



HAL
open science

Positive Semidefinite Relaxations for Imaging Science

Paul Catala

► **To cite this version:**

Paul Catala. Positive Semidefinite Relaxations for Imaging Science. Mathematics [math]. Ecole Normale Supérieure, 2020. English. NNT: . tel-03131464

HAL Id: tel-03131464

<https://theses.hal.science/tel-03131464v1>

Submitted on 4 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à l'Ecole Normale Supérieure

Relaxations Semi-Définies Positives pour l'Imagerie

Soutenue par

Paul CATALA

Le 01/10/2020

Ecole doctorale n° 386

**Sciences Mathématiques de
Paris Centre**

Spécialité

Mathématiques

Composition du jury :

Didier, HENRION Directeur de recherche CNRS, LAAS	<i>Président</i>
Laurent, CONDAT Research scientist, KAUST	<i>Rapporteur</i>
Carlos, FERNANDEZ-GRANDA Professeur, NYU	<i>Rapporteur</i>
Jean-François, CARDOSO Directeur de recherche CNRS, IAP	<i>Examineur</i>
Marianne, CLAUSEL Professeur, Université de Lorraine	<i>Examineur</i>
Georgina, HALL Professeur, INSEAD	<i>Examineur</i>
Vincent, DUVAL Chargé de recherche, Inria	<i>Co-directeur de thèse</i>
Gabriel, PEYRE Directeur de recherche CNRS, DMA	<i>Directeur de thèse</i>

Contents

Introduction	7
0.1 Motivations	9
0.2 Off-the-grid recovery	12
0.3 Contributions	20
1 Multivariate reconstruction from trigonometric moments	23
1.1 Introduction	24
1.2 Problem formulation	29
1.3 A multivariate extension of Prony	32
1.4 The positive case	41
1.5 ESPRIT	49
1.6 Extraction for non-flat data	53
2 Semidefinite relaxations for Beurling lasso	64
2.1 Introduction	65
2.2 Theoretical background on the Blasso	69
2.3 Spectral approximation of sensing operators	71
2.4 Lasserre’s hierarchy for the BLASSO	78
3 The Fourier-based Frank Wolfe algorithm	88
3.1 Introduction	89
3.2 Toeplitz relaxation	94
3.3 Fast Fourier Transform-based Frank-Wolfe (FFW)	102
3.4 Numerics	108
4 Semidefinite Approaches for Optimal transport	116
4.1 Introduction	117
4.2 Semidefinite formulations for OT	120
4.3 MMD unbalanced optimal transport	128
4.4 SDP transport in higher dimensions	135
4.5 Group-Blasso	138
4.6 Wasserstein-Blasso	141
Conclusion	148

Remerciements

Mes premiers et plus vifs remerciements vont à Gabriel et Vincent. Merci pour votre bienveillance, votre enthousiasme et votre soutien infaillible tout au long de ces quatre années. Entre ton dynamisme Gabriel et ta rigueur Vincent, je n' imagine pas de meilleurs conditions d'encadrement. Merci pour tout.

Je tiens à remercier mes rapporteurs Laurent Condat et Carlos Fernandez-Granda, ainsi que les membres du jury, Jean-François Cardoso, Marianne Clausel, Georgina Hall et Didier Henrion, pour l'honneur qu'ils me font de s'intéresser à ce travail.

Ces quatre années auraient été bien plus tristes sans les nombreuses personnes que j'ai eu la chance de rencontrer, entre Dauphine, l'ENS et l'Inria. Merci tout d'abord à l'incroyable équipe des doctorants du CEREMADE pour leur accueil chaleureux lors de mon stage de M2. Merci ensuite à toute l'équipe MOKAPLAN, et en particulier à Mélanie, Simon, Romain, Kathy, Gabriele et Andrea pour avoir égayé la plupart de mes mercredis. Merci enfin et surtout au DMA, et à toute la "team Gabriel", Augustin, Guillaume, Nicolas (merci pour les repas sans gluten cinq étoiles!), Arthur, Shuangjian, et bien sûr mes "frères" et "soeurs" de thèse, Thibault, Gwendoline, Aude, Lénaïc et Jonathan, que j'ai pris plaisir à côtoyer chaque jour.

Enfin, un immense merci à ma famille, Christine, Ivan, Claude, Flore et Aurore, à qui je dois essentiellement tout.

Introduction (English)

Summary of the thesis

The most common approach to imaging sciences is to discretize the problem over some predefined grid, and to constrain the object of interest (*e.g.* contours in an image, parameters of a distribution, etc.) to be restricted to this grid. In the last decade, a lot of effort has been put into off-the-grid models, that do not assume any discretization of the domain. These approaches lead to sharper analysis and stronger guarantees, but often come with higher computational costs. In the perspective of extending the scope of these methods to higher dimensional problems, the development of efficient algorithms remains to this day a crucial and a challenging problem.

This work studies the use of semidefinite programming for solving such off-the-grid inverse problems in imaging. We consider applications ranging from the sparse super-resolution problem, where one wishes to recover pointwise sources from low-resolution and noisy measurements, to optimal transport, which aims at finding an optimal mapping between several distributions. Such questions are ubiquitous in image processing and machine learning, and are involved in problems as diverse as deconvolution in medical and astronomical imaging, EEG/MEG reconstruction, compressed sensing, or mixture estimation.

Off-the-grid approaches may be divided into two main families of methods, namely parametric and variational. The algorithms developed in this thesis use ideas from both, and principally rely on the use of the so-called Lasserre’s hierarchy. The hierarchy gives a systematic way of approximating infinite-dimensional tasks with semidefinite programming (SDP) problems, by searching for the moments of the sought after measures rather than for the measures themselves. In many real-world applications however, the sizes of the involved semidefinite programs are prohibitive for usual solvers, often unable to take advantage of the structure of their variables.

The first contribution of this thesis, detailed in chapter 1, is a new parametric method to retrieve measures supported on curves or hyper-surfaces. The algorithm combines the traditional singular value decomposition methodology of parametric approaches with a new optimization-based joint-diagonalization step, and is able to recover a geometrically faithful discrete approximation of the continuous support. Our second contribution, which is explained in chapter 2, is a new semidefinite approximation in arbitrary dimension of the variational formulation of super-resolution, which is based on Lasserre’s hierarchy. Our third contribution, detailed in chapter 3, is a new solver for this moment approach, that unlocks the aforementioned computational bottleneck for the sparse recovery problem, and remains tractable in higher dimensions. Our approach takes advantage of the low-rank structure of its iterates to maintain computations with low memory and time costs. Finally, chapter 4 contains our last contribution, which is a new semidefinite solver for the optimal transport problem, based on the same moment approximation principle. We focus in particular on a relaxation of the group-BLASSO problem involving a Wasserstein penalization.

0.1 Motivations

Many problems in signal and image processing require to recover highly concentrated objects, *e.g.* pointwise sources or contour lines. While discretizing the search space in the recovery process allows for efficient solvers, these often suffer from severe accuracy limitations when trying to reconstruct such condensed signals. The main goal of this thesis is to develop efficient off-the-grid algorithms (*i.e.* requiring no discretization) for these tasks, targeting in particular applications in super-resolution and optimal transport. Both problems share the important feature of involving variables that may be modelled as Radon measures.

0.1.1 Sparse super-resolution

The sparse super-resolution problem aims at recovering pointwise sources from low-resolution and possibly noisy measurements. Point sources are an important physical model in signal processing and imaging science. They may represent molecules in microscopy, spectral lines in spectroscopy, neuron spikes in neuroscience, directions of arrival in antenna processing, or celestial objects in astronomical imaging. Beyond, they can also model statistical quantities, like the parameters of a Gaussian distribution in mixture estimation problems [Gribonval et al., 2017].

Recovering the fine details of a signal from which only partial information is accessible is an ubiquitous problem in image processing. In optical imaging for instance, the resolution of sensing devices is limited by diffraction, causing a point source to be imaged as an Airy pattern whose diameter is proportional to the wavelength of the illuminating light, hence typically preventing from observing subwavelength structures. The famous Rayleigh and Abbe criteria thus define minimal distances between objects that are necessary to resolve them. Super-resolution algorithms are used in the purpose of beating these limitations, and retrieving from a non-resolved image the true underlying signal with high accuracy.

Fluorescence microscopy. One of the most celebrated applications of these methods in the recent years concerns fluorescence microscopy. Fluorescence techniques for microscopy were developed around 2000 and earned their creators a Nobel prize in Chemistry in 2014. Among them, the single-molecule microscopy approaches [Moerner and Fromm, 2003], such as PALM [Hess et al., 2006, Betzig et al., 2006] or STORM [Rust et al., 2006], are of particular interest to us. To super-resolve a microscopic medium, fluorophores (either natural or chemical) may be alternatively activated and de-activated using a laser. Letting just a few of the molecules glow each time, one may thus enhance sparse subsets of the sought object. Assuming the glowing particles are sufficiently separated on each scan, the resulting images may then be super-imposed to obtain a dense super-image resolved at the nanometer level (where




Figure 1: Image of a focal adhesion with a classical optical system (left) and with the PALM system (right). Focal adhesions are molecular assemblies conveying exchanges between a cell and the extracellular medium. The red and green colors correspond to two different proteins. Images are taken from the [cell image library](#).

diffraction limited the resolution of optical systems to roughly 200 nanometers). Figure 1 shows an example of reconstruction with the PALM system.

In this process, super-resolution algorithms come into play to recover the particles' locations from each individual scan. Being able to recover accurately and efficiently subsets of particles as large as possible allows to diminish the time spent in the acquisition of snapshots, and hence improve the temporal resolution, allowing to observe microscopic phenomena occurring over short periods of time.

Radon measures modeling. A relevant mathematical model for pointwise sources are Dirac masses. A source located at a position x may thus be seen as the linear form δ_x over continuous functions, defined as the pointwise evaluation

$$\delta_x : \varphi \in \mathcal{C}_0 \mapsto \varphi(x).$$

This encourages to model our signals as Radon measures, *i.e.* continuous linear forms over the set of compactly supported continuous functions. The set of Radon measures (resp. nonnegative Radon measures) over a space X is denoted by $\mathcal{M}(X)$ (resp. $\mathcal{M}(X)_+$). In this study, we focus on the case where X is the d -dimensional torus, for some $d \in \mathbb{N}_*$, that is $X \equiv \mathbb{T}^d$ where $\mathbb{T} \stackrel{\text{def}}{=} \mathbb{R}/\mathbb{Z}$. Considering this periodic setting offers a lot of practical advantages, and many of the usual periodization artifacts can be handled easily, for instance by symmetrizing or zero-padding the domains.

Super-resolution tasks aim at retrieving fine scale details of a signal that are lost in the acquisition, due to resolution limitations, and might be corrupted by noise. In mathematical terms, the sensing process may generically be modeled as an integral operator over the space of Radon measures

$$\Phi : \mu \in \mathcal{M}(\mathbb{T}^d) \mapsto \int_{x \in \mathbb{T}^d} \varphi(x) d\mu(x) \in \mathcal{H} \quad (1)$$

where \mathcal{H} is a separable Hilbert space (typically, $\mathcal{H} \equiv \mathbb{C}^n$) and $\varphi : \mathbb{T}^d \rightarrow \mathcal{H}$ is a continuous function. This model encompasses a broad variety of scenarios, *e.g.* convolutions, Laplace transforms, etc., the most elementary one and most important to us being ideal low-pass filtering, which consists in returning the Fourier moments of the measure over a domain

$$\Omega_n \stackrel{\text{def.}}{=} \llbracket -n, n \rrbracket^d \quad (2)$$

for some $n \in \mathbb{N}^*$. We denote the corresponding operator by $\Phi \stackrel{\text{def.}}{=} \mathcal{F}_n$, *i.e.*

$$\mathcal{F}_n : \mu \mapsto \left(\int_{\mathbf{x} \in \mathbb{T}^d} e^{-2i\pi \langle \mathbf{k}, \mathbf{x} \rangle} d\mu(\mathbf{x}) \right)_{\mathbf{k} \in \Omega_n}, \quad (3)$$

and in this case, $\mathcal{H} \equiv \mathbb{C}^{|\Omega_n|}$ and $\varphi(\mathbf{x}) \stackrel{\text{def.}}{=}} (e^{-2i\pi \langle \mathbf{k}, \mathbf{x} \rangle})_{\mathbf{k} \in \Omega_n}$. The measurements are then of the form

$$\mathbf{y} \stackrel{\text{def.}}{=} \mathbf{y}_0 + w \stackrel{\text{def.}}{=} \Phi(\mu_0) + w \quad (4)$$

where $\mu_0 \in \mathcal{M}(\mathbb{T}^d)$ is the sought measure and $w \in \mathcal{H}$ accounts for measurement noise. Super-resolution thus consists in recovering μ_0 from the sole knowledge of \mathbf{y} .

We introduce in Chapter 3 a new solver for this problem, in arbitrary dimension. Our approach is based on semidefinite approximations of the so-called BLASSO (see section 0.2.2 below, and chapter 2). While the usual semidefinite programming algorithms cannot solve these relaxations in dimension higher than one due to their too large variables, our method reduces the complexity per iteration from $O(n^{2d})$ to $O(n^d \log n)$ by exploiting the low-rank structure of the problem, see section 2.4. Our contributions are further detailed in section 0.3 below.

0.1.2 Optimal transport

The optimal transport problem was initially stated in the 18th century by French mathematician Gaspard Monge, and has known several major breakthroughs during the 20th century, in particular thanks to the work of Kantorovitch (Kantorovich shared the Nobel Prize of Economy Sciences in 1975 with T. Koopmans for their work in optimum allocation of resources). Of particular interest for this thesis is the so-called Wasserstein distance $\mathcal{W}(\mu_1, \mu_2)$ between two probability distributions μ_1 and μ_2 . Informally, $\mathcal{W}(\mu_1, \mu_2)$ measures the amount of effort necessary to displace mass from μ_1 to μ_2 . For instance, if the mass of μ_1 is concentrated at a single point x and the mass of μ_2 at another location y , then $\mathcal{W}(\mu_1, \mu_2)$ is equal to some power of the actual ground distance between x and y .

Similarly to super-resolution, the densities may be modeled as Radon measures. Given two positive densities $\mu_1, \mu_2 \in \mathcal{M}(\mathbb{T}^d)_+$ of equal mass and a cost function $h : \mathbb{T}^d \times \mathbb{T}^d \rightarrow \mathbb{R}_+$ over \mathbb{T}^d , the optimal transport between μ_1 and μ_2 with respect to h may be retrieved as a solution of the optimization problem

$$\mathcal{W}_h(\mu_1, \mu_2) \stackrel{\text{def.}}{=} \inf_{\gamma \in \mathcal{M}(\mathbb{T}^d \times \mathbb{T}^d)_+} \int_{\mathbb{T}^d \times \mathbb{T}^d} h(\mathbf{x}, \mathbf{y}) d\gamma(\mathbf{x}, \mathbf{y}) \quad \text{s.t.} \quad \begin{cases} \pi_1(\gamma) = \mu_1 \\ \pi_2(\gamma) = \mu_2 \end{cases}, \quad (5)$$

where π_1 and π_2 extract the two marginals of the coupling γ . This is known as Kantorovitch’s formulation.

When $h(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \mathbf{y})^p$ for some distance d and integer $p \geq 1$, then $\mathcal{W}_h^{1/p}$ is the so-called Wasserstein distance. Compared to many other distances over measure spaces, Wasserstein distances have the main advantage of metrizing weak convergence, so that they take a fair account of the geometry of the densities. They are now popular in image processing for shape matching problems [Feydy et al., 2017] or supervised learning tasks on histograms [Bonneel et al., 2016], as well as in machine learning, *e.g.* for metric learning [Cuturi and Avis, 2014] or unsupervised learning, to determine parameters of distributions in generative models [Genevay et al., 2018].

Most of the methods to solve problem (OT) rely on a fixed discretization of the state space, where the marginals are approximated by discrete measures (with high cardinality), located on a predefined grid. In chapter 4, we consider a novel direction, where the domain is not discretized spatially but spectrally. Making use of Lasserre’s hierarchy, our method recovers the moments of the transport plan, under the constraint that a *finite* number of them satisfy the marginal equalities of (OT), see section 4.2.

0.2 Off-the-grid recovery

Whether it is for super-resolution or optimal transport problems, many solvers build their solutions over a predefined grid, which inevitably limits the accuracy of the reconstruction. In contrast, we focus in this thesis on “off-the-grid” methods, allowing for sharper recoveries, free from discretization artifacts, and whose numerical complexity is not dependent on the size of a grid.

0.2.1 Parametric approaches for sparse inverse problems

Parametric approaches are methods that retrieve the parameters of an *assumed* underlying model of the data. When the model is indeed satisfied, such methods have the ability to perform exact recovery of the sought signal. However, their high level of performance depends on whether this assumption indeed holds, and on how closely the mathematical model matches the actual physical process that produced the signals.

In our case, the signal model of interest is the harmonic model, consisting in s complex exponentials in noise

$$y_k = \sum_{j=1}^s a_j e^{-2i\pi k x_j} + w_k \stackrel{\text{def.}}{=} c_k + w_k, \quad k \in \Omega_n \quad (6)$$

where $\Omega_n = \llbracket -n, n \rrbracket$, $n \in \mathbb{N}_*$, and $x_j \in \mathbb{T}$ and $a_j \in \mathbb{C}_*$ are the parameters to retrieve. This is equivalent to the sparse spikes deconvolution problem: the

vector y may be seen as the (noisy) Fourier moments of the sparse measure

$$\mu_0 = \sum_{j=1}^s a_j \delta_{x_j} \in \mathcal{M}(\mathbb{T}),$$

from which we want to recover μ_0 .

Prony’s method All parametric approaches to solve this problem have as common foundation an idea that was brought by Prony in 1795 [R. de Prony, 1795], consisting in searching for the frequencies x_j as the root of the polynomial

$$P_0(z) = \prod_{j=1}^s (z - e^{2i\pi x_j}). \quad (7)$$

When the observations are not corrupted by noise, the coefficients $(p_{0,k})$ of this polynomial may be easily found as they satisfy the simple linear equations

$$\sum_{k=0}^s p_{0,k} c_{l-k} = 0, \quad \text{for } l = 0, \dots, n \quad (8)$$

Recovering the frequencies thus amounts to solving a linear system – assuming that n is sufficiently large for (8) to be solvable, followed by a root-finding task, which is typically performed by diagonalizing the companion matrix of P . The amplitudes may finally be recovered by plugging the recovered positions into the model (6). We give further details on this seminal method in section 1.3.1.

Subspace methods Extensions of Prony’s method capable of dealing with noisy data were developed in the eighties, mostly in the context of direction-of-arrivals (DOA) estimation in array processing. In such problems, spatially propagating waves, arriving from far-field sources, produce complex exponential signals that are collected by an array of sensors. The frequencies of the exponentials are determined by the angles of arrival of the impinging waves, so that the frequencies estimation problem is known as DOA estimation.

The simplest model of array is the uniform linear array (ULA), in which the sensors are equally spaced on a line. In this model, the waves reach each successive sensor with a constant delay (depending on the spacing between the sensors and on the impinging angle), which translates into regular phase modulations in the measured complex exponentials: at a given time, the array returns a snapshot $\mathbf{y}^{(k)} \in \mathbb{C}^M$ of M (the number of sensors) measurements, where each entry follows the harmonic model (6). Given a sequence of L such snapshots, subspace methods, such as Pisarenko’s method [Pisarenko, 1973], the MULTIPLE SIGNAL CLASSIFICATION algorithm [Schmidt, 1986] or the ESTIMATION OF SIGNAL PARAMETERS VIA ROTATIONAL INVARIANCE TECHNIQUES [Roy et al., 1986, Roy and Kailath, 1989], are able to determine a *signal subspace*, spanned by the vectors

$$v_M(x_j) \stackrel{\text{def.}}{=} [1 \quad e^{-2i\pi x_j} \quad \dots \quad e^{-2i\pi(M-1)x_j}]^\top \in \mathbb{C}^M, \quad j = 1, \dots, s, \quad (9)$$

where x_j are the sources locations (or angles). Most often, this subspace decomposition is revealed by the singular value decomposition of a specific matrix built from the snapshots. In Pisarenko and MUSIC, this decomposition is performed from an approximate correlation matrix

$$\widehat{R}(\mathbf{y}) \stackrel{\text{def.}}{=} \frac{1}{L} \sum_{k=1}^L \mathbf{y}^{(k)} (\mathbf{y}^{(k)})^H \in \mathbb{C}^{M \times M}, \quad (10)$$

while ESPRIT rather considers the data matrix

$$E(\mathbf{y}) \stackrel{\text{def.}}{=} (\mathbf{y}^{(0)} \quad \dots \quad \mathbf{y}^{(L-1)})^\top \in \mathbb{C}^{L \times M}. \quad (11)$$

The ESPRIT method is further detailed in section 1.5.

The performance and resolution capacity of these approaches obviously depend on the number of snapshots that are accessible. However, if the number M of sensors is sufficiently large, that is, *assuming that* $M \geq 2s$, where s is the number of sources to recover, then MUSIC and ESPRIT may be adapted to the single-snapshot case [Liao and Fannjiang, 2014, Fannjiang, 2016]. In that case, the matrix of interest may be defined from one snapshot $[y_{-n} \quad \dots \quad y_n]$ ($n \geq s$) as

$$T(\mathbf{y}) \stackrel{\text{def.}}{=} \begin{pmatrix} y_0 & y_{-1} & \dots & y_{-s+1} \\ y_1 & y_0 & \dots & y_{-s+2} \\ \vdots & \vdots & & \vdots \\ y_n & y_{n-1} & \dots & y_{n-s+1} \end{pmatrix}, \quad (12)$$

where it is assumed that the model order s is known. Note that in the noiseless case, $T(\mathbf{c})$ is exactly the matrix of the linear system (8) in Prony’s method. The matrix $T(\mathbf{c})$, which is a *truncated moment matrix* of the measure μ_0 representing \mathbf{c} , plays a crucial role throughout this dissertation.

Guarantees In noiseless regimes, and in one dimension, Prony’s method and its subspace-based extensions are able to recover exactly arbitrarily close atoms of signed amplitudes. This contrasts with convex methods, which require a minimal separation distance in the signed case to guarantee an exact recovery, see section 2.1. However, all parametric approaches require to know the number s of sources beforehand, while optimization-based methods do not.

In noisy regimes, the robustness of MUSIC and ESPRIT depends on the conditioning of Vandermonde matrices with nodes $(e^{-2i\pi x_j})_{j=1,\dots,s}$ that appears in the decomposition of $\widehat{R}(\mathbf{y})$ or $E(\mathbf{y})$. Stability results for these subspace methods usually take the form of bounds on the reconstruction error with respect to the minimum singular value of these Vandermonde matrices, see for instance [Liao and Fannjiang, 2014] for MUSIC, or [Fannjiang, 2016, Potts and Tasche, 2017] for ESPRIT. These bounds crucially depend on the geometry of the nodes. They are well-known in the well-separated regime [Moitra, 2015], *i.e.* when the minimal separation distance Δ is greater than $1/n$ (n being the number of accessible moments), but the “nearly-colliding”

regime on the other hand (*i.e.* when $\Delta < \frac{1}{n}$) is more involved and has been the focus of recent works, often requiring technical assumptions on the geometry of the nodes. In particular, many results focus on the case of clusters of colliding frequencies, see *e.g.* [Kunis and Nagel, 2018, Li and Liao, 2019, Li et al., 2019, Batenkov et al., 2020].

