Non-convex inverse problems

Irène Waldspurger

waldspurger@ceremade.dauphine.fr

Initial version: January to March 2023 This version: January to February 2025

Contents

1	Intr	Introduction		
	1.1	Inverse problems		6
		1.1.1	Definition	6
		1.1.2	Theoretical aspects	6
		1.1.3	Our focus: algorithms	9
	1.2	Conve	x vs non-convex	10
	1.3 Non-convex inverse problems: examples		onvex inverse problems: examples	11
		1.3.1	Sparse recovery - compressed sensing	11
		1.3.2	Low rank matrix recovery	14
		1.3.3	Other examples	19
2	Convexification			23
2.1 The basis: compressed sensing		asis: compressed sensing	25	
		2.1.1	Convexification: principle	25
		2.1.2	Intuition	25
		2.1.3	Tightness guarantees under restricted isometry	29
Α	Ren	ninders	s on symmetric and Hermitian matrices	31

Acknowledgements

Many thanks to Claire Boyer and Vincent Duval for sharing their lecture notes with me.

CONTENTS

4

Chapter 1

Introduction

What you should know / be able to do after this chapter

- Know the definition of "inverse problem", and a few examples.
- Understand what we call (in the context of this course) *theoretical aspects* and *algorithmic aspects* of an inverse problem. Know that the class will be about algorithmic aspects.
- Know the definition of "uniqueness" and "stability" in the context of inverse problems.
- For a linear problem, determine whether it is stable or not by looking at the singular values.
- With some guidance, be able to prove that a given inverse problem satisfies the uniqueness and stability properties (or not).
- Know our evaluation criteria for algorithms.
- Identify the main differences between convex and non-convex inverse problems.
- Be able to determine whether a given problem is convex or not.
- Identify the main common points and differences between sparse and low-rank recovery.
- Understand the change of variable which turns phase retrieval into a low-rank matrix recovery problem.

1.1 Inverse problems

1.1.1 Definition

An *inverse problem* consists in identifying a (possibly complicated) object from a set of observations¹. For instance, if we are given (two-dimensional) photographs of a building, viewed from different angles, reconstructing a three-dimensional model of the building is an inverse problem. Here, the "object" is the 3D shape of the building and the set of observations is the set of photographs.

Mathematically, these problems are formalized as follows. Let E be the set of possible *objects*, and F the set of possible *observations*. The *observation* procedure is described by a function $M : E \to F$. An inverse problem is, given some observation $y \in F$,

find
$$x \in E$$
 such that $M(x) = y$.

(Inverse)

Remark

The notion of *inverse problem* is often opposed to the notion of *direct problem*. A direct problem is the converse of an inverse problem: assuming the object and the observation procedure are known, compute the observations. For instance, if we are given a description of a fluid at some instant (viscosity, density, velocity at each point...), predicting how the fluid will be one minute later is a direct problem, which amounts to solving a specific partial differential equation. Here, the object is the fluid, and the observation procedure is "let it flow for one minute, then look at it".

1.1.2 Theoretical aspects

Problems of the form (Inverse) can be approached from two main angles.

• One can try to describe the properties of the solutions, without explicitly computing them. I will call this the *theoretical aspects*.

 $^{^1\}mathrm{Here,}$ we will call observation any procedure which, from the object, produces an outcome.

1.1. INVERSE PROBLEMS

• One can design algorithms to numerically solve the problem. I will call this the *algorithmic aspects*. 2

This class is about algorithmic aspects. However, it is difficult to design a sensible algorithm if one has no idea at all of the properties of the solution. Therefore, in this section, we give a very brief overview of the theoretical aspects.

When given a specific instance of Problem (Inverse), a first question that arises is the *existence* of solutions: for an arbitrary y, does there always exist a solution x to Problem (Inverse)? If we restrict ourselves to vectors y which are the outcome of a real measurement process (that is, of the form y = M(x)for some x), the answer is obviously yes. But if some errors have occured in the process, the answer may not be obvious anymore. For the problems we will consider in this class, existence will rarely be a problem, so we leave this question aside.

Assuming a solution exists, the other main two questions are *uniqueness* and *stability*.

• Uniqueness: Is the solution of Problem (Inverse) unique? This question is crucial, since, if the solution is not unique, it is impossible to recover the true object of interest with certainty.

More formally, we say that Problem (Inverse) satisfies the uniqueness property if and only if

 $\forall x_1, x_2 \in E \text{ such that } x_1 \neq x_2, M(x_1) \neq M(x_2).$

• Stability: If y is not exactly known, but only available up to some error, what will the solution(s) of Problem (Inverse) look like? Will it be close to the "true" solution, the one we would have obtained if there had been no error on y? This is also crucial: in real life, exact measurements are never available.

There are several sensible, but not equivalent, ways to translate this informal property to a formal one. A standard one is to say that Problem (Inverse) is stable if there exists a constant C > 0 "not too large" (say $C \leq 10$) such that

 $\forall x_1, x_2 \in E \text{ such that } x_1 \neq 0,$

 $^{^{2}}$ This choice of names does not mean that there is no "theory" behind algorithms. Actually, this class is about algorithmic aspects, but it will be mostly theoretical and rigorous.

CHAPTER 1. INTRODUCTION

$$\frac{||x_1 - x_2||_E}{||x_1||_E} \le C \frac{||M(x_1) - M(x_2)||_F}{||M(x_1)||_F}.$$
 (1.1)

Here, $||.||_E$ and $||.||_F$ are norms on E and $F.^3$

Example 1.1: finite-dimensional linear inverse problem

Let us assume that

- E, F are real finite-dimensional vector spaces: $E = \mathbb{R}^d$ and $F = \mathbb{R}^m$ for some $d, m \in \mathbb{N}^*$;
- $M: E \to F$ is linear, represented by some matrix $A \in \mathbb{R}^{m \times d}$.

