2. Compact operators on Hilbert space

There is a class of bounded linear transformation on a Hilbert space H that is closely analogous to linear transformations between finite-dimensional spaces — the compact operators. Throughout, we will take H to be separable and infinite dimensional. Recall that there is only 'one such H' up to unitary equivalence.

Let us define the closed unit ball $B \subset H$ to be

$$B = \{f \in H \mid \|f\| \leq 1\}.$$

Notice that B is not compact. Indeed, if $e_1, e_2, \ldots$ is an orthonormal basis of H, then this is a sequence in B with no convergent subsequence, since $\|e_i - e_j\| = \sqrt{2}$ if $i \neq j$.

Definition 2.1. A BLT $T : H \to H$ is compact if the closure of $T(B)$ is compact in H. Equivalently, T is compact if, for every bounded sequence f_n, Tf_n contains a convergent subsequence.

Thus, the identity operator on H is not compact. Here are some examples of compact operators:

• Finite rank operators. A BLT is said to be of finite rank if its range is finite dimensional. Let F be a finite rank BLT. Then $F(B)$ is a bounded set contained in a finite dimensional subspace of H. Its closure is therefore compact (since closed, bounded subsets of $\mathbb{C}^n$ are compact).

Exercise. If F is finite rank, let $n = \dim \operatorname{ran}(F)$. Show that there are n vectors $f_1, \ldots, f_n$ so that, with $S = \operatorname{span}(f_1, \ldots, f_n)$, $H = \operatorname{Ker}(F) \oplus S$. Hence, show that F has the form

$$Fk = \sum_{i=1}^{n} g_i(f_i, k).$$

• Integral operators. If $H = L^2([0,1])$, let the operator T be defined by

$$Tf(x) = \int_0^1 K(x,y) f(y)\, dy.$$

Then if $K(x,y)$ is L^2 on $[0,1]^2$, then T is compact. We will show this shortly.

Properties of compact operators

Proposition 2.2 (Proposition 6.1 of SS)**.** *Let T be a bounded linear operator on H.*

(i) If S is a compact operator on H, then ST and TS are compact.

(ii) Suppose that there exists a sequence T_n of compact operators such that $\|T - T_n\| \to 0$ as $n \to \infty$. Then T is compact.

(iii) Every compact operator T is the norm limit of a sequence of finite rank operators.

(iv) T is compact iff T^* is compact.

Remarks on the proof:

• (i) uses some standard point-set topology.

• (ii) uses a diagonal argument.

• (iii) is proved using a family of projection operators associated to an ONB.

• (iv) follows readily from (iii).

Proof: The proof of (i) is straightforward. Let (f_n) be a bounded sequence. Then Tf_n is another bounded sequence, and hence STf_n has a convergent subsequence. Hence ST is compact. Also, we note that Sf_n has a convergent subsequence Sf_{j_n}, and since T is continuous, TSf_{j_n} is convergent. Therefore TS is compact.

(ii) Let f_n be a bounded sequence. Then since T_1 is compact there is a subsequence $f_{1,k}$ such that $T_1 f_{1,k}$ converges. Since T_2 is compact there is a subsequence $f_{2,k}$ of $f_{1,k}$ such that $T_2 f_{2,k}$ converges. And so on: we thus generate a family of nested subsequences $f_{n,k}$. Let $g_k = f_{k,k}$. Then g_k is eventually a subsequence of the nth subsequence $f_{n,k}$ so $T_n g_k$ converges as $k \to \infty$ for each n. We now claim that Tg_k is a Cauchy sequence, and hence convergent. To see this we write

$$\|T(g_k - g_l)\| \leq \|Tg_k - T_m g_k\| + \|T_m(g_k - g_l)\| + \|T_m g_l - Tg_l\|$$

which is valid for any m. Let M be an upper bound on the $\|g_k\|$. Then the first and third terms are bounded by $M\|T - T_m\|$ which is

small provided m is chosen large enough. Fixing any sufficiently large m, the second term is small if k, l are large enough.

(iii) Choose an orthonormal basis $e_1, e_2, \ldots$ and let P_n be the orthogonal projection onto the span of the first n basis vectors, and $Q_n = \mathrm{Id} - P_n$. Then $\|Q_n T f\|$ is a nonincreasing function of n, so therefore $\|Q_n T\|$ is nonincreasing in n. If $\|Q_n T\| = \|P_n T - T\| \to 0$ then the statement is proved, so assume, for a contradiction, that $\|Q_n T\| \geq c$ for all n. Choose f_n, $\|f_n\| = 1$, such that $\|Q_n T f_n\| \geq c/2 \|f_n\|$ for each n. By compactness of T, there is a subsequence such that $T f_{k_n} \to g$ for some g. Then $\|Q_{k_n} T f_{k_n}\| \leq \|Q_{k_n} g\| + \|Q_{k_n}(g - T f_{k_n})\| \leq \|Q_{k_n} g\| + \|g - T f_{k_n}\|$ (since Q_{k_n} always has norm 1), and both the terms on the RHS converge to zero, which is our desired contradiction.

