T) = \{\pm 1\}$, with $f(1) = a$, $f(-1) = b$, we get a polynomial $p(x) = \frac{1}{2}((a-b)x + a + b)$ that agrees with f on $\sigma(T)$. Evaluating p on T gives

$$f(T) = p(T) = \frac{1}{2} \begin{pmatrix} a+b & (a-b)i \\ (b-a)i & a+b \end{pmatrix}.$$

Any other choice of polynomial will differ from p by a multiple of $x^2 - 1$, and so will yield the same matrix when applied to T.

Note that we have used very crucially the fact that T is self-adjoint (or at least normal). In general, the q that we constructed above will merely divide the minimal polynomial, so a function on $\sigma(T)$ may not give rise to a polynomial unique modulo m_T, only modulo q.

Example 2.4. Let $V = \mathbb{C}^2$, $T = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. The minimal polynomial of T is $m_T(x) = x^2$. The polynomials 0 and x both look the same on $\sigma(T) = \{0\}$, but yield very different values when evaluated at T!

Applying functions on the spectrum to self-adjoint operators is called the **functional calculus**. We've defined an isomorphism between a certain algebra of operators, namely polynomials in the self-adjoint operator T, with the algebra of functions on the spectrum. In other words, we have the following theorem.

Theorem 2.5 (Finite Dimensional Functional Calculus)**.** *Let V be a finite dimensional Hilbert space, and let T be a self-adjoint operator. Then there exists an isomorphism $\Phi : C(\sigma(T)) \xrightarrow{\cong} \mathbb{C}[T] \subset \mathrm{End}(V)$ of algebras, and this isomorphism satisfies $\Phi(p)(T) = p(T)$ for any polynomial function p.*

We used the spectral theorem to prove this. But in fact, if we had some independent proof of this theorem, it would easily imply the spectral theorem. For each

$\lambda \in \sigma(T)$, we define $f_\lambda : \sigma(T) \to \mathbb{C}$ by

$$f_\lambda(x) = \begin{cases} 1 & \text{if } x = \lambda, \\ 0 & \text{if } x \neq \lambda. \end{cases}$$

Note that

$$f_\lambda(x) f_\mu(x) = \begin{cases} f_\lambda(x) & \text{if } \lambda = \mu, \\ 0 & \text{if } \lambda \neq \mu. \end{cases}$$

The functional calculus allows us to apply f_λ to T. We can write the function $f(x) = x$ as $f(x) = \sum_{\lambda \in \sigma(T)} \lambda f_\lambda(x)$, so it follows that $T = \sum_{\lambda \in \sigma(T)} \lambda f_\lambda(T)$. Since $f_\lambda^2 = f_\lambda$, it follows that $f_\lambda(T)$ is projection onto some subspace V_λ of V. Note that $T|_{V_\lambda} = \lambda \operatorname{Id}$, since for $v \in V_\lambda$,

$$\begin{aligned} Tv &= T f_\lambda(T) v \\ &= \sum_{\mu \in \sigma(T)} \mu f_\mu(T) f_\lambda(T) v \\ &= \lambda f \lambda^2(T) v \\ &= \lambda v. \end{aligned}$$

Furthermore, $f_\lambda(T)$ is self-adjoint because f_λ can be represented by a polynomial with real coefficients, so it is an orthogonal projection. Since $f_\lambda f_\mu = 0$, it follows that $f_\lambda(T)$ and $f_\mu(T)$ are projections onto orthogonal spaces. Thus, we have written $V = \bigoplus_{\lambda \in \sigma(T)} V_\lambda$, so we have proved the spectral theorem!

Recall our fancier version of the spectral theorem, where we can find orthonormal eigenbases simultaneously for any family of commuting self-adjoint operators. How can this be interpreted in the language of functional calculus?

For a single self-adjoint operator T, the functional calculus gives an isomorphism $\Phi : C(\sigma(T)) \xrightarrow{\cong} \mathbb{C}[T] \subset \operatorname{End}(V)$. There is an obvious generalization of the right-hand side of this isomorphism to a family $\mathcal{T}$ of commuting operators; namely, we can take polynomials in any number of variables in the elements of $\mathcal{T}$. This forms a perfectly nice subalgebra $\mathbb{C}[\mathcal{T}]$ of $\operatorname{End}(V)$.

What about the other side? We need some sort of generalization of the notion of spectrum from the single operator T to the family of commuting operators $\mathcal{T}$. What our suggestive notation says we should try is to formulate $\sigma(T)$ in terms of the algebra $\mathbb{C}[T]$, and then apply whatever construction we used to $\mathbb{C}[\mathcal{T}]$.

