Notes for Inverse Problems
University of Cambridge, Lent 2022

Aditya Sengupta

June 13, 2022

Contents

Inverse Problems **Lent 2022**

Lecture 1: Introduction

Lecturer: Malena Sabaté Landman *20 January* *Aditya Sengupta*

Note: *LaTeX format adapted from template for lecture notes from CS 267, Applications of Parallel Computing, UC Berkeley EECS department.*

In an inverse problem, we have some object that goes through a forward process to create some data, and the inverse problem consists of reversing that process to recover that object. In this class, we'll talk mostly about continuous inverse problems, so we'll be using tools from functional analysis, but we'll use a lot of discrete examples. Discrete inverse problems make use of the tools of numerical linear algebra.

Learning goals:

1. What are inverse problems?
2. Why are they important?
3. Why are they challenging to solve?
4. What are their important properties?

We'll stick to classic approaches – no Bayesian solvers, for example, but we can email for requests for extra reviews. There's also no machine learning.

1.1 Well-posed and ill-posed problems

Consider a problem of the form $f = Au$, where A is an operator between Banach/Hilbert spaces: $A : \mathcal{X} \to \mathcal{Y}$. This is too simplistic of a model, so instead we consider

$$f_n = Au + n, \tag{1.1}$$

where n is noise with some possibly-not-nice characteristics. We'll see that this is actually applied mathematics: a lot of the not-niceness of inverse problems is because these noise terms aren't nice.

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

1. *There exists some solution (existence)*
2. *There exists at most one solution (uniqueness)*
3. *The solution depends continuously on the data (stability)*

The general version of the discrete problem $f_n = Au + n$ has the form $f_n \in \mathbb{R}^k, A \in \mathbb{R}^{d\times k}, u \in \mathbb{R}^d$.

Ill-posed discrete (linear algebraic) problems can therefore be characterised: you can construct a problem where $d < k$ and no solution exists. You can construct a problem where $d > k$ and there are many solutions.

We can't deal with the third condition here because linear systems are formally well-posed. This is because the error in the solution is bounded by a constant times the error in the right-hand side. These problems are still interesting, and there are ways to express instabilities even in discrete problems, but they're not our primary focus. The issue here comes in with the constant multiplier, which isn't bounded and scales with the condition number of the matrix. The problem here is said to be *ill-conditioned*, not strictly ill-posed, but for the purpose of an introduction the distinction isn't that useful.

Let's say $u, f_n \in \mathbb{R}^n$ and A is symmetric positive definite in $\mathbb{R}^{n\times n}$, meaning it admits an eigendecomposition. We could generalise this with the SVD, but let's go with this for simplicity for now.

$$A = V\Lambda V^{\intercal} = \sum_{j=1}^{n} \lambda_j v_j v_j^{\intercal}, \tag{1.2}$$

where these are related by $\Lambda = diag(\lambda_1, \ldots, \lambda_n)$ and WLOG $\lambda_1 \geq \lambda_2 \cdots \geq \lambda_n > 0$.

Say we have the noisy and noiseless versions of a problem:

$$Au_n = f_n = f + n \tag{1.3}$$
$$Au = f \tag{1.4}$$

Naively, the inverse noisy problem can be solved by

$$u_n = A^{-1}f + A^{-1}n = u + A^{-1}n. \tag{1.5}$$

If $A^{-1}n$ is well behaved, we're fine, but it often isn't. We can look at the difference between u and u_n in terms of the eigendecomposition of A:

$$u - u_n = \sum_{j=1}^{n} \lambda_j^{-1} v_j v_j^{\intercal}(f - f_n) \tag{1.6}$$

$$\begin{aligned} \|u - u_n\|_2^2 &= \sum_{j=1}^{n} \lambda_j^{-2} \|v_j\|_2^2 \|v_j^{\intercal}(f - f_n)\| \\ &\leq \lambda_n^{-2} \|f - f_n\|_2^2 \\ &\leq \frac{\lambda_1^2}{\lambda_n^2}\delta^2 \\ &= \kappa^2\delta^2. \end{aligned} \tag{1.7}$$

where we use the fact that $\|f - f_n\| \leq \lambda_1 \cdot \delta = \|A\|\delta$. This gives us a nice characterisation of a problem stability just in terms of its eigenvalues, using its condition number,

$$\kappa = \frac{\lambda_1}{\lambda_n} = \frac{\text{largest eigenvalue of } A}{\text{smallest eigenvalue of } A} \tag{1.8}$$

Example 1.1. The standard example in discrete inverse problems is this:

$$A = \begin{bmatrix} 1 & 1 \\ 1 & \frac{1001}{1000} \end{bmatrix} \tag{1.9}$$

$$f = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \tag{1.10}$$

This has a solution $u = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, and A is full rank, so there should be no problems, but that doesn't agree with what we want to say about it – it's very close to being degenerate. We can bring out these bad properties with numerical experiments.

A key property of ill-conditioned discrete inverse problems is the singular values decay very fast. Intuitively, this is because the eigenvectors are all relatively clustered together and point in about the same direction, meaning the eigenvalues for the principal directions away from this cluster will be really small. This directly gives us a high condition number, and we'll see the consequences of this in later lectures.

□

Inverse Problems **Lent 2022**

Lecture 2: Examples of inverse problems

Lecturer: Malena Sabaté Landman *25 January* *Aditya Sengupta*

2.1 Signal and image deblurring

If an image is aberrated by some process, what can we do to reverse that aberration?

The deblurring problem can be formulated by

$$f^{ex}(t_1) = \int_{-\infty}^{\infty} a(t_1 - s)\mu^{ex}(s)\mathrm{d}s \; \forall t_1 \in Y \tag{2.1}$$

so we can define

$$f^{ex}(t) = \int_{-\infty}^{\infty} a(t - s)\mu^{ex}(s)\mathrm{d}s. \tag{2.2}$$

This is a convolution, so we can invoke the convolution theorem,

$$\hat{f}^{ex}(\xi) = \hat{a}(\xi)\hat{\mu}^{ex}(\xi), \tag{2.3}$$

where $\hat{f}^{ex}$ is the Fourier transform of f^{ex},

$$\hat{f}^{ex}(\xi) = \int_{-\infty}^{\infty} e^{-ift} f^{ex}(t)\mathrm{d}t. \tag{2.4}$$

Therefore you get the inverse,

$$\hat{\mu}^{ex}(\xi) = \frac{\hat{f}^{ex}(\xi)}{\hat{a}(\xi)} \tag{2.5}$$

$$\mu^{ex}(s) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{is\xi}\frac{\hat{f}^{ex}(\xi)}{\hat{a}(\xi)}\mathrm{d}s. \tag{2.6}$$

In practice, we have to account for noise, though:

$$f_n(t) = f^{ex}(t) + n(t) \tag{2.7}$$

$$\mu_n(t) = \mu^{ex}(t) + \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{is\xi}\frac{\hat{n}(\xi)}{\hat{a}(\xi)}\mathrm{d}\xi, \tag{2.8}$$

and this additional term makes our lives difficult. In the real world, this term is usually not even well defined.

By defining a kernel matrix that's symmetric and close to diagonal, we can model blurring.

2.2 The heat equation

For another example, consider a torus with Dirichlet boundary conditions:

$$\begin{aligned} \frac{\mathrm{d}v}{\mathrm{d}t} - \triangle v = 0 \text{ on } \pi^d \times \mathbb{R}_+ \\ v(x,t) = 0 \text{ on } \partial\pi^d \times \mathbb{R}_+ \\ v(x,t) = f(x) \text{ on } \pi^d \\ v(x,0) = \mu(x) \text{ on } \pi^d. \end{aligned} \tag{2.9}$$

Here, the Laplacian $\triangle$ represents the difference between the average value of a function in the neighbourhood of a point and the value of the function at that point.

Suppose you have the heat profile at time t, and you want to reconstruct the original profile. Since this is a diffusion process, this is difficult. There can be a lot of high-frequency noise components that further complicate this.

2.3 Tomography

Tomography consists of computing reconstructions from projections.

$$f(\theta, s) = (Pu)(\theta, s) = \int_{\mathbb{R}} \mu(s\theta + t\theta^{\perp})\mathrm{d}t. \tag{2.10}$$

Inverse Problems **Lent 2022**

Lecture 3: Orthogonal decompositions, solutions

Lecturer: Malena Sabaté Landman *27 January* *Aditya Sengupta*

We'll look at one more example of an inverse problem.

Consider the evaluation of the derivative of a function between two specific spaces,

$$f \in L^2\left[0, \frac{\pi}{2}\right] \tag{3.1}$$

$$Df = f', D : L^2\left[0, \frac{\pi}{2}\right] \to L^2\left[0, \frac{\pi}{2}\right]. \tag{3.2}$$

Proposition 3.1. *D is unbounded from $L^2\left[0, \frac{\pi}{2}\right] \to L^2\left[0, \frac{\pi}{2}\right]$.*

Proof .
We have to find a sequence of functions that is bounded, but whose image under D is not. This is satisfied by, for example, $f_n(x) = \sin(nx)$ for $n \in \mathbb{N}$. We can show that $f_n \in L^2\left[0, \frac{\pi}{2}\right]$ for all n, and $\|f_n\| = \frac{\sqrt{\pi}}{4} < \infty$. However, $Df_n(x) = n\cos(nx)$ and this clearly goes to infinity when $n \to \infty$, so D is unbounded.

□

Differentiation can be seen as the inverse problem of solving the integral equation $f(x) = \int_0^x u(t)\mathrm{d}t$.

Having built the foundation of a few examples, let's start looking at how we build generalised solutions to inverse problems. To do this, we need to recap some tools from functional analysis.

Consider a linear operator between spaces $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. For convenience, we'll let $\mathcal{X}$ and $\mathcal{Y}$ be Hilbert spaces, which means they're equipped with scalar products $\langle \cdot, \cdot \rangle_{\mathcal{X}}, \langle \cdot, \cdot \rangle_{\mathcal{Y}}$.

We are particularly interested in the case where A is bounded and linear. This means

$$\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} = \sup_{u \in \mathcal{X} \setminus \{0\}} \frac{\|Au\|_{\mathcal{Y}}}{\|u\|_{\mathcal{X}}} = \sup_{u \mid \|u\|_{\mathcal{X}} = 1} \|Au\|_Y < \infty. \tag{3.3}$$

We'll denote the domain of A by $D(A) = \mathcal{X}$, the kernel by $\mathcal{N}(A) = \{u \in \mathcal{X} \mid Au = 0\}$, the range by $\mathcal{R}(A) = \{f \in \mathcal{Y} \mid f = Au, u \in \mathcal{X}\}$.

Definition 3.1. *A is continuous at $u \in \mathcal{X}$ if for all $\epsilon > 0$, there exists $\delta > 0$ such that if $\|Au - Av\|_{\mathcal{Y}} < \epsilon$ for some $v \in \mathcal{X}$, then v satisfies $\|u - v\|_{\mathcal{X}} < \delta$.*

Theorem 3.2. *Let A be a linear operator between normed spaces. Then A is bounded iff it is continuous.*

By "bounded", we mean that there exists c such that $\|Au\|_{\mathcal{Y}} \leq c\|u\|_{\mathcal{X}}$ for all u. In such a case we say $c = \|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})}$ when c is the minimal scalar for which this would hold.

Definition 3.2. *The adjoint A^* to the operator A is the operator such that $\langle Au, v\rangle_{\mathcal{Y}} = \langle u, A^* v\rangle_{\mathcal{X}}$ for all $u \in \mathcal{X}, v \in \mathcal{Y}$.*

Definition 3.3. *$u, v \in \mathcal{X}$ are orthogonal if $\langle u, v\rangle_{\mathcal{X}} = 0$.*

Definition 3.4. *The orthogonal complement of $\mathcal{X}' \subset \mathcal{X}$ is $X'^{\perp} := \{u \in \mathcal{X} \mid \langle u, v\rangle_{\mathcal{X}} = 0 \,\forall v \in X'\}$.*

Note that if $\mathcal{X}'^{\perp}$ is a closed space, then $\mathcal{X}^{\perp} = \{0\}$, and in general $\mathcal{X}' \subset (\mathcal{X}'^{\perp})^{\perp}$ with equality if $\mathcal{X}'$ is closed.

For closed subspaces $\mathcal{X}' \subset \mathcal{X}$ we have an orthogonal decomposition $\mathcal{X} = \mathcal{X}' \oplus \mathcal{X}'^{\perp}$. This implies that any $u \in \mathcal{X}$ can be written as $u = \underbrace{X}_{\in \mathcal{X}'} + \underbrace{X^{\perp}}_{\mathcal{X}'^{\perp}}$, which is a unique decomposition that will be useful later.

Definition 3.5. *The orthogonal projection on $\mathcal{X}'$ is a map*

$$\mathrm{proj}_{\mathcal{X}'} : u \to \mathcal{X}, \mathrm{proj}_{\mathcal{X}'} \in \mathcal{L}(\mathcal{X}, \mathcal{X}). \tag{3.4}$$

Lemma 3.3. *Let $\mathcal{X}' \subset \mathcal{X}$ be a closed subspace. Then $\mathrm{proj}_{X'}$ satisfies*

1. $\mathrm{proj}^*_{\mathcal{X}'} = \mathrm{proj}_{X'}$ *(self-adjoint)*
2. $\|\mathrm{proj}_{X'}\|_{\mathcal{L}(\mathcal{X},\mathcal{X})} = 1$ *if* $\mathcal{X}' \neq \{0\}$
3. $I - \mathrm{proj}_{X'} = \mathrm{proj}_{X'^{\perp}}$
4. $\|u - \mathrm{proj}_{X'} u\|_{\mathcal{X}} \leq \|u - v\|_{\mathcal{X}} \;\forall v \in \mathcal{X}'$.
5. $x = \mathrm{proj}_{X'} u \iff x \in \mathcal{X}'$ *and* $u - x \in \mathcal{X}'^{\perp}$.

Remark 3.4. *If $\mathcal{X}'$ is not closed, then $(\mathcal{X}'^{\perp})^{\perp} = \overline{\mathcal{X}'}$, so for $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$,*

- $\mathcal{R}(A)^{\perp} = \mathcal{N}(A^*)$
- $\mathcal{R}(A^*)^{\perp} = \mathcal{N}(A)$

This gives rise to the following orthogonal decompositions of $\mathcal{X}, \mathcal{Y}$:

$$\mathcal{X} = \mathcal{N}(A) \oplus \overline{\mathcal{R}(A^*)} \tag{3.5}$$
$$\mathcal{Y} = \mathcal{N}(A^*) \oplus \overline{\mathcal{R}(A)} \tag{3.6}$$

Lemma 3.5. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then $\overline{\mathcal{R}(A^*A)} = \overline{\mathcal{R}A^*}$.*

Proof .
We show they are subsets of each other. In the $\subseteq$ direction,

$$\overline{\mathcal{R}(A^*A)} = \overline{\mathcal{R}(A^*)\mathcal{R}(A)} \subseteq \overline{\mathcal{R}(A^*)}. \tag{3.7}$$

In the $\supseteq$ direction, take $u \in \overline{\mathcal{R}(A^*)}$. Then there exists $f \in \mathcal{N}(A^*)$ such that $\|A^* f - u\|_{\mathcal{X}} < \frac{\epsilon}{2}$. But $\mathcal{N}(A^*)^{\perp} = \overline{\mathcal{R}(A)}$, so there also exists $x \in \mathcal{X}$ such that $\|Ax - f\|_{\mathcal{Y}} < \frac{\epsilon}{2\|A\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})}}$. Then

$$\|A^*Ax - u\|_{\mathcal{X}} \leq \|A^*Ax - A^*f\|_{\mathcal{X}} + \|A^*f - u\|_{\mathcal{X}} \tag{3.8}$$
$$\leq \underbrace{\|A^*\|\|Ax - f\|}_{<\frac{\epsilon}{2}} + \underbrace{\|A^*f - u\|_{\mathcal{X}}}_{<\frac{\epsilon}{2}} < \epsilon. \tag{3.9}$$

□

With this foundation set, we can start looking into solutions for inverse problems. To start with, let's define what we mean by a solution.