Under these assumptions, Problem (Inverse) rewrites as

find $x \in \mathbb{R}^d$ such that Ax = y.

For a given y, assuming a solution x_* exists, it is *unique* if

$$\{x \in \mathbb{R}^d, Ax = y\} = \{x_*\},\$$

that is if and only if $Ker(A) = \{0\}$ (A is an injective matrix).

We now assume that the solution is unique. Is it *stable*? If the norms $||.||_E$ and $||.||_F$ in Equation (1.1) are the standard ℓ^2 -norms, then it is possible to show that the problem is *stable* if and only if the smallest and largest singular values of A satisfy

$$\frac{\lambda_{\max}(A)}{\lambda_{\min}(A)} \lesssim 10$$

The ratio $\frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}$ is called *condition number* of A. For more details, see the exercises.

As said before, these questions will not be the subject of the class. For each newly encountered problem, we will try to give conditions under which

³These norms must in principle be carefully chosen according to the physical structure of the concrete underlying problem. Some choices may reflect better than others the desired properties of the solutions.

the solution is unique and stable but we will not spend much time on it. When these questions are not mentionned, the reader can simply assume that the considered problem satisfies uniqueness and stability properties. However, in principle, when facing a new problem, these questions must be the starting point, otherwise we are at risk of working towards the conception of algorithms for solving problems which can actually not be solved.

1.1.3 Our focus: algorithms

In this class, we will be interested in algorithms which allow to solve inverse problems. Cambridge dictionary defines the word algorithm as

"a set of mathematical instructions or rules that, especially if given to a computer, will help to calculate an answer to a problem."

Following this definition, an *algorithm* can take many forms. In particular, although the class of *iterative algorithms* (that is, those that repeat a set of instructions until some stopping criterion is met) will be of particular importance to us, one must not imagine that all algorithms are iterative.

In applications, a "good" algorithm is an algorithm which

- works: given a problem, it must output a correct solution; we can tolerate the algorithm failing once in a while, but the failure rate must be as small as possible;
- uses as few computational resources as possible: it must be fast (not too many operations) and have a moderate memory footprint.

Here, we will be interested in algorithms for which, moreover,

• these good properties (especially the first one) can be rigorously proved.

This additional requirement tends to be in contradiction with the computational efficiency, in the sense that, oftentimes, the algorithms which work best in practice are difficult to study rigorously. As a consequence, the algorithms we will present in this class will in most cases not be the best ones for real applications. They must be considered as toy models for "really usable" algorithms, should ideally retain as many specificities of their "really usable" counterparts as possible, but will inevitably miss some.

Similarly, the hypotheses under which we will establish correctness guarantees for the algorithms will often be much stronger than what holds in real applications. It is an important but difficult research direction to weaken these hypotheses.

1.2 Convex vs non-convex

All inverse problems can be reformulated as *optimization problems*, that is problems of the following form:

minimize
$$f(x)$$

over all $x \in H$
such that $x \in C_1$, (Opt)
 \dots
 $x \in C_S$.

Here, $f : H \to \mathbb{R} \cup \{+\infty\}$ can be any *objective* function, over a real or complex vector space H, and C_1, \ldots, C_S are subsets of H which model the constraints imposed on the unknown x.

An optimization problem is called *convex* if f is a convex function and C_1, \ldots, C_S are convex sets. By extension, we say that an inverse problem is convex if it can be reformulated as a convex optimization problem.

Definition 1.2: convexity

A function $f : H \to \mathbb{R} \cup \{+\infty\}$ is *convex* if, for any $x_1, x_2 \in H$ and any $s \in [0, 1]$,

$$f((1-s)x_1 + sx_2) \le (1-s)f(x_1) + sf(x_2).$$
(1.2)

A set $C \subset H$ is *convex* if, for any $x_1, x_2 \in C$ and any $s \in [0, 1]$, the vector

 $(1-s)x_1 + sx_2$

is also an element of C.

In first approximation, we can say that convex problems admit efficient algorithms. This is not an absolute rule, since some convex sets or functions are quite difficult to manipulate. However, it is true that many algorithms exist for convex problems, with a behavior which is quite well understood. The situation is very different for the problems we will consider in this class, which are non-convex. For non-convex problems, the existence of algorithms both guaranteed to succeed and running in an reasonable amount of time is an exception.

Intuitively, convexity allows to deduce global information from local one. For instance, if one knows the values at a few points of a convex function f and its gradient, Inequality (1.2) makes it possible to compute upper and lower bounds on f, and hence obtain an approximation of its minimum. One can then query the values at other points to refine the approximation. This is illustrated on Figures 1.1a and 1.1b. But if the function is not convex, the knowledge of its values at a few points provides no information about the values at other points and, in particular, provides no information on its minimum. This is illustrated on Figures 1.1c and 1.1d. This is what makes non-convex optimization much more difficult than convex optimization.

This difficulty is a fundamental property of non-convex problems: if we do not have good algorithms able to solve any non-convex problem, it is not because we have not discovered these good algorithms yet. It is because good algorithms do not exist.⁴ As a consequence, in this class, we will not try to propose algorithms able to solve all problems of a given non-convex family: this is hopeless. At best, our algorithms will be able to solve "a large part" of problems of the family.

1.3 Non-convex inverse problems: examples

Let us now present a few examples of non-convex inverse problems.

1.3.1 Sparse recovery - compressed sensing

Our first example is called *sparse recovery* or *compressed sensing*. It consists in recovering a vector $x \in \mathbb{R}^d$ from linear measurements

$$y \stackrel{def}{=} Ax \in \mathbb{R}^m$$

where $A \in \mathbb{R}^{m \times d}$ is a known matrix, under the assumption that x is *sparse*. The word *sparse* means that x has a small number of non-zero coordinates:

⁴In particular, many families of non-convex problems have been proved to be NPdifficult. This means that, unless P=NP, there exists no algorithm able to solve all problems in the family with a time complexity at most polynomial in their dimension.