Higher dimensions Extensions of Prony-like methods to multidimensional settings are difficult to implement, and many different approaches have been proposed in the literature, see *e.g.* [Sacchini et al., 1993, Haardt and Nossek, 1998, Rouquette and Najim, 2001]. Multivariate recovery guarantees have more recently been investigated, by connecting Prony’s method to polynomial system solving problems [Kunis et al., 2016, Sauer, 2018]. In particular, these recent works reveal weaker guarantees than in one dimension, with for instance the appearance of a minimal separation distance even for noiseless measurements [Ehler et al., 2019]. We detail these results in section 1.3. Multivariate extensions of Prony also typically need $O(s^d)$ measurements to reconstruct s Dirac masses, unless some clever sampling scheme is used [Cuyt and Lee, 2018].

0.2.2 Variational formulations

Parametric approaches offer simple algorithms (often based on a single singular value decomposition) with strong recovery guarantees, under constraining assumptions on the signal model. In particular, they can only be applied to the case where the observations consist in a convolution, and to our knowledge there exists no extensions to more general measurements, such as for example sub-sampled convolution, or random/partial sampling of the moments in compressed sensing applications.

Variational approaches on the other hand are more flexible as they do not assume any particular structure on the measurements or the noise, nor require to know the number of sources. Generally speaking, they consist in searching among all [measures] the one that minimizes the sum of a term measuring the error with respect to the observations (the data-fidelity term) and a structure-inducing term (the regularization term). These methods allow for a fine analysis and understanding of the geometry of their solutions [Vaïter et al., 2015]. Examples of well-understood priors include sparsity [Donoho, 1992] or group-sparsity, piecewise smoothness [Chambolle et al., 2016], or low-rank (which generalizes sparsity to matrix-valued functions).

Regularization of measures Regularization approaches have become very popular for recovery tasks in Hilbert spaces, leading to famous and extensively studied problems such as ℓ^1 -minimization [Tibshirani, 1996, Chen et al., 2001, Candès et al., 2006] or total variation (of gradient) minimization [Chambolle, 2004, Condat, 2017, Abergel and Moisan, 2017]. On the other hand, many recent works have been interested in inverse problems in more general Banach spaces, such as the space of Radon measures

[Bredies and Pikkarainen, 2013]. These Banach extensions have led to several significant breakthroughs, in particular for sparse recovery tasks via convex methods [Candès and Fernandez-Granda, 2014]. In this thesis, we mainly focus on two types of off-the-grid regularizers: total variation (of measures), and Wasserstein distances.

Total variation regularization The total variation of measures is the infinite-dimensional counterpart of the ℓ^1 -norm of vectors. Given $\mu \in \mathcal{M}(\mathbb{T}^d)$, its total variation reads

$$|\mu|(\mathbb{T}^d) \stackrel{\text{def.}}{=} \sup \left\{ \int_{\mathbb{T}^d} \eta d\mu ; \eta \in \mathcal{C}(\mathbb{T}^d) \quad \text{and} \quad \|\eta\|_\infty \leq 1 \right\}.$$

To see how this generalizes the ℓ^1 -norm, note for instance that if $\mu = \sum a_i \delta_{x_i}$, then $|\mu|(\mathbb{T}^d) = \|a\|_1$. The super-resolution problem of recovering a measure μ_0 from measurements \mathbf{y} of the general form (4) may then be expressed as the optimization problem

$$\min_{\mu \in \mathcal{M}(\mathbb{T}^d)} \frac{1}{2\lambda} \|\Phi\mu - \mathbf{y}\|_{\mathcal{H}}^2 + |\mu|(\mathbb{T}^d) \quad (13)$$

where λ should be adapted to the noise level. In dimension $d = 1$, the noiseless case (*i.e.* $\lambda = 0$) was first investigated in [de Castro and Gamboa, 2012], for nonnegative measures, and then extended by the seminal paper [Candès and Fernandez-Granda, 2014], in which the authors show that a discrete measure with arbitrary amplitudes may be recovered exactly as a solution of the optimization assuming some minimal separation distance between the spikes (while none is needed in the nonnegative case). When $\lambda > 0$, problem (2.3) is known as BLASSO [Azaïs et al., 2015] (short for Beurling LASSO), and its super-resolution abilities are now well-understood in one dimension [Duval and Peyré, 2015, Denoyelle et al., 2017, Morgenshtern, 2020]. This problem has also received considerable attention for compressed sensing applications, in the case where Φ satisfies some randomness assumptions [Tang et al., 2013, Poon et al., 2020].

Wasserstein regularization Due to their good geometrical properties, Wasserstein distances may serve as priors in inverse problems, to exploit similarities between solutions and a reference for instance, or between solutions themselves. In multi-task sparse super-resolution for instance, one is concerned with recovering simultaneously several groups of sources, sharing similar supports. These sources may model positions on the brain surface of several patients in MEG imaging [Owen et al., 2009], or positions in genomes in genomics [Laurent, 2010]. Having access to measurements stemming from several subjects, one would like to constrain the positions recovered across the tasks to be close, and take advantage of the multiple data to improve the statistical recovery. This can be achieved via optimal transport regularization, enforcing the recovered measures to be close with respect to a Wasserstein distance. This idea was first introduced in a finite-dimensional setting in

[Janati et al., 2019b, Janati et al., 2019a], where the authors use Wasserstein priors in addition to ℓ^1 -priors, and can be naturally extended to an off-the-grid approach by considering the optimization problem

$$\min_{(\mu_i), \bar{\mu} \in \mathcal{M}(\mathbb{T}^d)} \frac{1}{2\lambda} \sum_i \|\Phi \mu_i - \mathbf{y}_i\|_{\mathcal{H}}^2 + \sum_i |\mu_i|(\mathbb{T}^d) + \sum_i \mathcal{W}(\mu_i, \bar{\mu}). \quad (14)$$

Guarantees for sparse recovery problems. In the noiseless regime, the seminal paper [Candès and Fernandez-Granda, 2014] shows that a sparse measure $\mu_0 \in \mathcal{M}(\mathbb{T})$ may be exactly recovered as a solution of (2.3) assuming some minimal separation distance of the order of $1/n$ between its atoms, thus not improving the natural Rayleigh limit (this requirement disappears when the measure is positive). This analysis has also been extended to noisy settings [Candès and Fernandez-Granda, 2013, Azaïs et al., 2015], but offer no super-resolution guarantee for signed measures. The separation distance is shown to be a necessary requirement in [Duval and Peyré, 2015, Tang, 2015, Ferreira Da Costa and Dai, 2018].

On the other hand, variational approaches allow for sharper analysis of their super-resolution abilities in the positive case. In noisy regimes, the study of how the signal-to-noise ratio should scale with respect to the spikes separation distance was initiated in [Donoho, 1992], further analyzed in [Demanet and Nguyen, 2015] and it is shown in [Denoyelle et al., 2017] that a scaling $\|w\| = O(\Delta^{2N-1})$, where N is the number of spikes clustered in a radius of Δ , ensures exact recovery under a non-degeneracy condition on the forward operator. This result was extended to the 2-D case in [Poon and Peyré, 2018], for $N = 2$. One of the main difficulties arising from this multidimensional setting is that support stability in that case depends on the geometric configuration of the sources.

Computational approaches. While the mathematical analysis of problems such as (2.3) or (14) leads to sharper results than discretized approaches, their numerical resolution on the other hand requires to optimize over a Banach space, which is a challenging task. Several numerical approaches to this problem are possible, which can be divided into three main categories.

Fixed spatial discretization. A first method is to look for the support of the target measure over a predefined grid. This leads back to the traditional finite-dimensional setting, benefiting from many well-studied solvers. For super-resolution, the Hilbert analogous of (2.3) is the LASSO [Tibshirani, 1996], for which the typical solvers are proximal methods, for instance: (accelerated) proximal gradient descent [Daubechies et al., 2004, Blumensath and Davies, 2008, Beck and Teboulle, 2009], primal-dual [Chambolle and Pock, 2011], forward-backward [Beck and Teboulle, 2010, Combettes and Pesquet, 2011], Alternating Direction Method of Multipliers [Boyd et al., 2011], or proximal coordinate descent [Fu, 1998, Friedman et al., 2010].

In optimal transport, there are mainly three classes of methods, depending on whether the measures transported are discrete or not. Discrete optimal transport (both marginals are discrete) is a simple linear program, and is usually solved using Sinkhorn iterations [Sinkhorn, 1964], for which very fast computations have been proposed in [Cuturi, 2013]. For the semidiscrete case (one marginal is discrete, one is continuous), a quasi-Newton solver is proposed in [Mérigot, 2011], and the use of a Newton solver is investigated in [De Goes et al., 2012] for computer graphics applications and in [Lévy, 2015] for 3-D volume processing. Finally, for continuous optimal transport (both marginals are continuous), the problem may be reduced to solving a Monge-Ampère equation, whose discretization is challenging [Benamou et al., 2015]. A solver is proposed for specific settings in [Benamou et al., 2016].

Partial or adaptive discretization. In the context of sparse optimization, several methods have been proposed that build their solutions atom by atom: at each step, a new point is added (or subtracted), and one may then work over the resulting grid, as well as optimize the positions of the current atoms – which usually requires resorting to non-convex optimization schemes. This encompasses Frank-Wolfe algorithms, which have been used to solve the super-resolution problem (2.3) in [Boyd et al., 2015, Denoyelle et al., 2020], and exchange algorithms, which may rather be used on the dual of the BLASSO [Flinth et al., 2019]. Such approaches have the advantage of requiring no Hilbertian structure on the optimization domain, making them good candidates to work in spaces of measures. Convergence results on these methods may be found in [Demyanov and Rubinov, 1970].

Spectral discretization. In [Tang et al., 2013], the authors show that the BLASSO with $\Phi \stackrel{\text{def.}}{=} \mathcal{F}_n$ is equivalent in one dimension to the semidefinite program

$$\min \frac{1}{2\lambda} \|z - y\|^2 + \frac{1}{2}(c_0 + \tau) \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} R & z \\ z^* & \tau \end{bmatrix} \succeq 0 \\ R = T_n(c) \end{cases}. \quad (15)$$

The main idea behind this reformulation is that one has replaced the variable $\mu \in \mathcal{M}(\mathbb{T})$ of (2.3) by truncated moment sequences of $|\mu|$ and μ , resulting in the variable c and z respectively. Such a change of variable is valid in dimension one thanks to the Carathéodory-Toeplitz theorem [Carathéodory, 1911, Toeplitz, 1911], which ensures that any sequence c such that $T_n(c) \succeq 0$ admits a positive representing measure, *i.e.* a measure μ satisfying

$$c_k = \int_{\mathbb{T}^d} e^{-2i\pi kx} d\mu(x). \quad (16)$$

Thus, in one dimension, the infinite-dimensional problem (2.3) may be cast as the simple, finite-dimensional semidefinite program $(\mathcal{P}_\lambda^{(\ell)}(y))$, which operates in the spectral domain and recovers the moments of the sought measure.

Having retrieved these denoised moments, a Prony step is typically applied to extract the solutions. The dual formulation of $(\mathcal{P}_\lambda^{(\ell)}(y))$ was introduced earlier in [Candès and Fernandez-Granda, 2014] as a reformulation of the dual of the BLASSO, which reads

$$\sup \langle y, p \rangle - \frac{\lambda}{2} \|p\|^2 \quad \text{s.t.} \quad \|\Phi^* p\|_\infty \leq 1. \quad (17)$$

In that case, an equivalent semidefinite program is obtained by replacing the positivity constraint $1 - \|\Phi^* p\|_\infty \geq 0$ by the constraint that the polynomial $1 - |\Phi^* p|^2$ be sum-of-squares (sos), which can be expressed as a semidefinite program. Again, this trick is possible in one dimension due to the Fejér-Riesz theorem [Fejér, 1916, Riesz, 1916], which states that any *univariate* nonnegative trigonometric polynomial admits a sum of squares decomposition.

Such explicit, SDP-representable characterizations of the moment cone on the one hand (*i.e.* the cone of sequences admitting a nonnegative representing measure) and of the cone of nonnegative polynomials on the other hand are not available in higher dimensions. The moment/sos approach may nonetheless be applied in these settings by resorting to Lasserre’s hierarchy [Lasserre, 2001], which provides SDP approximations of these cones, involving sequences of semidefinite programs of increasing size. In the multivariate case however, there are no systematic guarantees on whether the hierarchy converges at a finite order of relaxation. These questions of finite convergence are central in semidefinite approximations methods, and are connected to the so-called *flatness property* of positive sequences [Curto and Fialkow, 1996, Curto and Fialkow, 1998, Laurent, 2005], that characterize the existence and uniqueness of a representing measure for a given sequence (see section 1.4).

0.2.3 Beyond sparse recovery

The problem of reconstructing measures supported on curves, or on more general algebraic domains, rather than points naturally appears when considering more complex geometrical objects, such as edges in an image, or transport plans, which are typically graphs of functions. In 1-D, the finite-rate-of-innovation (FRI) framework extends Prony’s method to a much wider class of signals, such as piecewise polynomials or nonuniform splines, and has been actively investigated in signal processing [Vetterli et al., 2002, Maravic and Vetterli, 2005, Dragotti et al., 2005, Blu et al., 2008].

FRI in higher dimensions. Several extensions of these methods to multi-dimensional signals have been studied, in particular for the recovery of polygonal shapes [Golub et al., 1999, Milanfar et al., 1995, Shukla and Dragotti, 2007, Chen et al., 2012]. These “shape from moments” inverse problems share many similarities with sparse recovery tasks, since only a finite set of vertices has to be retrieved. More recently, [Pan et al., 2014, Ongie and Jacob, 2016] considered the problem of recovering 2-D piecewise constant images, whose discontinuities are located on the zero set of a trigonometric polynomial. Similarly to

Prony’s method, their approach relies on the annihilation of the coefficients of this trigonometric curve by convolution with the Fourier transform of the *derivatives* (in the distribution sense) of the image. However, their study is limited to the recovery of curves embedded in 2-D domains, and cannot be extended to more general hypersurfaces.

The Christoffel function. In another line of work, many recent studies have focused on the use of Christoffel-Darboux polynomials, or their inverse, known as Christoffel functions, to extract relevant geometrical information on the support or density of measures [Lasserre and Pauwels, 2016, Lasserre and Pauwels, 2019, Pauwels et al., 2020, Marx et al., 2019], see section 1.6.4. Christoffel-Darboux polynomials are sos polynomials whose coefficients depend only on the moment of the underlying measure and whose degree depend on how many moments are considered. They are well-known in interpolation and approximation theory [Dunkl and Xu, 2001]. Their asymptotic behavior, *i.e.* when their degree goes to infinity, provides important information on the geometry of the underlying measure [Gustafsson et al., 2009, Kroó and Lubinsky, 2012, Lasserre and Pauwels, 2019], and the recent works [Lasserre and Pauwels, 2016, Lasserre and Pauwels, 2019, Pauwels et al., 2020] shed a new light on their applicability in data science problems, for instance to estimate the shape of a point cloud or the presence of outliers.

0.3 Contributions

Chapter 1 A natural extension of Prony’s algorithm to the multivariate case consists in encoding the points to retrieve as the *common* zeros of a few multivariate polynomials. Section 1.3 is a review of an algebraic approach to solve this multivariate recovery problem, whose theoretical analysis is deeply connected to the trigonometric moment problem.

Our first contribution, detailed in section 1.4, is a proof of the flatness property for the truncated trigonometric moment problem (see Theorem 6). The flatness property characterizes trigonometric sequences that admit a sparse representing measure. Our proof for the trigonometric case is in continuity of the seminal works [Curto and Fialkow, 1996] and [Laurent, 2010], in the complex and real cases respectively, and highlights a clear connection between the three settings. This result brings a new perspective on the guarantee of success of the multivariate Prony’s algorithm in the positive case, without any minimal separation assumption when there is no noise, as discussed in section 0.2.1 above. It is also central for the semidefinite approximations scheme discussed above, that we use in chapter 2.

The nucleus of the multivariate extraction method consists in the joint diagonalization of a few matrices. In section 1.6, we introduce our second contribution, which is an extension of the recovery procedure to the case of non-sparse measures. Our approach is based on a gradient descent algorithm to perform the joint diagonalization step, replacing the usual random linear combination

method that is used in sparse cases. We show in section 1.6.2 the limits of the linear combination approach in the case of non-sparse recovery problems. Section 1.6.3 then introduces our optimization-based joint diagonalization algorithm, that allows to recover with good accuracy measures supported on generic algebraic varieties. We give numerical illustrations of our method in section 1.6.4.

Chapter 2 Prony’s approaches are sensitive to noise and cannot be applied when the sensing operator is not a convolution. The first contribution of this chapter is a spectral approximation scheme of the super-resolution sensing operator, detailed in section 2.3, that is useful for later applying semidefinite hierarchies to the BLASSO. We provide a method to approximate any forward model by an equivalent spectral operator (Definition 8), that satisfies the structural requirements of Lasserre’s hierarchy and that we show can be arbitrarily close to the original operator (Proposition 5). Our approach gives a novel means to tackle complex, non-translation-invariant measurements, such as subsampled convolution or foveation, see section 2.3.2.

Section 2.4 introduces our second contribution, which is a semidefinite approximation of the BLASSO and its dual in arbitrary dimension. Our formulations generalize those of [Tang et al., 2013] in the uni-dimensional case, or of [Candès and Fernandez-Granda, 2014] for the dual. In comparison to the adaptive discretization approaches detailed before in section 0.2.2, they operate over a finite-dimensional space while involving no discretization of the reconstruction domain. We prove in section 2.4.3 a connection between the polar decomposition of the measure solving the BLASSO and the solutions of *lowest rank* of these semidefinite approximations, in the case of finite convergence of the hierarchy. This low-rank property is at the heart of the numerical solver we describe in chapter 3.

Chapter 3 For a measurement vector consisting of n trigonometric moments of the sought measure (or more general linear measurements), the size of the semidefinite approximations for the BLASSO is typically $O(n^{2d})$, where d is the dimension, and usual interior points solver are quickly inefficient when n or d increase.

The first contribution of this chapter is introduced in section 3.2 and consists in a penalized approximation of the semidefinite formulations detailed in section 2.4, that drastically simplifies their numerical resolution. We highlight in section 3.2.4 the many similarities that this penalization shares with augmented Lagrangian methods. We give in section 3.2.3 numerical evidence of the stability of its solutions with respect to those of the original problem, in particular in terms of their rank.

Our second contribution is the Fourier-based Frank Wolfe (FFW) algorithm, detailed in section 3.3. Our approach is at the crossroad of two popular methodologies for improving the efficiency of semidefinite optimization. On the one hand, it targets low-rank solutions by maintaining iterates as convex combinations of a few matrices of rank one, in the spirit of Frank-Wolfe algorithm

[Frank and Wolfe, 1956] – which is one major representant of the off-the-grid solvers described above in section 0.2.2. Each of these matrices is the solution of a linear oracle that requires low complexity computations. On the other hand, the variable are stored in the factorized form UU^H throughout all the iterations, in the spirit of Burer-Monteiro [Burer and Monteiro, 2003]. This induces a considerable gain in spatial memory, and further enables extremely efficient Fast Fourier Transforms based computations, that we detail in section 3.3.2. The resulting algorithm reaches a complexity of $O(n^d \log n)$ per iteration. We benchmark in section 3.4 the performance of our method both on synthetic data and real data. In particular, we provide examples of results on data coming from the Single-Molecule Localization Microscopy (SMLM) challenge [Group, 2013].

Chapter 4 This chapter investigates the use of semidefinite hierarchies for optimal transport problems. Our first contribution introduces the semidefinite relaxations of optimal transport and its dual, see section 4.2. In comparison with the usual Lasserre’s framework, described in section 0.2.2 and used on the BLASSO in chapter 2, the optimal transport setting requires a few simple additional technical details to deal with the infinite number of moment constraints, that we provide. These spectral formulations give a new way of computing Wasserstein distances without discretizing the state space.

Our second contribution is an unbalanced formulation of the optimal transport problem, involving maximum mean discrepancy (MMD) norms, that we present in section 4.3. Our formulation is an alternative to the usual unbalanced formulations that employ divergences between measures, and has the main advantage of being well suited for semidefinite relaxations. We highlight the strong similarities that this unbalanced problem shares with the BLASSO, and provide numerical evidence of the instability of the solutions as the MMD penalization goes to infinity.

Semidefinite approximations offer a new off-the-grid method to solve optimal transport, for which usual solvers often rely on a discretization of the domain. However, computing a transport between d -dimensional measures requires to retrieve a $2d$ -dimensional transport plan, thus leading to semidefinite variables of size $O(n^{2d})$, far too large for usual solvers. We show in section 4.4 numerical illustrations of the use of the FFW algorithm introduced in chapter 3 to compute Wasserstein distances between measures. The FFW algorithm retrieves the moment a transport plan, non necessarily sparse, on which we apply the extraction procedure of section 1.6 to recover the coupling. Finally, our third and last contribution is a new solver for the multi-task super-resolution problem, that generalizes the recent work [Janati et al., 2019b] to an off-the-grid setting. We use a Wasserstein prior as a coupling between several BLASSO to simultaneously recover the support of several close measures. We provide numerical illustrations of our semidefinite-based implementation.

Chapter 1

Multivariate reconstruction from trigonometric moments

Abstract

The guiding thread of this dissertation is to design and study off-the-grid super-resolution and optimal transport methods. Instead of resorting to a fixed spatial discretization, one rather searches for the moments of the objects to be recovered. While chapters 2 and 3 focus on efficiently optimizing over moment sequences, this chapter is dedicated to the remaining fundamental matter of reconstructing a Radon measure from a few of its Fourier coefficients. These reconstruction techniques are central in semidefinite programming recovery algorithms, where problems are approximated by semidefinite hierarchies, that precisely produce low order moments of the sought measures. In the univariate case, Prony's methods solve the problem and perfectly recover the sparsest (*i.e.* with smallest support) measure matching the moments, when these are not corrupted by noise. The analysis of the multivariate case on the other hand is more involved. This chapter first reviews some of the existing methods, and their guarantees in the noiseless case, for both signed and positive measures. In the positive case in particular, we show that exact recovery is achieved assuming a flatness property on the trigonometric sequence. Our proof adapts to the trigonometric setting a well-known result in the complex and real cases. From the numerical standpoint, several generalizations of Prony-like methods to arbitrary dimension have been proposed, which often rely on a joint diagonalization procedure. We propose as a second contribution a new recovery algorithm, based on the optimization of a co-diagonalization energy using multiplicative updates. We illustrate the improvement offered by this approach over simpler projection-based diagonalization. We also show that it can be applied to recover sparse approximation of non-sparse measures, which shows the robustness of the approach to approximate co-diagonalization.

Contents

1.1	Introduction	24
1.1.1	Multivariate parametric approaches	25
1.1.2	Previous works	26
1.1.3	Contributions	28
1.2	Problem formulation	29
1.3	A multivariate extension of Prony	32
1.3.1	The univariate case	32
1.3.2	Multivariate recovery via spectral analysis	35
1.3.3	The simpler theory of infinite matrices	38
1.3.4	Back to truncated matrices	40
1.4	The positive case	41
1.4.1	The truncated moment problem	41
1.4.2	Proof of Theorem 7	43
1.4.3	Guarantees for Prony	49
1.5	ESPRIT	49
1.5.1	Univariate ESPRIT	50
1.5.2	Multivariate ESPRIT	52
1.6	Extraction for non-flat data	53
1.6.1	Rank of moment matrices	53
1.6.2	Joint diagonalization by random linear combination	54
1.6.3	Joint diagonalization by optimization	57
1.6.4	Numerical illustrations	60

This chapter is dedicated to the study of a multivariate matrix-pencil scheme, boiling down to the joint diagonalization of a few non-Hermitian matrices. We improve the usual method by resorting to an approximate joint diagonalization scheme, based on optimization.

