Part III: Inverse Problems

Lecture notes, Michaelmas 2024
University of Cambridge

Ferdia Sherry
ferdiasherry.com

October 15, 2024

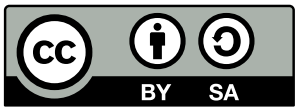

Contents

These lecture notes are for Part III "Inverse Problems", taught by Ferdia Sherry in Michaelmas 2024 at the University of Cambridge. These notes are based on the notes for Part III "Inverse Problems" taught by Yury Korolev and Jonas Latz in Michaelmas 2020 at the University of Cambridge [1]. Complementary material can be found in the following books, lecture notes and review papers:

1. Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*. Springer, 1996.

2. Otmar Scherzer, Markus Grasmair, Harald Grossauer, Markus Haltmeier and Frank Lenzen. *Variational Methods in Imaging*. Springer, 2008.

3. Kristian Bredies and Dirk Lorenz. *Mathematical Image Processing*. Springer, 2018

4. Martin Benning and Martin Burger. *Modern regularization methods for inverse problems*. Acta Numerica, 2018.

5. Bryan P. Rynne and Martin A. Youngson, *Linear Functional Analysis*, Springer, 2008.

6. Masoumeh Dashti and Andrew M. Stuart, *The Bayesian approach to inverse problems*, Handbook of Uncertainty Quantification, 2016.

7. Jari Kaipio and Erkki Somersalo, *Statistical and computational inverse problems*, vol. 160 of Applied Mathematical Sciences, 2005.

8. O. Kallenberg, *Foundations of modern probability theory*, Springer, 1997.

9. Andrew M. Stuart, *Inverse problems: a Bayesian perspective*, Acta Numerica, 2010.

These lecture notes are under constant redevelopment and might contain typos or errors. I would very much appreciate if you report any mistakes found to fs436@cam.ac.uk. Thanks!

[1] https://www.damtp.cam.ac.uk/research/cia/inverse-problems-michaelmas-2020

Chapter 1

Introduction to Inverse Problems

Inverse problems arise from the need to gain information about an unknown object of interest from given indirect measurements. Inverse problems have several applications varying from medical imaging and industrial process monitoring to ozone layer tomography and modelling of financial markets. The common feature for inverse problems is the need to understand indirect measurements and to overcome extreme sensitivity to noise and modelling inaccuracies. In this course we employ both deterministic and probabilistic approach to inverse problems to find stable and meaningful solutions that allow us quantify how inaccuracies in the data or model affect the obtained estimate.

1.1 Well-posed and ill-posed problems

We start by considering the problem of finding $u \in \mathbf{R}^d$ that satisfies the equation

$$f = Au, \tag{1.1}$$

where $f \in \mathbf{R}^k$ is given. We refer to f as observed data or measurement and u as an unknown. The physical phenomena that relates the unknown and the measurement is modelled by a matrix $A \in \mathbf{R}^{k\times d}$. In real life, the perfect data given in (1.1) is perturbed by noise and we observe measurements

$$f_\eta = Au + \eta, \tag{1.2}$$

where $\eta \in \mathbf{R}^k$ represents the observational noise.

We are interested in ill-posed inverse problems, where the inverse problem is more difficult to solve than the direct problem of finding f_η when u is given. To explain this, we first need to introduce well-posedness as defined by Jacques Hadamard [4]:

Definition 1.1. *A problem is called well-posed if*

1. *There exists at least one solution. (Existence)*
2. *There is at most one solution. (Uniqueness)*
3. *The solution depends continuously on data. (Stability)*

The direct or forward problem is assumed to be well-posed. The inverse problems are ill-posed and break at least one of the above conditions.

1. Assume that $d < k$ and $A : \mathbf{R}^d \to \mathcal{R}(A) \subsetneq \mathbf{R}^k$, where the range of A is a proper subset of $\mathbf{R}^k$. Furthermore, we assume that A has a unique inverse $A^{-1} : \mathcal{R}(A) \to \mathbf{R}^k$. Because of the noise in the measurement $f_\eta \notin \mathcal{R}(A)$ so that simply inverting A with the data given in (1.2) is not possible. Note that usually only the statistical properties of the noise n are known so we cannot just subtract it.

2. Assume next that $d > k$ and $A : \mathbf{R}^d \to \mathbf{R}^k$, in which case the system is underdetermined. We then have more unknowns than equations which means that there are several possible solutions.

3. Consider next the case $d = k$, in which there exists $A^{-1} : \mathbf{R}^k \to \mathbf{R}^k$. The condition number $\kappa = \lambda_1/\lambda_k$, where λ_1 and λ_k are the biggest and smallest eigenvalues of A, may be very large. Such a matrix is said to be ill-conditioned and is almost singular. In this case the problem is sensitive even to smallest errors in the measurement. Hence the naïve reconstruction $\widetilde{u} = A^{-1}f_\eta = u + A^{-1}\eta$ does not produce a meaningful solution but will be dominated by $A^{-1}\eta$. Note that $\|A^{-1}\eta\|_2 \approx \|\eta\|_2/\lambda_k$ can be arbitrarily large.

The last part illustrates one of the key questions of inverse problem theory: how can we stabilise the reconstruction process while maintaining acceptable accuracy?

A deterministic way of achieving a unique and stable solution for the problem (1.2) is to use regularisation theory. In the classical Tikhonov regularisation a solution is attained by solving

$$\min_{u\in\mathbf{R}^d} \left(\left\| Au - f_\eta \right\|^2 + \alpha \|Lu\|^2 \right). \tag{1.3}$$

Above, α acts as a tuning parameter balancing the effect of the data fidelity term $\left\|Au - f_\eta\right\|^2$ and the stabilising regularisation term $\|u\|^2$. The first half of the course will concentrate on regularisation theory.

Another way of tackling problems arising from ill-posedness is Bayesian inversion. The idea of statistical inversion methods is to rephrase the inverse problem as a question of statistical inference. We then consider the problem

$$f_\eta = Au + \eta, \tag{1.4}$$

where the measurement, unknown and noise are now modelled as random variables. This approach allows us to model the noise through its statistical properties. We can also encode our *a priori* knowledge of the unknown in form of a probability distribution that assigns higher probability to those values of u we expect to see. The solution to (1.4) is so-called *posterior distribution*, which is the conditional probability distribution of u given a measurement f_η. This distribution can then be used to obtain estimates that are most likely in some sense. We will return to the Bayesian approach to inverse problems in the second half of the course

In this course we will concentrate on continuous inverse problems where in (1.1) and (1.2) $A : \mathcal{X} \to \mathcal{Y}$ is a linear or non-linear forward operator acting between some spaces $\mathcal{X}$ and $\mathcal{Y}$, typically Hilbert or Banach spaces, the measured data $f_\eta \in \mathcal{Y}$ is a function and $u \in \mathcal{X}$ is the quantity we want to reconstruct from the data. Linear inverse problems include such important applications as computer tomography, magnetic resonance imaging and image deblurring in microscopy or astronomy. In other important applications, such as seismic imaging, the forward operator is non-linear (e.g., parameter identification problems for PDEs). Next we will take a look at some examples of linear and non-linear inverse problems to see what kind of challenges we face when trying to solve them.

1.2 Examples of inverse problems

1.2.1 Signal deblurring

The deblurring (or deconvolution) problem of recovering an input signal u from an observed signal

$$f_\eta(t) = \int_{-\infty}^{\infty} a(t-s)u(s)\,\mathrm{d}s + \eta(t)$$

occurs in many imaging, and image- and signal processing applications. Here the function a is known as the blurring kernel.

The noiseless data is given by $f(t) = \int_{-\infty}^{\infty} a(t-s)u(s)\,\mathrm{d}s$ and its Fourier transform is $\widehat{f}(\xi) = \int_{-\infty}^{\infty} \exp(-i\xi t)f(t)dt$. The convolution theorem implies

$$\widehat{f}(\xi) = \widehat{a}(\xi)\widehat{u}(\xi),$$

and hence by inverse Fourier transform

$$u(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \exp(it\xi)\frac{\widehat{f}(\xi)}{\widehat{a}(\xi)}\,\mathrm{d}\xi.$$

However, we can only observe noisy measurements and hence we have, in the frequency domain, $\widehat{f_\eta}(\xi) = \widehat{a}(\xi)\widehat{u}(\xi) + \widehat{\eta}(\xi)$. The estimate u_{est} based on the convolution theorem is given by

$$u_{\text{est}}(t) = u(t) + \frac{1}{2\pi}\int_{-\infty}^{\infty} \exp(it\xi)\frac{\widehat{\eta}(\xi)}{\widehat{a}(\xi)}\,\mathrm{d}\xi,$$

which is often not even well defined, since usually the kernel a decreases exponentially (or has compact support), making the denominator small, whereas the Fourier transform of the noise will be non-zero.

