1

Spectral Problems in Infinite Dimensions

Throughout this book, our goal is to develop *optimal algorithms* for *infinite-dimensional* spectral problems. The spectrum of a linear operator A on a separable Hilbert space $\mathcal{H}$ is the following set of scalars:

$$\mathrm{Sp}(A) = \left\{ z \in \mathbb{C} \colon (A - zI)^{-1} \text{ does not exist as a bounded operator on } \mathcal{H} \right\},$$

where I denotes the identity operator. Spectra generalise the familiar notion of eigenvalues of finite matrices. The central question of this book is how to compute $\mathrm{Sp}(A)$, or more generally, other spectral properties of A. The feasibility of such computations, and the methods employed, depend crucially on the information accessible to algorithms and on the class of operators or spectral property under consideration.

Our programme has two kinds of results:

- establishing *lower bounds* (on problem difficulty) by proving impossibility theorems;
- establishing *upper bounds* (on problem difficulty) by providing convergent algorithms.

Combined, these two types of results enable us to precisely classify the difficulty of computational spectral problems in infinite dimensions and to develop optimal algorithms that represent the boundaries of computational feasibility.

This chapter highlights the interplay between finite and infinite dimensions that underpins these objectives, and underscores the need to classify infinite-dimensional spectral problems. We begin with four key challenges associated with discretisation, which will recur throughout this book. We then introduce the classical computational spectral problem, showing how the resolvent $(A - zI)^{-1}$ emerges as a central tool for algorithm design (including Algorithm 1.1, which converges for general operators while avoiding discretisation pitfalls). Finally, we present several examples illustrating the varying difficulty of problems, with some surprising outcomes. We use $\langle \cdot, \cdot \rangle$ to denote the inner product on $\mathcal{H}$, linear in the first argument, and $\| \cdot \|$ for the induced norm.

1.1 Dangerous Pitfalls of Discretisation and a Glimpse of Their Remedies

The most common approach to infinite-dimensional spectral problems involves discretising to a finite matrix and computing eigenvalues. See [78, 170, 304, 450, 459, 1043] for a small sample, or the in-depth discussion in the notes section of this chapter.

For example, consider the eigenvalues of $\mathcal{P}A\mathcal{P}^*$, where $\mathcal{P}$ is an orthogonal projection onto a finite-dimensional subspace of $\mathcal{D}(A) \subset \mathcal{H}$. A projection is an operator $\mathcal{P}$ with $\mathcal{P}^2 = \mathcal{P}$. It is orthogonal precisely when the associated projection on $\mathcal{H}$ is self-adjoint. In this book,

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

however, we adopt a slightly different notation to keep track of various spaces. Throughout, $\mathcal{P}$ will denote the operator from $\mathcal{H}$ to its range $\mathrm{ran}(\mathcal{P})$, and $\mathcal{P}^*$ the operator from $\mathrm{ran}(\mathcal{P})$ to $\mathcal{H}$. This means that in various places we will write things like "$\mathcal{P}_n^*\mathcal{P}_n$ converges strongly to the identity" where a reader might instead expect "$\mathcal{P}_n$ converges strongly to the identity". The computation of the eigenvalues of $\mathcal{P}A\mathcal{P}^*$ is known as the *finite section method*, and it is the most intensely studied computational method for approximating $\mathrm{Sp}(A)$.

More broadly, following a discretisation, one compares outputs across different discretisation sizes or methods to strengthen one's belief in the accuracy of the results. While such an approach is successful for many problems, it is well known that it can fail catastrophically. It often remains heuristic, lacking convergence guarantees.[1] Some challenges faced include:

- **Spectral pollution.** This term describes the accumulation of false eigenvalues at points not in $\mathrm{Sp}(A)$, as discretisation sizes increase. More generally, we define it as follows.

  Definition 1.1.1 (Spectral pollution) Let $S_n(A) \subset \mathbb{C}$ be approximations of $\mathrm{Sp}(A)$ for $n \in \mathbb{N}$. We say that the sequence $\{S_n(A)\}$ suffers from spectral pollution if there exists $\lambda \in \mathbb{C} \setminus \mathrm{Sp}(A)$ such that $\liminf_{n\to\infty} \mathrm{dist}(\lambda, S_n(A)) = 0$.

  Spectral pollution is a stronger and more dangerous phenomenon than spurious eigenvalues for a fixed discretisation. In practice, it often lacks visible indicators (e.g., non-convergence or instability) that would distinguish it from approximations of the true spectrum.

  Spectral pollution is prevalent and problematic for many operators encountered in applications, such as the Dirac and Schrödinger operators, operators in magnetohydrodynamics, condensed matter physics, and photonic waveguides.

  For example, an operator A is *self-adjoint* if $A = A^*$, where A^* denotes the adjoint of A. For self-adjoint operators, the essential spectrum is the part of the spectrum that does not consist of isolated eigenvalues of finite multiplicity. Later, we will see that spectral pollution can occur in each gap of an extended version of this set.
- **Spectral invisibility.** Discretising an operator can cause us to miss parts of its spectrum, even as the size of the discretisation increases.

  Definition 1.1.2 (Spectral invisibility) Let $S_n(A) \subset \mathbb{C}$ be approximations of $\mathrm{Sp}(A)$ for $n \in \mathbb{N}$. We say that the sequence $\{S_n(A)\}$ suffers from spectral invisibility if there exists $\lambda \in \mathrm{Sp}(A)$ such that $\limsup_{n\to\infty} \mathrm{dist}(\lambda, S_n(A)) > 0$.

  The absence of spectral invisibility is sometimes called *spectral inclusion*.
- **Continuous and essential spectra.** Though $\mathrm{Sp}(A)$ includes the eigenvalues of A, it is often richer. Informally, the *continuous spectrum* can be viewed as a continuum of spectral values where the operator acts continuously with respect to a spectral parameter. We use the term 'continuous spectrum' to mean the continuous part in the spectral type decomposition

[1] In the context of spectral methods for differential equations, Boyd calls this approach to computing spectra the 'no-brain method' in his excellent textbook (self-described as 'algorithms-guidelines-rules-of-thumb'), and highlights its pitfalls [229, Chapter 7]. He writes that understanding how to overcome these pitfalls is *"as valuable as a lifevest and a waterproof radio after your code has been torpedoed by one of the difficulties."*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

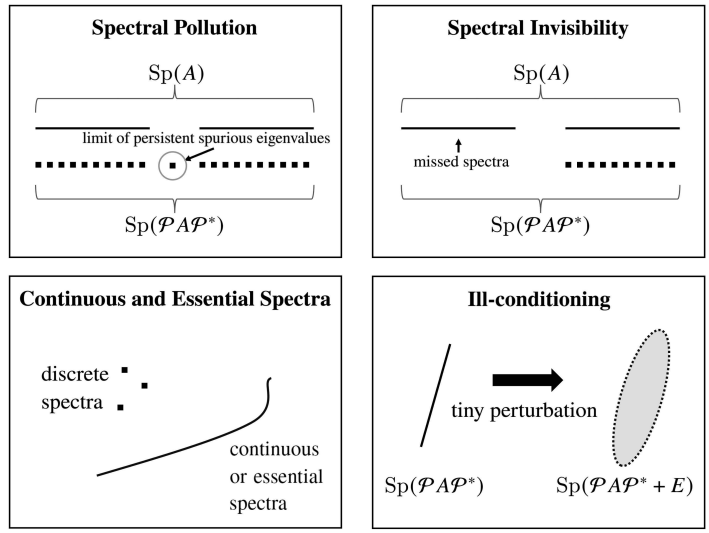


Figure 1.1 Four examples of dangerous pitfalls of discretisation.

for self-adjoint and unitary operators (Section 6.1). If such an operator has a continuous spectral component, its eigenvectors do not form a basis of the Hilbert space.

For general operators, the *essential spectrum* refers to points z in the spectrum where $A - zI$ fails to be invertible in an 'infinite-dimensional' manner (Section 8.1). A compact operator on a Hilbert space is one that can be approximated, in operator norm, by finite-rank operators. One common definition of the essential spectrum is the portion of the spectrum that remains unchanged under compact perturbations of the operator. A full discussion is postponed since there are at least five definitions of essential spectra for non-normal operators! The *discrete spectrum* consists of isolated points z in the spectrum that are eigenvalues with finite multiplicity, reflecting a 'finite-dimensional' failure of the invertibility of $A - zI$.

Discretising to a finite matrix results in a finite, discrete set of eigenvalues. However, many operators encountered in practice possess continuous or essential spectra. *How, then, can we recover these spectra and, in the case of self-adjoint and unitary operators, achieve an analogue of diagonalisation in infinite dimensions?*

- **Ill-conditioning.** A *normal* operator A is one that commutes with its adjoint A^*. It has a well-conditioned spectrum in the sense that $\mathrm{Sp}(A + E) \subset \mathrm{Sp}(A) + B_{\|E\|}(0)$, where $\|E\|$ is the operator norm of E. Nevertheless, discretisation can wreak havoc and cause the spectrum to become severely ill-conditioned, even for normal operators. For a perturbation $E\colon \mathrm{ran}(\mathcal{P}) \to \mathrm{ran}(\mathcal{P})$, $\mathrm{Sp}(\mathcal{P}A\mathcal{P}^* + E)$ and $\mathrm{Sp}(\mathcal{P}A\mathcal{P}^*)$ may differ by much more than $\|E\|$. In such cases, not even stable finite-dimensional solvers can remedy the issue after discretisation.

Another issue is the intrinsic ill-conditioning of an operator's spectrum prior to discretisation. We shall see that some infinite-dimensional spectral problems are unsolvable—even with infinitely precise computations—due to this effect!

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Health warning: Discretisation, such as in the finite section method, can cause both spectral pollution and invisibility, even if A is self-adjoint with compact resolvent (and hence has a well-conditioned, discrete spectrum).

These pitfalls are illustrated in Figure 1.1 and we now present discretisations of normal operators that exhibit each phenomenon. These examples are by no means exhaustive, and readers experienced in infinite-dimensional spectral computations will likely have encountered such issues before. This section is not intended to paint a gloomy picture; these challenges (and others) will each be addressed throughout the book, and we provide a glimpse of their remedies below. Once the necessary skills are acquired, a wide range of possibilities opens up. Extensive code, practice exercises, and detailed examples are included throughout this book to support this development.

1.1.1 Spectral Pollution

Schrödinger operator with constant magnetic field

Our first example of spectral pollution involves a two-dimensional Schrödinger operator on $\mathcal{H} = L^2(\mathbb{R}^2)$ with a constant unit strength magnetic field:

$$T = \left(\mathrm{i}\frac{\partial}{\partial x} + \frac{y}{2}\right)^2 + \left(\mathrm{i}\frac{\partial}{\partial y} - \frac{x}{2}\right)^2.$$

One can show that T is unitarily equivalent to a direct sum $\tilde{T}$ of countably many copies of the harmonic oscillator $-\mathrm{d}^2/\mathrm{d}x^2 + x^2$ [683, Section 10.4]. In particular,

$$\mathrm{Sp}(T) = \mathrm{Sp}(\tilde{T}) = \{2k-1 \colon k \in \mathbb{N}\} = \{1, 3, 5, \ldots\}.$$

The eigenfunctions of the harmonic oscillator are the Hermite functions:

$$\psi_m(x) = \frac{1}{\sqrt{2^m m!\sqrt{\pi}}} e^{-x^2/2} H_m(x), \quad m \in \mathbb{Z}_{\geq 0}, \tag{1.1}$$

where $H_m(x) = (-1)^m e^{x^2} \frac{\mathrm{d}^m}{\mathrm{d}x^m} e^{-x^2}$ is the mth Hermite polynomial. The Hermite functions form an orthonormal basis of $L^2(\mathbb{R})$ and we consider the basis $\{\psi_m(x)\psi_n(y) \colon m, n \in \mathbb{Z}_{\geq 0}\}$ of $L^2(\mathbb{R}^2)$. If we truncate the operator $\tilde{T}$ to act on the subspace

$$\mathcal{H}_N = \mathrm{span}\{\psi_m(x)\psi_n(y) \colon m, n = 0, \ldots, N-1\},$$

there is no problem: we obtain the eigenvalues $1, 3, 5, \ldots, 2N-1$ repeated N times according to multiplicity. Of course, this is cheating, as it assumes prior knowledge of an eigenbasis! In practice, such a basis is almost never known in advance, and many operators do not even possess a basis of eigenfunctions.

In a more realistic scenario, suppose we truncate the original operator T (instead of $\tilde{T}$) by restricting it to the same set of functions in $\mathcal{H}_N$. Let $\mathcal{P}_{\mathcal{H}_N}$ denote the orthogonal projection

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

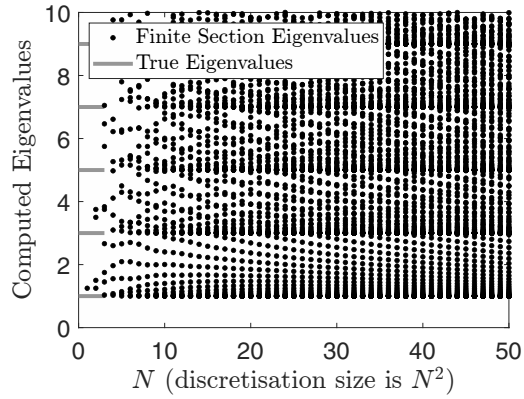


Figure 1.2 Spectral pollution for the two-dimensional magnetic Schrödinger operator T. Finite section eigenvalues are shown as dots; the true eigenvalues (odd integers $1, 3, 5, \ldots$) appear as grey horizontal bars on the left. Each vertical slice shows the eigenvalues for a given truncation parameter N.

onto $\mathcal{H}_N$. The Hermite functions satisfy the following recurrence relations for $m \in \mathbb{N}$:

$$\psi'_m(x) = \sqrt{\frac{m}{2}}\,\psi_{m-1}(x) - \sqrt{\frac{m+1}{2}}\,\psi_{m+1}(x), \tag{1.2}$$

$$x\psi_m(x) = \sqrt{\frac{m}{2}}\,\psi_{m-1}(x) + \sqrt{\frac{m+1}{2}}\,\psi_{m+1}(x). \tag{1.3}$$

Thus, constructing the matrix for the truncated operator $\mathcal{P}_{\mathcal{H}_N} T \mathcal{P}^*_{\mathcal{H}_N}$ is straightforward. Figure 1.2 plots the eigenvalues of this matrix for various values of N, with the true eigenvalues indicated by grey horizontal bars on the left. There is severe spectral pollution in the gaps between these eigenvalues.

Bounded aperiodic operator

Spectral pollution also occurs for bounded operators. As an example, we consider an operator associated with quasicrystals. Quasicrystals extend the notion of a periodic crystal to structures with quasiperiodic translational order.[2] They exhibit symmetries (like five-fold rotational symmetry) that are impossible in traditional periodic crystals. As examples of aperiodic systems, they have attracted attention for their distinct spectral properties [377, 1239], which pose challenges in spectral approximation due to their aperiodicity. They are less well understood than periodic systems, as discussed further in Chapter 7.

We consider a Penrose tiling [563] illustrated in Figure 1.3 (left), which serves as a standard two-dimensional model for quasicrystals [889, 1366]. In this model, the tiling is treated as an infinite graph. A lattice 'site' exists at each vertex (represented by the black dots), with tunnelling bonds along the tile edges. A charged particle can occupy and tunnel between these sites via the bonds. When a perpendicular magnetic field of strength B is applied, it

[2] They were discovered in 1982 by Shechtman [1186]. A natural quasicrystal, icosahedrite, was discovered in 2009 [157]. Shechtman was subsequently awarded the 2011 Nobel Prize in Chemistry for *"a paradigm shift within chemistry."*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

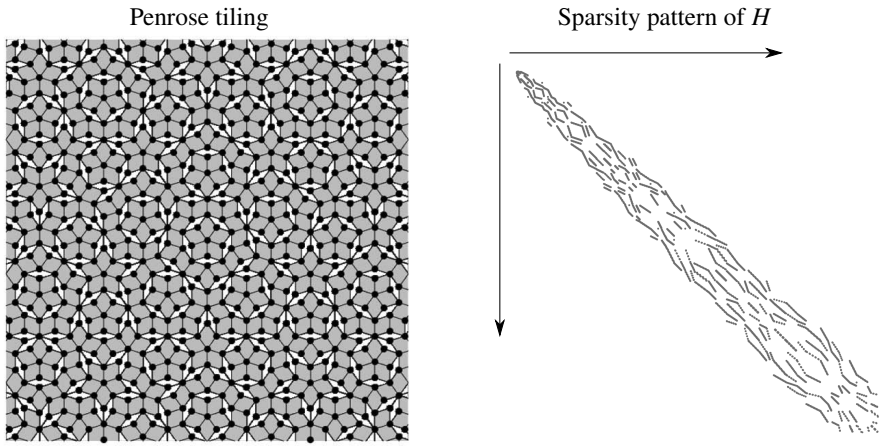


Figure 1.3 Left: Infinite aperiodic Penrose tiling. The model consists of the vertices (sites) and edges between adjacent sites. Right: The corresponding sparsity pattern (nonzero entries) of the infinite matrix of the operator H.

alters the tunnelling strengths. The resulting self-adjoint operator is given by [673]

$$[H\psi](i) = \sum_{i\sim j} e^{\mathrm{i}\alpha_{ji}B}\psi(j), \quad \alpha_{ij} = -\alpha_{ji} \in \mathbb{R}. \tag{1.4}$$

In this equation, i, j label the vertices or sites. The Hilbert space $\mathcal{H}$ corresponds to $\ell^2(V)$, where V is the set of vertices, and a vector $\psi = \{\psi(j)\}_{j\in V} \in \mathcal{H}$ represents the wave function. The notation $i \sim j$ in the sum indicates that an edge connects vertices i and j. Since the degrees of the vertices are uniformly bounded, H is a bounded operator. The phase factor α_{ji} between sites i and j is a line integral from i to j proportional to $\int_i^j (0, x) \cdot \mathrm{d}\mathbf{r}$. We pick a base vertex i_0 as our origin and list the vertices in an order corresponding to radial distance from i_0. This induces an orthonormal basis of $\ell^2(V)$ consisting of δ_j with $\delta_j(i) = \delta_{ij}$. Figure 1.3 (right) displays a finite portion of the resulting infinite sparse matrix that represents H.