1.1 Introduction

A recurrent problem throughout this thesis is to retrieve a discrete measure (*i.e.* a sum of Dirac masses), given only a few of its moments. We describe in this chapter a parametric method which recovers the parameters of the sought measures (*i.e.* positions and amplitudes) via simple linear algebra. This may serve for sparse recovery problems in microscopy or astronomical imaging for instance, where the sought objects are indeed pointwise sources, see Chapter 3. In another line of work, such recovery algorithms may be used for finding good sparse approximations of continuous objects, *e.g.* edges in images, or transport maps. Some applications of the recovery process to optimal transport are discussed in Chapter 4.

1.1.1 Multivariate parametric approaches

Parametric approaches for the recovery of a sparse measure from some low-order moments (trigonometric or else) all derive from Prony’s method, a technique that dates back to the 18th century and bears the name of its inventor, a French engineer. These algorithms were extensively studied from the seventies to the nineties, owing in particular to their strong super-resolving properties – we refer to [Krim and Viberg, 1996] for a thorough overview.

Prony and super-resolution: the univariate case. Given $2n + 1$ consecutive Fourier coefficients of a Radon measure μ , the theory of super-resolution seeks to provide recovery guarantees when μ is composed of spikes separated by a distance less than $1/n$, the classical Nyquist limit. Convex minimization methods, which are the focus of Chapter 2, are known to achieve super-resolution when μ is non-negative [Candès and Fernandez-Granda, 2014, Denoyelle et al., 2017], but may fail in the other cases, even in the noiseless regime, unless some minimal separation distance is satisfied [Candès and Fernandez-Granda, 2014]. On the other hand, Prony’s method provides exact recovery of closely spaced atoms regardless of the signs of the amplitudes, when there is no noise.

Difficulties arise when the observations are corrupted by noise, to which the method is highly sensitive, or even in the noiseless case when one seeks to extend the approach to multidimensional settings. Numerical instability to noise has led to numerous robust extensions, such as MUSIC [Schmidt, 1986], ESPRIT [Roy and Kailath, 1989] or matrix pencils [Hua and Sarkar, 1989] to name a few, which achieved major breakthrough in DOA estimation, see [Krim and Viberg, 1996] for a review, and are now widely used in applications, ranging from brain imaging [Mosher et al., 1993] to personal communications [Ström et al., 1994].

Multivariate extensions. On the other hand, in spite of their efficiency and strong guarantees in one dimension, Prony-based approaches remain non-trivial to extend to higher dimensions, and multivariate extraction procedures have been an active path of research for the last twenty years [Hua, 1992, Haardt and Nossek, 1998, Sauer, 2017, Ehler et al., 2019]. Moreover, in contrast to the univariate case, the analysis of these extensions is more difficult when no assumption are made on the sign of the underlying measure.

This chapter focuses on a multivariate approach based on joint eigenvalue analysis, which has its roots in the field of polynomial system solving, see in particular [Möller and Stetter, 1995]. We refer to it as a multivariate Prony’s method, or also, following [Ehler et al., 2019], as a matrix pencil method, since its central step consist in the joint diagonalization of a few matrices. We review some of the main tools from both algebraic geometry (Section 1.3) and moment theory (Section 1.4) that are necessary for understanding and analyzing this extraction procedure. In particular, we insist on the theoretical gap existing between the case of signed measures, where success depends on geometrical

properties of the underlying measure, and the case of positive measures, where many shortcomings disappear. This can be explained in light of the truncated moment problem (Section 1.4), which characterizes the cases of existence and determinacy of sparse *positive* representing measures for a given sequence, and thus provides sufficient conditions for Prony’s method and related approaches to work. This theory is also at the core of Lasserre’s hierarchy, which is discussed in Chapter 2.

General context Our exposition in this chapter focuses on the noiseless setting. Indeed, in the framework exposed in Section 3.3, Prony’s extraction is used on top of a preliminary optimization-based recovery step, which computes “denoised” moments of the absolute value of the sought measure. This first variational step is detailed in Chapter 2. We thus circumvent the signed and noisy settings for which the analysis of Prony is difficult, by employing the method in the most favorable framework, where strong recovery guarantees may be derived. Nonetheless, in order to bring out the important differences between the positive and signed cases, we consider measures of arbitrary signs until Section 1.4.

1.1.2 Previous works

Numerous variants of Prony’s method in the multivariate case exist, and there is no general agreement on a canonical extension, even in 2-D. Nonetheless, approaches might be divided into two main lines of works: on the one hand, solving the multivariate problem by splitting it into several univariate ones, and on the other, via a joint analysis. While Prony’s method only requires $2s$ moments to recover s spikes in 1-D, this linear dependence often does not hold in the multivariate setting, and many existing methods require of the order of $(2s)^d$ samples to work, unless some specific (but maybe unpractical) sampling scheme is used.

Splitting into univariate problems. First attempts to extend the extraction to multi-dimensional signals consisted in solving several single eigenproblem to estimate each component separately, either successively [Sacchini et al., 1993] or independently, with an additional optimization-based matching step [Hua, 1992, Rouquette and Najim, 2001], used to correctly pair the coordinates altogether. In a similar fashion, multivariate extraction may be performed by projecting the data along a finite number of sampling lines and solving the corresponding uni-dimensional problems [Potts and Tasche, 2013, Diederichs and Iske, 2015], again followed by a matching step. This approach has the benefit that much less samples may be needed for the reconstruction, depending on the choice of lines along which the sampling is done. This was further explored in [Cuyt and Lee, 2018], where an astute sampling scheme allows to achieve reconstruction by combining univariate Prony techniques with a few linear systems solving, using only $(d + 1)s$ samples, in theory.

Joint eigenvalue analysis. Prony’s method works by encoding the points to recover as the roots of some polynomial. A more faithful extension of this principle to the multivariate case is to look for the sources as the common roots of a few multivariate polynomials. One is thus left with solving a multivariate polynomial system. While strongly nonlinear, this task may be reduced to a weakly nonlinear eigenvalue problem [Stetter, 1996]. The joint eigenvalue algorithm we discuss has thus been known by researchers in polynomial optimization since the end of the eighties, in particular through the seminal works of Stetter and Möller [Auzinger and Stetter, 1988, Möller, 1993, Möller and Stetter, 1995]. It was later harnessed by Henrion and Lasserre [Henrion and Lasserre, 2005] as an extraction step for Lasserre’s hierarchy, and its usefulness for building explicit solutions of the moment problem has been extensively described since, see e.g. [Laurent, 2010].

In the signal processing literature, the approach to our knowledge first appeared as a multivariate extension of ESPRIT [Haardt and Nosssek, 1998], and its theoretical analysis, in terms of correct pairing and recovery, has been the subject of many works in the last years [Kunis et al., 2016, Harmouch et al., 2017, Sauer, 2017, Ehler et al., 2019]. Contrarily to the univariate case, the theoretical guarantees regarding the behavior and performance of the multivariate algorithm in the generic case involve a minimal separation distance between the points to recover, see e.g. [Sauer, 2017, Remark 2] or [Ehler et al., 2019, Theorem 1]. This bottleneck disappears in the case of positive measures (of prime interest to us) which is explained by results from the truncated trigonometric moment problem, see Section 1.4.

Joint-diagonalization algorithms. Although several variants may be found, the essence of the methods relies on a singular value decomposition of the initial moment matrix to build a few (non Hermitian) co-diagonalizable matrices, called multiplication matrices, whose common spectrum reveals the sought parameters.

The joint diagonalization step is most often performed by diagonalizing a random linear combination of the matrices [Stetter, 1996, Corless et al., 1997, Henrion and Lasserre, 2005, Harmouch et al., 2017, Ehler et al., 2019], which is efficient as long as the eigenspaces of the resulting matrix are all of dimension one. This holds with probability one in the case of sparse measures; on the other hand, the approach is not robust at all when applied on data derived from continuous measures (supported on curves for instance), as we aim to do in the last part of this thesis. To this end, we study in Section 1.6 the use of an optimization-based joint diagonalization step.

Usual co-diagonalization schemes assume that the matrices to diagonalize are Hermitian (or symmetric) [Flury and Gautschi, 1986, Matsuoka et al., 1995, Cardoso and Souloumiac, 1993, Cardoso and Souloumiac, 1996], whereas multiplication matrices are not, and hence cannot be diagonalized in an orthogonal basis. A similar obstacle is tackled in [Haardt and Nosssek, 1998] by rather jointly trigonalizing the matrices using an optimization-based Schur decompo-

sition algorithm. On the other hand, co-diagonalization algorithms for non-symmetric have also been studied in [Pham, 2001, Yeredor, 2002]; however, although they drop the orthogonality assumption, these works still involve transformation of the form TXT^H , and hence can only perform approximate joint diagonalization even if the matrices are indeed co-diagonalizable [Pham, 2001].

Extensions to non-atomic measures. Several alternative to Prony’s algorithm have been proposed to tackle the recovery of more general structures than pointwise sources. Among them, finite rate of innovation techniques play an important role [Vetterli et al., 2002], in particular for the recovery of curves [Chen et al., 2012, Pan et al., 2014, Ongie and Jacob, 2016]. For specific curves featuring a finite number of degrees of freedom, such as polygons [Chen et al., 2012], or the zero locus of a multivariate polynomial [Ongie and Jacob, 2016], these methods achieves off-the-grid recovery using Prony’s annihilating equations on an adequate set of moments (*e.g.* moments of the derivative for piecewise constant functions [Ongie and Jacob, 2016]).

In a different line of work, approximations of non-sparse supports may be computed using the Christoffel function [Lasserre and Pauwels, 2016, Lasserre and Pauwels, 2019, Pauwels et al., 2020]. This function is built directly from a finite number of moments of the measure, and its level sets contain the sought support with high probability. The accuracy of the estimation improves as a greater number of moments are involved [Pauwels et al., 2020]. As it is an important alternative to Prony’s method, we describe it in more details in Section 1.6.4.

On the other hand, there are few studies on the use and limits of sparse recovery techniques for non-atomic measures. Let us only mention the work [Klep et al., 2018], that studies the robustness of the sparse extraction procedure in a non-sparse regime, for positive measures. In particular, the authors are able to describe an “almost sparse” regime in which the recovery works satisfactorily.

1.1.3 Contributions

We state and prove in Section 1.4.2 results on existence and unicity of positive representing measures under some rank assumptions on moment matrices. These are well-known results in the theory of moments, but whose proof can only be found in the case of real or complex polynomial moments in the literature. The trigonometric scenario is similar but still requires some additional details that we provide. In particular, we shed a new light on the equivalence between the three settings.

Section 1.6 introduces the main contribution of this chapter, which is a novel approach of the joint diagonalization step in the multivariate recovery procedure. Section 1.6.2 analyzes the limitations of the diagonalization by random projections in the non-sparse setting. We then introduce our method in Section 1.6.3, which consists in a gradient descent with multiplicative updates to minimize a co-diagonalization energy, thus contrasting with the usual random linear combinations. Casting the recovery in a variational framework improves

significantly the performance when approximating non-sparse measures. Indeed, in this setting, the matrices only approximately commute, which necessitate a more robust numerical scheme. We give numerical illustrations of the method.

1.2 Problem formulation

We consider the problem of identifying a d -dimensional discrete Radon measure $\mu \in \mathcal{M}(\mathbb{T}^d)$ given a finite number of its Fourier coefficients (or trigonometric moments)

$$c_{\mathbf{k}}(\mu) \stackrel{\text{def.}}{=} \int_{\mathbb{T}^d} e^{-2i\pi(\mathbf{k}, \mathbf{x})} d\mu(\mathbf{x}) \quad (1.1)$$

where \mathbf{k} covers the regular domain

$$\Omega_n \stackrel{\text{def.}}{=} \{\mathbf{k} \in \mathbb{Z}^d ; \|\mathbf{k}\|_\infty \leq n\} \subset \mathbb{Z}^d. \quad (1.2)$$

for some $n \in \mathbb{N}$. We denote by $N(n) \stackrel{\text{def.}}{=} (2n+1)^d$, or by N when there is no ambiguity, the number of indices in Ω_n (the dimension d should always be clear from the context). This inverse problem is ill-posed, since the operator

$$\begin{aligned} \mathcal{F}_n : \mathcal{M}(\mathbb{T}^d) &\rightarrow \mathbb{C}^{N(n)} \\ \mu &\mapsto (c_{\mathbf{k}}(\mu))_{\mathbf{k} \in \Omega_n}, \end{aligned} \quad (1.3)$$

mapping an infinite-dimensional space onto $\mathbb{C}^{N(n)}$, is not injective. Given a truncated moment sequence $\mathbf{c} \in \text{Im } \mathcal{F}_n$, any Radon measure ν whose moments match \mathbf{c} is called a representing measure for \mathbf{c} .

Remark 1. All the results of this Chapter can be readily extended to the case where the frequency domain is of the form $\llbracket -n_1, n_1 \rrbracket \times \dots \times \llbracket -n_d, n_d \rrbracket$ (instead of $\llbracket -n, n \rrbracket^d$), simply at the cost of more heavy notations. For clarity, we hold to the conventional framework.

Monomial orderings When considering multi-indexed sequences, which we signal by a bold subscript, the d -tuples are assumed to be ordered following the colexicographical ordering, unless otherwise specified.

Definition 1 (Colexicographical ordering). *Given d -tuples $\mathbf{i}, \mathbf{j} \in \mathbb{Z}^d$, one has*

$$\begin{aligned} (\mathbf{i}_1, \dots, \mathbf{i}_d) \preceq_{\text{colex}} (\mathbf{j}_1, \dots, \mathbf{j}_d) \\ \iff \\ (\mathbf{i}_d < \mathbf{j}_d) \text{ or } ((\mathbf{i}_d = \mathbf{j}_d) \text{ and } (\mathbf{i}_1, \dots, \mathbf{i}_{d-1}) \preceq_{\text{colex}} (\mathbf{j}_1, \dots, \mathbf{j}_{d-1})). \end{aligned}$$

The choice of an ordering has important implications: it determines for example the leading term of a polynomial, as well as its vectorial representation.

Polynomials Sparse recovery problems are intimately related to polynomial interpolation problems [Sauer, 2018, Poon and Peyré, 2018]. For instance, Prony’s method works by encoding the support of the sought measure as the roots of a well-chosen trigonometric polynomial.

A Laurent polynomial is defined as

$$P(\mathbf{z}) = \sum_{\mathbf{k} \in \Omega} p_{\mathbf{k}} \mathbf{z}^{\mathbf{k}}$$

where Ω is a finite subset of \mathbb{Z}^d . If $\Omega \subset \mathbb{N}^d$, we say that P is causal. We define the degree of P as its maximum degree, *i.e.*

$$\deg(P) = \deg_{\infty}(P) \stackrel{\text{def.}}{=} \max \{ \|\mathbf{k}\|_{\infty} ; p_{\mathbf{k}} \neq 0, \mathbf{k} \in \Omega \}$$

We denote by $\mathbb{C}[\mathbf{z}, \mathbf{z}^{-1}]$ the ring of Laurent polynomials, and by $\mathbb{C}[\mathbf{z}]$ the ring of causal polynomials. For $n \in \mathbb{N}$, the Laurent polynomials (resp. causal polynomials) of degree at most n are denoted by $\mathbb{C}_n[\mathbf{z}, \mathbf{z}^{-1}]$ (resp. $\mathbb{C}_n[\mathbf{z}]$). When working with finite-degree causal polynomials, it is natural to consider the domain

$$\Omega_n^+ \stackrel{\text{def.}}{=} \{ \mathbf{k} \in \mathbb{N}^d ; \|\mathbf{k}\|_{\infty} \leq n \} \subset \mathbb{N}^d, \quad (1.4)$$

of cardinality $N_+(n) \stackrel{\text{def.}}{=} (n+1)^d$. The set $\mathbb{C}_n[\mathbf{z}, \mathbf{z}^{-1}]$ (resp. $\mathbb{C}_n[\mathbf{z}]$) may be identified with the vector space $\mathbb{C}^{N(n)}$ (resp. $\mathbb{C}^{N_+(n)}$).

Example 1. With respect to the colexicographical order, the leading term of the 2-variate (causal) polynomial $P(\mathbf{z}) = a - bz_2^2 + cz_1z_2 - dz_1^2z_2 \in \mathbb{C}[\mathbf{z}]$ is $(-bz_2^2)$, and its coefficient vector in $\mathbb{C}_2[\mathbf{z}]$ is $p = (a \ 0 \ 0 \ 0 \ c \ -d \ -b \ 0 \ 0) \in \mathbb{C}^9$.

Given $P \in \mathbb{C}[\mathbf{z}, \mathbf{z}^{-1}]$ (or $\mathbb{C}[\mathbf{z}]$), we denote by $P^{\mathbb{T}}$ its trigonometric analogue, *i.e.*

$$\forall \mathbf{x} \in \mathbb{T}, \quad P^{\mathbb{T}}(\mathbf{x}) \stackrel{\text{def.}}{=} P(e^{2i\pi\mathbf{x}})$$

Note that if P is Hermitian, *i.e.* $p_{-\mathbf{k}} = \overline{p_{\mathbf{k}}}$, then $P^{\mathbb{T}}$ is real-valued.

Vandermonde matrices They play a crucial role in problems of moments and polynomial interpolation.

Definition 2. Given finite sets $\mathbf{X} \subset \mathbb{T}^d$ and $\Omega \subset \mathbb{Z}^d$, the Vandermonde matrix $V(\mathbf{X}, \Omega)$ is the matrix

$$V(\mathbf{X}, \Omega) \stackrel{\text{def.}}{=} \left(e^{-2i\pi \langle \mathbf{k}, \mathbf{x} \rangle} \right)_{\mathbf{k} \in \Omega, \mathbf{x} \in \mathbf{X}}, \quad (1.5)$$

whose rows are ordered following the colexicographical order.

For a finite set of points $\{\mathbf{x}_1, \dots, \mathbf{x}_s\} \subset \mathbb{T}^d$, the Vandermonde matrix

$$V_n(\mathbf{x}) \stackrel{\text{def.}}{=} V(\mathbf{x}, \Omega_n^+) \in \mathcal{M}_{N_+, s}(\mathbb{C}) \quad (1.6)$$

plays an important role in Prony's problem. For simplicity, we sometimes denote it by V_n when there is no ambiguity. As before, we also denote by $v_n(\mathbf{x}_j)$ (or simply $v_{n,j}$) the j -th column of this matrix, *i.e.*

$$v_n(\mathbf{x}_j) \stackrel{\text{def.}}{=} (e^{-2i\pi\langle \mathbf{k}, \mathbf{x}_j \rangle})_{\mathbf{k} \in \Omega_n^+} \in \mathbb{C}^{N_n^+}.$$

Similarly, V_n^\pm (resp. $v_{n,j}^\pm$) designates the matrix of $\mathcal{M}_{N,s}(\mathbb{C})$ (resp. vector of \mathbb{C}^N) whose rows are indexed by Ω_n , instead of Ω_n^+ .

Moment matrices Radon measures can be alternatively represented by (infinite) moment sequences, which may themselves be seen as (infinite) moment matrices. But truncated moment matrices sometimes contain enough information to fully recover the underlying measure. In this chapter, we review how the spectral properties of such matrices are connected to the structure of representing measures. Finite moment matrices also plays a central role in Chapter 2, as optimization variables.

Given $\mathbf{c} \in \text{Im } \mathcal{F}_n$, and μ a representing measure, the matrix

$$T_n(\mathbf{c}) \stackrel{\text{def.}}{=} (c_{\mathbf{k}-\mathbf{l}})_{\mathbf{k}, \mathbf{l} \in \Omega_n^+} \in \mathcal{M}_{N_+(n)}(\mathbb{C}), \quad (1.7)$$

is called moment matrix of μ . Again, its rows and columns are assumed to be sorted following the colexicographical order. When $d = 1$, such matrices are said to be Toeplitz, meaning that they are constant along their diagonals. We sometimes refer to the operator $T_n : \mathbb{C}^N \rightarrow \mathcal{M}_{N_+}(\mathbb{C})$ as the Toeplitz operator.

Example 2. In 1-d, the moment matrix $T_3(c)$ for instance reads

$$\begin{pmatrix} c_0 & c_{-1} & c_{-2} & c_{-3} \\ c_1 & c_0 & c_{-1} & c_{-2} \\ c_2 & c_1 & c_0 & c_{-1} \\ c_3 & c_2 & c_1 & c_0 \end{pmatrix},$$

and is Toeplitz.

In 2-d, the moment matrix $T_2(\mathbf{c})$ reads

$$\begin{pmatrix} c_{0,0} & c_{1,0} & c_{2,0} & | & c_{0,-1} & c_{-1,-1} & c_{-2,-1} & | & c_{0,-2} & c_{1,-2} & c_{2,-2} \\ c_{1,0} & c_{0,0} & c_{1,0} & | & c_{1,-1} & c_{0,-1} & c_{-1,-1} & | & c_{1,-2} & c_{0,-2} & c_{1,-2} \\ c_{2,0} & c_{1,0} & c_{0,0} & | & c_{2,-1} & c_{1,-1} & c_{0,-1} & | & c_{2,-2} & c_{1,-2} & c_{0,-2} \\ \hline c_{0,1} & c_{-1,1} & c_{-2,1} & | & c_{0,0} & c_{-1,0} & c_{-2,0} & | & c_{0,-1} & c_{-1,-1} & c_{-2,-1} \\ c_{1,1} & c_{0,1} & c_{-1,1} & | & c_{1,0} & c_{0,0} & c_{-1,0} & | & c_{1,-1} & c_{0,-1} & c_{-1,-1} \\ c_{2,1} & c_{1,1} & c_{0,1} & | & c_{2,0} & c_{1,0} & c_{0,0} & | & c_{2,-1} & c_{1,-1} & c_{0,-1} \\ \hline c_{0,2} & c_{-1,2} & c_{-2,2} & | & c_{0,1} & c_{-1,1} & c_{-2,1} & | & c_{0,0} & c_{-1,0} & c_{-2,0} \\ c_{1,2} & c_{0,2} & c_{-1,2} & | & c_{1,1} & c_{0,1} & c_{-1,1} & | & c_{1,0} & c_{0,0} & c_{-1,0} \\ c_{2,2} & c_{1,2} & c_{0,2} & | & c_{2,1} & c_{1,1} & c_{0,1} & | & c_{2,0} & c_{1,0} & c_{0,0} \end{pmatrix}.$$

Note that here the Toeplitz-block Toeplitz structure is due to the choice of the colexicographical order on the duplets. More generally, moment matrices in dimension d have a d -level Toeplitz structure, for the colexicographical order.

A moment matrix $T_n(\mathbf{c})$ defines a bilinear form $\langle \cdot, \cdot \rangle_{\mathbf{c}}$ over \mathbb{C}^{N_+} as

$$\forall \mathbf{p}, \mathbf{q} \in \mathbb{C}^{N_+}, \quad \langle \mathbf{p}, \mathbf{q} \rangle_{\mathbf{c}} \stackrel{\text{def.}}{=} \langle \mathbf{p}, T_n(\mathbf{c})\mathbf{q} \rangle = \mathbf{p}^H T_n(\mathbf{c})\mathbf{q},$$

or equivalently over $\mathbb{C}_n[\mathbf{z}]$ as

$$\forall P, Q \in \mathbb{C}_{N_+}[\mathbf{z}], \quad \langle P, Q \rangle_{\mathbf{c}} \stackrel{\text{def.}}{=} \int_{\mathbb{T}^d} \overline{P^{\mathbb{T}}(\mathbf{x})} Q^{\mathbb{T}}(\mathbf{x}) d\mu(\mathbf{x}).$$

In particular, one has $c_{\mathbf{k}-\mathbf{l}} = \langle \mathbf{z}^{\mathbf{k}}, \mathbf{z}^{\mathbf{l}} \rangle_{\mathbf{c}}$.