1.2.2 Heat equation

Next, we study the problem of recovering the initial condition u of the heat equation from a noisy observation f_n of the solution at some time $T > 0$. We consider the heat equation on a torus $\mathbf{T}^d = (\mathbf{R}/\mathbf{Z})^d$, with Dirichlet boundary conditions

$$\begin{cases} \frac{\partial v}{\partial t} - \Delta v = 0 & \text{on } \mathbf{T}^d \times \mathbf{R}_+ \\ v(x,t) = 0 & \text{on } \partial\mathbf{T}^d \times \mathbf{R}_+ \\ v(x,T) = f(x) & \text{on } \mathbf{T}^d \\ v(x,0) = u(x) & \text{on } \mathbf{T}^d \end{cases}$$

where Δ denotes the Laplace operator and $\mathcal{D}(\Delta) = H_0^1(\mathbf{T}^d) \bigcap H^2(\mathbf{T}^d)$. Note that the operator $-\Delta$ is positive and self-adjoint on Hilbert space $\mathcal{H} = L^2(\mathbf{T}^d)$.

Given a function $u \in L^2(\mathbf{T}^d)$ we can decompose it as a Fourier series

$$u(x) = \sum_{n\in\mathbf{Z}^d} \widehat{u_n}\exp(2\pi i\langle n, x\rangle),$$

where $\widehat{u_n} = \langle u, \exp(2\pi i\langle n, \cdot\rangle)\rangle_{L^2(\mathbf{T}^d)}$ are the Fourier coefficients, and the identity holds for almost every $x \in \mathbf{T}^d$. The L^2-norm of u is given by the Parseval's identity $\|u\|_{L^2}^2 = \sum |u_n|^2$. Remember that

the Sobolev space $H^s(\mathbf{T}^d)$, $s \in \mathbf{N}$, consist of all $L^2(\mathbf{T}^d)$ integrable functions whose α^{th} order weak derivatives exist and are $L^2(\mathbf{T}^d)$ integrable for all $|\alpha| \leqslant s$. The fractional Sobolev space $H^s(\mathbf{T}^d)$ is given by the subspace of functions $u \in L^2(\mathbf{T}^d)$, such that

$$\|u\|_{H^s}^2 = \sum_{n \in \mathbf{Z}^d} (1 + 4\pi^2|n|^2)^s |u_n|^2 < \infty. \tag{1.5}$$

Note that for a positive integer s, the above definition agrees with the definition given using the weak derivatives. For $s < 0$, we define $H^s(\mathbf{T}^d)$ via duality or as the closure of $L^2(\mathbf{T}^d)$ under the norm (1.5). The resulting spaces are separable for all $s \in \mathbf{R}$.

The eigenvectors of $-\Delta$ in $\mathbf{T}^d$ form the orthonormal basis of $L^2(\mathbf{T}^d)$ and the eigenvalues are given by $4\pi^2|n|^2$, $n \in \mathbf{Z}^d$. We can also work on real-valued functions where the eigenfunctions $\{\varphi_j\}_{j=1}^{\infty}$ comprise sine and cosine functions. The eigenvalues of $-\Delta$, when ordered as a sequence, then satisfy $\lambda_j \asymp j^{2/d}$. The notation $\asymp$ means that there exist constants $C_1, C_2 > 0$, such that $C_1 j^{2/d} \leqslant \lambda_j \leqslant C_2 j^{2/d}$.

The solution to the forward heat equation can be written as

$$v(t) = \sum_{j=1}^{\infty} u_j \exp(-\lambda_j t)\varphi_j.$$

We notice that

$$\|v(t)\|_{H^s}^2 \asymp \sum_{j=1}^{\infty} j^{\frac{2s}{d}} \exp(-2\lambda_j t)|u_j|^2 \asymp t^{-s} \sum_{j=1}^{\infty} (\lambda_j t)^s \exp(-2\lambda_j t)|u_j|^2 \leqslant Ct^{-s} \sum_{j=1}^{\infty} |u_j|^2 = Ct^{-s}\|u\|_{L^2}$$

which implies that $v(t) \in H^s(\mathbf{T}^d)$ for all $s > 0$.

We now have the observation model

$$f_\eta = Au + \eta,$$

where $A = \exp(T\Delta)$ and η is the observational noise. The noise is not usually smooth (the often assumed white noise is not even an L^2 function), and hence the measurement f_η is not in the image space $\operatorname{im}(\exp(T\Delta)) \subset \cap_{s>0} H^s(\mathbf{T}^d)$.

1.2.3 Differentiation

Consider the problem of evaluating the derivative of a function $f \in L^2[0, \pi/2]$. Let

$$Df = f',$$

where $D\colon L^2[0, \pi/2] \to L^2[0, \pi/2]$.

Proposition 1.2. *The operator D is unbounded from $L^2[0, \pi/2] \to L^2[0, \pi/2]$.*

Proof. Take a sequence $f_n(x) = \sin(nx)$, $n = 1, \ldots, \infty$. Clearly, $f_n \in L^2[0, \pi/2]$ for all n and $\|f_n\| = \sqrt{\frac{\pi}{4}}$. However, $Df_n(x) = n\cos(nx)$ and $\|Df_n\| = n \to \infty$ as $n \to \infty$. Therefore, D is unbounded. □

This shows that differentiation is ill posed when considered as an operator from L^2 to L^2. It does not mean that it can not be well-posed in other spaces. For instance, it is well-posed from H^1 (the Sobolev space of L^2 functions whose derivatives are also in L^2) to L^2. Indeed, $\forall u \in H^1$ we get

$$\|Df\|_{L^2} = \|f'\|_{L^2} \leqslant \|f\|_{H^1} \asymp \|f\|_{L^2} + \|f'\|_{L^2}.$$

However, since in practice we typically deal with functions corrupted by non-smooth noise, the setting of L^2 is relevant to practice, while the H^1 setting is not. Differentiation can be written as an inverse problem for an integral equation. For instance, the derivative u of some function $f \in L^2[0,1]$ with $f(0) = 0$ satisfies

$$f(x) = \int_0^x u(t)\,\mathrm{d}t,$$

which can be written as an operator equation $Au = f$ with $(A\cdot)(x) := \int_0^x \cdot(t)\,\mathrm{d}t$.

1.2.4 Matrix inversion

In finite dimensions, the inverse problem (1.1) is a linear system. Linear systems are formally well posed, in the sense that the error in the solution is bounded by some constant times the error in the right-hand side. However, this constant depends on the condition number of the matrix A and can get arbitrary large for matrices with large condition numbers. In this case, we speak of *ill-conditioned* problems.

Consider the problem (1.1) with $u \in \mathbf{R}^n$ and $f \in \mathbf{R}^n$ being n-dimensional vectors with real entries and $A \in \mathbf{R}^{n\times n}$ being a matrix with real entries. Assume further that A is symmetric and positive definite. We know from the spectral theory of symmetric matrices that there exist eigenvalues $\lambda_1 \geqslant \lambda_2 \geqslant \ldots \geqslant \lambda_n > 0$ and corresponding (orthonormal) eigenvectors $a_j \in \mathbf{R}^n$ for $j \in \{1,\ldots,n\}$ such that A can be written as

$$A = \sum_{j=1}^{n} \lambda_j a_j a_j^\top. \tag{1.6}$$

It is well known from numerical linear algebra that the condition number $\kappa = \lambda_1/\lambda_n$ is a measure of how stably (1.1) can be solved, which we will illustrate in what follows.

We assume that we measure f_δ rather than f, with $\|f - f_\delta\|_2 \leqslant \delta\|A\| = \delta\lambda_1$, where $\|\cdot\|_2$ denotes the Euclidean norm of $\mathbf{R}^n$ and $\|A\|$ the operator norm of A (which equals the largest eigenvalue of A). Then, if we further denote with u_δ the solution of $Au_\delta = f_\delta$, the difference between u_δ and the solution u to (1.1) is

$$u - u_\delta = \sum_{j=1}^{n} \lambda_j^{-1} a_j a_j^\top (f - f_\delta).$$

Therefore, we can estimate

$$\|u - u_\delta\|_2^2 = \sum_{j=1}^{n} \lambda_j^{-2} \underbrace{\|a_j\|_2^2}_{=1} |a_j^\top (f - f_\delta)|^2 \leqslant \lambda_n^{-2} \|f - f_\delta\|_2^2,$$

due to the orthonormality of eigenvectors, the Cauchy–Schwarz inequality, and $\lambda_n \leqslant \lambda_j$. Thus, taking square roots on both sides yields the estimate

$$\|u - u_\delta\|_2 \leqslant \lambda_n^{-1} \|f - f_\delta\|_2 \leqslant \kappa\delta.$$

Hence, we observe that in the worst case an error δ in the data y is amplified by the condition number κ of the matrix A. A matrix with large κ is therefore called *ill-conditioned*. Let us demonstrate the effect of this error amplification with a small example.