Figure 1.1: (a) Representation of the values and derivatives of a function $f : \mathbb{R} \to \mathbb{R}$ at a few points. (b) Upper and lower bounds on f (respectively orange and red lines) one can deduce from the knowledge of these values and derivatives if f is convex. Observe that it gives a reasonably tight approximation of f, its minimum and minimizer. (c) A non-convex function compatible with these values and derivatives. (d) Another non-convex function the compatible with these values and derivatives. Observe that the minimum and minimizer are significantly different from 1.1c.

for some $k \in \mathbb{N}^*$ much smaller than d,

$$||x||_0 \le k,$$

where $||x||_0 = \text{Card}\{i \leq d, x_i \neq 0\}$. (This quantity is often called the ℓ^0 -norm, although it is not a norm, since it is not homogeneous.)

Note that, if $m \ge d$ and A is injective, then this problem can be solved by inverting A; it is not necessary to use the sparsity assumption. This problem is only interesting when m is much smaller than d, in which case A is not injective and, if we were to ignore the sparsity assumption, y would not uniquely determine x.

Assuming that k is known, the problem can be written as

recover
$$x \in \mathbb{R}^d$$

such that $Ax = y$,
and $||x||_0 \le k$. (CS)

It is non-convex because the set $\{x, ||x||_0 \le k\}$ is non-convex.

Sometimes, the unknown x is not directly sparse, but only sparse when represented in some adequate basis, or after some adequate linear transformation. In this case, the condition " $||x||_0 \leq k$ " must be replaced with " $||\Phi x||_0 \leq k$ ", where Φ encodes the basis or linear transformation.

This problem is notably natural in image processing, since many natural images enjoy a sparsity structure. Photos, for instance, are well-known to be approximately sparse when represented in a *wavelet basis*.

For compressed sensing, uniqueness of the reconstruction can be guaranteed through a condition on the kernel of A.

Proposition 1.3: unique recovery for compressed sensing

We assume that $\operatorname{Ker}(A)$ does not contain a vector X such that $||X||_0 \leq 2k$. Then, if Problem (CS) has a solution, this solution is unique.

Proof. Let us assume, by contradiction, that Problem (CS) has two distinct solutions $X_1, X_2 \in \mathbb{R}^d$. Then

$$A(X_1 - X_2) = AX_1 - AX_2 = y - y = 0,$$

so $X_1 - X_2$ belongs to Ker(A). And

$$||X_1 - X_2||_0 \le ||X_1||_0 + ||X_2||_0 \le 2k,$$

which contradicts the assumption.

From this proposition, one can show that, if $m \ge 2k$, then almost all matrices A guarantee unique recovery of the underlying sparse vector. Under a stronger condition on A, one can also establish stability recovery guarantees (see for instance the introductory article [Candès and Wakin, 2008]).

1.3.2 Low rank matrix recovery

In low-rank matrix recovery, the goal is also to recover an object from linear measurements. This time, the "object" is a matrix $X \in \mathbb{R}^{d_1 \times d_2}$ (or $X \in \mathbb{C}^{d_1 \times d_2}$). As in the case of compressed sensing, there are not enough linear measurements to uniquely determine X without additinal information, but we do have some additional information on X: it is low-rank. This yields the problem

recover $X \in \mathbb{R}^{d_1 \times d_2}$ such that $\mathcal{L}(X) = y$, and $\operatorname{rank}(X) \leq r$.

(Low rank)

Here, $\mathcal{L} : \mathbb{R}^{d_1 \times d_2} \to \mathbb{R}^m$ is the linear measurement operator and r is a given upper bound on the rank of the matrix. Given that any $d_1 \times d_2$ matrix has rank at most $\min(d_1, d_2)$, the rank constraint is only useful if $r < \min(d_1, d_2)$. In some applications, it is relevant to assume that $d_1 = d_2$ and X is semidefinite positive: $X \succeq 0$.

This problem is sometimes called *matrix sensing*, especially when \mathcal{L} is a random operator. A uniqueness result similar to Proposition (1.3) holds.

Proposition 1.4: uniqueness for low-rank matrix recovery We assume that $\text{Ker}(\mathcal{L})$ does not contain a matrix X such that

 $\operatorname{rank}(X) \le 2r.$

Then, if Problem (Low rank) has a solution, this solution is unique.

The proof of the proposition is identical to Proposition 1.3. From this proposition, one can show (but it is not easy) that the solution of Problem (Low rank), when it exists, is unique, for almost all operators \mathcal{L} , provided that

$$m \ge 2r(d_1 + d_2 - 2r) \quad \text{if } 2r \le \min(d_1, d_2), \\ \ge d_1 d_2 \quad \text{if } \min(d_1, d_2) < 2r < 2\min(d_1, d_2).$$

When r is small (of order 1, for instance), this shows that we can hope to recover the "true" matrix X with a number of linear measurements much smaller than what we would need if we did not know X to be low-rank (in this case, we would need $m \ge \dim(\mathbb{R}^{d_1 \times d_2}) = d_1 d_2$, which is much larger than $2r(d_1 + d_2 - 2r)$ if $r \ll \min(d_1, d_2)$).

Matrix completion Several special cases of Problem (Low rank) are of particular interest, and form subfamilies of inverse problems with their own applications and theoretical characteristics. The first one is *matrix completion*. In this case, the linear measurements available on X are a subset of coefficients:

recover $X \in \mathbb{R}^{d_1 \times d_2}$ such that $X_{ij} = y_{ij}, \forall (i, j) \in \Omega$ and rank $(X) \leq r$.