Discrete representing measures If a sequence $\mathbf{c} \in \text{Im } \mathcal{F}_n$ admits a discrete representing measure of the form

$$\mu_0 \stackrel{\text{def.}}{=} \sum_{j=1}^s a_j \delta_{\mathbf{x}_j} \in \mathcal{M}(\mathbb{T}^d),$$

with $\mathbf{a} \in (\mathbb{C}_*)^r$ and $\{\mathbf{x}_1, \dots, \mathbf{x}_s\} \subset \mathbb{T}^d$, all distinct, equation (1.1) then simply becomes

$$c_{\mathbf{k}} = \sum_{j=1}^s a_j e^{-2i\pi \langle \mathbf{k}, \mathbf{x}_j \rangle}, \quad \forall \mathbf{k} \in \Omega_n \quad (1.8)$$

or, in matrix form,

$$\mathbf{c} = V_n^{\pm}(\mathbf{x})\mathbf{a}. \quad (1.9)$$

In that case, its moment matrix may be factorized as

$$T_n(\mathbf{c}) = V_n(\mathbf{x}) \text{Diag}(\mathbf{a}) (V_n(\mathbf{x}))^H = \sum_{j=1}^s a_j v_n(\mathbf{x}_j) (v_n(\mathbf{x}_j))^H. \quad (1.10)$$

Thus, one may see the problem of finding a discrete representing measure for a moment sequence $\mathbf{c} \in \text{Im } \mathcal{F}_n$ as the problem of finding a Vandermonde decomposition of $T_n(\mathbf{c})$. Once the Vandermonde matrix $V_n(\mathbf{x})$ is known, and hence the support \mathbf{x} , it is easy to recover $V_n^{\pm}(\mathbf{x})$, and from there the vector of amplitudes \mathbf{a} by solving the system (1.9).

1.3 A multivariate extension of Prony

1.3.1 The univariate case

We recall in this section the main ideas behind Prony's method, which allows one to recover a s -sparse measure $\mu_0 \stackrel{\text{def.}}{=} \sum_{j=1}^s a_j \delta_{x_j} \in \mathcal{M}(\mathbb{T})$ (where $a_j \in \mathbb{C}^*$ and $x_j \in \mathbb{T}$, for all j) from the sole knowledge of at least $(2s+1)$ moments.

For $n \in \mathbb{N}$, let $\mathbf{c} = (c_k)_{k \in \Omega_n}$ be the vector of Fourier coefficients of μ_0 , given by (1.8) (with $d=1$). As noted in [Condat, 2020, Remark 3], there is an infinity of measures supported on at most $2n$ points that represent the sequence \mathbf{c} . Indeed, for any set of $(2n)$ points $\{\omega_1, \dots, \omega_{2n}\} \subset \mathbb{T}$, all distinct,

the Vandermonde matrix $V_n^\pm(\omega)$ is invertible, and the linear system (1.9) admits a solution a_ω , so that the discrete measure $\sum (a_\omega)_j \delta_{\omega_j}$ explains the sequence \mathbf{c} . On the other hand, only one representing measure may be supported on less or as many as n points [Condat, 2020, Remark 2], and that is the one retrieved by Prony's method, assuming one has enough measurements (*i.e.* $n \geq s$).

Prony's method In a paper of 1795 [R. de Prony, 1795], Gaspard R. de Prony, a French engineer, addressed the following question: how can one determine a positive integer $s \in \mathbb{N}^*$, distinct parameters $\chi_{0,j} \in]-\infty, 0] + i[0, 2\pi[$, and complex coefficients $a_j \in \mathbb{C}^*$ ($j = 1, \dots, r$) explaining the process

$$f(t) = \sum_{j=1}^s a_j e^{-\chi_j t}$$

given only sampled data $c_k = f(k)$, for $k = -n, \dots, n$, with $n \geq s$? Restricting the problem to the case where the parameters χ_j are pure imaginary $\chi_j = 2i\pi x_j$, we fall back on our primary problem of recovering the Radon measure $\sum_{j=1}^s a_j \delta_{x_j}$ from a few trigonometric moments.

Let $\zeta_j \stackrel{\text{def.}}{=} e^{2i\pi x_j}$ for $j = 1, \dots, s$, and $P_0 \in \mathbb{C}[z]$ be defined as

$$P_0(z) = \prod_{j=1}^s (z - \zeta_j) \stackrel{\text{def.}}{=} z^s + \sum_{k=0}^{s-1} p_{0,k} z^k, \quad \forall z \in \mathbb{C}$$

Prony's method exploits the fact that the coefficients of P_0 are solution of a simple linear system.

Lemma 1. *The coefficients of P_0 satisfy*

$$\sum_{k=0}^s p_{0,k} c_{l-k} = 0 \tag{1.11}$$

for all $-n + s \leq l \leq n$.

Proof. We have

$$\sum_{k=0}^s p_{0,k} c_{l-k} = \sum_{j=1}^s a_j e^{-2i\pi l x_j} \sum_{k=0}^s p_{0,k} e^{2i\pi k x_j} = \sum_{j=1}^s a_j e^{-2i\pi l x_j} P_0(\zeta_j) = 0$$

□

Defining the rectangular matrix

$$R_s(\mathbf{c}) \stackrel{\text{def.}}{=} \begin{pmatrix} c_{-n+s} & \dots & c_{-n} \\ \vdots & & \vdots \\ c_n & \dots & c_{n-s} \end{pmatrix} \in \mathcal{M}_{2n-s+1, s+1}(\mathbb{C})$$




Figure 1.1: Modulus of Prony's polynomial, over the trigonometric circle ($|P_0^\mathbb{T}|$, left), or the complex plane ($|P_0|$, right, in logarithmic scale)

the equations (1.11) show that p_0 solves the linear system

$$R_s(\mathbf{c})p = 0. \quad (1.12)$$

In fact, $R_s(\mathbf{c})$ is of rank s [Potts and Tasche, 2011, Lemma 2.1], and \mathbf{p}_0 spans its kernel. In practice, the number of sources s , which is needed to build the matrix $R_s(\mathbf{c})$, is often unknown. However, having access to the noiseless moments, one may retrieve it as the rank of the moment matrix $T_n(\mathbf{c})$ (1.7).

Lemma 2. *If \mathbf{c} satisfies (1.8) with $d = 1$, then $\text{rank } T_n(\mathbf{c}) = s$, for all $n \geq s$.*

Proof. This follows from the Vandermonde decomposition (1.10) of $T_n(\mathbf{c})$. \square

Now reasoning in reverse, given the moment matrix $T_n(\mathbf{c})$, one may thus build the linear system (1.12), and from there retrieve the polynomial P_0 . In the literature, P_0 is sometimes called Prony's polynomial, or annihilating polynomial. It is displayed in Figure 1.1.

Remark 2. Lemma 2 does not hold in the multivariate case. Given a (truncated) moment sequence coming from a s -sparse measure over \mathbb{T}^d , $d > 1$, one may have $\text{rank } T_n(\mathbf{c}) < s$, even if $n \geq s$. For this reason, we make a clear distinction in our notations between the rank of moment matrices, denoted by r , and the sparsity of underlying measures, denoted by s .

Remark 3. Note that $\text{rank } T_n(\mathbf{c}) = r$ does not necessarily imply that $\text{rank } T_r(\mathbf{c}) = r$, which is why we consider the matrix $R_s(\mathbf{c})$ rather than $T_s(\mathbf{c})$.

Remark 4. Given $\mathbf{c} \in \mathbb{C}^{2n+1}$, let $\Delta_n(\mathbf{c}) \stackrel{\text{def.}}{=} \det T_n(\mathbf{c})$. If $\Delta_{s-1}(\mathbf{c}) \neq 0$, the polynomial

$$W_s(z) \stackrel{\text{def.}}{=} \frac{1}{\Delta_{s-1}(\mathbf{c})} \begin{vmatrix} c_0 & c_{-1} & \dots & c_{-s} \\ c_1 & c_0 & \dots & c_{-s+1} \\ \vdots & \vdots & & \vdots \\ c_{s-1} & c_{s-2} & \dots & c_{-1} \\ 1 & z & \dots & z^s \end{vmatrix} = z^s + \dots$$

is well-defined, and is, in fact, Prony’s polynomial, as one can see by applying Cramer’s rule to the system (1.12). The polynomials W_n play an important role in the trigonometric moment problem [Carathéodory, 1911, Szego, 1954], and the behavior of their roots is well understood when the sequence \mathbf{c} is positive, that is when $T_k(\mathbf{c}) \succ 0$ for $k = 0, \dots, s$ [Akhiezer and Krein, 1938]. There are also related to Christoffel polynomials, see section 1.6.4.

The support of the sought measure is thus encoded in the roots of the annihilating polynomial P_0 , computed from the sequence \mathbf{c} . Finding the roots of P_0 is typically performed by diagonalizing its companion matrix,

$$C_0 \stackrel{\text{def.}}{=} \begin{pmatrix} 0 & -p_{0,0} \\ I_s & \vdots \\ & -p_{0,s} \end{pmatrix} \in \mathcal{M}_{r+1}(\mathbb{C}). \quad (1.13)$$

The amplitudes are then obtained by plugging the result into the harmonic model (1.8), and solving the linear (Vandermonde) system (1.9). Algorithm 1 summarizes the main steps of Prony’s method.

Algorithm 1: 1-D Prony’s method

Input: $c_k, k = -n, \dots, n$

- 1 with $r = \text{rank } T_n(\mathbf{c})$, find $\mathbf{p}_0 \in \mathbb{C}^{r+1}$ such that $R_r(\mathbf{c})\mathbf{p}_0 = 0$
- 2 compute the eigenvalues $\{\zeta_1, \dots, \zeta_r\}$ of companion matrix C_0
- 3 with $x_j = \arg \zeta_j / 2\pi$, compute $\{a_1, \dots, a_r\}$ as a solution to (1.9)

Output: $r, \{x_j\}, \{a_j\}$

1.3.2 Multivariate recovery via spectral analysis

We now consider measures over the d -dimensional torus \mathbb{T}^d , for $d \in \mathbb{N}^*$. Let $s \in \mathbb{N}^*$, and $\mu_0 \stackrel{\text{def.}}{=} \sum_{j=1}^s a_j \delta_{\mathbf{x}_j} \in \mathcal{M}(\mathbb{T}^d)$ be s -sparse. As before, we assume that we are given a vector $\mathbf{c} = (c_{\mathbf{k}})_{\mathbf{k} \in \Omega_n}$ of moments of μ_0 , *i.e.* of the form (1.8).

In one dimension, Prony’s method extracts a polynomial from the kernel of the Toeplitz matrix $T_n(\mathbf{c})$, whose roots determine the support of the discrete representing measure. However, zero sets of non trivial polynomials in several variables are rarely discrete. Multivariate generalizations of Prony’s method rather look for the support of the measure as the intersection of zero loci of a few annihilating polynomials. Similarly to the uni-dimensional case, these polynomials lie in the kernel of the moment matrix $T_n(\mathbf{c})$ (1.7), or more precisely in the ideal generated by $\text{Ker } T_n(\mathbf{c})$.

Kernel of moment matrices. To simplify notations, we use the same letter in this section to denote a polynomial or its vector of coefficients. The meaning should always be clear from the context.

Definition 3. Given a set $E \subset \mathbb{C}[\mathbf{z}]$ of polynomials, the ideal of $\mathbb{C}[\mathbf{z}]$ generated by E , which we denote by $\langle E \rangle$, is

$$\langle E \rangle \stackrel{\text{def.}}{=} \left\{ \sum_{j \in I} p_j q_j ; I \subset \mathbb{N}^d \text{ finite, } p_j \in E, \quad q_j \in \mathbb{C}[\mathbf{z}] \right\} \quad (1.14)$$

Definition 4. Given an ideal \mathcal{I} of $\mathbb{C}[\mathbf{z}]$, the algebraic variety

$$\mathcal{Z}(\mathcal{I}) \stackrel{\text{def.}}{=} \{ \mathbf{z} \in \mathbb{C}^d ; p(\mathbf{z}) = 0, \quad \forall p \in \mathcal{I} \} \quad (1.15)$$

is called the variety of \mathcal{I} .

Let $\zeta_j \stackrel{\text{def.}}{=} e^{-2i\pi \mathbf{x}_j}$, $j = 1, \dots, s$, where $\mathbf{x} \in (\mathbb{T}^d)^s$ is the support of the measure explaining \mathbf{c} . Analogously to the uni-dimensional case, there is a link between the variety of $\langle \text{Ker } T_n(\mathbf{c}) \rangle$ and the points ζ_j . Theorem 3.1 in [Kunis et al., 2016] states that the two sets are equal for any order $n \geq s$. Although generic, this last condition may pose difficulties, as it means that at least $(2s)^d$ observations should be needed to recover s spikes in dimension d . It can be improved, at the cost of introducing a minimal separation distance.

Theorem 1 ([Kunis et al., 2018, Corollary 2.10]). Let μ_0 be a s -sparse measure with minimal separation distance Δ , and let $\mathbf{c} \in \mathbb{C}^{N(n)}$ be its truncated moment sequence (1.8). If

$$n > 1 + \frac{3 + 2 \log d}{\Delta}, \quad (1.16)$$

then

$$\mathcal{Z}(\langle \text{Ker } T_n(\mathbf{c}) \rangle) = \{ \zeta_1, \dots, \zeta_s \}. \quad (1.17)$$

An ideal whose variety is a finite set is said to be zero-dimensional. Given that (1.17) holds, we are thus left with the problem of computing the variety of a zero-dimensional ideal, which is a well-known problem in algebraic geometry. We discuss the case of infinite varieties in section 1.6.

Since any ideal of $\mathbb{C}[\mathbf{z}]$ admits a finite set of generators, finding the complex exponentials ζ_j ultimately amounts to solve a system of polynomial equations. Practical ways of solving such systems are numerous and have been extensively studied in the algebraic geometry literature [Sturmfels, 2002]. We present the Stetter-Möller matrix method [Möller and Stetter, 1995], which connects the exponentials ζ_j to the spectrum of so-called *multiplication operators* associated to the quotient space $\mathbb{C}[\mathbf{z}]/\langle \text{Ker } T_n \rangle$. This is also known as the Gelfand-Naimark-Segal construction [Klep et al., 2018]. This method does not require to know the polynomials generating $\langle \text{Ker } T_n \rangle$. In the following, and until section 1.6, we always assume that (1.17) holds.

The quotient space \mathcal{T}_n . We consider the quotient space

$$\mathcal{T}_n \stackrel{\text{def.}}{=} \mathbb{C}[\mathbf{z}]/\langle \text{Ker } T_n(\mathbf{c}) \rangle,$$

which is a \mathbb{C} -vector space with addition $[f] + [h] = [f + h]$ and scalar multiplication $\lambda[f] = [\lambda f]$, and an algebra with multiplication $[f][h] = [fh]$, for $f, h \in \mathbb{C}[\mathbf{z}]$ and $\lambda \in \mathbb{C}$. Its dual space \mathcal{T}_n^* may be identified with the set of linear forms over $\mathbb{C}[\mathbf{z}]$ that vanish on $\langle \text{Ker } T_n(\mathbf{c}) \rangle$.

Let $\tau_n \stackrel{\text{def.}}{=} \dim(\mathcal{T}_n)$ and $\sigma_n \stackrel{\text{def.}}{=} \#\mathcal{Z}(\langle \text{Ker } T_n(\mathbf{c}) \rangle)$. We have the following important result.

Theorem 2 ([Elkadi and Mourrain, 2007, Theorem 4.3]). *For any $n \in \mathbb{N}$, one has $\sigma_n \leq \tau_n$, and $\tau_n < \infty$ if and only if $\sigma_n < \infty$. Furthermore, if $T_n(\mathbf{c}) \succeq 0$, then $\tau_n = \sigma_n$.*

Thus in our case, assuming that (1.17) holds, the quotient space \mathcal{T}_n is finite-dimensional. However, note that its dimension may be strictly greater than the number of points in the variety of $\langle \text{Ker } T_n(\mathbf{c}) \rangle$.

Multiplication operators For $p \in \mathbb{C}[\mathbf{z}]$, the *multiplication operator* \mathcal{X}_p is defined over \mathcal{T}_n as

$$\begin{aligned} \mathcal{X}_p : \mathcal{T}_n &\rightarrow \mathcal{T}_n \\ [h] &\mapsto [hp] \end{aligned} \quad (1.18)$$

Let $\mathcal{X}_{(i)}$ denote the multiplication by z_i , which plays a central role in the following. The operators $\mathcal{X}_{(i)}$ commute pairwise, and for a polynomial $p \in \mathbb{C}[\mathbf{z}]$, one has

$$\mathcal{X}_p = \sum_{\mathbf{k} \in \mathbb{Z}^d} p_{\mathbf{k}} \mathcal{X}_{(1)}^{k_1} \dots \mathcal{X}_{(d)}^{k_d} = p(\mathcal{X}_{(1)}, \dots, \mathcal{X}_{(d)}).$$

Note also that $\mathcal{X}_p \equiv 0$ if and only if $p \in \langle \text{Ker } T_n \rangle$.

The adjoint of the multiplication operator reads

$$\begin{aligned} \mathcal{X}_p^* : \mathcal{T}_n^* &\rightarrow \mathcal{T}_n^* \\ \Lambda &\mapsto \Lambda \circ \mathcal{X}_p \end{aligned}$$

Spectrum of multiplication operators For $\zeta \in \mathbb{C}^d$, the linear form

$$e_{\zeta} : p \in \mathbb{C}[\mathbf{z}] \mapsto p(\zeta)$$

is called evaluation at ζ . If $\zeta \in \mathcal{Z}(\langle \text{Ker } T_n \rangle)$, then $e_{\zeta} \in \mathcal{T}_n^*$.

Theorem 3 ([Elkadi and Mourrain, 2007, Theorem 4.23]). *The eigenvalues of $\mathcal{X}_{(i)}$ are $\zeta_{1,i}, \dots, \zeta_{s,i}$, i.e. the i -th coordinates of the points ζ_j in the variety $\mathcal{Z}(\langle \text{Ker } T_n \rangle)$. In particular, one has*

$$\forall j = 1, \dots, r, \quad \mathcal{X}_{(i)}^*(e_{\zeta_j}) = \zeta_{j,i} e_{\zeta_j}, \quad (1.19)$$

so that the evaluations $e_{\zeta_1}, \dots, e_{\zeta_r}$ are eigenvectors common to all $\mathcal{X}_{(i)}^*$.

Proof. For any $h \in \mathbb{C}[\mathbf{z}]$, one has

$$\mathcal{X}_{(i)}^*(e_{\zeta_j})([h]) = e_{\zeta_j}([z_i h]) = \zeta_{j,i} h(\zeta_j) = \zeta_{j,i} e_{\zeta_j}([h])$$

and hence

$$\mathcal{X}_{(i)}^*(e_{\zeta_j}) = \zeta_{j,i} e_{\zeta_j}.$$

Thus, the $\zeta_{j,i}$ are eigenvalues of $\mathcal{X}_{(i)}$, and the evaluations e_{ζ_j} are eigenvectors common to all $\mathcal{X}_{(i)}^*$. It remains to prove that $\mathcal{X}_{(i)}$ has no other eigenvalues. The polynomial

$$p_i(\mathbf{z}) = \prod_j (z_i - \zeta_{j,i})$$

vanishes over $\mathcal{Z}(\langle \text{Ker } T_n \rangle)$, hence, by Hilbert's Nullstellensatz, it belongs to the radical of $\langle \text{Ker } T_n \rangle$, *i.e.* there exists $k \in \mathbb{N}$ such that $p_i^k \in \langle \text{Ker } T_n \rangle$. Thus,

$$\mathcal{X}_{p_i^k} \equiv 0 \equiv p_i^k(\mathcal{X}_{(1)}, \dots, \mathcal{X}_{(d)}) \equiv \prod_j (\mathcal{X}_{(i)} - \zeta_{j,i})^k$$

which implies that the minimal polynomial of $\mathcal{X}_{(i)}$ divides the (univariate) polynomial $\prod (z - \zeta_{j,i})^k$. Hence, any eigenvalue of $\mathcal{X}_{(i)}$ is of the form $\zeta_{j,i}$. \square

The coordinates of each point in the measure may thus be retrieved as eigenvalues of multiplication operators, and the shared eigenstructure reveals their correct pairing. Thus, we may recover the parameters ζ_j via simple linear algebra operations, assuming that we know how to compute the matrices of multiplication operators. This first step remains difficult in the fully general case, in particular when the dimension of \mathcal{T}_n is not equal to the cardinality of the variety $\mathcal{Z}(\langle \text{Ker } T_n \rangle)$. Building multiplication matrices requires to know some basis of the space \mathcal{T}_n , and though it is always possible to compute such a basis by resorting to the theory of Gröbner bases of ideals, see e.g. [Faugère et al., 1993, Corless, 1996], such approaches are far out of the scope of this thesis. To develop the theory as far as we can while keeping the exposition simple, and give an intuitive overview of the mechanics of the extraction algorithm (Algorithm 2), we first consider the simpler case where one has access to infinite moment matrices.

1.3.3 The simpler theory of infinite matrices

We explain how to build multiplication matrices when one has access to the infinite sequence of moments. This section mostly follows the exposition of [Harmouch et al., 2017]. Given an infinite moment sequence $\mathbf{c} \in \mathbb{C}^{\mathbb{Z}}$, let $L_{\mathbf{c}}$ be the linear functional over $\mathbb{C}[\mathbf{z}, \mathbf{z}^{-1}]$

$$L_{\mathbf{c}} : \mathbf{p} = \sum_{\mathbf{k} \in \mathbb{Z}^d} p_{\mathbf{k}} \mathbf{z}^{\mathbf{k}} \mapsto \sum p_{\mathbf{k}} c_{\mathbf{k}}$$

which is known as the Riesz functional. The infinite moment matrix associated to \mathbf{c} , denoted by $T_{\infty}(\mathbf{c})$, defines as before a bilinear form over $\mathbb{C}[\mathbf{z}]$ (or $\mathbb{C}^{(\mathbb{Z}_+)^d}$) as

$$\forall \mathbf{p}, \mathbf{q} \in \mathbb{C}[\mathbf{z}], \quad \langle \mathbf{p}, \mathbf{q} \rangle_{\mathbf{c}} = \langle \mathbf{p}, T_{\infty}(\mathbf{c})\mathbf{q} \rangle = L_{\mathbf{c}}(\overline{\mathbf{p}}\mathbf{q}).$$

We denote by \mathcal{T}_{∞} the quotient space $\mathbb{C}[\mathbf{z}] / \text{Ker } T_{\infty}$.

Lemma 3. *The set $\text{Ker } T_\infty \stackrel{\text{def.}}{=} \{q \in \mathbb{C}[\mathbf{z}] ; \forall p \in \mathbb{C}[\mathbf{z}], \langle p, T_\infty(\mathbf{c})q \rangle = 0\}$ is an ideal of $\mathbb{C}[\mathbf{z}]$. Moreover, if $\text{rank } T_\infty = r < \infty$, then $\dim \mathcal{T}_\infty = r$.*

Proof. Let $q \in \text{Ker } T_\infty$ and $q' \in \mathbb{C}[\mathbf{z}]$. Then, for all $p \in \mathbb{C}[\mathbf{z}]$, $\langle p, T_\infty qq' \rangle = L_{\mathbf{c}}(\bar{p}qq') = L_{\mathbf{c}}(\bar{p}q'q) = \langle \bar{p}q', T_\infty q \rangle = 0$, so that $qq' \in \text{Ker } T_\infty$. From the first isomorphism theorem, $\mathbb{C}[\mathbf{z}]/\text{Ker } T_\infty$ and $\text{Im } T_\infty$ are isomorphic, so that $\dim \mathcal{T}_\infty = \text{rank } T_\infty$. \square

In the following, we assume that $T_\infty(\mathbf{c})$ is of finite rank s – and hence $\dim \mathcal{T}_\infty = s$. The next lemma gives a simple criterion for linear independency in \mathcal{T}_∞ .