Example 1.1. Consider the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & \frac{1001}{1000} \end{pmatrix},$$

which has eigenvalues $\lambda_j = 1 + \frac{1}{2000} \pm \sqrt{1 + \frac{1}{2000^2}}$, condition number $\kappa \approx 4002 \gg 1$, and operator norm $\|A\| \approx 2$. For given data $f = (1,1)^\top$ the solution to $Au = f$ is $u = (1,0)^\top$. Now let us instead consider perturbed data $f_\delta = (99/100, 101/100)^\top$. The solution u_δ to $Au_\delta = f_\delta$ is then $u_\delta = (-19.01, 20)^\top$. Let us reflect on the amplification of the measurement error. By our initial assumption we find that $\delta = \|f - f_\delta\|/\|A\| \approx \|(0.01, -0.01)^\top\|/2 = \sqrt{2}/200$. Moreover, the norm of the error in the reconstruction is then $\|u - u_\delta\| = \|(20.01, 20)^\top\| \approx 20\sqrt{2}$. As a result, the amplification due to the perturbation is $\|u - u_\delta\|/\delta \approx 4000 \approx \kappa$.

1.2.5 Tomography

In almost any tomography application, the underlying inverse problem is either the inversion of the Radon transform[1] or of the X-ray transform. For $u \in C_0^\infty(\mathbf{R}^n)$, $s \in \mathbf{R}$, and $\theta \in S^{n-1}$ the *Radon transform* $R : C_0^\infty(\mathbf{R}^n) \to C^\infty(S^{n-1} \times \mathbf{R})$ can be defined as the integral operator

$$\begin{aligned} f(\theta, s) = (Ru)(\theta, s) &= \int_{x \cdot \theta = s} u(x)\, \mathrm{d}x \\ &= \int_{\theta^\perp} u(s\theta + y)\, \mathrm{d}y, \end{aligned} \tag{1.7}$$

which, for $n = 2$, coincides with the X-ray transform,

$$f(\theta, s) = (Pu)(\theta, s) = \int_{\mathbf{R}} u(s\theta + t\theta^\perp)\, \mathrm{d}t,$$

for $\theta \in S^{n-1}$ and $\theta^\perp$ being the vector orthogonal to θ. Hence, the X-ray transform (and therefore also the Radon transform in two dimensions) integrates the function u over lines in $\mathbf{R}^n$, see Fig. 1.1.

Example 1.2. Let $n = 2$. Then S^{n-1} is simply the unit sphere $S^1 = \{\theta \in \mathbf{R}^2 \mid \|\theta\| = 1\}$. We can choose for instance $\theta = (\cos(\varphi), \sin(\varphi))^\top$, for $\varphi \in [0, 2\pi)$, and parametrise the Radon transform in terms of φ and s, i.e.

$$f(\varphi, s) = (Ru)(\varphi, s) = \int_{\mathbf{R}} u(s\cos(\varphi) - t\sin(\varphi), s\sin(\varphi) + t\cos(\varphi))\, \mathrm{d}t. \tag{1.8}$$

Note that—with respect to the origin of the reference coordinate system—φ determines the angle of the line along one wants to integrate, while s is the offset from that line from the centre of the coordinate system. It can be shown that the Radon transform is linear and continuous, i.e. $R \in \mathcal{L}(L^2(B), L^2(Z))$, and even compact.

In **X-ray Computed Tomography (CT)**, the unknown quantity u represents a spatially varying density that is exposed to X-radiation from different angles, and that absorbs the radiation according to its material or biological properties. The basic modelling assumption for the intensity decay of an X-ray beam is that within a small distance Δt it is proportional to the intensity itself, the density, and the distance, i.e.

$$\frac{I(x + (t + \Delta t)\theta) - I(x + t\theta)}{\Delta t} = -I(x + t\theta)u(x + t\theta) + O(\Delta t),$$

[1] Named after the Austrian mathematician Johann Karl August Radon (16 December 1887 – 25 May 1956).

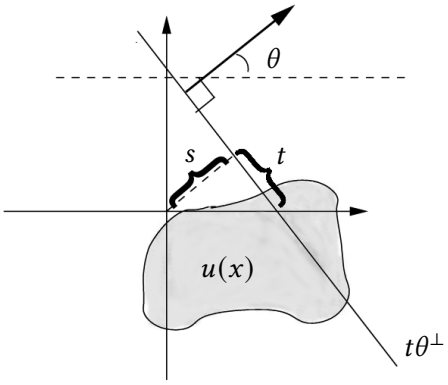

Figure 1.1: Visualisation of the Radon transform in two dimensions[2] (which coincides with the X-ray transform). The function u is integrated over the ray parametrised by θ and s.

for $x \in \theta^\perp$. By taking the limit $\Delta t \to 0$ we end up with the ordinary differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t} I(x + t\theta) = -I(x + t\theta)u(x + t\theta), \tag{1.9}$$

Let $r > 0$ be the radius of the domain of interest centred at the origin. Then, we integrate (1.9) from $t = -\sqrt{r^2 - \|x\|_2^2}$, the position of the emitter, to $t = \sqrt{r^2 - \|x\|_2^2}$, the position of the detector, and obtain

$$\int_{-\sqrt{r^2-\|x\|_2^2}}^{\sqrt{r^2-\|x\|_2^2}} \frac{\frac{\mathrm{d}}{\mathrm{d}t} I(x + t\theta)}{I(x + t\theta)} \,\mathrm{d}t = -\int_{-\sqrt{r^2-\|x\|_2^2}}^{\sqrt{r^2-\|x\|_2^2}} u(x + t\theta) \,\mathrm{d}t.$$

Note that, since $\mathrm{d}/\mathrm{d}x \log(f(x)) = f'(x)/f(x)$, the left hand side in the above equation simplifies to

$$\int_{-\sqrt{r^2-\|x\|_2^2}}^{\sqrt{r^2-\|x\|_2^2}} \frac{\frac{\mathrm{d}}{\mathrm{d}t} I(x + t\theta)}{I(x + t\theta)} \,\mathrm{d}t = \log\left(I\left(x + \sqrt{r^2 - \|x\|_2^2}\theta\right)\right) - \log\left(I\left(x - \sqrt{r^2 - \|x\|_2^2}\theta\right)\right).$$

As we know the radiation intensity at both the emitter and the detector, we therefore know $f(x, \theta) = \log(I(x - \theta\sqrt{r^2 - \|x\|_2^2})) - \log(I(x + \theta\sqrt{r^2 - \|x\|_2^2}))$ and we can write the estimation of the unknown density u as the inverse problem of the X-ray transform (1.8) (if we further assume that u can be continuously extended to zero outside of the circle of radius r).

1.2.6 Groundwater flow/hydraulic tomography

One goal in hydraulic tomography is to estimate the permeability of a groundwater reservoir. The permeability describes the conductivity of the groundwater reservoir and is, e.g., used to estimate the travel time of toxic or radioactive particles in the groundwater. To estimate the permeability, the water pressure in several positions within the reservoir is measured. Pressure head and permeability are linked through Darcy's law and the (assumed) incompressibility of water.

[2] Figure adapted from Wikipedia `https://commons.wikimedia.org/w/index.php?curid=3001440`, by Begemotv2718, CC BY-SA 3.0.

Let $D \subseteq \mathbf{R}^d$ $(d = 1, 2, 3)$ be an open, bounded, connected set with smooth boundary representing the groundwater reservoir. Let $a : \overline{D} \to (0, \infty)$ be a continuously differentiable function representing the permeability and let $s : \overline{D} \to \mathbf{R}$ be a continuous function representing the water sources in the reservoir. Furthermore, assume that the water pressure is 0 outside of D. Darcy's law states that the pressure $p : D \to \mathbf{R}$, the flux $\vec{q} : D \to \mathbf{R}^d$, and the permeability in the reservoir are related as follows:

$$\vec{q}(x) = -a(x)\nabla p(x) \qquad (x \in D).$$

Incompressibility on the other hand requires that the divergence of the flux is fully controlled by in- and outflow given through the source term s:

$$\nabla \cdot \vec{q}(x) = s(x) \qquad (x \in D).$$

Finally, we can combine these assertions and obtain the elliptic partial differential equation

$$\begin{aligned} -\nabla \cdot (a(x)\nabla p(x)) &= s(x) && (x \in D) \\ p(x) &= 0 && (x \in \partial D). \end{aligned}$$

In the described set-up, we now observe the pressure p in several positions $x_1, \ldots, x_I \in D$, e.g., we observe $f_\eta = (p(x_i) : i = 1, \ldots, I) + \eta$. We consider the inverse problem consisting in the estimation of the permeability a using the pressure measurements f_η. Indeed, using noisy point evaluations of the solution of the partial differential equation, we try to estimate its diffusion coefficient. Note that the map $a \mapsto (p(x_i) : i = 1, \ldots, I)$ is non-linear. Hence, this inverse problem is a non-linear inverse problem.

Chapter 2

Generalised Solutions

Functional analysis is the basis of the theory that we will cover in this course. We cannot recall all basic concepts of functional analysis and instead refer to popular textbooks that deal with this subject, e.g., [2, 8, 7]. Nevertheless, we will recall a few important definitions that will be used in this lecture.

