FREDHOLM OPERATORS
(NOTES FOR FUNCTIONAL ANALYSIS)

CHRIS WENDL

We assume throughout the following that X and Y are Banach spaces over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ unless otherwise specified. For a bounded linear operator $T \in \mathscr{L}(X, Y)$, we denote its transpose, also known as its *dual operator*, by $T^* : Y^* \to X^* : \Lambda \mapsto \Lambda \circ T$.

Contents

1. Definitions and examples

A bounded linear operator $T : X \to Y$ is called a **Fredholm operator** if

$$\dim \ker T < \infty \qquad \text{and} \qquad \dim \operatorname{coker} T < \infty,$$

where by definition the **cokernel** of T is

$$\operatorname{coker} T := Y \big/ \operatorname{im} T,$$

so the second condition means that the image of T has finite codimension. The **Fredholm index** of T is then the integer

$$\operatorname{ind}(T) := \dim \ker T - \dim \operatorname{coker} T \in \mathbb{Z}.$$

Fredholm operators arise naturally in the study of linear PDEs, in particular as certain types of differential operators for functions on compact domains (often with suitable boundary conditions imposed).

Example 1.1. For periodic functions of one variable $x \in S^1 = \mathbb{R}/\mathbb{Z}$ with values in a finite-dimensional vector space V, the derivative $\partial_x : C^k(S^1) \to C^{k-1}(S^1)$ is a Fredholm operator with index 0 for any $k \in \mathbb{N}$. Indeed,

$$\ker \partial_x = \left\{\text{constant functions } S^1 \to V\right\} \subset C^k(S^1),$$

and

$$\operatorname{im} \partial_x = \left\{ g \in C^{k-1}(S^1) \;\middle|\; \int_{S^1} g(x)\, dx = 0 \right\};$$

the latter follows from the fundamental theorem of calculus since the condition $\int_{S^1} g(x)\, dx = 0$ ensures that the function $f(x) := \int_0^x g(t)\, dt$ on $\mathbb{R}$ is periodic. The surjective linear map

$$C^{k-1}(S^1) \to V : g \mapsto \int_{S^1} g(x)\, dx$$

thus has $\operatorname{im} \partial_x$ as its kernel, so it descends to an isomorphism $\operatorname{coker} \partial_x \to V$, implying $\operatorname{ind}(\partial_x) = \dim V - \dim V = 0$.

Example 1.2. For the same reasons as explained in Example 1.1, $\partial_x : C^{k,\alpha}(S^1) \to C^{k-1,\alpha}(S^1)$ is Fredholm with index 0 for every $k \in \mathbb{N}$ and $\alpha \in (0,1]$.

Exercise 1.3. Use Fourier series to show that the unique extension of $\partial_x : C^\infty(S^1) \to C^\infty(S^1)$ to a bounded linear operator $H^{s+1}(S^1) \to H^s(S^1)$ is also Fredholm with index 0 for every $s \geqslant 0$.

Exercise 1.4. Show that for functions taking values in a vector space V of dimension n, the derivative $\partial_x : C^k([0,1]) \to C^{k-1}([0,1])$ is a surjective Fredholm operator with index n, but imposing the boundary condition $f(0) = f(1) = 0$ produces an injective Fredholm operator

$$\left\{ f \in C^k([0,1]) \;\middle|\; f(0) = f(1) = 0 \right\} \xrightarrow{\partial_x} C^{k-1}([0,1])$$

with index $-n$.

Exercise 1.5. Show that for $n \geqslant 2$ and each $j = 1, \dots, n$, the bounded linear operators $\partial_j : C^k(\mathbb{T}^n) \to C^{k-1}(\mathbb{T}^n)$, $\partial_j : C^{k,\alpha}(\mathbb{T}^n) \to C^{k-1,\alpha}(\mathbb{T}^n)$ and $\partial_j : H^{s+1}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$ have infinite-dimensional kernels and are thus *not* Fredholm.

Example 1.6. The Laplacian $\Delta := \sum_{j=1}^n \partial_j^2$ on fully periodic functions of n variables valued in a finite-dimensional vector space V defines a Fredholm operator

$$\Delta : H^{s+2}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$$

with index 0 for each $s \geqslant 0$. Indeed, if $u \in H^{s+2}(\mathbb{T}^n)$ and $f = \Delta u$, then f has Fourier coefficients

$$\widehat{f}_k = \sum_{j=1}^n \widehat{\partial_j^2 u}_k = -4\pi^2 |k|^2 \widehat{u}_k \in V,$$

thus

$$\begin{aligned} \|f\|_{H^s}^2 &= \sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^s |\widehat{f}_k|^2 = 16\pi^4 \sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^s |k|^4 |\widehat{u}_k|^2 \leqslant 16\pi^4 \sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^{s+2} |\widehat{u}_k|^2 \\ &= 16\pi^4 \|u\|_{H^{s+2}}^2, \end{aligned}$$

proving that Δ is bounded from H^{s+2} to H^s. If $\Delta u = 0$, then $\widehat{f}_k = 0$ for all $k \in \mathbb{Z}^n$, implying $\widehat{u}_k = 0$ for all $k \in \mathbb{Z}^n \backslash \{0\}$, but there is no condition on the coefficient $\widehat{u}_0 \in V$, thus $\ker \Delta$ is the space of functions in $H^{s+2}(\mathbb{T}^n)$ whose only nonvanishing Fourier coefficient is $\widehat{u}_0$, also known as the constant functions $\mathbb{T}^n \to V$. Similarly, the equation $\Delta u = f$ can be solved for a given $f \in H^s(\mathbb{T}^n)$ by writing $\widehat{u}_k = -\frac{1}{4\pi^2|k|^2} \widehat{f}_k$ for all $k \in \mathbb{Z}^n \backslash \{0\}$, but this is only possible if $\widehat{f}_0 = 0$, thus $\operatorname{im} \Delta = \left\{ f \in H^s(\mathbb{T}^n) \;\middle|\; \widehat{f}_0 = 0 \right\}$, and the surjective linear map $H^s(\mathbb{T}^n) \to V : f \mapsto \widehat{f}_0$ therefore descends to an isomorphism $\operatorname{coker} \Delta \to V$. We conclude $\operatorname{ind} \Delta = \dim V - \dim V = 0$.

Exercise 1.7. Show that the wave operator $\partial_t^2 - \partial_x^2 : H^{s+2}(\mathbb{T}^2) \to H^s(\mathbb{T}^2)$ for fully periodic functions of two variables $(t,x) \in \mathbb{R}^2$ has infinite-dimensional kernel, so it is not Fredholm.
Hint: Consider functions of the form $(t,x) \mapsto f(t \pm x)$.

Example 1.8. On any bounded open domain $\Omega \subset \mathbb{R}^n$, the Laplacian defines bounded linear operators $C^{k+2}(\bar{\Omega}) \to C^k(\bar{\Omega})$, $C^{k+2,\alpha}(\Omega) \to C^{k,\alpha}(\Omega)$ for each $k \geqslant 0$ and $\alpha \in (0,1]$, as well as $W^{k+2,p}(\Omega) \to W^{k,2}(\Omega)$ for each $p \in [1,\infty]$, but none of these operators are Fredholm. The reason is that all smooth solutions to the equation $\Delta u = 0$ on $\mathbb{R}^n$ (these are called **harmonic** functions) belong to the kernels of these operators, and there is an infinite-dimensional space of such solutions. This is especially easy to see in the case $n = 2$, where one can identify $\mathbb{R}^2 = \mathbb{C}$ and extract harmonic functions from the real parts of holomorphic functions $\mathbb{C} \to \mathbb{C}$.