(iv) This follows from parts (ii) and (iii), and from the identity $\|A\| = \|A^*\|$ for all BLT A. □

Corollary 2.3. *Let E be a measurable subset of $\mathbb{R}^n$. Let T be an integral operator on $L^2(E)$ with kernel $K(x, y)$. Assume that $K \in L^2(E^2)$. Then T is a bounded operator with $\|T\| \leq \|K\|_{L^2(E^2)}$. Moreover, T is compact.*

Sketch: Approximate K by linear combinations of functions $\chi_A(x)\chi_B(y)$ for A and B measurable sets in E. The corresponding integral operators are finite rank, and approximate T. □

The book gives a different proof.

An abstraction of this class of operators is the class of Hilbert-Schmidt operators; see SS, p187. A Hilbert-Schmidt operator is one with finite "Hilbert-Schmidt" norm,

$$\|A\|_{HS}^2 = \sum_i \|Ae_i\|^2 .$$

Later we'll be able to show that for every Hilbert-Schmidt operator $T : H \to H$, there is a measure space E, an integral operator S in $L^2(E)$, and a unitary $U : H \to L^2(E)$ so that

$$T = U^* T_K U.$$

Spectral theorem for compact operators

The following important theorem is a direct analogue of the spectral theorem for real symmetric matrices. Before stating it we give some more examples.

Example. Diagonal or 'multiplier' operators. Let $e_1, e_2, \ldots$ be an orthonormal basis of a Hilbert space H, and let $\lambda_1, \lambda_2, \ldots$ be a bounded sequence of complex numbers. Define (if possible) the operator T by $Te_i = \lambda_i e_i$ for all i. Show that

(1) there is a unique bounded operator T with this property, and $\|T\| = \sup_i |\lambda_i|$.

(2) Show that T is compact iff $\lambda_i \to 0$ as $i \to \infty$.

Definition 2.4. We say that an operator $T : H \to H$ is self-adjoint, or symmetric, if $T = T^*$, or equivalently, if $(Tf, g) = (f, Tg)$ for all f, g.

Example. Orthogonal projections are self-adjoint. The operator on $L^2([0, 1])$ mapping $f(x)$ to $xf(x)$ is self-adjoint. The operator mapping $f(x)$ to $e^{ix} f(x)$ is not self-adjoint. Nor is $f(x) \mapsto \int_0^x f(s) ds$ self-adjoint. (What are the adjoints?)

Exercise. Let $K(x, y)$ be a continuous function on $[a, b] \times [a, b]$. Show that the integral operator on $L^2([a, b])$

$$f(x) \mapsto \int_a^b K(x, y) f(y)\, dy$$

is self-adjoint exactly if $K(x, y) = \overline{K(y, x)}$. (If $K(x, y)$ is only bounded and measurable, then the same result holds for a.e. (x, y).)

It turns out that for self-adjoint compact operators, the diagonal example above is in fact the general case:

Theorem 2.5. *Let T be a compact self-adjoint operator on H. Then there is an orthonormal basis $e_1, e_2, \ldots$ of H consisting of eigenvectors of T. Thus $Te_i = \lambda_i e_i$, and we have $\lambda_i \in \mathbb{R}$ and $\lambda_i \to 0$ as $i \to \infty$.*

• This is the analogue in infinite dimensions of the fact that a real symmetric matrix is diagonalizable via an orthogonal matrix.

Steps in the proof:

1. Show that $\|T\| = \sup_{\|f\|=1} |(Tf, f)|$.
2. Show that the quantity on the RHS takes a maximum value at some f which is an eigenvector of T.
3. Eigenspaces of T corresponding to distinct eigenvalues are orthogonal.
4. The operator T restricts to a compact self-adjoint operator $T|_{V^\perp}$ whenever V is an eigenspace, or direct sum of eigenspaces.
5. Thus the direct sum of all eigenspaces must be the whole space.