We will focus on inverse problems with *bounded linear operators* A, i.e. $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ with

$$\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} := \sup_{u \in \mathcal{X}\setminus\{0\}} \frac{\|Au\|_{\mathcal{Y}}}{\|u\|_{\mathcal{X}}} = \sup_{\|u\|_{\mathcal{X}} \leqslant 1} \|Au\|_{\mathcal{Y}} < \infty.$$

For $A\colon \mathcal{X} \to \mathcal{Y}$ we further denote by

1. $\operatorname{dom}(A) := \mathcal{X}$ the domain of A,
2. $\ker(A) := \{u \in \mathcal{X} \mid Au = 0\}$ the kernel of A,
3. $\operatorname{im}(A) := \{f \in \mathcal{Y} \mid \exists u \in \mathcal{X},\ f = Au\}$ the range of A.

We say that A is continuous at $u \in \mathcal{X}$ if for all $\varepsilon > 0$ there exists $\delta > 0$ with

$$\|Au - Av\|_{\mathcal{Y}} \leqslant \varepsilon \text{ for all } v \in \mathcal{X} \text{ with } \|u - v\|_{\mathcal{X}} \leqslant \delta.$$

For linear K it can be shown that continuity is equivalent to boundedness, i.e. the existence of a constant $C > 0$ such that

$$\|Au\|_{\mathcal{Y}} \leqslant C\|u\|_{\mathcal{X}}$$

for all $u \in \mathcal{X}$. Note that the optimal constant C actually equals the operator norm $\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})}$.

In this Chapter we only consider $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ with $\mathcal{X}$ and $\mathcal{Y}$ being Hilbert spaces. Every Hilbert space $\mathcal{U}$ is equipped with an *inner product*, which we are going to denote by $\langle \cdot, \cdot \rangle_{\mathcal{U}}$ (or simply $\langle \cdot, \cdot \rangle$, whenever the space is clear from the context). In analogy to the transpose of a matrix, this inner product structure together with the theorem of Fréchet-Riesz [8, Section 2.10, Theorem 2.E] allows us to define the *adjoint operator* of A, denoted with $A^* \in \mathcal{L}(\mathcal{Y}, \mathcal{X})$, as the unique solution to the following identity:

$$\langle Au, v \rangle_{\mathcal{Y}} = \langle u, A^* v \rangle_{\mathcal{X}}, \text{ for all } u \in \mathcal{X}, v \in \mathcal{Y}.$$

In addition to this, the inner product is used to define orthogonality. Two elements $u, v \in \mathcal{X}$ are said to be *orthogonal* if $\langle u, v \rangle = 0$. For a subset $\mathcal{X}' \subset \mathcal{X}$ the *orthogonal complement* of $\mathcal{X}'$ in $\mathcal{X}$ is defined as

$$\mathcal{X}'^{\perp} := \left\{ u \in \mathcal{X} \mid \langle u, v \rangle_{\mathcal{X}} = 0 \text{ for all } v \in \mathcal{X}' \right\}.$$

One can show that $\mathcal{X}'^{\perp}$ is a closed subspace and that $\mathcal{X}^{\perp} = \{0\}$. Moreover, we have that $\mathcal{X}' \subset (\mathcal{X}'^{\perp})^{\perp}$. If $\mathcal{X}'$ is a closed subspace then we even have $\mathcal{X}' = (\mathcal{X}'^{\perp})^{\perp}$. In this case, we can give an *orthogonal decomposition* of $\mathcal{X}$:

$$\mathcal{X} = \mathcal{X}' \oplus \mathcal{X}'^{\perp}.$$

By this notation, we mean that every element $u \in \mathcal{X}$ can uniquely be represented as

$$u = x + x^{\perp} \text{ with } x \in \mathcal{X}' \text{ and } x^{\perp} \in \mathcal{X}'^{\perp},$$

see for instance [8, Section 2.9, Corollary 1]. The mapping $u \mapsto x$ defines a linear operator $P_{\mathcal{X}'} \in \mathcal{L}(\mathcal{X}, \mathcal{X})$, which is called the *orthogonal projection* on $\mathcal{X}'$.

Lemma 2.1 (cf. [6, Section 5.16])**.** *Let $\mathcal{X}' \subset \mathcal{X}$ be a closed subspace. The orthogonal projection onto $\mathcal{X}'$, $P_{\mathcal{X}'}$ satisfies the following conditions:*

1. *$P_{\mathcal{X}'}$ is self-adjoint, i.e. $P^*_{\mathcal{X}'} = P_{\mathcal{X}'}$,*
2. *$\|P_{\mathcal{X}'}\|_{\mathcal{L}(\mathcal{X},\mathcal{X})} = 1$ if $\mathcal{X}' \neq \{0\}$,*
3. *$I - P_{\mathcal{X}'} = P_{\mathcal{X}'^{\perp}}$,*
4. *$\|u - P_{\mathcal{X}'}u\|_{\mathcal{X}} \leqslant \|u - v\|_{\mathcal{X}}$ for all $v \in \mathcal{X}'$,*
5. *$x = P_{\mathcal{X}'}u$ if and only if $x \in \mathcal{X}'$ and $u - x \in \mathcal{X}'^{\perp}$.*

Remark 2.2. Note that for a non-closed subspace $\mathcal{X}'$ we only have $(\mathcal{X}'^{\perp})^{\perp} = \overline{\mathcal{X}'}$. For $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ we therefore have

- $\operatorname{im}(A)^{\perp} = \ker(A^*)$ and thus $\ker(A^*)^{\perp} = \overline{\operatorname{im}(A)}$,
- $\operatorname{im}(A^*)^{\perp} = \ker(A)$ and thus $\ker(A)^{\perp} = \overline{\operatorname{im}(A^*)}$.

Hence, we can deduce the following orthogonal decompositions

$$\mathcal{X} = \ker(A) \oplus \overline{\operatorname{im}(A^*)} \text{ and } \mathcal{Y} = \ker(A^*) \oplus \overline{\operatorname{im}(A)}.$$

2.1 Generalised Inverses

Recall the inverse problem

$$Au = f, \tag{2.1}$$

where $A\colon \mathcal{X} \to \mathcal{Y}$ is a linear bounded operator and $\mathcal{X}$ and $\mathcal{Y}$ are Hilbert spaces.

Definition 2.3 (Minimal-norm solutions)**.** *An element $u \in \mathcal{X}$ is called*

- *a least-squares solution of* (2.1) *if*

$$\|Au - f\|_{\mathcal{Y}} = \inf\{\|Av - f\|_{\mathcal{Y}} | v \in \mathcal{X}\},$$

- *a minimum-norm solution of* (2.1) *(and is denoted by $u^{\dagger}$) if it is a least-squares solution, and*

$$\|u^{\dagger}\|_{\mathcal{X}} \leqslant \|v\|_{\mathcal{X}} \quad \text{for all least-squares solutions } v.$$

Remark 2.4. Since $\mathrm{im}(A)$ is not closed in general (it is never closed for a compact operator, unless the range is finite-dimensional), a least-squares solution may not exist. If it exists, then the minimum-norm solution is unique (it is the orthogonal projection of the zero element onto the non-empty closed convex set defined by $\|Au - f\|_{\mathcal{Y}} = \min\{\|Av - f\|_{\mathcal{Y}} | v \in \mathcal{X}\}$).

In numerical linear algebra it is a well known fact that the normal equations can be used to compute least-squares solutions. The same holds true in the infinite-dimensional case.

Theorem 2.5. *Let $f \in \mathcal{Y}$ and $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then, the following three assertions are equivalent for any $u \in \mathcal{X}$.*

1. $Au = P_{\overline{\mathrm{im}(A)}} f$.
2. *u is a least-squares solution of the inverse problem* (2.1).
3. *u solves the* normal equation
$$A^* Au = A^* f. \tag{2.2}$$

Remark 2.6. The name normal equation is derived from the fact that for any solution u its residual $Au - f$ is orthogonal (normal) to $\mathrm{im}(A)$. This can be readily seen, as we have for any $v \in \mathcal{X}$ that

$$0 = \langle v, A^*(Au - f)\rangle_{\mathcal{X}} = \langle Av, Au - f\rangle_{\mathcal{Y}}$$

which shows $Au - f \in \mathrm{im}(A)^\perp$.

Proof of Theorem 2.5. For $1 \Rightarrow 2$: Let $u \in \mathcal{X}$ such that $Au = P_{\overline{\mathrm{im}(A)}} f$ and let $v \in \mathcal{X}$ be arbitrary. With the basic properties of the orthogonal projection, Lemma 2.1, point 4, we have

$$\|Au - f\|_{\mathcal{Y}} = \|P_{\overline{\mathrm{im}(A)}} f - f\|_{\mathcal{Y}} \leqslant \inf_{g \in \overline{\mathrm{im}(A)}} \|g - f\|_{\mathcal{Y}} \leqslant \inf_{g \in \mathrm{im}(A)} \|g - f\|_{\mathcal{Y}} = \inf_{v \in \mathcal{X}} \|Av - f\|_{\mathcal{Y}},$$

which shows that u is a least-squares solution.