2. The Sobolev spaces $H_0^1(\Omega)$ and $H^{-1}(\Omega)$

The Laplacian Δ is the most popular example of an *elliptic* operator; in contrast to the wave operator of Exercise 1.7, it has the right properties to produce a Fredholm operator in suitable functional-analytic settings, as demonstrated by Example 1.6. The problem with Example 1.8

turns out to be not the operator Δ itself, but the fact that it is being considered on a bounded domain without imposing any boundary condition.[1]

To discuss the Laplacian with boundary conditions, it is useful to introduce a few new variations on the usual Sobolev spaces $H^s(\mathbb{R}^n)$. We shall assume in the following that all functions take values in a fixed finite-dimensional complex inner product space $(V, \langle\ ,\ \rangle)$ unless otherwise noted. Recall that $H^s(\mathbb{R}^n)$ is defined for each $s \geqslant 0$ as the space of functions $f \in L^2(\mathbb{R}^n)$ with the property that the product of the Fourier transform $\widehat{f} : \mathbb{R}^n \to V$ with the function $\mathbb{R}^n \to \mathbb{R} : p \mapsto (1+|p|^2)^{s/2}$ is also in $L^2(\mathbb{R}^n)$. The same definition does not quite make sense for $s < 0$ since in that case, $(1+|p|^2)^{s/2}\widehat{f}$ could very well be of class L^2 without $\widehat{f}$ itself being of class L^2, in which case one should not require $H^s(\mathbb{R}^n)$ to be a subspace of $L^2(\mathbb{R}^n)$. The remedy is to define $H^s(\mathbb{R}^n)$ for $s < 0$ as a space of *tempered distributions* rather than functions. In fact, the resulting definition also makes sense for $s \geqslant 0$, but reduces then to the previous definition since tempered distributions whose Fourier transforms are L^2-functions can always be represented by L^2-functions.

Definition 2.1. For any $s \in \mathbb{R}$, we define $H^s(\mathbb{R}^n) \subset \mathscr{S}'(\mathbb{R}^n)$ as the space of all tempered distributions Λ whose Fourier transforms are represented by functions of the form $\widehat{\Lambda}(p) = (1+|p|^2)^{-s/2} f(p)$ for some $f \in L^2(\mathbb{R}^n)$. The H^s-norm is then defined via the inner product

$$\langle \Lambda, \Lambda' \rangle_{H^s} := \left\langle (1+|p|^2)^{s/2}\widehat{\Lambda}, (1+|p|^2)^{s/2}\widehat{\Lambda}' \right\rangle_{L^2} = \int_{\mathbb{R}^n} (1+|p|^2)^s \langle \widehat{\Lambda}(p), \widehat{\Lambda}'(p) \rangle \, dp.$$

It is easy to see that $H^s(\mathbb{R}^n)$ is a Hilbert space, as it admits a natural unitary isomorphism to $L^2(\mathbb{R}^n)$, defined by taking Fourier transforms and multiplying by $(1+|p|^2)^{s/2}$.

Exercise 2.2. Assuming $s \in \mathbb{R}$, prove:

(a) A distribution $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$ is in $H^{-s}(\mathbb{R}^n)$ if and only if it satisfies a bound $|\Lambda(\varphi)| \leqslant c\|\varphi\|_{H^s}$ for all test functions $\varphi \in \mathscr{S}(\mathbb{R}^n)$.

(b) The space of vector-valued Schwartz-class functions $\mathscr{S}(\mathbb{R}^n)$ is dense in $H^s(\mathbb{R}^n)$.

(c) The pairing $\mathscr{S}(\mathbb{R}^n) \times \mathscr{S}(\mathbb{R}^n) \to \mathbb{C} : (\varphi, \psi) \mapsto \langle \varphi, \psi \rangle_{L^2}$ extends to a continuous real-bilinear pairing

$$\langle\ ,\ \rangle_s : H^{-s}(\mathbb{R}^n) \times H^s(\mathbb{R}^n) \to \mathbb{C},$$

$$\langle \Lambda, f \rangle_s := \left\langle (1+|p|^2)^{-s/2}\widehat{\Lambda}, (1+|p|^2)^{s/2}\widehat{f} \right\rangle_{L^2} = \int_{\mathbb{R}^n} \langle \widehat{\Lambda}(p), \widehat{f}(p) \rangle \, dp,$$

such that the real-linear map $\Lambda \mapsto \langle \Lambda, \cdot \rangle_s$ sends $H^{-s}(\mathbb{R}^n)$ isomorphically to the dual space of $H^s(\mathbb{R}^n)$.

Definition 2.3. For each $s \in \mathbb{R}$ and an open subset $\Omega \subset \mathbb{R}^n$, we identify $C_0^\infty(\Omega)$ with the space of smooth functions $\mathbb{R}^n \to V$ that have compact support in Ω, and define the closed subspace

$$\widetilde{H}^s(\Omega) \subset H^s(\mathbb{R}^n)$$

as the closure of $C_0^\infty(\Omega)$ in the H^s-norm. We also define the closed subspace

$$H^s_{\Omega^c}(\mathbb{R}^n) := \left\{ \Lambda \in H^s(\mathbb{R}^n) \;\middle|\; \Lambda(\varphi) = 0 \text{ for all } \varphi \in \mathscr{D}(\Omega) \right\} \subset H^s(\mathbb{R}^n),$$

i.e. $H^s_{\Omega^c}(\mathbb{R}^n)$ is the space of distributions in $H^s(\mathbb{R}^n)$ whose supports are contained in $\Omega^c := \mathbb{R}^n \backslash \Omega$. Finally, we define the quotient Banach space

$$H^s(\Omega) := H^s(\mathbb{R}^n) \big/ H^s_{\Omega^c}(\mathbb{R}^n).$$

[1] The reader might protest at this point that in Exercise 1.4, we saw an example of a differential operator for functions on a bounded domain that was Fredholm despite no mention of any boundary condition. The domain in that example, however, had dimension 1, and since boundaries of 1-dimensional domains are isolated points at which functions take values in a finite-dimensional space, the difference between having a boundary condition or not in this situation is finite dimensional. Moreover, differential equations for functions of one variable are governed by the Picard-Lindelöf theorem on local existence and uniqueness of solutions. The situation for PDEs in more than one variable is radically different.

Exercise 2.4. Given $s \in \mathbb{R}$, let $I : H^{-s}(\mathbb{R}^n) \to (H^s(\mathbb{R}^n))^*$ denote the natural real-linear isomorphism in Exercise 2.2, and assume $\Omega \subset \mathbb{R}^n$ is an open set. Prove:

(a) I maps $H^{-s}_{\Omega^c}(\mathbb{R}^n) \subset H^{-s}(\mathbb{R}^n)$ onto the annihilator of $\widetilde{H}^s(\Omega)$, i.e. the space of bounded linear functionals on $H^s(\mathbb{R}^n)$ that vanish on $\widetilde{H}^s(\Omega)$.

(b) The map $H^{-s}(\mathbb{R}^n) \to \big(\widetilde{H}^s(\Omega)\big)^* : \Lambda \mapsto I(\Lambda)|_{\widetilde{H}^s(\Omega)}$ descends to the quotient $H^{-s}(\Omega) = H^{-s}(\mathbb{R}^n)/H^{-s}_{\Omega^c}(\mathbb{R}^n)$ to define a Banach space isomorphism $H^{-s}(\Omega) \to \big(\widetilde{H}^s(\Omega)\big)^*$.