As B varies, the spectrum of H changes, forming a fractal, self-similar pattern known as a Hofstadter butterfly [723]. The most common approach for computing $\mathrm{Sp}(H)$ involves truncating the infinite matrix to a finite $n \times n$ matrix, an example of the finite section method. The top of Figure 1.4 shows the resulting spectrum approximations for $n \approx 5000$ at various magnetic field strengths. Large portions of this plot correspond to significant spectral pollution. A physicist might refer to these as 'edge states'. For this model, these states have physical significance and are experimentally important. Often, there is a desire to distinguish them from 'bulk states', which correspond to approximations of $\mathrm{Sp}(H)$. This distinction is closely related to the topological nature of the eigenstates of the truncated operator [1097, 1323].

Glimpse of remedy: In Chapter 3, we introduce a general technique that computes spectra with error control by bounding the spectral distance $\mathrm{dist}(z, \mathrm{Sp}(A))$. The output of this method for this example is presented at the bottom of Figure 1.4. There is now guaranteed convergence, an explicit error bound on points in the output, and no spectral pollution.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

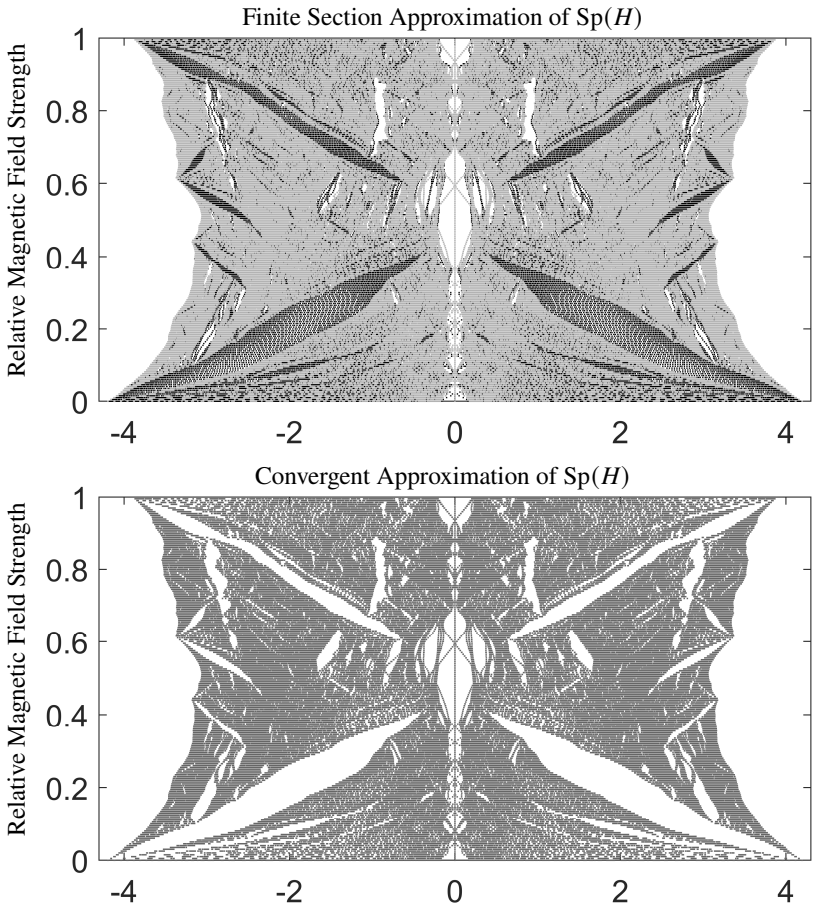


Figure 1.4 Top: Finite section approximation of the spectrum of H. Each horizontal slice corresponds to a spectrum for a fixed value of relative magnetic field B/B_0. Severe spectral pollution (coloured black) occurs in gaps of the essential spectrum. Bottom: Convergent spectrum approximation. There is no spectral pollution, and every point in the output lies within a distance of at most 0.01 from the spectrum.

Where does spectral pollution occur?

Spectral pollution is captured by essential numerical ranges. The *numerical range* of A is

$$W(A) = \{\langle Ax, x\rangle \colon x \in \mathcal{D}(A), \|x\| = 1\}\,. \tag{1.5}$$

The *essential numerical range* of A is [184, 1240]

$$W_e(A) = \bigcap_{B \text{ compact}} \mathrm{Cl}\left(W(A+B)\right). \tag{1.6}$$

It represents the portion of the numerical range that remains stable under compact perturbations of the operator, and it can be empty if A is unbounded.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Self-adjoint operators

If A is self-adjoint ($A = A^*$), then $W_e(A)$ is characterised by the essential spectrum of A. The discrete spectrum $\mathrm{Sp}_{\mathrm{d}}(A)$ consists of isolated eigenvalues of finite multiplicity, and the essential spectrum is $\mathrm{Sp}_{\mathrm{ess}}(A) = \mathrm{Sp}(A) \setminus \mathrm{Sp}_{\mathrm{d}}(A)$. The *extended essential spectrum* is

$$\widehat{\mathrm{Sp}_{\mathrm{ess}}}(A) = \begin{cases} \mathrm{Sp}_{\mathrm{ess}}(A), & \text{if } \mathrm{Sp}(A) \text{ is bounded,} \\ \mathrm{Sp}_{\mathrm{ess}}(A) \cup \{+\infty\}, & \text{if } -\infty < \inf \mathrm{Sp}(A) \text{ and } \sup \mathrm{Sp}(A) = +\infty, \\ \mathrm{Sp}_{\mathrm{ess}}(A) \cup \{-\infty\}, & \text{if } -\infty = \inf \mathrm{Sp}(A) \text{ and } \sup \mathrm{Sp}(A) < +\infty, \\ \mathrm{Sp}_{\mathrm{ess}}(A) \cup \{\pm\infty\}, & \text{if } \mathrm{Sp}(A) \text{ is unbounded above and below.} \end{cases}$$

It should be thought of as a subset of $\mathbb{R} \cup \{\pm\infty\}$. If A is self-adjoint, then $W_e(A) = \mathrm{conv}(\widehat{\mathrm{Sp}_{\mathrm{ess}}}(A)) \setminus \{\pm\infty\}$ [184, Theorem 3.8], where $\mathrm{conv}(S)$ denotes the convex hull of a set S. The following theorem shows that spectral pollution can occur anywhere on $W_e(A) \setminus \mathrm{Sp}(A)$.

Theorem 1.1.3 ([891, Theorem 2.1, Remark 5.1]) *Let A be self-adjoint on $\mathcal{H}$. Then for every $z \in W_e(A) \setminus \mathrm{Sp}_{\mathrm{ess}}(A)$, there exists a sequence of increasing finite-rank orthogonal projections $\{\mathcal{P}_n\}_{n\in\mathbb{N}}$, each with range in the domain of A, such that $\mathcal{P}_n^*\mathcal{P}_n$ converges strongly to the identity and $z \in \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)$ for all $n \in \mathbb{N}$.*

Many self-adjoint operators encountered in applications have $W_e(A) \setminus \mathrm{Sp}(A) \neq \emptyset$. For instance, if $\mathrm{Sp}(A)$ is unbounded above and below, $W_e(A) = \mathbb{R}$ and spectral pollution can occur at every point in $\mathbb{R} \setminus \mathrm{Sp}(A)$. Hence, severe spectral pollution can occur even when the spectrum is discrete and the operator has a compact resolvent.

In Figure 1.2, we observe spectral pollution on $[1, \infty) \setminus \mathrm{Sp}(T)$ for the Schrödinger operator T. This is consistent with Theorem 1.1.3, as the eigenvalues of T have infinite multiplicity so lie in the essential spectrum, leading to $W_e(T) = [1, \infty)$.

In the case of the quasicrystal example and the top of Figure 1.4, we see spectral pollution in some, but not all, gaps of the essential spectrum. This illustrates that, for a given operator A and sequence of finite-rank orthogonal projections, spectral pollution may not occur over the entirety of $W_e(A) \setminus \mathrm{Sp}(A)$.

General operators

For general bounded operators, the following theorem addresses the behaviour of the finite section method outside and inside $W_e(A)$.

Theorem 1.1.4 ([1074, Theorem 1 and Corollary 2]) *Let A be a bounded operator on $\mathcal{H}$ and $\{\mathcal{P}_n\}$ finite-rank orthogonal projections such that $\mathcal{P}_n^*\mathcal{P}_n$ converges strongly to I.*

- *If $z \notin W_e(A)$, then $z \in \mathrm{Sp}(A)$ if and only if $\liminf_{n\to\infty} \mathrm{dist}(z, \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)) = 0$ (and if this occurs $\liminf$ can be replaced by $\lim$).*
- *If $S \subset W_e(A)$ is compact, then there exist finite-rank orthogonal projections $\{\mathcal{Q}_n\}$ such that $\mathrm{ran}(\mathcal{P}_n) \subset \mathrm{ran}(\mathcal{Q}_n)$ and*

$$\lim_{n\to\infty} \sup_{x \in \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*) \cup S} \mathrm{dist}(x, \mathrm{Sp}(\mathcal{Q}_n A \mathcal{Q}_n^*)) = \lim_{n\to\infty} \sup_{x \in \mathrm{Sp}(\mathcal{Q}_n A \mathcal{Q}_n^*)} \mathrm{dist}(x, \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*) \cup S) = 0.$$

That is, $\mathrm{Sp}(\mathcal{Q}_n A \mathcal{Q}_n^*)$ and $\mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*) \cup S$ approach each other as $n \to \infty$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Hence, spectral pollution for a bounded operator A is confined to $W_e(A) \setminus \mathrm{Sp}(A)$, where it can be arbitrarily severe. For general unbounded operators, similar theorems show that $W_e(A)$ captures spectral pollution when A is approximated either by projection methods or by domain truncation methods; see [184, Sections 6–7] and [183].

1.1.2 Spectral Invisibility

As an example of spectral invisibility, we consider the following convection–diffusion operator on $L^2(\mathbb{R})$ [409, 413, 1108, 1325]:

$$[\mathcal{L}u](x) = -\frac{\mathrm{d}^2u}{\mathrm{d}x^2}(x) - 2\frac{\mathrm{d}u}{\mathrm{d}x}(x), \quad \mathcal{D}(\mathcal{L}) = H^2(\mathbb{R}). \tag{1.7}$$

A standard approach for such operators is to truncate the physical domain $\mathbb{R}$ to a finite interval, discretise the operator, and compute eigenvalues of the resulting matrix.

As a first attempt, note that the domain of the operator $\mathcal{D}(\mathcal{L})$ imposes decay of functions at $\pm\infty$. Hence, we might truncate to the interval $[-n, n]$ and apply *Dirichlet boundary conditions*. We write $u(x) = e^{-x}v(x)$ to see that

$$[(\mathcal{L} - \lambda I)u](x) = e^{-x}\left(-\frac{\mathrm{d}^2v}{\mathrm{d}x^2}(x) + (1-\lambda)v(x)\right).$$

Since e^{-x} is bounded above and bounded below away from zero on $[-n, n]$, the domain-truncated operator is similar to the differential operator in brackets and has a complete set of eigenfunctions and eigenvalues:

$$u_m(x) = e^{-x}\sin\left(\frac{m\pi}{2}\left(\frac{x}{n}+1\right)\right), \quad \lambda_m = 1 + \frac{m^2\pi^2}{4n^2}, \quad m \in \mathbb{N}. \tag{1.8}$$

This set of eigenvalues converges to $[1, +\infty)$ as $n \to \infty$. However, we immediately observe a potential issue: as $n \to \infty$, the exponential factor e^{-x} in the eigenfunctions grows without bound and any nonzero limit of these eigenfunctions fails to lie in $L^2(\mathbb{R})$. The similarity transformation corresponding to multiplication by e^{-x} is not uniformly bounded. Figure 1.5 shows several eigenfunctions together with the condition number of the eigenbasis, illustrating the resulting exponential ill-conditioning as $n \to \infty$.

To determine the spectrum of $\mathcal{L}$, one might instead look for (non-normalisable but bounded) generalised eigenfunctions of the form $e^{-\mathrm{i}kx}$, where $k \in \mathbb{R}$. In this case, the spectral parameter is given by $\lambda = k^2 + 2\mathrm{i}k$. This can be made rigorous using the Fourier transform, exploiting the fact that $\mathcal{L}$ has constant coefficients, to see that $\mathrm{Sp}(\mathcal{L}) = \{k^2 + 2\mathrm{i}k : k \in \mathbb{R}\}$ (Exercise 1.2). Figure 1.6 plots the spectrum of $\mathcal{L}$ together with the eigenvalues of the operator with Dirichlet boundary conditions on the truncated domain for $n = 2$ and $n = 10$. This picture demonstrates both spectral pollution and spectral invisibility of domain truncation.

Given the use of the Fourier transform above, it is tempting to truncate to $[-n, n]$ and apply *periodic boundary conditions*. The standard Fourier basis on $[-n, n]$ provides eigenfunctions with eigenvalues $\{k^2 + 2\mathrm{i}k : k = m\pi/n, m \in \mathbb{Z}\}$, which converge to $\mathrm{Sp}(\mathcal{L})$ as $n \to \infty$.

This example is simple enough for us to fix the discretisation issues by adjusting the boundary conditions. However, similar strategies cannot always be applied, e.g., for variable coefficient operators. Unfortunately, no universal method currently exists for choosing boundary conditions in domain truncations that guarantees convergence to the correct spectrum.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

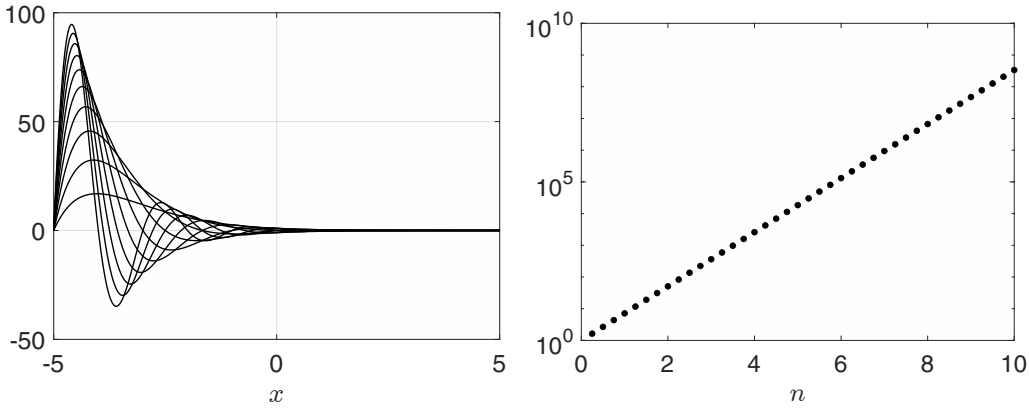


Figure 1.5 Left: The first ten eigenfunctions in (1.8) for $n = 5$. The exponential envelope e^{-x} is clearly visible. Right: The condition number of the set of the first 100 eigenfunctions, each normalised to have unit norm, for different values of n. The basis is exponentially ill-conditioned in n.

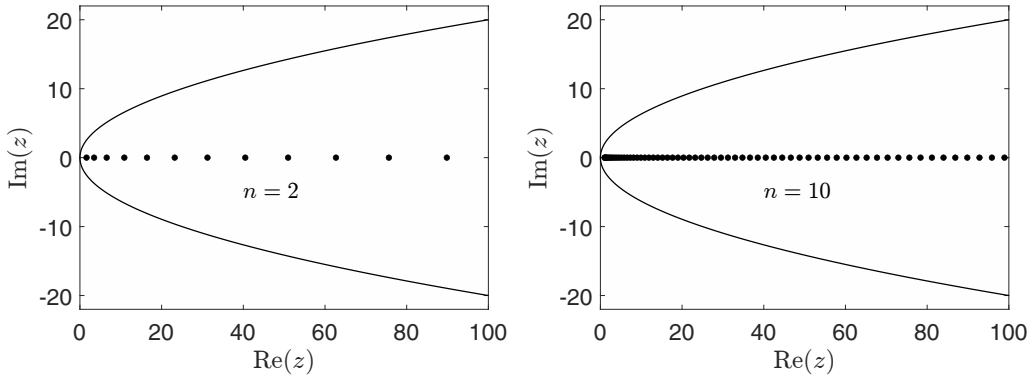


Figure 1.6 The spectrum of the operator $\mathcal{L}$ in (1.7) is the parabola shown in the figure. The spectra of the operators on the truncated domain $[-n, n]$ with Dirichlet boundary conditions are shown as dots for $n = 2$ (left) and $n = 10$ (right). As $n \to \infty$, these eigenvalues fill the interval $[1, +\infty)$.

Unthinkingly discretising an operator may entirely miss parts of the spectrum, a common issue with non-self-adjoint operators. For example, Figure 3.15 presents results for an operator for which neither Dirichlet nor periodic boundary conditions recover the full spectrum.

Glimpse of remedy: In Section 3.4.2, we develop general methods for differential operators on unbounded domains, using the approach that produced the bottom panel of Figure 1.4.

Where does spectral invisibility occur?

While the essential numerical range characterises spectral pollution, no analogous characterisation for spectral invisibility (Definition 1.1.2) is currently known.

However, the following is known. Suppose that $\{\mathcal{P}_n\}$ is a sequence of orthogonal projections such that $\mathcal{P}_n^* \mathcal{P}_n$ converges strongly to I. If A is a bounded self-adjoint operator, then

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

spectral invisibility cannot occur for the approximation sequence $\{\mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)\}$ [52]:

$$\limsup_{n\to\infty} \mathrm{dist}(z, \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)) = 0 \quad \forall z \in \mathrm{Sp}(A).$$

Spectral invisibility is also avoided for unbounded self-adjoint operators under suitable conditions (Corollary 4.2.7 and Lemma 4.2.14). Nevertheless, the following theorem highlights the delicate nature of selecting appropriate projections for the finite section method. It says that under mild conditions, there exist projections that are arbitrarily bad yet arbitrarily close to an original choice of projections.