Let's start with the $C(\sigma(T))$ side. We can characterize the elements of $\sigma(T)$ as the "places where we can evaluate functions on $\sigma(T)$." Namely, for each $\lambda \in \sigma(T)$, we get a linear functional $\operatorname{ev}_\lambda : C(\sigma(T)) \to \mathbb{C}$ that takes a function $f : \sigma(T) \to \mathbb{C}$ to $f(\lambda)$. Checking dimensions, we find that the $\operatorname{ev}_\lambda$ form a basis for the dual space to $\mathbb{C}(\sigma(T))$. In fact, these are more than just linear functionals; they are *multiplicative* linear functionals: this means that $\operatorname{ev}_\lambda(fg) = \operatorname{ev}_\lambda(f) \operatorname{ev}_\lambda(g)$, and $\operatorname{ev}_\lambda(1) = 1$.

Exercise 2.6. *The $\operatorname{ev}_\lambda$ are the only multiplicative linear functionals on $C(\sigma(T))$.*

Since $C(\sigma(T))$ and $\mathbb{C}[T]$ are isomorphic as algebras, this means we can identify $\sigma(T)$ as the multiplicative linear functionals on $\mathbb{C}[T]$. This suggests that we define $\sigma(\mathcal{T})$ as the set of multiplicative linear functionals on $\mathbb{C}[\mathcal{T}]$.

How can we interpret this definition as something sensible? Suppose $\mathcal{T} = \{T_1, T_2, \ldots, T_n\}$. A multiplicative linear functional $\mathbb{C}[\mathcal{T}] \to \mathbb{C}$ will be determined by where it sends each of the T_i. Where can it send T_i? If λ is not an eigenvalue of

T_i, then $\lambda - T_i$ is invertible, so the functional had better send $\lambda - T_i$ to something nonzero. This means that T_i cannot be sent to λ unless λ is an eigenvalue of T_i! Thus, a multiplicative linear functional can be described as an ordered n-tuple $(\lambda_1, \dots, \lambda_n)$, where λ_i is an eigenvalue of T_i.

But will any ordered n-tuple give rise to a multiplicative functional? No! For example, take $\mathcal{T} = (T_1, T_2) = \left(\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, \begin{pmatrix} c & 0 \\ 0 & d \end{pmatrix} \right)$, where $a, b, c, d \in \mathbb{R}$ are distinct. Can we get the tuple (a, d)? Well, $(T_1 - b)(T_2 - c) = 0$, so it has to be sent to zero under any functional. But this implies that either T_1 has to be sent to b, or T_2 has to be sent to c, so (a, d) is out! An identical argument shows that (b, c) is impossible. What about (a, c) and (b, d)? These are easy: $e_1 = (1, 0)$ and $e_2 = (0, 1)$ are simultaneous eigenvectors for T_1 and T_2 with these eigenvalues. Furthermore, they are simultaneous eigenvectors for everything in $\mathbb{C}[\mathcal{T}]$, so we can define $\Lambda_1 : \mathbb{C}[\mathcal{T}] \to \mathbb{C}$ and $\Lambda_2 : \mathbb{C}[\mathcal{T}] \to \mathbb{C}$ to take an operator $T \in \mathbb{C}[\mathcal{T}]$ to the eigenvalue corresponding to e_1 and e_2, respectively.

Exercise 2.7. *Check that Λ_1 and Λ_2 are multiplicative functionals.*

In general, the allowed n-tuples $(\lambda_1, \dots, \lambda_n)$ will be precisely the sets of eigenvalues corresponding to simultaneous eigenvectors for the $T_1, \dots, T_n$.

Exercise 2.8. *Prove this.*

Now that we understand what $\sigma(\mathcal{T})$ is, we are ready to generalize the functional calculus.

Theorem 2.9 (Finite Dimensional Functional Calculus II)**.** *Let V be a finite dimensional Hilbert space, and let $\mathcal{T}$ be a commuting family of self-adjoint operators. Then there exists an isomorphism $\Phi : C(\sigma(\mathcal{T})) \xrightarrow{\cong} \mathbb{C}[\mathcal{T}] \subset \mathrm{End}(V)$ of algebras, and if we identify a maximal linearly independent set $T_1, \dots, T_n$ in $\mathcal{T}$, this isomorphism satisfies $\Phi(p)(T_1, \dots, T_n) = p(T_1, \dots, T_n)$ for any polynomial function p.*