Lemma 4 ([Harmouch et al., 2017, Lemma 3.4]). *Let T_∞ be an infinite moment matrix, such that $\text{rank } T_\infty = s < \infty$. Let $u = \{u_1, \dots, u_k\} \subset \mathbb{C}[\mathbf{z}]$ and $v = \{v_1, \dots, v_k\} \subset \mathbb{C}[\mathbf{z}]$. If the (finite) matrix $M = (\langle u_i, T_\infty v_j \rangle)_{i,j}$ is invertible, then $[u]$ and $[v]$ are linearly independent in \mathcal{T}_∞ .*

Proof. Assume that $q = \sum \lambda_j v_j \in \text{Ker } T_\infty$. Then for any $p \in \mathbb{C}[\mathbf{z}]$, $L_{\mathbf{c}}(\bar{p}q) = 0$, so in particular $\sum \lambda_j \langle u_i, v_j \rangle_{\mathbf{c}} = 0$ for any $i \in \llbracket 1, s \rrbracket$, which reads in matrix form $M\lambda = 0$, and hence $\lambda = 0$. Therefore v is linearly independent. The same reasoning may be applied for u . \square

The next result gives an explicit formula for multiplication matrices.

Proposition 1. *Assume that $\text{rank } T_\infty(\mathbf{c}) = s$, and let T_m be a submatrix of $T_\infty(\mathbf{c})$, such that $\text{rank } T_m = s$. Let $T_m(\mathbf{c}) = U\Sigma V^H$ be the reduced singular value decomposition of $T_m(\mathbf{c})$, and let us denote by $u = (u_1, \dots, u_s)$ and $v = (v_1, \dots, v_s)$ the columns of U and V respectively. Then $[v]$ is a basis of \mathcal{T}_∞ , in which the matrix of $\mathcal{X}_{(i)}$ reads*

$$X_{(i)} = \Sigma^{-1} U^H T_m^{(i)} V \in \mathbb{C}^{s \times s} \quad (1.20)$$

where $T_m^{(i)}$ is the shifted matrix

$$T_m^{(i)} = (c_{\mathbf{k}-\mathbf{l}+\delta_i})_{\mathbf{k}, \mathbf{l} \in \Omega_m^+}. \quad (1.21)$$

Proof. For $1 \leq i, j \leq r$, one has

$$\langle u_i, T_\infty v_j \rangle = u_i^H T_n v_j = \Sigma_{i,j}.$$

Since Σ is invertible, it follows from Lemma 4 that $[v]$ (resp. $[u]$) is linearly independent in \mathcal{T}_∞ . As $\dim \mathcal{T}_\infty = r$, $[v]$ is a basis of \mathcal{T}_∞ . Let $X_{(i)} = (x_{j,k}^{(i)})$ be the matrix of $\mathcal{X}_{(i)}$ in this basis. The k -th column of $X_{(i)}$ contains the residue of the polynomial $z_i v_k$ modulo $\text{Ker } T_\infty$ with respect to v , that is,

$$z_i v_k - \sum_{j=1}^s x_{j,k}^{(i)} v_j \in \text{Ker } T_\infty.$$

which yields

$$T_\infty(z_i v_k) = \sum_{j=1}^s x_{j,k}^{(i)} T_\infty v_j.$$

Since $\deg v_j \leq m$, we have

$$T_m^{(i)} v_k = \sum_{j=1}^s x_{j,k}^{(i)} T_n v_j = \sum_j x_{j,k}^{(i)} \sigma_j u_j.$$

which leads to the desired relation (1.20). \square

Remark 5. When one only has access to a truncated moment matrix $T_m(\mathbf{c})$, the shifted matrix $T_m^{(i)}$ cannot be determined, and further assumptions are needed to carry out the procedure. Typically, if $\text{rank } T_{m-1} = \text{rank } T_m = s$, then one may build the multiplication matrices from $T_{m-1}^{(i)}$. For positive matrices, such conditions are known as *flat extension* conditions, and are central in the trigonometric moment problem, see Section 1.4.

1.3.4 Back to truncated matrices

In practice, we only have access to the truncated moment matrix $T_n(\mathbf{c})$. When $T_n(\mathbf{c}) \succeq 0$, assuming that n is sufficiently large with respect to s and that some flatness property [Curto and Fialkow, 1996] is satisfied, it is known that the multiplication matrices (1.20) are jointly diagonalizable [Laurent, 2010], allowing to recover the support of the underlying measure regardless of its geometry – we further detail this case in the next section. In the generic case however, the theory is more difficult, and most of the results we know make further geometrical assumptions on the support to ensure recovery. In particular, reconstruction is possible if some minimal separation distance condition is satisfied.

Theorem 4 ([Ehler et al., 2019, Theorem 2.1]). *Let $\mathbf{c} \in \mathbb{C}^{N(n)}$ be the moment sequence (1.8) of the s -sparse measure μ_{a_0, \mathbf{x}_0} , and let Δ be the minimal separation distance in μ_{a_0, \mathbf{x}_0} . If*

$$n > 2 + \frac{3 \log d}{\Delta} \tag{1.22}$$

then the matrices $X_{(i)}$ (1.20) are jointly diagonalizable. Specifically, there exists $p \in \mathcal{M}_s(\mathbb{C})$ invertible such that

$$X_{(i)} = P^{-1} \text{Diag}(\zeta_{1,i}, \dots, \zeta_{s,i}) P \tag{1.23}$$

for all $i = 1, \dots, d$.

Algorithm 2: Multivariate Prony's method

Input: $T_n(\mathbf{c})$, for $\mathbf{c} \in \mathbb{C}^{N(n)}$
1 with $r = \text{rank } T_{n-1}(\mathbf{c})$, compute the reduced svd $T_{n-1}(\mathbf{c}) = U\Sigma V^H$
2 **for** $i = 1, \dots, d$ **do**
3 compute the shifted matrix $T_{n-1}^{(i)}$
4 compute the multiplication matrix $X_{(i)} = \Sigma^{-1}U^H T_{n-1}^{(i)}V$
5 **end**
6 jointly diagonalize $X_{(1)}, \dots, X_{(d)}$, as $X_{(i)} = P^{-1}\text{Diag}(\zeta_{1,i}, \dots, \zeta_{s,i})P$
7 compute $\mathbf{x}_j = (\arg \zeta_{j,1}/2\pi, \dots, \arg \zeta_{j,d}/2\pi)$, $j = 1, \dots, r$
8 compute $\{a_1, \dots, a_r\}$ as a solution to
Output: $r, \{\mathbf{x}_j\}, \{a_j\}$

1.4 The positive case

Given a sequence $\mathbf{c} \in \mathbb{C}^N$, Prony's method can only work when \mathbf{c} derives from some sufficiently sparse measure, and the performance of the algorithm therefore depends on whether such a measure exists and is unique. Although not much can be said about these questions in the general case, they are on the contrary much more widely understood when one looks for a positive representing measure. This falls within the scope of the celebrated trigonometric moment problem (TMP), and we recall in this section several fundamental results from this theory. The reader will find in Figure 1.2 a partial graphical representation of the results of this section, displaying some of the correspondences between moment matrices and representing measures.

1.4.1 The truncated moment problem

In all this section, we consider sequences $(c_{\mathbf{k}})$ that are Hermitian, *i.e.* that satisfies $c_{-\mathbf{k}} = \overline{c_{\mathbf{k}}}$. In particular, the matrix $T_n(\mathbf{c})$ is Hermitian. A first necessary condition for such a sequence to admit a positive representing measure is that the matrix $T_n(\mathbf{c})$ be positive semidefinite.

Proposition 2. *If there is a positive measure explaining \mathbf{c} , then $T_n(\mathbf{c}) \succeq 0$.*

Proof. Let μ be a positive Radon measure explaining \mathbf{c} . Since μ is real-valued, $T_n(\mathbf{c})$ is Hermitian. Moreover, for any $p \in \mathbb{C}^{N_+(n)}$,

$$p^H T_n(\mathbf{c}) p = \int_{\mathbb{T}^d} |P^\mathbb{T}(\mathbf{x})|^2 d\mu(\mathbf{x}) \geq 0.$$

□

In the following, we denote by \mathcal{H}_N (resp. \mathcal{H}_N^+) the set of Hermitian (resp. Hermitian positive semidefinite) matrices of size $N \in \mathbb{N}$.

The TMP in one dimension The very first result on the truncated trigonometric moment problem, which can be traced back to the work of Carathéodory and Toeplitz [Carathéodory, 1907, Carathéodory, 1911, Toeplitz, 1911], gives a converse result in one dimension.

Theorem 5 (Carathéodory-Toeplitz). *A (Hermitian) sequence $(c_k)_{k=-n,\dots,n}$ admits a positive representing measure supported on s points $\zeta_1, \dots, \zeta_s \in \mathbb{U}$ ($s \leq n+1$) if and only if $T_n(c) \succeq 0$ and $\text{rank } T_n(c) = s$. This measure is unique if $s \leq n$.*

This result has been revisited several times with various approaches, ranging from operator theory [Toeplitz, 1911, Stone, 1932] or orthogonal polynomials [Carathéodory, 1911, Akhiezer and Krein, 1938], to more recently flat extensions [Curto and Fialkow, 1991]

Remark 6. Note that a sequence c satisfying $T_n(c) \succeq 0$ and $\text{rank } T_n(c) = r$ may also admit a *non positive* representing measure supported on strictly more than r points.

The multi-dimensional case: flat extensions In one dimension, if the moment matrix $T_n(\mathbf{c})$ is rank deficient of rank r , not only does it admit a unique r -sparse positive representing measure, but also, as noted in [Condat, 2020], there are no sparser representing measures. These properties do not generalize as is to the multi-dimensional case. Existence and determinacy of positive representing measures in that case are governed by the so-called *flatness property*, formalized by Curto and Fialkow [Curto and Fialkow, 1991, Curto and Fialkow, 1996, Curto and Fialkow, 1998], which is more restrictive than rank-deficiency.

Given a Hermitian matrix $Y \in \mathcal{H}_{N_+(n)}$, we write the block decomposition

$$Y = \begin{array}{c} \Omega_{n-1}^+ \\ \Omega_n^+ \setminus \Omega_{n-1}^+ \end{array} \begin{pmatrix} \Omega_{n-1}^+ & \Omega_n^+ \setminus \Omega_{n-1}^+ \\ X & B \\ B^* & C \end{pmatrix}. \quad (1.24)$$

Definition 5 (Flat extension). *Let $Y \in \mathcal{H}_{N_+(n)}$. In the decomposition (1.24), we say that Y is flat (or that Y is a flat extension of X) if $\text{rank}(Y) = \text{rank}(X)$.*

Theorems 6 and 7 below, due to Curto and Fialkow [Curto and Fialkow, 1996], generalize Carathéodory-Toeplitz to any dimension.

Theorem 6 ([Curto and Fialkow, 1996, Theorem 5.13]). *A sequence $\mathbf{c} \in \mathbb{C}^{N(n)}$ admits a positive, $\text{rank } T_n(\mathbf{c})$ -sparse representing measure if and only if $T_n(\mathbf{c}) \succeq 0$, and $T_n(\mathbf{c})$ admits a positive flat extension of the form $T_{n+1}(\tilde{\mathbf{c}})$.*

Admitting a flat extension guarantees the existence of a positive representing measure, but one only has unicity when T_n is itself flat. The following result adapts Corollary 5.14 in [Curto and Fialkow, 1996] to our setting. As it is not exactly the one used in [Curto and Fialkow, 1996], we detail its proof in the next section.

Theorem 7 ([Curto and Fialkow, 1996, Corollary 5.14]). *If $T_n(\mathbf{c})$ is flat and $T_n(\mathbf{c}) \succeq 0$, then \mathbf{c} admits a unique positive representing measure, and this measure is rank $T_n(\mathbf{c})$ -sparse.*

1.4.2 Proof of Theorem 7

The proof of this result essentially relies on the ideas of Curto and Fialkow [Curto and Fialkow, 1991, Curto and Fialkow, 2002, Curto and Fialkow, 1996] for the truncated K -moment problem (see also the variant [Laurent, 2005] by Laurent) though our problem does not exactly fit into the framework of the aforementioned articles. While experts in this topic should have no difficulty in filling the gap, we present here some detail which might help the non-specialists.

In the monograph [Curto and Fialkow, 1996], the authors consider moment matrices with multivariate polynomials in (z, \bar{z}) . Seeing the torus \mathbb{T} as the complex unit circle makes it possible to reformulate our problem to their setting, except that their moment matrices are indexed by monomials with total degree less than ℓ ,

$$\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d, \quad \text{and} \quad \deg_1(\bar{\mathbf{z}}^{\mathbf{j}_1} \mathbf{z}^{\mathbf{j}_2}) \stackrel{\text{def.}}{=} j_{1,1} + j_{2,1} + \dots + j_{1,d} + j_{2,d} \leq \ell, \quad (1.25)$$

whereas the indices in our framework are naturally selected by their maximum degree,

$$\mathbf{i} \in \mathbb{Z}^d, \quad \text{and} \quad \deg_\infty(\mathbf{z}^{\mathbf{i}}) \stackrel{\text{def.}}{=} |\mathbf{i}|_\infty \stackrel{\text{def.}}{=} \max(|i_1|, \dots, |i_d|) \leq \ell. \quad (1.26)$$

On the contrary, [Laurent et al., 2009] handles moment matrices indexed with more general sets of indices, but their analysis is given in the real case.

In the following, in order to derive Theorem 7, we combine the ideas of [Laurent et al., 2009] and [Curto and Fialkow, 2002]. We consider a positive Borel measure ν defined on the torus, and we see it alternatively as a measure in the complex plane or in the real plane, going through the following different moment matrices.

Trigonometric moment matrices The moment matrices considered in this paper are mainly trigonometric moment matrices. Given a positive measure on \mathbb{T}^d , its (trigonometric) moment matrix $M_{\mathbb{T}}$ has entries

$$(M_{\mathbb{T}})_{\mathbf{i}, \mathbf{j}} = \int_{\mathbb{T}^d} e^{-2i\pi\langle \mathbf{i}, t \rangle} e^{2i\pi\langle \mathbf{j}, t \rangle} d\nu_{\mathbb{T}}(t) = \int_{\mathbb{T}^d} e^{2i\pi\langle \mathbf{j} - \mathbf{i}, t \rangle} d\nu_{\mathbb{T}}(t), \quad (1.27)$$

for $\mathbf{i}, \mathbf{j} \in \mathbb{Z}^d$, $|\mathbf{i}|_\infty, |\mathbf{j}|_\infty \leq \ell$. Such matrices are *generalized Toeplitz* in the sense that for all admissible multi-indices,

$$(M_{\mathbb{T}})_{\mathbf{i} + \mathbf{s}, \mathbf{j}} = (M_{\mathbb{T}})_{\mathbf{i}, \mathbf{j} - \mathbf{s}}. \quad (1.28)$$

For fixed \mathbf{j} , as the column $(M_{\mathbb{T}})_{\cdot, \mathbf{j}}$ contains the moments of the measure $e^{2i\pi\langle \mathbf{j}, t \rangle} d\nu_{\mathbb{T}}(t)$, we say that it corresponds to the monomial $\mathbf{z}^{\mathbf{j}}$. The flatness property (Definition 5) is equivalent to the fact that for all $\mathbf{j} \in \Omega_\ell \setminus \Omega_{\ell-1}$ the column $\mathbf{z}^{\mathbf{j}}$ is a linear combination of the set of columns $\{\mathbf{z}^{\mathbf{j}'}; \mathbf{j}' \in \Omega_{\ell-1}\}$.

Moment matrices in the complex plane. Given a measure $\nu_{\mathbb{C}}$ defined on \mathbb{C}^d , one may consider its moments against the variable \mathbf{z} and its conjugate $\bar{\mathbf{z}}$, namely

$$(M_{\mathbb{C}})_{(\mathbf{i}_1, \mathbf{i}_2), (\mathbf{j}_1, \mathbf{j}_2)} = \int_{\mathbb{C}^d} (\bar{\mathbf{z}}^{\mathbf{i}_2} \mathbf{z}^{\mathbf{i}_1}) (\bar{\mathbf{z}}^{\mathbf{j}_1} \mathbf{z}^{\mathbf{j}_2}) d\nu_{\mathbb{C}}(\mathbf{z}) = \int_{\mathbb{C}^d} \bar{\mathbf{z}}^{\mathbf{j}_1 + \mathbf{i}_2} \mathbf{z}^{\mathbf{i}_1 + \mathbf{j}_2} d\nu_{\mathbb{C}}(\mathbf{z}), \quad (1.29)$$

for $\mathbf{i}_1, \mathbf{i}_2, \mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d$ such that $\max(\mathbf{i}_1 + \mathbf{i}_2) \leq \ell$ and $\max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell$. Such matrices have a structure property recalling that of Hankel matrices,

$$(M_{\mathbb{C}})_{(\mathbf{i}_1 + \mathbf{r}_1, \mathbf{i}_2 + \mathbf{r}_2), (\mathbf{j}_1, \mathbf{j}_2)} = (M_{\mathbb{C}})_{(\mathbf{i}_1, \mathbf{i}_2), (\mathbf{j}_1 + \mathbf{r}_1, \mathbf{j}_2 + \mathbf{r}_2)}. \quad (1.30)$$

Similarly as above, we note that the columns $(M_{\mathbb{C}})_{\cdot, (\mathbf{j}_1, \mathbf{j}_2)}$ correspond to the monomial $\bar{\mathbf{z}}^{\mathbf{j}_1} \mathbf{z}^{\mathbf{j}_2}$. We say that $M_{\mathbb{C}}$ is *flat* if the columns corresponding to $\bar{\mathbf{z}}^{\mathbf{j}_1} \mathbf{z}^{\mathbf{j}_2}$, where $\max(\mathbf{j}_1 + \mathbf{j}_2) = \ell$, are linear combinations of the columns $\bar{\mathbf{z}}^{\mathbf{j}'_1} \mathbf{z}^{\mathbf{j}'_2}$ where $\max(\mathbf{j}'_1 + \mathbf{j}'_2) \leq \ell - 1$.

Moment matrices in the real plane. Given a measure $\nu_{\mathbb{R}^2}$ defined on $(\mathbb{R}^2)^d$, we consider its moments against the variables \mathbf{x} and \mathbf{y} , namely

$$\begin{aligned} (M_{\mathbb{R}^2})_{(\mathbf{i}_1, \mathbf{i}_2), (\mathbf{j}_1, \mathbf{j}_2)} &= \int_{(\mathbb{R}^d)^2} (\mathbf{x}^{\mathbf{i}_1} \mathbf{y}^{\mathbf{i}_2}) (\mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2}) d\nu_{\mathbb{R}^2}(\mathbf{x}, \mathbf{y}) \\ &= \int_{(\mathbb{R}^d)^2} (\mathbf{x}^{\mathbf{i}_1 + \mathbf{j}_1} \mathbf{y}^{\mathbf{i}_2 + \mathbf{j}_2}) d\nu_{\mathbb{R}^2}(\mathbf{x}, \mathbf{y}) \end{aligned} \quad (1.31)$$

for $\mathbf{i}_1, \mathbf{i}_2, \mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d$ such that $\max(\mathbf{i}_1 + \mathbf{i}_2) \leq \ell$ and $\max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell$. Such moment matrices have the *generalized Hankel* property, that is

$$(M_{\mathbb{R}^2})_{(\mathbf{i}_1 + \mathbf{r}_1, \mathbf{i}_2 + \mathbf{r}_2), (\mathbf{j}_1, \mathbf{j}_2)} = (M_{\mathbb{R}^2})_{(\mathbf{i}_1, \mathbf{i}_2), (\mathbf{j}_1 + \mathbf{r}_1, \mathbf{j}_2 + \mathbf{r}_2)}. \quad (1.32)$$

For fixed $\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d$, the column $(M_{\mathbb{R}^2})_{\cdot, (\mathbf{j}_1, \mathbf{j}_2)}$ corresponds to the monomial $\mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2}$. As above, we say that $M_{\mathbb{R}^2}$ is *flat* if the columns corresponding to $\mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2}$, where $\max(\mathbf{j}_1 + \mathbf{j}_2) = \ell$, are linear combinations of the columns $\mathbf{x}^{\mathbf{j}'_1} \mathbf{y}^{\mathbf{j}'_2}$ where $\max(\mathbf{j}'_1 + \mathbf{j}'_2) \leq \ell - 1$.

From the torus to the complex plane Given a positive Borel measure $\nu \stackrel{\text{def.}}{=} \nu_{\mathbb{T}}$ on \mathbb{T}^d , we may see it as a measure in \mathbb{C}^d by considering its image measure by T , $\nu_{\mathbb{C}} \stackrel{\text{def.}}{=} T_{\#} \nu_{\mathbb{T}}$, where

$$T : \mathbb{T}^d \rightarrow \mathbb{C}^d, \quad (t_1, \dots, t_d) \mapsto (e^{2i\pi t_1}, \dots, e^{2i\pi t_d}). \quad (1.33)$$

The resulting measure has support in $(\mathbb{S}^1)^d$, where $\mathbb{S}^1 \stackrel{\text{def.}}{=} \{z \in \mathbb{C}; |z| = 1\}$. We recall that the image measure $\nu_{\mathbb{C}} = T_{\#} \nu_{\mathbb{T}}$ is characterized by $\nu_{\mathbb{C}}(B) = \nu_{\mathbb{T}}(T^{-1}(B))$ for all Borel sets, so that for all $\nu_{\mathbb{C}}$ -summable function ψ ,

$$\int_{\mathbb{C}^d} \psi(z) d\nu_{\mathbb{C}}(z) = \int_{\mathbb{T}^d} \psi(T(t)) d\nu_{\mathbb{T}}(t). \quad (1.34)$$

Obviously, the moment matrices $M_{\mathbb{T}}$ and $M_{\mathbb{C}}$ have different sizes. However, the following is a first step in relating them.

Lemma 5. *Let $\ell \geq 2$, and let $M_{\mathbb{C}}$ be a moment matrix representing some positive Borel measure $\nu_{\mathbb{C}}$ on \mathbb{C}^d . Then $\text{Supp } \nu_{\mathbb{C}} \subset (\mathbb{S}^1)^d$ if and only if the columns corresponding to $\bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2}$ and $\bar{z}^{\mathbf{j}'_1} z^{\mathbf{j}'_2}$ are equal for all multi-indices such that $\mathbf{j}_1 - \mathbf{j}_2 = \mathbf{j}'_1 - \mathbf{j}'_2$, where $|\mathbf{j}_1 + \mathbf{j}_2|_{\infty} \leq \ell$, $|\mathbf{j}'_1 + \mathbf{j}'_2|_{\infty} \leq \ell$.*

Proof. The column $\bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2}$ contains elements of the form $\int_{\mathbb{C}^d} (\bar{z}^{\mathbf{i}_1} z^{\mathbf{i}_2}) \bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2} d\nu_{\mathbb{C}}(z)$ for $\mathbf{i}_1, \mathbf{i}_2 \in \mathbb{N}^d$. If $\text{Supp } \nu_{\mathbb{C}} \subset (\mathbb{S}^1)^d$, then $\bar{z}^{\mathbf{j}_1, k} z^{\mathbf{j}_2, k} = \bar{z}^{\mathbf{j}'_1, k} z^{\mathbf{j}'_2, k}$ for all z in the integration domain and all $k \in \llbracket 1, d \rrbracket$, hence

$$\int_{\mathbb{C}^d} (\bar{z}^{\mathbf{i}_1} z^{\mathbf{i}_2}) \bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2} d\nu_{\mathbb{C}}(z) = \int_{\mathbb{C}^d} (\bar{z}^{\mathbf{i}_1} z^{\mathbf{i}_2}) \bar{z}^{\mathbf{j}'_1} z^{\mathbf{j}'_2} d\nu_{\mathbb{C}}(z), \quad (1.35)$$

and the two columns are equal.