Theorem 1.1.5 ([891, Theorem 2.3]) *Let A be a self-adjoint operator on $\mathcal{H}$ that is unbounded above, let $\{\mathcal{P}_n\}_{n\in\mathbb{N}}$ be a sequence of finite-rank orthogonal projections each with range in the domain of A, and let $\{\epsilon_n\}$ and $\{R_n\}$ be positive sequences with $\epsilon_n \downarrow 0$ and $R_n \uparrow \infty$. Then there exists a sequence of orthogonal projections $\{\mathcal{P}'_n\}_{n\in\mathbb{N}}$ (each with range in the domain of A) such that*

$$\|\mathcal{P}_n - \mathcal{P}'_n\| \leq \epsilon_n \quad \text{and} \quad \mathrm{Sp}(\mathcal{P}'_n A \mathcal{P}'^*_n) \subset [R_n, +\infty) \quad \forall n \in \mathbb{N}.$$

A similar statement holds if A is unbounded below by considering $-A$.

Beyond self-adjoint operators, avoiding spectral invisibility becomes more murky. Spectral invisibility can occur even for finite sections of bounded normal operators, as exemplified by the Laurent operator studied in Section 1.1.4.

1.1.3 Continuous Spectra

Operators in applications often have a continuous spectral component, which prevents them from being diagonalised solely through eigenvalues and eigenvectors. For example, the operator $-\mathrm{d}^2/\mathrm{d}x^2$ on $\mathcal{H} = L^2(\mathbb{R})$ has a continuum of non-normalisable generalised eigenfunctions $e^{\pm \mathrm{i}kx} \notin \mathcal{H}$. The associated spectral parameters are $k^2 \in [0, +\infty)$, corresponding to a continuous spectrum. However, one can still obtain a generalisation of 'diagonalisation' by expanding functions in terms of $e^{\pm \mathrm{i}kx}$ using the Fourier transform.

We add a potential to this example and consider the following Schrödinger operator:

$$[\mathcal{L}u](x) = -\frac{\mathrm{d}^2 u}{\mathrm{d}x^2}(x) + V(x)u(x), \quad V(x) = -\frac{3}{2(1+x^2)}. \tag{1.9}$$

The specific potential V has the effect that, in addition to the continuous spectrum $[0, +\infty)$, $\mathcal{L}$ has negative eigenvalues. The corresponding eigenfunctions are termed 'bound states' in quantum mechanics. The decay of this potential V implies that $\mathcal{L}$ has no positive eigenvalues [1111, Theorem XIII.58], but infinitely many negative eigenvalues [292].

To discretise $\mathcal{L}$, we use the Malmquist–Takenaka functions (with scaling $L = 1/10$),

$$\phi_m(x) = \sqrt{L/\pi}(1 + \mathrm{i}Lx)^m (1 - \mathrm{i}Lx)^{-(m+1)}, \quad m \in \mathbb{Z}, \quad L = 1/10. \tag{1.10}$$

These functions form an orthonormal basis of $L^2(\mathbb{R})$, and their approximation properties are discussed in [227, 1395]. There are several attractive features of this basis. Expansion coefficients can be efficiently computed using the fast Fourier transform [361], and differentiation

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

and multiplication correspond to sparse and approximately sparse matrices, respectively. We use the finite section method with subspace $\text{span}\{\phi_m\}_{m=-n}^{n}$, corresponding to a discretisation size of $N = 2n + 1$, and denote the eigenvalues and a corresponding orthonormal set of eigenfunctions of the discretised operator by $\{\lambda_j\}_{j=1}^{N}$ and $\{u_j\}_{j=1}^{N}$, respectively.

Given a function $v(x) = x^2\sqrt{9/\pi}/(1+x^6)$, which has been normalised to have $\|v\|_{L^2(\mathbb{R})} = 1$, we consider the so-called *spectral measure*:

$$\mu_{v_N} = \frac{1}{\sum_{j=1}^{N} |u_j^* v|^2} \sum_{j=1}^{N} |u_j^* v|^2 \delta_{\lambda_j} = \frac{1}{\sum_{j=1}^{N} |u_j^* v_N|^2} \sum_{j=1}^{N} |u_j^* v_N|^2 \delta_{\lambda_j}.$$

Here, v_N is the orthogonal projection of v onto $\text{span}\{\phi_m\}_{m=-n}^{n}$ and δ_z denotes a Dirac delta distribution supported at $z \in \mathbb{C}$. The measure μ_{v_N} indicates how the total energy content of v_N (its squared norm) is distributed across the eigenspaces of the discretised operator. The normalisation factor in front of the sum ensures that μ_{v_N} is a probability measure ($\|v_N\|_{L^2(\mathbb{R})} \neq 1$). The spectral measures μ_{v_N} are shown in Figure 1.7 for $N = 101, 1001$, and 10001. Since each μ_{v_N} is a sum of Dirac delta distributions, we represent the positions of these distributions and their corresponding coefficients $|u_j^* v|^2/\|v_N\|_{L^2(\mathbb{R})}^2$ as vertical lines.

Two features stand out as N increases:

(i) Part of μ_{v_N} approximates the discrete spectrum of $\mathcal{L}$ (the negative eigenvalues), along with the inner products between v and the bound states. The strengths of these Dirac delta distributions converge to non-negative, finite values;

(ii) The strengths of the Dirac delta distributions in the continuous spectrum $[0, +\infty)$ decrease, while their numbers increase, effectively filling out the continuous spectrum.

Do the measures μ_{v_N} converge in some manner as $N \to \infty$? If so, how can we compute this limit, and how is it related to the infinite-dimensional problem?

Glimpse of remedy: These questions are addressed in Section 4.2. For now, we plot the cumulative distribution functions of the spectral measures in Figure 1.8. The vertical jumps correspond to eigenvalues. The figure suggests that there is indeed a limit as $N \to \infty$.

The spectral theorem (Theorem 4.1.2) states that a normal operator can be realised as multiplication by a spectral parameter via an integral with respect to a *projection-valued measure*, extending the notion of spectral measures to infinite dimensions. This may be viewed as a generalisation of diagonalisation. A corresponding probability measure μ_v exists, which in this example has a density function on $[0, \infty)$ [7, Theorem 6.1] and a discrete part on $(-\infty, 0)$. Chapters 4 and 5 cover computing these measures, and thus achieving generalisations of diagonalisation in the presence of continuous spectra. This includes methods that integrate μ_v against a smoothing kernel via computing the action of the resolvent $(A - zI)^{-1}$ on appropriate vectors, which we use to compute the density in the bottom panel of Figure 1.7. Computational methods for determining spectral type (e.g., continuous spectra versus eigenvalues) are discussed in Chapter 6, and computational problems involving the essential spectrum are the topic of Chapter 8.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

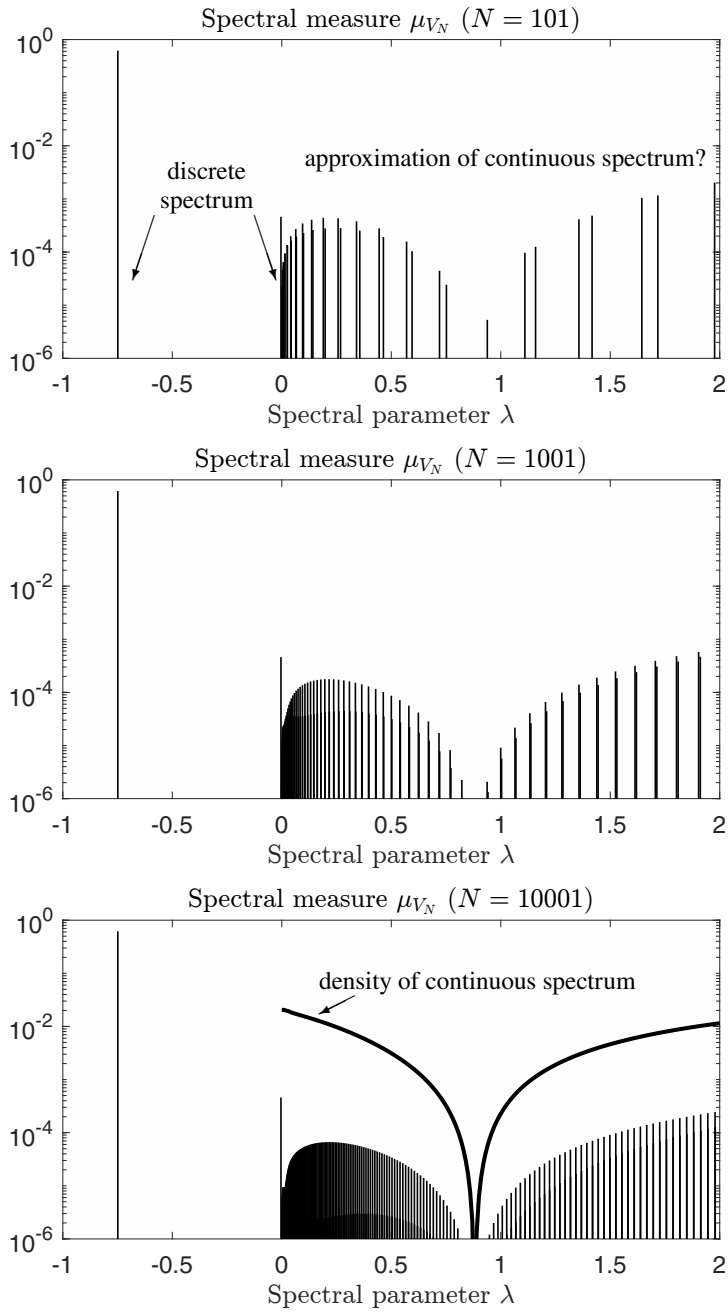


Figure 1.7 The discretised spectral measures μ_{v_N} for different discretisation sizes N used for the Schrödinger operator $\mathcal{L}$ in (1.9). Each vertical line corresponds to a Dirac delta distribution at the corresponding λ, with the strength given by the height of the line. The top plot labels the discrete spectrum and the region of continuous spectrum. The bottom plot overlays the density of the spectral measure of $\mathcal{L}$ with respect to v, computed using the methods of Chapter 4.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

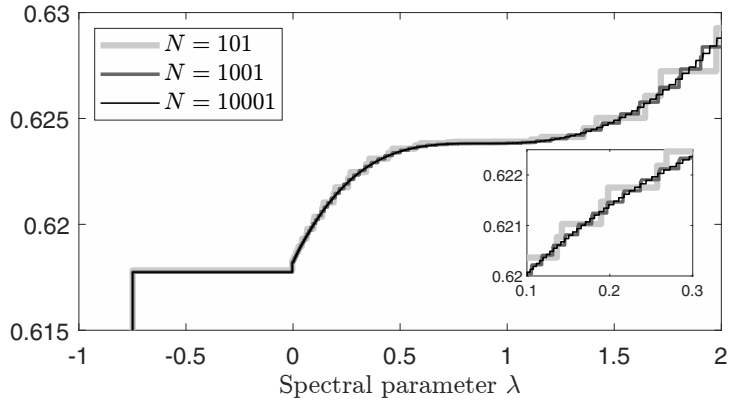


Figure 1.8 Cumulative distribution functions of the measures shown in Figure 1.7.

1.1.4 Ill-Conditioning

Discretisation can introduce significant ill-conditioning into the spectrum of discretised operators, even when the spectrum of the underlying infinite-dimensional operator is well-conditioned. This effect arises, for example, in spectral method discretisations of eigenvalue problems (even for self-adjoint operators) [727, Figure 1.1], and is common in applications, such as the analysis of hydrodynamic (in)stability [450]. We have already encountered such behaviour in the case of the convection–diffusion operator in Figure 1.5.

We now examine a bounded operator with similar behaviour and investigate the effect of discretisations on computed eigenvalues. The following beautiful example is from the front cover of the book by Böttcher and Silbermann [207]. A *Laurent operator* on $\ell^2(\mathbb{Z})$ is a bi-infinite matrix with constant diagonals. The 'whale' Laurent operator is:

$$L = \begin{pmatrix} \ddots & \ddots & \ddots & \ddots & \ddots & & \\ \ddots & 0 & 1 & \mathrm{i} & -3-2\mathrm{i} & -1 & \\ \ddots & 10 & 0 & 1 & \mathrm{i} & -3-2\mathrm{i} & \ddots \\ \ddots & 3+\mathrm{i} & 10 & 0 & 1 & \mathrm{i} & \ddots \\ \ddots & 4 & 3+\mathrm{i} & 10 & 0 & 1 & \ddots \\ & \mathrm{i} & 4 & 3+\mathrm{i} & 10 & 0 & \ddots \\ & & \ddots & \ddots & \ddots & \ddots & \ddots \end{pmatrix}, \tag{1.11}$$

where the main diagonal corresponds to the sequence of 0s, and diagonals not explicitly shown are zero. This matrix is banded with bandwidth 4. Define the complex function

$$a(z) = -z^{-4} - (3+2\mathrm{i})z^{-3} + \mathrm{i}z^{-2} + z^{-1} + 10z + (3+\mathrm{i})z^2 + 4z^3 + \mathrm{i}z^4,$$

which is termed the symbol of L. The operator L is normal, and $\mathrm{Sp}(L) = \{a(z) : |z| = 1\}$, which can be derived using Fourier series (Exercise 1.3). By employing Fourier series as

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

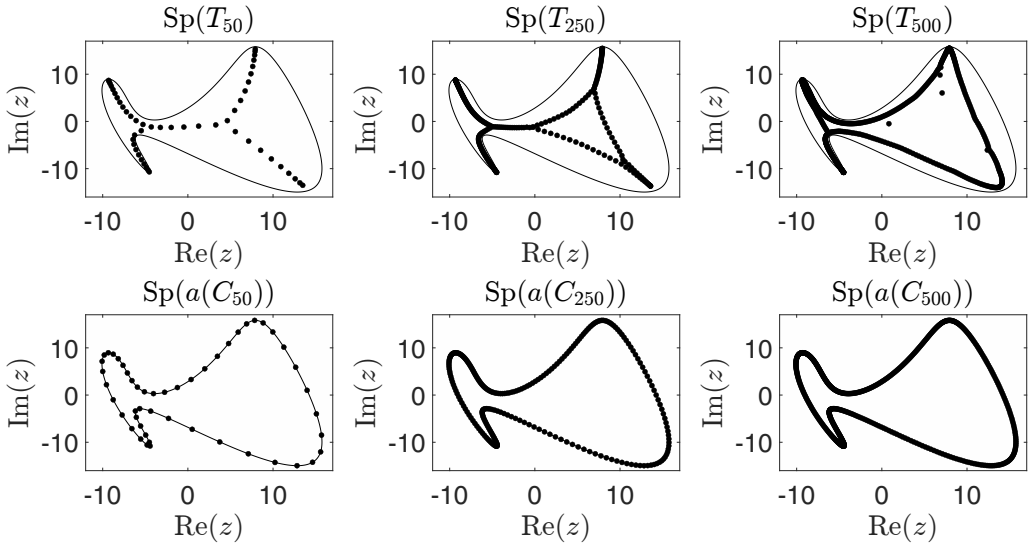


Figure 1.9 Top row: 'Eigenvalues' of T_n computed using standard double precision for various n. As $n \to \infty$, the true eigenvalues of T_n converge to a connected finite union of analytic arcs, visible in the leftmost plot. In practice, the computed eigenvalues cluster according to pseudospectra (Definition 1.1.6). Bottom row: Eigenvalues of $a(C_n)$ computed using standard double precision for various n. Here, the discretised eigenvalue problem is well-conditioned and the eigenvalues converge to $\mathrm{Sp}(L)$.

opposed to the Fourier transform, this example can be seen as a discrete analogue of the differential operator discussed in Section 1.1.2.

The finite sections of L are *Toeplitz matrices* $T_n \in \mathbb{C}^{n\times n}$, corresponding to projections onto $\mathrm{span}\{e_{j+1}, \ldots, e_{j+n}\}$ for every $j \in \mathbb{Z}$ (each j gives the same matrix). The top row of Figure 1.9 shows eigenvalues computed using standard double precision for various n. One can show that the set of exact eigenvalues for this example has a limit $\lim_{n\to\infty} \mathrm{Sp}(T_n)$, which is a connected, finite union of analytic arcs in the interior of the curve $\{a(z)\colon |z| = 1\}$ [1166]. In practice, the ill-conditioning of $\mathrm{Sp}(T_n)$ means that, in finite precision arithmetic, the computed eigenvalues of T_n deviate from this limit curve for large n.

A structured discretisation trick: To better approximate $\mathrm{Sp}(L)$, we consider the shift operator S_l on $\ell^2(\mathbb{Z})$, where $S_l e_j = e_{j+1}$, so that $L = a(S_l)$. The shift operator can be approximated by truncation with periodic boundary conditions. The resulting unitary circulant matrix

$$C_n = \begin{pmatrix} 0 & & & 1 \\ 1 & 0 & & \\ & \ddots & \ddots & \\ & & 1 & 0 \end{pmatrix} \in \mathbb{C}^{n\times n}$$

has eigenvalues $\{\exp(2\pi \mathrm{i} j/n)\}_{j=1}^n$, which approximate $\mathrm{Sp}(S_l) = \mathbb{T}$ as $n \to \infty$. We consider replacing T_n by its circulant counterpart $a(C_n)$. The bottom row of Figure 1.9 shows eigenvalues of $a(C_n)$ computed using standard double precision for various n. Now, the discretisation leads to a well-conditioned eigenvalue problem, and $\mathrm{Sp}(L)$ is approximated as $n \to \infty$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

This approximation has another interpretation. Namely, the circulant matrix C_n belongs to the set of unitary matrices that minimise the Frobenius norm distance to the $n \times n$ truncation of the shift operator S_l. This type of structured discretisation will be explored further in Section 5.2.3, in relation to computing spectral measures of unitary operators.