It is straightforward to show that the obvious inclusion

$$H^s(\mathbb{R}^n) \hookrightarrow \mathscr{S}'(\mathbb{R}^n)$$

is continuous, as the topology defined on $H^s(\mathbb{R}^n)$ by the H^s-norm is stronger than the usual weak*-topology on the space of tempered distributions. Given an open set $\Omega \subset \mathbb{R}^n$, the continuous inclusion $\mathscr{D}(\Omega) \hookrightarrow \mathscr{S}(\mathbb{R}^n)$ dualizes to define a continuous restriction map $\mathscr{S}'(\mathbb{R}^n) \to \mathscr{D}'(\Omega) : \Lambda \mapsto \Lambda|_{\mathscr{D}(\Omega)}$, so that composing this with the inclusion above yields a natural continuous linear map

$$H^s(\mathbb{R}^n) \to \mathscr{D}'(\Omega).$$

The kernel of this map is $H^s_{\Omega^c}(\mathbb{R}^n)$, thus it descends to the quotient $H^s(\Omega) = H^s(\mathbb{R}^n)/H^s_{\Omega^c}(\mathbb{R}^n)$ as a natural continuous linear injection

$$H^s(\Omega) \hookrightarrow \mathscr{D}'(\Omega) : [\Lambda] \mapsto \Lambda|_{\mathscr{D}(\Omega)}. \tag{2.1}$$

It is useful to keep this injection in mind and regard elements of $H^s(\Omega)$ as distributions on Ω: from this perspective, $H^s(\Omega)$ is precisely the space of distributions on Ω that arise as restrictions to Ω of distributions in $H^s(\mathbb{R}^n)$. For elements $[\Lambda] \in H^s(\Omega)$ such that Λ can be represented by a locally integrable function $f : \mathbb{R}^n \to V$ (as is for instance always possible when $s \geqslant 0$), the corresponding distribution on Ω is represented by $f|_\Omega$, which is uniquely determined up to equality almost everywhere on Ω. In particular, $H^s(\Omega)$ for each $s \geqslant 0$ is identified in this way with a linear subspace of $L^2(\Omega)$.

Exercise 2.5. For any integer $k \geqslant 0$ and open subset $\Omega \subset \mathbb{R}^n$, show that the the map $[f] \mapsto f|_\Omega$ defines an injective bounded linear operator $H^k(\Omega) \to W^{k,2}(\Omega)$ with norm at most 1.
Hint: Recall that $H^k(\mathbb{R}^n) = W^{k,2}(\mathbb{R}^n)$, so the H^k-norm for functions on $\mathbb{R}^n$ is equivalent to a norm written in terms of weak derivatives instead of Fourier transforms.

Remark 2.6. The injection $H^k(\Omega) \hookrightarrow W^{k,2}(\Omega)$ is also surjective, and thus a Banach space isomorphism, whenever it can be shown that every $f \in W^{k,2}(\Omega)$ admits an extension over $\mathbb{R}^n$ that belongs to $W^{k,2}(\mathbb{R}^n) = H^k(\mathbb{R}^n)$. This is not true in general, but it is trivially true for $k = 0$, thus giving a natural isomorphism $H^0(\Omega) = L^2(\Omega)$. By a standard extension result in the theory of Sobolev spaces, it is also true for every $k \in \mathbb{N}$ if the boundary of Ω satisfies certain regularity assumptions, e.g. it is true whenever Ω is bounded and its closure in $\mathbb{R}^n$ is a smooth submanifold with boundary.

As a closed subspace of a Hilbert space, $\widetilde{H}^s(\Omega)$ is also a Hilbert space, and a Hilbert space structure can also be assigned to $H^s(\Omega)$ by identifying it with the H^s-orthogonal complement of $H^s_{\Omega^c}(\mathbb{R}^n) \subset H^s(\mathbb{R}^n)$, though for our purposes, it will usually suffice to regard $H^s(\Omega)$ as a Banach space with the natural quotient norm. Notice that since $C_0^\infty(\Omega) \subset H^s(\mathbb{R}^n)$ for every $s \in \mathbb{R}^n$, there is a natural inclusion

$$C_0^\infty(\Omega) \hookrightarrow H^s(\Omega) : f \mapsto [f].$$

The following definition and subsequent proposition are not strictly necessary for our exposition in these notes, but we include them in order to make our notation consistent with what is found in most textbooks.

Definition 2.7. For $s > 0$ and an open subset $\Omega \subset \mathbb{R}^n$, the closed subspace

$$H^s_0(\Omega) \subset H^s(\Omega)$$

is defined as the H^s-closure of $C_0^\infty(\Omega) \subset H^s(\Omega)$.

Remark 2.8. The reason to restrict to $s > 0$ in Definition 2.7 is that for $s \leqslant 0$, $C_0^\infty(\Omega)$ is already dense in $H^s(\Omega)$, so the definition would in those cases give nothing new. The density of $C_0^\infty(\Omega) \subset H^s(\Omega)$ for $s \leqslant 0$ is an easy consequence of the case $s = 0$, for which Exercise 2.5 and Remark 2.6 give a natural identification $H^k(\Omega) = L^2(\Omega)$ and we can appeal to the fact that $C_0^\infty(\Omega)$ is dense in $L^2(\Omega)$. One can then deduce from the density of $\mathscr{S}(\mathbb{R}^n) \subset H^s(\mathbb{R}^n)$ that $L^2(\Omega)$ is dense in $H^s(\Omega)$ for every $s < 0$.

Proposition 2.9. *For any open set $\Omega \subset \mathbb{R}^n$ and each $k \in \mathbb{N}$, there is a natural isomorphism $\widetilde{H}^k(\Omega) \to H_0^k(\Omega) : f \mapsto [f]$.*

Proof. The quotient projection $H^k(\mathbb{R}^n) \to H^k(\Omega) : f \mapsto [f]$ restricts to an injection on the subspace $C_0^\infty(\Omega) \subset H^k(\mathbb{R}^n)$, and since it is bounded with respect to the H^k-norm, it has a unique extension to a bounded linear map $\Phi : \widetilde{H}^k(\Omega) \to H^k(\Omega)$. The closure of the image of Φ is $H_0^k(\Omega)$ by definition, thus our task is to prove that Φ is injective with closed image. To see this, we use the $W^{k,2}$-inner product as a substitute for the H^k-inner product and claim that every $f \in C_0^\infty(\Omega)$ is $W^{k,2}$-orthogonal to the subspace $H^k_{\Omega^c}(\mathbb{R}^n)$. Indeed, elements of $H^k_{\Omega^c}(\mathbb{R}^n)$ can be regarded as functions $g \in L^2(\mathbb{R}^n)$ that vanish almost everywhere on Ω and have weak derivatives $\partial^\alpha g \in L^2(\mathbb{R}^n)$ for every multi-index α of order $|\alpha| \leqslant k$, and it follows that these weak derivatives also vanish almost everywhere on Ω. We thus have

$$\langle f, g\rangle_{W^{k,2}} = \sum_{|\alpha| \leqslant k} \langle \partial^\alpha f, \partial^\alpha g\rangle_{L^2} = 0,$$

as the integrand in each of these L^2-inner products vanishes almost everywhere. Since the $W^{k,2}$- and H^k-norms are equivalent, it follows that the closure $\widetilde{H}^k(\Omega)$ is contained in the $W^{k,2}$-orthogonal complement of $H^k_{\Omega^c}(\mathbb{R}^n) \subset H^k(\mathbb{R}^n)$. The quotient projection $H^k(\mathbb{R}^n) \to H^k(\Omega)$ restricts to the latter space as a Banach space isomorphism, and the restriction of that isomorphism to the smaller closed subspace $\widetilde{H}^k(\Omega)$ is Φ, whose injective image is therefore also a closed subspace. □