Conversely, if the above-mentioned columns are equal, let $k \in \llbracket 1, d \rrbracket$, and $\mathbf{j} \in \mathbb{N}^d$ such that $j_{k'} = 1$ for $k' = k$, 0 otherwise. Then,

$$\begin{aligned} \int_{\mathbb{C}^d} (1 - |z_k|^2)^2 d\nu_{\mathbb{C}}(z) &= \int_{\mathbb{C}^d} 1 d\nu_{\mathbb{C}}(z) + \int_{\mathbb{C}^d} (\bar{z}^{\mathbf{j}} z^{\mathbf{j}}) (\bar{z}^{\mathbf{j}} z^{\mathbf{j}}) d\nu_{\mathbb{C}}(z) - 2 \int_{\mathbb{C}^d} (\bar{z}^{\mathbf{j}} z^{\mathbf{j}}) d\nu_{\mathbb{C}}(z) \\ &= \int_{\mathbb{C}^d} 1 d\nu_{\mathbb{C}}(z) - \int_{\mathbb{C}^d} (\bar{z}^{\mathbf{j}} z^{\mathbf{j}}) d\nu_{\mathbb{C}}(z) \\ &= 0, \end{aligned} \quad (1.36)$$

using the equality between the columns 1 and $\bar{z}^{\mathbf{j}} z^{\mathbf{j}}$ and the corresponding relations between their entries. Since the integrand is nonnegative, we deduce that ν charges only points where $\bar{z}^{\mathbf{j}} z^{\mathbf{j}} = 1$, i.e. $\bar{z}_k z_k = 1$. As a result, $\text{Supp } \nu \subset \bigcap_{k=1}^d \{z \in \mathbb{C}^d ; \bar{z}_k z_k = 1\} = (\mathbb{S}^1)^d$. \square

Using Lemma 5, we see that for all $\mathbf{j} \in \mathbb{Z}^d$ such that $|\mathbf{j}|_{\infty} \leq \ell$, all the columns $(M_{\mathbb{C}})_{\cdot, (\mathbf{j}_1, \mathbf{j}_2)}$ such that $\mathbf{j}_2 - \mathbf{j}_1 = \mathbf{j}$ (and $\max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell$) are equal. In fact, from (1.34) we see that those columns are obtained by "repeating" the column \mathbf{j} of $M_{\mathbb{T}}$ at all indices such that $\mathbf{j}_2 - \mathbf{j}_1 = \mathbf{j}$ (and similarly for the rows).

More precisely, given a maximal degree ℓ , recall our notation for the set of Laurent polynomials

$$\mathbb{C}_{\ell}[\mathbf{z}, \mathbf{z}^{-1}] \stackrel{\text{def.}}{=} \text{Span} \{ \mathbf{z}^{\mathbf{j}} ; \mathbf{j} \in \mathbb{Z}^d, |\mathbf{j}|_{\infty} \leq \ell \}. \quad (1.37)$$

Let us also denote the set of polynomials by

$$\mathbb{C}_{\ell}[\mathbf{z}, \bar{\mathbf{z}}] \stackrel{\text{def.}}{=} \text{Span} \{ \bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2} ; \mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d, \max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell \} \quad (1.38)$$

and by J the matrix of the operator

$$\mathbb{C}_{\ell}[\mathbf{z}, \bar{\mathbf{z}}] \rightarrow \mathbb{C}_{\ell}[\mathbf{z}, \mathbf{z}^{-1}], \quad \bar{z}^{\mathbf{j}_1} z^{\mathbf{j}_2} \mapsto \frac{1}{c(\mathbf{j}_2 - \mathbf{j}_1)} z^{\mathbf{j}_2 - \mathbf{j}_1}, \quad (1.39)$$

where $c(\mathbf{j}) = \text{Card} \{ (\mathbf{j}'_1, \mathbf{j}'_2) ; \mathbf{j}'_1, \mathbf{j}'_2 \in \mathbb{N}^d, \max(\mathbf{j}'_1 + \mathbf{j}'_2) \leq \ell, \mathbf{j}'_2 - \mathbf{j}'_1 = \mathbf{j} \}$. Then, from (1.34), the following relation holds,

$$M_{\mathbb{C}} = J^* M_{\mathbb{T}} J. \quad (1.40)$$

We immediately deduce:

Lemma 6. *The relation $\nu_{\mathbb{C}} = T_{\#}\nu_{\mathbb{T}}$ defines a one-to-one correspondence between (positive Borel) measures $\nu_{\mathbb{T}}$ on \mathbb{T}^d and measures on \mathbb{C}^d supported on $(\mathbb{S}^1)^d$. That correspondence preserves the cardinality of the support and one has $T(\text{Supp } \nu_{\mathbb{T}}) = \text{Supp } \nu_{\mathbb{C}}$. Moreover their moment matrices are related by (1.40).*

Conversely, let $M_{\mathbb{T}}$ and $M_{\mathbb{C}}$ be matrices (indexed by $\mathbb{C}_{\ell}[\mathbf{z}, \mathbf{z}^{-1}]$ and $\mathbb{C}_{\ell}[\mathbf{z}, \bar{\mathbf{z}}]$ respectively) such that (1.40) holds. Then,

1. $M_{\mathbb{T}}$ satisfies (1.28) if and only if $M_{\mathbb{C}}$ satisfies (1.30),
2. $M_{\mathbb{C}} \succeq 0$ if and only if $M_{\mathbb{T}} \succeq 0$,
3. $\text{rank } M_{\mathbb{T}} = \text{rank } M_{\mathbb{C}}$. Moreover $M_{\mathbb{T}}$ is flat if and only if $M_{\mathbb{C}}$ is flat.

From the complex plane to the real plane Now, given a positive Borel measure $\nu_{\mathbb{C}}$ on \mathbb{C}^d , we see it as a measure on $(\mathbb{R}^2)^d$ by considering its image measure by S , $\nu_{\mathbb{R}^2} \stackrel{\text{def.}}{=} S_{\#}\nu_{\mathbb{C}}$, where

$$S : \mathbb{C}^d \rightarrow (\mathbb{R}^2)^d, \quad (z_1, \dots, z_d) \mapsto \left(\frac{z_1 + \bar{z}_1}{2}, \frac{z_1 - \bar{z}_1}{2i}, \dots, \frac{z_d + \bar{z}_d}{2}, \frac{z_d - \bar{z}_d}{2i} \right). \quad (1.41)$$

We consider the following subspace of polynomials in the variables $X_1, Y_1, \dots, X_d, Y_d$,

$$\mathbb{C}_{\ell}[\mathbf{x}, \mathbf{y}] \stackrel{\text{def.}}{=} \text{Span} \{ \mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2} ; \mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d, \max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell \}. \quad (1.42)$$

Obviously, $\mathbb{C}_{\ell}[\mathbf{z}, \bar{\mathbf{z}}]$ and $\mathbb{C}_{\ell}[\mathbf{x}, \mathbf{y}]$ are isomorphic as vector spaces. We are interested in the relations between the moment matrices $M_{\mathbb{R}^2}$ and $M_{\mathbb{C}}$ when changing variables with S , that is

$$\forall k \in \{1, \dots, d\}, \quad z_k \stackrel{\text{def.}}{=} x_k + iy_k, \quad \bar{z}_k \stackrel{\text{def.}}{=} x_k - iy_k, \quad \text{or conversely,} \quad (1.43)$$

$$\forall k \in \{1, \dots, d\}, \quad x_k \stackrel{\text{def.}}{=} \frac{1}{2}(z_k + \bar{z}_k), \quad y_k \stackrel{\text{def.}}{=} \frac{1}{2i}(z_k - \bar{z}_k). \quad (1.44)$$

We note that for all $k \in \{1, \dots, d\}$, given some indices i_k, j_k ,

$$\bar{z}_k^{j_{1,k}} z_k^{j_{2,k}} = (x_k + iy_k)^{j_{1,k}} (x_k - iy_k)^{j_{2,k}} = \sum_{r_1, r_2} c_{r_1, r_2} x_k^{r_1} y_k^{r_2},$$

where $c_{r_1, r_2} \in \mathbb{C}$ and the sum is over all the indices $r_1, r_2 \in \mathbb{N}$ such that $r_1 + r_2 = j_{1,k} + j_{2,k}$. As a result, given $\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d$,

$$\bar{\mathbf{z}}^{\mathbf{j}_1} \mathbf{z}^{\mathbf{j}_2} = \prod_{k=1}^d (x_k + iy_k)^{j_{1,k}} (x_k - iy_k)^{j_{2,k}} = \sum_{\mathbf{r}_1, \mathbf{r}_2} c_{\mathbf{r}_1, \mathbf{r}_2} \mathbf{x}^{\mathbf{r}_1} \mathbf{y}^{\mathbf{r}_2}, \quad (1.45)$$

where the sum is over all the multi-indices $\mathbf{r}_1, \mathbf{r}_2 \in \mathbb{N}^d$ such that, for all k , $r_{1,k} + r_{2,k} = j_{1,k} + j_{2,k}$.

As a result, the change of variable (1.43) induces a linear map $L : \mathbb{C}_\ell[\mathbf{x}, \mathbf{y}] \rightarrow \mathbb{C}_\ell[\mathbf{z}, \bar{\mathbf{z}}]$ which admits a block decomposition, mapping surjectively (hence bijectively), for each $\mathbf{t} \in \mathbb{N}^d$ with $\max(\mathbf{t}) \leq \ell$, the space

$$\begin{aligned} & \text{Span} \{ \mathbf{x}^i \mathbf{y}^j ; \mathbf{i}, \mathbf{j} \in \mathbb{N}^d, \mathbf{i} + \mathbf{j} = \mathbf{t} \} \\ \text{onto} & \quad \text{Span} \{ \bar{\mathbf{z}}^i \mathbf{z}^j ; \mathbf{i}, \mathbf{j} \in \mathbb{N}^d, \mathbf{i} + \mathbf{j} = \mathbf{t} \} \end{aligned}$$

Confronting (1.29) and (1.31), in view of the change of variable formula

$$\int_{(\mathbb{R}^d)^2} \psi(\mathbf{x}, \mathbf{y}) d\nu_{\mathbb{R}^2}(\mathbf{x}, \mathbf{y}) = \int_{\mathbb{C}^d} \psi(S(\mathbf{z})) d\nu_{\mathbb{C}}(\mathbf{z}), \quad (1.46)$$

we deduce that

$$M_{\mathbb{R}^2} = L^* M_{\mathbb{C}} L, \quad (1.47)$$

where L is invertible.

Lemma 7. *The relation $\nu_{\mathbb{R}^2} = S_{\#} \nu_{\mathbb{C}}$ defines a one-to-one correspondence between (positive Borel) measures $\nu_{\mathbb{C}}$ on \mathbb{C}^d which represent $M_{\mathbb{C}}$ and measures $\nu_{\mathbb{R}^2}$ on $(\mathbb{R}^d)^2$ which represent $M_{\mathbb{R}^2}$. That correspondence preserves the cardinality of the support and $S(\text{Supp } \nu_{\mathbb{C}}) = \text{Supp } \nu_{\mathbb{R}^2}$.*

Conversely, let $M_{\mathbb{C}}$ and $M_{\mathbb{R}^2}$ be matrices (indexed by $\mathbb{C}_\ell[\mathbf{z}, \bar{\mathbf{z}}]$ and $\mathbb{C}_\ell[\mathbf{x}, \mathbf{y}]$ respectively) such that (1.47) holds. Then,

1. $M_{\mathbb{C}}$ satisfies (1.30) if and only if $M_{\mathbb{R}^2}$ satisfies (1.32),
2. $M_{\mathbb{R}^2} \succeq 0$ if and only if $M_{\mathbb{C}} \succeq 0$,
3. $\text{rank } M_{\mathbb{C}} = \text{rank } M_{\mathbb{R}^2}$. Moreover $M_{\mathbb{C}}$ is flat if and only if $M_{\mathbb{R}^2}$ is flat.

Existence of a representing measure The rest of the proof of Theorem 7 relies on the results of [Laurent et al., 2009]. The matrix R is a positive semidefinite matrix indexed by the set of monomials $\{ \mathbf{z}^j ; \mathbf{j} \in \mathbb{Z}^d, |\mathbf{j}|_\infty \leq \ell \}$ which satisfies (1.28).

Applying successively (1.40) and (1.47), we obtain a matrix $N \stackrel{\text{def.}}{=} L^* J^* R J L$, which satisfies (1.32), indexed by the set of monomials

$$\mathcal{C}_\ell \stackrel{\text{def.}}{=} \{ \mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2} ; \mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d, \max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell \}.$$

We note that $\mathcal{C}_{\ell-1}$ is closed under taking divisors, hence connected to 1 (see [Laurent et al., 2009, Sec. 1.3] for the definition). Moreover, its closure,

$$\mathcal{C}_{\ell-1}^+ \stackrel{\text{def.}}{=} \mathcal{C}_{\ell-1} \cup \left(\bigcup_{k=1}^d x_k \mathcal{C}_{\ell-1} \right) \cup \left(\bigcup_{k=1}^d y_k \mathcal{C}_{\ell-1} \right) \quad (1.48)$$

$$= \{ m, x_1 m, \dots, x_d m, y_1 m, \dots, y_d m ; m \in \mathcal{C}_{\ell-1} \} \quad (1.49)$$

is a subset of \mathcal{C}_ℓ . The matrix $N|_{\mathcal{C}_\ell}$ being a flat extension of $N|_{\mathcal{C}_{\ell-1}}$ (by Lemma 6 and 7), this implies that $N|_{\mathcal{C}_{\ell-1}^+}$ is a flat extension of $N|_{\mathcal{C}_{\ell-1}}$.

The theorem [Laurent et al., 2009, Th. 3.2] by Laurent and Mourrain then ensures that there exists a representing measure $\nu_{\mathbb{R}^2}$ for $N|_{\mathcal{C}_{\ell-1}^+}$. As the theorem does not say anything about the rows and columns in $\mathcal{C}_{\ell} \setminus \mathbb{C}_{\ell-1}^+$, we show below that $\nu_{\mathbb{R}^2}$ actually represents all the entries of N .

Lemma 8. *The measure $\nu_{\mathbb{R}^2}$ represents N .*

Proof. Let us write

$$R = \begin{array}{c} \mathcal{C}_{\ell-1} \\ \mathcal{C}_{\ell} \setminus \mathcal{C}_{\ell-1} \end{array} \begin{array}{cc} \mathcal{C}_{\ell-1} & \mathcal{C}_{\ell} \setminus \mathcal{C}_{\ell-1} \\ \left(\begin{array}{cc} A & B \\ B^* & C \end{array} \right) \stackrel{\text{def.}}{=} N. \end{array} \quad (1.50)$$

A first step is to show that $B = AQ$ for some matrix Q which only depends on $N|_{\mathcal{C}_{\ell-1}^+}$, the restriction of N to $\mathcal{C}_{\ell-1}^+ \times \mathcal{C}_{\ell-1}^+$. Then, since $N \succeq 0$ is flat, we deduce by [Curto and Fialkow, 1996, Prop. 2.2] that C is uniquely determined from A and B , as $C = Q^*AQ$.

In a second step, we consider the moment matrix $M_{\mathbb{R}^2}$ of $\nu_{\mathbb{R}^2}$ on \mathcal{C}_{ℓ} . As it is flat, positive semidefinite and generalized Hankel, its blocks satisfy a similar property as those of N , involving some matrix \tilde{Q} which depends on $M_{\mathbb{R}^2}|_{\mathcal{C}_{\ell-1}^+}$. The key point is that since $M_{\mathbb{R}^2}$ and N coincide in $\mathcal{C}_{\ell-1}^+ \times \mathcal{C}_{\ell-1}^+$, the matrices Q and \tilde{Q} are equal, hence the whole matrices N and $M_{\mathbb{R}^2}$ are equal.

The main point is therefore to prove that $B = AQ$ (the argument for $M_{\mathbb{R}^2}$ being similar). To lighten the notation, we write $\tilde{\mathbf{x}}^{\mathbf{j}} \stackrel{\text{def.}}{=} \mathbf{x}^{\mathbf{j}_1} \mathbf{y}^{\mathbf{j}_2}$ with $\mathbf{j} = (\mathbf{j}_1, \mathbf{j}_2)$. For $k \in \{1, \dots, d\}$, we denote by $\mathbf{e}^{(k)} = (\mathbf{e}_1^{(k)}, \mathbf{e}_2^{(k)}) \in \mathbb{N}^d \times \mathbb{N}^d$ any multi-index such that

$$\mathbf{e}_{1,n}^{(k)} = \mathbf{e}_{2,n}^{(k)} = 0 \quad \text{for } n \neq k, \quad \text{and} \quad (\mathbf{e}_1^{(k)}, \mathbf{e}_2^{(k)}) = (1, 0) \text{ or } (0, 1). \quad (1.51)$$

The elements of $\mathcal{C}_{\ell-1}^+ \setminus \mathcal{C}_{\ell-1}$ are of the form $\tilde{\mathbf{x}}^{\mathbf{j}+\mathbf{e}^{(k)}}$ with $\max(\mathbf{j}_1 + \mathbf{j}_2) = \ell - 1$, whereas the elements of $\mathcal{C}_{\ell} \setminus \mathcal{C}_{\ell-1}$ have the form $\tilde{\mathbf{x}}^{\mathbf{j}+\mathbf{e}^{(k_1)}+\dots+\mathbf{e}^{(k_n)}}$ where k_1, \dots, k_n are distinct elements of $\{1, \dots, d\}$.

Let $\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}^d$ with $\max(\mathbf{j}_1 + \mathbf{j}_2) \leq \ell - 1$. Since $N|_{\mathcal{C}_{\ell-1}^+}$ is a flat extension of $N|_{\mathcal{C}_{\ell-1}}$, any column of $N|_{\mathcal{C}_{\ell-1}^+}$ of the form $\tilde{\mathbf{x}}^{\mathbf{j}+\mathbf{e}^{(k_1)}}$ is a linear combination of the columns in $\mathbb{C}_{\ell-1}$. Let $q_{\mathbf{j}, \mathbf{e}^{(k_1)}}$ (resp. $q_{\mathbf{j}, \mathbf{e}^{(k_1)}}(\mathbf{x}, \mathbf{y})$) denote the coefficients of that combination (resp. the corresponding polynomial¹). In other words, the polynomial $f^{(1)} \stackrel{\text{def.}}{=} (\tilde{\mathbf{x}})^{\mathbf{j}+\mathbf{e}^{(k_1)}} - q_{\mathbf{j}, \mathbf{e}^{(k_1)}}(\mathbf{x}, \mathbf{y})$ is in $\ker N|_{\mathcal{C}_{\ell-1}^+}$. Since $N \succeq 0$, we deduce that in fact $f^{(1)} \in \ker N$.

Let $f^{(n)} \stackrel{\text{def.}}{=} (\tilde{\mathbf{x}})^{\mathbf{j}+\mathbf{e}^{(k_1)}+\dots+\mathbf{e}^{(k_n)}} - \tilde{\mathbf{x}}^{\mathbf{e}^{(k_2)}+\dots+\mathbf{e}^{(k_n)}} q_{\mathbf{j}, \mathbf{e}^{(k_1)}}(\mathbf{x}, \mathbf{y}) \in \mathcal{C}_{\ell}$, for $2 \leq n \leq d$. As the k_i 's are pairwise distinct, we see that for $\mathbf{i} \in \mathcal{C}_{\ell-1}$, we have

¹In the following we identify polynomials with the vectors of their coefficients, so that we can “apply a moment matrix to a polynomial” with suitable degree

$(\tilde{\mathbf{x}})^{\mathbf{j}+e^{(k_1)}+\dots+e^{(k_n)}} \in \mathcal{C}_\ell$. Hence, proceeding as in [Laurent, 2010, Lem. 5.7], we get

$$(Nf^{(n)})_{\mathbf{i}} = (Nf^{(1)})_{\mathbf{i}+e^{(k_2)}+\dots+e^{(k_n)}} = 0, \quad (1.52)$$

since $f^{(1)} \in \ker N$.

From (1.52) and the definition of $f^{(n)}$, we deduce that the column $(\tilde{\mathbf{x}})^{\mathbf{j}+e^{(k_1)}+\dots+e^{(k_n)}}$ of B is a linear combination of columns of the form $(\tilde{\mathbf{x}})^{\mathbf{j}'+e^{(k'_1)}+\dots+e^{(k'_{n-1})}}$ with $\max(\mathbf{j}') \leq \ell - 1$. By an easy induction we obtain that the column $(\tilde{\mathbf{x}})^{\mathbf{j}+e^{(k_1)}+\dots+e^{(k_n)}}$ of N is thus a linear combination of columns of the form $(\tilde{\mathbf{x}})^{\mathbf{j}'+e^{(k'_1)}}$. By flatness of $\mathcal{C}_{\ell-1}^+$, it is thus a combination of columns corresponding to $\mathcal{C}_{\ell-1}$.

As a result, there exists a matrix Q such that $B = AQ$, and Q is uniquely determined from the polynomials $q_{\mathbf{j},e^{(k_1)}}$ for $\mathbf{j} \in \mathcal{C}_{\ell-1}$, $1 \leq k_1 \leq d$. Since the same polynomials $q_{\mathbf{j},e^{(k_1)}}$ can be used for $M_{\mathbb{R}^2}$, we obtain that $\tilde{Q} = Q$ hence $M_{\mathbb{R}^2} = N$. \square

Now, we may go back to the torus. We observe that for $1 \leq k \leq d$, the polynomial $x_k^2 + y_k^2 - 1$ is in the kernel of $N = M_{\mathbb{R}^2}$, hence $\nu_{\mathbb{R}^2}$ is supported in $\bigcap_{k=1}^d \{(x, y) \in (\mathbb{R}^d)^2 ; x_k^2 + y_k^2 = 1\}$. Applying Lemma 7 and 6 above, we obtain the existence of a measure $\nu_{\mathbb{T}}$ such that R is the moment matrix of $\nu_{\mathbb{T}}$ on \mathbb{T}^d , which is the claimed result.

1.4.3 Guarantees for Prony

For positive measures, flatness is sufficient to ensure the success of the multivariate procedure described in Section 1.3. Indeed, assuming that $T_n(\mathbf{c})$ is flat (and positive), and denoting by s its rank, it admits a unique representing measure μ , which is s -sparse, and it can be extended flatly to an infinite moment matrix $T_\infty(\mathbf{c})$ of μ is of rank s (Theorem 6). Thus, since $T_n(\mathbf{c})$ is flat, $\text{rank } T_{n-1}(\mathbf{c}) = s$ and $T_n(\mathbf{c})$ is therefore a submatrix of $T_\infty(\mathbf{c})$ with same rank, so that the formula (1.20) for the multiplication matrices hold. Furthermore, since $T_n(\mathbf{c})$ is positive, the dimension of the quotient space \mathcal{T}_n is equal to the cardinality s of the variety $\mathcal{Z}(\langle \text{Ker } T_n(\mathbf{c}) \rangle)$ by Theorem 2. Then, Theorem 3 ensures that the multiplication operators (and hence multiplication matrices given the above) admits a set of s common eigenvectors, in a space of dimension s . This immediately leads to the following conclusion.