Finally, we investigate the finite dimensional functional calculus for normal operators. The one place above where we crucially used the fact that T was self-adjoint was showing that $f_\lambda(T)$ is self-adjoint. If T is only normal, this is not at all obvious without an additional application the spectral theorem. Instead, we can replace $\mathbb{C}[T]$ by $\mathbb{C}[T, T^*] \subset \mathrm{End}(V)$, that is, polynomials in both T and T^*. In fact, this is the same algebra as before; T^* can be written as a polynomial in T. So this algebra is still isomorphic to $C(\sigma(T))$. But now both of these algebras have an additional interesting structure: they are **$*$-algebras**. This means that they carry an anti-linear map $x \mapsto x^*$ such that $x^{**} = x$ (i.e., $*$ is an **involution**), $(xy) = y^*x^*$, and $1^* = 1$. On the $C(\sigma(T))$ side, the involution just takes a function to its complex conjugate. On the $\mathbb{C}[T, T^*]$ side, it switches T and T^* and conjugates all coefficients.

Exercise 2.10. *Verify that the isomorphism $\Phi : C(\sigma(T)) \to \mathbb{C}[T, T^*]$ is a **$*$-isomorphism**, i.e., that $\Phi(f^*) = \Phi(f)^*$.*

Exercise 2.11. *Using the fact that Φ is a $*$-isomorphism, show that $f_\lambda(T)$ is self-adjoint.*

With this reformulation, we are now ready to think about spectral theory for noncompact operators. Instead of trying to generalize the version of the spectral theorem that says we can find orthonormal eigenbases for self-adjoint operators, we will generalize the functional calculus.

3. Operator Algebras

The heart of the finite dimensional functional calculus was the identification of the algebra of functions on the spectrum with a certain algebra of operators. What needs to be changed in the infinite dimensional setting? Perhaps most glaring difference is that in the finite dimensional setting, the spectrum was always a discrete set, while in infinite dimensions, the spectrum can have nontrivial topology. This suggests that we could consider continuous functions on the spectrum in this setting; this point of view leads to the **continuous functional calculus**. There are other choices here; we could also work with Borel functions to get the **Borel functional calculus**, and since the spectrum is contained in $\mathbb{C}$, we might consider holomorphic functions and work with a **holomorphic functional calculus**. But continuous functions are in some sense the simplest, so we will work with these.

In the finite dimensional setting, the functional calculus was an isomorphism of two commutative $*$-algebras. But in the infinite dimensional world, the algebra of continuous functions on the spectrum and algebras of operators on the Hilbert space carry additional structure; namely, they have a topology. We start by describing this additional structure abstractly.

Definition 3.1. A **unital Banach algebra** is a Banach space B that is also a unital algebra, such that $\|vw\| \leq \|v\|\|w\|$ for all $v, w \in B$, and $\|1\| = 1$.

Definition 3.2. A **unital C^*-algebra** is a Banach algebra $\mathcal{A}$ that is a $*$-algebra, such that $\|x^*x\| = \|x\|^2$.

Exercise 3.3. *Let V be a Hilbert space. Show that $B(V)$, the algebra of bounded operators on V, is a unital C^*-algebra.*

Exercise 3.4. *Let X be a compact Hausdorff space. Show that $C(X)$ is a commutative unital C^*-algebra.*

Note that any closed sub-$*$-algebra of a C^*-algebra is also a C^*-algebra. So it is natural to try to generalize our one-operator version of the functional calculus as follows.

Theorem 3.5 (Continuous Functional Calculus)**.** *Let V be a Hilbert space, and let $T : V \to V$ be a bounded normal operator. Let $\mathcal{A}_T$ be the closure of $\mathbb{C}[T, T^*]$ inside $B(V)$. Then there is an isomorphism $\Phi : C(\sigma(T)) \to \mathcal{A}_T$ of C^*-algebras such that Φ sends $p \in \mathbb{C}[x, x^*]$ to $p(T, T^*)$.*

We should also expect a generalization to commuting families of bounded normal operators. Recall that in the finite dimensional setting, we formulated this version of the functional calculus by abstractly defining the "spectrum" associated with a commuting family of operators. In fact, this spectrum did not depend on the particular commuting family, but was described entirely in terms of the algebra it generated. This suggests the following definition of spectrum of an abstract commutative unital C^*-algebra.

Definition 3.6. Let $\mathcal{A}$ be a commutative unital C^*-algebra. A **multiplicative functional** on $\mathcal{A}$ is a continuous linear function $x : \mathcal{A} \to \mathbb{C}$ such that $x(ab) = x(a)x(b)$, $x(1) = 1$, and $x(a^*) = x(a)^*$.