Definition 3.6. *An element $u \in \mathcal{X}$ is*

- *a least-squares solution of $Au = f$ if $\|Au - f\|_{\mathcal{Y}} = \inf\{\|Av - f\|_{\mathcal{Y}}, v \in \mathcal{Y}\}$*
- *a minimal-norm solution of $Av = f$, which we denote $u^\dagger$, if $\|u^\dagger\|_{\mathcal{X}} \leq \|v\|_{\mathcal{X}}$ for all least-squares solutions v of $Av = f$.*

Least-squares solutions are not necessarily unique, but we can obtain uniqueness for the minimal-norm solution.

Note that if $\mathcal{R}(A)$ is not closed, a least-squares solution may not exist. If there exists at least one least-squares solution, then the minimal-norm solution is unique. This is because the orthogonal projection of the zero element onto the subspace defined by $\|Au - f\|_{\mathcal{Y}} = \min\{\|Av - f\|_{\mathcal{Y}}, v \in \mathcal{X}\}$ characterises the minimal-norm solution uniquely.

Inverse Problems **Lent 2022**

Lecture 4: Least-squares, pseudoinverses

Lecturer: Malena Sabaté Landman *1 February* *Aditya Sengupta*

We've been talking about least-squares solution existence and uniqueness.

If a least-squares solution exists, then $u^\dagger$ is unique and is the orthogonal projection of the zero element onto the affine subspace defined by

$$\|Au - f\|_y = \min\{\|Av - f\|_y, v \in X\} = S. \tag{4.1}$$

If $Au = f$, then we can take u_0 such that $Au_0 = f$ and $W = \{v \in X \mid Av = 0\} = \mathbb{N}(A)$. Here, S is given by $S = \{u_0 + w \mid w \in W\}$.

If $Av = b$ is not consistent, then the least-squares solution fulfills $A^*Au = A^*f$.

Theorem 4.1. *Let $f \in Y, A \in \mathcal{L}(X, Y)$. Then the following statements are equivalent:*

1. *$u \in X$ satisfies $Au = \mathcal{P}_{\overline{\mathcal{R}(A)}}f$.*
2. *u is a least-squares solution of $Au = f$.*
3. *u solves the normal equations $A^*Au = A^*f$.*

The part of this that might be unintuitive is: where did the normal equations come from? What are they telling us?

We can rearrange them to say $A^*(Au - f) = 0$, so we can verify that if $Au = f$, we have a solution. We haven't yet shown where this comes from or what the significance of the A^* is, though – let's prove it.

Proof .
We show 1 implies 2. Take $u \in \mathcal{X}$ satisfying $Au = \mathcal{P}_{\overline{\mathcal{R}(A)}}f$. Look at the residual:

$$\|Au - f\|_{\mathcal{Y}} = \left\|\mathcal{P}_{\overline{\mathcal{R}(A)}}f - f\right\|_Y = \inf_{g \in \overline{\mathcal{R}(A)}} \|g - f\|_Y \leq \inf_{g \in \mathcal{R}(A)} = \inf_{v \in \mathcal{X}} \|Av - f\|_{\mathcal{Y}}, \tag{4.2}$$

which is the definition of a least-squares solution.
We show 2 implies 3. If u is a least-squares solution, we want to show it satisfies the normal equations exactly. Let $F : \mathbb{R} \to \mathbb{R}$ be defined by $F(\lambda) = \|A(\mu + \lambda v) - f\|^2_{\mathcal{Y}}$. This can be rewritten as

$$F(\lambda) = \lambda^2\|Av\|^2_{\mathcal{Y}} - 2\lambda\langle Av, f - Au\rangle_{\mathcal{Y}} + \|f - Au\|^2_Y. \tag{4.3}$$

Take a derivative and note that we want $F'(\lambda)\big|_{\lambda=0} = 0$. This bcomes equivalent to the cross term being 0.

$$\langle Av, f - Au\rangle_{\mathcal{Y}} = 0 \tag{4.4}$$
$$\langle v, A^*(f - Au)\rangle_{\mathcal{X}} = 0 \tag{4.5}$$
$$A^*(f - Au) = 0. \tag{4.6}$$

We show 3 implies 1. Let u be such that $A^*(f - Au) = 0$. Then $f - Au \in \mathcal{R}(A)^\perp = \overline{\mathcal{R}(A)}^\perp$. We also see that $Au \in \mathcal{R}(A)$, so

$$Au = \mathcal{P}_{\overline{\mathcal{R}(A)}} f \iff \begin{cases} Au \in \overline{\mathcal{R}(A)} \\ f - Au \in (\overline{\mathcal{R}(A)})^\perp \end{cases}. \tag{4.7}$$

This is the definition of a projection, and it follows directly from the normal equations.

□

Lemma 4.2. *Let $f \in \mathcal{Y}$ and let S be the set of least-squares solutions of $Au = f$. Then S is nonempty iff $f \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$.*

Note that this doesn't require finite dimension (otherwise this would be trivial).

Proof .
In the forward direction, suppose we have a nonempty space S. There exists $u \in S$ such that $A^*(Au - f) = 0$, so $f = \underbrace{Au}_{\in \mathcal{R}(A)} + \underbrace{(f - Au)}_{\in \mathcal{R}(A)^\perp} \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$.
In the backward direction, let $f \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$. Then $f = \underbrace{Au}_{\in \mathcal{R}(A)} + \underbrace{g}_{\in \mathcal{R}(A)^\perp = \overline{\mathcal{R}(A)}^\perp}$. This gives us $\mathcal{P}_{\overline{\mathcal{R}(A)}} f = Au$, and this implies a least-squares solution exists by the theorem.

□

Note that if $\mathcal{X}$ and $\mathcal{R}(A)$ are finite, then $\mathcal{R}(A)$ is closed, so a least-squares solution always exists. We can further formalise this.

Theorem 4.3. *Let $f \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$. Then there exists a unique minimal-norm solution $u^\dagger$ to $Au = f$ and all least-squares solutions can be given by $\{u^\dagger\} + \mathcal{N}(A)$.*

Proof .
From Lemma 4.2, we know that S is nonempty. Let $s \in S$. We know s satisfies the normal equations.

$$A(s - u^+) = As - Au^+ = \mathcal{P}_{\overline{\mathcal{R}(A)}} f - \mathcal{P}_{\overline{\mathcal{R}(A)}} f = 0 \implies s - u^+ \in \mathcal{N}(A). \tag{4.8}$$

□

If a minimum-norm solution exists, it can be characterised by the Moore-Penrose pseudoinverse.

Definition 4.1. *Let $A : \mathcal{X} \to \mathcal{Y}$ be linear (or $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$) and*

$$\tilde{A} := A\Big|_{N(A)^\perp} : N(A)^\perp \to \mathcal{R}(A). \tag{4.9}$$

Then the Moore-Penrose pseudoinverse $A^\dagger$ is defined as the unique linear extension of $\tilde{A}^{-1}$ to $D(A^\dagger) = \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$ satisfying $N(A^\dagger) = \mathcal{R}(A)^\perp$.

The restriction $\tilde{A}$ is injective and surjective, so it has an inverse. Since $D(A^\dagger) = \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$, we can always decompose an element $f \in D(A^\dagger)$ into two terms, $f = f_1 + f_2$. Then

$$A^\dagger f = A^\dagger f_1 + A^\dagger f_2 = A^\dagger f_1 = \tilde{A}^{-1} \mathcal{P}_{\mathcal{R}(A)} f. \tag{4.10}$$

Orthogonal complements are always closed, so $\overline{D(A^\dagger)} = \overline{\mathcal{R}(A)} + \overline{\mathcal{R}(A)^\perp} = \mathcal{Y}$, so $D(A^\dagger)$ is dense in Y. If $\mathcal{R}(A)$ is closed, then $D(A^\dagger) = \mathcal{Y}$: the domain of the pseudoinverse is the whole space. This goes both ways, so if the domain is the whole space we can conclude $\mathcal{R}(A)$ is closed. This more or less only happens in the finite-dimensional setting, so it doesn't come up much in inverse problems, which are mostly formulated in infinite dimensions.

Theorem 4.4. *$A^\dagger$ satisfies $\mathcal{R}(A^\dagger) = \mathcal{N}(A)^\perp$ and*

1. $A^\dagger A = \mathcal{P}_{N(A)^\perp}$
2. $AA^\dagger = \mathcal{P}_{\overline{\mathcal{R}(A)}}\Big|_{D(A^\perp)}$
3. $AA^\dagger A = A$
4. $A^\dagger A A^\dagger = A^\dagger$.

Example 4.1. Let $A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$. The pseudoinverse is $A^\dagger = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$, and it satisfies all of the above. For instance, $A^\dagger A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ which is the projection matrix onto $\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^\intercal$, i.e. the projection matrix onto $\mathcal{N}(A)^\perp$.

□

Inverse Problems **Lent 2022**

Lecture 5: Compact operators, the SVD

Lecturer: Malena Sabaté Landman *3 February* *Aditya Sengupta*

5.1 Compact operators

A compact operator is a linear operator that maps bounded sets to compact sets.

Definition 5.1. *Let $A = \mathcal{L}(X, Y)$. Then A is compact if for any bounded set $B \subset X$, $\overline{A(B)}$ is compact in Y.*

An alternative definition is: for every sequence $\subset \{u_j\}_{j \in \mathbb{N}}$, the image set $\subset \{Au_j\}_{j \in \mathbb{N}} \subset Y$ has a convergent subsequence $\{Au_{j_k}\}_{k \in \mathbb{N}} \subset Y$.

The space of all compact operators is denoted by $K(X, Y)$.

Theorem 5.1. *Let $A \in K(X, Y)$ with an infinite-dimensional range. Then the Moore-Penrose pseudoinverse is discontinuous.*

Proof .
Since $\dim \mathcal{R}(A) = \infty$, we must also have that $\dim \mathcal{N}(A)^\perp = \infty$. Therefore, we can construct infinitely many orthonormal elements. Let $\{u_j\}_j \subset \mathcal{N}(A)^\perp$ be such that $\|u_j\| = 1$ and $\langle u_j, u_k \rangle = \delta_{jk}$. Since A is a compact operator, the sequence $\{Au_j : j \in \mathbb{N}\}$ has a convergent subsequence, so for all $\delta > 0$ we can find j, k such that $\|Au_j - Au_k\| < \delta$. However,

$$\left\|A^\dagger A u_j - A^\dagger A u_k\right\|^2 = \left\|\mathrm{proj}_{\mathcal{N}(A)^\perp} u_j - \mathrm{proj}_{\mathcal{N}(A)^\perp} u_k\right\|^2 = \|u_j - u_k\|^2 = \|u_j\|^2 + \|u_k\|^2 - 2\langle u_j, u_k \rangle = 2, \tag{5.1}$$

so the Moore-Penrose pseudoinverse is discontinuous.

□

5.2 Singular value decomposition

Theorem 5.2. *Let $\mathcal{X}$ be a Hilbert space and $A \in K(X, X)$ be self-adjoint. Then there exists an orthonormal basis $\{X_j\}_{j \in \mathbb{N}} \subset \mathcal{X}$ of $\overline{\mathcal{R}(A)}$ and a sequence of eigenvalues $\{\lambda_j\}_{j \in \mathbb{N}} \subset \mathbb{R}$ with $|\lambda_1| \geq |\lambda_2| \geq \cdots > 0$ such that for all $u \in \mathcal{X}$ we have*

$$Au = \sum_{j=1}^{\infty} \lambda_j \langle u, x_j \rangle_x x_j. \tag{5.2}$$

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

Theorem 5.3. *Let $A \in K(X,Y)$. Then there exists*

1. *a not-necessarily infinite null-sequence $\{\sigma_j\}_{j\in\mathbb{N}}$ with $\sigma_1 \geq \sigma_2 \geq \cdots > 0$*

2. *an orthogonal basis $\{x_j\}_{j\in\mathbb{N}} \subset \mathcal{X}$ of $\mathcal{N}(A)^\perp$*

3. *an orthogonal basis $\{y_j\}_{j\in\mathbb{N}} \subset \mathcal{Y}$ of $\overline{\mathcal{R}(A)}$*

such that

$$Ax_j = \sigma_j y_j \tag{5.3}$$
$$A^* y_j = \sigma_j x_j. \tag{5.4}$$

Moreover,

$$Au = \sum_{j=1}^{\infty} \sigma_j \langle u, x_j \rangle y_j \ \forall u \tag{5.5}$$
$$A^* f = \sum_{j=1}^{\infty} \sigma_j \langle f, y_j \rangle x_j \ \forall f \tag{5.6}$$

Inverse Problems **Lent 2022**

Lecture 6: SVD, Picard criterion

Lecturer: Malena Sabaté Landman *8 February* *Aditya Sengupta*

Last time, we covered the decomposition theorem for compact operators on infinite-dimensional vector spaces. We can extend this:

Theorem 6.1. *Let $A \in K(\mathcal{X}, \mathcal{Y})$ with SVD $\{(\sigma_j, x_j, y_j)\}_{j \in \mathbb{N}}$ and $f \in D(A^\dagger)$. Then*

$$A^\dagger f = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle x_j. \tag{6.1}$$

Proof .
Since $f \in D(A^\dagger)$, $u^\dagger = A^\dagger f$ solves $A^* A u^\dagger = A^* f$. Therefore

$$A^* A u^\dagger = A^* \sum_{j=1}^{\infty} \sigma_j \langle u^\dagger, x_j \rangle x_j = \sum_{j=1}^{\infty} \sigma_j^2 \langle u^\dagger, x_j \rangle x_j \tag{6.2}$$

$$A^* f = \sum_{j=1}^{\infty} \sigma_0 \langle f, y_j \rangle x_j. \tag{6.3}$$

These imply

$$\langle u^\dagger, x_j \rangle = \sigma_j^{-1} \langle f, y_j \rangle. \tag{6.4}$$

Since $u^\dagger \in \mathcal{N}(A)^\perp$, we get

$$u^\dagger = \sum_{j=1}^{\infty} \langle u^\dagger, x_j \rangle x_j = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle x_j = A^\dagger f. \tag{6.5}$$

$\square$

Note that $A^\dagger$ is unbounded if $\mathcal{R}(A)$ is infinite dimensional. Think about the sequence $\{y_j\}$, where $\|y_j\| = 1$, so $\left\|A^\dagger y_j\right\| \cdot \sigma_j^{-1} \to \infty$.

Definition 6.1. *We say that f satisfies the Picard condition if*

$$\left\|A^\dagger f\right\|^2 = \sum_{j=1}^{\infty} \frac{|\langle f, y_j \rangle|^2}{\sigma_j^2} < \infty. \tag{6.6}$$

Observe that as $\sigma_r \to 0$, this is a condition on $\langle f, y_j \rangle$ decaying sufficiently fast.