Proof: 1. Claim: For any self-adjoint operator T (compact or not),

$$\|T\| = \sup_{\|f\|=1} |(Tf, f)|.$$

To see this, we use the characterization

$$\|T\| = \sup_{\|f\|,\|g\|=1} |(Tf, g)|.$$

So, with $M = \sup_{\|f\|=1} |(Tf, f)|$, we have $\|T\| \geq M$. To prove $\|T\| \leq M$, we write using the self-adjointness of T

$$4\,\mathrm{Re}(Tf, g) = (T(f+g), f+g) - (T(f-g), f-g).$$

Then, we get

$$4|\,\mathrm{Re}(Tf, g)| \leq M(\|f+g\|^2 + \|f-g\|^2),$$

and the 'parallelogram law' gives

$$\|f+g\|^2 + \|f-g\|^2 = 2(\|f\|^2 + \|g\|^2) = 4.$$

So $|\,\mathrm{Re}(Tf, g)| \leq M$. Replacing g by $e^{i\theta}g$ we can make $|\,\mathrm{Re}(Tf, g)| = |(Tf, g)|$ and the proof is complete.

2. Therefore, either $T = 0$, in which case the theorem is trivial, or $|(Tf, f)| > 0$ for some f with $\|f\| = 1$. By replacing T with $-T$ if necessary, we can assume that there exists f with $(Tf, f) > 0$ (note that by self-adjointness, (Tf, f) is real for all f).

Consider the problem of maximizing (Tf, f) as f ranges over the unit ball of H. By 1., the set of values (Tf, f), $f \in B$, has a supremum $\mu = \|T\| > 0$, so we may take a sequence f_n, $\|f_n\| = 1$ with

$$(Tf_n, f_n) \to \mu.$$

I claim that

$$\|(T - \mu)f_n\| \to 0.$$

To see this, we square the LHS and compute

$$0 \leq \|(T-\mu)f_n\|^2 = \|Tf_n\|^2 - 2\mu(Tf_n, f_n) + \mu^2$$
$$\leq 2\mu(\mu - (Tf_n, f_n)) \to 0$$

which verifies the claim. Now we exploit compactness of T: the sequence (Tf_n) has a subsequence converging, say to μf. Passing to the subsequence we may assume that the sequence (Tf_n) itself converges. Then f_n converges to f, since

$$\|f_n - f\| \leq \mu^{-1}\Big(\|(T-\mu)f_n\| + \|Tf_n - \mu f\|\Big) \to 0.$$

By continuity of T, $Tf = \lim_n T(f_n) = \mu f$.

Thus we have found an eigenvector f of T.

3. If $Tv = \lambda v$, and $Tw = \mu w$, then we have

$$(v, w) = \lambda^{-1}(Tv, w) = \lambda^{-1}(v, Tw) = \mu\lambda^{-1}(v, w),$$

so $(v, w) = 0$ unless $\lambda = \mu$.

4. Whenever a self-adjoint operator preserves a subspace, i.e $T(v) \in V$ for every $v \in V$, then it also preserves the orthogonal complement, since $(Tx, v) = (x, Tv)$. Certainly T preserves each eigenspace, and thus T restricts to an operator on the orthogonal complement of all eigenspaces. It's easy to see that it is still compact and self-adjoint there.

5. Finally, we see that the orthogonal complement of all eigenspaces must be the zero subspace: otherwise, by the above, T restricts to it and has eigenvector there! □

Example. Be warned: the situation for nonself-adjoint compact operators is quite different. For example, consider the operator UT where T maps e_i to e_i/i and U maps e_i to e_{i+1}. This is compact, but it has no eigenvectors at all.

Applications of the spectral theorem

There are many applications of this result. One I want to mention here is to showing that orthonormal sets are actually bases. For example, suppose we want to show that the orthonormal set of functions

$$(2\pi)^{-1/2}e^{i(n+1/2)\theta}, \quad n \in \mathbb{Z},$$

as elements of $L^2([0,2\pi])$, form an ONB. We can do this by manufacturing a compact self-adjoint operator T for which these functions are the eigenfunctions! Which operator? You might think of $T = id/d\theta$, but this doesn't work because it is not bounded, let alone compact. Instead, we use integration.

Check that the operator

$$f(\theta) \mapsto \frac{i}{2}\Big(\int_0^\theta f(s)\,ds - \int_\theta^{2\pi} f(s)\,ds\Big)$$

is compact and self-adjoint, and that its eigenfunctions are precisely the set $(2\pi)^{-1/2}e^{i(n+1/2)\theta}$, $n \in \mathbb{Z}$.

Sturm-Liouville operators

A Sturm-Liouville operator is an operator $L : C^2([a,b]) \to C([a,b])$ of the form

$$Lf(x) = -f''(x) + q(x)f(x)$$

where $q(x)$ is a continuous function. Here we will assume that $q(x) \geq 0$. We will prove that there is a complete set of eigenfunctions of L in $L^2([a,b])$, that is, functions $\phi_n(x)$ such that

$$L\phi_n(x) = \mu_n\phi_n(x).$$

Notice that if $q(x) \equiv 1$, and $[a,b] = [0,\pi]$, then a complete set of eigenfunctions is the set $\sin nx$, $n = 1, 2, \ldots$. The result can then be viewed as a generalized, 'variable coefficient' version of Fourier series.

