

Lecture 5: October 14, 2024

Lecturer: Avrim Blum (notes based on notes from Madhur Tulsiani)

1 The Real Spectral Theorem

In this lecture, we will prove the “real spectral theorem” for self-adjoint operators $\varphi : V \rightarrow V$ (so named because the eigenvalues of a self-adjoint operator are real, not because other spectral theorems are fake!) We will show that any such operator is not only diagonalizable (has a basis of eigenvectors) but is in fact *orthogonally diagonalizable* i.e., has an *orthonormal* basis of eigenvectors. This gives a very convenient way of thinking about the action of such operators. In particular, let $\dim(V) = n$ and $\{w_1, \dots, w_n\}$ form an orthonormal basis of eigenvectors for φ , with corresponding eigenvalues $\lambda_1, \dots, \lambda_n$. Then for any vector v expressible in this basis as (say) $v = \sum_{i=1}^n c_i \cdot w_i$, we can think of the action of φ as

$$\varphi(v) = \varphi\left(\sum_{i=1}^n c_i \cdot w_i\right) = \sum_{i=1}^n c_i \cdot \lambda_i \cdot w_i.$$

Of course, we can also think of the action of φ in this way as long as w_1, \dots, w_n form a basis (not necessarily orthonormal). However, this view is particularly useful when they form an orthonormal basis. As we will later see, this also provides the “right” basis to think about many matrices, such as the adjacency matrices of graphs (where such decompositions are the subject of spectral graph theory). To prove the spectral theorem, we will need the following statement (which we’ll prove later).

Proposition 1.1 *Let V be a finite-dimensional inner product space (over \mathbb{R} or \mathbb{C}) and let $\varphi : V \rightarrow V$ be a self-adjoint linear operator. Then φ has at least one eigenvalue.*

Using the above proposition, we will prove the spectral theorem below for finite dimensional vector spaces. The proof below can also be made to work for Hilbert spaces (using the axiom of choice). The above proposition, which gives the existence of an eigenvalue is often proved differently for finite and infinite-dimensional spaces, and the proof for infinite-dimensional Hilbert spaces requires additional conditions on the operator φ . We first prove the spectral theorem assuming the above proposition.

Proposition 1.2 (Real spectral theorem) *Let V be a finite-dimensional inner product space and let $\varphi : V \rightarrow V$ be a self-adjoint linear operator. Then φ is orthogonally diagonalizable.*

Proof: By induction on the dimension of V . Let $\dim(V) = 1$. Then by the previous proposition φ has at least one eigenvalue, and hence at least one eigenvector, say w . Then $w / \|w\|$ is a unit vector which forms a basis for V .

Let $\dim(V) = k + 1$. Again, by the previous proposition φ has at least one eigenvector, say w . Let $W = \text{Span}(\{w\})$ and let $W^\perp = \{v \in V \mid \langle v, w \rangle = 0\}$. Check the following:

- W^\perp is a subspace of V .
- $\dim(W^\perp) = k$.
- W^\perp is invariant under φ i.e., $\forall v \in W^\perp, \varphi(v) \in W^\perp$.

Thus, we can consider the operator $\varphi' : W^\perp \rightarrow W^\perp$ defined as

$$\varphi'(v) := \varphi(v) \quad \forall v \in W^\perp.$$

Then, φ' is a self-adjoint (check!) operator defined on the k -dimensional space W^\perp . By the induction hypothesis, there exists an orthonormal basis $\{w_1, \dots, w_k\}$ for W^\perp such that each w_i is an eigenvector of φ . Thus $\left\{w_1, \dots, w_k, \frac{w}{\|w\|}\right\}$ is an orthonormal basis for V , comprising of eigenvectors of φ . ■

2 Existence of eigenvalues

We now prove Proposition 1.1, which shows that a self-adjoint operator must have at least one eigenvalue. Let us begin by considering an easier case, where V is an inner product space over \mathbb{C} . In this case we don't need self-adjointness to guarantee an eigenvalue.