For any multi-index α of order $|\alpha| = m$, it is straightforward to show that the differential operator ∂^α defines bounded linear maps

$$\partial^\alpha : H^{s+m}(\mathbb{R}^n) \to H^s(\mathbb{R}^n)$$

for each $s \in \mathbb{R}$, and since ∂^α preserves the space of distributions with support disjoint from any given open set $\Omega \subset \mathbb{R}^n$, ∂^α descends to a bounded linear map of the quotients

$$\partial^\alpha : H^{s+m}(\Omega) \to H^s(\Omega).$$

For $f \in H^{s+m}(\Omega)$ and $g \in H^s(\Omega)$, each represented via (2.1) as distributions on Ω, the meaning of the relation $\partial^\alpha f = g$ can now be understood as follows. By definition, f and g each admit extensions $\widetilde{f} \in H^{s+m}(\mathbb{R}^n)$ and $\widetilde{g} \in H^s(\mathbb{R}^n)$ whose restrictions to Ω are f and g respectively, and $\partial^\alpha f = g$ then holds in $H^s(\Omega) = H^s(\mathbb{R}^n)/H^s_{\Omega^c}(\mathbb{R}^n)$ if and only if $\partial^\alpha \widetilde{f} - \widetilde{g}$ is in $H^s_{\Omega^c}(\mathbb{R}^n)$, which projects trivially to the quotient. This means that $\partial^\alpha \widetilde{f} - \widetilde{g}$ vanishes on all test functions supported in Ω, which is the same as saying that $\partial^\alpha f$ and g are identical distributions in $\mathscr{D}'(\Omega)$. If f and g are locally integrable functions, this means exactly what one would expect: g is equal to a weak derivative $\partial^\alpha f$ over the domain Ω. (Note that the notion of weak differentiation depends on the choice of domain, and no claim is being made here about the relation $\partial^\alpha \widetilde{f} = \widetilde{g}$ holding on any domain larger than Ω.)

Allowing linear combinations of such differential operators, the remarks of the previous paragraph apply in particular to the Laplace operator Δ with $m = 2$. The following can then be regarded as an existence and uniqueness result for distributional solutions of the Poisson equation $\Delta u = f$ with boundary condition $u|_{\partial\bar{\Omega}} \equiv 0$; in particular, it provides for every $f \in L^2(\Omega) \subset H^{-1}(\Omega)$ a unique weak solution u in the space $H^1_0(\Omega)$.

Theorem 2.10. *For any bounded open set $\Omega \subset \mathbb{R}^n$, the Laplacian defines a Fredholm operator*
$$H_0^1(\Omega) \xrightarrow{\Delta} H^{-1}(\Omega)$$
with trivial kernel and cokernel, i.e. it is a Banach space isomorphism.

This result is the first step in the study of the *Dirichlet problem*, which seeks solutions to the Laplace equation $\Delta u = 0$ on a bounded domain $\Omega \subset \mathbb{R}^n$ with prescribed boundary values $u|_{\partial\bar{\Omega}}$ (see for instance [Tay96, §5.1]). We will prove the theorem via a series of exercises in §6, by showing that Δ can be written as the sum of an isomorphism with a compact operator; as a "compact perturbation" of an index 0 Fredholm operator, it is therefore an index 0 Fredholm operator. A straightforward integration by parts argument then shows that the operator is injective, and since its index is 0, surjectivity follows immediately.

3. Main theorems

The simplest examples of Fredholm operators come from finite-dimensional linear algebra: every linear map $A : \mathbb{K}^n \to \mathbb{K}^m$ is Fredholm, and the fact that A descends to an isomorphism $\mathbb{K}^n/\ker A \to \operatorname{im} A$ reveals that the index of A is
$$\dim \ker A - (m - \dim \operatorname{im} A) = n - (n - \ker A) + \dim \operatorname{im} A - m = n - m.$$
Notice that this result depends only on the dimensions of the domain and target of A, not on A itself. For a Fredholm operator $T : X \to Y$ in infinite dimensions, one cannot so readily extract information from the isomorphism $X/\ker T \cong \operatorname{im} T$ since both sides are infinite dimensional. The remarkable fact is that the *index* of T, while dependent on more data than merely the spaces X and Y, still does not change under small perturbations or continuous deformations of T through families of Fredholm operators.

Theorem 3.1. *The set* $\operatorname{Fred}(X, Y) \subset \mathscr{L}(X, Y)$ *of Fredholm operators from X to Y is open, and the function*
$$\operatorname{ind} : \operatorname{Fred}(X, Y) \to \mathbb{Z}$$
is continuous, i.e. it is locally constant.

Corollary 3.2. *For any continuous map* $[0, 1] \to \operatorname{Fred}(X, Y) : s \mapsto T_s$*,* $\operatorname{ind}(T_s)$ *is independent of s.* ☐

Our second main result is the following theorem on "compact perturbations," proved in §6.

Theorem 3.3. *If* $T \in \operatorname{Fred}(X, Y)$ *and* $K \in \mathscr{L}(X, Y)$ *is a compact operator, then $T + K$ is also Fredholm.*

Notice that in the setting of Theorem 3.3, the operators $tK : X \to Y$ are also compact for every $t \in [0, 1]$, giving rise to a continuous family of Fredholm operators $T_s := T + sK$. Corollary 3.2 thus implies
$$\operatorname{ind}(T) = \operatorname{ind}(T + K) \qquad \text{whenever } T \text{ is Fredholm and } K \text{ is compact.}$$
This applies in particular to all operators of the form $\mathbb{1} - K$ for compact $K \in \mathscr{L}(X)$: they are Fredholm with index 0 since the isomorphism $\mathbb{1} : X \to X$ also is. The following consequence of this observation is known as the *Fredholm alternative*, and should remind you of finite-dimensional linear algebra.

Corollary 3.4 (the Fredholm alternative)**.** *For any compact operator $K : X \to X$, exactly one of the following holds:*

(i) The linear homogeneous equation $x - Kx = 0$ has a nontrivial finite-dimensional space of solutions;
(ii) The linear inhomogeneous equation $x - Kx = y$ has a unique solution for every $y \in X$.

Proof. Since $\operatorname{ind}(\mathbb{1} - K) = 0$, it satisfies $\dim \ker(\mathbb{1} - K) = \dim \operatorname{coker}(\mathbb{1} - K) < \infty$. The first alternative occurs when these dimensions are positive, and the second when they are zero. ☐

As a concrete example of Corollary 3.4, consider the Banach space isomorphism $\Delta : H^1_0(\Omega) \to H^{-1}(\Omega)$ for a bounded open subset $\Omega \subset \mathbb{R}^n$. We will see in Exercise 6.2 that the natural inclusion $j : H^1_0(\Omega) \hookrightarrow H^{-1}(\Omega)$ is compact, thus for any scalar $\lambda \in \mathbb{C}$, the operator $\lambda j \Delta^{-1} : H^{-1}(\Omega) \to H^{-1}(\Omega)$ is the composition of a bounded operator with a compact operator and thus compact. The equation

$$\left(\mathbb{1} - \lambda j \Delta^{-1}\right) u = 0$$

is then satisfied if and only if u is a function in $H^1_0(\Omega)$ with $\Delta u = \lambda u$, i.e. it is an eigenfunction of Δ over the domain Ω, with eigenvalue λ. Corollary 3.4 thus guarantees among other things that for any eigenvalue λ of Δ on a bounded domain Ω with the boundary condition $u|_{\partial\bar{\Omega}} \equiv 0$, the corresponding eigenspace is finite dimensional.

Exercise 3.5. Show that for any finite-dimensional vector space V and any continuous periodic functions $g : S^1 \to \mathbb{R}$ and $h : S^1 \to V$ satisfying $\int_{S^1} g(x)\,dx \neq 0$, the linear inhomogeneous differential equation

$$f' + gf = h$$

has a unique periodic solution $f : S^1 \to V$.
Hint: Show that $T : C^1(S^1) \to C^0(S^1) : f \mapsto f' + gf$ is a compact perturbation of the Fredholm operator in Example 1.1.