The **Gelfand spectrum** of $\mathcal{A}$ is the set $\operatorname{Spec} \mathcal{A}$ of multiplicative functionals on $\mathcal{A}$.

There is a topology on the Gelfand spectrum such that if $x_n \to x$ in $\operatorname{Spec}\mathcal{A}$, then $x_n(a) \to x(a)$ in $\mathbb{C}$ for every $a \in \mathcal{A}$. Using the theory of nets, this can essentially be taken as the definition of the topology, but we will sweep the topological details under the rug. What is important is that the Gelfand spectrum, in this topology, is a compact Hausdorff space.

We know that continuous functions on compact Hausdorff spaces also form a C^*-algebra. Given $a \in \mathcal{A}$, we get a continuous function $\widehat{a}$ on $\operatorname{Spec}\mathcal{A}$ by $\widehat{a}(x) = x(a)$.

Exercise 3.7. *Show that $\widehat{a}$ is continuous, and that the assignment $a \mapsto \widehat{a}$ is a map $\mathcal{A} \to C(\operatorname{Spec}(\mathcal{A}))$ of unital C^*-algebras.*

This map is called the **Gelfand representation** of $\mathcal{A}$. It can be shown that the Gelfand representation is in fact an isomorphism, although we would have to understand the topology better to prove this.

Even more incredibly, if we start with an arbitrary compact Hausdorff space X, take its C^*-algebra $C(X)$ of continuous functions, and take its Gelfand spectrum $\operatorname{Spec} C(X)$, the map $X \to \operatorname{Spec} C(X)$ that takes $x \in X$ to the multiplicative functional "evaluation at x" is a homeomorphism!

These results, which are sometimes referred to as the **commutative Gelfand-Naimark theorem** or the **Gelfand representation theorem**, imply that we can pass freely between commutative unital C^*-algebras and compact Hausdorff spaces. On the one hand, we can take a commutative unital C^*-algebra and get its spectrum, and on the other hand, we can take a compact Hausdorff space and get its algebra of continuous functions. These constructions are inverse to each other. (In more precise language, there is a contravariant equivalence of categories between the category of commutative unital C^*-algebras and the category of compact Hausdorff spaces.)

The commutative Gelfand-Naimark theorem allows us to define the continuous functional calculus abstractly, without any reference to the underlying Hilbert space. Namely, given a commutative unital C^*-algebra $\mathcal{A}$ (which we can take to be an algebra of commuting normal operators on a Hilbert space, if we wish), we get an isomorphism $\Phi : C(\sigma(\mathcal{A})) \to \mathcal{A}$ as the inverse to the Gelfand representation $\mathcal{A} \to C(\sigma(\mathcal{A}))$.

The commutative Gelfand-Naimark theorem immediately implies every version of the functional calculus we've formulated, except for the part about taking polynomials to polynomials. In order to make sense of this part, we need to choose a generating set $\{a_i\}_{i\in I}$ for our C^*-algebra $\mathcal{A}$. We get an embedding $\operatorname{Spec}\mathcal{A} \hookrightarrow \mathbb{C}^I$ via $x \mapsto (x(a_i))_{i\in I}$. A multiplicative functional is determined by its values on a generating set, so this is indeed an embedding. Now we note that if $p \in \mathcal{A}$ is a polynomial in the a_i, say, $p((a_i)_{i\in I}) = \sum_{(J_i)_{i\in I}\in\mathbb{N}^I} b_{(J)} \prod_{i\in I} a_i^{J_i}$, then its image in $C(\operatorname{Spec}(\mathcal{A}))$ under the Gelfand representation will be

$$\begin{aligned} x(p((a_i)_{i\in I})) &= x\left(\sum_{(J_i)_{i\in I}\in\mathbb{N}^I} b_{(J)} \prod_{i\in I} a_i^{J_i}\right) \\ &= \sum_{(J_i)_{i\in I}\in\mathbb{N}^I} b_{(J)} \prod_{i\in I} x(a_i)^{J_i} \\ &= p((x(a_i))_{i\in I}), \end{aligned}$$

which is precisely the polynomial p applied to the image of x under the above embedding of $\operatorname{Spec}\mathcal{A}$ in $\mathbb{C}^I$.

References

[Zim90] Robert J Zimmer, *Essential results of functional analysis*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1990.

Department of Mathematics, University of Chicago, Chicago, IL 60637
E-mail address: ejenkins@math.uchicago.edu