For $2 \Rightarrow 3$: Let $u \in \mathcal{X}$ be a least-squares solution and let $v \in \mathcal{X}$ an arbitrary element. We define the quadratic polynomial $F\colon \mathbf{R} \to \mathbf{R}$,

$$F(\lambda) := \|A(u + \lambda v) - f\|_{\mathcal{Y}}^2 = \lambda^2 \|Av\|_{\mathcal{Y}}^2 - 2\lambda \langle Av, f - Au\rangle_{\mathcal{Y}} + \|f - Au\|_{\mathcal{Y}}^2 .$$

A necessary condition for $u \in \mathcal{X}$ to be a least-squares solution is $F'(0) = 0$, which leads to $\langle v, A^*(f - Au)\rangle_{\mathcal{X}} = 0$. As v was arbitrary, it follows that the normal equation (2.2) must hold.

For $3 \Rightarrow 1$: From the normal equation it follows that $A^*(f - Au) = 0$, which is equivalent to $f - Au \in \mathrm{im}(A)^\perp$, see Remark 2.6. Since $\mathrm{im}(A)^\perp = \left(\overline{\mathrm{im}(A)}\right)^\perp$ and $Au \in \mathrm{im}(A) \subset \overline{\mathrm{im}(A)}$, the assertion follows from Lemma 2.1, point 5:

$$Au = P_{\overline{\mathrm{im}(A)}} f \Leftrightarrow Au \in \overline{\mathrm{im}(A)} \text{ and } f - Au \in \left(\overline{\mathrm{im}(A)}\right)^\perp .$$

□

Lemma 2.7. *Let $f \in \mathcal{Y}$ and let $\mathbf{L}$ be the set of least-squares solutions to the inverse problem* (2.1). *Then, $\mathbf{L}$ is non-empty if and only if $f \in \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$.*

Proof. Let $u \in \mathbf{L}$. It is easy to see that $f = Au + (f - Au) \in \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$ as the normal equations are equivalent to $f - Au \in \mathrm{im}(A)^\perp$.

Consider now $f \in \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$. Then there exists $u \in \mathcal{X}$ and $g \in \mathrm{im}(A)^\perp = \left(\overline{\mathrm{im}(A)}\right)^\perp$ such that $f = Au + g$ and thus $P_{\overline{\mathrm{im}(A)}} f = P_{\overline{\mathrm{im}(A)}} Au + P_{\overline{\mathrm{im}(A)}} g = Au$ and the assertion follows from Theorem 2.5, point 1. □

Remark 2.8. If $\mathrm{im}(A)$ is finite-dimensional, then $\mathrm{im}(A)$ is closed, i.e. $\overline{\mathrm{im}(A)} = \mathrm{im}(A)$. Thus, when the measurements are finite-dimensional, there always exists a least-squares solution.

Theorem 2.9. *Let $f \in \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$. Then there exists a unique minimum-norm solution $u^\dagger$ to the inverse problem* (2.1) *and all least-squares solutions are given by $\{u^\dagger\} + \ker(A)$.*

Proof. From Lemma 2.7 we know that there exists a least-squares solution. As noted in Remark 2.4, in this case the minimum-norm solution is unique. Let φ be an arbitrary least-squares solution. Using Theorem 2.5 we get

$$A(\varphi - u^\dagger) = A\varphi - Au^\dagger = P_{\overline{\mathrm{im}(A)}} f - P_{\overline{\mathrm{im}(A)}} f = 0, \tag{2.3}$$

which shows that $\varphi - u^\dagger \in \ker(A)$, hence the assertion. □

If a least-squares solution exists for a given $f \in \mathcal{Y}$, then the minimum-norm solution can be computed (at least in theory) using the Moore–Penrose generalised inverse.

Definition 2.10. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ and let*

$$\widetilde{A} \coloneqq A|_{\ker(A)^\perp} : \ker(A)^\perp \to \mathrm{im}(A)$$

denote the restriction of A to $\ker(A)^\perp$. The Moore–Penrose inverse *$A^\dagger$ is defined as the unique linear extension of $\widetilde{A}^{-1}$ to*

$$\mathrm{dom}(A^\dagger) = \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$$

with

$$\ker(A^\dagger) = \mathrm{im}(A)^\perp.$$

Remark 2.11. Due to the restriction to $\ker(A)^\perp$ and $\mathrm{im}(A)$ we have that $\widetilde{A}$ is injective and surjective. Hence, $\widetilde{A}^{-1}$ exists and is linear and — as a consequence — $A^\dagger$ is well-defined on $\mathrm{im}(A)$. Moreover, due to the orthogonal decomposition $\mathrm{dom}(A^\dagger) = \mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$, there exist for arbitrary $f \in \mathrm{dom}(A^\dagger)$ elements $f_1 \in \mathrm{im}(A)$ and $f_2 \in \mathrm{im}(A)^\perp$ with $f = f_1 + f_2$. Therefore, we have

$$A^\dagger f = A^\dagger f_1 + A^\dagger f_2 = A^\dagger f_1 = \widetilde{A}^{-1} f_1 = \widetilde{A}^{-1} P_{\overline{\mathrm{im}(A)}} f\,, \tag{2.4}$$

where we have used that $f_2 \in \mathrm{im}(A)^\perp = \ker(A^\dagger)$. Thus, $A^\dagger$ is well-defined on the entire domain $\mathrm{dom}(A^\dagger)$.

Remark 2.12. As orthogonal complements are always closed we get that

$$\overline{\mathrm{dom}(A^\dagger)} = \overline{\mathrm{im}(A)} \oplus \mathrm{im}(A)^\perp = \mathcal{Y},$$

and hence, $\mathrm{dom}(A^\dagger)$ is dense in $\mathcal{Y}$. Thus, if $\mathrm{im}(A)$ is closed it follows that $\mathrm{dom}(A^\dagger) = \mathcal{Y}$ and on the other hand, $\mathrm{dom}(A^\dagger) = \mathcal{Y}$ implies $\mathrm{im}(A)$ is closed. We note that for ill-posed problems $\mathrm{im}(A)$ is usually not closed; for instance, if A is compact then $\mathrm{im}(A)$ is closed if and only if it is finite-dimensional [1, Ex.1 Section 7.1].

If A is bijective, we have that $A^\dagger = A^{-1}$. We also highlight that the extension $A^\dagger$ is not necessarily continuous.

Example 2.13. To illustrate the definition of the Moore–Penrose inverse we consider a simple example in finite dimensions. Let the linear operator $A\colon \mathbf{R}^3 \to \mathbf{R}^2$ be given by

$$Ax = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 2x_1 \\ 0 \end{pmatrix}.$$

It is easy to see that $\operatorname{im}(A) = \{f \in \mathbf{R}^2 \mid f_2 = 0\}$ and $\ker(A) = \{x \in \mathbf{R}^3 \mid x_1 = 0\}$. Thus, $\ker(A)^\perp = \{x \in \mathbf{R}^3 \mid x_2, x_3 = 0\}$. Therefore, $\widetilde{A}\colon \ker(A)^\perp \to \operatorname{im}(A)$, given by $x \mapsto (2x_1, 0)^\top$, is bijective and its inverse $\widetilde{A}^{-1}\colon \operatorname{im}(A) \to \ker(A)^\perp$ is given by $f \mapsto (f_1/2, 0, 0)^\top$.

To get the Moore–Penrose inverse $A^\dagger$, we need to extend $\widetilde{A}^{-1}$ to $\operatorname{im}(A) \oplus \operatorname{im}(A)^\perp$ in such a way that $A^\dagger f = 0$ for all $f \in \operatorname{im}(A)^\perp = \{f \in \mathbf{R}^2 \mid f_1 = 0\}$. It is easy to see that the Moore–Penrose inverse $A^\dagger\colon \mathbf{R}^2 \to \mathbf{R}^3$ is given by the following expression

$$A^\dagger f = \begin{pmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \end{pmatrix} = \begin{pmatrix} f_1/2 \\ 0 \\ 0 \end{pmatrix}.$$

Let us consider data $\widetilde{f} = (8, 1)^\top \notin \operatorname{im}(A)$. Then, $A^\dagger \widetilde{f} = A^\dagger (8, 1)^\top = (4, 0, 0)^\top$.

Let us show that $A^\dagger$ can be characterised by the Moore–Penrose equations:

Theorem 2.14 ([3, Prop. 2.3])**.** *The Moore–Penrose inverse $A^\dagger$ satisfies $\operatorname{im}(A^\dagger) = \ker(A)^\perp$ and the Moore–Penrose equations*

1. $A^\dagger A = P_{\ker(A)^\perp}$,
2. $AA^\dagger = P_{\overline{\operatorname{im}(A)}}\big|_{\operatorname{dom}(A^\dagger)}$,
3. $AA^\dagger A = A$,
4. $A^\dagger A A^\dagger = A^\dagger$,

where $P_{\ker(A)}$ and $P_{\overline{\operatorname{im}(A)}}$ denote the orthogonal projections onto $\ker(A)$ and $\overline{\operatorname{im}(A)}$, respectively.