Proposition 2.1 *Let V be a finite dimensional inner product space over \mathbb{C} and let $\varphi : V \rightarrow V$ be a linear operator. Then φ has at least one eigenvalue.*

Proof: Let $\dim(V) = n$. Let $v \in V \setminus 0_V$ be any non-zero vector. Consider the set of $n + 1$ vectors $\{v, \varphi(v), \varphi^2(v), \dots, \varphi^n(v)\}$ where $\varphi^i(v) = \varphi(\varphi^{i-1}(v))$. Since the dimension of V is n , there must exist $c_0, \dots, c_n \in \mathbb{C}$ not all 0 such that

$$c_0 \cdot v + c_1 \cdot \varphi(v) + \dots + c_n \varphi^n(v) = 0_V.$$

For convenience, assume that $c_n \neq 0$, otherwise we can instead consider the sum to the largest i such that $c_i \neq 0$. What we want to do now is to factor the expression above into a product of degree-1 terms. This is where working over \mathbb{C} will be useful.

Let $P(x)$ denote the polynomial $c_0 + c_1x + \dots + c_nx^n$. Then the above can be written as $(P(\varphi))(v) = 0$, where $P(\varphi) : V \rightarrow V$ is a linear operator defined as

$$P(\varphi) := c_0 \cdot \text{id} + c_1 \cdot \varphi + \dots + c_n \varphi^n,$$

with id used to denote the identity operator. Since P is a degree- n polynomial over \mathbb{C} , it can be factored into n linear factors, and we can write $P(x) = c_n \prod_{i=1}^n (x - \lambda_i)$ for $\lambda_1, \dots, \lambda_n \in \mathbb{C}$. This means that we can write

$$P(\varphi) = c_n(\varphi - \lambda_n \cdot \text{id}) \cdots (\varphi - \lambda_1 \cdot \text{id}).$$

Let $w_0 = v$ and define $w_i = \varphi(w_{i-1}) - \lambda_i \cdot w_{i-1}$ for $i \in [n]$. That is, we are working through the computation of $P(\varphi)(v)$ from right to left. Note that $w_0 = v \neq 0_V$ and $w_n = P(\varphi)(v) = 0_V$. Let i^* denote the largest index i such that $w_i \neq 0_V$. Then, we have

$$0_V = w_{i^*+1} = \varphi(w_{i^*}) - \lambda_{i^*+1} \cdot w_{i^*}.$$

This implies that w_{i^*} is an eigenvector with eigenvalue λ_{i^*+1} . ■

We now consider the case that V is a finite dimensional inner product space over \mathbb{R} rather than over \mathbb{C} . In this case, we can no longer necessarily factor P into linear terms, but we *can* factor P into linear and irreducible quadratic terms. What we now need to show is that when we run the argument in the proof above, we hit 0 in one of the linear terms and not one of the irreducible quadratic terms. Specifically, we want to show that we don't get an equation of the form $0_V = \varphi^2(w_{i^*}) + b\varphi(w_{i^*}) + cw_{i^*}$ where $b^2 < 4c$. This is where self-adjointness comes in. In particular, we can write:

$$\begin{aligned} \langle w_{i^*}, \varphi^2(w_{i^*}) + b\varphi(w_{i^*}) + cw_{i^*} \rangle &= \langle w_{i^*}, \varphi^2(w_{i^*}) \rangle + b\langle w_{i^*}, \varphi(w_{i^*}) \rangle + c\langle w_{i^*}, w_{i^*} \rangle \\ &= \langle \varphi(w_{i^*}), \varphi(w_{i^*}) \rangle + b\langle w_{i^*}, \varphi(w_{i^*}) \rangle + c\langle w_{i^*}, w_{i^*} \rangle \\ &= \|\varphi(w_{i^*})\|^2 + b\langle w_{i^*}, \varphi(w_{i^*}) \rangle + c\|w_{i^*}\|^2 \\ &\geq \|\varphi(w_{i^*})\|^2 - |b| \|w_{i^*}\| \|\varphi(w_{i^*})\| + c\|w_{i^*}\|^2 \\ &= \left(\|\varphi(w_{i^*})\| - \frac{|b| \|w_{i^*}\|}{2} \right)^2 + \left(c - \frac{b^2}{4} \right) \|w_{i^*}\|^2 \\ &> 0. \end{aligned}$$

So, the quadratic term can't be equal to 0.