(Matrix completion)

Here, $\Omega \subset \{1, \ldots, d_1\} \times \{1, \ldots, d_2\}$ contains the indices of available coefficients.

The most popular application is the so-called "Netflix problem".⁵ In this application, X represents the opinion of users on films: the coefficient X_{ij} is an "affinity score" between User *i* and Film *j* (it represents how much User *i* would like Film *j*). It is reasonable to assume that X is low-rank:⁶ this models the similarities between the users, and between the films (e.g. if User 1 and 2 have the same opinion on Films 1, 2, 3, 4, it is plausible that they also have essentially the same opinion on Film 5). The available coefficients X_{ij}

⁵asked by Netflix in 2006, with a 1,000,000\$ prize, and declared solved in 2009

⁶Keep however in mind that this assumption is only approximately satisfied by the "true" Netflix affinity scores matrix. On the other hand, the true matrix has additional structure that can be exploited to solve the problem.

correspond to pairs (i, j) for which User *i* has watched Film *j* and sent the corresponding score to the film distribution platform. The other coefficients are not available, but the platform would like to guess them, so as to be able to propose relevant film suggestions to their users. Guessing the non-available coefficients exactly amounts to solving Problem (Matrix completion).

Phase retrieval Another special case of Problem (Low rank) which we will discuss in length in this course is *phase retrieval*.

At first sight, phase retrieval problems have nothing to do with matrices and low-rankness. They are problems of the following general form

> recover $x \in \mathbb{C}^d$ such that $|L_j(x)| = y_j, \forall j \leq m$.

(Phase retrieval)

Here, $L_1, \ldots, L_m : \mathbb{C}^d \to \mathbb{C}$ are known linear operators, the notation "|.|" stands for the usual complex modulus, and y_1, \ldots, y_m are given.

The main motivations for studying phase retrieval come from the field of imaging. Indeed, it is much easier to record the intensity (that is, the modulus, in an adequate mathematical model) of an electromagnetic wave than its phase. It is therefore frequent to have to recover an object from modulus-only measurements. Oftentimes, these measurements can specifically be described by a Fourier transform (because, under some assumptions, the diffraction pattern of an object is the Fourier transform of its characteristic function), but not always. Phase retrieval is also of interest for audio processing.

Remark

For any $x \in \mathbb{C}^d$ and $u \in \mathbb{C}$ such that |u| = 1, it holds

$$|L_j(ux)| = |uL_j(x)| = |u| |L_j(x)| = |L_j(x)|, \quad \forall j \le m.$$

Therefore, the sole knowledge of $(y_j = |L_j(x)|)_{j \le m}$ can never allow to exactly recover x. There is always a global phase ambiguity: x cannot be distinguished from ux.

This is in general not harmful in applications, and we will be satisfied if we can recover x up to a global phase.

Given specific linear forms L_j , it is in general difficult to determine if the (Phase retrieval) problem satisfies the uniqueness and stability properties. However, it is known that uniqueness holds "in principle" as soon as m is larger than (roughly) 4d.

Proposition 1.5: [Conca, Edidin, Hering, and Vinzant, 2015]

Let us assume that $m \geq 4d - 4$. Then, for almost all linear maps $L_1, \ldots, L_m : \mathbb{C}^d \to \mathbb{C}$, it holds that, for all $x, x' \in \mathbb{C}^d$,

$$(|L_j(x)| = |L_j(x')|, \forall j \le m) \Rightarrow (\exists u \in \mathbb{C}, |u| = 1, x = ux').$$

With a slightly larger m, stability also "generically" holds.

Let us now explain why phase retrieval is a special case of low-rank matrix recovery. Readers which are not perfectly comfortable with the notions of Hermitian matrices and of semidefinite positive matrices should first read Appendix A.

The crucial ingredient is an adequate change of variable: instead of recovering $x \in \mathbb{C}^d$ up to a global phase, let us try to recover

$$X \stackrel{def}{=} xx^* = \begin{pmatrix} |x_1|^2 & x_1\overline{x}_2 & \dots & x_1\overline{x}_d \\ x_2\overline{x}_1 & |x_2|^2 & \dots & x_2\overline{x}_d \\ \vdots & \ddots & \vdots \\ x_d\overline{x}_1 & \dots & |x_d|^2 \end{pmatrix}.$$

Remark

A matrix $X\in \mathbb{C}^{d\times d}$ can be written as $X=xx^*$ for some $x\in \mathbb{C}^d$ if and only if

 $X \succeq 0$ and $\operatorname{rank}(X) \leq 1$.

When these conditions hold, x is equal, up to a global phase, to

$$\sqrt{\lambda_1}z_1,$$

where λ_1 is the largest eigenvalue of X, and z_1 any unit-normed eigenvector for this eigenvalue.

Proof. For any $x \in \mathbb{C}^d$, the matrix xx^* is Hermitian, and semidefinite positive:

$$\forall z \in \mathbb{C}^d, \quad \langle z, xx^*z \rangle = z^*(xx^*)z = |z^*x|^2 \ge 0.$$

It has rank at most 1 because $\operatorname{Range}(xx^*) = \operatorname{Vect}\{x\}$.

Conversely, if $X \succeq 0$ and rank $(X) \le 1$, then X can be diagonalized in an orthogonal basis (z_1, \ldots, z_d) :

$$X = \sum_{k=1}^{d} \lambda_k z_k z_k^* \quad \text{with } \lambda_1 \ge \dots \ge \lambda_d \text{ the eigenvalues}.$$

All the eigenvalues are nonnegative, since $X \succeq 0$. Since rank $(X) \le 1$, they are all 0, except possibly the first one, so

$$X = \lambda_1 z_1 z_1^* = (\sqrt{\lambda_1} z_1) (\sqrt{\lambda_1} z_1)^*,$$

so it can be written as $X = xx^*$ with $x = \sqrt{\lambda_1}z_1$. This proves the first part of the remark.