Proof. First, by the definition of the Moore–Penrose inverse we have for any $u \in \mathcal{X}$

$$A^\dagger A u = A^\dagger A (P_{\ker(A)} u + P_{\ker(A)^\perp} u) = A^\dagger A P_{\ker(A)^\perp} u = \widetilde{A}^{-1} A P_{\ker(A)^\perp} u = P_{\ker(A)^\perp} u,$$

which proves 1. Now, for any $f \in \operatorname{dom}(A^\dagger)$ we have (see (2.4))

$$AA^\dagger f = A\widetilde{A}^{-1} P_{\overline{\operatorname{im}(A)}} f = P_{\overline{\operatorname{im}(A)}} f,$$

which proves 2. Applying A to 1., we get 3., and applying $A^\dagger$ to 2., we get 4., which completes the proof. □

Corollary 2.15. The Moore–Penrose inverse is uniquely characterised by points 1 and 2 of Theorem 2.14. That is, if a linear operator $B\colon \operatorname{im}(A) \oplus \operatorname{im}(A)^\perp \to \ker(A)^\perp$ satisfies $BA = P_{\ker(A)^\perp}$ and $AB = P_{\overline{\operatorname{im}(A)}}$ then $B = A^\dagger$.

Proof. First we show that $B|_{\mathrm{im}(A)} = \widetilde{A}^{-1}$. Indeed, let $f = Au \in \mathrm{im}(A)$, where $u \in \ker(A)^\perp$. Then

$$Bf = BAu = P_{\ker(A)^\perp} u = u = \widetilde{A}^{-1} f,$$

where the last equality holds since $\widetilde{A}$ is bijective and hence uniquely invertible.

Now we prove that $B|_{\mathrm{im}(A)^\perp} = 0$. Indeed, for any $f \in \mathrm{im}(A)^\perp$ we have

$$ABf = P_{\overline{\mathrm{im}(A)}} f = 0.$$

Since $Bf \in \ker(A)^\perp$ and A is injective on $\ker(A)^\perp$, we conclude that $Bf = 0$. Therefore, B is an extension of $\widetilde{A}^{-1}$ to $\mathrm{im}(A) \oplus \mathrm{im}(A)^\perp$ with $\ker(B) = \mathrm{im}(A)^\perp$. Since such an extension is unique, $B = A^\dagger$. □

Remark 2.16. If an operator B satisfies only $ABA = A$ (resp. $BAB = B$), it is called the *inner inverse* (resp. *outer inverse*) of A.

Theorem 2.17 ([3, Prop. 2.4])**.** *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then $A^\dagger$ is continuous, i.e. $A^\dagger \in \mathcal{L}(\mathrm{dom}(A^\dagger), \mathcal{X})$, if and only if $\mathrm{im}(A)$ is closed.*

Remark 2.18. Theorem 2.17 is a relatively straightforward consequence of the *closed graph theorem*: using the Moore-Penrose equations (Theorem 2.14) it can be shown that the graph of $A^\dagger$, $\{(f, A^\dagger f) | f \in \mathrm{dom}(A^\dagger)\} \subset \mathrm{dom}(A^\dagger) \times \ker(A)^\perp$ is a closed set.

The next theorem shows that minimum-norm solutions can indeed be computed using the Moore–Penrose generalised inverse.

Theorem 2.19. *For each $f \in \mathrm{dom}(A^\dagger)$, the minimum-norm solution $u^\dagger$ to the inverse problem* (2.1) *is given via*

$$u^\dagger = A^\dagger f.$$

Proof. As $f \in \mathrm{dom}(A^\dagger)$, we know from Theorem 2.9 that the minimum-norm solution $u^\dagger$ exists and is unique. With $u^\dagger \in \ker(A)^\perp$, Lemma 2.14, and Theorem 2.5 we conclude that

$$u^\dagger = P_{\ker(A)^\perp} u^\dagger = A^\dagger A u^\dagger = A^\dagger P_{\overline{\mathrm{im}(A)}} f = A^\dagger A A^\dagger f = A^\dagger f.$$

□

As a consequence of Theorem 2.19 and Theorem 2.5, we find that the minimum-norm solution $u^\dagger$ of $Au = f$ is a minimum-norm solution of the normal equation (2.2), i.e.

$$u^\dagger = (A^* A)^\dagger A^* f.$$

Thus, in order to compute $u^\dagger$ we can equivalently consider finding the minimum-norm solution of the normal equation.

2.2 Compact Operators

Definition 2.20. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then A is said to be* compact *if for any bounded set $B \subset \mathcal{X}$ the closure of its image $\overline{A(B)}$ is compact in $\mathcal{Y}$. We denote the space of compact operators by $\mathcal{K}(\mathcal{X}, \mathcal{Y})$.*

Remark 2.21. We can equivalently define an operator A to be compact if the image of a bounded sequence $\{u_j\}_{j \in \mathbb{N}} \subset \mathcal{X}$ contains a convergent subsequence $\{Au_{j_k}\}_{k \in \mathbb{N}} \subset \mathcal{Y}$.

Compact operators are very common in inverse problems. In fact, almost all (linear) inverse problems involve the inversion of a compact operator. As the following result shows, compactness of the forward operator is a major source of ill-posedness.

Theorem 2.22. *Let $A \in \mathcal{K}(\mathcal{X}, \mathcal{Y})$ have an infinite-dimensional range. Then, the Moore–Penrose inverse of A is discontinuous.*

Proof. As the range $\operatorname{im}(A)$ is of infinite dimension, we can conclude that $\mathcal{X}$ and $\ker(A)^\perp$ are also infinite dimensional. We can therefore find a sequence $\{u_j\}_{j\in\mathbf{N}}$ with $u_j \in \ker(A)^\perp$, $\|u_j\|_\mathcal{X} = 1$ and $\langle u_j, u_k\rangle_\mathcal{X} = 0$ for $j \neq k$. Since A is a compact operator the sequence $f_j = Au_j$ has a convergent subsequence, hence, for all $\delta > 0$ we can find j, k such that $\|f_j - f_k\|_\mathcal{Y} < \delta$. However, we also obtain

$$\begin{aligned}\|A^\dagger f_j - A^\dagger f_k\|_\mathcal{X}^2 &= \|A^\dagger A u_j - A^\dagger A u_k\|_\mathcal{X}^2 \\ &= \|u_j - u_k\|_\mathcal{X}^2 = \|u_j\|_\mathcal{X}^2 - 2\langle u_j, u_k\rangle_\mathcal{X} + \|u_k\|_\mathcal{X}^2 = 2,\end{aligned}$$

which shows that $A^\dagger$ is discontinuous. Here, the second identity follows from Lemma 2.14, point 1, and the fact that $u_j, u_k \in \ker(A)^\perp$. □

To have a better understanding of when we have $f \in \overline{\operatorname{im}(A)} \setminus \operatorname{im}(A)$ for compact operators A, we will consider the singular value decomposition of compact operators.

Singular value decomposition of compact operators

Theorem 2.23 ([5, p. 225, Theorem 9.16])**.** *Let $\mathcal{X}$ be a Hilbert space and $A \in \mathcal{K}(\mathcal{X}, \mathcal{X})$ be self-adjoint. Then there exists an orthonormal basis $\{x_j\}_{j\in\mathbf{N}} \subset \mathcal{X}$ of $\overline{\operatorname{im}(A)}$ and a sequence of eigenvalues $\{\lambda_j\}_{j\in\mathbf{N}} \subset \mathbf{R}$ with $|\lambda_1| \geqslant |\lambda_2| \geqslant \ldots > 0$ such that for all $u \in \mathcal{X}$ we have*

$$Au = \sum_{j=1}^{\infty} \lambda_j \langle u, x_j\rangle_\mathcal{X} x_j\,.$$

The sequence $\{\lambda_j\}_{j\in\mathbf{N}}$ is either finite or we have $\lambda_j \to 0$.

Remark 2.24. The notation in the theorem above only makes sense if the sequence $\{\lambda_j\}_{j\in\mathbf{N}}$ is infinite. For the case that there are only finitely many λ_j the sum has to be interpreted as a finite sum. Moreover, as the eigenvalues are sorted by absolute value $|\lambda_j|$, we have $\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{X})} = |\lambda_1|$.

If A is not self-adjoint, the decomposition in Theorem 2.23 does not hold any more. Instead, we can consider the so-called *singular value decomposition* of a compact linear operator. To prove its existence, we will use the following lemma on the relationship between the ranges of A^* and A^*A:

Lemma 2.25. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then $\overline{\operatorname{im}(A^*A)} = \overline{\operatorname{im}(A^*)}$.*

Proof. It is clear that $\overline{\operatorname{im}(A^*A)} = \overline{\operatorname{im}(A^*|_{\operatorname{im}(A)})} \subseteq \overline{\operatorname{im}(A^*)}$, so we are left to prove that $\overline{\operatorname{im}(A^*)} \subseteq \overline{\operatorname{im}(A^*A)}$.