Theorem 6.2. *Let $A \in K(\mathcal{X}, \mathcal{Y})$ with SVD $\{(\sigma_j, x_j, y_j)\}_{j\in\mathbb{N}}$ and $f \in \overline{\mathcal{R}(A)}$. This measurement belongs to the range of A (not just the closure) if and only if it fulfils the Picard criterion.*

Proof .
Forward direction: let $f \in \mathcal{R}(A)$. Then there exists $u \in \mathcal{X}$ such that $Au = f$.

$$\langle f, y_j \rangle_{\mathcal{Y}} = \langle Au, y_j \rangle_{\mathcal{Y}} = \langle u, A^* y_j \rangle_{\mathcal{X}} = \sigma_j \langle u, x_j \rangle_{\mathcal{X}}. \tag{6.7}$$

Therefore

$$\sum_{j=1}^{\infty} \sigma_j^{-2} |\langle f, y_j \rangle_{\mathcal{Y}}|^2 = \sum_{j=1}^{\infty} |\langle u, x_j \rangle_{\mathcal{X}}|^2 \leq \|u\|_{\mathcal{X}}^2 < \infty. \tag{6.8}$$

Backward direction: start with the Picard criterion, and let

$$u = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle_Y x_j \in \mathcal{X} \tag{6.9}$$

$$Au = \sum_{j=1}^{\infty} \sigma_j^{-1} \langle f, y_j \rangle_{\mathcal{Y}} A x_j = \sum_{j=1}^{\infty} \langle f, y_j \rangle y_j = \mathcal{P}_{\overline{\mathcal{R}(A)}} f = f, \tag{6.10}$$

so $f \in \mathcal{R}(A)$.

$\square$

Definition 6.2. *A problem is mildly ill-posed if the singular values decay at most with polynomial speed, i.e. three exists $\gamma, c > 0$ such that $\sigma_j \geq c j^{-\gamma}$ for all j.*

Definition 6.3. *A problem is severely ill-posed if the singular values decay faster than polynomial speed, i.e. there exist $\gamma, c > 0$ such that $\sigma_j \leq c j^{-\gamma}$ for all j big enough.*

Example 6.1. Consider the differentiation operator, $A : L^2([0,1]) \to L^2([0,1])$,

$$(Au)(t) = \int_0^t u(s) \mathrm{d}s = \int_0^1 K(s,t) u(s) \mathrm{d}s, \tag{6.11}$$

where $K(s,t) = \begin{cases} 1 & s \leq t \\ 0 & \text{otherwise} \end{cases}$. Then

$$(A^* f)(s) = \int_0^1 K(t,s) f(t) \mathrm{d}t = \int_s^1 v(t) \mathrm{d}t. \tag{6.12}$$

We are interested in the SVD of A^*A, so we want $\sigma^2, x \in L^2([0,1])$ such that

$$\sigma^2 x(s) = (A^*Ax)(s) = \int_s^1 \int_0^t x(r)\mathrm{d}r\mathrm{d}t, \tag{6.13}$$

which implies $x(1) = 0$.

Also

$$\sigma^2 x'(s) = \frac{\mathrm{d}}{\mathrm{d}s} \int_s^1 \int_0^t x(r)\mathrm{d}r\mathrm{d}t = -\int_0^s x(r)\mathrm{d}r \implies x'(0) = 0. \tag{6.14}$$

Also

$$\sigma^2 x''(s) + x(s) = 0, \tag{6.15}$$

so we solve this as we usually solve a second-order ODE,

$$x(s) = \sqrt{2}\cos\left(\frac{2}{(2j-1)\pi} s\right) \tag{6.16}$$

for $j \in \mathbb{N}$, where we chose the leading $\sqrt{2}$ to be a normalising constant. We can use the x_js to find the y_js, by computing $Ax_j = \sigma_j y_j$.

$$y_j(s) = \sqrt{2}\sin\left(\left(j - \frac{1}{2}\right)\pi s\right). \tag{6.17}$$

□

Inverse Problems **Lent 2022**

Lecture 7: Regularisation theory

Lecturer: Malena Sabaté Landman *10 February* *Aditya Sengupta*

Regularisation consists of replacing a problem with another one that's easier to solve, and whose solution is close to that of the original problem. We'll have to show that the problem and its regularisation are related in some way, but that the regularised one is easier to solve.

If we have a measurement f_δ (i.e. $\|f - f_\delta\| \leq \delta$), then $A^\dagger f_\delta \nrightarrow A^\dagger f$ as $\delta \to 0$ in general. So how do we find an operator that does this and is close to $A^\dagger$?

We take a family of operators R_α, parameterised by some $\alpha(\delta, f_\delta)$ such that $R_{\delta,f_\delta}(f_\delta) \to A^\dagger f$ as $\delta \to 0$, for all $f \in D(A^\dagger)$ and $f_\delta \in \mathcal{Y}$ such that $\|f - f_\delta\|_{\mathcal{Y}} \leq \delta$.

Definition 7.1. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ be bounded. A family $\{R_\alpha\}_{\alpha>0}$ of continuous operators is called a regularisation of $A^\dagger$ if $R_\alpha f \to A^\dagger f = u^\dagger$ for all $f \in D(A^\dagger)$ as $\alpha \to 0$.*

Observe that since $A^\dagger$ may not be continuous, R_α might not be bounded when $\alpha \to 0$.

Theorem 7.1. *Let $\mathcal{X}, \mathcal{Y}$ be Hilbert spaces and $\{A_j\}_{j\in\mathbb{N}} \subset \mathcal{L}(\mathcal{X}, \mathcal{Y})$ be a family of pointwise bounded operators, that is, for all $u \in \mathcal{X}$, there exists $c(u) > 0$ such that if $\sup_{j\in\mathbb{N}} \|A_j u\| \leq c(u)$, then $\sup_{j\in\mathbb{N}} \|A_j\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} < \infty$.*

Corollary 7.2. *Let $\mathcal{X}, \mathcal{Y}$ be Hilbert spaces and let $\{A_j\}_{j\in\mathbb{N}} \subset \mathcal{L}(\mathcal{X}, \mathcal{Y})$. The following are equivalent:*

1. *There exists $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ such that $Au = \lim_{j\to\infty} A_j u$ for all $u \in \mathcal{X}$.*
2. *There exists a dense subspace $\mathcal{X}' \subset \mathcal{X}$ such that $\lim_{j\to\infty} A_j u$ exists for all $u \in \mathcal{X}'$ and $\sup_{j\in\mathbb{N}} \|A_j\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})}$.*

Theorem 7.3. *Let $\mathcal{X}, \mathcal{Y}$ be Hilbert spaces, let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ and let $\{R_\alpha\}_{\alpha>0}$ be a regularisation of $A^\dagger$. If $A^\dagger$ is not continuous, $\{R_\alpha\}_{\alpha>0}$ cannot be uniformly bounded. In particular, there exists $f \in \mathcal{Y}$ and a sequence $\alpha_j \to 0$ such that $\left\|R_{\alpha_j} f\right\|_{j\to\infty} \to \infty$.*

Proof .
We proceed by contradiction. Assume $\{R_\alpha\}_{\alpha>0}$ is uniformly bounded. Then, there exists C such that $\|R_\alpha\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} \leq C$ for all $\alpha > 0$. Given the definition of regularisation, $R_{\alpha_j} \to A^\dagger$ on $D(A^\dagger)$ for all $\alpha_j \to 0$, since $D(A^\dagger)$ is dense in $\mathcal{Y}$. So if the limit exists in the whole space $\mathcal{Y}$, then we can extend $A^\dagger$ to the whole $\mathcal{Y}$ as this limit, so $A^\dagger$ is a bounded operator on $\mathcal{L}(\mathcal{X}, \mathcal{Y})$. But this is a contradiction as $A^\dagger$ is not continuous on Y.
In the other direction, we also proceed by contradiction. Assume that for all $f \in \mathcal{Y}$ and any sequence $\alpha_j \to 0$, we have $\sup_{j\in\mathbb{N}} \left\|R_{\alpha_j} f\right\| \leq c(f) < \infty$. Then we would have $\sup_{j\in\mathbb{N}} \left\|R_{\alpha_j}\right\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} \leq C < \infty$, which is a contradiction.

□

Theorem 7.4. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ and let $\{R_\alpha\}_{\alpha>0}$ be a linear regularisation of $A^\dagger$. If $\sup_{\alpha>0} \|AR_\alpha\|_{\mathcal{L}(\mathcal{X},\mathcal{X})} < \infty$, then $\|R_\alpha f\|_{\mathcal{X}} \to \infty$ for all $f \in D(A^\dagger)$.*

Proof .
Define $u_\alpha = R_\alpha f$ for $f \notin D(A^\dagger) = R(A) \oplus R(A)^\perp$. Assume there exists a sequence $\alpha_k \to 0$ such that $\|u_{\alpha_k}\|_\mathcal{X}$ is uniformly bounded.
At this point, we'll need a few functional analysis tools:

- Bounded sets in Hilbert spaces are weakly precompact, i.e. there exists a weakly convergent subsequence $u_{\alpha_{k_l}}$ with limit $u \in \mathcal{X}$.
- Continuous operators are weakly continuous.

Putting these two together, we can conclude that $Au_{\alpha_{k_l}} \to Au$.
In the reverse direction, for any $g \in D(A^\dagger)$ we have that $AR_{\alpha_{k_l}} g = AA^\dagger g = \mathcal{P}_{\overline{\mathcal{R}(A)}} g$. Since $D(A^\dagger)$ is dense in Y, using the corollary, we can conclude that $AR_{\alpha_{k_l}} f \to \mathcal{P}_{\overline{\mathcal{R}(A)}} f$, and (as we did it in the first part of the proof) we know that $Au_{\alpha_{k_l}} \rightharpoonup Au$, so we get $Au = \mathcal{P}_{\overline{\mathcal{R}(A)}} f \in \mathcal{R}(A)$. But $\mathcal{Y} = \overline{\mathcal{R}(A)} \oplus \mathcal{R}(A)^\perp$, so $f \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp = D(A^\dagger)$.

$\square$

So far, we've worked with an arbitrary α, but for concrete problems we'd like a particular way to choose it, preferably as a function of the noise δ. We'd like to set things up such that $\alpha \to 0$ as $\delta \to 0$.

Consider

$$\begin{aligned} \left\|R_\alpha f_\delta - u^\dagger\right\|_\mathcal{X} &\le \|R_\alpha f_\delta - R_\alpha f\|_\mathcal{X} + \left\|R_\alpha f - u^\dagger\right\|_\mathcal{X} && (7.1) \\ &\le \delta \|R_\alpha\|_{\mathcal{L}(\mathcal{X},\mathcal{X})} + \left\|R_\alpha f - A^\dagger f\right\|_\mathcal{X}. && (7.2) \end{aligned}$$

The first term could be complicated; we saw above that it isn't necessarily uniformly bounded. These represent a regularisation tradeoff: the more you regularise, the more the magnitude of the first term goes down and the more the magnitude of the second term goes up. As $\alpha \to 0$, the data error goes to infinity and the approximation error goes to 0, and as $\alpha \to \infty$, the reverse happens.

Inverse Problems **Lent 2022**

Lecture 8: Parameter choice

Lecturer: Malena Sabaté Landman *15 February* *Aditya Sengupta*

8.1 Parameter choice rules

Definition 8.1. *A function $\alpha : R_{>0} \times \mathcal{Y} \to \mathbb{R}_{>0}, (\delta, f_\delta) \to \alpha(\delta, f_\delta)$ is a parameter choice rule.*

In particular, we can have

- *a priori* parameter choice rules, in which the rule depends only on the noise level δ and not on the measurements f_δ;
- *a posteriori* parameter choice rules, in which the rule depends on both δ and f_δ; and,
- an heuristic parameter choice rule, in which the rule depends only on the measurement.

Definition 8.2. *Let $\{R_\alpha\}_{\alpha>0}$ be a regularisation of $A^\dagger$. If $f \in \mathrm{dom}(A^\dagger)$, there exists a parameter choice rule such that*

$$\lim_{\delta\to 0}\left(\sup_{f_\delta \mid \|f - f_\delta\|_{\mathcal{Y}} < \delta} \left\|R_\alpha f_\delta - A^\dagger f\right\|_{\mathcal{X}}\right) = 0. \tag{8.1}$$

and

$$\lim_{\delta\to 0}\left(\sup_{f_\delta \mid \|f - f_\delta\|_{\mathcal{Y}} < \delta} \alpha(\delta, f_\delta)\right) = 0. \tag{8.2}$$

In this case, we say (R_α, α) is a convergent regularisation.

In words, we're taking the supremum over a noise level δ, i.e. we're taking the worst case, and ensuring we can still control it, i.e. the (regularised function on noisy data) and (pseudoinverse on true data) still match up for sufficiently small δ.

8.1.1 *A priori* parameter choice rules

Theorem 8.1. *Let $\{R_\alpha\}_{\alpha>0}$ be a regularisation of $A^\dagger$, where $A \in \mathcal{L}(\mathcal{X}, Y)$. There is always a parameter choice rule depending only on δ such that we have a convergent regularisation.*

Theorem 8.2 (Linear convergent regularisation)**.** *Let $\{R_\alpha\}_{\alpha>0}$ be a linear regularisation and $\alpha : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be an a priori parameter choice rule. Then (R_α, α) is a convergent regularisation method if and only if $\lim_{\delta\to 0} \alpha(\delta) \to 0$ and $\lim_{\delta\to 0} \delta\left\|R_{\alpha(\delta)}\right\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} = 0$.*

Proof .
In the backward direction, we assume the two limits hold. For any $f \in \operatorname{dom}(A^\dagger)$ and any $f_\delta \in \mathcal{Y}$ such that $\|f - f_\delta\|_{\mathcal{Y}} \leq \delta$, using the two limits

$$\begin{aligned}\lim_{\delta\to 0} \sup_{f_\delta : \|f - f_\delta\| < \delta} \left\|R_{\alpha(\delta)} f_\delta - A^\dagger f_\delta\right\| \leq & \lim_{\delta\to 0} \sup_{f_\delta : \|f - f_\delta\| < \delta} \left\|R_{\alpha(\delta)} f_\delta - R_{\alpha(\delta)} f\right\| \\ & + \lim_{\delta\to 0} \left\|R_{\alpha(\delta)} f - A^\dagger f\right\| \\ & + \lim_{\delta\to 0} \sup_{f_\delta : \|f - f_\delta\| < \delta} \|f - f_\delta\|.\end{aligned} \tag{8.3}$$

The first term goes to 0 by the operator norm on R_α. The second term goes to 0 by α going to 0 with δ and by R_α being a regularisation. The third term goes to 0 because of the condition on the supremum. Therefore we have

$$\left\|R_{\alpha(\delta)} - A^\dagger f\right\|_{\mathcal{X}} \xrightarrow{\delta\to 0} 0, \tag{8.4}$$

and so (R_α, α) is a convergent regularisation method.
In the forward direction, we take the contrapositive: we need to show that if *either* limit does not hold individually, then (R_α, α) is not a convergent regularisation method. If $\lim_{\delta\to 0} \alpha(\delta) \to 0$ does not hold, then the definition of being a convergent regularisation is violated, so assume it does hold, and consider the case where $\lim_{\delta\to 0} \delta \left\|R_{\alpha(\delta)}\right\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} = 0$ does not hold. Then there exists a null sequence $\{\delta_k\}_{k\in\mathbb{N}}$ such that $\delta_k \|R_\alpha(\delta_k)\|_{\mathcal{L}(\mathcal{X},\mathcal{Y})} \geq c > 0$ for some c.
This implies we can find a sequence $\{g_k\}_{k\in\mathbb{N}} \subset \mathcal{Y}$ with $\|g_k\|_{\mathcal{Y}} = 1$ and $\delta_k \left\|R_{\alpha(\delta_k)} g_k\right\|_{\mathcal{X}} \geq \tilde{c}$. Let $f \in \operatorname{dom}(A^\dagger)$ and define $f_k = f + \delta_k g_k$. Then

$$\|f - f_k\| \leq \delta_k \tag{8.5}$$

$$R_{\alpha(\delta_k)} f_k - A^\dagger f = R_{\alpha(\delta_k)} f - A^\dagger f + \underbrace{\delta_k R_{\alpha(\delta_k)} g_k}_{\text{bounded from below}} \not\to 0. \tag{8.6}$$

Therefore, this is not a convergent regularisation.