For the second part, let us assume that $X = xx^*$ for some $x \in \mathbb{C}^d$. We have just seen that X is also equal to $\tilde{x}\tilde{x}^*$ for $\tilde{x} = \sqrt{\lambda_1}z_1$. We must simply show that x and \tilde{x} are equal up to a global phase. As

$$\operatorname{Vect}\{x\} = \operatorname{Range}(X) = \operatorname{Vect}\{\tilde{x}\},\$$

it holds that x and \tilde{x} are collinear: there exists $u \in \mathbb{C}$ such that $x = u\tilde{x}$. In addition,

$$||x||^2 = \operatorname{Tr}(X) = ||\tilde{x}||^2$$

hence x and \tilde{x} have the same norm. As $||x|| = |u| ||\tilde{x}||$, this implies that |u| = 1: x and \tilde{x} are equal up to a global phase.

From the previous remark, it is equivalent to recover x up to a global phase or X. Indeed, X can be computed from x (even up to a global phase: $(ux)(ux)^* = u\overline{u}xx^* = xx^*$ if |u| = 1) and x can be computed up to a global phase from X by extracting the only eigenvector of X with non-zero eigenvalue.

In addition, for any j, knowing $|L_j(x)|$ is equivalent to knowing $|L_j(x)|^2$. Denoting v_j the vector such that $L_j = \langle v_j, . \rangle$, we have

$$|L_j(x)|^2 = \langle v_j, x \rangle \overline{\langle v_j, x \rangle} = (v_j^* x) (x^* v_j) = v_j^* X v_j.$$

18

Consequently, Problem (Phase retrieval) is equivalent to

recover $X \in \mathbb{C}^{d \times d}$ such that $v_j^* X v_j = y_j^2, \forall j \le m,$ $X \succeq 0,$ rank $(X) \le 1.$ (Matrix PR)

This is, as announced, a low rank matrix recovery problem.

1.3.3 Other examples

These examples will not be covered in class (except for super-resolution, but later); they are provided for curious readers only.

Machine learning In a machine learning task, the goal is to predict some output y given some input x. For instance, the input can be a photograph, and the output the name of the objects represented on the photograph, or the input can be a low-quality audio signal and the output the corresponding high-quality signal. We denote P the "perfect" prediction function, which to an input x maps the correct

$$y = P(x)$$

The predictor P is unknown and must be learned from the available inputoutput examples $(x_1, y_1), \ldots, (x_n, y_n)$. This leads to the problem

find
$$P \in \mathcal{H}$$

such that $P(x_k) = y_k, \forall k \le n$, (ML)

where \mathcal{H} is a well-chosen class of functions (\mathcal{H} can for instance be the set of linear maps, or the set of neural networks with a given architecture).

The questions raised by Problem (ML) are quite different from the ones raised by the other inverse problems we have seen. Indeed, it often happens that the perfect predictor P is not in the chosen set \mathcal{H} , in which case the problem may not have an exact solution, only an approximate one. In addition, if \mathcal{H} is a bit sophisticated, there are typically several (and even many) elements $P \in \mathcal{H}$ such that $P(x_k) = y_k$ for all k (in other words, the uniqueness property does not hold). All these elements P yield the same predictions for the available inputs x_1, \ldots, x_n , but may differ significantly on unseen examples. It is therefore important to choose, among these P, the one which has the best chances to perform well on unseen examples (i.e. which generalizes best)

Dictionary learning In this problem, one is given a set of "interesting" signals $y_1, \ldots, y_m \in \mathbb{R}^d$ (e.g. patches of natural photographs or of medical images), and the goal is to learn a good "representation" for them, under the form of a *dictionary*. A dictionary is a set of elements $a_1, \ldots, a_M \in \mathbb{R}^d$, usually called *atoms*, such that any signal y_k can be written as a linear combination of a small number of atoms:

$$y_k = \sum_{l=1}^M \lambda_l^{(k)} a_l$$
 such that $||\lambda^{(k)}||_0$ is small.

We write the dictionary in matricial form by concatenating the atoms into a single matrix:

$$A = \begin{pmatrix} a_1 & a_2 & \dots & a_M \end{pmatrix}$$

Finding the dictionary A consists in solving the following problem

find $A \in \mathbb{R}^{d \times M}, \lambda^{(1)}, \dots, \lambda^{(m)} \in \mathbb{R}^{M}$ such that $A\lambda^{(k)} = y_k, \forall k \leq m,$ $||\lambda^{(k)}||_0 \leq S,$

(Dictionary learning)

where S is an a priori bound on the number of atoms involved in the decomposition of each signal y_k .

Super-resolution Super-resolution is a general term which covers all problems where one tries to recover a "sharp" signal from a "blurred" version. In this paragraph, we present the simplest possible model for such a problem.

The signal we aim at identifying is a collection of point masses in [0; 1[. The positions of the masses are τ_1, \ldots, τ_S and their weights are a_1, \ldots, a_S . This signal can be represented by a measure

$$\mu = \sum_{s=1}^{S} a_s \delta_{\tau_s} \in \mathcal{M}([0;1[),$$

where $\mathcal{M}([0; 1[) \text{ is the set of signed (or even complex-valued, if <math>a_1, \ldots, a_s$ are complex) finite Borel measures on [0; 1[and, for any s, δ_{τ_s} is the dirac at position τ_s .⁷

The "blurred" version of the signal is modelled as the set of low-frequency coefficients of the Fourier transform of μ : for all $k = -N, \ldots, N$, we have access to

$$\hat{\mu}[k] = \int_0^1 e^{-2\pi i k t} d\mu(t) \left(= \sum_{s=1}^S a_s e^{-2\pi i k \tau_s} \right).$$

If we call y_{-N}, \ldots, y_N the known Fourier coefficients, the problem can be written as

find $\mu \in \mathcal{M}([0;1[))$ such that $\hat{\mu}[k] = y_k, \forall k = -N, \dots, N$, and μ is a sum of S diracs.