Let $u \in \overline{\operatorname{im}(A^*)}$ and let $\varepsilon > 0$. Then, there exists $f \in \ker(A^*)^\perp = \overline{\operatorname{im}(A)}$ with $\|A^*f - u\|_\mathcal{X} < \varepsilon/2$ (recall the orthogonal decomposition in Remark 2.2). As $\ker(A^*)^\perp = \overline{\operatorname{im}(A)}$, there exists $x \in \mathcal{X}$ such that $\|Ax - f\|_\mathcal{Y} < \varepsilon/(2\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})})$. Putting these together we have

$$\begin{aligned}\|A^*Ax - u\|_\mathcal{X} &\leqslant \|A^*Ax - A^*f\|_\mathcal{X} + \|A^*f - u\|_\mathcal{X} \\ &\leqslant \underbrace{\|A^*\|_{\mathcal{L}(\mathcal{Y},\mathcal{X})}\|Ax - f\|_\mathcal{Y}}_{<\varepsilon/2} + \underbrace{\|A^*f - u\|_\mathcal{X}}_{<\varepsilon/2} < \varepsilon\end{aligned}$$

which shows that $u \in \overline{\operatorname{im}(A^*A)}$ and thus also $\overline{\operatorname{im}(A^*)} \subseteq \overline{\operatorname{im}(A^*A)}$. □

Theorem 2.26. *Let $A \in \mathcal{K}(\mathcal{X}, \mathcal{Y})$. Then there exists*

1. *a not-necessarily infinite null sequence $\{\sigma_j\}_{j\in\mathbf{N}}$ with $\sigma_1 \geqslant \sigma_2 \geqslant \ldots > 0$,*
2. *an orthonormal basis $\{x_j\}_{j\in\mathbf{N}} \subset \mathcal{X}$ of $\ker(A)^\perp$,*
3. *an orthonormal basis $\{y_j\}_{j\in\mathbf{N}} \subset \mathcal{Y}$ of $\overline{\operatorname{im}(A)}$ with*

$$Ax_j = \sigma_j y_j, \quad A^* y_j = \sigma_j x_j, \quad \text{for all } j \in \mathbf{N}. \tag{2.5}$$

Moreover, for all $u \in \mathcal{X}$ we have the representation

$$Au = \sum_{j=1}^{\infty} \sigma_j \left\langle u, x_j \right\rangle \, y_j. \tag{2.6}$$

The sequence $\{(\sigma_j, x_j, y_j)\}$ is called singular system *or* singular value decomposition (SVD) *of A. For the adjoint operator A^* we have the representation*

$$A^* f = \sum_{j=1}^{\infty} \sigma_j \left\langle f, y_j \right\rangle \, x_j \quad \forall f \in \mathcal{Y}. \tag{2.7}$$

Proof. Consider $B = A^*A$ and $C = AA^*$. Both B and C are compact, self-adjoint and positive semidefinite, so that by Theorem 2.23 both admit a spectral representation and, by positive semidefiniteness, their eigenvalues are positive. Therefore, we can write

$$Cf = \sum_{j=1}^{\infty} \sigma_j^2 \left\langle f, y_j \right\rangle y_j \quad \forall f \in \mathcal{Y},$$

where $\{y_j\}$ is an orthonormal basis of $\overline{\operatorname{im}(AA^*)} = \overline{\operatorname{im}(A)}$ (Lemma 2.25), $\sigma_j > 0$ for all j and $\sigma_j \to 0$ as $j \to \infty$. Now consider the element $A^* y_j \in \mathcal{X}$. Since σ_j^2 is an eigenvalue of C for the eigenvector y_j, we get that

$$\sigma_j^2 A^* y_j = A^*(\sigma_j^2 y_j) = A^* C y_j = A^* A A^* y_j = B A^* y_j$$

and therefore σ_j^2 is also an eigenvalue of B (for the eigenvector $A^* y_j$). Now we will show that the system $\{A^* y_j/\sigma_j\}_{j\in\mathbf{N}}$ forms an orthonormal basis of $\overline{\operatorname{im}(A^*)} = \ker(A)^\perp$. Indeed, we have

$$\left\langle \frac{A^* y_j}{\sigma_j}, \frac{A^* y_k}{\sigma_k} \right\rangle = \frac{1}{\sigma_j \sigma_k} \left\langle y_j, AA^* y_k \right\rangle = \frac{1}{\sigma_j \sigma_k} \left\langle y_j, \sigma_k^2 y_k \right\rangle = \begin{cases} 1, & \text{if } j = k, \\ 0, & \text{otherwise.} \end{cases}$$

Hence, $\{A^* y_j/\sigma_j\}_{j\in\mathbf{N}}$ are orthonormal. It is also clear that they are dense in $\overline{\operatorname{im}(A^*)} = \ker(A)^\perp$, hence they form a basis. Therefore, we can choose $\{x_j\}_{j\in\mathbf{N}} = \{A^* y_j/\sigma_j\}_{j\in\mathbf{N}}$, i.e.

$$x_j = \sigma_j^{-1} A^* y_j$$

and we get (by construction) that

$$A^* y_j = \sigma_j x_j.$$

We also observe that

$$Ax_j = \sigma_j^{-1} AA^* y_j = \sigma_j^{-1} \sigma_j^2 y_j = \sigma_j y_j,$$

which proves (2.5). Extending the basis $\{x_j\}$ of $\overline{\operatorname{im}(A^*)}$ to a basis $\{\widetilde{x}_j\}$ of $\mathcal{X}$, we expand an arbitrary $u \in \mathcal{X}$ as $u = \sum_{j=1}^{\infty} \langle u, \widetilde{x}_j \rangle \widetilde{x}_j$. Applying A and using the fact that $\mathcal{X} = \ker(A) \oplus \overline{\operatorname{im}(A^*)}$ (Remark 2.2), we obtain the singular value decomposition (2.6) (and also (2.7) in a similar manner)

$$Au = \sum_{j=1}^{\infty} \sigma_j \langle u, x_j \rangle y_j \quad \forall u \in \mathcal{X}, \quad A^* f = \sum_{j=1}^{\infty} \sigma_j \langle f, y_j \rangle x_j \quad \forall f \in \mathcal{Y}.$$

□

We can now derive a representation of the Moore–Penrose inverse in terms of the singular value decomposition.

Theorem 2.27. *Let $A \in \mathcal{K}(\mathcal{X}, \mathcal{Y})$ with singular system $\{(\sigma_j, x_j, y_j)\}_{j \in \mathbb{N}}$ and $f \in \operatorname{dom}(A^\dagger)$. Then the Moore–Penrose inverse of A can be written as*

$$A^\dagger f = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle x_j \,. \tag{2.8}$$

Proof. We know that, since $f \in \operatorname{dom}(A^\dagger)$, $u^\dagger = A^\dagger f$ solves the normal equations

$$A^* A u^\dagger = A^* f.$$

From Theorem 2.26 we know that

$$A^* A u^\dagger = \sum_{j=1}^{\infty} \sigma_j^2 \langle u^\dagger, x_j \rangle x_j, \quad A^* f = \sum_{j=1}^{\infty} \sigma_j \langle f, y_j \rangle x_j, \tag{2.9}$$

which implies that

$$\langle u^\dagger, x_j \rangle = \sigma_j^{-1} \langle f, y_j \rangle$$

Expanding $u^\dagger \in \ker(A)^\perp$ in the basis $\{x_j\}$, we get

$$u^\dagger = \sum_{j=1}^{\infty} \langle u^\dagger, x_j \rangle x_j = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle x_j = A^\dagger f.$$

□

The representation (2.8) makes it clear again that the Moore–Penrose inverse is unbounded if $\operatorname{im}(A)$ is infinite dimensional. Indeed, taking the sequence y_j we note that $\|A^\dagger y_j\| = \sigma_j^{-1} \to \infty$, although $\|y_j\| = 1$. The unboundedness of the Moore–Penrose inverse is also reflected in the fact that the series in (2.8) may not converge for a given f. The convergence criterion for the series is called the *Picard criterion*:

Definition 2.28. *We say that the data f satisfy the Picard criterion, if*

$$\|A^\dagger f\|^2 = \sum_{j=1}^{\infty} \frac{\left|\langle f, y_j \rangle\right|^2}{\sigma_j^2} < \infty. \tag{2.10}$$

Remark 2.29. The Picard criterion is a condition on the decay of the coefficients $\langle f, y_j \rangle$. As the singular values σ_j decay to zero as $j \to \infty$, the Picard criterion is only met if the coefficients $\langle f, y_j \rangle$ decay sufficiently fast. In case the singular system is given by the Fourier basis, then the coefficients $\langle f, y_j \rangle$ are just the Fourier coefficients of f. Therefore, the Picard criterion is a condition on the decay of the Fourier coefficients which is equivalent to the smoothness of f.