As preparation for the proof of Theorem 3.1, the following result is in some sense dual to the fact that all finite-dimensional subspaces of a Banach space are complemented.

Lemma 3.6. *If $T \in \mathscr{L}(X, Y)$ has finite-codimensional cokernel, then $\operatorname{im} T \subset Y$ is closed.*

Proof. Choose $w_1, \ldots, w_n \in Y$ such that the equivalence classes $[w_1], \ldots, [w_n]$ form a basis of $\operatorname{coker} T = Y / \operatorname{im} T$, and define the linear injection

$$\Phi : \mathbb{R}^n \hookrightarrow Y : (\lambda_1, \ldots, \lambda_n) \mapsto \sum_{j=1}^{n} \lambda_j w_j.$$

We can use this to define a surjective bounded linear operator

$$\Psi : X \oplus \mathbb{R}^n \to Y : (x, z) \mapsto Tx + \Phi(z),$$

whose kernel is $\ker T \oplus \{0\} \subset X \oplus \mathbb{R}^n$. The surjectivity of this operator implies that it has closed image, so by Exercise 3.7 below, there exists a constant $c > 0$ such that

$$\|\Psi(x, z)\| \geqslant c \cdot \inf_{v \in \ker T} \|(x + v, z)\|$$

for all $(x, z) \in X \oplus \mathbb{R}^n$. In particular, setting $z = 0$ in this estimate yields

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x + v\|,$$

which by Exercise 3.7 implies that $\operatorname{im} T \subset Y$ is closed. □

Exercise 3.7. Prove:

(a) An injective operator $T \in \mathscr{L}(X, Y)$ has closed image if and only if it satisfies a bound of the form $\|Tx\| \geqslant c\|x\|$ for some constant $c > 0$ independent of $x \in X$.
(b) Every $T \in \mathscr{L}(X, Y)$ descends to a bounded linear operator $X / \ker T \to Y : [x] \mapsto Tx$.
(c) An operator $T \in \mathscr{L}(X, Y)$ has closed image if and only if it satisfies a bound of the form

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x + v\|$$

for some constant $c > 0$ independent of $x \in X$.

Remark 3.8. Some sources explicitly include the condition that $\operatorname{im} T \subset Y$ is closed as part of the definition of a Fredholm operator $T : X \to Y$. Lemma 3.6 shows that this is unnecesary, but it makes little difference in practice, as the standard ways of proving that T is a Fredholm operator (e.g. Lemma 5.2 below) typically include an explicit proof that T has closed image.

Since finite-dimensional and finite-codimensional closed subspaces always admit closed complements, we now obtain the following useful picture of an arbitrary Fredholm operator $T_0 : X \to Y$. Let us abbreviate

$$K := \ker T_0 \subset X, \qquad W := \operatorname{im} T_0 \subset Y,$$

and choose closed subspaces $V \subset X$ and $C \subset Y$ such that

$$X = V \oplus K \qquad \text{and} \qquad Y = W \oplus C, \tag{3.1}$$

keeping in mind that $\dim K < \infty$ and $\dim C = \dim(Y/W) < \infty$. The restriction $A_0 := T_0|_V : V \to \operatorname{im} T_0 = W$ is not a bounded linear bijection, so by the inverse mapping theorem, it is a Banach space isomorphism, meaning its inverse $A_0^{-1} : W \to V$ is also bounded. In block form with respect to the splittings (3.1), T_0 now takes the form

$$T_0 = \begin{pmatrix} A_0 & 0 \\ 0 & 0 \end{pmatrix} : V \oplus K \to W \oplus C. \tag{3.2}$$

We can of course use the same splittings to write any other operator $T \in \mathscr{L}(X, Y)$ in a similar block form

$$T = \begin{pmatrix} A & B \\ C & D \end{pmatrix} : V \oplus K \to W \oplus C \tag{3.3}$$

for bounded linear operators $A : V \to W$, $B : K \to W$, $C : V \to C$ and $D : K \to C$, e.g. A is the composition $\Pi_W T \iota_V$ where $\iota_V : V \hookrightarrow X$ is the continuous inclusion and $\Pi_W : Y \to W$ is the continuous projection along C, and so forth. The most useful observation will be that since the space of Banach space isomorphisms $V \to W$ is an open subset of $\mathscr{L}(V, W)$, the term $A : V \to W$ will remain invertible whenever T is sufficiently close to T_0.

Proof of Theorem 3.1. Given $T_0 \in \operatorname{Fred}(X, Y)$, choose splittings as in (3.1) with $K = \ker T_0$ and $W = \operatorname{im} T_0$ in order to write each $T \in \mathscr{L}(X, Y)$ in block from as in (3.3). Since the block $A \in \mathscr{L}(V, W)$ depends continuously on T and the set of invertible bounded linear maps is open, we can define an open neighborhood $\mathcal{U} \subset \mathscr{L}(X, Y)$ of T_0 by

$$\mathcal{U} := \left\{ T \in \mathscr{L}(X, Y) \bigm| A : V \to W \text{ is invertible} \right\}.$$

We claim that every $T \in \mathcal{U}$ is Fredholm, with

$$\dim \ker T \leqslant \dim \ker T_0, \quad \dim \operatorname{coker} T \leqslant \dim \operatorname{coker} T_0, \quad \text{and} \quad \operatorname{ind} T = \operatorname{ind} T_0.$$

To see this, we can associate to each $T \in \mathcal{U}$ a pair of Banach space isomorphisms $\Phi \in \mathscr{L}(X)$ and $\Psi \in \mathscr{L}(Y)$, expressed in block form with respect to the splittings $X = V \oplus K$ and $Y = W \oplus C$ as

$$\Phi := \begin{pmatrix} \mathbb{1} & -A^{-1}B \\ 0 & \mathbb{1} \end{pmatrix}, \qquad \Psi := \begin{pmatrix} \mathbb{1} & 0 \\ -CA^{-1} & \mathbb{1} \end{pmatrix}.$$

That these are both Banach space isomorphisms is straightforward to check: their inverses are namely

$$\Phi^{-1} := \begin{pmatrix} \mathbb{1} & A^{-1}B \\ 0 & \mathbb{1} \end{pmatrix}, \qquad \Psi^{-1} := \begin{pmatrix} \mathbb{1} & 0 \\ CA^{-1} & \mathbb{1} \end{pmatrix}.$$

The linear map $T : X \to Y$ is thus conjugate to

$$T' := \Psi T \Phi = \begin{pmatrix} A & 0 \\ 0 & T^{\text{red}} \end{pmatrix},$$

where we define the "reduced" operator

$$T^{\text{red}} := D - CA^{-1}B \in \mathscr{L}(K, C).$$

There are two crucial things to observe about the block-diagonal operator T': its top left block is invertible, and its bottom right block is a linear map between *finite-dimensional* vector spaces. We thus have

$$\ker T = \Phi(\ker T') = \Phi(\{0\} \oplus \ker T^{\text{red}}),$$

implying $\dim\ker T = \dim\ker T^{\mathrm{red}} \leqslant \dim K = \dim\ker T_0$. Similarly, Ψ maps $\operatorname{im} T' = W \oplus \operatorname{im} T^{\mathrm{red}}$ isomorphically to $\operatorname{im} T$ and thus descends to an isomorphism $\Psi : \operatorname{coker} T' \to \operatorname{coker} T$, where

$$\operatorname{coker} T' = (W \oplus C)/(W \oplus \operatorname{im} T^{\mathrm{red}}) \cong C \oplus \operatorname{im} T = \operatorname{coker} T^{\mathrm{red}},$$

which gives $\dim\operatorname{coker} T = \dim\operatorname{coker} T^{\mathrm{red}} \leqslant \dim C = \dim\operatorname{coker} T_0$. Observe finally that as an operator between finite-dimensional spaces, the index of $T^{\mathrm{red}} : K \to C$ depends only on the spaces theselves, so it is the same as the index of the zero map $K \to C$, giving

$$\operatorname{ind} T = \operatorname{ind} T^{\mathrm{red}} = \operatorname{ind}\left(K \xrightarrow{0} C\right) = \dim K - \dim C = \operatorname{ind} T_0.$$

□

4. Some preparatory results

The results of this and the next section will serve as preparation for the proof of Theorem 3.3 on compact perturbations.