(Super-resolution)

This problem can be seen as a continuous version of compressed sensing (Problem (CS)). The unknown, instead of a finite-dimensional vector, is a measure on [0; 1[, but it must still be recovered from linear measurements, and satisfies a sparsity constaint (it is the sum of at most S diracs).

⁷that is to say, δ_{τ_s} is the measure such that, for any measurable $E \subset [0; 1[, \mu(E) = 1$ if $\tau_s \in E$ and $\mu(E) = 0$ otherwise.

Chapter 2

Convexification

What you should know / be able to do after this chapter

- Understand the general principle of convexification, and what "tightness" means.
- Be able to suggest convex relaxations of non-convex problems, based notably on the « convex hull » reasoning which provides intuition in the cases of compressed sensing and low-rank recovery.
- Using a 2-dimensional picture, explain why (Basis Pursuit) can be expected to be a tight relaxation of compressed sensing.
- Know the definition of « restricted isometry ».
- Know the proof technique for establishing tightness guarantees which relies on restricted isometry (in particular, know the statements of Theorems 2.4 and ??).
- Know that restricted isometry holds true for the simplest cases of random linear operators.
- Explain the limitations of this technique: restricted isometry does *not* hold for some more "structured" operators.
- Prove that it does not hold in the example of super-resolution for Fourier measurements.



Figure 2.1: Principle of convexified algorithms, when relaxation is tight.

- Sketch the proof technique for establishing tightness guarantees using dual certificates.
- Understand (i.e. be able to do it again alone, with minimal help) the derivation of the dual problem of TV minimization.

As discussed in the introduction, non-convexity is a major hurdle for numerically solving inverse problems. Simple local search algorithms are at risk of getting stuck in poor local optima. A possible strategy to overcome this difficulty is to approximate the non-convex problem with a convex one. This convex approximation is called a *convex relaxation*. Since numerically solving a convex problem is in general doable, we can in general solve the approximation. At first sight, there is no reason why solving this approximation would provide useful information towards solving the non-convex problem. But surprisingly, it turns out that, in many situations, the convex approximation has the same solution as the original non-convex problem! One then says that relaxation is *tight*. When this happens, it yields a convenient method for solving the non-convex problem. This general scheme is depicted on Figure 2.1.

2.1 The basis: compressed sensing

2.1.1 Convexification: principle

The model example for this chapter, which serves as a basis for other problems, is compressed sensing.

recover
$$x \in \mathbb{R}^d$$

such that $Ax = y$,
and $||x||_0 \le k$. (CS)

When the problem has a unique solution, it is the vector with minimal ℓ^0 norm among all vectors x such that Ax = y. This allows to reformulate the
problem as

minimize
$$||x||_0$$

for $x \in \mathbb{R}^d$ $(\ell^0\text{-min})$
such that $Ax = y$.

The set $\{x \in \mathbb{R}^d, Ax = y\}$ is convex. The non-convex part of the problem is the objective function $||.||_0$. To make the problem convex, we replace the ℓ^0 -norm with the ℓ^1 -norm:

$$||x||_1 = \sum_{i=1}^d |x_i|,$$

which leads to the following *convex* problem:

minimize $||x||_1$ for $x \in \mathbb{R}^d$ such that Ax = y.

(Basis Pursuit)

2.1.2 Intuition

An intuitive reason for using the ℓ^1 -norm as a convex approximation of the ℓ^0 -norm is that the unit ℓ^1 -ball is the smallest convex set which contains the "maximally sparse" vectors of norm 1.

Proposition 2.1: ℓ^1 -ball as a convex hull

Let S be the set of vectors with exactly one non-zero coordinate, equal to -1 or 1.

The unit ℓ^1 -ball $\{x \in \mathbb{R}^d, ||x||_1 \le 1\}$ is the convex hull of \mathcal{S} .

Proof. This proposition is a consequence of Proposition 2.2. Indeed, the unit ℓ^1 -ball is a closed compact subset of \mathbb{R}^d . It is therefore the convex hull of its extremal points (from the Krein-Milman theorem, see for instance [Barvinok, 2002, Chapter II, thm 3.3]), that it is the convex hull of \mathcal{S} .

The next proposition states a stronger, but similar, result, which is crucial in explaining the success of (Basis Pursuit) (i.e. why it is oftentimes a *tight* convex relaxation).

Proposition 2.2: extremal points of the l^1 -ball

The extremal points^{*a*} of the unit ℓ^1 -ball $\{x \in \mathbb{R}^d, ||x||_1 \leq 1\}$ are the vectors with exactly one non-zero coordinate, equal to -1 or 1.

^aAn *extremal point* of a convex set C is a point y which cannot be written as

$$y = (1 - \theta)z_1 + \theta z_2$$

for $z_1, z_2 \in C$ different from y and $\theta \in [0; 1]$.

Proof. Let S be the set of vectors with exactly one non-zero coordinate, equal to -1 or 1. Let B_{ℓ^1} be the unit ℓ^1 -ball.

First, we show that the elements of S are extremal points of B_{ℓ^1} . Let $y \in S$ be fixed. Let *i* be its unique non-zero coordinate. Let us assume for simplicity that $y_i = 1$ (the reasoning can be adapted for $y_i = -1$). Let $z_1, z_2 \in B_{\ell^1}, \theta \in [0; 1]$ be such that

$$y = (1 - \theta)z_1 + \theta z_2.$$

We must show that $z_1 = y$ or $z_2 = y$. If $\theta = 0$, then $z_1 = y$, and if $\theta = 1$, then $z_2 = y$, so we can assume $\theta \neq 0, 1$.