Theorem 8 ([Laurent, 2010]). *Let $\mathbf{c} \in \mathbb{C}^{N(n)}$ such that $T_n(\mathbf{c}) \succeq 0$. If $T_n(\mathbf{c})$ is flat, then the matrices $X_{(i)}$ are co-diagonalizable as in (1.23).*

1.5 ESPRIT

We present in this section the ESPRIT method and its multivariate extension, whose principle shares striking similarities with the multivariate Prony's approach discussed previously.




Figure 1.2: Partial visualization of the correspondences between (generalized) Toeplitz matrices and discrete measures. In the absence of the positivity and flatness assumptions, the matrix T_n may not admit a unique representing measure, and it may not always be clear what solution is returned by Prony’s algorithm. However the algorithm should not return a solution corresponding to a dashed line on the drawing.

1.5.1 Univariate ESPRIT

The Estimation of Signal Parameters via Rotational Invariance Techniques (ESPRIT) method was developed by Roy and Kailath [Roy et al., 1986], in the context of array processing and direction of arrivals (DOA) problems. The method is based on the decomposition of the space into a signal and a noise subspace, and further exploits deterministic invariance properties between overlapping subsets of the data. Unlike MUSIC [Schmidt, 1986], the signal subspace is not estimated from a correlation matrix but directly from the data matrix

$$E(\mathbf{y}) \stackrel{\text{def.}}{=} \begin{pmatrix} y_0 & y_1 & \cdots & y_{M-1} \\ y_1 & y_2 & \cdots & y_M \\ \vdots & \vdots & & \vdots \\ y_{L-1} & y_L & \cdots & y_{M+L-2} \end{pmatrix} = E(\mathbf{c}) + E(\mathbf{w}) \quad (1.53)$$

formed from L snapshots $\mathbf{y}^{(k)} \stackrel{\text{def.}}{=} [y_k \ \dots \ y_{k+M-1}]^\top$, $k = 0, \dots, L-1$, where each entry y_k follows the harmonic model (6). As a consequence, $E(\mathbf{y})$ may be decomposed as

$$E(\mathbf{y}) = \sum_{j=1}^s a_j \mathbf{v}_L(x_j) \mathbf{v}_M^\top(x_j) + E(\mathbf{w}), \quad (1.54)$$

where

$$\mathbf{v}_L(x) \stackrel{\text{def.}}{=} [1 \ e^{-2i\pi x} \ \dots \ e^{-2i\pi(L-1)x}]^\top, \quad (1.55)$$

which exhibits a signal subspace $\text{Span}(\mathbf{v}_L(x_1), \dots, \mathbf{v}_L(x_s))$. Let $J_1(L)$ and $J_2(L)$ be the selection matrices defined as

$$\begin{aligned} J_1(L) &\stackrel{\text{def.}}{=} (I_{L-1} \ \mathbf{0}) \in \mathbb{C}^{(L-1) \times L} \\ J_2(L) &\stackrel{\text{def.}}{=} (\mathbf{0} \ I_{L-1}) \in \mathbb{C}^{(L-1) \times L}. \end{aligned} \quad (1.56)$$

The ESPRIT algorithm takes advantage of the fact that the Vandermonde matrix $V_L(\mathbf{x}) \stackrel{\text{def.}}{=} [\mathbf{v}_L(x_1) \ \dots \ \mathbf{v}_L(x_s)]$ satisfies the invariance property

$$J_1(L)V_L(\mathbf{x})\Phi(\mathbf{x}) = J_2(L)V_L(\mathbf{x}) \quad (1.57)$$

where

$$\Phi(\mathbf{x}) \stackrel{\text{def.}}{=} \text{diag}(e^{-2i\pi x_1}, \dots, e^{-2i\pi x_s}) \in \mathbb{C}^{s \times s}. \quad (1.58)$$

In practice, $V_L(\mathbf{x})$ is unknown and cannot be used to recover $\Phi(\mathbf{x})$, but the singular value decomposition of $E(\mathbf{c})$ gives access to a similar matrix, *assuming that the number of sources s is known*.

Lemma 9. *Let $U = [U_s \ \mathbf{0}] \in \mathbb{C}^{L \times M}$ be the left singular vectors of $E(\mathbf{c})$, where U_s contains the singular vectors corresponding to the singular values with the s largest magnitudes. Then $\text{Span } U_s = \text{Span } V_L(\mathbf{x})$, and there exists a matrix Ψ , similar to $\Phi(\mathbf{x})$, such that*

$$J_1(L)U_s\Psi = J_2(L)U_s. \quad (1.59)$$

Proof. Since $\text{Span } U_s = \text{Span } V_L(\mathbf{x})$, there exists $M \in \mathcal{M}_s(\mathbb{C})$ invertible such that $V_L = U_s M$. From (1.57), we get that $J_1(L)U_s M \Phi = J_2(L)U_s M$, and hence $J_1(L)U_s (M \Phi M^{-1}) = J_2(L)U_s$. \square

A similar result holds for the decomposition $[U_s \ U_n]$ of $E(\mathbf{y})$, where the “signal” eigenvectors can be more or less easily identified depending on the size of the spectral gap in $E(\mathbf{y})$, which depends on the noise level. Thus, Ψ may be estimated from (1.59) in the least-square sense as

$$\Psi = (J_1(L)U_s)^\dagger J_2(L)U_s \in \mathcal{M}_s(\mathbb{C}),$$

and the exponentials $\{e^{-2i\pi x_j}\}$ recovered by diagonalization. When s is not known, it can be estimated as the rank of the matrix $T_n(\mathbf{y})$, or via more robust techniques.

Although the method is originally supposed to involve multiple snapshots, ESPRIT adapts readily to the case of a single-snapshot $\mathbf{y} = [y_{-n} \ \dots \ y_n]$, assuming that the length of the snapshot is sufficiently large, that is larger than twice the number s of sources to recover (*i.e.* $n \geq s$). In that case, the matrix $E(\mathbf{y})$ is simply replaced with $T_n(\mathbf{y})$ defined in (1.7), which admits the similar decomposition $T_n(\mathbf{y}) = \sum_{j=1}^s a_j \mathbf{v}_n(x_j) \mathbf{v}_n(x_j)^H + T_n(\mathbf{w})$. Algorithm 3 summarizes the basic steps of this single-snapshot procedure.

Algorithm 3: Single-snapshot least-square ESPRIT

- Input:** $(y_k)_{k=-n, \dots, n-1}$ satisfying (6)
- 1 with $r = \text{rank } T_n$, compute the SVD $T_n = A_n \Sigma_r B_n^H$, with $A_n \in \mathcal{M}_{n,r}(\mathbb{C})$
 - 2 compute $\Psi = (J_1(n)A_n)^\dagger J_2(n)A_n$
 - 3 compute eigenvalues $\{\zeta_1, \dots, \zeta_r\}$ of Ψ
 - 4 with $x_{0,j} = -\arg \zeta_j / 2\pi$, compute $\{a_1, \dots, a_r\}$ as a solution to (1.9)
- Output:** $\{x_j\}, \{a_j\}$
-

Remark 7 (Connection with Prony). Let $T_{r,r+1}(\mathbf{c})$ be the truncated moment matrix of \mathbf{c} of size $r \times (r+1)$. Considering the noiseless case, and using the recurrence relation (1.11), note that $J_1(r)T_{r,r+1}(\mathbf{c}) = J_2(r)T_{r,r+1}(\mathbf{c})C_0$, where C_0 is the companion matrix (1.13). On the other hand, let $T_{r,r+1} = U_r \Sigma_r Z_{r+1}^H$ be the singular value decomposition of $T_{r,r+1}$. Then

$$\begin{aligned} J_1(r)T_{r,r+1} &= J_1(r)U_r \Sigma_r Z_{r+1}^H \\ &= J_2(r)U_r \Psi^H \Sigma_r Z_{r+1}^H \quad \text{from (1.59)} \\ &= J_2(r)T_{r,r+1}(\Sigma_r Z_{r+1}^H)^{-1} \Psi^H \Sigma_r Z_{r+1}^H \end{aligned}$$

from which we deduce that $C_0 = (\Sigma_r Z_{r+1}^H)^{-1} \Psi^H \Sigma_r Z_{r+1}^H$, which gives an explicit correspondence between the two methods (in the noiseless case).

1.5.2 Multivariate ESPRIT

The multivariate extension of ESPRIT we describe in this section was first proposed in [Haardt and Nossek, 1998]. For $n \in \mathbb{N}$, we define d pairs of selection matrices as

$$\begin{aligned} J_{1,i}(n) &\stackrel{\text{def.}}{=} I_{n^{d-i}} \otimes J_1(n) \otimes I_{n^{i-1}} \in \mathbb{C}^{(n-1)n^{d-1} \times n^d} \\ J_{2,i}(n) &\stackrel{\text{def.}}{=} I_{n^{d-i}} \otimes J_2(n) \otimes I_{n^{i-1}} \in \mathbb{C}^{(n-1)n^{d-1} \times n^d} \end{aligned} \quad (1.60)$$

for $i = 1, \dots, d$, where $J_1(n)$ and $J_2(n)$ are defined in (1.56). The invariance equations (1.57) then generalize as the following d invariance equations (one for each dimension):

$$J_{2,i}V_n(\mathbf{x}) = J_{1,i}V_n(\mathbf{x})\Phi_i, \quad i = 1, \dots, d$$

where

$$\Phi_i = \text{Diag}(e^{-2i\pi x_{1,i}}, \dots, e^{-2i\pi x_{s,i}}) \in \mathbb{C}^{s \times s}.$$

The following Lemma extends Lemma 9.

Lemma 10. *Let U_s be the left singular vectors of T_n , associated to the s largest singular values. There exists an invertible matrix $T \in \mathbb{C}^{r \times r}$ such that, for $i = 1, \dots, d$, the matrices $\Psi_i \stackrel{\text{def.}}{=} T\Phi_i T^{-1}$ satisfy*

$$J_{2,i} U_s \Psi_i = J_{1,i} U_s. \tag{1.61}$$

Thus, in a similar fashion, Lemma 10 ensures that the matrices Ψ_i defined by (1.61) may be jointly diagonalized to recover the matrices Φ_i , and thus the correctly matched locations $\mathbf{x}_1 \dots, \mathbf{x}_s$. As in Remark 7, the matrices Ψ_i are directly connected to the multiplication matrices of the multivariate matrices described in Section 1.3.3.

1.6 Extraction for non-flat data

Prony’s method is designed to recover a discrete measure, by encoding its support as the (finite) set of roots of a few trigonometric polynomials. The previous sections make the assumption that the quotient space \mathcal{T}_n is of finite dimension (*i.e.* that the ideal $\langle \text{Ker } T_n \rangle$ is zero-dimensional), so that in particular the matricial representations (1.20) for the multiplication operators may hold. When \mathcal{T}_n is of infinite dimension on the other hand, the matrices (1.20) may still be built, but are unlikely to commute, hence to be co-diagonalizable. However, one may wonder how the extraction procedure would behave in this setting, assuming that some approximate joint diagonalization scheme is used. This section studies (mainly numerically) this computational framework. The main motivation of this study is the application of Algorithm 2 to data deriving from a measure supported on a continuous curve, or more general hypersurfaces, which we then use in applications to optimal transport in Chapter 4.

1.6.1 Rank of moment matrices

A crucial parameter in Prony’s method is the rank of the input matrix, which determines the cardinality of the reconstructed support. We discuss in this section the link between the rank of a moment matrix T_n and the geometry of the variety of $\mathcal{I}_n \stackrel{\text{def.}}{=} \langle \text{Ker } T_n \rangle$. When \mathcal{I}_n is zero-dimensional, we have seen that, for n large enough (for instance, having flatness suffices), the rank of T_n is constant and equal to the number of points in the variety of \mathcal{I}_n . In the general case, the rank of T_n is a polynomial in n , whose degree is given by the dimension of the variety \mathcal{I}_n [Pauwels et al., 2020, proposition 4]. This is illustrated in Figure 1.3. The function $n \mapsto \text{rank } T_n$ is known as the Hilbert function (or Hilbert polynomial) associated to the variety of \mathcal{I}_n .

For sparse measures, the rank of the moment matrix can be robustly estimated from its singular value decomposition, which features a significant and

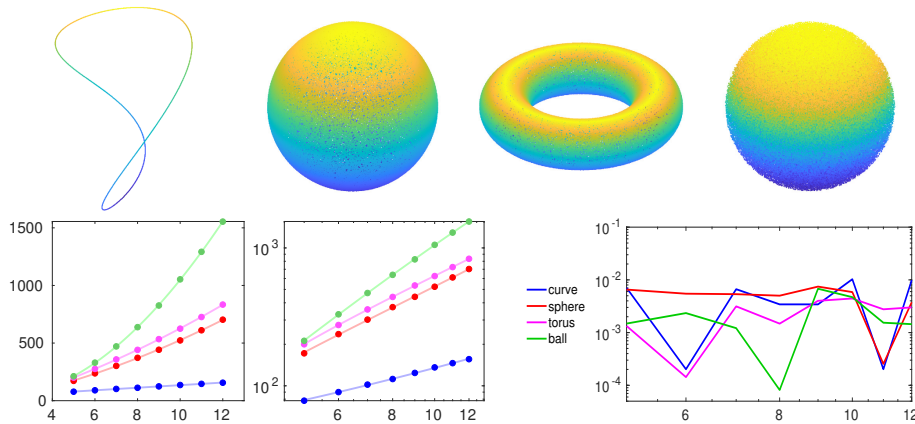


Figure 1.3: Top: visualization of the varieties of \mathbb{R}^3 considered in this example, *i.e.* a curve (1D), a sphere (2D), a torus (2D) and a ball (3D). Bottom: the left plot shows the evolution of the rank of the moment matrix $T_n(\mathbf{c})$ with respect to n , shown also in log-log scale in the middle. Dots represent the measured rank and lines the interpolation of degree 1, 2 or 3 according to the dimension of the underlying variety. The last plot shows the residuals between the measured rank and the interpolation for each case.

easily detectable gap, see Figure 1.4. In other cases however, although the moment matrix may remain rank-deficient, this gap tends to disappear, see Figure 1.4. To make the extraction more robust in this setting, we propose to regularize the problem by simply truncating the singular value decomposition in the first step of the algorithm, with respect to some tolerance. We present in Figure 1.7 some results on the impact of the threshold on the reconstructed signals.

1.6.2 Joint diagonalization by random linear combination

Assuming that the matrices $X_{(1)}, \dots, X_{(d)}$ are jointly diagonalizable, typical implementations uncover the common eigenstructure by diagonalizing a random linear combination $\sum \lambda_i X_{(i)}$ [Stetter, 1996, Corless et al., 1997]. If the sought points ζ_j are all distinct, this produces a matrix whose eigenvalues $\sum_i \lambda_i \zeta_{j,i}$ are all simple with probability one (X_λ is said to be non-derogatory in that case), and hence whose eigenspaces automatically reveal the common eigenvectors. In particular, these random linear combinations resolve the problem of multiple eigenvalues which may appear for some $X_{(i)}$, for instance as soon as two points have the same i -th coordinate.

These approaches rely on the fact that there is a *finite* number of points ζ_j , *i.e.* that the ideal \mathcal{I}_n is zero-dimensional. On the other hand, the picture is fairly different when dealing with positive dimensional ideals. We present in this section a theoretical and experimental analysis of the failure modes of the random projection method.




Figure 1.4: Singular values of the moment matrix of discrete random measures on the sphere. As the sampling refines ($s = 50, 100, 250, 500$), the spectral gap diminishes until it disappears. The first step of the multivariate extraction in this case is to threshold the singular values. The chosen threshold determines the sparsity of the reconstructed measure (see Figure 1.7 for instance).

Let $\lambda \in \mathbb{C}_*^d$ and let $X_{(i)}$, $i = 1, \dots, d$ be co-diagonalizable matrices, with respective eigenvalues $\{\zeta_{1,i}, \dots, \zeta_{s,i}\}$, such that $PX_{(i)} = \text{Diag}(\zeta_{1,i}, \dots, \zeta_{s,i})P$ for some invertible matrix P . Let $X_\lambda \stackrel{\text{def.}}{=} \sum \lambda_i X_{(i)}$; the eigenvalues of X_λ are of the form $\sum \lambda_i \zeta_{j,i}$. The random linear combination method assumes that X_λ is non-derogatory, which fails when at least two distinct eigenvalues ζ_k and ζ_l satisfy $\sum \lambda_i \zeta_{k,i} = \sum \lambda_i \zeta_{l,i}$. Although this happens with probability zero when considering a finite set of points $\{\zeta_1, \dots, \zeta_s\}$, such relations are likely to hold if this set is infinite. We detail these failing cases when $d = 2$.

In 2-D, let $\lambda = (\lambda_1, \lambda_2) \in \mathbb{C}^2 \setminus \{0\}$, and let $\zeta = (\zeta_1, \zeta_2) \in \mathbb{U}^2$ be given, with $\zeta_i = e^{-2i\pi x_i}$, for $x_i \in \mathbb{T}$ ($i = 1, 2$). We search for points $\xi \in \mathbb{U}^2$ that satisfy

$$\lambda_1 \xi_1 + \lambda_2 \xi_2 = \lambda_1 \zeta_1 + \lambda_2 \zeta_2, \quad (1.62)$$

or equivalently, for points $\mathbf{t} \stackrel{\text{def.}}{=} (t_1, t_2) \in \mathbb{T}^2$ such that

$$\lambda_1 e^{-2i\pi t_1} + \lambda_2 e^{-2i\pi t_2} = \lambda_1 e^{-2i\pi x_1} + \lambda_2 e^{-2i\pi x_2}. \quad (1.63)$$

In the following, let

$$Z_\lambda : \mathbb{T}^2 \rightarrow \mathbb{C}, \quad \mathbf{t} \mapsto \lambda_1 e^{-2i\pi t_1} + \lambda_2 e^{-2i\pi t_2}.$$

For $z \in \mathbb{C}$, let $\theta(z) \in \mathbb{T}$ denote the argument of z , i.e. $\theta(z) \stackrel{\text{def.}}{=} (-\arg z / 2\pi) \in \mathbb{T}$. We introduce the following mapping on \mathbb{T}^2 :

$$h_\lambda : \mathbb{T}^2 \rightarrow \mathbb{T}^2, \quad \mathbf{t} \mapsto (2\theta(Z_\lambda(\mathbf{t})) - 2\theta(\lambda_1) - t_1, \quad 2\theta(Z_\lambda(\mathbf{t})) - 2\theta(\lambda_2) - t_2). \quad (1.64)$$

The next lemma explicitly builds a point satisfying (1.63).

Lemma 11. *Given $\mathbf{x} \in \mathbb{T}^2$ and $\lambda \in \mathbb{C}^2$, the point*

$$\mathbf{t}_0 \stackrel{\text{def.}}{=} h_\lambda(\mathbf{x}) \quad (1.65)$$

satisfies (1.63).

Proof. Let $\theta_i \stackrel{\text{def.}}{=} \theta(\lambda_i) \in \mathbb{T}$ for $i = 1, 2$ and $\alpha \stackrel{\text{def.}}{=} \theta(Z_\lambda(\mathbf{x})) \in \mathbb{T}$. We simply verify that

$$\begin{aligned} \lambda_1 e^{-2i\pi t_{0,1}} + \lambda_2 e^{-2i\pi t_{0,2}} &= \lambda_1 e^{4i\pi\theta_1} e^{-2i\pi(2\alpha-x_1)} + \lambda_2 e^{4i\pi\theta_2} e^{-2i\pi(2\alpha-x_2)} \\ &= e^{-4i\pi\alpha} (\overline{\lambda_1} e^{2i\pi x_1} + \overline{\lambda_2} e^{2i\pi x_2}) \\ &= e^{-4i\pi\alpha} \overline{Z_\lambda(\mathbf{x})} = Z_\lambda(\mathbf{x}) \end{aligned}$$

□

Let $\xi_0 = e^{-2i\pi t_0} \in \mathbb{T}^2$. In fact, when $Z_\lambda(\mathbf{x}) \neq 0$, $\xi = \zeta$ or $\xi = \xi_0$ are the only two solutions to (1.62).

Proposition 3. *Let $\xi \in \mathbb{U}^2$ satisfying (1.62). If $Z_\lambda(\mathbf{x}) \neq 0$, then $\xi = \zeta$ or $\xi = \xi_0$.*

Proof. Consider the system of equations

$$\begin{cases} \lambda_1 \xi_1 + \lambda_2 \xi_2 = Z_\lambda(\mathbf{x}) \\ |\xi_1| = 1 \\ |\xi_2| = 1 \end{cases} \quad (1.66)$$

where $\xi_1, \xi_2 \in \mathbb{C}$. Let $\delta : \xi \in \mathbb{C}^2 \mapsto \lambda_1 \xi_1 + \lambda_2 \xi_2$. Since δ is a non null linear form, the set $\{\xi \in \mathbb{C}^2 ; \delta(\xi) = Z_\lambda(\mathbf{x})\}$ is of the form

$$\{\zeta + z\mathbf{v} ; z \in \mathbb{C}\}$$

where $\delta(\zeta) = Z_\lambda(\mathbf{x})$ and $\mathbf{v} \in \text{Ker } \delta \setminus \{0\}$, *i.e.*

$$\mathbf{v} = t \begin{pmatrix} \lambda_2 \\ -\lambda_1 \end{pmatrix}$$

for some $t \in \mathbb{C} \setminus \{0\}$. Hence $\xi \in \mathbb{C}^2$ is a solution of (1.66) if and only if $\xi = \mathbf{a} + t\mathbf{v}$ with

$$\begin{cases} |\zeta_1 + z(t\lambda_2)| = 1 \\ |\zeta_2 - z(t\lambda_1)| = 1 \end{cases}$$

The point $z \in \mathbb{C}$ above is therefore in the intersection of two circles of centers $-\zeta_1/(t\lambda_2)$ and $\zeta_2/(t\lambda_1)$, and of radii $1/|t\lambda_2|$ and $1/|t\lambda_1|$, respectively. Note that if these circles have the same center, which corresponds to the case where $Z_\lambda(\mathbf{x}) = 0$, then they necessarily have the same radius and system (1.66) has an infinite number of solutions. Otherwise, it admits at most two solutions, and at least one since ζ always solves the system. If $\xi_0 \neq \zeta$, they are the two only solutions. □




Figure 1.5: Example of failure in the discrete case, for given $(\lambda_1, \lambda_2) \in \mathbb{C}^2$ (different values have been used for each of the case displayed). Points in all color but blue have been generated in pair such that they satisfy relation (1.67), leading to errors in the reconstruction. The method remains however partly robust, and sometimes manage to extract correct co-diagonalizing eigenvectors from X_λ despite the fact that it admits eigenspaces of dimension greater than one (in these synthetic examples, X_λ admits eigenvalues differing by less than 10^{-14}).

Thus, given $\lambda \in \mathbb{C}^2 \setminus \{0\}$, if the support of μ contains at least two points \mathbf{x}_j and \mathbf{x}_k satisfying

$$\mathbf{x}_k = h_\lambda(\mathbf{x}_j), \quad (1.67)$$

then diagonalizing $\lambda_1 X^{(1)} + \lambda_2 X^{(2)}$ may not suffice to jointly diagonalize $X^{(1)}$ and $X^{(2)}$. Figure 1.5 shows synthetic examples of failure in the discrete case. Note that for measures supported on a curve, the configuration (1.67) is very likely to be reached. Figure 1.6 shows examples of the image of a circle by the transformation h_λ .