Pseudospectra

To help understand potential ill-conditioning of spectra, we turn to pseudospectra. For $\epsilon > 0$, the ϵ-pseudospectrum of a finite matrix $B \in \mathbb{C}^{n\times n}$ is the set

$$\mathrm{Sp}_\epsilon(B) = \left\{z \in \mathbb{C}\colon \|(B - zI)^{-1}\|^{-1} \leq \epsilon\right\} = \mathrm{Cl}\left(\left\{z \in \mathbb{C}\colon \|(B - zI)^{-1}\|^{-1} < \epsilon\right\}\right). \quad (1.12)$$

We adopt the convention that $\|(B - zI)^{-1}\|^{-1} = 0$ if $z \in \mathrm{Sp}(B)$. Some authors instead define the ϵ-pseudospectrum as $\{z \in \mathbb{C}\colon \|(B - zI)^{-1}\|^{-1} < \epsilon\}$, but we use '$\leq \epsilon$' to ensure that $\mathrm{Sp}_\epsilon(B)$ is a closed subset of $\mathbb{C}$.

An equivalent characterisation is given by [1336, Theorem 2.1]:

$$\mathrm{Sp}_\epsilon(B) = \bigcup_{E\in\mathbb{C}^{n\times n}, \|E\|\leq\epsilon} \mathrm{Sp}(B + E). \quad (1.13)$$

Thus, the ϵ-pseudospectrum of B is the region in the complex plane describing how far an ϵ-sized perturbation of the matrix B can move its eigenvalues.

In infinite dimensions, pseudospectra must be defined with greater care since the resolvent norm $\|(A - zI)^{-1}\|$ of an unbounded operator on a Hilbert space can be constant on open sets (Exercise 1.5). This means that the second equality in (1.12) need not hold.

Definition 1.1.6 (Pseudospectra) Let A be a closed and densely defined operator acting on a separable Hilbert space $\mathcal{H}$. For each $\epsilon > 0$, the ϵ-pseudospectrum of A is the set

$$\mathrm{Sp}_\epsilon(A) = \mathrm{Cl}\left(\left\{z \in \mathbb{C}\colon \left\|(A - zI)^{-1}\right\|^{-1} < \epsilon\right\}\right),$$

i.e., the closure of the set of points with a resolvent norm greater than $1/\epsilon$.

For every $\epsilon > 0$, the pseudospectrum $\mathrm{Sp}_\epsilon(A)$ of a closed and densely defined operator is non-empty (Exercise 9.8). An analogue of (1.13) also holds (Exercise 1.4):

$$\mathrm{Sp}_\epsilon(A) = \mathrm{Cl}\left(\bigcup_{E\in\mathscr{B}(\mathcal{H})\colon \|E\|<\epsilon} \mathrm{Sp}(A + E)\right). \quad (1.14)$$

If A is a normal operator, the spectrum is well-conditioned and (Exercise 1.8)

$$\mathrm{Sp}_\epsilon(A) = \mathrm{Sp}(A) + B_\epsilon(0),$$

where $B_\epsilon(0)$ is the closed ball of radius ϵ centred at 0. However, in general, $\mathrm{Sp}_\epsilon(A)$ can be much larger than $\mathrm{Sp}(A) + B_\epsilon(0)$, which corresponds to an ill-conditioning of the spectrum.

Throughout this book, we visualise pseudospectra by plotting contours of the function $z \mapsto \|(A - zI)^{-1}\|^{-1}$ on a logarithmic scale. For example, returning to the whale Laurent operator, the left panel of Figure 1.10 shows $\mathrm{Sp}_\epsilon(T_{50})$, whereas the right shows $\mathrm{Sp}_\epsilon(L)$. The finite section method leads to severe ill-conditioning of the discretised problem, which grows

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

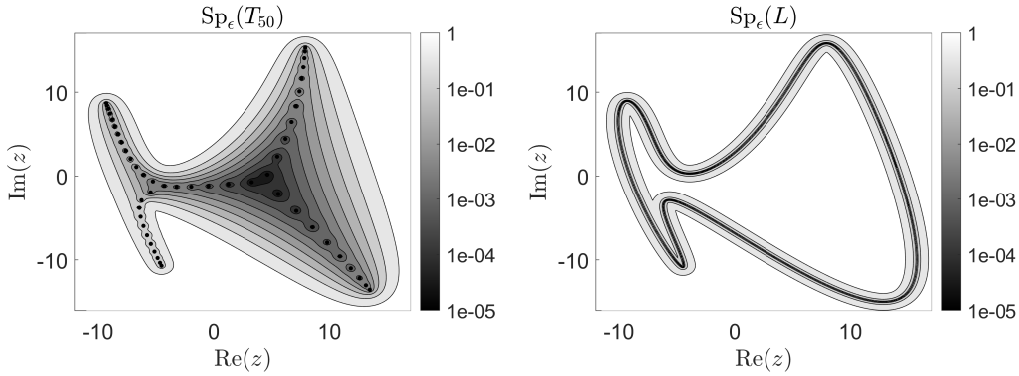


Figure 1.10 Left: Pseudospectra of 50×50 truncation, T_{50}, of the operator L in (1.11). Eigenvalues of T_{50} are shown as dots. Right: Pseudospectra of L.

exponentially with n [1336, Theorem 7.2]. The eigenvalues of T_n computed in Figure 1.9 tend to accumulate according to boundaries of the pseudospectra of T_n.

Analysing the sensitivity of spectra under perturbations is far from the only application of pseudospectra. For example, as we will see in Section 3.2, concepts related to pseudospectra enable the computation of spectra. Pseudospectra are also used to analyse the transient behaviour of dynamical systems [1337]. In such cases, if the underlying operator is non-normal, spectra can be highly misleading, as discussed in Sections 9.1 and 9.2.

1.1.5 Other Issues

The above difficulties are by no means the only challenges associated with discretisation. Even when a method converges as a discretisation parameter increases, typically only a subset of the computed eigenvalues are reliable for a fixed discretisation size [1396, 1448]. Furthermore, attempts to compare eigenvalues across different discretisation sizes can lead to erroneous conclusions [228]. Consequently, it becomes imperative to determine the extent to which we can trust the output for a finite discretisation size and whether *verification*—computing error bounds—is feasible.

Beyond $\mathrm{Sp}(A)$, other infinite-dimensional spectral properties are frequently of interest, such as the spectral type (Chapter 6); fractal dimensions of $\mathrm{Sp}(A)$ (Chapter 7); or geometric quantities such as radii, abscissas, and gaps (Chapter 9). In some problems and applications, spectral features like pseudospectra (Definition 1.1.6) may be more informative than the spectrum itself.

1.2 The Classical Computational Spectral Problem

When computing the spectral properties of operators in infinite dimensions, we inevitably have to discretise or truncate at some point. Hence, the above pitfalls are alarming. How can we address these challenges? Here, we discuss their resolution in the context of the *classical computational spectral problem*.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

1.2.1 Operators as Matrices

It is often helpful to view operators as infinite matrices. A Hilbert space $\mathcal{H}$ is separable if and only if it admits a countable orthonormal basis. Suppose that A is a closed, densely defined operator on such a space. Let $\{e_n\}_{n\in\mathbb{N}}$ be an orthonormal basis of $\mathcal{H}$ such that $\operatorname{span}\{e_n : n \in \mathbb{N}\}$ forms a core of A (which always exists). As detailed in Appendix A.3, a core C of A is a subset of $\mathcal{D}(A)$ such that the closure of $A\restriction_C$ is A. We can represent A as an infinite matrix with respect to this basis:

$$A = \begin{pmatrix} A_{11} & A_{12} & A_{13} & \dots \\ A_{21} & A_{22} & A_{23} & \dots \\ A_{31} & A_{32} & A_{33} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \quad A_{ij} = \langle Ae_j, e_i\rangle. \tag{1.15}$$

In a slight but standard abuse of notation, we use the same symbol A for both the operator and its matrix representation. If A is bounded, then

$$Av = \sum_{i=1}^{\infty}\left(\sum_{j=1}^{\infty} A_{ij}v_j\right)e_i \quad \forall v = \sum_{n=1}^{\infty} v_n e_n \in \mathcal{H}.$$

More generally, by the definition of a core, a given infinite matrix $\{A_{ij}\}$ can represent at most one closed and densely defined operator A for which $\operatorname{span}\{e_n : n \in \mathbb{N}\}$ is a core.

By the usual identification $e_i = \{\delta_{ij}\}_{j=1}^{\infty}$, where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise, $\mathcal{H}$ is isomorphic to the space of square summable sequences:

$$\ell^2(\mathbb{N}) = \left\{x = \{x_j\}_{j=1}^{\infty} : \sum_{j=1}^{\infty}|x_j|^2 < \infty\right\}, \text{ with inner product } \langle x, y\rangle = \sum_{j=1}^{\infty} x_j\overline{y_j},\ x, y \in \ell^2(\mathbb{N}).$$

One may think of $\ell^2(\mathbb{N})$ as 'coefficient space' with respect to a countable orthonormal basis, and the infinite matrix representing A as acting on $\ell^2(\mathbb{N})$. Therefore, much of our effort will go into understanding operators on $\ell^2(\mathbb{N})$ since results applicable here can be transferred to other separable Hilbert spaces by choosing an appropriate basis.

To reduce notational burden, we assume that $\operatorname{span}\{e_n : n \in \mathbb{N}\}$ forms a core of A and A^*. Hence, A^* is represented by the infinite matrix with (i, j) entry $\langle A^*e_j, e_i\rangle = \langle e_j, Ae_i\rangle = \overline{A_{ji}}$. However, our discussion, and all the results of this book, can be easily extended to using different bases to represent A and A^*. We thus consider the following class of operators.

Definition 1.2.1 (Operators on $\ell^2(\mathbb{N})$ defined through matrices) Let $\Omega_{\ell^2(\mathbb{N})}$ denote the class of closed and densely defined operators A on $\ell^2(\mathbb{N})$ for which the subspace $\operatorname{span}\{e_n : n \in \mathbb{N}\}$ forms a core of both A and its adjoint A^*.

1.2.2 The Canonical Example

Given $A \in \Omega_{\ell^2(\mathbb{N})}$, we aim to compute $\operatorname{Sp}(A)$ from its matrix elements $\langle Ae_j, e_i\rangle$. Owing to the infinite amount of information encoded in the matrix of A, defining a 'computational spectral problem' and an 'algorithm' in infinite dimensions requires caution. Precise definitions are provided in Section 2.1. Suppose, however, that an algorithm uses a finite number of matrix

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

elements, which it can select adaptively, and yields an output $\Gamma_n(A) \subset \mathbb{C}$.[3] We denote the algorithm with a subscript n, indicating a sequence of approximations, or possibly different algorithms, that converge as $n \to \infty$. For instance, in finite dimensions, n might correspond to the number of iterations of the QR algorithm, which converges to the set of eigenvalues of a matrix under certain conditions. In infinite dimensions, n could represent some form of discretisation or truncation parameter. However, in general, the number of matrix elements selected by an algorithm can be adaptive, and the subscript n may have nothing to do with the number of elements viewed by the algorithm. It simply indexes a sequence of algorithms.

Key question: Can we design algorithms $\{\Gamma_n\}_{n=1}^{\infty}$ so that $\Gamma_n(A)$ converges to $\mathrm{Sp}(A)$?

The correct metric spaces

To avoid the problems of spectral pollution and spectral invisibility discussed in Section 1.1, any reliable algorithm for approximating $\mathrm{Sp}(A)$ should produce approximations that converge uniformly on every compact subset of $\mathbb{C}$. This mode of convergence ensures that no spurious limit points arise outside $\mathrm{Sp}(A)$ and that all genuine spectral points are captured. We refer to this as *locally uniform convergence* (a term also used for the convergence of functions). To formalise this notion for spectral sets, we work within two distinct metric spaces.

The Hausdorff metric [116, page 85] is suitable when $\mathrm{Sp}(A)$ is bounded and non-empty, and corresponds to uniform convergence of non-empty compact sets.

Definition 1.2.2 (Hausdorff metric) Let $\mathcal{M}$ be the set of non-empty compact subsets of a metric space $(\mathcal{M}', d')$. The Hausdorff metric is

$$d_{\mathrm{H}}(X,Y) = \max\left\{\sup_{x\in X}\inf_{y\in Y} d'(x,y), \sup_{y\in Y}\inf_{x\in X} d'(x,y)\right\}, \quad X,Y \in \mathcal{M}. \tag{1.16}$$

When $\mathcal{M}' = \mathbb{C}$, we take d' to be the standard Euclidean metric and denote the resulting metric space by $(\mathcal{M}_{\mathrm{H}}, d_{\mathrm{H}})$.

The Attouch–Wets metric [116, page 79] is suitable when $\mathrm{Sp}(A)$ is unbounded.

Definition 1.2.3 (Attouch–Wets metric) Let $\mathcal{M}$ be the set of non-empty closed subsets of a metric space $(\mathcal{M}', d')$. The Attouch–Wets metric is

$$d_{\mathrm{AW}}(X,Y) = \sum_{n=1}^{\infty} 2^{-n} \min\left\{1, \sup_{d'(x_0,x)\leq n} |d'(x,X) - d'(x,Y)|\right\}, \quad X,Y \in \mathcal{M}. \tag{1.17}$$

Here, $x_0 \in \mathcal{M}'$ is some fixed point and $d'(x,Y) = \inf_{y\in Y} d'(x,y)$. When $\mathcal{M}' = \mathbb{C}$ with its usual metric, we take $x_0 = 0$ and denote the resulting metric space by $(\mathcal{M}_{\mathrm{AW}}, d_{\mathrm{AW}})$.

The Attouch–Wets metric generalises the Hausdorff metric and induces the same notion of convergence when restricted to non-empty, compact, and uniformly bounded sets (Exercise 1.10). The following lemma, proven in Exercise 1.10, shows why d_{AW} is the correct

[3] The term 'lazy evaluation' is often used to describe the approach of computing with infinite data structures by deferring the use of information until necessary.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

metric when striving to avoid spectral pollution (Definition 1.1.1) and spectral invisibility (Definition 1.1.2): it captures the notion of locally uniform convergence.

Lemma 1.2.4 (Characterisation of Attouch–Wets metric) *For non-empty closed sets* $C, C_n \subset \mathbb{C}$, $\lim_{n\to\infty} d_{\mathrm{AW}}(C_n, C) = 0$ *if and only if* $\lim_{n\to\infty} d_K(C_n, C) = 0$ *for every compact* $K \subset \mathbb{C}$. *Here,*

$$d_K(C_1, C_2) = \max\left\{ \sup_{a\in C_1\cap K} \operatorname{dist}(a, C_2), \sup_{b\in C_2\cap K} \operatorname{dist}(b, C_1) \right\}, \tag{1.18}$$

where, by convention, the supremum over the empty set is zero. Convergence occurs if and only if for each $\delta > 0$ *and compact* K, *there exists an* N *such that, for* $n > N$, $C_n \cap K \subset C + B_\delta(0)$ *and* $C \cap K \subset C_n + B_\delta(0)$. *Moreover, it suffices to consider sets* K *of the form* $B_m(0)$ *for large* $m \in \mathbb{N}$.

We can extend Lemma 1.2.4 to deal with the empty set, an essential consideration as the spectrum of unbounded operators may be empty. Similarly, other spectral sets of bounded operators, such as discrete spectra, can also be empty. We say that $F_n \to \emptyset$ as $n \to \infty$ if and only if $F_n \cap B_m(0) = \emptyset$ eventually, for each $m \in \mathbb{N}$. With an abuse of terminology, we continue to refer to the Attouch–Wets metric space $(\mathcal{M}_{\mathrm{AW}}, d_{\mathrm{AW}})$ as accommodating the empty set.

The classical computational spectral problem

The classical computational spectral problem is to develop algorithms that compute the spectra of operators in $\Omega_{\ell^2(\mathbb{N})}$ with convergence in the Attouch–Wets and Hausdorff metrics. The problem dates back to Szegő's work on finite section approximations [1282] and to Schwinger's foundational work in quantum mechanics for Schrödinger operators [1174, 1175]. It has been a challenging area, one of interest for mathematicians for decades; see the notes section for additional historical context. Arveson, pivotal in merging spectral computations with functional analysis [50–52], highlighted in 1994 the (then) unresolved challenge of computing spectra for general self-adjoint operators [53]:

"Most operators that arise in practice are not presented in a representation in which they are diagonalized, and it is often very hard to locate even a single point in the spectrum ... This raises the question of how to implement the methods of finite dimensional numerical linear algebra to compute the spectra of infinite dimensional operators. Unfortunately, there is a dearth of literature on this basic problem and, so far as we have been able to tell, there are no proven techniques."

A compelling reason underlies Arveson's lament:

Health warning: For many infinite-dimensional spectral problems, there does not exist a convergent sequence of algorithms $\{\Gamma_n\}_{n=1}^{\infty}$, regardless of numerical precision, computational power, or computational model.

However, all hope is not lost! We can address this issue by broadening our computational

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

paradigm, paving the way for an intriguing classification theory. The next two chapters explore this problem in more detail.

1.2.3 A 'Three-Limit' Algorithm for Spectra

We can address the classical computational spectral problem by relaxing our demand for a sequence of algorithms Γ_n that converge in a single limit $n \to \infty$. Instead, we now outline an algorithm, Γ_{n_3,n_2,n_1}, that converges to the spectrum of general operators through three *successive limits*:

$$\lim_{n_3\to\infty} \lim_{n_2\to\infty} \lim_{n_1\to\infty} \Gamma_{n_3,n_2,n_1}(A) = \mathrm{Sp}(A) \quad \forall A \in \Omega_{\ell^2(\mathbb{N})}. \tag{1.19}$$

Here, convergence is in the Attouch–Wets metric defined in (1.17). Hence, our discussion surrounding Lemma 1.2.4 shows that the convergence in (1.19) avoids the pitfalls of spectral pollution (Definition 1.1.1) and spectral invisibility (Definition 1.1.2).