□

8.1.2 *A posteriori* parameter choice rules

Most parameter choice rules used in practice fall into this category.

Note that if $\alpha = \alpha(\delta)$ defines a convergent regularisation method, then so does $\tilde{\alpha} = \alpha(c\delta)$ for any $c > 0$. Asymptotically they are equivalent, but for a given δ they might give very different results. This adds in a level of degeneracy we might not want, that might be corrected by adding in direct dependence on the measurements.

The idea behind *a posteriori* parameter choice rules is as follows. We have $f \in \operatorname{dom}(A^\dagger)$ and $f_\delta \in \mathcal{Y}$ such that $\|f - f_\delta\| \leq \delta$, and we want to look at the residual associated to $u_\alpha = R_\alpha f_\delta$.

If we define

$$\mu = \inf\{\|Au - f\|, u \in \mathcal{X}\} = \|Au^* - f\|, \tag{8.7}$$

then

$$\left\|Au^\dagger - f_\delta\right\| \leq \left\|Au^\dagger - f\right\| + \|f_\delta - f\| \leq \mu + \delta. \tag{8.8}$$

In a lot of cases, this is a pretty tight bound.

Definition 8.3 (Morozov's discrepancy principle)**.** *Let $u_\alpha = R_\alpha f_\delta$ with $\alpha(\delta, f_\delta)$ chosen by*

$$\alpha(\delta, f_\delta) = \sup\{\alpha > 0 \mid \|Au_\alpha - f_\delta\| \leq \eta\delta\}, \tag{8.9}$$

where $\eta > 1$ is a fixed constant representing a "safety threshold".

Then we say $u_{\alpha(\delta,f_\delta)} = R_{\alpha(\delta,f_\delta)} f_\delta$ satisfies the discrepancy principle.

Note that the discrepancy principle yields a convergent regularisation method.

8.1.3 Heuristic parameter choice rules

In some settings, you don't have a good estimate of the error and would like parameter choices to be based only on the data itself. This has a massive drawback, though.

Theorem 8.3 (Bakushinskii veto)**.** *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ and $\{R_\alpha\}$ be a regularisation where we let $\alpha = \alpha(f_\delta)$. Then if (R_α, α) is a convergent regularisation of $A^\dagger$, then $A^\dagger$ is continuous from $\mathcal{Y}$ to $\mathcal{X}$.*

Inverse Problems **Lent 2022**

Lecture 9: Parameter choice rules, Tikhonov regularisation

Lecturer: Malena Sabaté Landman *17 February* *Aditya Sengupta*

9.1 Parameter choice

Recall

$$A^\dagger f = \sum_{j=1}^{\infty} \frac{1}{\sigma_j}\langle f, y_j\rangle x_j, \tag{9.1}$$

where $\{\sigma_j, x_j, y_j\}$ is the SVD of A. Let us consider the regularization

$$R_\alpha f = \sum_{j=1}^{\infty} g_\alpha(\sigma_j)\langle f, y_j\rangle x_j \;\forall f \in \mathcal{Y}, \tag{9.2}$$

where $g_\alpha : \mathbb{R}_+ \to \mathbb{R}_+$, and for all $\sigma > 0$, $g_\alpha(\sigma) \xrightarrow{\sigma\to 0} \frac{1}{\sigma}$, and $g_\alpha(\sigma) \leq c_\alpha$ for all $\sigma \in \mathbb{R}_{>0}$.

Theorem 9.1. *Let $g_\alpha : \mathbb{R}_+ \to \mathbb{R}_+$ be a piecewise continuous function such that*

1. $g_\alpha(\sigma) \leq c_\alpha \;\forall \sigma \in \mathbb{R}_+$
2. $\lim\limits_{\sigma\to 0} g_\alpha(\sigma) = \frac{1}{\sigma}$
3. $\sup_{\alpha,\sigma} \sigma g_\alpha(\sigma) \leq \gamma$ *for some* $\gamma > 0$.

Then $R_\alpha f = \sum_{j=1}^{\infty} g_\alpha(\sigma_j)\langle f, y_j\rangle x_j$ satisfies $R_\alpha f \xrightarrow{\alpha\to 0} A^\dagger f$.

It's important that this is piecewise continuous.

Proof .
This proof is based on the spectral decomposition, and because we've defined both the components of $A^\dagger$ and R_α in this decomposition, this basically just comes down to comparing basis components. Using the SVD of $A^\dagger$,

$$R_\alpha f - A^\dagger f = \sum_{j=1}^{\infty}\left(g_\alpha(\sigma_j) - \frac{1}{\sigma_j}\right)\langle f, y_j\rangle_{\mathcal{Y}} x_j \tag{9.3}$$

$$= \sum_{j=1}^{\infty}(\sigma_j g_\alpha(\sigma_j - 1)\langle u^+, x_j\rangle_{\mathcal{X}} x_j \tag{9.4}$$

So we can look at the norm-squared,

$$\left\|R_\alpha f - A^{-1} f\right\|_{\mathcal{X}}^2 = \sum_{j=1}^{\infty} (\sigma_j g_\alpha(\sigma_j) - 1)^2 \left|\langle u^+, x_j \rangle_{\mathcal{X}}\right|^2. \tag{9.5}$$

From the third condition in the theorem, we have $(\sigma_j g_\alpha(\sigma_j) - 1)^2 \leq (1+\gamma^2)$, and further we can say $\sum_{j=1}^{\infty}(1+\gamma^2)|\langle u^+, x_j\rangle_{\mathcal{X}}|^2 = (1+\gamma^2)\|u^+\|^2 < \infty$. Therefore

$$\lim_{\alpha \to 0} \sup \left\|R_\alpha f - A^\dagger f\right\|_{\mathcal{X}}^2 = \lim_{\alpha \to 0} \sup \sum_{j=1}^{\infty} (\sigma_j g_\alpha(\sigma_j) - 1)^2 \left(\langle u^+, x_j \rangle_{\mathcal{X}}\right)^2 \tag{9.6}$$

$$\leq \sum_{j=1}^{\infty} \left(\lim_{\alpha \to 0} \sup (\sigma_j g_\alpha(\sigma_j))^2 \left|\langle u^+, x_j\rangle_{\mathcal{X}}\right|^2 \right) = 0, \tag{9.7}$$

where the $\leq$ is done due to the reverse Fatou lemma and the equality to 0 comes via the pointwise convergence of $g_\alpha(\sigma_j)$ to $\frac{1}{\sigma_j}$.

$\square$

Theorem 9.2 (Parameter choice rules for filtering methods)**.** *Let the assumptions of the previous theorem hold. Let $\alpha = \alpha(\delta)$, i.e. we take a regularisation parameter that just depends on the noise level. Then $(R_{\alpha(\delta)}, \alpha(\delta))$ is a convergent regularisation method if*

$$\lim_{\delta \to 0} \delta C_{\alpha(\delta)} = 0. \tag{9.8}$$

The proof follows from $\left\|R_{\alpha(\delta)}\right\|_{\alpha(\mathcal{Y},\mathcal{Y})} \leq C_{\alpha(\delta)}$ and the previous theorem.

Example 9.1. Truncated singular value decomposition: this is pretty self-explanatory. We've seen a decay rate (for σ_j against the number of singular values) for true problems that isn't replicated by regularisers, which tend to just level out. An easy way to avoid this is to just cut off the series.

$$R_\alpha f = \sum_{\sigma_j \geq \alpha} \frac{1}{\sigma_j} \langle f, y_j \rangle_{\mathcal{Y}} x_j = \sum_{j=1}^{\infty} g_\alpha(\sigma_j) \langle f, y_j \rangle_Y x_j, \tag{9.9}$$

where $g_\alpha(\sigma) = \begin{cases} \frac{1}{\sigma} & \sigma \geq \alpha \\ 0 & \sigma < \alpha \end{cases}$.

We can look at the conditions on the spectral regularisation theorem:

1. $C_\alpha = \frac{1}{\alpha}$

2. $\lim_{\alpha \to 0} g_\alpha(\sigma) = \frac{1}{\sigma}$ for all σ

3. If $\gamma = 1$, this holds.

So TSVD is a convergent regularisation method if $\lim\limits_{\delta\to 0} \frac{\delta}{\alpha(\delta)} = 0$. It remains to be seen how to choose this threshold.

Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ with an SVD $\{(\sigma_j, x_j, y_j)\}_{j\in\mathbb{N}}$ and choose for $j > 0$ an index function $j^* : \mathbb{R}_+ \to \mathbb{N}$, such that $j^*(\delta) \xrightarrow{\delta\to 0} \infty$ and $\lim_{\delta\to 0} \frac{\delta}{\sigma_{j^*(\delta)}} = 0$. Then $\alpha(\delta) = \sigma_{j^*(\delta)}$ as a parameter choice rule leads to a convergent regularisation method.

□

9.2 Tikhonov regularisation

We can think of Tikhonov regularisation as a shift in the eigenvalues of the normal equations. We have $A^*Au = A^*f$, and A^*A admits a decomposition to $X\Sigma X^\intercal$. If we know the singular values don't behave very well, we can approximate this expansion by a problem where they do. We take $X\Sigma X^\intercal \approx X(\Sigma + \alpha I)X^\intercal$. We can prove this approximation corresponds to

$$R_\alpha f = \sum_{j=1}^{\infty} \frac{\sigma_j}{\sigma_j^2 + \alpha} \langle f, y_j\rangle_{\mathcal{Y}} x_j, \tag{9.10}$$

which is the standard regularisation form for $g_\alpha(\sigma) = \frac{\sigma}{\sigma^2+\alpha}$.

Let's look at the conditions of the theorem of spectral regularisation, and see if they're satisfied.

1. $0 \le (\sigma - \sqrt{\alpha})^2 = \sigma^2 - 2\sigma\sqrt{\alpha} + \alpha \implies \sigma^2 + \alpha \ge 2\sigma\sqrt{\alpha}$, so $\frac{1}{2\sqrt{\alpha}} \ge \frac{\sigma}{\sigma^2+\alpha}$, so condition 1 holds for $C_\alpha = \frac{1}{2\sqrt{\alpha}}$.
2. Taking the limit is trivial: $\lim_{\alpha\to 0} g_\alpha(\sigma) = \frac{1}{\sigma} \to 0$.
3. This is satisfied for $\gamma = 1$.

Therefore Tikhonov regularisation is a convergent regularisation method if $\lim_{\delta\to 0} \frac{\delta}{\sqrt{\alpha}} = 0$.

Applying $A^*A + \alpha I$ to $u_\alpha = R_\alpha f$, we have

$$(A^*A + \alpha I)u_\alpha = \sum_{j=1}^{\infty} (\sigma_j^2 + \alpha\langle u_\alpha, x_j\rangle_{\mathcal{X}} x_j = \sum_{j=1}^{\infty} (\sigma_j^2 + \alpha) \frac{\sigma_j}{\sigma_j^2 + \alpha} \langle f, y_j\rangle_{\mathcal{Y}} x_j = A^*f. \tag{9.11}$$

This proves that computing a Tikhonov solution is equivalent to solving the original equation.

$$(A^*A + \alpha I)u_\alpha = A^*f, \tag{9.12}$$

or

$$A^*(Au_\alpha - f) + \alpha u_\alpha = 0, \tag{9.13}$$

which looks like a first-order optimality condition of

$$\min_{u \in \mathcal{X}} \left\{ \|Au - f\|^2 + \alpha \|u\|^2 \right\}. \tag{9.14}$$

Inverse Problems **Lent 2022**

Lecture 10: Variational regularisation

Lecturer: Malena Sabaté Landman *22 February* *Aditya Sengupta*

Variational regularisation can be written as the following minimisation problem:

$$\min_{u \in X} \frac{1}{2} \|Au - f_\delta\|^2 + \alpha J(u), \tag{10.1}$$

where the first term is the "fit-to-data" or fidelity term, and the second is the regulariser, which penalises unwanted behaviour in the solution.

Variational regularisation has a nice link to Bayesian statistics, as the fit-to-data term is related to the noise and the regularisation term is related to the prior.

The regularisation operator is defined as

$$R_\alpha f_\delta \in \arg\min_{u \in X} \frac{1}{2} \|Au - f_\delta\|^2 + \alpha J(u). \tag{10.2}$$

In general, the minimiser does not have to be unique, so we express this with $\in$ not $=$.

We'll go into some background on functional analysis first.

A Banach space is a complete normed vector space. They aren't necessarily inner product spaces.

Definition 10.1. *The dual space $\mathcal{X}^*$ to a Banach space $\mathcal{X}$ is the space of all continuous linear functionals $f \in \mathcal{L}(\mathcal{X}, \mathbb{R})$.*

Note that given $u \in \mathcal{X}$ and $p \in X^*$, we usually write the dual product as $\langle p, u \rangle$ instead of $p(u)$.

We can use this to define the adjoint operator of a linear operator: given $A \in \mathcal{L}(\mathcal{X}, Y)$, there exists a unique operator $A^* : \mathcal{Y}^* \to X^*$ called the adjoint of A, such that for all $u \in \mathcal{X}, p \in Y^*$,

$$\langle A^* p, u \rangle = \langle p, Au \rangle. \tag{10.3}$$

The dual space of a Banach space $\mathcal{X}$ can be equipped with a norm

$$\|p\|_{X^*} := \sup_{u \in \mathcal{X}, \|u\|_X \leq 1} \langle p, u \rangle. \tag{10.4}$$

With this norm, $\mathcal{X}^*$ is a Banach space. We can do this again, to get the bidual space X^{**} which is also a Banach space.

Every $u \in \mathcal{X}$ defines a continuous linear mapping on the dual space X^* defined as $\langle E(u), p \rangle := \langle p, u \rangle$. The mapping $E : \mathcal{X} \to X^{**}$ is well defined. It can be shown that E is linear, continuous, and an isometry (and therefore injective). We call this mapping the *canonical embedding* from $\mathcal{X}$ to $\mathcal{X}^{**}$. We also say that if the canonical embedding E is surjective, then $\mathcal{X}$ is reflexive: that is, the space $\mathcal{X}$ coincides with its bidual.

Examples of reflexive Banach spaces include Hilbert spaces, L^p spaces (spaces of functions with norms that are generalisations of $p-$norms in infinite dimensions) or ℓ^p spaces (spaces of sequences with associated norms which are generalisations of p-norms).

Definition 10.2. *A Banach space $\mathcal{X}$ is separable if $\exists \mathcal{X}' \subset \mathcal{X}$ of at most countable cardinality such that $\overline{\mathcal{X}'} = \mathcal{X}$.*

A problem with infinite-dimensional spaces is that bounded sequences might not have convergent subsequences.

Example 10.1. Take $\{u^k\}_{k\in\mathbb{N}} \subset \ell^2$ such that $u^k_j = \delta_{kj}$. This is a sequence of sequences like $(1,0,0,\dots),(0,1,0,\dots),(0,0,1,\dots),\dots$. Then $\left\|u^k\right\|_{\ell^2} = 1$, but there is no $u \in \ell^2$ such that $u^k \to u$. This is a bounded sequence without a convergent subsequence.