We have

$$1 = y_i = (1 - \theta)(z_1)_i + \theta(z_2)_i.$$

2.1. THE BASIS: COMPRESSED SENSING

Observe that $(z_1)_i \leq |(z_1)_i| \leq ||z_1||_1 \leq 1$ and, similarly, $(z_2)_i \leq 1$. These two inequalities must be equalities, otherwise $1 = (1 - \theta)(z_1)_i + \theta(z_2)_i < (1 - \theta) + \theta = 1$.

Now that we know that $(z_1)_i = 1$, we can say that

$$\sum_{j \neq i} |(z_1)_j| = ||z_1||_1 - |(z_1)_i| = ||z_1||_1 - 1 \le 0,$$

hence $(z_1)_j = 0$ for all $j \neq i$. This shows $z_1 = y$, and concludes the proof that y is an extremal point of B_{ℓ^1} .

Conversely, we show that every extremal point of B_{ℓ^1} is in \mathcal{S} . Let $y \in B_{\ell^1}$ be extremal.

First, we note that $||y||_1 = 1$. Indeed, if $||y||_1 < 1$, we can write, for any vector $\epsilon \in \mathbb{R}^d \setminus \{0\}$,

$$y = \frac{1}{2}(y+\epsilon) + \frac{1}{2}(y-\epsilon).$$

When ϵ is close enough to zero, it holds $||y + \epsilon||_1, ||y - \epsilon||_1 \le ||y||_1 + ||\epsilon||_1 \le 1$, so $y + \epsilon, y - \epsilon$ belong to B_{ℓ^1} and are different from y, which contradicts the extremality of y.

Now, we show that y has only one non-zero coordinate. Let i be such that $y_i \neq 0$. By contradiction, we assume that not all other coordinates are zero. Let \tilde{y} be the vector which is equal to y, except that the *i*-th coordinate y_i has been replaced with 0; it is not the null vector. Let $e \in \mathbb{R}^d$ be the vector such that

$$e_i = \operatorname{sign}(y_i),$$

$$e_j = 0, \quad \forall j \neq i$$

Then

$$y = |y_i|e + ||\tilde{y}||_1 \frac{\tilde{y}}{||\tilde{y}||_1} = |y_i|e + (1 - |y_i|) \frac{\tilde{y}}{||\tilde{y}||_1},$$

which contradicts the extremality. (The last equality is true because $||\tilde{y}||_1 = \sum_{j \neq i} |y_j| = ||y||_1 - |y_i| = 1 - |y_i|$.)

Let us give an intuitive explanation, based on the previous proposition, of why we can expect (Basis Pursuit) to be a *tight* relaxation of Problem (CS), at least in some situations.

If the vector x_* we are trying to recover through Problem (CS) is sparse, then it is a linear combination of a small number of "maximally sparse" nonzero vectors. From Proposition 2.2, it is therefore a linear combination of a



Figure 2.2: Representation of B_{ℓ^1,x_*} and $\{x \in \mathbb{R}^2, Ax = Ax_*\}$ for A = (1 - 3) in two situations: (a) when $x_* = (0, 1)$ is sparse ; (b) when $x_* = (\frac{1}{4}, \frac{3}{4})$ is not sparse. Observe that $B_{\ell^1,x_*} \cap \{x, Ax = Ax_*\}$ is a singleton in the first case, but not in the second one.

small number of extremal points of the ℓ^1 -ball. This can be geometrically interpreted as the fact that x_* belongs to a "corner" of the ℓ^1 -ball

$$B_{\ell^1, x_*} \stackrel{def}{=} \{ x \in \mathbb{R}^d, ||x||_1 \le ||x_*||_1 \}.$$

The convex approximation (Basis Pursuit) has a unique minimizer equal to x_* if and only if

$$\nexists x \in \mathbb{R}^d \text{ such that } Ax = y = Ax_* \text{ and } ||x||_1 \le ||x_*||_1$$

$$\iff B_{\ell^1, x_*} \cap \{x \in \mathbb{R}^d, Ax = Ax_*\} = \{x_*\}$$

$$\iff B_{\ell^1, x_*} \cap (\{x_*\} + \operatorname{Ker}(A)) = \{x_*\}.$$

And the intersection of B_{ℓ^1,x_*} and an affine space containing x_* has much more chances to be the singleton $\{x_*\}$ if x_* is in a "corner" of B_{ℓ^1,x_*} (very crudely, if x_* is in a "corner", then, in the neighborhood of x_* , B_{ℓ^1,x_*} occupies only a small fraction of the space; it is therefore easier not to intersect it when considering an affine space going through x_*). This is depicted on Figure 2.2.

2.1.3 Tightness guarantees under restricted isometry

The convex problem (Basis Pursuit) can be traced back to at least the 70's. Since then, many researchers have proposed conditions on x_* and A under which the relaxation is tight (that is, the solutions of (Basis Pursuit) and (CS) are the same). A major progress (due notably to Candès, Donoho, Romberg and Tao) on this subject was, around twenty years ago, the introduction of the so-called *Restricted Isometry Property*, which is a simple assumption on A under which it is possible to guarantee tightness without imposing stringent conditions on x_* .

```
Definition 2.3: restricted isometry
```

Let $A \in \mathbb{R}^{m \times d}$ be a matrix. For any $k \in \{1, \ldots, d\}$, we define the k-restricted isometry constant δ_k of A as the smallest real number such that

 $(1 - \delta_k)||z||_2 \le ||Az||_2 \le (1 + \delta_k)||z||_2$

for all vectors $z \in \mathbb{R}^d$ with at most k non-zero coordinates.