We first describe the ingredients of the algorithm for a general separable Hilbert space $\mathcal{H}$. For a closed operator A with domain $\mathcal{D}(A) \subset \mathcal{H}$, its injection modulus is defined as

$$\sigma_{\inf}(A) = \inf\{\|Ax\| \colon x \in \mathcal{D}(A), \|x\| = 1\} = \sup\{\tau \geq 0 \colon \|Ax\| \geq \tau\|x\| \, \forall x \in \mathcal{D}(A)\}. \tag{1.20}$$

The injection modulus is a quantitative measure of 'how injective' an operator is, and generalises the notion of the smallest singular value of a finite matrix. We prefer to use the term 'injection modulus' over 'smallest singular value' not only for its conciseness, but also because it subtly underscores the fact that we often deal with operators in infinite-dimensional spaces. Such operators can exhibit behaviours quite distinct from those of finite matrices. For example, even if A is bounded, there may not be a unit-norm vector x with $\|Ax\| = \sigma_{\inf}(A)$. It is immediate from the triangle inequality that the injection modulus is 1-Lipschitz continuous with respect to bounded perturbations in the operator norm.

The following function of $z \in \mathbb{C}$ is a fundamental tool of this book:

$$\gamma(z, A) = \min\left\{\sigma_{\inf}(A - zI), \sigma_{\inf}(A^* - \overline{z}I)\right\}. \tag{1.21}$$

We can characterise the resolvent norm $\|(A - zI)^{-1}\|$ in terms of γ. This characterisation is instrumental because we can approximate injection moduli through the smallest singular values of appropriate finite matrices.

Lemma 1.2.5 (Injection moduli and resolvent norms) *For each closed, densely defined operator A on $\mathcal{H}$, $\gamma(z, A) = \|(A - zI)^{-1}\|^{-1}$, where we adopt the convention that $\|(A - zI)^{-1}\|^{-1} = 0$ if $z \in \mathrm{Sp}(A)$. Moreover, if $z \notin \mathrm{Sp}(A)$, then $\sigma_{\inf}(A - zI) = \sigma_{\inf}(A^* - \overline{z}I)$.*

Proof We use some basic properties of operators that are listed in Appendix A.3. We first deal with the case $z \notin \mathrm{Sp}(A)$, where we prove that $\sigma_{\inf}(A - zI) = \sigma_{\inf}(A^* - \overline{z}I) = \|(A - zI)^{-1}\|^{-1}$. We show this for $\sigma_{\inf}(A - zI)$, and the other case is similar using the fact that $\|(A - zI)^{-1}\| = \|(A^* - \overline{z}I)^{-1}\|$. Let $x \in \mathcal{D}(A)$ with $\|x\| = 1$, then

$$1 = \|x\| = \left\|(A - zI)^{-1}(A - zI)x\right\| \leq \left\|(A - zI)^{-1}\right\| \|(A - zI)x\|.$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Upon taking the infimum over x, $\sigma_{\inf}(A - zI) \geq \|(A - zI)^{-1}\|^{-1}$. Conversely, let $x_n \in \mathcal{H}$ such that $\|x_n\| = 1$ and $\lim_{n\to\infty} \|(A - zI)^{-1}x_n\| = \|(A - zI)^{-1}\|$. Then

$$1 = \|x_n\| = \left\|(A - zI)(A - zI)^{-1}x_n\right\| \geq \sigma_{\inf}(A - zI)\left\|(A - zI)^{-1}x_n\right\|.$$

Letting $n \to \infty$, we obtain $\sigma_{\inf}(A - zI) \leq \|(A - zI)^{-1}\|^{-1}$.

Now suppose that $z \in \mathrm{Sp}(A)$. We prove that $\gamma(z, A) = 0$ by assuming that $\sigma_{\inf}(A - zI) > 0$ and $\sigma_{\inf}(A^* - \overline{z}I) > 0$ and deriving a contradiction. Since $\sigma_{\inf}(A^* - \overline{z}I) > 0$, it follows that $\ker(A^* - \overline{z}I) = \{0\}$. Since $\mathrm{ran}(A - zI)^{\perp} = \ker(A^* - \overline{z}I)$, the range of $(A - zI)$ is dense in $\mathcal{H}$. Furthermore, $(A - zI)$ is injective on its domain. Therefore, we can define $(A - zI)^{-1}$ on the dense subspace $\mathrm{ran}(A - zI)$. This operator is bounded since $\sigma_{\inf}(A - zI) > 0$, allowing us to extend it to a bounded operator on the whole of $\mathcal{H}$. The closedness of A implies that $(A - zI)(A - zI)^{-1} = I$. Clearly, $(A - zI)^{-1}(A - zI)x = x$ for all $x \in \mathcal{D}(A)$. Hence, $z \notin \mathrm{Sp}(A)$, a contradiction. □

At this point, the reader may wonder why we need both $\sigma_{\inf}(A - zI)$ and $\sigma_{\inf}(A^* - \overline{z}I)$ in the definition of γ. In a finite-dimensional Hilbert space, we always have $\sigma_{\inf}(A - zI) = \sigma_{\inf}(A^* - \overline{z}I)$. However, in infinite dimensions, it is possible to have $z \in \mathrm{Sp}(A)$ and $\sigma_{\inf}(A - zI) \neq \sigma_{\inf}(A^* - \overline{z}I)$ (Exercise 1.12), though this cannot happen for normal operators.

Lemma 1.2.5 means that we can characterise spectra and pseudospectra as

$$\mathrm{Sp}(A) = \{z \in \mathbb{C} \colon \gamma(z, A) = 0\}, \quad \mathrm{Sp}_{\epsilon}(A) = \mathrm{Cl}(\{z \in \mathbb{C} \colon \gamma(z, A) < \epsilon\}). \tag{1.22}$$

Hence, we should expect that γ gives a route to computing $\mathrm{Sp}(A)$.

To compute γ, we now suppose that $A \in \Omega_{\ell^2(\mathbb{N})}$. For $n_1, n_2 \in \mathbb{N}$, let

$$\gamma_{n_2,n_1}(z, A) = \min\left\{\sigma_{\inf}(\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*), \sigma_{\inf}(\mathcal{P}_{n_1}(A^* - \overline{z}I)\mathcal{P}_{n_2}^*)\right\}, \tag{1.23}$$

$$\gamma_{n_2}(z, A) = \min\left\{\sigma_{\inf}((A - zI)\mathcal{P}_{n_2}^*), \sigma_{\inf}((A^* - \overline{z}I)\mathcal{P}_{n_2}^*)\right\}, \tag{1.24}$$

where $\mathcal{P}_n$ denotes the orthogonal projection onto $\mathrm{span}\{e_1, \ldots, e_n\}$. We may think of the operators in the definition of $\gamma_{n_2}(z, A)$ as infinitely tall matrices with a finite number of columns. Computing $\gamma_{n_2,n_1}(z, A)$ consists of computing the smallest singular values of the two *finite rectangular* matrices $\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*$ and $\mathcal{P}_{n_1}(A^* - \overline{z}I)\mathcal{P}_{n_2}^*$.

The functions γ_{n_2,n_1} have the following beautiful and useful monotonicity properties.

Lemma 1.2.6 (Monotonicity of injection moduli) *For every $A \in \Omega_{\ell^2(\mathbb{N})}$ and $z \in \mathbb{C}$,*

(i) $\lim_{n_1\to\infty} \gamma_{n_2,n_1}(z, A) = \gamma_{n_2}(z, A)$, *where the convergence is monotonic from below, and uniform on compact subsets of $\mathbb{C}$;*

(ii) $\lim_{n_2\to\infty} \gamma_{n_2}(z, A) = \gamma(z, A)$, *where the convergence is monotonic from above, and uniform on compact subsets of $\mathbb{C}$.*

Proof Consider the functions $z \mapsto \sigma_{\inf}(\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*)$. As n_1 increases, the range of $\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*$ increases and hence its injection modulus cannot decrease. We also have $\sigma_{\inf}(\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*) \leq \sigma_{\inf}((A - zI)\mathcal{P}_{n_2}^*)$. If we can prove that $\sigma_{\inf}(\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*)$ converges pointwise to $\sigma_{\inf}((A - zI)\mathcal{P}_{n_2}^*)$, then since all relevant functions are continuous, Dini's theorem applies: a monotone sequence of continuous functions on a compact space that converges pointwise to a continuous limit actually converges uniformly. Hence, in this case, $\sigma_{\inf}(\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^*)$ converges locally uniformly to $\sigma_{\inf}((A - zI)\mathcal{P}_{n_2}^*)$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

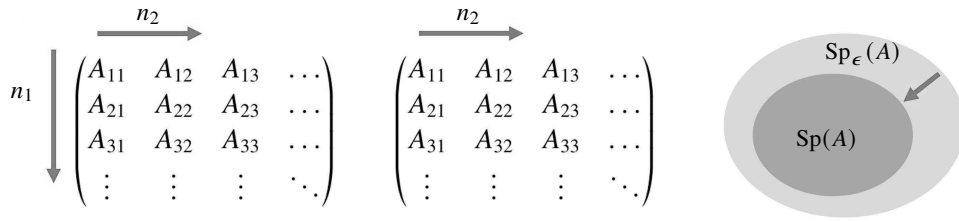


Figure 1.11 The key idea of the three-limit algorithm for computing spectra. The first parameter n_1 controls the number of rows of the truncation, and the second parameter n_2 controls the number of columns. As $n_1 \to \infty$, we move from the left panel to the middle panel, and the computed minimum singular values increase. As $n_2 \to \infty$, the computed singular values decrease. The final parameter corresponds to $\epsilon = 1/n_3$ and controls the convergence of $\mathrm{Sp}_\epsilon(A)$ to $\mathrm{Sp}(A)$, shown in the right panel.

To that end, fix $z \in \mathbb{C}$. Since $(A - zI)\mathcal{P}_{n_2}^*$ is finite-rank and $\mathcal{P}_{n_1}^*\mathcal{P}_{n_1}$ converges strongly to the identity,

$$\lim_{n_1\to\infty} \|\mathcal{P}_{n_1}^*\mathcal{P}_{n_1}(A - zI)\mathcal{P}_{n_2}^* - (A - zI)\mathcal{P}_{n_2}^*\| = 0.$$

Since $\sigma_{\inf}$ is continuous with respect to the operator norm topology when restricted to bounded operators, the convergence follows. We can argue similarly for $\sigma_{\inf}(\mathcal{P}_{n_1}(A^* - \overline{z}I)\mathcal{P}_{n_2}^*)$ and $\sigma_{\inf}((A^* - \overline{z}I)\mathcal{P}_{n_2}^*)$. Part (i) now follows.

For part (ii), we use the fact that the domains of $(A - zI)\mathcal{P}_{n_2}^*$ and $(A^* - \overline{z}I)\mathcal{P}_{n_2}^*$ are increasing in n_2, and hence $\gamma_{n_2}(z, A)$ is decreasing. Hence, analogously, it suffices to prove pointwise convergence. To that end, fix $z \in \mathbb{C}$ and $\epsilon \in (0, 1)$. We may choose $x \in \mathcal{D}(A)$ of unit norm with $\|(A - zI)x\| \leq \sigma_{\inf}(A - zI) + \epsilon$. Since $\mathrm{span}\{e_n : n \in \mathbb{N}\}$ forms a core of A, we may choose $x' \in \mathrm{span}\{e_n : n \in \mathbb{N}\}$ with $\|x'\| \in [1 - \epsilon, 1 + \epsilon]$ such that $\|(A - zI)(x - x')\| \leq \epsilon$. Since $x' \in \mathrm{span}\{e_1, \ldots, e_{n_2}\}$ for sufficiently large n_2, it follows that

$$\begin{aligned}\limsup_{n_2\to\infty} \sigma_{\inf}((A - zI)\mathcal{P}_{n_2}^*) &\leq (1 - \epsilon)^{-1}\|(A - zI)x'\| \leq (1 - \epsilon)^{-1}\|(A - zI)x\| + \epsilon(1 - \epsilon)^{-1} \\ &\leq (1 - \epsilon)^{-1}\sigma_{\inf}(A - zI) + 2\epsilon(1 - \epsilon)^{-1}.\end{aligned}$$

Since $\epsilon > 0$ was arbitrary, the required convergence holds. We can argue similarly for $\sigma_{\inf}((A^* - \overline{z}I)\mathcal{P}_{n_2}^*)$ and $\sigma_{\inf}(A^* - \overline{z}I)$. Part (ii) now follows. □

Algorithm 1.1 outlines a three-limit algorithm that achieves the convergence in (1.19). Using Lemma 1.2.6, one can show that

$$\lim_{n_2\to\infty} \lim_{n_1\to\infty} \Gamma_{n_3,n_2,n_1}(A) = \mathrm{Sp}_{\frac{1}{n_3}}(A). \tag{1.25}$$

Since $\lim_{n_3\to\infty} \mathrm{Sp}_{1/n_3}(A) = \mathrm{Sp}(A)$ (Exercise 1.11), the convergence in (1.19) now follows. The idea is sketched in Figure 1.11:

- The first limit increases the number of rows in the truncation;
- The second limit increases the number of columns (and the grid resolution);
- The third and final limit decreases pseudospectra to the spectrum.

The reader may be surprised that we compute spectra via injection moduli (singular values) rather than eigenvalues. This has two key advantages. First, injection moduli are directly

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Algorithm 1.1 A three-limit algorithm for computing the spectrum without spectral pollution or invisibility. Exercise 1.15 modifies this sketch. Different grid choices, such as dyadic rationals or adaptive grids, may be more efficient.

Input: Operator $A \in \Omega_{\ell^2(\mathbb{N})}$, $n_1, n_2, n_3 \in \mathbb{N}$.

1: Set $\mathrm{Grid}(n_2) = \left\{ z \in \left(\frac{1}{n_2}\mathbb{Z} + \frac{\mathrm{i}}{n_2}\mathbb{Z} \right) : |z| \le n_2 \right\}$.

2: For each $z \in \mathrm{Grid}(n_2)$, compute $\gamma_{n_2,n_1}(z, A)$ (defined in (1.23)).

3: Set $\Gamma_{n_3,n_2,n_1}(A) = \left\{ z \in \mathrm{Grid}(n_2) : \gamma_{n_2,n_1}(z, A) + \frac{1}{n_2} \le \frac{1}{n_3} \right\}$.

Output: $\Gamma_{n_3,n_2,n_1}(A)$, a finite subset of $\mathbb{C}$ with $\lim_{n_3\to\infty} \lim_{n_2\to\infty} \lim_{n_1\to\infty} \Gamma_{n_3,n_2,n_1}(A) = \mathrm{Sp}(A)$.

related to the resolvent norm through Lemma 1.2.5. Second, they can be approximated for arbitrary operators such as $A\mathcal{P}_n^*$, whereas eigenvalues of such operators are not well defined.

In summary, we have developed an algorithm that computes spectra without suffering from spectral pollution or invisibility. This is encouraging, and readers may wonder why further discussion is needed. Yet a dilemma arises: convergence of a three-limit algorithm cannot be realised on a finite machine. *How, then, can we ensure the reliability of numerical computations? Are three successive limits truly necessary, or might there exist a different algorithm requiring fewer? And if three limits are indeed essential, for which classes of operators can spectra be computed with fewer limits?* Answering such questions for different spectral problems will be a key theme of this book, and it is helpful to first reflect on what we can learn from the above algorithm.

1.3 What Can We Learn from This Algorithm?

Algorithm 1.1 reveals two key insights. First, it introduces a hierarchy for measuring computational difficulty and algorithm optimality. Second, it highlights the resolvent's importance in addressing infinite-dimensional spectral problems. We briefly discuss each insight, and conclude this chapter with three example classifications in the resulting hierarchy.

1.3.1 The SCI Hierarchy, Our Tool for Foundations

The study of the above three-limit algorithm led to the *solvability complexity index* (SCI). Informally, the SCI counts the number of successive limits required to solve a computational problem, providing a measure of its intrinsic difficulty and giving rise to the *SCI hierarchy*. Chapter 2 introduces and analyses this hierarchy. In particular, Theorem 2.3.10 shows that the three successive limits in (1.19) are optimal for the full class $\Omega_{\ell^2(\mathbb{N})}$ (indeed, even for the subclass of all bounded operators). In other words, it is impossible to compute spectra of completely general operators using two successive limits (let alone one)! This result holds for all models of computation, even if one allows exact arithmetic over $\mathbb{R}$.

Here, we are assuming that a fixed algorithm can only use a finite amount of information about the operator before producing its output; see Section 2.1.2. In this case, the information corresponds to sampling the matrix entries of the operator. If we modify the information available to our algorithms, it may be possible to avoid three limits.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Moreover, the fact that three limits are sharp does not preclude algorithms converging for different *subclasses* of operators. As an analogy, recall the halting problem: given an algorithm and an input, decide whether the algorithm will eventually terminate on that input. This problem is undecidable, meaning that no universal algorithm exists that solves the problem for all possible program-input pairs. However, one can design a universal algorithm for certain programs. That is, if we restrict the class of program-input pairs, the halting problem can become decidable.

Similarly, not all spectral problems present the same level of difficulty and require three limits. For instance, the algorithm used to generate the bottom panel of Figure 1.4 (quasicrystal spectra) converges with a single limit and offers verification. In Chapter 3, we shall see that spectra can be computed with error control for a broad class of operators. This process merges the first two limits in Figure 1.11 while controlling the final limit.