$\square$

To work around this, we'll introduce weak convergence.

Definition 10.3. *A sequence $\{u^k\}_{k\in\mathbb{N}} \subset \mathcal{X}$ converges weakly to $u \in \mathcal{X}$, denoted by $u^k \rightharpoonup u$ if for all $p \in X^*$, the sequence of real numbers $\{\langle p, u^k\rangle\}_{u\in\mathbb{N}}$ converges to $\langle p, u\rangle$.*

Definition 10.4. *A sequence $\{p^k\}_{k\in\mathbb{N}} \subset \mathcal{X}^*$ converges weakly-* to $p \in \mathcal{X}$, denoted by $p^k \rightharpoonup^* \mathcal{X}^*$ if $\langle p^k, u\rangle \to \langle p, u\rangle$ for all $u \in \mathcal{X}$.*

Theorem 10.1 (Banach-Alaoglu)**.** *Let X^* be the dual space of a Banach space $\mathcal{X}$. Then the unit ball $B_{\mathcal{X}} = \{p \in \mathcal{X}^* \mid \|p\| \le 1\}$ is compact in the weak-* topology. If $\mathcal{X}$ is separable, then the weak-* topology is metrisable on bounded sets, and every bounded sequence $\{u^k\}_{u\in\mathbb{N}} \subset \mathcal{X}$ has a weak-* convergent subsequence.*

Theorem 10.2. *Each bounded sequence $\{u^k\}_{k\in\mathbb{N}}$ in a reflexive, separable Banach space $\mathcal{X}$ has a weakly convergent subsequence.*

Definition 10.5. *Let $\mathcal{X}$ be a Banach space with a given topology $\tau_{\mathcal{X}}$. Then the functional $E : \mathcal{X} \to \overline{\mathbb{R}} = \mathbb{R}\cup\{-\infty,\infty\}$ is said to be sequentially lower semi-continuous with respect to $\tau_{\mathcal{X}}$ at $u \in \mathcal{X}$ if*

$$E(u) \le \lim_{j\to\infty}\inf E(u_j), \tag{10.5}$$

for all sequences $\{u_j\}_{j\in\mathbb{N}} \subset \mathcal{X}$ such that $u_j \to u$ in the topology $\tau_{\mathcal{X}}$.

Note that if a topology is induced by a metric, then continuity and sequential continuity are equivalent.

Example 10.2. The functional $\|\cdot\|_1 : \ell^2 \to \overline{\mathbb{R}}$ defined as

$$\|u\|_1 = \begin{cases} \sum_{j=1}^{\infty} |u_j| & u \in \ell^1 \\ \infty & u \notin \ell^1 \end{cases} \tag{10.6}$$

is weakly lower semi-continuous in ℓ^2.

□

Proof .
Take a weakly convergent sequence $\{u_j\}_{j\in\mathbb{N}} \subset \ell^2$, and let its limit be u. If we take the element in the dual $\delta_k : \ell^2 \in \overline{\mathbb{R}}$ such that

$$\langle \delta_k, v\rangle = v_k \ \forall k \in \mathbb{N}, \tag{10.7}$$

then we have that $u^j_k = \langle \delta_k, u^j\rangle \to \langle \delta_k, u\rangle = u_k$. We can now see that the 1-norm $\|\cdot\|_1$ is lower-semi continuous, because

$$\|u\|_1 = \sum_{k=1}^{\infty} |u_k| = \sum_{k=1}^{\infty} \lim_{j\to\infty} \left|u^j_k\right| \underbrace{\leq}_{\text{Fatou's lemma}} \liminf_{j\to\infty} \sum_{k=1}^{\infty} \left|u^j_k\right| = \liminf_{j\to\infty} \left\|u^j\right\|_1. \tag{10.8}$$

Note that it is not clear whether either side of this inequality is finite.

□

The other bit of background we'll pull from is convex analysis. We'll look at functions $E : \mathcal{X} \to \overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$, where we treat $\pm\infty$ as elements of $\overline{\mathbb{R}}$ such that $x < +\infty$ and $x > -\infty$ for all $x \in \mathbb{R}$.

We define operations with these elements (where $x \in \mathbb{R}$ and $\lambda > 0$) as follows:

$$x \pm \infty = \pm\infty + x = \pm\infty \tag{10.9}$$
$$\lambda \cdot (\pm\infty) = (\pm\infty) \cdot \lambda = \pm\infty \tag{10.10}$$
$$-1 \cdot (\pm\infty) = \mp\infty \tag{10.11}$$
$$x/(\pm\infty) = 0 \tag{10.12}$$
$$+\infty + \infty = \infty, -\infty - \infty = -\infty \tag{10.13}$$

Note that we don't define $+\infty - \infty$, or $(\pm\infty)(\pm\infty)$.

This is a useful language with which to model constraints. For example, if we want to minimise

$$E : [-1, \infty) \to \mathbb{R} \tag{10.14}$$
$$x \to x^2, \tag{10.15}$$

we can model this as

$$\tilde{E} : \mathbb{R} \to \mathbb{R} \tag{10.16}$$
$$\tilde{E}(x) = \begin{cases} x^2 & x \geq -1 \\ \infty & \text{else} \end{cases}. \tag{10.17}$$

This makes theoretical arguments easier, as $\tilde{E}(x+y)$ is always well defined, and it makes practical implementations easier, as unconstrained optimisation is in general easier than constrained optimisation. This has the drawback of making it so that $\tilde{E}$ is not everywhere differentiable.

Definition 10.6. *The characteristic function $X_C : \mathcal{X} \to \overline{\mathbb{R}}$ of a set $C \subset \mathcal{X}$*

$$X_C(u) = \begin{cases} 0 & u \in C \\ \infty & \textit{else} \end{cases} \tag{10.18}$$

We can use this to write constrained optimisation problems as unconstrained ones.

$$\min_{u \in C} E(u) \leftrightarrow \min_{u \in X} E(u) + X_C(u). \tag{10.19}$$

With this, we get the drawback of not having a well-defined domain, so we make a definition to deal with that.

Definition 10.7. *Let $\mathcal{X}$ be a vector space and let $E : \mathcal{X} \to \overline{\mathbb{R}}$. Then the effective domain of E is*

$$\mathrm{dom}(E) = \{u \in \mathcal{X} \mid E(u) < \infty\}. \tag{10.20}$$

Definition 10.8. *A functional E is proper if its effective domain is not empty.*

Inverse Problems **Lent 2022**

Lecture 11: Convexity, Bregman distance

Lecturer: Malena Sabaté Landman *24 February* *Aditya Sengupta*

11.1 Convexity

Definition 11.1. *Let $\mathcal{X}$ be a vector space. A subset $C \subset \mathcal{X}$ is called convex if $\lambda u + (1-\lambda)v \in C$ for all $u, v \in C$ and $\lambda \in (0,1)$.*

We can extend this concept to functionals.

Definition 11.2. *A functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ is called convex if*

$$E(\lambda u + (1-\lambda)v) \leq \lambda E(u) + (1-\lambda)E(v), \tag{11.1}$$

for all $u, v \in dom(E)$ with $u \neq v$ and for all $\lambda \in (0,1)$. It is called strictly convex if the inequality is strict, and it is called strongly convex with constant θ if $E(u) - \theta\|u\|^2$ is convex.

We have that strong convexity implies strict convexity implies convexity.

Example 11.1. $f : \mathbb{R} \to \mathbb{R}, x \mapsto |x|$ is convex.

□

Example 11.2. $g : \mathbb{R} \to \mathbb{R}, x \mapsto x^4$ is strictly convex, but not strongly convex as $x^4 - \theta x^2$ has roots at $\pm\sqrt{\theta}$ and at 0, creating non-convex bumps.

□

Example 11.3. $h : \mathbb{R} \to \mathbb{R}, x \mapsto x^2$ is strongly convex.

□

Lemma 11.1. *Let $\alpha \geq 0$ and $E, F : \mathcal{X} \to \overline{\mathbb{R}}$ be two convex functionals. Then $E + \alpha F : \mathcal{X} \to \overline{\mathbb{R}}$ is convex. Moreover, if $\alpha > 0$ and F is strictly convex, then $E + \alpha F$ is strictly convex.*

The second property is really useful: even if we start with something that is just convex (not strictly), we can perturb it a bit with something that's strictly convex to get a strictly convex result.

Definition 11.3. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a functional. The Fenchel conjugate of E is defined as*

$$E^* : \mathcal{X}^* \to \overline{\mathbb{R}}, E^*(p) = \sup_{u \in \mathcal{X}} \left[\langle p, u \rangle - E(u) \right]. \tag{11.2}$$

Theorem 11.2. *For any functional $E : \mathcal{X} \to \overline{\mathbb{R}}$, the following inequality holds:*

$$E^{**}(u) = (E^*)^*(u) \leq E(u) \; \forall u \in \mathcal{X}. \tag{11.3}$$

If E is proper, lower semicontinuous, and convex, then $E^{**} = E$.

11.2 Subgradients

Definition 11.4. *A functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ is called subdifferentiable at $u \in \mathcal{X}$ if there exists an element $p \in \mathcal{X}^*$ such that*

$$E(v) \geq E(u) + \langle p, v - u \rangle \; \forall v \in \mathcal{X}. \tag{11.4}$$

We call p the subgradient at u.

Definition 11.5. *The collection of subgradients of E at u is called the subdifferential of E at u:*

$$\partial E(u) = \{ p \in X^* \mid E(v) \geq E(u) + \langle p, v - u \rangle \; \forall v \in \mathcal{X} \}. \tag{11.5}$$

The intuitive idea here is that a straight line always has to go below the function we're looking at. A straight line is a (possibly bad) approximation from below.

For a differentiable functional, the subdifferential consists of just one element, so this isn't a very interesting idea (for example, $x \mapsto x^2$ has just the slope $2x$ with the appropriate intercept at whatever point you choose).

For nondifferentiable functionals, the subdifferential is multivalued. For example, consider $E : \mathbb{R} \to \mathbb{R}, u \mapsto |u|$. The subdifferential is

$$\partial E(u) = \begin{cases} \{1\} & u > 0 \\ [-1, 1] & u = 0 \\ \{-1\} & u < 0 \end{cases} \tag{11.6}$$

because you can put any line between $y = x$ and $y = -x$ down and have it lower-bound $|x|$ at $x = 0$.

If a convex functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ is proper, then for all $u \notin \mathrm{dom}(E)$, the subdifferential is empty.

Proposition 11.3. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a convex functional and let $u \in \mathrm{dom}(E)$ such that E is continuous at u. Then $\partial E(u) \neq \varnothing$.*

Theorem 11.4. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a proper convex functional and $u \in \mathrm{dom}(E)$. Then $\partial E(u)$ is a weak-* compact convex subset of $\mathcal{X}^*$.*

Theorem 11.5. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ and $F : \mathcal{X} \to \overline{\mathbb{R}}$ be proper, lsc convex and we have $u \in \mathrm{dom}(E) \cap \mathrm{dom}(F)$ such that E is continuous at u. Then $pd(E + F) = \partial E + \partial F$.*

Theorem 11.6. *An element $u \in \mathcal{X}$ is a minimiser of the functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ iff $0 \in \partial E(u)$.*

Proof .
By definition, $0 \in \partial E(u) \iff$ for all $v \in \mathcal{X}, E(v) \geq E(u) + \langle 0, v - u\rangle = E(u)$.

□

Theorem 11.7. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a convex function and $E^* : \mathcal{X}^* \to \overline{\mathbb{R}}$ be its convex conjugate. Then*

$$p \in \partial E(u) \iff E(u) + E^*(p) == \langle p, u\rangle. \tag{11.7}$$

11.3 Bregman distances

Definition 11.6. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a convex functional. Moreover, let $u, v \in \mathcal{X}$ such that $E(v) < \infty$ and $q \in \partial E(v)$. Then the (generalised) Bregman distance of E between u and v is defined as*

$$D_E^q(u, v) = E(u) - E(v) - \langle q, u - v\rangle. \tag{11.8}$$

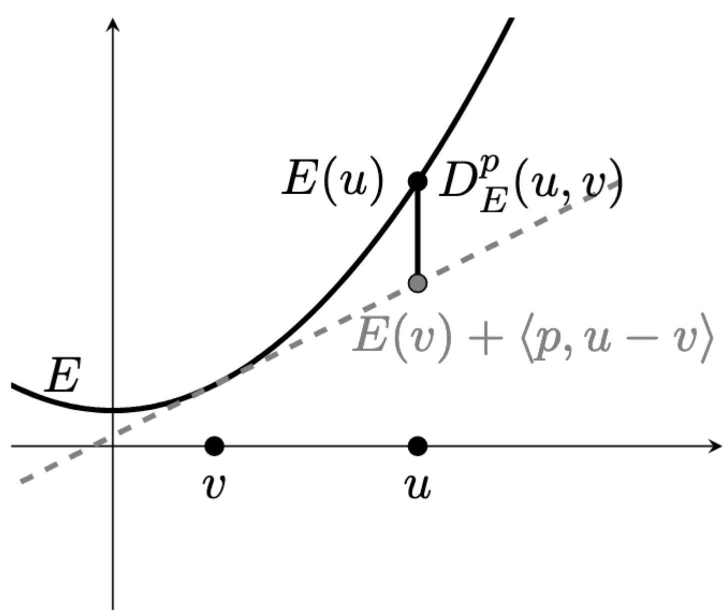

Note that this looks almost like a metric, because it's always greater than or equal to 0, and the distance between a point and itself is 0. However, in general, $D_E^q(u,v) = 0 \not\Rightarrow u = v$. Further, this isn't symmetric.

However, for some specific cases, the Bregman distance is actually a metric. For example, for $E(u) = \frac{1}{2}\|u\|_{\mathcal{X}}^2$, if $\mathcal{X}$ is a Hilbert space, we have a familiar metric:

$$D_E^q(u,v) = \frac{1}{2}\|u-v\|_{\mathcal{X}}^2. \tag{11.9}$$

To fix this in general, we introduce the *symmetric Bregman distance*.

Definition 11.7. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a convex functional and let $u, v \in \mathcal{X}$. Let $E(u) < \infty, E(v) < \infty, q \in \partial E(v), p \in \partial E(u)$. Then the symmetric Bregman distance of E between u and v is defined as*

$$D_E^{symm}(u,v) := D_E^q(u,v) + D_E^p(v,u) = \langle p - q, u - v \rangle. \tag{11.10}$$

Inverse Problems **Lent 2022**

Lecture 12: Direct method for optimisation

Lecturer: Malena Sabaté Landman *1 March* *Aditya Sengupta*

Definition 12.1. *Let $E : \mathcal{X} \to \mathbb{R}$ be a convex functional and let $u, v \in \mathcal{X}$, $E(u) < \infty, E(v) < \infty$. Let $q \in \partial E(v)$ and $p \in \partial E(u)$. Then, the symmetric Bregman distance of E between u and v is defined as*

$$D_E^{symm}(u, v) = D_E^q(u, v) + D_E^p(v, u) = \langle p - q, u - v\rangle. \tag{12.1}$$

Definition 12.2. *A functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ is called absolutely one-homogeneous if $E(\lambda u) = |\lambda| E(u)$ for all $\lambda \in \mathbb{R}, x \in \mathcal{X}$.*

Note that absolutely one-homogeneous convex functionals have $E(0) = 0$.