As before, the operator L cannot be bounded on L^2, since it involves derivatives. The idea is to construct the inverse operator to L. This can be done is a surprisingly explicit way. What we do is look for two solutions $\phi_-(x)$ and $\phi_+(x)$ of the equation $L\phi = 0$. These are specified by their initial conditions: we require that $\phi_-(a) = 0$, $\phi'_-(a) = 1$, while $\phi_+(b) = 0$, $\phi'_+(b) = 1$. I claim that $\phi_-(b) \neq 0$. Otherwise, compute

$$\begin{aligned} 0 &= \int_a^b \phi_0(x)L\phi_-(x)\,dx \\ &= \int_a^b \phi_-(x)\Big(-\phi''_-(x) + q(x)\phi_-(x)\Big)\,dx \\ &= \int_a^b (\phi'_-(x))^2 + q(x)(\phi_-(x))^2\,dx. \end{aligned}$$

Here we integrated by parts and used the boundary conditions, $\phi_-(a) = \phi_-(b) = 0$ to eliminate the boundary term (which is $\phi_-(b)\phi'_-(b) - \phi_-(a)\phi'_-(a)$). Because we assumed that $q \geq 0$, this can only be if ϕ_- is identically zero, which contradicts the condition $\phi'_-(a) = 1$.

We next conclude that ϕ_- and ϕ_+ are linearly independent; otherwise $\phi_-(b) = 0$.

Recall from ODE theory that the Wronskian,

$$W(x) = \phi_+(x)\phi'_-(x) - \phi_-(x)\phi'_+(x)$$

is constant in x. Evaluating at $x = b$ we see that it is nonzero. We write $W = W(b)$.

Now I claim that the integral operator T with kernel

$$K(x,y) = \begin{cases} \phi_-(x)\phi_+(y)/W, & x \leq y \\ \phi_+(x)\phi_-(y)/W, & x \geq y \end{cases}$$

is an bounded operator on $L^2([a,b])$. An interesting computation shows that, for all continuous $f \in C([a,b])$, Tf is C^2 and

$$L(Tf) = f.$$

(Do it!) However, T is a self-adjoint compact operator, and hence has a complete set of eigenfunctions $\phi_n(x)$ such that $T\phi_n(x) = \lambda_n\phi_n(x)$. It is not hard to check that the range of T consists of continuous functions, so each $\phi_n(x)$ is continuous, and hence C^2. It follows that $L\phi_n(x) = \lambda_n^{-1}\phi_n(x)$. This shows that L has a complete set of eigenfunctions, as claimed.

Remark. Sturm-Liouville operators, and the corresponding differential equations, are very important in physics and applied mathematics. As an example, the (time independent) Schrödinger equation describing the quantum mechanical behaviour of a particle moving on a interval with potential $q(x)$ is exactly the equation $-f''(x) + q(x)f(x) = \lambda f(x)$. For a particle moving in $\mathbb{R}$ rather than a bounded interval $[a, b]$, the analysis above does not apply, and in general the eigenvectors do not form a basis.

Remark. Even though L is not bounded, it can still be understood as a self-adjoint operator on $L^2([a, b])$. There are two technicalities: first, we must restrict the domain to functions whose second derivative lies in L^2, and second, for self-adjointness, we must impose suitable boundary conditions on the functions.

Notice that, from the ODE point of view, in order to solve $Lu = f$ uniquely for a given f, say in $C([a, b])$, we need to specify two values of u, since there are two arbitrary constants in the solution of a second order ODE. You might think it would be natural to specify say $u(a)$ and $u'(a)$, but this does not give a self-adjoint problem. Instead we impose one condition at $x = a$ and one at $x = b$.

To see this, compute for smooth enough functions

$$\begin{aligned}
&(Lf, g) - (f, Lg) \\
&\quad = \int_a^b \Big(\text{-}f''(x) + q(x)f(x)\Big)\overline{g(x)} - \\
&\qquad - f(x)\Big(\text{-}\overline{g}''(x) + q(x)\overline{g}(x)\Big)\, dx \\
&\quad = \int_a^b \Big(- f''(x)\overline{g(x)} + f(x)\overline{g}''(x)\Big)\, dx \\
&\quad = f'(x)\overline{g(x)} - f(x)\overline{g'(x)}\Big|_a^b \\
&\quad = f'(b)\overline{g(b)} - f(b)\overline{g'(b)} - f'(a)\overline{g(a)} + f(a)\overline{g'(a)}
\end{aligned}$$

This vanishes, for example, if we require that f and g vanish at a and b. (Another suitable condition is that f' and g' vanish at a and b.)