There exists a class Σ_1 in the SCI hierarchy tailored to verified computations. Recall the Hausdorff metric from Definition 1.2.2 for non-empty compact sets $X, Y \subset \mathbb{C}$:

$$d_{\mathrm{H}}(X,Y) = \max\left\{\sup_{x\in X}\inf_{y\in Y}|x-y|, \sup_{y\in Y}\inf_{x\in X}|x-y|\right\}. \tag{1.26}$$

A computational problem lies in Σ_1 if there exists a convergent algorithm Γ_n that also computes a sequence $E_n \to 0$ such that $\Gamma_n(A)$ is contained in the E_n-neighbourhood of the correct answer. In the spectral setting, this means we can control the first supremum on the right-hand side of (1.26) with $X = \Gamma_n(A)$ and $Y = \mathrm{Sp}(A)$. Being in Σ_1 thus guarantees not only convergence to the true spectrum, but also that every computed point for a given n is reliable. However, it does not ensure that all parts of the spectrum have been approximated for a fixed n. That is, the second supremum in (1.26) may remain uncontrolled when $X = \Gamma_n(A)$ and $Y = \mathrm{Sp}(A)$. Numerous spectral problems belong to Σ_1, though not all.

Throughout this book, the SCI hierarchy serves as our toolkit for understanding the foundations of infinite-dimensional spectral computations. The need for multiple limits also arises in many other areas of computational mathematics; see Chapter 2. Broadly speaking, algorithms often depend on limiting values of several parameters. Balancing these parameters can occasionally reduce the number of successive limits required, but often it cannot. The SCI hierarchy captures this phenomenon precisely. Sharp classifications within it delineate the boundaries of what is computationally achievable and provide a rigorous measure of a problem's intrinsic difficulty, preventing wasted effort on algorithms that cannot exist.

1.3.2 The Resolvent, Our Algorithm-Building Tool

A recurring theme in this book is the reduction of infinite-dimensional spectral problems to a sequence of practical, finite-dimensional ones that are guaranteed to converge. Importantly, these finite-dimensional problems are not limited to matrix eigenvalue problems!

The most important tool in this book for constructing and analysing algorithms is the resolvent operator:

$$(A,z) \mapsto (A - zI)^{-1} \quad z \notin \mathrm{Sp}(A).$$

It provides a robust way to access the spectral properties of infinite-dimensional operators.

Typically, we use the resolvent in one or both of the following ways:

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Table 1.1 *Examples of the use of the resolvent in this book. The use of flavour (b) in Chapter 10 is one of the most classical applications of the resolvent in infinite-dimensional spectral computations.*

Chapter	Flavour (a): Injection moduli and related	Flavour (b): Action of $(A-zI)^{-1}$ on vectors
3	Computing spectra and pseudospectra, and testing if these sets intersect a region.	*not applicable*
4	*not applicable*	Computing spectral measures and functions of self-adjoint operators.
5	*not applicable*	Computing spectral measures and functions of unitary operators.
6	*not applicable*	Computing decompositions of spectral measures and spectra into different types for self-adjoint and unitary operators.
7	Computing size of spectra: Lebesgue measure, fractal dimension, number of connected components, and capacity.	*not applicable*
8	Computing essential and discrete spectra, multiplicities, and related problems.	*not applicable*
9	Computing spectral radii and abscissas, and spectral gaps.	*not applicable*
10	Computing spectra and pseudospectra for nonlinear spectral problems, verifying the output of contour methods.	Computing contour integrals (e.g., spectral projections) for eigenvalues and eigenvectors of nonlinear spectral problems.
11	Data-driven computations of spectra of Koopman operators for (nonlinear) dynamical systems.	Data-driven computations of spectral measures of Koopman operators for (nonlinear) dynamical systems.

(a) Computation of injection moduli or related quantities. Lemma 1.2.5 forms the basis of this technique, and injection moduli are also explicit in the definition of pseudospectra (Definition 1.1.6). The resolvent norm always satisfies (Exercise 1.8)

$$\operatorname{dist}(z, \operatorname{Sp}(A)) \geq \|(A - zI)^{-1}\|^{-1}.$$

On the other hand, if we can bound $\operatorname{dist}(z, \operatorname{Sp}(A))$ above by quantities involving $\|(A - zI)^{-1}\|^{-1}$, we can compute spectra without passing through pseudospectra.

(b) Computation of the action of $(A - zI)^{-1}$ on vectors in $\mathcal{H}$. Here, the main computational task is solving linear systems, one of the most developed areas of computational mathematics. The challenge, however, is that the linear systems are infinite-dimensional.

We have already hinted at two examples of this approach in Sections 1.1.1 and 1.1.3, for flavours (a) and (b), respectively. Rather than using a 'discretise-then-solve' approach, we develop algorithms that use the resolvent and discretise at the last moment. We can think of these techniques as a 'solve-then-discretise' approach since they are typically composed of adaptive computations involving the resolvent. This is a tool for directly studying the full infinite-dimensional operator, even on a finite computer. Table 1.1 lists some of the spectral problems where each of these two flavours is employed.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

1.3.3 Three Warm-Up Classifications

Concluding this chapter, we discuss three example classifications in the SCI hierarchy. These illustrate that problems involving different classes of operators may have varying computational difficulties. In the next chapter, we pick up this thread and present the SCI hierarchy as a toolkit for precisely classifying spectral computations in infinite dimensions.

Diagonal operators

As a simple example of a problem in Σ_1, consider the class Ω_{D} of bounded self-adjoint diagonal operators on $\ell^2(\mathbb{N})$. These are operators of the form

$$A = \begin{pmatrix} a_1 & & \\ & a_2 & \\ & & \ddots \end{pmatrix}, \quad \{a_j\}_{j=1}^{\infty} \subset \mathbb{R}, \quad \sup_{j\in\mathbb{N}} |a_j| < \infty. \tag{1.27}$$

Letting $\Gamma_n(A) = \{a_j : j = 1, \ldots, n\} \subset \mathbb{R}$ be the leading subset of the diagonal of A, $\lim_{n\to\infty} \Gamma_n(A) = \mathrm{Sp}(A)$ in the Hausdorff metric. Additionally, $\Gamma_n(A)$ satisfies

$$\Gamma_n(A) \subset \mathrm{Sp}(A) \quad \forall n \in \mathbb{N} \quad \forall A \in \Omega_{\mathrm{D}}. \tag{1.28}$$

Namely, we have convergence 'from below'. This statement is much stronger than just convergence since we can bound the distance of every point of $\Gamma_n(A)$ to the spectrum (in this case, this bound is zero). This property is important since it means that the output comes with a guarantee of reliability. Diagonal operators may seem trivial, but they provide a clear example of Σ_1 convergence, formally defined in Section 2.2.3.

Additionally, no sequence of (even adaptive) algorithms $\{\tilde{\Gamma}_n\}$ can produce outputs $\tilde{\Gamma}_n(A)$ accompanied by error bounds $E_n(A)$ for the full Hausdorff metric satisfying:

$$d_{\mathrm{H}}(\tilde{\Gamma}_n(A), \mathrm{Sp}(A)) \leq E_n(A), \quad \lim_{n\to\infty} E_n(A) = 0 \quad \forall A \in \Omega_{\mathrm{D}}. \tag{1.29}$$

This is impossible because an algorithm $\tilde{\Gamma}_n$ can only process a finite subset of the $\{a_j\}$ for each n. Consequently, extreme diagonal elements may lie beyond those observed.

Comparing (1.28) and (1.29) with the Hausdorff metric in (1.16), we see that for diagonal operators, we can control precisely one of the suprema that define the distance d_{H}. Lemma 2.2.1 provides a formal statement of these results in the language of the SCI hierarchy.

In specific cases, the stronger form of error control in (1.29) is possible, notably for finite-rank perturbations of self-adjoint tridiagonal Toeplitz operators [1392], periodic operators [124], and certain aperiodic operators (Sections 7.2.5 and 7.2.6). Achieving these results necessitates substantial 'global' knowledge of the structure of the operators.

Jacobi operators

Remarkably, one can achieve essentially the same type of error control for many operators as for the seemingly trivial case of diagonal operators. For example, consider the class of

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

bounded Jacobi operators

$$J = \begin{pmatrix} b_1 & a_1 & & \\ a_1 & b_2 & a_2 & \\ & a_2 & b_3 & \ddots \\ & & \ddots & \ddots \end{pmatrix}, \quad \{a_j, b_j\}_{j=1}^{\infty} \subset \mathbb{R}, \quad a_j > 0, \quad \text{such that} \quad J \in \mathscr{B}(\ell^2(\mathbb{N})).$$

Significant research has been dedicated to these operators [426, 1303]. Yet, spectral pollution makes computing their spectra a serious challenge.

We can build an algorithm Γ_n with

$$\lim_{n\to\infty} d_{\mathrm{H}}(\Gamma_n(J), \mathrm{Sp}(J)) = 0,$$

such that the algorithm also computes an error bound $E_n(J; z)$ with

$$\mathrm{dist}(z, \mathrm{Sp}(J)) \leq E_n(J; z) \ \forall z \in \Gamma_n(J) \quad \text{and} \quad \lim_{n\to\infty} \sup_{z\in\Gamma_n(J)} E_n(J; z) = 0. \tag{1.30}$$

This is the Σ_1 notion of error control alluded to above. Chapter 3 discusses convergence to spectra with this type of verified error control for a broad class of operators.

For Jacobi operators, the description of the method is relatively simple. Essentially, we can bypass two of the limits depicted in Figure 1.11 as follows:

- The fact that J is tridiagonal implies that $\gamma_{n_2}(z, J) = \gamma_{n_2,n_2+1}(z, J)$ (i.e., set $n_1 = n_2 + 1$);
- The self-adjointness of J implies that $\gamma(z, J) = \mathrm{dist}(z, \mathrm{Sp}(J))$ (so any candidate point z can be verified as approximating the spectrum by checking $\gamma(z, J)$ itself).

We will see later how to use these facts to compute $\gamma_{n_2}(z, J)$ and then apply a local optimisation routine to compute $\mathrm{Sp}(J)$.

Self-adjoint compact operators

The spectrum of a self-adjoint compact operator A comprises zero and, at most, countably many real eigenvalues, which can only accumulate at zero. This result, often introduced in undergraduate courses on infinite-dimensional spaces, represents the simplest form of infinite-dimensional spectra. Moreover, the finite section method converges for spectra of compact operators. For example, on $\ell^2(\mathbb{N})$ (Exercise 1.20):

$$\lim_{n\to\infty} d_{\mathrm{H}}(\mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*), \mathrm{Sp}(A)) = 0, \tag{1.31}$$

and the multiplicities of nonzero eigenvalues also converge. Surely, then, computing the spectrum of such operators is at least as easy as Jacobi operators?

The answer is no! The following theorem shows that error control for general self-adjoint compact operators is impossible. Exercise 1.22 guides the reader through the proof as preparation for the impossibility theorems in later chapters.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Theorem 1.3.1 (Error control for general self-adjoint compact operators is impossible) *There does not exist a sequence of algorithms* $\{\Gamma_n\}$ *that takes as input each compact self-adjoint operator* A *acting on* $\ell^2(\mathbb{N})$ *and computes a non-empty compact set* $\Gamma_n(A)$ *from the matrix entries of* A *such that*

$$\lim_{n\to\infty} d_{\mathrm{H}}(\Gamma_n(A), \mathrm{Sp}(A)) = 0 \quad \textit{and} \quad \sup_{z\in\Gamma_n(A)} \mathrm{dist}(z, \mathrm{Sp}(A)) \leq 2^{-n}.$$

For the 'other side' of the Hausdorff metric, neither does there exist a sequence of algorithms $\{\Gamma_n\}$ *that takes as input each compact self-adjoint operator* A *acting on* $\ell^2(\mathbb{N})$ *and computes a non-empty compact set* $\Gamma_n(A)$ *from the matrix entries of* A *such that*

$$\lim_{n\to\infty} d_{\mathrm{H}}(\Gamma_n(A), \mathrm{Sp}(A)) = 0 \quad \textit{and} \quad \sup_{z\in\mathrm{Sp}(A)} \mathrm{dist}(z, \Gamma_n(A)) \leq 2^{-n}.$$

In both statements, $\{2^{-n}\}$ *can be replaced by any positive sequence converging to zero.*

Take-home messages of this chapter

- **Handling infinite-dimensional spectral problems:** Discretising infinite-dimensional spectral problems requires caution due to common pitfalls such as spectral pollution (persistent erroneous limit points in an approximation sequence), spectral invisibility (omitted regions or points of the spectrum), the challenge of continuous and essential spectra, and ill-conditioning.
- **Solving the classical computational spectral problem:** Operators on separable Hilbert spaces can be represented by infinite matrices. The classical computational spectral problem on $\ell^2(\mathbb{N})$ is to compute their spectra. It can be solved by an algorithm that achieves convergence through three successive limits (Algorithm 1.1). This algorithm uses smallest singular values, instead of eigenvalues, due to connections with the resolvent. It represents a shift beyond traditional methods in how we address the computational challenges posed by infinite-dimensional spectral problems.
- **Importance of the resolvent:** The resolvent operator stands out as the most important tool for developing convergent algorithms for infinite-dimensional spectral problems.
- **Classifying problem difficulty through the SCI:** The difficulty of computational problems varies across different classes of operators and the spectral properties we wish to compute. The Solvability Complexity Index (SCI) hierarchy allows us to categorise these problems based on the number of successive limits needed for convergence, and to prove the optimality of algorithms.

Next chapter...

We formally introduce the SCI hierarchy as our toolkit for classifying computational problems, culminating in the proof that (and reasons why) three successive limits are necessary for computing spectra of general operators.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Notes

Spectral pollution and invisibility

The term *spectral pollution* was introduced in a 1977 paper by Rappaz [1104], although the phenomenon had been observed earlier. Spectral pollution arises across many fields, including elasticity, electromagnetism, hydrodynamics, magnetohydrodynamics, quantum chemistry and physics, relativistic computations, Sturm–Liouville operators, waveguides, among many others [2, 219, 222, 403, 440, 458, 476, 623, 635, 770, 859, 1105, 1106, 1180, 1241, 1256, 1257]. Theorem 1.1.4 is due to Pokrzywa [1074], with related results by Descloux [439]. The essential numerical range and its role in characterising spectral pollution are discussed further in Section 8.4. If the resolvent can be computed, spectral projections and related generalisations can be applied in regions away from the essential spectrum to avoid pollution; see, for example, the analysis of contour methods in Section 10.3.

Significant effort has been devoted to designing specific bases that mitigate spectral pollution for particular classes of operators. For quantum mechanics, see Lewin and Séré [893]; for finite element spaces in magnetohydrodynamics, see Gruber and Rappaz [635]; and for edge elements in Maxwell's eigenvalue problem, see Boffi [170, Section 20] (analysed in Chapter 3). It is often difficult to select bases that avoid pollution unless the operator admits a convenient decomposition or there is substantial prior knowledge of its spectrum.

For differential operators on unbounded domains or those with singularities, spectral pollution arising from domain truncation has been extensively studied [84, 244, 944, 1257]. Pollution can sometimes be avoided by imposing suitable boundary conditions on the truncated domain [797]. However, devising such 'non-reflecting' boundary conditions can be as difficult and problem-specific as designing non-polluting bases [948]. An alternative is the complex scaling method [8], reappearing as the perfectly matched layer method [136]. This latter technique replaces a self-adjoint problem with a non-self-adjoint one.

Another strategy for differential operators, proposed by Marletta [945], involves perturbing a self-adjoint operator. The perturbation lifts eigenvalues from gaps in the essential spectrum into the complex plane [15, 711, 946, 948, 1263]. A closely related technique is the dissipative barrier method [16], which likewise introduces perturbations to produce a sequence of non-self-adjoint spectral problems. These techniques originate in computational chemistry, where they are often referred to as complex absorbing potential methods, and have also been used to study quantum mechanical resonances [1028, 1029].

Several software packages have been developed for singular Sturm–Liouville problems, posed on either infinite intervals or finite intervals with coefficients exhibiting singular behaviour near at least one endpoint [87, 88, 947, 1089]. The standard approach involves using interval truncation to regularise the problem. Rigorous justification of these approaches, including establishing the absence of spectral pollution and invisibility, has been achieved for certain self-adjoint [84] and non-self-adjoint [241] Sturm–Liouville problems.

The term *spectral exactness*, meaning approximations are free from spectral pollution and invisibility, was introduced by Bailey, Everitt, Weidmann and Zettl [84] in the context of regular approximations of singular self-adjoint Sturm–Liouville problems. Some results regarding spectral invisibility for dissipative perturbations can be found in [946].

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Resolvent convergence

When studying spectral exactness, one often analyses the convergence of operator resolvents [785, Sections IV.2, VIII.1], [1113, Theorems VIII.23–25], [1400, Section 9.3]. This can involve varying the underlying spaces by embedding them into a larger one (e.g., Laplacians on varying domains) and examining the convergence of the corresponding lifted resolvents. For example, A_n converges to A in the generalised strong resolvent sense if there exists some z such that the lifted resolvents $(A_n - zI)^{-1}$ converge to the lifted resolvent $(A - zI)^{-1}$. Convergence in the generalised norm resolvent sense occurs when this convergence holds in norm [179, Definition 2.1]. These notions of convergence will appear in Section 4.2.

Generalised norm resolvent convergence guarantees the absence of spectral pollution [179, Theorem 2.4], whereas generalised strong resolvent convergence does not, even if all the operators are self-adjoint with compact resolvent [220, Example 5]. For self-adjoint operators, generalised strong resolvent convergence ensures the absence of spectral invisibility [179, Theorem 2.4]; however, in the non-self-adjoint case, even norm resolvent convergence may not prevent spectral invisibility [785, Example IV.3.8].

Bögli showed that generalised strong resolvent convergence, together with compact resolvents forming a discretely compact sequence, implies spectral exactness [179]. See [180] for local convergence results for operators with non-empty essential spectrum. The notion of a discretely compact sequence was introduced by Stummel, who established a spectral convergence theory for bounded operators [1269, 1270]. Anselone and Palmer obtained similar results [27] as did Osborn [1044] for the related notion of collectively compact sets of bounded operators.