Proposition 12.1. *Let E be a convex absolutely one-homogeneous functional and let $p \in \partial E(u)$. Then $E(u) = \langle p, u\rangle$.*

Note that $D_E^p(v, u) = E(v) - \langle p, v\rangle$.

Proposition 12.2. *Let $E : \mathcal{X} \to \mathbb{R}$ be a proper, convex, l.s.c. and abs one-hom functional. Then the Fenchel conjugate E^* is the characteristic function of the convex set $\partial E(0)$.*

Proposition 12.3. *For any $u \in \mathcal{X}, p \in \partial E(u)$ iff $p \in \partial E(0)$ and $E(u) = \langle p, u\rangle$.*

12.1 Minimisers

Definition 12.3. *Let $X \to \overline{\mathbb{R}}$ be a functional. We say that $u^* \in \mathcal{X}$ solves the minimisation problem $\min_{u \in \mathcal{X}} E(u)$ when $E(u^*) < \infty$ and $E(u^*) \leq E(u)$ for all $u \in \mathcal{X}$. We call u^* the minimiser of E.*

Definition 12.4. *A functional $E : \mathcal{X} \to \overline{\mathbb{R}}$ is bounded from below if there exists a constant $C > -\infty$ such that for all $u \in \mathcal{X}$, $E(u) \geq C$.*

For minimiser existence, it's important to place certain "sanity conditions", because we can create convergent sequences whose image under E aren't necessarily bounded. One such condition is *coercivity*.

Definition 12.5. *$E : \mathcal{X} \to \overline{\mathbb{R}}$ is coercive if for all sequences $\{u_j\}_{j \in \mathbb{N}}$ with $\|u_j\|_{\mathcal{X}} \to \infty$ we have $E(u_j) \to \infty$.*

Equivalently, a coercive functional is one such that if $\{E(u_j)\}_{j \in \mathbb{N}} \subset \mathbb{R}$ then so is $\{u_j\}_{j \in \mathbb{N}} \subset \mathcal{X}$.

Lemma 12.4. *Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a proper, concave functional bounded from below. Then, the infimum $\inf_{u \in \mathcal{X}} E(u)$ exists in $\mathbb{R}$, there are minimising sequences (i.e. $\{u_j\}_{j \in \mathbb{N}} \subset \mathcal{X}$ with $E(u_j) \to \inf_{u \in \mathcal{X}} E(u)$) and all minimising sequences are bounded.*

Proof .
As E is proper and bounded from below, there exists $C_1 > 0$ such that $-\infty < -C_1 < \inf_u E(u) < \infty$. Then there exist minimising sequences.

$\square$

Let $\{v_j\}_{j\in\mathbb{N}}$ be any minimising sequence. Then there exists $j_0 \in \mathbb{N}$ such that for $v_{j>j_0}$ we have that $E(v_j) \leq \underbrace{\inf E(u) + 1}_{C_2} < \infty$. Therefore $|E(v_j)| < \max\{C_1, C_2\}$ for all $j > j_0$ and from the coercivity it follows that $\{v_j\}_{j\in j_0}$ is bounded, and so $\{u_j\}_{j\in\mathbb{N}}$ is bounded.

Theorem 12.5 ("Direct method", or the fundamental theorem of optimisation)**.** *Let $\mathcal{X}$ be a Banach space and τ_X be a topology on $\mathcal{X}$ (not necessarily the one induced by the norm), such that bounded sequences have τ_X-convergent subsequences. Let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be proper, bounded from below, coercive, and τ_X lsc. Then E has a minimiser.*

Proof .
We know that $\inf_{u\in\mathcal{X}} E(u)$ is finite, and that minimising sequences exist and that they are finite. Let $\{u_j\}_{j\in\mathbb{N}} \subset \mathcal{X}$ be a minimising sequence. From the assumptions on the topology, there exist $\{u_{j_k}\}_{k\in\mathbb{N}}$ and $u^* \in \mathcal{X}$, such that $u_{j_k} \xrightarrow{\tau_x, k\to\infty} u^*$. From the l.s.c. property of E,

$$E(u^*) \leq \liminf_{k\to\infty} E(u_{j_k}) = \lim_{j\to\infty} E(v_j) = \inf_{u\in\mathcal{X}} E(u) < \infty. \tag{12.2}$$

So $E(u^*) < \infty$, $E(u^*) \leq E(u)$ for all $u \in \mathcal{X}$, and therefore u^* is the minimiser of E.

$\square$

There are some cases where these assumptions aren't really satisfied.

Corollary 12.6. *Let $\mathcal{X}$ be a reflexive Banach space, and let $E : \mathcal{X} \to \overline{\mathbb{R}}$ be a functional which is proper, bounded from below, coercive, and l.s.c. with respect to the weak topology. Then the functional E has a minimiser.*

This is just the direct method without the constraint on τ_X sequences, because that's provided by $\mathcal{X}$ being reflexive.

A convex function is lower semi-continuous with respect to the weak topology iff it's lower semi-continuous with respect to the strong topology.

Basically, all we've done is apply the Bolzano-Weierstrass theorem from $\mathbb{R}$ to general Banach spaces, by saying that bounded sequences need to have convergent subsequences for unique minimisers to work.

Theorem 12.7. *Assume $E : \mathcal{X} \to \overline{\mathbb{R}}$ is strictly convex and has at least one minimiser. Then the minimiser is unique.*

Proof .
Let u, v be two minimisers of E, $u \neq v$. Then

$$E(u) \leq E\left(\frac{1}{2}u + \frac{1}{2}v\right) < \frac{1}{2}E(u) + \frac{1}{2}E(v) \leq E(u), \tag{12.3}$$

as $E(v) \leq E(u)$ (wlog).

$\square$

Inverse Problems **Lent 2022**

Lecture 13: Duality

Lecturer: Malena Sabaté Landman *3 March* *Aditya Sengupta*

13.1 Duality in convex optimisation

Consider the following optimisation problem: let $E : \mathcal{Y} \to \overline{\mathbb{R}}, F : \mathcal{X} \to \overline{\mathbb{R}}$ be proper, convex, and lsc. Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ be linear and bounded.

$$P : \inf_{uin\mathcal{X}} (E(Au) + F(u)). \tag{13.1}$$

We call P the *primal* problem.

Since E is convex and lower semi-continuous, the conjugate of the conjugate is the functional itself: $E = (E^*)^*$. As a reminder, the Fenchel conjugate of E is given by

$$E^*(y) = \sup_{y \in \mathcal{Y}} (\langle p, y \rangle - E(y)). \tag{13.2}$$

This allows us to rewrite E,

$$E(y) = (E^*)^*(y) = \sup_{\eta \in \mathcal{Y}^*} (\langle \eta, y \rangle - E^*(\eta)), \tag{13.3}$$

and using this we can also rewrite P:

$$P : \inf_{u \in \mathcal{X}} \left(\sup_{\eta \in Y^*} (\langle \eta, Au \rangle - E^*(\eta)) \right) + F(u) \tag{13.4}$$

$$P : \inf_{u \in \mathcal{X}} \sup_{\eta \in Y^*} (\langle \eta, Au \rangle - E^*(\eta) + F(u)). \tag{13.5}$$

We call this form of P the *saddle point problem*.

Optimising over two spaces at once is awkward, so we'd like to change it. We can use the property that $\inf \sup \geq \sup \inf$ to bound P by some other problem that may be more tractable.

$$\begin{aligned}
\inf_{u \in \mathcal{X}} (E(Au) + F(u)) &\geq \sup_{\eta \in \mathcal{Y}^*} \inf_{u \in \mathcal{X}} (\langle \eta, Au \rangle - E^*(\eta) + F(u)) && (13.6)\\
&= \sup_{\eta \in \mathcal{Y}^*} \inf_{u \in \mathcal{X}} (\langle A^* \eta, u \rangle - E^*(\eta) + F(u)) && (13.7)\\
&= \sup_{\eta \in Y^*} (-E^*(\eta) - \sup [\langle -A^* \eta, u \rangle - F(u)]), && (13.8)
\end{aligned}$$

where in the last step we've used the property $\inf x = -\sup(-x)$, along with the lack of dependence of $E^*(\eta)$ on the supremum over u.

We end up with the Fenchel conjugate of F, so we can rewrite to get rid of the inner sup and get

$$D: \sup_{\eta \in \mathcal{Y}^*} (-E^*(\eta) - F^*(A^*\eta)).$$

This is the *dual problem*. Due to the $\geq$ in the middle, the optimal value of the primal problem is always greater than or equal to the optimal value of the dual problem; if it's less, we have *weak duality*, but if they're equal we have *strong duality*. The difference between the two is called the *dual gap*.

Theorem 13.1. *Suppose that*

(i) the functional $E(Au) + F(u) : \mathcal{X} \to \overline{\mathbb{R}}$ is open, convex, lsc, and coercive.

(ii) there exists some $u_0 \in \mathcal{X}$ such that $F(u_0) < \infty, E(Au_0) < \infty$, and $E(y)$ is continuous at $y = Au_0$.

Then

1. *the dual problem D has at least one solution $\hat{\eta}$;*
2. *there is no dual gap; and*
3. *if P has an optimal solution $\hat{u}$, then $-A^*\hat{\eta} \in \partial F(\hat{u})$, $\hat{\eta} \in \partial E(A\hat{u})$.*

These conditions show the existence of a solution to D and characterise a solution to P if it exists, but they do not ensure a solution to P exists.

13.2 Well-posedness and regularisation properties

We are interested in the properties of the variational regularisation minimisation problem:

$$\min_{u \in \mathcal{X}} \|Au - f\|^2 + \alpha J(u) \tag{13.9}$$

Specifically, we want to find when this is a convergent minimiser of $Au = f$. Here, $A : \mathcal{X} \to \mathcal{Y}$ is a linear bounded operator, $\mathcal{Y}$ is a Banach space, and $\mathcal{X}$ is the dual of a separable Banach space. We'll look at the existence of minimisers and at parameter choice rules for α, and how the minimisers converge to a regularised solution to the inverse problem.

Definition 13.1. *Let $u_J^\dagger$ be a least-squares solution to $Au = f$. That is, $\left\|Au_J^\dagger - f\right\|_{\mathcal{Y}} = \inf_{v \in \mathcal{X}} \|Av - f\|_{\mathcal{Y}}$. If $J(u_J^\dagger) \leq J(\tilde{u})$ for all least-squares solutions $\tilde{u}$, we say that $u_J^\dagger$ is a J-minimising solution of $Au = f$.*

We'll usually assume there is a least-squares solution having a finite value of J. It may not be unique, so we may have to select one using some selection operator later, but we won't worry about this too much.

First, let's clarify what kinds of functionals we're thinking about.

Lemma 13.2. *Let $J(u) = \sum_{i=1}^n J_i(u)$, where each J_i is convex and p_i-homogeneous, meaning that for $p_i > 0$, we have*

$$J_i(\lambda u) = |\lambda|^{p_i} J_i(u)\ \forall u \in \mathcal{X}, \lambda \in \mathbb{R} \tag{13.10}$$

Then the null set $N(J) := \{u \in \mathcal{X} \mid J(u) = 0\}$ is a linear subspace of $\mathcal{X}$

We'll always end up looking at regularising functionals of this kind.

Proof .
First, we can see that $J_i(u) \geq 0$ identically as follows:

$$\begin{aligned} 0 &= J_i(0) \quad p_i \text{ homogeneity with } \lambda = 0 \\ &= J_i\left(\frac{1}{2}u - \frac{1}{2}u\right) \\ &\leq \frac{1}{2}J_i(u) + \frac{1}{2}J_i(-u) \quad \text{convexity} \\ &= J_i(u) \quad p_i \text{ homogeneity with } \lambda = -1 \quad . \end{aligned} \tag{13.11}$$

We use this to show that if $u, v \in N(J)$, so is $\lambda u + v$. Let $u, v \in N(J)$. Then $J_i(u) = 0$, $J_i(v) = 0$ for $i = 1, \ldots, n$. Therefore

$$\begin{aligned} 0 &\leq J_i(\lambda u + v) \quad \text{as above} \\ &= 2^{p_i} J_i\left(\frac{\lambda u}{2} + \frac{v}{2}\right) \quad p_i \text{ homogeneity with } \lambda = 2 \\ &\leq 2^{p_i}\left(\frac{1}{2}J_i(\lambda u) + \frac{1}{2}J_i(v)\right) \quad \text{convexity} \\ &= 2^{p_i - 1}(J_i(\lambda u) + J_i(v)) \quad p_i \text{ homogeneity} \\ &= 2^{p_i - 1}(|\lambda|^{p_i} J_i(u) + J_i(v)) = 0. \end{aligned} \tag{13.12}$$

Therefore $J_i(\lambda u + v) = 0$ for all i, so $J(\lambda u + v) = 0$. $N(J)$ therefore satisfies superposition and scaling, so it is a linear subspace, which was what we wanted.

□

Lemma 13.3. *Let J be defined as above. Suppose $u \in \mathcal{X}$ and $v \in N(J)$. Then $J(u + v) = J(u)$.*

If $N(J)$ is finite-dimensional, then it is complemented in $\mathcal{X}$, meaning there exists a closed subspace $\mathcal{X}_0 \subset \mathcal{X}$ such that $\mathcal{X} = \mathcal{X}_0 \oplus N(J)$ and $\mathcal{X}_0 \cap N(J) = \{0\}$.

This is building towards an idea where we decompose elements of $\mathcal{X}$ into those that do and don't vanish under J. Let's add more lemmas to make this easier:

Lemma 13.4. *Suppose $J : \mathcal{X} \to \overline{\mathbb{R}}_+$ as defined above is proper (and from above, convex and the sum of p_i-homogeneous terms). Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ be a bounded linear operator. Further, suppose* *(i)* $\dim N(J) < \infty$ *and* J *is coercive on* $\mathcal{X}_0$ *(where* $\mathcal{X}_0$ *is defined as above)*

(ii) the null spaces of A *and* J *intersect only at 0.*

Then, the function

$$\phi_\alpha(u) = \frac{1}{2}\|Au - f\|_{\mathcal{Y}}^2 + \alpha J(u) \tag{13.13}$$

is coercive on $\mathcal{X}$ *for any* $\alpha > 0$.

Coercivity is one of the key tools that we saw earlier to guarantee a minimiser exists, so we're building up these conditions to ensure we have that.

Proof .
Recall that coercivity ensures that if $\{\phi_\alpha(u_j)\}_{j\in\mathbb{N}}$ is bounded, then $\{u_j\}_{j\in\mathbb{N}} \subset \mathcal{X}$ have bounded norms.
Let $\{u_j\}_{j\in\mathbb{N}} \subset \mathcal{X}$. We use the decomposition $\mathcal{X} = \mathcal{X}_0 + N(J)$: for each u_j there are unique $u_j^0 \in \mathcal{X}_0, u_j^N \in N(J)$ such that $u_j = u_j^0 + u_j^N$.
Next, suppose $\{\phi_\alpha(u_j)\}_{j\in\mathbb{N}}$ is bounded, that is, there is some C such that $\phi_\alpha(u_j) \leq C$ for all j. This gives us $J(u_j) \leq C$. We know that J is coercive in $\mathcal{X}_0$, so we can say $\|u_j^0\| \leq C'$.
It remains to show the boundedness of u_j^N. To do this, we need to consider A. We restrict our linear operator A to the relevant region by working with the restriction $\tilde{A} = A\big|_{N(J)}$, so that its domain is $N(J)$ and its codomain is the range $AN(J) \subset \mathcal{Y}$. By this restriction, A is surjective, and since the null spaces of A and J have a trivial intersection, A is also injective. Therefore $\tilde{A}^{-1}$ exists and is bounded. For convenience, let $\left\|\tilde{A}^{-1}\right\| = \tilde{C}$.
We can now show that u_j^N has bounded norm:

$$\|u_j^N\| = \left\|\tilde{A}^{-1}(\tilde{A}u_j^N)\right\| \leq \tilde{C}\|Au_j^N\| \tag{13.14}$$

So it suffices to prove that Au_j^N is bounded:

$$\|Au_j^N\| = \|Au_j^N + Au_j^0 - f - (Au_j^0 - f)\| \leq \|Au_j - f\| + \|Au_j^0 - f\| \leq C + \|A\|\|u_j^0\| + \|f\| \leq C'' \tag{13.15}$$

where in the last step we used the boundedness of A and u_j^0, and that $\|Au_j - f\|$ is upper-bounded by the condition from the start on ϕ_α.