Proposition 4.1. *A normed vector space is finite dimensional if and only if the closed unit ball about the origin is compact.*

Proof. One direction of the statement follows from first-year analysis, since all closed and bounded subsets of finite-dimensional vector spaces are compact. For the converse, assume X is a normed vector space and $\overline{B_1(0)} \subset X$ is compact. We will give an argument that, with minor modifications,[2] also applies to arbitrary topological vector spaces, proving that the finite-dimensional vector spaces are the only *locally compact* topological vector spaces.

Let $\mathcal{U} := B_1(0) \subset X$ and assume its closure $\bar{\mathcal{U}}$ is compact. Observe that for each $x \in X$, the set $x + \frac{1}{2}\mathcal{U} \subset X$ is a neighborhood of x, so compactness implies

$$\mathcal{U} \subset \bar{\mathcal{U}} \subset \bigcup_{i=1}^{n}\left(x_i + \frac{1}{2}\mathcal{U}\right) \tag{4.1}$$

for some finite set $x_1, \dots, x_n \in X$. We will show that the finite-dimensional subspace $V \subset X$ spanned by $x_1, \dots, x_n$ is in fact X. Indeed, (4.1) implies $\mathcal{U} \subset V + \frac{1}{2}\mathcal{U}$, and rescaling then implies $\frac{1}{2}\mathcal{U} \subset V + \frac{1}{4}\mathcal{U}$ since V is a linear subspace, and thus

$$\mathcal{U} \subset V + \frac{1}{2}\mathcal{U} \subset V + \frac{1}{4}\mathcal{U}.$$

Repeating this argument finitely many times produces

$$\mathcal{U} \subset V + \frac{1}{2^n}\mathcal{U}$$

for every $n \in \mathbb{N}$. It follows that every $x \in \mathcal{U}$ belongs for each $n \in \mathbb{N}$ to the ball of radius $1/2^n$ about some point in V, and is therefore in the closure of V. Since $\dim V < \infty$, V is already closed, so this implies $\mathcal{U} \subset V$. For an arbitrary $x \in X$, we can now choose $\epsilon > 0$ so that $\epsilon x \in \mathcal{U}$, and it follows that $x = \frac{1}{\epsilon}\epsilon x \in V$. □

Remark 4.2. Another popular proof of Proposition 4.1 (which however does not generalize to topological vector spaces) uses a basic geometric result called the *Riesz lemma*, which states that for any closed proper subspace V in a normed vector space X,

$$\sup_{x \in X,\ \|x\|=1} \operatorname{dist}(x, V) = 1.$$

If $\dim X = \infty$, one can use this to construct for any $\delta \in (0,1)$ a sequence $x_n \in X$ that satisfies $\|x_n\| = 1$ for all n but $\|x_n - x_m\| \geqslant \delta$ for all $m \neq n$, so that no subsequence can be Cauchy. (See e.g. [BS18, §2.2].) If X is an inner product space, then one can do better and achieve $\delta = \sqrt{2}$ by constructing x_n to be orthonormal.

[2] see in particular https://terrytao.wordpress.com/2011/05/24/locally-compact-topological-vector-spaces/

Proposition 4.3. *If $K \in \mathscr{L}(X, Y)$ is compact, then so is $K^* \in \mathscr{L}(Y^*, X^*)$.*

Proof. Assume $K : X \to Y$ is compact and $\Lambda_n \in Y^*$ is a sequence satisfying $\|\Lambda_n\| \leqslant C$ for some constant $C > 0$. Letting $\overline{B_1(0)} \subset X$ denote the closed unit ball in X, the set

$$M := \overline{K(\overline{B_1(0)})} \subset Y$$

is then compact. The functions $\Lambda_n|_M : M \to \mathbb{K}$ then satisfy

$$|\Lambda_n(y)| \leqslant C \cdot \max_{y \in M} \|y\|$$

and are thus uniformly bounded; they also satisfy the Lipschitz condition

$$|\Lambda_n(y) - \Lambda_n(y')| \leqslant C \cdot \|y - y'\|$$

for $y, y' \in M$, so they are equicontinuous. It now follows from the Arzelà-Ascoli theorem that after replacing Λ_n with a subsequence, the sequence $\Lambda_n|_M : M \to \mathbb{K}$ is uniformly convergent. (Note that in applying the Arzelà-Ascoli theorem, we are using the fact that $M \subset Y$ is compact, which follows from the compactness of K.) This implies that the sequence $K^*\Lambda_n|_{\overline{B_1(0)}} = \Lambda_n \circ K|_{\overline{B_1(0)}} : \overline{B_1(0)} \to \mathbb{K}$ also converges uniformly, hence it is uniformly Cauchy, implying that $K^*\Lambda_n$ is also a Cauchy sequence and therefore convergent in X^*. □

Remark 4.4. We will not need to use this, but the converse of Proposition 4.3 is also true; see [BS18, Theorem 4.28(iii)].

The **annihilator** of a subset $V \subset X$ is defined by

$$V^\perp := \left\{ \Lambda \in X^* \;\middle|\; \Lambda|_V = 0 \right\} \subset X^*,$$

and similarly, the **pre-annihilator** of a set of dual vectors $V \subset X^*$ is

$$^\perp V := \left\{ x \in X \;\middle|\; \Lambda(x) = 0 \text{ for all } \Lambda \in V \right\} \subset X.$$

In other words, $^\perp V = J^{-1}(V^\perp) \subset X$ for the canonical inclusion $J : X \to X^{**}$. It is easy to check that whenever V is a linear subspace of X or X^*, $V^\perp$ or $^\perp V$ respectively is a closed linear subspace.

Exercise 4.5. For a closed subspace $V \subset X$ with inclusion map $i : V \hookrightarrow X$ and quotient projection $\pi : X \to X/V$, prove:

(1) The map $i^* : X^* \to V^*$ descends to a Banach space isomorphism $X^*/V^\perp \to V^*$.
(2) The map $\pi^* : (X/V)^* \to X^*$ defines a Banach space isomorphism onto $V^\perp \subset X^*$.

Proposition 4.6. *For any $T \in \mathscr{L}(X, Y)$,*

$$(\operatorname{im} T)^\perp = \ker T^* \quad \text{and} \quad {}^\perp(\operatorname{im} T^*) = \ker T.$$

If additionally $\operatorname{im} T \subset Y$ is closed, then $\operatorname{im} T^* \subset X^*$ is also closed, and

$$\operatorname{im} T = {}^\perp(\ker T^*) \quad \text{and} \quad \operatorname{im} T^* = (\ker T)^\perp.$$

Proof. The first two equalities are readily verified from the definitions, where in the second case, one needs to use the fact that $y \in Y$ vanishes if and only if $\Lambda(y) = 0$ for every $\Lambda \in Y^*$, which follows from the Hahn-Banach theorem. It is similarly straightforward to verify the inclusions $\operatorname{im} T \subset {}^\perp(\ker T^*)$ and $\operatorname{im} T^* \subset (\ker T)^\perp$.