Pseudospectra

The primary reference for pseudospectra is Trefethen and Embree's extensive book [1336], which includes a detailed subject history in Chapter 6. Interpreting pseudospectra as spectra of perturbed operators is in the spirit of backward error analysis as pioneered by Wilkinson [1419]. In one of his last papers, Wilkinson [1420] considered pseudospectra of matrices under the tentative name 'fundamental inclusion domains'. The term 'pseudospectrum' was introduced by Trefethen [1324]. Further references from the operator point of view include Davies [415] and Helffer [683]; see also the examples in [1325]. A lucid discussion of the subtleties in different definitions of pseudospectra is given in [293]. Pseudospectra have many uses, particularly for non-normal operators, including spectral sensitivity analysis and quantifying transient behaviour. Beyond the standard definition, one can consider variants restricting the perturbations in (1.14) to particular structures [1336, Chapter 50], [712, 1143, 1314], or extensions to nonlinear spectral problems [154, 704, 975] (see Chapter 10).

As with spectra, great care is needed when computing pseudospectra via discretisation methods, as emphasised in the review by Trefethen [1326]. Nevertheless, pseudospectra often exhibit better discretisation behaviour. For example, let $T(a)$ be a Toeplitz operator on $\ell^2(\mathbb{N})$ with symbol $a \in L^\infty(\mathbb{T})$. That is, $T(a)$ has the matrix representation $[T(a)]_{jk} = a_{j-k}$, where $\{a_n\}_{n\in\mathbb{Z}}$ are the Fourier coefficients of a. Suppose that a is piecewise continuous. While $\mathrm{Sp}(\mathcal{P}_n T(a)\mathcal{P}_n^*)$ can be very different from $\mathrm{Sp}(T(a))$ (see Section 1.1.4), $\mathrm{Sp}_\epsilon(\mathcal{P}_n T(a)\mathcal{P}_n^*)$ converges to $\mathrm{Sp}_\epsilon(T(a))$ as $n \to \infty$ in the Hausdorff metric for every $\epsilon > 0$ [197]. For

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

convergence rates of the resolvent norm, see [200]. Truncation methods for pseudospectra of differential operators are discussed in [409].

The one-dimensional convection–diffusion operator discussed in Section 1.1.2 was an important early operator example in pseudospectral analysis. Reddy and Trefethen showed that, although the spectrum of the truncated problem with Dirichlet boundary conditions lies on the real axis, its pseudospectra occupy large approximately parabolic regions and the resolvent norm may grow exponentially with the Péclet number [1108]. This fits into the broader Toeplitz/differential-operator picture described in [1336, Parts II and III]: in the constant-coefficient cases, pseudomodes are typically exponentially decaying boundary modes [1114, 1324], whereas variable coefficients lead to interior wave-packet pseudomodes governed by a twist condition [407, 408, 1327, 1334].

Some properties of pseudospectra are explored in Exercises 1.4–1.8 and 1.11. For example, Shargorodsky showed that the resolvent norm of an unbounded operator on a separable Hilbert space can be constant on open sets [1183]. This shows that the pseudospectrum can jump discontinuously as ϵ is varied. However, Globevnik showed that this cannot happen for bounded operators [598], and it is straightforward to show that it cannot happen for normal (even unbounded) operators. Moreover, it cannot happen for operators on Hilbert spaces with compact resolvent [417] or operators that generate a strongly continuous semigroup [1184]. The resolvent norm $\|(A - zI)^{-1}\|$ is subharmonic on $\mathbb{C} \setminus \mathrm{Sp}(A)$, which can sometimes be useful for proving the absence of spectral invisibility. Finally, as $\epsilon \downarrow 0$, $\mathrm{Sp}_\epsilon(A)$ converges to $\mathrm{Sp}(A)$ in the Attouch–Wets topology (see Exercise 1.11).

Finite section methods

Throughout this book, we use the term *finite section* to refer to spectral approximation via projection onto subspaces. For instance, given an operator $A \in \Omega_{\ell^2(\mathbb{N})}$, we can consider the eigenvalues of $\mathcal{P}_n A \mathcal{P}_n^*$, where $\mathcal{P}_n$ is the orthogonal projection onto $\mathrm{span}\{e_1, \ldots, e_n\}$. It is well known that the finite section method can lead to the pitfalls discussed in Section 1.1. Even in cases where it converges to the spectrum, it generally fails to provide the Σ_1 error control discussed in Section 1.3.1.

Arveson [49–53], Böttcher [197, 198], and Brown [245–247] pioneered finite section spectral computations from a C^*-algebraic perspective, both for general spectral problems and for Schrödinger operators. This line of work traces back at least to Böttcher and Silbermann [205]. Arveson [51–53] and Böttcher and Silbermann [207, 208] also explored spectral computation in terms of limiting distributions of eigenvalues, which relates to Szegő's work [1282] on finite section approximations. Laptev and Safarov have also obtained similar results [870]. Hansen [663] discusses the history of the finite section method and how to generalise various finite-dimensional techniques to infinite dimensions.

The term 'finite section' also applies to solving systems of linear equations, such as $Au = v$. One may approximate u by the solutions u_n of $\mathcal{P}_n A \mathcal{P}_n^* u_n = \mathcal{P}_n v$. The stability of $\{\mathcal{P}_n A \mathcal{P}_n^*\}$ plays an important role here, meaning that $\mathcal{P}_n A \mathcal{P}_n^*$ are invertible with uniformly bounded inverses for sufficiently large n. The limit operator method, which has its roots in the work of Favard [509], can be used to study Fredholmness and invertibility at infinity through certain asymptotic matrices related to A [909, 1098]. The books of Lindner [909] and Chandler–Wilde and Lindner [296] contain overviews and connections with spectra.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

The classes of operators most analysed in relation to the finite section method are Laurent and Toeplitz operators; standard references for these operators are the books of Böttcher and Silbermann and their collaborators [202, 203, 206–208, 1193], and classical references include [454, 599, 1415]. Methods devised for Toeplitz operators find applications in convolution operators, band-dominated operators, integral equations, and more [102, 196, 204, 648, 649, 781, 1088, 1099, 1126–1128, 1179]. For algorithms that compute the limiting set of eigenvalues of finite sections of Toeplitz operators, see [108, 201].

The classical computational spectral problem

The three-limit algorithm described in Section 1.2.3 is based on Hansen's algorithm [665] that provided the first universal algorithm for computing spectra and the pseudospectra of bounded operators on $\ell^2(\mathbb{N})$. See [664] for its development in the context of numerical linear algebra. This fundamental problem in operator theory goes back to Szegő's work on finite section approximations [1282] and was directly influenced by the above-mentioned work of Arveson, Böttcher and Brown on spectral computations in Hilbert spaces.

Computational approaches to obtaining spectral information date back to eminent physicists and mathematicians such as Schrödinger [1171], Kato [782], and Schwinger [1174, 1175]. Goldstine, Murray and von Neumann [607] were among the first to establish rigorous convergence results, yielding a Δ_1 classification in the SCI hierarchy (convergence in one limit with error control) for specific self-adjoint finite-dimensional problems. Schwinger introduced finite-dimensional approximations for quantum systems in infinite-dimensional spaces [1174, 1175], enabling spectral computations. These ideas were present in the work of Weyl [1414]. Digernes, Varadarajan and Varadhan [444] proved the convergence of spectra of Schwinger's finite-dimensional discretisation matrices for a class of Schrödinger operators with compact resolvent. This result is a crucial motivation for the results of Section 3.4.2.

The method described in Section 1.2.3 gave rise to the Solvability Complexity Index (SCI), or the number of limits required to compute an object, which we discuss in detail in Chapter 2. It was unknown for several years whether the three limits were sharp, and this was proven in [122], thus showing that the SCI hierarchy does not collapse.

Compact operators

Some of the first convergence results for spectral approximation in infinite dimensions concerned compact operators or those with compact resolvent. We refer the reader to the classical monograph of Babuška and Osborn [78], the early papers of Vainikko [1356–1358], Chatelin [300, 301, 303], and the books of Anselone [28] and Krasnoselskii, Vainikko, Zabreyko, Ruticki and Stet'senko [838]. For the general framework of discrete-regular convergence, see the work of Grigorieff [631–633]. For a closed operator, under the general hypothesis of strong stability, error bounds are given in Chatelin and Lemordant [305] and Chatelin [302]. For early results on non-self-adjoint eigenvalue approximations, see Bramble and Osborn [232] and Osborn [1044]. For lower bounds on eigenvalue approximation errors, see Kolata [825]. The book of Chatelin [304] studies spectral approximations of closed operators in a neighbourhood of an isolated eigenvalue with finite algebraic multiplicity. A major motivation for studying compact operators is finite element approximation of eigenvalue problems posed in variational form [74, 79], [170, Section 9]. Another important class of examples is

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

integral operators. For instance, see the work of Spence on the Nyström method [1235–1237]. In the exercises of this chapter, we walk the reader through a proof that the finite section method converges for computing spectra of general compact operators.

Solve-then-discretise

Many of the methods in this book aim to address the underlying operator as directly as possible—a 'solve-then-discretise' approach, in contrast to discretising first and then computing spectral properties. This is most evident in the computation of spectral measures (Chapters 4 and 5) and contour methods (Section 10.3). Other examples, not covered here, include the infinite Arnoldi method for infinite-dimensional eigenvalue problems [754–756], infinite-dimensional extensions of the QL and QR algorithms [340, 663, 1391] (see also [1322] for factorisations of quasimatrices), Toda and Hamiltonian flows [427, 428] (which inspired the infinite-dimensional QR algorithm), operator analogues of the Lanczos process for computing pseudospectra [438], and spectral measures via connection coefficients [1392]. The guiding principle is to postpone discretisation until the last possible moment. This perspective has proven highly effective across many areas of numerical analysis. Below, we list just a few examples.

For inverse problems, formulating Bayes' formula in infinite dimensions provides a framework for well-posedness, including robustness to data and approximations of the forward model. It also connects to regularisation theory and motivates new algorithms that exploit structure. See Stuart's article [1267] and the lecture notes of Dashti and Stuart [402].

The method of lines solves PDEs by discretising all but one dimension, typically time, yielding a system of ODEs. See the book of Schiesser [1156]. Conversely, Rothe's method (the method of lines transpose) discretises time, producing a sequence of boundary value problems [1137]. Another example is Kantorovich's extension of Newton's method to functional spaces [779]; see the survey [1079] and Deuflhard's book [441]. A linearise-then-discretise strategy can also significantly improve convergence for nonlinear integral equations [621, 622]. Directly handling operators has been applied in contour methods for time-dependent PDEs (including time-fractional cases) [328, 334], where, after the solve step, the task reduces to computing an infinite-dimensional inverse Laplace transform.

The Chebfun software for computing with functions [461] employs a solve-then-discretise approach. It combines spectral methods with lazy evaluation [460]; see [1328], the ODE examples in [1333], and block structures at the continuous level in [62]. For Trefethen's personal account, see [1330, Section 21]. Related work includes infinite-dimensional numerical linear algebra and computations with infinite data structures [1038, 1040].

Krylov subspace methods for infinite-dimensional operators lead to preconditioners that behave uniformly with respect to model and discretisation parameters; see Kirby [808] and Mardal and Winther [938]. They can also be built for differential operators using spectral methods [590]. In discontinuous Galerkin solutions of the Navier–Stokes equations, a discretisation is *adjoint consistent* if the discrete adjoint problem is a consistent discretisation of the continuous adjoint problem [670, 1173].

Probably the most popular delayed discretisation strategy is the 'optimise-then-discretise' approach in PDE-constrained optimisation. For example, see the work of Allaire and colleagues [17] in the field of shape and topology optimisation. The idea is to form optimality

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

conditions using variational calculus and then discretise these optimality conditions using finite elements. This follows the 'differentiate-then-discretise' philosophy discussed in Gunzburger's monograph [640], which also discusses and compares with the reverse approach. An alternative approach is to discretise the original PDEs and then derive optimality conditions. These two approaches are contrasted in the book of Hinze, Pinnau, Ulbrich and Ulbrich [716, Chapter 3]. Just like the spectral problems of this book, PDE-constrained optimisation problems are posed in infinite-dimensional spaces.

The finite-dimensional case and condition numbers

Golub and van der Vorst give an overview of developments in finite-dimensional eigenvalue problems during the 20th century [608]. Standard textbooks on matrix eigenvalue problems include Demmel [436], Golub and Van Loan [609], Householder [730], Parlett [1051], Saad [1146], Trefethen and Bau [1332], and Wilkinson [1419]. The QR algorithm [541, 853], ranked among the top ten algorithms of the 20th century [449], is guaranteed to converge for self-adjoint matrices when used with Wilkinson shifts [1051, Theorem 8.10.1].

Several key differences distinguish finite- from infinite-dimensional spectral problems. Fundamentally, many infinite-dimensional problems do not admit a convergent one-limit algorithm, motivating the development of the SCI. Algorithms that avoid the discretisation issues of Section 1.1 often have a different flavour from finite-dimensional algorithms. For example, in Chapter 3 we use singular values of truncations, rather than eigenvalues, to compute spectra. Much of this contrast stems from the presence of essential spectra: the infinite-dimensional QR algorithm [340, 663], for instance, tends to get 'stuck' at the essential spectrum, typically failing to recover the full spectrum or exploit shifts. A key theme of this book is to transform the infinite-dimensional problem into a sequence of finite-dimensional ones, designed so that the resulting methods converge and are practical.

Condition numbers for spectral projections can be defined for regions of discrete spectra [626, 785, 1249]. For non-isolated spectral points, however, one typically resorts to pseudospectral analysis. A natural candidate for conditioning the spectrum (measuring and bounding its sensitivity) will be introduced in Section 3.2.2. Even then, care is needed. For bounded self-adjoint operators A and B on the same Hilbert space [784, page 291],

$$d_{\mathrm{H}}(\mathrm{Sp}(A), \mathrm{Sp}(B)) \leq \|A - B\|,$$

which suggests that self-adjoint spectra, viewed as subsets of $\mathbb{R}$, are 'perfectly conditioned'. Yet, it is impossible to compute the spectrum of self-adjoint operators in one limit (Theorem 3.1.3). Even under additional assumptions such as bounded dispersion (a generalisation of bandedness; see Definition 3.3.1), the spectrum cannot be computed with full error control,[4] but rather Σ_1 error control. This also explains why classical complexity theory plays a lesser role in infinite dimensions: complexity is meaningful only for problems admitting full error control. Complexity estimates can be obtained for the finite problems arising from discretisation, but they often do not extend to the underlying infinite-dimensional problem.

[4] In the language of the SCI hierarchy, we are referring here to Δ_1; see Definition 2.1.16. These are the problems for which there exists a sequence of algorithms Γ_n such that for some known $c_n \to 0$, the error between $\Gamma_n(A)$ and the desired answer is bounded by c_n.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Exercises

Throughout these exercises, unless otherwise specified, $\mathcal{H}$ is a separable Hilbert space.

1.1 Spectral pollution eigenvectors converge weakly to zero. Let A be a bounded operator on $\mathcal{H}$. Suppose that $\{\mathcal{P}_n\}$ is a sequence of orthogonal projections onto finite-dimensional subspaces such that $\mathcal{P}_n^*\mathcal{P}_n$ converges strongly to the identity, and that there exists $\lambda_n \in \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)$ with $\lambda_n \to \lambda \notin \mathrm{Sp}(A)$. Let v_n be normalised eigenvectors of $\mathcal{P}_n A \mathcal{P}_n^*$ corresponding to λ_n. Prove that $\mathcal{P}_n^* v_n$ converges weakly to zero as $n \to \infty$.

1.2 Convection–diffusion operator. Consider the convection–diffusion operator $\mathcal{L}$ in (1.7). Using the Fourier transform, show that $\mathrm{Sp}(\mathcal{L}) = \{k^2 + 2\mathrm{i}k \colon k \in \mathbb{R}\}$. What happens if you approximate $\mathrm{Sp}(\mathcal{L})$ using the finite section method with a basis of Hermite functions?

1.3 The whale operator. Consider the Laurent operator L in (1.11). Using Fourier series, show that $\mathrm{Sp}(L) = \{a(z) \colon z \in \mathbb{C}, |z| = 1\}$. Plot the eigenvalues of the $n \times n$ truncations T_n and $a(C_n)$ for various values of n. You should see that, for small to medium n, the eigenvalues of T_n appear to cluster along arcs disjoint from $\{a(z) \colon z \in \mathbb{C}, |z| = 1\}$, whereas, for larger n, the eigenvalues are unstable as you vary n.

This example demonstrates spectral pollution, spectral invisibility, and ill-conditioning.

1.4 Equivalent definition of pseudospectra. Prove that if A is a closed and densely defined operator on $\mathcal{H}$, then

$$\mathrm{Sp}_\epsilon(A) = \mathrm{Cl}\left(\left\{z \in \mathbb{C} \colon \left\|(A - zI)^{-1}\right\|^{-1} < \epsilon\right\}\right) = \mathrm{Cl}\left(\bigcup_{E \in \mathscr{B}(\mathcal{H}) \colon \|E\| < \epsilon} \mathrm{Sp}(A + E)\right) \quad \forall \epsilon > 0.$$

1.5 The resolvent norm can be constant on open sets. Consider the Hilbert space

$$\mathcal{H} = \left\{u \in L^2([-\pi, \pi]) \colon \int_{-\pi}^{\pi} u(\tau)\, \mathrm{d}\tau = 0\right\}.$$

Let $\mathcal{T}$ denote the periodic Hilbert transform

$$[\mathcal{T}u](t) = \frac{1}{2\pi}\, \mathrm{p.v.} \int_{-\pi}^{\pi} u(\tau) \cot\left(\frac{t - \tau}{2}\right) \mathrm{d}\tau.$$

For smooth periodic functions $u \in \mathcal{H}$, define

$$[\mathcal{A}u](t) = u'(t) + u'(-t) + \frac{1}{2}[\mathcal{T}u](t) - \frac{1}{2}[\mathcal{T}u](-t) + \frac{1}{4}\int_{-t}^{t} u(\tau)\, \mathrm{d}\tau.$$

Show that $\mathcal{A}$ admits a closure that is unitarily equivalent to the operator

$$A = \mathrm{diag}(A_1, A_2, \ldots), \quad A_n = \begin{pmatrix} 0 & 2n \\ 1 + \frac{1}{2n} & 0 \end{pmatrix}.$$

Prove that $\|(A - zI)^{-1}\| = 1$ for $|z| < (\sqrt{5} - 1)/2$. It follows that

$$\mathrm{Cl}\left(\left\{z \in \mathbb{C} \colon \left\|(A - zI)^{-1}\right\|^{-1} < 1\right\}\right) \neq \left\{z \in \mathbb{C} \colon \left\|(A - zI)^{-1}\right\|^{-1} \leq 1\right\}.$$

Explain why this result shows that the map $\mathbb{R}_{>0} \ni \epsilon \mapsto \mathrm{Sp}_\epsilon(A) \in \mathcal{M}_{\mathrm{AW}}$ is not continuous. Show, however, that it is always left-continuous.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Hint: Show that $\{\cos(nt)/\sqrt{\pi}, \sin(nt)/\sqrt{\pi} : n \in \mathbb{N}\}$ is an orthonormal basis of $\mathcal{H}$ with $\mathcal{T}\cos(nt) = \sin(nt)$ and $\mathcal{T}\sin(nt) = -\cos(nt)$.