□

Inverse Problems **Lent 2022**

Lecture 14: Solution convergence, regularisers

Lecturer: Malena Sabaté Landman *8 March* *Aditya Sengupta*

Finally, we can show there is a $J-$minimising solution and a regularised solution for any choice of $\alpha > 0$.

Theorem 14.1. *Let $\mathcal{X}, \mathcal{Y}$ be Banach spaces, and let $\tau_{\mathcal{X}}, \tau_{\mathcal{Y}}$ be any topologies on them (not necessarily the ones induced by the norm). Suppose*

(i) bounded sequences in $\mathcal{X}$ converge in $\tau_{\mathcal{X}}$;

(ii) $J : \mathcal{X} \to \overline{\mathbb{R}}_+$ is proper, convex, $\tau_{\mathcal{X}}$ l.s.c., and satisfies the prior assumptions (is a sum of convex p_i-homogeneous components);

(iii) $A : \mathcal{X} \to \mathcal{Y}$ is continuous in $\tau_{\mathcal{X}}, \tau_{\mathcal{Y}}$ (open sets in τ_X are mapped to open sets in $\tau_{\mathcal{Y}}$);

(iv) the norm $\|\cdot\|_{\mathcal{Y}}$ is lower semi-continuous in $\tau_{\mathcal{Y}}$.

Then there exists a $J-$minimising solution $u_J^\dagger$ of the inverse problem. Further, for any $\alpha > 0$ and $f \in \mathcal{Y}$, there exists a minimiser

$$u^\alpha \in \arg\min_{u \in \mathcal{X}} \frac{1}{2}\|Au - f\|_{\mathcal{Y}}^2 + \alpha J(u). \tag{14.1}$$

Proof .
We'll write the set of least-squares solutions as a sublevel set:

$$L = \{u \in \mathcal{X} \mid \|Au - f\|_{\mathcal{Y}} \leq \mu\},$$
$$\mu = \inf\{\|Av - f\|_{\mathcal{Y}} \mid v \in \mathcal{X}\}.$$

This rewriting seems a bit redundant, as we can only have equality with μ and not $<$, but it actually allows us to make use of the continuity of A. Since A is $\tau_{\mathcal{X}} \to \tau_{\mathcal{Y}}$ continuous and $\|\cdot\|_{\mathcal{Y}}$ is lower semi-continuous in $\tau_{\mathcal{Y}}$, we have that L is closed in $\tau_{\mathcal{X}}$.
For some intuition on this step, first note that level sets of lower semi-continuous functions are closed. A level set of $\|Au - f\|_{\mathcal{Y}}$ will either pass through a continuous point of the norm, in which case it is closed as it includes the limit point μ, or it will pass through a discontinuity, in which case lower semi-continuity means we include the limit point below. Without lower semi-continuity, the $\leq \mu$ may not have included a limit point, and the set may not be closed.
Consider the following infimum:

$$\inf_{u \in L} J(u) = \inf_{u \in \mathcal{X}} J(u) + \chi_L(u) \tag{14.2}$$

where χ_L is an indicator function for L. Assuming this is feasible, there exists a candidate solution $u \in L$ with $J(u) < \infty$. Since Equation 14.2 is bounded from below (by 0, at least), we proceed as in the proof of coercivity of ϕ_α to show that it is coercive.

Next, since L is closed in $\tau_{\mathcal{X}}$, the indicator function χ_L is l.s.c. (try an example after going back to the definition of lower semi-continuity). Since we've taken J to also be l.s.c., we can conclude that Equation 14.2 is l.s.c. in $\tau_{\mathcal{X}}$. Therefore, by the direct method, a minimiser exists.
Next, we show that ϕ_α has a minimiser. By the preceding lemma we know that ϕ_α is coercive. Since it is a combination of a $\tau_{\mathcal{X}}$ l.s.c. function (J) and a $\tau_{\mathcal{Y}}$ l.s.c. function $\|\cdot\|_{\mathcal{Y}}$ composed with a $\tau_{\mathcal{X}} \to \tau_Y$ continuous function A, we can see that ϕ_α is l.s.c., so the direct method applies and a minimiser exists.

□

Having developed the theory for a general regularising parameter α, we can finally go back to a realistic scenario. We consider $f \to f_\delta$ where δ parameterises the level of noise, and relate α to a noise level δ in the measurements.

In this case, we aim to solve

$$\min_{u\in\mathcal{X}} \frac{1}{2}\|Au - f_\delta\|_{\mathcal{Y}}^2 + \alpha(\delta)J(u). \tag{14.3}$$

We are interested in studying the behaviour as $\delta \to 0$, because we want to ensure that everything is well-behaved in the limit of no noise.

For simplicity, we'll assume our least-squares solutions are actual solutions.

Theorem 14.2. *Let the conditions of Theorem 14.1 hold and suppose that* $\inf\{\|Av - f\|_{\mathcal{Y}} \mid v \in \mathcal{X}\} = 0$. *Let* $\alpha = \alpha(\delta)$ *satisfy*

$$\lim_{\delta\to 0} \alpha(\delta) = 0 \quad \textit{and} \quad \limsup_{\delta\to 0} \frac{\delta^2}{\alpha(\delta)} = 0. \tag{14.4}$$

Then $u_\delta := u_\delta^{\alpha(\delta)} \xrightarrow{\tau_{\mathcal{X}}} u_J^\dagger$ *as* $\delta \to 0$ *(possibly along a subsequence) and* $J(u_\delta) \to J(u_J^\dagger)$*, where* $u_J^\dagger$ *is a* $J-$*minimising solution.*

The proof essentially consists of checking conditions for previous theorems and applying them. This theorem proves convergence in $\tau_{\mathcal{X}}$, but this might not be a strong topology and in general isn't the one induced by the norm. Howver, we can get convergence in the norm topology if we require that J satisfy the Radon-Riesz property, which says that if $u_j \xrightarrow{\tau_{\mathcal{X}}} u$ and $J(u_j) \to J(u)$ then $\|u_j - u\| \to 0$.

For example, the norm in a Hilbert space, where $\tau_{\mathcal{X}}$ is the weak topology, satisfies this property.

Examples of regularisers include the norm in a Hilbert space $J(u) = \|u\|^2$. This is weakly l.s.c., and we know (from just after Banach-Alaoglu) that norm-bounded sequences have weakly convergent subsequences, so the regularised solutions are weakly convergent. Further, since the norm has the Radon-Riesz property, this implies convergence in the norm (strong convergence). This is sometimes referred to as the *smoothing functional*, as when we let the Hilbert space be H^1 (continuous and weakly differentiable, with continuous derivative) it ensures that regularised solutions are H^1 as well.

Another example is the ℓ^1 regulariser, $J(u) = \sum_i |u_i|$. This operates on the Banach space $\mathcal{X} = \ell^2$ of square-summable sequences. This is helpful in sparse settings, i.e. if $u \in \ell^2$ has finite support. Geometrically, we can see the relationship between these two as follows.

Let $f = Au, u \in \mathbb{R}^2$, and let $Au = f$ be indeterminate, with infinite solutions. These form an affine subspace, say a line $L \subset \mathbb{R}^2$. If we wanted to find a solution $u \in L$ that minimised the 2-norm $\|u\|_2 = \sqrt{u_1^2 + u_2^2}$, we would conceptually start with $r = 0$ and increase r until the circle $x^2 + y^2 = r^2$ intersected L. This is easily generalised past $\mathbb{R}^2$. However, this approach doesn't favour sparsity, which we may care about for physical reasons. We could also look at the "0-norm", which just counts how many components are nonzero, but this is NP-hard. Geometrically, we look for intersections of L with the coordinate axes, but it's hard to scale this up past two dimensions. Can we make a relaxation that includes both? Yes: this is the ℓ_1 norm. This works similarly to ℓ_2, but expanding outwards in diamonds $|x| + |y| = r$. In the $\mathbb{R}^2$ case, we can see that this approach finds a sparse solution (intersecting a coordinate axis) without the computational difficulty of looking for "zero hyperplanes".

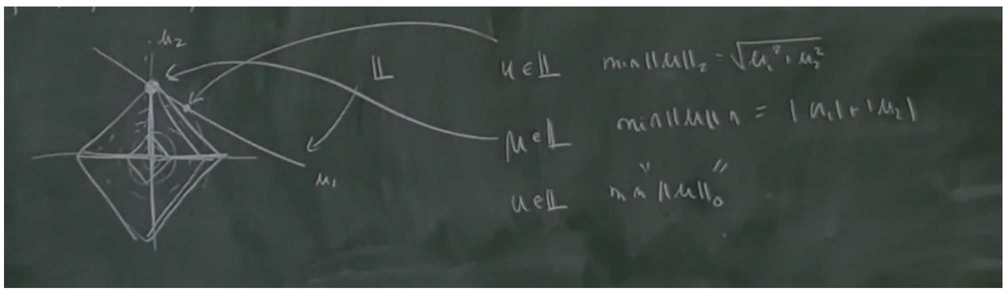

We can also take combinations of these regularisers as long as they are individually p_i-homogeneous. For instance, *elastic regularisation* is characterised by

$$J(u) = \alpha\|u\|_1 + \beta\|u\|_2^2. \tag{14.5}$$

Inverse Problems **Lent 2022**

Lecture 15: Total variation

Lecturer: Malena Sabaté Landman *10 March* *Aditya Sengupta*

In this lecture, we'll look at a particular choice of regulariser, the *total variation*. The reason why we're interested in it is it allows for reasonably regular reconstructions on top of underlying data that may be discontinuous.

Consider the set of vector-valued test functions on a domain $\Omega \subset \mathbb{R}^n$,

$$D(\Omega) = \{\varphi \in C_0^\infty(\Omega) \mid \sup_{x\in\Omega} \|\varphi(x)\|_2 \leq 1\}. \tag{15.1}$$

Definition 15.1. *The total variation of $u \in L^1(\Omega)$ is*

$$\mathrm{TV}(u) = \sup_{\varphi \in D(\Omega)} \int_\Omega u(x) \operatorname{div} \varphi(x) \mathrm{d}x. \tag{15.2}$$

This seems kind of awkward to work with, but we can make it more natural by adding some assumptions. We won't always make these assumptions; we're just introducing them for now to make TV seem more intuitive.

Suppose $u \in W^{1,1}(\Omega)$; that is, u is in $L^1(\Omega)$ and has a weak derivative ∇u that is also in $L^1(\Omega)$. This lets us integrate by parts:

$$\mathrm{TV}(u) = \sup_{\varphi \in D(\Omega)} \int_\Omega -\langle \nabla u(x), \varphi(x)\rangle \mathrm{d}x. \tag{15.3}$$

Recall that φ is vector-valued and so is the gradient ∇u of the scalar function u, so this is a vector dot product. By Cauchy-Schwarz, we can say

$$|\nabla u(x), \varphi(x)| \leq \|\nabla u(x)\|_2 \|\varphi(x)\|_2 \leq \|\nabla u(x)\|_2, \tag{15.4}$$

almost everywhere for $x \in \Omega$. We can make a choice of φ that achieves equality:

$$\varphi(x) = -\frac{\nabla u(x)}{\|\nabla u(x)\|_2} \tag{15.5}$$

This may not live in D, but we can proceed as if it does by approximation. Therefore, the supremum is attained at this value, and the TV becomes

$$\mathrm{TV}(u) = \int_\Omega \|\nabla u(x)\|_2 \mathrm{d}x := \|\nabla u\|_1. \tag{15.6}$$

So we see that TV enforces sparsity on the gradient. A sparse gradient would be 0 in most places, but with occasional big jumps, so regularising by TV favours, for example, piecewise constant solutions.

Note that the space of functions with bounded total variation is actually bigger than the Sobolev space $W^{1,1}$. For example, functions with discontinuities aren't in $W^{1,1}$, but they do have bounded total variations.

Proposition 15.1. *TV is a proper, convex, and absolutely 1-homogeneous functional $L^1(\Omega) \to \overline{\mathbb{R}}$. For any constant function $\mathbf{c}(x) = c \in \mathbb{R}\ \forall x \in \mathbb{R}$ and any $u \in L^1(\Omega)$,*

$$\mathrm{TV}(\mathbf{c}) = 0 \tag{15.7}$$

$$\mathrm{TV}(u + c) = \mathrm{TV}(u). \tag{15.8}$$

Note that $\mathrm{TV}(u) = 0$ implies u is constant, so this goes both ways the null space of TV is exactly the set of constant functions.

Definition 15.2. *The functions $u \in L^1(\Omega)$ with a finite TV form a normed space called the space of bounded variation (the BV-space):*

$$BV(\Omega) := \{u \in L^1(\Omega) \mid \|u\|_{BV} := \|u\|_1 + TV(u) < \infty\} \tag{15.9}$$

It can be shown that this space has some nice properties. It is the dual of a separable Banach space. Further, weak-* convergence $u_n \rightharpoonup^* u$ in BV is equivalent to strong convergence $u_n \to u$ in L^1 and convergence of $\mathrm{TV}(u_n) \to \mathrm{TV}(u)$.

There are a few slightly more technical results: that BV(Ω) is compactly embedded in $L^1(\Omega)$. We can show this using Rellich-Kondrachov: let $\Omega \subset \mathbb{R}^n$ be a nonempty, open, connected domain with Lipschitz boundary. Let $p, k \in \mathbb{N}$. Let

$$p^* := \begin{cases} \frac{np}{n-kp} & \text{if } n > kp \\ \infty & \text{if } n \leq kp. \end{cases} \tag{15.10}$$

Then the embedding $W^{k,p}(\Omega) \to L^q(\Omega)$ is continuous for all $1 \leq q \leq p^*$ and compact for all $1 \leq q < p^*$.

Some functions from BV(Ω) can be approximated by functions in the Sobolev space $W^{1,1}(\Omega)$, and we can embed this in higher L^p spaces using Rellich-Kondrachov. The threshold value here for (integrability 1, differentiability 1) is 1, so for n at least 2, we have compactness. For (integrability 2, differentiability 1), we have a threshold of 2, so for $n = 2$, we get that $p^* = 2$.

Corollary 15.2 (of Rellich-Kondrachov)**.** *For any bounded Lipschitz domain $\Omega \subset \mathbb{R}^n$, the embedding $\mathrm{BV}(\Omega) \Subset L^1(\Omega)$ is compact for any $n \geq 2$, and the embedding $\mathrm{BV}(\Omega) \hookrightarrow L^2(\Omega)$ is continuous for $n = 2$.*

Remember that a compact embedding is also continuous: being compact is a stronger condition.

We make use of this to show what we need to in order to show that TV admits a minimiser.