3 Rayleigh quotients: eigenvalues as optimization

Definition 3.1 Let $\varphi : V \rightarrow V$ be a self-adjoint linear operator and $v \in V \setminus \{0_V\}$. The Rayleigh quotient of φ at v is defined as

$$\mathcal{R}_\varphi(v) := \frac{\langle v, \varphi(v) \rangle}{\|v\|^2}.$$

We can equivalently write $\mathcal{R}_\varphi(v) = \langle \hat{v}, \varphi(\hat{v}) \rangle$ for $\hat{v} = v / \|v\|$.

Proposition 3.2 Let $\dim(V) = n$ and let $\varphi : V \rightarrow V$ be a self-adjoint operator with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. Then,

$$\lambda_1 = \max_{v \in V \setminus \{0_V\}} \mathcal{R}_\varphi(v) \quad \text{and} \quad \lambda_n = \min_{v \in V \setminus \{0_V\}} \mathcal{R}_\varphi(v)$$

So, the unit-length vector \hat{v} such that φ applied to it has the largest projection onto itself is the eigenvector of largest eigenvalue, and likewise the one for which φ applied to it has the smallest (or most negative) projection onto itself is the eigenvector of smallest eigenvalue.

Using the above, Rayleigh quotients can be used to prove the spectral theorem for Hilbert spaces, by showing that the above maximum¹ is attained at a point in the space, and defines an eigenvalue if the operator φ is “compact”. A proof can be found in these notes by Paul Garrett [Gar12].

Proposition 3.3 (Courant-Fischer theorem) Let $\dim(V) = n$ and let $\varphi : V \rightarrow V$ be a self-adjoint operator with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. Then,

$$\lambda_k = \max_{\substack{S \subseteq V \\ \dim(S)=k}} \min_{v \in S \setminus \{0_V\}} \mathcal{R}_\varphi(v) = \min_{\substack{S \subseteq V \\ \dim(S)=n-k+1}} \max_{v \in S \setminus \{0_V\}} \mathcal{R}_\varphi(v).$$

Definition 3.4 Let $\varphi : V \rightarrow V$ be a self-adjoint operator. φ is said to be positive semidefinite if $\mathcal{R}_\varphi(v) \geq 0$ for all $v \neq 0$. φ is said to be positive definite if $\mathcal{R}_\varphi(v) > 0$ for all $v \neq 0$.

Proposition 3.5 Let $\varphi : V \rightarrow V$ be a self-adjoint linear operator. Then the following are equivalent:

1. $\mathcal{R}_\varphi(v) \geq 0$ for all $v \neq 0$.
2. All eigenvalues of φ are non-negative.
3. There exists $\alpha : V \rightarrow V$ such that $\varphi = \alpha^* \alpha$.

The decomposition of a positive semidefinite operator in the form $\varphi = \alpha^* \alpha$ is known as the Cholesky decomposition of the operator. Note that if we can write φ as $\alpha^* \alpha$ for any $\alpha : V \rightarrow W$, then this in fact also shows that φ is self-adjoint and positive semidefinite.

¹Strictly speaking, we should write sup and inf instead of max and min until we can justify that max and min are well defined. The difference is that sup and inf are defined as limits while max and min are defined as actual maximum and minimum values in a space, and these may not always exist while we are at looking infinitely many values. Thus, while $\sup_{x \in (0,1)} x = 1$, the quantity $\max_{x \in (0,1)} x$ does not exist. However, in the cases we consider, the max and min will always exist (since our spaces are closed under limits) and we will use max and min in the class to simplify things.

References

- [Gar12] Paul Garrett, *Compact operators on Hilbert space*, 2012,
http://www.math.umn.edu/~garrett/m/fun/Notes/04b_cpt_ops_hsp.pdf.
4