We claim moreover that $\operatorname{im} T$ is always dense in ${}^\perp(\ker T^*)$. Indeed, consider a bounded linear functional $\Lambda : {}^\perp(\ker T^*) \to \mathbb{K}$ such that $\Lambda|_{\operatorname{im} T} = 0$, and use the Hahn-Banach theorem to extend Λ to a bounded linear functional on Y. Then $\Lambda \in (\operatorname{im} T)^\perp = \ker T^*$, thus $\Lambda(y) = 0$ for all $y \in {}^\perp(\ker T^*)$ by definition, implying that the original unextended functional was trivial. The density of $\operatorname{im} T$ now follows from the Hahn-Banach theorem. This proves that $\operatorname{im} T = {}^\perp(\ker T^*)$ if and only if $\operatorname{im} T$ is closed.

It remains to prove that if $\operatorname{im} T$ is closed, then $(\ker T)^\perp \subset \operatorname{im} T^*$; this will imply that $\operatorname{im} T^*$ is also closed since $(\ker T)^\perp$ always is. As a first step, Exercise 3.7 gives an estimate

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x + v\| \tag{4.2}$$

for some constant $c > 0$. Now suppose $\Lambda \in (\ker T)^\perp \subset X^*$, so $\Lambda(v) = 0$ for all $v \in \ker T$ and thus

$$|\Lambda(x)| = |\Lambda(x + v)| \leqslant \|\Lambda\| \cdot \|x + v\|$$

for all $x \in X$ and $v \in \ker T$. Taking the infimum over $v \in \ker T$ and combining this with (4.2) gives

$$|\Lambda(x)| \leqslant \frac{1}{c}\|\Lambda\| \cdot \|Tx\|. \tag{4.3}$$

To show that $\Lambda \in \operatorname{im} T^*$, observe that there exists a unique bounded linear functional $\lambda_0 : \operatorname{im} T \to \mathbb{K}$ such that

$$\lambda_0(Tx) = \Lambda(x) \quad \text{for all} \quad x \in X;$$

indeed, the value of $\Lambda(x)$ is independent of the choice of $x \in T^{-1}(Tx)$ since $\Lambda(x) = 0$ whenever $Tx = 0$, and the estimate (4.3) implies that this functional is bounded. Extending λ_0 to $\lambda \in Y^*$ via the Hahn-Banach theorem, we now have $\Lambda = \lambda_0 \circ T = \lambda \circ T = T^*\lambda$. □

5. The semi-Fredholm property

A bounded linear map $T : X \to Y$ is said to be **semi-Fredholm** if

$$\dim \ker T < \infty \qquad \text{and} \qquad \operatorname{im} T \text{ is closed.}$$

This condition often turns out to be a convenient stepping stone toward proving the Fredholm property.

Lemma 5.1. *The following conditions on an operator $T \in \mathscr{L}(X, Y)$ are equivalent:*

(1) T and T^* are both semi-Fredholm.
(2) T is Fredholm.
(3) T^* is Fredholm.

Moreover, when these conditions hold,

$$\dim \ker T^* = \dim \operatorname{coker} T \quad \textit{and} \quad \dim \operatorname{coker} T^* = \dim \ker T,$$

hence $\operatorname{ind} T^* = -\operatorname{ind} T$.

Proof. Assume T and T^* are both semi-Fredholm, so $\ker T$ and $\ker T^*$ are both finite dimensional and $\operatorname{im} T$ and $\operatorname{im} T^*$ are both closed. Using the isomorphisms from Exercise 4.5 and Proposition 4.6, we have

$$(\ker T)^* \cong X^*/(\ker T)^\perp = X^*/\operatorname{im} T^* = \operatorname{coker} T^*,$$

and

$$(\operatorname{coker} T)^* = (Y/\operatorname{im} T)^* \cong (\operatorname{im} T)^\perp = \ker T^*,$$

so the finite-dimensionality of $\ker T$ and $\ker T^*$ implies that $\operatorname{coker} T$ and $\operatorname{coker} T^*$ are also both finite dimensional. This proves (1) ⇒ (2) and (1) ⇒ (3), along with the stated relations between dimensions.

If we instead assume T is Fredholm, then Lemma 3.6 and Proposition 4.6 imply that $\operatorname{im} T$ and $\operatorname{im} T^*$ are both closed, so the isomorphisms above still hold, and the finite-dimensionality of $\ker T$ and $\operatorname{coker} T$ implies the same for $\ker T^*$ and $\operatorname{coker} T^*$, proving (2) ⇒ (3). The proof that (3) implies (1) or (2) requires an additional argument to show that $\operatorname{im} T$ is closed whenever $\operatorname{im} T^*$ is closed, but we shall omit this since it is not needed in the sequel. (A proof may be found in [BS18, Theorem 4.16].) □

Lemma 5.2. *An operator $T \in \mathscr{L}(X, Y)$ is semi-Fredholm if and only if there exists a Banach space Z and compact operator $K \in \mathscr{L}(X, Z)$ satisfying*

$$\|x\| \leqslant c\,(\|Tx\| + \|Kx\|) \tag{5.1}$$

for all $x \in X$ and some constant $c > 0$ independent of x.

Proof. Assume $T : X \to Y$ is semi-Fredholm. Its finite-dimensional kernel $Z := \ker T \subset X$ then admits a closed complement, so there is a continuous linear projection map $K : X \to Z$. This operator is clearly compact since it has finite rank. The linear map

$$X \to Y \oplus Z : x \mapsto (Tx, Kx)$$

is then bounded and injective, with image the closed subspace $\operatorname{im} T \oplus Z \subset Y \oplus Z$, so Exercise 3.7 gives the estimate

$$\|(Tx, Kx)\| = \|Tx\| + \|Kx\| \geqslant c\|x\|$$

for some constant $c > 0$ independent of $x \in X$.

Conversely, suppose a compact operator $K : X \to Z$ is given such that the estimate $\|x\| \leqslant c\|Tx\| + c\|Kx\|$ is satisfied. We will show that the closed unit ball in $\ker T$ is compact, implying via Proposition 4.1 that $\dim \ker T < \infty$. Indeed, if $x_n \in \ker T$ is a sequence satisfying $\|x_n\| \leqslant 1$, then after reducing to a subsequence, we can assume Kx_n converges in Z, due to the compactness of K. In particular, Kx_n is a Cauchy sequence, and since $Tx_n = 0$ for all n, applying the estimate (5.1) to $x_n - x_m$ yields

$$\|x_n - x_m\| \leqslant c\|Kx_n - Kx_m\|.$$

This proves that x_n is a Cauchy sequence in X, so x_n converges, and $\ker T$ is therefore finite dimensional.

To prove that $\operatorname{im} T$ is closed, we first simplify the situation by restricting T to a closed subspace $V \subset X$ that is complementary to $\ker T$; such a subspace necessarily exists since finite-dimensional subspaces are always complemented, and the restricted operator $T|_V : V \to Y$ is now injective but has the same image as T. Now if $x_n \in V$ is a sequence such that $Tx_n \to y \in Y$, we claim that x_n must be bounded. If not, then after restricting to a subsequence, we can assume $\|x_n\| \to \infty$ and thus $T(x_n/\|x_n\|) \to 0$, while the boundedness of $x_n/\|x_n\|$ implies without loss of generality that $K(x_n/\|x_n\|)$ converges. Arguing as in the previous paragraph via Cauchy sequences, we now conclude from (5.1) that $x_n/\|x_n\|$ converges to some $x_\infty \in V$ with $\|x_\infty\| = 1$ but $Tx_\infty = 0$, and that is impossible since $\ker T \cap V = \{0\}$. But now that we know x_n is bounded, Kx_n must in turn have a convergent subsequence, while Tx_n converges by assumption, so another application of (5.1) to Cauchy sequences proves that x_n has a subsequence convergent to some element $x \in V$, which must then satisfy $Tx = y$, proving that $\operatorname{im} T$ is closed. □

6. Compact perturbations

Let us now restate and prove Theorem 3.3.