Theorem 15.3. *Let $\Omega \subset \mathbb{R}^n$ be open and bounded. Then the total variation is l.s.c. in $L^1(\Omega)$.*

Proof .
Let $\{u_j\}_{j\in\mathbb{N}} \subset \mathrm{BV}(\Omega)$ converge in $L^1(\Omega)$ to $u \in L^1(\Omega)$. We want to show that $TV(u) \leq \liminf TV(u_j)$, which we can do by taking a test function $\varphi \in D(\Omega)$ and using the fact that strong convergence implies weak convergence:

$$\int_\Omega u_j(x) \operatorname{div} \varphi(x) \mathrm{d}x \to \int_\Omega u(x) \operatorname{div} \varphi(x) \mathrm{d}x. \tag{15.11}$$

Therefore,

$$TV(u) = \sup_{\varphi \in D(\Omega)} \int_\Omega u(x) \operatorname{div} \varphi(x) \mathrm{d}x \tag{15.12}$$

$$= \sup_{\varphi \in D} \lim_{j\to\infty} \int_\Omega u_j(x) \operatorname{div} \varphi(x) \mathrm{d}x \tag{15.13}$$

Now, we'd like to swap the supremum and the limit, but the limit of the supremum over all sequences may not exist. To fix this, we use the lim inf. The reasoning for this is as follows. We want to work with the sequence $t_j := \sup_\varphi \int u_j(x) \operatorname{div} \varphi(x) \mathrm{d}x$, where we fix u_j and let φ vary element-to-element. t_j may not converge, but it has subsequential limits (that may be infinity). How do these subsequential limits relate to $\mathrm{TV}(u)$? We'd like it if they were all greater than $\mathrm{TV}(u)$, because then in particular we could say their greatest lower bound ($\liminf_{j\to\infty} t_j$) was greater than $\mathrm{TV}(u)$, which would show lower semi-continuity. We *can* in fact say this, with the following reasoning: say the test function optimising the TV argument is φ^*. Because each t_j gets to optimise over the entire space of test functions, we can't have any subsequence (t_{j_k}) converging to something less than $\mathrm{TV}(u)$. If we did, we could replace the choice of test function φ_{j_k} made by each t_{j_k} with φ^*, and we'd have a convergent sequence that was at least $\mathrm{TV}(u)$. So the only subsequential limits that can remain are $\mathrm{TV}(u)$ or more, and therefore the lim inf is at least $\mathrm{TV}(u)$, so we can say

$$\mathrm{TV}(u) \leq \liminf_{j\to\infty} \sup_{\varphi\in D} \int_\Omega u_j(x) \operatorname{div} \varphi(x) \mathrm{d}x \tag{15.14}$$

$$= \liminf_{j\to\infty} \mathrm{TV}(u_j). \tag{15.15}$$

Note that neither side is necessarily finite, because u isn't necessarily in $\mathrm{BV}(\Omega)$, and t_j may diverge to infinity.

□

TV is not actually coercive on L^1, because its null space is nontrivial; specifically, we already saw that it consists of the set of constant functions. This means you could take a sequence converging in the functional, and add arbitrarily large constants to it to create an unbounded sequence whose image converges.

This wasn't really our intent when we were requiring coercivity; we meant to capture more pathological behaviour. In this case, we can characterise the null space really well, so we can "mod out" constant functions and work with some notion of almost-coercivity. We formalise this as follows:

Proposition 15.4. *Let $\Omega \subset \mathbb{R}^n$ be a bounded Lipschitz domain. Then there exists a constant $C > 0$ such that for all $u \in \mathrm{BV}(\Omega)$, the Poincaré inequality*

$$\|u - \overline{u}\|_{L^1} \leq C\,\mathrm{TV}(u) \tag{15.16}$$

holds, where $\overline{u} = \frac{1}{|\Omega|}\int_\Omega u(x)\mathrm{d}x$.

This lets us consider the subspace of zero-mean functions $\mathrm{BV}_0(\Omega) \subset \mathrm{BV}(\Omega)$, defined by

$$\mathrm{BV}_0(\Omega) := \{u \in \mathrm{BV}(\Omega) \mid \int_\Omega u(x)\mathrm{d}x = 0\}. \tag{15.17}$$

The above proposition then lets us show that TV is coercive on this subspace, as it is within a constant factor of the ℓ^1 norm.

It isn't really a practical problem to make this restriction: we just compute the mean of the function in some other way, subtract it off, work with the zero-mean version, and add it back later. TV also has a finite-dimensional null space, so it allows the decomposition $L^1 = L^1_0 \oplus N(\mathrm{TV})$, where $L^1_0 = \{u \in L^1 \mid \int_\Omega u(x)\mathrm{d}x = 0\}$. This is the formal way of saying we can split any integrable function into its zero-mean component and its constant mean component.

Now, we can show that minimisers exist as we want.

Theorem 15.5. *Let $\mathcal{X} = L^1(\Omega)$ where $\Omega \subset \mathbb{R}^n$ is bounded Lipschitz, and let $\mathcal{Y}$ be a Banach space. Let $A : L^1 \to \mathcal{Y}$ be a linear bounded operator such that $A(\mathbf{1}) \neq 0$, where $\mathbf{1}$ is the function that is 1 everywhere. Then minimisers of the problem*

$$\min_{u \in L^1(\Omega)} \frac{1}{2}\|Au - f_\delta\|^2_{\mathcal{Y}} + \alpha(\delta)\,\mathrm{TV}(u) \tag{15.18}$$

exist and converge strongly in L^1 to a TV-minimising solution as $\delta \to 0$ if $\alpha(\delta)$ is chosen as required by the relevant preceding theorem (lecture 14).

The only unusual condition here is the requirement that $A\mathbf{1} \neq 0$, which comes about because we have to enforce no nontrivial intersections of the kernels of A and TV. Since the kernel of TV is constant functions, A needs to not vanish when given a constant function, so we enforce this extra condition.

Inverse Problems **Lent 2022**

Lecture 16: Convex Duality

Lecturer: Malena Sabaté Landman *15 March* *Aditya Sengupta*

In the previous chapter, we established the necessary conditions for a regularised solution u_δ to converge to a $J-$minimising solution $u_J^\dagger$. Here, we ask: how fast does this happen? We want the convergence to be relatively fast, so that these methods are useful in practice.

The usual way of looking at these is via the Bregman distances associated with J. We'd like to establish that

$$D_J^{symm}(u_\delta, u_J^\dagger) \leq \psi(\delta) \xrightarrow{\delta \to 0} 0, \tag{16.1}$$

where $\psi : \mathbb{R}_+ \to \mathbb{R}_+$ is some known function.

16.1 The dual problem

We know that u_δ solves the primal problem

$$\min_{u \in \mathcal{X}} \underbrace{\frac{1}{2}\|Au - f_\delta\|_{\mathcal{Y}}^2}_{E(Au)} + \underbrace{\alpha J(u)}_{F(u)} \tag{16.2}$$

with all the usual assumptions ($\alpha = \alpha(\delta)$, A is linear and bounded between Banach spaces, J satisfies the required assumptions, and both functionals go to the extended real line.) We'll also assume that J is absolutely one-homogeneous, i.e. $J(\lambda u) = |\lambda| J(u)$, and that least-squares solutions exactly satisfy the unregularised noiseless problem, i.e. $Au_J^\dagger = f$.

Lemma 16.1. *Let J be an absolutely one-homogeneous, convex, proper, lsc functional. Then its convex (Fenchel) conjugate J^* is the characteristic function of the convex set $\partial J(0)$.*

Lemma 16.2. *Let $\mathcal{X}$ be a Banach space and $\mathcal{X}^*$ its dual. Let $\varphi(x) = \frac{1}{2}\|x\|_{\mathcal{X}}^2$. Then the convex conjugate of φ is $\varphi^*(\xi) = \frac{1}{2}\|\xi\|_{\mathcal{X}^*}^2$ where $\xi \in \mathcal{X}^*$.*

Using a result from the preceding lectures, we know that $p \in \partial E$ if and only if $E(u) + E^*(p) = \langle p, u \rangle$. So for $x \in \mathcal{X}, \xi \in \partial\varphi(x)$, we get

$$\frac{1}{2}\|x\|_{\mathcal{X}}^2 + \frac{1}{2}\|\xi\|_{\mathcal{X}^*}^2 = \langle \xi, x \rangle \leq \|\xi\|_{\mathcal{X}^*}\|x\|_{\mathcal{X}}, \tag{16.3}$$

and this has the form $\frac{1}{2}(a^2 + b^2) \leq ab$, which is only possible if $a = b$, so this gives us

$$\|\xi\|_{\mathcal{X}^*} = \|x\|_{\mathcal{X}}. \tag{16.4}$$

So we can use this to write the dual of the first term of the primal problem:

$$E^*(\eta) = \sup_{y \in \mathcal{Y}} \langle \eta, y \rangle - \frac{1}{2}\|y - f\|_{\mathcal{Y}}^2 = \langle \eta, f \rangle - \sup_{z \in \mathcal{Y}} \left(\langle \eta, z \rangle - \frac{1}{2}\|z\|_{\mathcal{Y}}^2 \right) = \langle \eta, f \rangle + \frac{1}{2}\|\eta\|_{\mathcal{Y}^*} \tag{16.5}$$

and so, putting everything together, we get the dual problem

$$\sup_{\eta \in \mathcal{Y}^*} -\langle \eta, f \rangle - \frac{1}{2}\|\eta\|_{\mathcal{Y}^*}^2 - \chi_{\partial J(0)}\left(\frac{-A^*\eta}{\alpha} \right). \tag{16.6}$$

For convenience, we write $\mu = -\frac{\eta}{\alpha} \in \mathcal{Y}^*$, and we get rid of the χ by just restricting the range of the sup:

$$\sup_{\mu \in \mathcal{Y}^*, A^*\mu \in \partial J(0)} \alpha\Big(\langle \mu, f \rangle - \frac{\alpha}{2}\|\mu\|_{\mathcal{Y}}^2 \Big) \tag{16.7}$$

From here, we can check all the assumptions we need to ensure the existence of a solution (the conditions on strong duality). We need to check coercivity, which is guaranteed by a finite-dimensional null space, coercivity of the regularising function (the $\mathcal{Y}$-norm-squared term), and a trivial intersection of the null spaces of $\langle \cdot, f \rangle$ and the regularising function. We also need continuity at a point, which is satisfied at $u_0 = 0$. We also have the existence of a primal solution u_δ.

This tells us that for any $\delta > 0$, there exists a solution μ_δ of the dual problem, and further that strong duality holds:

$$\frac{1}{2}\|A\mu_\delta - f_\delta\|_{\mathcal{Y}}^2 + \alpha J(u_\delta) = \alpha \langle \mu_\delta, f_\delta \rangle - \frac{\alpha^2}{2}\|\mu_\delta\|_{\mathcal{Y}}^2. \tag{16.8}$$

We also have the optimality conditions

- $A^*\mu_\delta \in \partial J(u_\delta)$
- $-\alpha\mu_\delta \in \partial\Big(\frac{1}{2}\|\cdot\|_{\mathcal{Y}}^2\Big)(Au_\delta - f_\delta)$.

We can conclude that $\|\alpha\mu_\delta\|_{\mathcal{Y}^*} = \|Au_\delta - f_\delta\|_{\mathcal{Y}}$.

Further, using the definition of a subgradient at $\frac{1}{2}\|\cdot\|_{\mathcal{Y}}^2$ and looking at the subgradient $-\alpha\mu_\delta$ at $(Au_\delta - f_\delta)$, we obtain

$$0 \geq \frac{1}{2}\|Au_\delta - f_\delta\|_{\mathcal{Y}}^2 + \langle -\alpha\mu_\delta, 0 - (Au_\delta - f_\delta) \rangle. \tag{16.9}$$

This givs us the estimate

$$\langle \alpha\mu_\delta, Au_\delta - f_\delta \rangle \leq -\frac{1}{2}\|Au_\delta - f_\delta\|_{\mathcal{Y}}^2. \tag{16.10}$$

So we've done all this work. What does it have to do with convergence rates?

16.2 Source condition and convergence rates

We consider the formal limits when $\delta = 0$ of the primal and dual problems: for the primal, we have

$$\inf_{u|Au=f} J(u) = \inf_{u\in\mathcal{X}} \chi_{\{f\}}(Au) + J(u) \tag{16.11}$$

and for the dual, we have

$$\sup_{\mu:A^*\mu\in\partial J(0)} \langle \mu, f\rangle = \sup_{\mu:A^*\mu\in\partial J(0)} \langle \mu, Au_J^\dagger\rangle \tag{16.12}$$

$$= \sup_{\mu:A^*\mu\in\partial J(0)} \langle A^*\mu, u_J^\dagger\rangle = \sup_{v\in\mathcal{R}(A^*)\cap\partial J(0)} \langle v, u_J^\dagger\rangle. \tag{16.13}$$

Since $\chi_{\{f\}}$ is continuous nowhere, we can't guarantee that a solution to the dual problem exists in this limit. So when is there a noiseless dual solution?

Theorem 16.3. *Let $\mathcal{X}, \mathcal{Y}$ be Banach spaces and let $\mathcal{Y}$ be separable. Let all the conditions exist that ensure the existence of a $J-$minimising solution $u_J^\dagger$ of the primal problem, and let $\alpha(\delta)$ be chosen appropriately. If the dual solution μ_δ is bounded uniformly in δ,then $u_J^\dagger$ satisfies the source condition, which means that there exists $\mu^\dagger \in \mathcal{Y}^*$ such that $A^*\mu^\dagger \in \partial J(u_J^\dagger)$ (belongs to the subdifferential of J at the $J-$minimising solution.)*

Another way of saying this is that the range of A^* has at least one point in common with the subdifferential of J at the $J-$minimising solution.

This is a necessary condition for the boundedness of dual solutions, but we can also show it is a sufficient condition, meaning we have an if-and-only-if.

Theorem 16.4. *Let $\mathcal{X}, \mathcal{Y}$ be Banach spaces and let $\mathcal{Y}$ be separable. Let all the conditions exist that ensure the existence of a $J-$minimising solution $u_J^\dagger$ of the primal problem, and let $\alpha(\delta)$ be chosen appropriately. If $u_J^\dagger$ satisfies the source condition, then μ_δ is bounded uniformly in δ. Moreover, $\mu_\delta \rightharpoonup^* \mu^\dagger$ in $\mathcal{Y}^*$ as $\delta \to 0$ (maybe up to a subsequence), where $\mu^\dagger$ is the solution of the dual limit problem with minimal norm.*

Finally, we can see the relationship to convergence rates!

Theorem 16.5. *Let the source condition be satisfied at a $J-$minimising solution $u_J^\dagger$ and let u_δ be a regularised solution of the primal problem. Then we have*

$$D_J^{p_\delta,p^\dagger}(u_\delta, u_J^\dagger) \le \frac{1}{4\alpha}\big(\delta + \alpha\|\mu^\dagger\|\big)^2 + \delta\|\mu^\dagger\|. \tag{16.14}$$

where $p_\delta = A^*\mu_\delta \in \partial J(u_\delta)$, $p^\dagger = A^*\mu^\dagger \in \partial J(u_J^\dagger)$. Further, for the optimal choice of α, given by $\alpha(\delta) = \frac{\delta}{\|\mu^\dagger\|}$, ths simplifies to

$$D_J^{p_\delta,p^\dagger}(u_\delta, u_J^\dagger) \le 3\delta\|\mu^\dagger\|. \tag{16.15}$$

We don't actually know $\|\mu^\dagger\|$ up front, but it's still a useful limit to know. Further, we have a linear bound on the Bregman distance with the noise level, and we get an upper limit on what the error should be.