Tightness of the convex relaxation (Basis Pursuit) under a restricted isometry condition is guaranteed by the following theorem.

Theorem 2.4

Let $A \in \mathbb{R}^{m \times d}$ be a matrix. For some $k \in \{1, \ldots, d\}$, we assume that its 4k-restricted isometry constant satisfies

$$\delta_{4k} < \frac{1}{4}.\tag{2.4}$$

For any $x_* \in \mathbb{R}^d$ with at most k non-zero coordinates, Problem (Basis Pursuit) with $y = Ax_*$ has a unique solution, which is x_* .

Under the same condition, it is moreover possible to prove a stability result for the convex relaxation: if y is "close" to Ax_* , then the solution of a slight modification of (Basis Pursuit) is "close" to x_* . The proof of Theorem 2.4 is the subject of an exercise, which follows [Candès, Romberg, and Tao, 2006].

Let us keep in mind that the restricted isometry property is a *sufficient* but not *necessary* condition for the correctness of the basis pursuit approach:

there are matrices A for which condition (2.4) does *not* hold and, nevertheless, Problems (CS) and (Basis Pursuit) have the same solution. However, it turns out that many natural matrices A satisfy the condition, hence Theorem 2.4 explains the success of the basis pursuit approximation in several interesting situations. The following theorem provides the simplest example of matrices with the restricted isometry property: matrices chosen at random according to a normal distribution (with high probability).

Theorem 2.5: [Candès and Tao, 2005]

Let c > 0 be some explicit constant, whose value we will not give here. We assume that $A \in \mathbb{R}^{m \times d}$ is generated at random according to a normal distribution^{*a*}. If

$$ck \log(d/k) \le m$$

Condition (2.4) holds with high probability.^b

^{*a*}that is, each coefficient of A is chosen independently at random according to a normal law $\mathcal{N}(0, 1/m)$.

^b With high probability means that it holds with probability at least $1 - e^{-\alpha m}$ for some constant $\alpha > 0$.

This theorem, combined with Theorem 2.4, shows that convexification allows to recovery k-sparse vectors from $O(k \log(d/k))$ linear measurements. This is surprisingly few. Indeed, Problem (CS) is only interesting when the number of measurements is at least O(k) (otherwise, the solution is not unique). At this threshold, solving this problem is a priori impossible with a polynomial time algorithm, but we see that it suffices to increase the number of measurements by a logarithmic factor so that polynomial time recovery becomes possible, through convexification.

Appendix A

Reminders on symmetric and Hermitian matrices

Let $d \in \mathbb{N}^*$ be fixed.

For any matrix $M \in \mathbb{C}^{d \times d}$, we denote M^* the transpose conjugate of M, that is the $d \times d$ matrix such that, for all $i, j \leq d$,

$$M_{ij}^* = \overline{M_{ji}}.$$

The notation " $\langle ., . \rangle$ " stands for the standard dot product when applied to real vectors: for all $a, b \in \mathbb{R}^d$,

$$\langle a,b\rangle = \sum_{i=1}^d a_i b_i$$

When applied to complex vectors, it denotes the standard Hermitian product: for all $a, b \in \mathbb{C}^d$,

$$\langle a,b\rangle = \sum_{i=1}^{d} \overline{a_i} b_i = a^* b.$$

Definition A.1: Hermitian matrices

A matrix $M \in \mathbb{C}^{d \times d}$ is Hermitian if $M = M^*$, that is if, for all $i, j \leq d$, $M_{ij} = \overline{M_{ji}}.$ Equivalently, M is Hermitian if and only if, for all $x, y \in \mathbb{C}^d$,

 $\langle Mx, y \rangle = \langle x, My \rangle \,.$

Definition A.2: semidefinite positive matrices

Let \mathbb{K} be \mathbb{R} or \mathbb{C} .

A matrix $M \in \mathbb{K}^{d \times d}$ is *semidefinite positive* if and only if it is symmetric (if $\mathbb{K} = \mathbb{R}$) or Hermitian (if $\mathbb{K} = \mathbb{C}$) and, for all $x \in \mathbb{K}^d$,

$$\langle x, Mx \rangle \in \mathbb{R}^+.$$

It is *definite positive* if and only if it is semidefinite positive and, for all $x \in \mathbb{K}^d \setminus \{0\}$,

 $\langle x, Mx \rangle > 0.$

Proposition A.3: diagonalization of symmetric / Hermitian matrices

Let \mathbb{K} be \mathbb{R} or \mathbb{C} . Let $M \in \mathbb{K}^{d \times d}$ be a symmetric or Hermitian matrix. It can be diagonalized in an orthogonal basis, with real eigenvalues: there exist $\lambda_1, \ldots, \lambda_d$ in \mathbb{R} and (z_1, \ldots, z_d) an orthonormal basis of \mathbb{K}^d such that

$$X = \begin{pmatrix} z_1 \dots z_d \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 \dots & 0 \\ 0 & \lambda_2 & \vdots \\ \vdots & \ddots & \ddots \\ \vdots & \ddots & \lambda_d \end{pmatrix} (z_1 \dots z_d)^* = \sum_{k=1}^d \lambda_k z_k z_k^*.$$

Matrix M is semidefinite positive if and only if $\lambda_k \ge 0$ for all $k \le d$. It is definite positive if and only if $\lambda_k > 0$ for all $k \le d$.

Bibliography

- A. Barvinok. A course in convexity, volume 54. American Mathematical Society, 2002.
- E. J. Candès and T. Tao. Decoding by linear programming. IEEE transactions on information theory, 51(12):4203–4215, 2005.
- E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2):21–30, 2008.
- E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- A. Conca, D. Edidin, M. Hering, and C. Vinzant. Algebraic characterization of injectivity in phase retrieval. *Applied and Computational Harmonic Analysis*, 32(2):346–356, 2015.