Theorem 6.1. *If $T \in \mathscr{L}(X, Y)$ is Fredholm, then so is $T + K$ for every compact operator $K \in \mathscr{L}(X, Y)$.*

Proof. We claim first that if T is semi-Fredholm, then so is $T + K$. We use the characterization of the semi-Fredholm condition in Lemma 5.2: assume T satisfies an estimate of the form $\|x\| \leqslant c\|Tx\| + c\|K_0x\|$ for some compact operator $K_0 : X \to Z$. The perturbed operator $T + K : X \to Y$ then satisfies

$$\|x\| \leqslant c\|Tx\| + c\|K_0x\| \leqslant c\|(T+K)x\| + c\|Kx\| + c\|K_0x\| = c\,(\|(T+K)x\| + \|K_1x\|)\,,$$

where we define the operator $K_1 : X \to X \oplus Z : x \mapsto (Kx, K_0x)$, and K_1 is compact since both K_0 and K are compact. Lemma 5.2 then implies that $T + K$ is semi-Fredholm.

Now if T is Fredholm, Lemma 5.1 implies that $T^* : Y^* \to X^*$ is also Fredholm, while $K^* : Y^* \to X^*$ is also compact due to Proposition 4.3. The result of the previous paragraph thus implies that $T^* + K^* = (T + K)^*$ is also semi-Fredholm, so by Lemma 5.1, $T + K$ is also Fredholm. $\square$

The next four exercises furnish the proof of Theorem 2.10, which stated that the Laplace operator defines a Banach space isomorphism

$$\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$$

for any bounded open subset $\Omega \subset \mathbb{R}^n$. It will be convenient to assume in the following that

$$\bar{\Omega} \subset (0,1)^n, \tag{6.1}$$

though it should be clear via scaling and translation that if the result is true for this special case, then it is true in general.

Exercise 6.2. Assume (6.1) holds.

(a) Associate to each $f \in C_0^\infty(\Omega)$ the unique function $F \in C^\infty(\mathbb{T}^n)$ such that $f(x) = F(x)$ for $x \in (0,1)^n$. Show that the map $C_0^\infty(\Omega) \to C^\infty(\mathbb{T}^n) : f \mapsto F$ extends to bounded linear injections

$$L^2(\Omega) \hookrightarrow L^2(\mathbb{T}^n) \quad \text{and} \quad \widetilde{H}^1(\Omega) \hookrightarrow H^1(\mathbb{T}^n)$$

whose images are closed.
Hint: Avoid Fourier analysis here by replacing the usual H^1-norm with the equivalent norm $\|u\|_{W^{1,2}} := \sum_{|\alpha| \leqslant 1} \|\partial^\alpha u\|_{L^2}$. This works equally well on $\mathbb{R}^n$ or $\mathbb{T}^n$.

(b) Deduce via the natural isomorphism $\widetilde{H}^1(\Omega) \to H_0^1(\Omega)$ from Proposition 2.9 and the compactness of the inclusion $H^1(\mathbb{T}^n) \hookrightarrow L^2(\mathbb{T}^n)$ that the linear injection $H_0^1(\Omega) \hookrightarrow L^2(\Omega) : [f] \mapsto f|_\Omega$ is also compact.

(c) Deduce that the natural inclusion $H_0^1(\Omega) \hookrightarrow H^{-1}(\Omega)$ is compact by presenting it as a composition of bounded linear operators in which at least one is compact.

Remark: This result is a special case of the Rellich-Kondrachov compactness theorem. Notice that the boundedness of Ω plays an essential role in the proof; by contrast, the inclusion $H^1(\mathbb{R}^n) \hookrightarrow H^{-1}(\mathbb{R}^n)$ for instance is not compact.

Exercise 6.3. Consider the bounded linear operator

$$\widetilde{\Phi} : H^1(\mathbb{R}^n) \to H^{-1}(\mathbb{R}^n) : u \mapsto u - \frac{1}{4\pi^2}\Delta u,$$

which descends to quotients to define a bounded linear operator $\Phi : H^1(\Omega) \to H^{-1}(\Omega)$. Prove:

(a) $\widetilde{\Phi} : H^1(\mathbb{R}^n) \to H^{-1}(\mathbb{R}^n)$ is a unitary isomorphism.

(b) Let $I : H^{-1}(\Omega) \to \left(\widetilde{H}^1(\Omega)\right)^*$ denote the natural real-linear isomorphism from Exercise 2.4, and denote by $Q : H^{-1}(\mathbb{R}^n) \to H^{-1}(\Omega) = H^{-1}(\mathbb{R}^n)/H^{-1}_{\Omega^c}(\mathbb{R}^n)$ the quotient projection. Then the map

$$I \circ Q \circ \widetilde{\Phi}|_{\widetilde{H}^1(\Omega)} : \widetilde{H}^1(\Omega) \to \left(\widetilde{H}^1(\Omega)\right)^*$$

is an isometric real-linear isomorphism. Deduce from this that $\Phi : H^1(\Omega) \to H^{-1}(\Omega)$ restricts to $H_0^1(\Omega)$ as an isomorphism $H_0^1(\Omega) \to H^{-1}(\Omega)$.
Hint: Write down an explicit formula for $IQ\widetilde{\Phi}(u)f$ for $u, f \in \widetilde{H}^1(\Omega)$.

Exercise 6.4. Writing $j : H_0^1(\Omega) \hookrightarrow H^{-1}(\Omega)$ for the natural inclusion, deduce from the formula $\Delta = 4\pi^2(j - \Phi) : H_0^1(\Omega) \to H^{-1}(\Omega)$ and the previous exercises that Δ is a Fredholm operator with index 0.

Exercise 6.5. Prove:

(a) Every $u \in C_0^\infty(\mathbb{R}^n)$ satisfies $-\int_{\mathbb{R}^n} \langle u, \Delta u \rangle \, dm = \int_{\mathbb{R}^n} |\nabla u|^2 \, dm.$

(b) For any open subset $\Omega \subset \mathbb{R}^n$, the relation in part (a) extends to $-\langle \Delta u, u \rangle_1 = \|\nabla u\|_{L^2}^2$ for every $u \in \widetilde{H}^1(\Omega)$, where $\langle\ ,\ \rangle_1 : H^{-1}(\mathbb{R}^n) \times H^1(\mathbb{R}^n) \to \mathbb{C}$ denotes the duality pairing in Exercise 2.2.

(c) The operator $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ is injective.
Caution: This is not difficult, but since $H^{-1}(\Omega)$ is a quotient, it is slightly more complicated than just assuming $u \in \widetilde{H}^1(\Omega)$ satisfies $\Delta u = 0$ and applying part (b).

(d) The operator $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ is also surjective if $\Omega \subset \mathbb{R}^n$ is bounded.

References

[BS18] T. Bühler and D. A. Salamon, *Functional analysis*, Graduate Studies in Mathematics, vol. 191, American Mathematical Society, Providence, RI, 2018.

[Tay96] M. E. Taylor, *Partial differential equations. I*, Applied Mathematical Sciences, vol. 115, Springer-Verlag, New York, 1996.

Institut für Mathematik, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

E-mail address: wendl@math.hu-berlin.de