The Picard criterion gives us another way to characterise elements in the range of the forward operator:

Theorem 2.30. *Let $A \in \mathcal{K}(\mathcal{X}, \mathcal{Y})$ with singular system $\{(\sigma_j, x_j, y_j)\}_{j\in\mathbf{N}}$, and $f \in \overline{\operatorname{im}(A)}$. Then $f \in \operatorname{im}(A)$ if and only if the* Picard criterion

$$\sum_{j=1}^{\infty} \frac{|\langle f, y_j\rangle_{\mathcal{Y}}|^2}{\sigma_j^2} < \infty$$

is met.

Proof. Let $f \in \operatorname{im}(A)$, so that there is a $u \in \mathcal{X}$ such that $Au = f$. It is easy to see that we have

$$\langle f, y_j\rangle_{\mathcal{Y}} = \langle Au, y_j\rangle_{\mathcal{Y}} = \langle u, A^* y_j\rangle_{\mathcal{X}} = \sigma_j \langle u, x_j\rangle_{\mathcal{X}}$$

and therefore

$$\sum_{j=1}^{\infty} \sigma_j^{-2} |\langle f, y_j\rangle_{\mathcal{Y}}|^2 = \sum_{j=1}^{\infty} |\langle u, x_j\rangle_{\mathcal{X}}|^2 \leqslant \|u\|_{\mathcal{X}}^2 < \infty\,.$$

Now let the Picard criterion (2.10) hold and define $u := \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j\rangle_{\mathcal{Y}} x_j \in \mathcal{X}$. It is well-defined by the Picard criterion (2.10) and we conclude

$$Au = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j\rangle_{\mathcal{Y}} Ax_j = \sum_{j=1}^{\infty} \langle f, y_j\rangle_{\mathcal{Y}} y_j = P_{\overline{\operatorname{im}(A)}} f = f\,,$$

which shows that $f \in \operatorname{im}(A)$. □

Although all ill-posed problems are not easy to solve, some are worse than others, depending on how fast the singular values decay to zero.

Definition 2.31. *We say that an ill-posed inverse problem* (2.1) *is* mildly ill-posed *if the singular values decay at most with polynomial speed, i.e. there exist $\gamma, C > 0$ such that $\sigma_j \geqslant C j^{-\gamma}$ for all j. We call the ill-posed inverse problem* severely ill-posed *if its singular values decay faster than with polynomial speed, i.e. for all $\gamma, C > 0$ one has that $\sigma_j \leqslant C j^{-\gamma}$ for j sufficiently large.*

Example 2.32. Let us consider the example of differentiation again, as introduced in Section 1.2.3. The forward operator $A\colon L^2([0,1]) \to L^2([0,1])$ in this problem is given by

$$(Au)(t) = \int_0^t u(s)\,\mathrm{d}s = \int_0^1 K(s,t)u(s)\,\mathrm{d}s,$$

with $K\colon [0,1]\times[0,1] \to \mathbf{R}$ defined as

$$K(s,t) := \begin{cases} 1 & s \leqslant t \\ 0 & \text{else} \end{cases}.$$

This is a special case of the integral operators as introduced in Section 1.2.1. Since the kernel K is square integrable, A is compact. The adjoint operator A^* is given by

$$(A^* f)(s) = \int_0^1 K(t,s) f(t)\,\mathrm{d}t = \int_s^1 v(t)\,\mathrm{d}t. \tag{2.11}$$

Now we want to compute the eigenvalues and eigenvectors of A^*A, i.e. we look for σ^2 and $x \in L^2([0,1])$ with

$$\sigma^2 x(s) = (A^*Ax)(s) = \int_s^1 \int_0^t x(r)\,\mathrm{d}r\,\mathrm{d}t.$$

We immediately observe that $x(1) = 0$ and further

$$\sigma^2 x'(s) = \frac{\mathrm{d}}{\mathrm{d}s} \int_s^1 \int_0^t x(r)\,\mathrm{d}r\,\mathrm{d}t = -\int_0^s x(r)\,\mathrm{d}r,$$

from which we conclude $x'(0) = 0$. Taking the derivative another time yields the ordinary differential equation

$$\sigma^2 x''(s) + x(s) = 0,$$

for which solutions are of the form

$$x(s) = c_1 \sin(\sigma^{-1}s) + c_2 \cos(\sigma^{-1}s),$$

with some constants c_1, c_2. In order to satisfy the boundary conditions $x(1) = c_1 \sin(\sigma^{-1}) + c_2 \cos(\sigma^{-1}) = 0$ and $x'(0) = c_1/\sigma = 0$, we choose $c_1 = 0$ and σ such that $\cos(\sigma^{-1}) = 0$. Hence, we have

$$\sigma_j = \frac{2}{(2j-1)\pi} \text{ for } j \in \mathbf{N},$$

and by choosing $c_2 = \sqrt{2}$ we obtain the following normalised representation of x_j:

$$x_j(s) = \sqrt{2} \cos\left(\left(j - \frac{1}{2}\right)\pi s\right).$$

According to (2.5) we further obtain

$$y_j(s) = \sigma_j^{-1}(Ax_j)(s) = \left(j - \frac{1}{2}\right)\pi \int_0^s \sqrt{2} \cos\left(\left(j - \frac{1}{2}\right)\pi t\right) \mathrm{d}t = \sqrt{2} \sin\left(\left(j - \frac{1}{2}\right)\pi s\right),$$

and hence, for $f \in L^2([0,1])$ the Picard criterion becomes

$$2\sum_{j=1}^{\infty} \sigma_j^{-2} \left(\int_0^1 f(s) \sin\left(\sigma_j^{-1}s\right)\,\mathrm{d}s\right)^2 < \infty.$$

Expanding f in the basis $\{y_j\}$

$$f(t) = 2\sum_{j=1}^{\infty} \left(\int_0^1 f(s) \sin\left(\sigma_j^{-1}s\right)\,\mathrm{d}s\right) \sin\left(\sigma_j^{-1}t\right),$$

and formally differentiating the series, we obtain

$$f'(t) = 2\sum_{j=1}^{\infty} \sigma_j^{-1} \left(\int_0^1 f(s) \sin\left(\sigma_j^{-1}s\right)\,\mathrm{d}s\right) \cos\left(\sigma_j^{-1}t\right).$$

Therefore, the Picard criterion is nothing but the condition for the legitimacy of such differentiation, i.e. for the differentiability of the Fourier series by differentiating its components, and it holds if f is differentiable and $f' \in L^2([0,1])$.From the decay of the singular values we see that this inverse problem is mildly ill posed.

Example 2.33 (Heat equation)**.** Consider the problem of recovering the initial condition u of the heat equation from an observation f of the solution at some time $T > 0$ (see Section 1.2.2). We consider the heat equation on $(0, \pi) \times \mathbf{R}_+$, with Dirichlet boundary conditions

$$\begin{cases} v_t - v_{xx} = 0 & \text{on } (0, \pi) \times \mathbf{R}_+, \\ v(0, t) = v(\pi, t) = 0 & \text{on } \mathbf{R}_+, \\ v(x, T) = f(x) & \text{on } (0, \pi), \\ v(x, 0) = u(x) & \text{on } (0, \pi). \end{cases}$$

The solution to the forward problem (determine f given u) is given by

$$f = Au := \sum_{j=1}^{\infty} \exp(-j^2 T) \widehat{u}_j \sin(jx),$$

where $\widehat{u}_j = \langle u, \sin(j\cdot) \rangle$ are Fourier coefficients of u. Hence, singular values of A are given by

$$\sigma_j = \exp(-j^2 T), \quad j \in \mathbf{N},$$

and

$$\frac{1}{\sigma_j} = \exp(j^2 T).$$

The singular values of A decay exponentially and the inverse problem is severely (exponentially) ill-posed.

Bibliography

[1] Y. A. Abramovich and C. D. Aliprantis, *An Invitation to Operator Theory*, Graduate Studies in Mathematics, American Mathematical Society, 2002.

[2] B. Bollobás, *Linear Analysis: An Introductory Course*, Cambridge University Press, Cambridge, second ed., 1999.

[3] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of inverse problems*, vol. 375, Springer Science & Business Media, 1996.

[4] J. Hadamard, *Sur les problèmes aux dérivés partielles et leur signification physique*, Princeton University Bulletin, 13 (1902), pp. 49–52.

[5] J. Hunter and B. Nachtergaele, *Applied Analysis*, World Scientific Publishing Company Incorporated, 2001.

[6] A. W. Naylor and G. R. Sell, *Linear Operator Theory in Engineering and Science*, Springer Science & Business Media, 2000.

[7] K. Saxe, *Beginning Functional Analysis*, Springer, 2002.

[8] E. Zeidler, *Applied Functional Analysis: Applications to Mathematical Physics*, vol. 108 of Applied Mathematical Sciences Series, Springer, 1995.