This example is due to Shargorodsky [1183].

1.6 Properties of pseudospectra I. Let $S \subset \mathbb{C}$ be a domain (non-empty connected open set) and $u : S \to \mathbb{R}$. We say that u is *subharmonic* if u is upper semicontinuous ($\limsup_{z \to z_0} u(z) \leq u(z_0)$) and for every $z \in S$ and $r > 0$ such that $B_r(z) \subset S$,

$$u(z) \leq \frac{1}{2\pi} \int_0^{2\pi} u(z + re^{\mathrm{i}\theta}) \, \mathrm{d}\theta.$$

If $u : S \to \mathbb{R}$ is subharmonic and attains a maximum in S, prove that it must be constant.

Suppose that A is a closed and densely defined operator on $\mathcal{H}$ and $z_0 \notin \mathrm{Sp}(A)$. Show that there is a domain containing z_0 on which the function $z \mapsto \|(A - zI)^{-1}\|$ is subharmonic. Deduce that if $\epsilon > 0$ and S is a bounded component of the set $\{z \in \mathbb{C} : \|(A - zI)^{-1}\|^{-1} < \epsilon\}$, then $S \cap \mathrm{Sp}(A) \neq \emptyset$.

1.7 When is the pseudospectrum continuous? Use Exercise 1.6 to show that if A is a closed and densely defined operator on $\mathcal{H}$, and if there is no open subset of $\mathbb{C} \setminus \mathrm{Sp}(A)$ on which the resolvent norm $\|(A - zI)^{-1}\|$ is constant, then

$$\mathrm{Cl}\left(\left\{z \in \mathbb{C} : \left\|(A - zI)^{-1}\right\|^{-1} < \epsilon\right\}\right) = \left\{z \in \mathbb{C} : \left\|(A - zI)^{-1}\right\|^{-1} \leq \epsilon\right\} \quad \forall \epsilon > 0.$$

Deduce that the map $\mathbb{R}_{>0} \ni \epsilon \mapsto \mathrm{Sp}_\epsilon(A) \in \mathcal{M}_{\mathrm{AW}}$ is continuous in this case.

The above hypothesis holds for normal operators and bounded operators [598].

1.8 Properties of pseudospectra II. Let A be a closed and densely defined operator on $\mathcal{H}$. Prove that the function $z \mapsto \|(A - zI)^{-1}\|^{-1}$ is Lipschitz continuous with Lipschitz constant bounded by one. Deduce that $\mathrm{dist}(z, \mathrm{Sp}(A)) \geq \|(A - zI)^{-1}\|^{-1}$. Show that if A is a normal operator, then $\|(A - zI)^{-1}\|^{-1} = \mathrm{dist}(z, \mathrm{Sp}(A))$ and $\mathrm{Sp}_\epsilon(A) = \mathrm{Sp}(A) + B_\epsilon(0)$.

Hint: For the final part, consider the spectral radius of the resolvent.

1.9 How many eigenvalues can we trust? Even if we avoid the pitfalls outlined in Section 1.1, only a portion of the eigenvalues computed for a fixed discretisation size are reliable. For the Hilbert space $\mathcal{H} = L^2([-1, 1])$, consider the (unbounded) operator

$$[\mathcal{L}u](x) = -\frac{\mathrm{d}}{\mathrm{d}x}\left[(1 + x^2)\frac{\mathrm{d}u}{\mathrm{d}x}\right](x), \quad u(\pm 1) = 0. \tag{1.32}$$

To build a discretisation of $\mathcal{L}$, consider the Legendre polynomials

$$P_0(x) = 1, \quad P_1(x) = x, \quad P_{n+1}(x) = \frac{1}{n+1}\left((2n+1)xP_n(x) - nP_{n-1}(x)\right) \text{ for } n \in \mathbb{N}.$$

These form an orthogonal basis of $L^2([-1, 1])$ and satisfy

$$P_n'(x) = \sum_{j=0}^{\lfloor \frac{n-1}{2} \rfloor} (2(n - 1 - 2j) + 1)P_{n-1-2j}(x).$$

To satisfy the boundary conditions at $x = \pm 1$, we consider the functions

$$f_n(x) = P_{n+1}(x) - P_{n-1}(x), \quad n \in \mathbb{N}.$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

Find infinite matrices $A = \{A_{ij}\}_{i,j\in\mathbb{N}}$ and $B = \{B_{ij}\}_{i,j\in\mathbb{N}}$ such that

$$\mathcal{L}f_j = \sum_{i=1}^{\infty} A_{ij}P_{i-1}, \qquad f_j = \sum_{i=1}^{\infty} B_{ij}P_{i-1}.$$

Argue why the eigenvalue problem associated with $\mathcal{L}$ is formally equivalent to $(A-\lambda B)v = 0$. Code up various truncations of this (generalised) eigenvalue problem and compare the computed eigenvalues across different truncations. Roughly speaking, what portion of the computed eigenvalues are reliable?

Often we wish to compute error bounds for eigenvalues or spectra. In this example, such bounds can be obtained as follows. Suppose that we have a candidate eigenvector

$$v(x) = \sum_{n=1}^{N} v_n f_n(x) = \sum_{n=1}^{N} v_n \sum_{i=1}^{\infty} B_{in}P_{i-1}(x)$$

with candidate eigenvalue λ. Using the fact that A and B have only finitely many nonzero entries in each column, express both v and $\mathcal{L}v - \lambda v$ as a finite sum of Legendre polynomials. Using the fact that

$$\int_{-1}^{1} P_m(x)P_n(x)\,\mathrm{d}x = \frac{2}{2n+1}\delta_{mn},$$

write a piece of code that computes $\|\mathcal{L}v - \lambda v\|$ and $\|v\|$. Using Exercise 1.8, explain why

$$\operatorname{dist}(\lambda, \operatorname{Sp}(\mathcal{L})) \le \|\mathcal{L}v - \lambda v\|/\|v\|,$$

and use your code to compute eigenvalue error bounds.

For further reading, see [1396, 1448].

1.10 Attouch–Wets vs. Hausdorff metrics. Prove the characterisation of convergence in the Attouch–Wets metric given in Lemma 1.2.4. Suppose that $C, C_n \subset \mathbb{C}$ are compact, non-empty sets. Prove that if $\lim_{n\to\infty} d_{\mathrm{H}}(C_n, C) = 0$, then $\lim_{n\to\infty} d_{\mathrm{AW}}(C_n, C) = 0$. Conversely, prove that if $\lim_{n\to\infty} d_{\mathrm{AW}}(C_n, C) = 0$ and $\cup_n C_n$ is contained in a bounded ball $B_R(0)$, then $\lim_{n\to\infty} d_{\mathrm{H}}(C_n, C) = 0$. What happens if we drop the assumption that $\cup_n C_n$ is contained in a ball?

A generalisation of Lemma 1.2.4 for general $(\mathcal{M}', d')$ is given in [116, Theorem 3.1.7].

1.11 Pseudospectra converge to spectra. Let A be a closed and densely defined operator on $\mathcal{H}$. Prove that $\operatorname{Sp}_\epsilon(A)$ converges to $\operatorname{Sp}(A)$ in the Attouch–Wets topology as $\epsilon \downarrow 0$.

1.12 We need the adjoint in the definition of γ. The approximate point spectrum of a closed and densely defined operator A on $\mathcal{H}$ is the set

$$\operatorname{Sp}_{\mathrm{ap}}(A) = \{z \in \mathbb{C} \colon \sigma_{\inf}(A - zI) = 0\}.$$

Show that an equivalent definition of the approximate point spectrum is

$$\operatorname{Sp}_{\mathrm{ap}}(A) = \left\{z \in \mathbb{C} \colon \exists\{x_n\} \subset \mathcal{D}(A) \text{ s.t. } \|x_n\| = 1, \lim_{n\to\infty} \|(A - zI)x_n\| = 0\right\}.$$

Give an example of an operator $A \in \Omega_{\ell^2(\mathbb{N})}$ and $z \in \mathbb{C}$ for which $\sigma_{\inf}(A - zI) \neq \sigma_{\inf}(A^* - \overline{z}I)$. Show that this cannot happen if $\operatorname{Sp}(A)$ has empty interior. (The interior of a subset S of a topological space X is the union of all subsets of S that are open in X.)

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

1.13 Dissecting the spectrum. Let A be a bounded linear operator on $\mathcal{H}$ and let $z \in \mathbb{C}$. A version of the open mapping theorem states that if $A - zI$ is surjective, then it is an open map (maps open sets to open sets). Using this, show that if $A - zI$ is a bijection on $\mathcal{H}$, then $z \notin \mathrm{Sp}(A)$. This allows us to decompose $\mathrm{Sp}(A)$ into three parts:

- *Point spectrum:* the set of eigenvalues of A (z for which $A - zI$ is not injective).
- *Continuous spectrum:* the set of z for which $A - zI$ is injective and has dense range, but is not bijective. (Not to be confused with the different definition used in Chapter 6.)
- *Residual spectrum:* the set of z for which $A - zI$ is injective but does not have dense range.

Show that if z is in either the point spectrum or continuous spectrum, then $\sigma_{\inf}(A - zI) = 0$. Is the converse true? Comment on the link between the residual spectrum and the need to include $\sigma_{\inf}(A^* - \overline{z}I)$ in the formula $\|(A - zI)^{-1}\|^{-1} = \min\{\sigma_{\inf}(A - zI), \sigma_{\inf}(A^* - \overline{z}I)\}$.

1.14 Adjoints, limit inferior, and limit superior. Suppose that for a sequence of operators $\{A_n\}$, $\mathcal{D}(A) \subset \mathcal{D}(A_n)$ and $\lim_{n\to\infty} \|A_n v - Av\| = 0$ for all $v \in \mathcal{D}(A)$. Show that

$$\limsup_{n\to\infty} \sigma_{\inf}(A_n - zI) \leq \sigma_{\inf}(A - zI) \quad \forall z \in \mathbb{C}.$$

Provide an example where $\sigma_{\inf}(A^* - \overline{z}I) < \liminf_{n\to\infty} \sigma_{\inf}(A_n^* - \overline{z}I)$.

1.15 Rigorous three limits. Prove (1.25), explaining the necessity of the '$+1/n_2$' term and '$\leq$' (instead of strict inequality '$<$') in the definition of $\Gamma_{n_3,n_2,n_1}(A)$. Assuming we can only compute the smallest singular value of a finite rectangular matrix to within any desired *nonzero* approximation error, adapt the algorithm in Section 1.2.3 to provide a rigorous method for computing $\mathrm{Sp}(A)$ in three successive limits from A's matrix elements.

1.16 Finite rank approximates compact. Recall from Appendix A.3 that an operator A on $\mathcal{H}$ is said to be compact if the closure of $\{Ax : \|x\| \leq 1\}$ is compact in $\mathcal{H}$. Show that every bounded, finite-rank operator is compact. Let A be compact and $\{e_n : n \in \mathbb{N}\}$ be an orthonormal basis of $\mathcal{H}$. Define $A_n = \mathcal{P}_n^* \mathcal{P}_n A \mathcal{P}_n^* \mathcal{P}_n$, where $\mathcal{P}_n : \mathcal{H} \to \mathrm{span}\{e_1, \ldots, e_n\}$ is the corresponding orthogonal projection. Prove that $\|A - A_n\| \to 0$ as $n \to \infty$.

1.17 Spectra cannot suddenly expand. Let A be a closed and densely defined operator on $\mathcal{H}$ such that $0 \notin \mathrm{Sp}(A)$. Let B be a bounded operator on $\mathcal{H}$ with $\|B\|\|A^{-1}\| < 1$. Prove that

$$0 \notin \mathrm{Sp}(A + B) \quad \text{and that} \quad \|(A + B)^{-1}\| \leq \frac{\|A^{-1}\|}{1 - \|B\|\|A^{-1}\|}.$$

Show that if A is bounded and $\epsilon > 0$, then there exists $\delta > 0$ such that

$$\sup_{z \in \mathrm{Sp}(A+B)} \mathrm{dist}(z, \mathrm{Sp}(A)) < \epsilon \quad \text{whenever} \quad \|B\| < \delta.$$

Hence, the spectrum is upper semicontinuous with respect to the operator norm.

For a discussion on the continuity points of the mapping $A \mapsto \mathrm{Sp}(A)$, see [358].

1.18 Approximating components of spectra. Let A be a closed and densely defined operator on $\mathcal{H}$ and let X be a bounded connected component of $\mathrm{Sp}(A)$ separated from the rest of the spectrum. Let B_n be a sequence of bounded operators with $\|B_n\| \to 0$ as $n \to \infty$. Show that given $\epsilon > 0$, there exists a $N \in \mathbb{N}$ such that if $n \geq N$, then $\inf_{z \in \mathrm{Sp}(A+B_n)} \mathrm{dist}(z, X) < \epsilon$.

Hint: Use Exercise 1.6 and Exercise 1.17.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003

1.19 Where did the spectrum go? Let A_ϵ act on $\ell^2(\mathbb{Z})$ by

$$A_\epsilon e_n = \begin{cases} \epsilon e_1, & n = 0, \\ e_{n+1}, & n \neq 0. \end{cases}$$

Show that $\mathrm{Sp}(A_0) = \{z\colon |z| \leq 1\}$, whereas $\mathrm{Sp}(A_\epsilon) = \{z\colon |z| = 1\}$ if $\epsilon \neq 0$. What does this example say about the discontinuity of the spectrum as a function of operators?

1.20 Finite section converges for compact operators. Let A be a compact operator on $\ell^2(\mathbb{N})$. The spectral theorem for compact operators says that $0 \in \mathrm{Sp}(A)$, $\mathrm{Sp}(A)$ is at most countable, and every nonzero element of $\mathrm{Sp}(A)$ is an eigenvalue (of finite algebraic multiplicity) [355, Section VII.7]. Use Exercises 1.16 and 1.18 to show that

$$\lim_{n\to\infty} \mathrm{dist}(z, \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)) = 0 \quad \forall z \in \mathrm{Sp}(A) \setminus \{0\}.$$

Noting that $\mathrm{Sp}(\mathcal{P}_n^* \mathcal{P}_n A \mathcal{P}_n^* \mathcal{P}_n)$ need not equal $\mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)$, adapt your argument to show that we may also take $z = 0$. Use Exercise 1.17 to show that

$$\lim_{n\to\infty} \sup_{z \in \mathrm{Sp}(\mathcal{P}_n A \mathcal{P}_n^*)} \mathrm{dist}(z, \mathrm{Sp}(A)) = 0.$$

Conclude that the eigenvalues of $\mathcal{P}_n A \mathcal{P}_n^*$ converge to $\mathrm{Sp}(A)$ in the Hausdorff metric.

This exercise shows that for compact operators, the finite section method avoids the pitfalls of spectral pollution and invisibility discussed in Section 1.1.

1.21 Finite section can fail for self-adjoint operators with compact resolvent. Let A be a closed and densely defined operator on $\mathcal{H}$. Show that if $z \in \mathbb{C} \setminus \mathrm{Sp}(A)$ is such that $(A - zI)^{-1}$ is compact, then $(A - wI)^{-1}$ is compact for every $w \in \mathbb{C} \setminus \mathrm{Sp}(A)$. In this case, we say that A has *compact resolvent*. Show that if A has compact resolvent, then its spectrum consists of isolated points that are eigenvalues. Find an example of a tridiagonal, self-adjoint $A \in \Omega_{\ell^2(\mathbb{N})}$ with compact resolvent for which the finite section method does not produce convergent approximations of its spectrum.

1.22 Proving a lower bound. Prove Theorem 1.3.1. For the first statement, assume that such an algorithm exists and arrive at the following contradiction. Consider an operator of the type

$$A = A_k \oplus \mathrm{diag}(0, 0, \ldots) \quad \text{with} \quad A_k = \begin{pmatrix} 1 & & & & 1 \\ & 0 & & & \\ & & \ddots & & \\ & & & 0 & \\ 1 & & & & 1 \end{pmatrix} \in \mathbb{C}^{k\times k},$$

where we choose the value of k later. Let $C = \mathrm{diag}(1, 0, 0, \ldots)$, then $\mathrm{Sp}(C) = \{0, 1\}$ and clearly A is compact with $\mathrm{Sp}(A) = \{0, 2\}$. Show how one may choose k to trick the supposed algorithm and arrive at a contradiction. A similar argument deals with the second statement of Theorem 1.3.1. Extend these results to the problem of computing pseudospectra. Finally, show that the first part of Theorem 1.3.1 no longer holds for compact self-adjoint operators if, in addition to the matrix entries of A, we are given access to the matrix entries of A^2.

The last part is an example of 'operator folding', which we shall meet in Section 3.4.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 08:18:44, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009382502.003