1.6.3 Joint diagonalization by optimization

Rather than relying on a generic combination of the multiplication matrices, we propose to perform the joint diagonalization via an optimization scheme. Although this may be computationally more challenging, it is also more robust, and extends to cases where the multiplication matrices are only approximately jointly diagonalizable. We use this approximation scheme to perform recovery of non-sparse measures from their (non-flat) moment matrices.

Joint diagonalization criterion Let $X_1, \dots, X_d \in \mathcal{M}_s(\mathbb{C})$ be the matrices to co-diagonalize, of size $s \times s$. We consider the problem of retrieving an invertible matrix $T \in \text{GL}_s(\mathbb{C})$ (not necessarily orthogonal) minimizing the criterion

$$\mathcal{O}(T) = \sum_q \sum_{i \neq j} (TX_q T^{-1})_{ij}^2. \quad (1.68)$$

This diagonalization criterion is similar to the Froebenius criterion introduced in the seminal works [Flury and Gautschi, 1986, Cardoso and Souloumiac, 1993], for orthogonal transformations T . Unlike

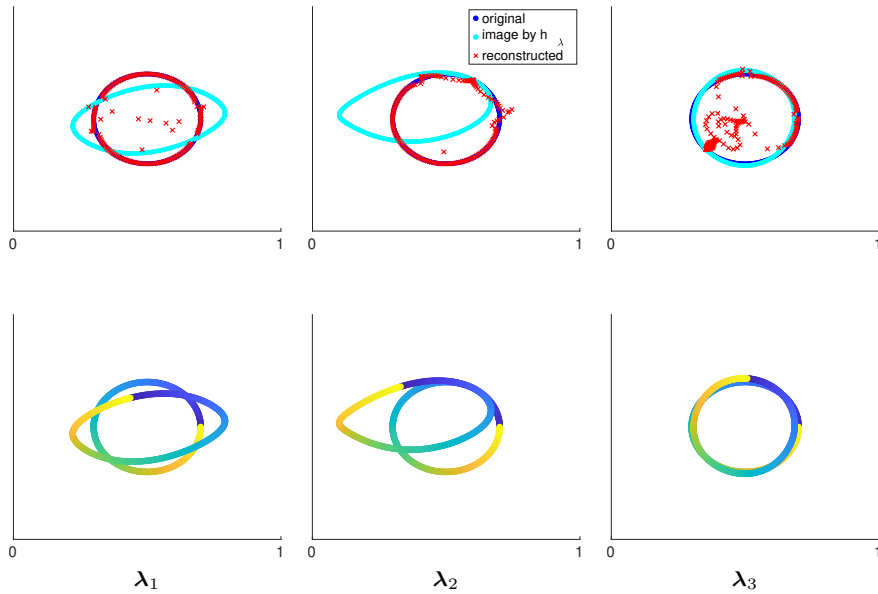


Figure 1.6: *Top row*: Reconstruction of a circle via multivariate Prony's algorithm (Algorithm 2), from $n = 100$ moments. The original measure is composed of 1000 Diracs. The joint diagonalization step is performed by random linear combination, for 3 random values λ_1, λ_2 and $\lambda_3 \in \mathbb{C}^2$. We represent in cyan the image of the circle by h_λ . We see that major errors occurs at the intersections of the two curves, which is expected. *Bottom row*: color display of the mapping $\mathbf{x} \mapsto h_\lambda(\mathbf{x})$. We see that some intersections correspond to cases where $h_\lambda(\mathbf{x}) = \mathbf{x}$, which does not result in a multiple eigenvalue for X_λ and hence does not affect the recovery, as observed.

usual joint diagonalization approaches (see Section 1.1.2), we do not assume that T is orthogonal, and consider transformations of the form TXT^{-1} rather than TXT^H . When the matrices X_q are co-diagonalizable, the solution to this problem is not unique, as given a solution T , any matrix $T' = \Lambda T$ for an invertible diagonal matrix Λ is also solution. Furthermore, there is always a global scale ambiguity, since any solution multiplied by a non null scalar remains a solution. These ambiguities may pose practical difficulties as they can lead to degenerate solutions. Typical approaches to remedy this problem consist in imposing further constraint on the solution, such as a specific normalization of its lines [Absil and Gallivan, 2006, Pham and Congedo, 2009, Tichavský and Yeredor, 2009]. Another solution, introduced in [Amari et al., 2000] in the context of blind source separation, is to exploit the geometry of the equivalence class $\{\Lambda T ; \Lambda \text{ diagonal and invertible}\}$ to cancel the action of diagonal scaling.

Invertibility may be enforced during the iteration of the algorithm by using multiplicative updates on T , taking the form

$$T_{k+1} = (I + \mathcal{E}_k)T_k. \quad (1.69)$$

Such updates are common in most blind source separation algorithms [Cardoso and Laheld, 1996, Amari, 1998, Akuzawa and Murata, 2001]. They may be interpreted as Riemannian descent steps over $\text{GL}_s(\mathbb{C})$, endowed with a right-invariant metric [Amari, 1998]. The first order Taylor expansion of $\mathcal{O}((I + \mathcal{E})T)$ yields

$$\mathcal{O}((I + \mathcal{E})T) = \mathcal{O}(T) + \langle G(T), \mathcal{E} \rangle + o(\|\mathcal{E}\|) \quad (1.70)$$

where $G(T) \in \mathcal{M}_s(\mathbb{C})$ is the *relative gradient* at T [Cardoso and Laheld, 1996]. Note that $G(T)$ is linked to the standard gradient by the simple relation

$$G(T) = \nabla \mathcal{O}(T)T^H.$$

Proposition 4. *Given $X \in \mathcal{M}_s(\mathbb{C})$, let $\underline{X} \stackrel{\text{def.}}{=} X - \text{Diag}(X)$. Then*

$$G(T) = \sum_q \underline{Y}_q Y_q^H - Y_q^H \underline{Y}_q \quad (1.71)$$

where $Y_q \stackrel{\text{def.}}{=} TX_q T^{-1}$.

Proof. For $T \in \text{GL}_s(\mathbb{C})$, one has $\mathcal{O}(T) = \sum_q \|TX_q T^{-1}\|^2 - \sum_q \sum_i (TX_q T^{-1})_{ii}^2$. Using the expansion $(I + \mathcal{E})^{-1} = I - \mathcal{E} + o(\|\mathcal{E}\|)$ yields on the one hand

$$\|(I + \mathcal{E})Y_q(I - \mathcal{E} + o(\|\mathcal{E}\|))\|^2 = \|Y_q\|^2 + 2\langle Y_q Y_q^H - Y_q^H Y_q, \mathcal{E} \rangle + o(\|\mathcal{E}\|).$$

On the other, we have

$$\sum_i [(I + \mathcal{E})Y_q(I - \mathcal{E} + o(\|\mathcal{E}\|))]_{ii}^2 = \sum_i (Y_q)_{ii}^2 + 2\langle \text{Diag}(Y_q)Y_q^H - Y_q^H \text{Diag}(Y_q), \mathcal{E} \rangle + o(\|\mathcal{E}\|)$$

Thus,

$$\mathcal{O}((I + \mathcal{E})T) = \mathcal{O}(T) - \langle 2 \sum_q (\underline{Y}_q Y_q^H - Y_q^H \underline{Y}_q), \mathcal{E} \rangle + o(\|\mathcal{E}\|).$$

□

The algorithm To minimize $\mathcal{O}(T)$, we perform a simple gradient descent algorithm, with descent direction $-G(T)$, using the multiplicative updates (1.69). As our descent direction does not guarantee that the iteration $T \leftarrow (I - G(T))T$ decreases the criterion, we use a backtracking linesearch to find a step $\alpha > 0$ ensuring that $\mathcal{O}((I - \alpha G(T))T) \leq \mathcal{O}(T)$. This base algorithm may be improved using Hessian pre-conditioning and conjugate descent directions, which is the subject of ongoing work in collaboration with Jean-François Cardoso. In this dissertation, we keep the exposition focused on the main ideas of the approach, and do not consider these additional implementation details. The main steps of the algorithm are summarized in Algorithm 4.

Algorithm 4: Approximate joint diagonalization

Input: $X_1, \dots, X_d \in \mathcal{M}_s(\mathbb{C})$, number of iterations N

- 1 Initialize $T_0 = I_s$
- 2 **for** $k = 1, \dots, N$ **do**
- 3 compute the descent direction $\delta_k = -G(T_k)$, where G is the relative gradient (1.71)
- 4 do a backtracking linesearch $\alpha_k = \operatorname{argmin}_\alpha \mathcal{O}((I + \alpha \delta_k + \frac{\alpha^2}{2} \delta_k^2)T_k)$
- 5 update $T_{k+1} \leftarrow (I + \alpha_k \delta_k + \frac{\alpha_k^2}{2} \delta_k^2)T_k$
- 6 **end**

Output: T_{N+1}

1.6.4 Numerical illustrations

One of the main advantages of performing the joint diagonalization step via optimization is that it may still be used in cases where the multiplication matrices are not co-diagonalizable – whereas this assumption is necessary when resorting to the random combination approach. We propose to perform the joint diagonalization step 6 of Prony’s algorithm 2 with the method introduced in previous section, all other steps remaining identical. We present in this section results obtained with our novel approach for the recovery of non-sparse measures from their moments.

Singular values threshold Due to the absence of spectral gap in moment matrices deriving from non-sparse measures (cf. Section 1.6.1), it is necessary to threshold their singular value decomposition in order to apply Prony’s algorithm 2. In our implementation, given a tolerance ε , and the singular values

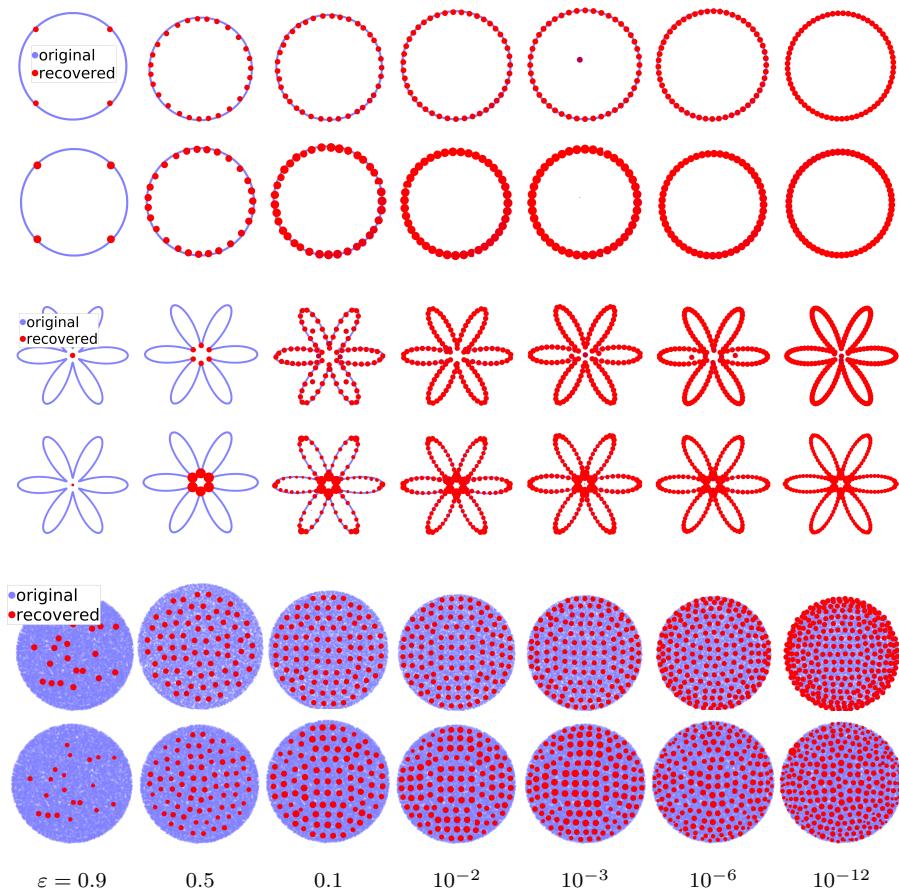


Figure 1.7: Examples of recovery for different thresholds on the singular values of the moment matrix, displayed each time with (2nd row) and without (1st row) amplitudes. The recovery is performed from $n = 25$ moments. Moments are computed from 5000 Diracs for the circle and treffle, and 20000 Diracs for the disc. Each recovered point is computed from two eigenvalues of the respective multiplication matrices; for each point, the color is proportional to the mean of the moduli of these eigenvalues (red is close to 1). There can be outliers in the reconstruction, but we see on the amplitude-less plots that these may be identified using this modulus criterion. The outliers have an almost null amplitude in the final reconstruction.

$\sigma_1, \dots, \sigma_s$, any singular values σ such that $\sigma \leq \varepsilon \max(\sigma_1, \dots, \sigma_s)$ is put to zero. The impact of the choice of threshold ε is shown in Figure 1.7.

Comparison with Christoffel polynomials Given a positive measure $\mu \in \mathcal{M}_+(\mathbb{T}^d)$, the Christoffel polynomial associated to μ is built from its moments and provides a lot of relevant information on its support. It is defined as follows.

Definition 6 (Christoffel polynomial). *Let $\mathbf{c} \in \mathbb{C}^{N(n)}$ be the vector of moments of μ up to order n , and let*

$$\psi_n(\mathbf{x}) \stackrel{\text{def.}}{=} (e^{-2i\pi\langle \mathbf{k}, \mathbf{x} \rangle})_{\mathbf{k} \in \Omega_n^+} \in \mathbb{C}^{N_+(n)} \quad (1.72)$$

be the vector of monomials of degree at most n (sorted in the colexicographical order). Assuming that $T_n(\mathbf{c})$ is positive definite, the Christoffel polynomial of degree n associated to μ is given by

$$Q_{\mathbf{c},n}(\mathbf{x}) \stackrel{\text{def.}}{=} \psi_n(\mathbf{x})^H (T_n(\mathbf{c}))^{-1} \psi_n(\mathbf{x}). \quad (1.73)$$

Its inverse $\Lambda_{\mathbf{c},n} = (Q_{\mathbf{c},n})^{-1}$ is called Christoffel function.

Note that in 1-D, for a s -sparse measure, the polynomial $Q_{\mathbf{c},s}$ is the same as Prony's polynomial (see Remark 4). Christoffel polynomials, or functions, associated to μ have the remarkable property of encoding with good accuracy the support of μ in their level sets: given μ supported on a set $S \subset \mathbb{R}^d$, with moments \mathbf{c} , $\alpha_n \in (\mathbb{R}_+)^{\mathbb{N}}$, and defining $S_n \stackrel{\text{def.}}{=} \{\mathbf{x} \in \mathbb{R}^d ; \Lambda_{\mathbf{c},n}(\mathbf{x}) \geq \alpha_n\}$, Theorem 3.12 in [Lasserre and Pauwels, 2019] shows that, under mild assumptions on S and a well-chosen sequence α_n , one has $\lim_{n \rightarrow \infty} d_H(S_n, S) = 0$, where d_H is the Hausdorff distance, *i.e.*

$$d_H(X, Y) \stackrel{\text{def.}}{=} \max \left\{ \sup_{x \in X} \inf_{y \in Y} d(x, y), \sup_{y \in Y} \inf_{x \in X} d(x, y) \right\}, \quad \forall X, Y \subset \mathbb{R}^d.$$

In practice, as the matrix $T_n(\mathbf{c})$ is often non invertible, it is common to regularize the inversion, for instance by adding a small positive constant to all eigenvalues of $T_n(\mathbf{c})$. Figure 1.8 gives an illustration of the convergence of the Christoffel polynomial as n increases. Note that while Christoffel polynomials provide an “outer” approximation of the sought support, giving access to a larger set, our extraction procedure on the other hand gives an “inner”, sparse approximation, returning a measure whose support is more localized.

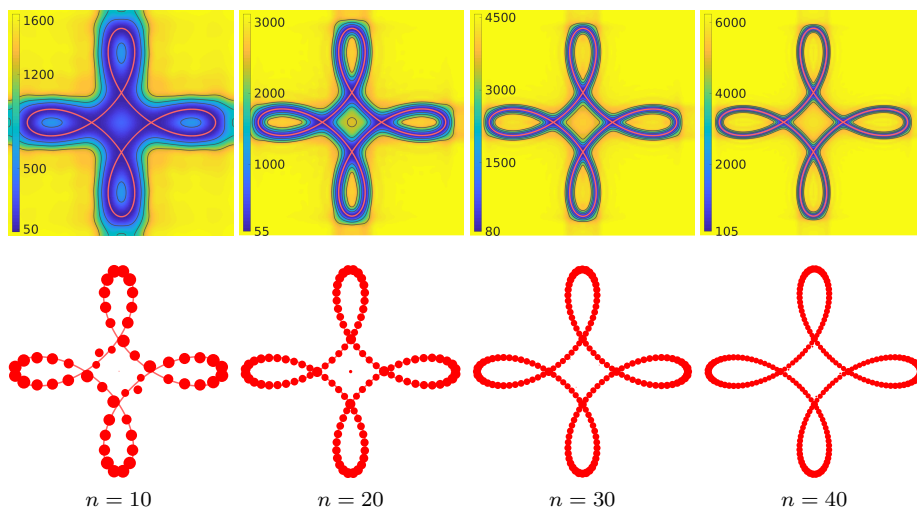


Figure 1.8: Christoffel polynomial versus Prony reconstruction. The moment matrix T_n is computed from a discrete approximation of the measure composed of 5000 spikes, of equal amplitudes $1/5000$. For Prony's extraction, the tolerance for the svd is set to 10^{-5} .

Chapter 2

Semidefinite relaxations for Beurling lasso

Abstract

From the sole knowledge of the low frequency trigonometric moments of a Radon measure, Prony's method gives a straightforward procedure to recover exactly its support and amplitudes. In practice though, one must often deal with much less favorable situations: the frequencies may be randomly sampled, like in compressed sensing, the observations may be corrupted by noise, or most importantly one may have access to other types of linear observations, not based on Fourier measurements. Although extensions of Prony's method exist to deal with these settings, they are specific to each type of measurement and might be difficult to implement. In contrast, variational approaches bring more modelling flexibility, at the cost of a heavier computational burden, as such methods rely on optimizing over the Banach space of Radon measures. We focus in this chapter on a spectral approach for solving the BLASSO, known as Lasserre's hierarchy, that recovers the moments of the sought measure via a hierarchy of semidefinite programs. The main contribution is a relaxation of the forward operator as a spectral operator, which allows to approximate the original hierarchy to arbitrary precision, while enabling the use of fast optimization solvers, detailed in the following chapter.

Contents

2.1 Introduction	65
2.1.1 The Beurling LASSO	65
2.1.2 Off-the-grid approaches	67
2.1.3 Contributions	69
2.2 Theoretical background on the Blasso	69
2.2.1 Radon measures and total variation	69
2.2.2 Dual problem and certificates	70
2.3 Spectral approximation of sensing operators	71
2.3.1 Spectral approximation operator	72
2.3.2 Examples	75
2.3.3 Discussion	78
2.4 Lasserre’s hierarchy for the Blasso	78
2.4.1 Problem statement	79
2.4.2 Moment relaxations for the BLASSO	80
2.4.3 Finite convergence	83
2.4.4 Dual sum-of-squares strengthening	86

The super-resolution problem, which we solve using Prony’s type methods in Chapter 1, may alternatively be formulated as an optimization problem over the space of Radon measures, whose numerical resolution may be challenging. One possible approach, that we investigate in this chapter, is to resort to Lasserre’s hierarchy, that allows to recover the moments of the sought measure.

2.1 Introduction

Parametric methods such as Prony’s method require strong structural assumptions on the measurements, and many of the aforementioned algorithms (cf 0.2 and chapter 1) may only be applied when the observations indeed consist of low-frequency Fourier moments of the signal to recover. Even though some extensions to slightly more general frameworks have been proposed [Peter et al., 2015], they remain limited to deconvolution problems.

Variational approaches on the other hand are much more flexible as they do not require any particular structure on the forward operator, and adapt easily to any type of noise.

2.1.1 The Beurling Lasso

Variational problems search for their solutions as the minimizers of some energy, combining a data fidelity term, which ensures that the recovered signal matches the observations, and a regularization term, which induces some structure on the solution, and may remedy the ill-posedness of the general problem.

Broadly speaking, given a vector space E (e.g. the space of Radon measures $\mathcal{M}(\mathbb{T}^d)$) and a sensing operator $\Phi : E \rightarrow \mathbb{C}^n$, $n \in \mathbb{N}$, they take the form

$$\inf_{v \in E} f(\Phi v) + R(v) \quad (2.1)$$

where $R : E \rightarrow \mathbb{R} \cup \{\infty\}$ is a convex and lower semi-continuous functional (the regularizer), and $f : E \rightarrow \mathbb{R}$ is arbitrary, but usually smooth (the data fidelity term). In all this thesis, f measures the quadratic error between the model Φv and the observations.

The analysis of programs of the form (2.1), as well as their solvers, differ depending on whether E is a Hilbert space or a Banach space. The Hilbertian setting is ubiquitous in image processing and machine learning, celebrated examples being Tikhonov regularization or ℓ^1 -optimization. Optimization problems over Banach spaces on the other hand have recently gained a strong interest because of so-called off-the-grid methods, which allow to solve (2.1) with very high accuracy.

Sparse super-resolution in Hilbert spaces, Lasso. The use of convex methods for recovering sparse signals on a grid has a long history, dating back to the seventies. It was pioneered by geophysicists [Claerbout and Muir, 1973, Levy and Fullagar, 1981, Santosa and Symes, 1986], who made use of the ℓ^1 -norm to extract from seismograms the locations of underground density changes, which may be modeled as sparse sums of Dirac masses. ℓ^1 -optimization then gained a lot of popularity in signal processing as well as in statistics thanks to the works of Donoho [Donoho, 1992, Chen et al., 2001] on Basis Pursuit and Tibshirani [Tibshirani, 1996] on the LASSO. Given linear measurements of the form $y = Ax_0 + w \in \mathbb{R}^m$, where $A \in \mathbb{R}^{m \times n}$, x_0 is an unknown sparse vector of \mathbb{R}^n (i.e. with few non-zero entries) and $w \in \mathbb{R}^m$ some random noise, the LASSO searches for x_0 as the solution of the following program,

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|Ax - y\|_2^2 + \lambda \|x\|_1, \quad (2.2)$$

where $\lambda \in \mathbb{R}_+$ is a parameter that should be adapted to the noise level. The vector x_0 may be interpreted as the sparse measure $\sum_{k \in \mathcal{G}} x_{0,k} \delta_k$, for some grid $\mathcal{G} \subset \mathbb{R}^d$, so that the LASSO may only recover signals over this pre-defined grid.

When A satisfies some uncertainty principle – the so-called Restricted Isometry Property – the theory of compressed sensing ensures that (2.2) stably recovers a good estimate of x_0 [Candès et al., 2006], even in regimes where $n \gg m$. The super-resolution picture is different since in that case A models a low-pass filter and may be ill-conditioned, even more so as the grid stepsize Δ goes to zero [Donoho, 1992], which is unfortunately the setting of interest if one seeks to recover extremely close spikes with high precision. Thus, computing the LASSO over thin grids, beyond being a numerical challenge, leads to well-understood but unavoidable mismatches between the sought signal and its estimate [Chi et al., 2011, Stoica and Babu, 2012]. In particular, in a low noise regime, the LASSO recovers twice as many spikes as the original number [Duval and Peyré, 2017].

Extension to Radon measures, Blasso. These instabilities drove researchers to consider more general optimization problems, defined over the space of Radon measures. Unifying a wide range of inverse problems under this generic Banach setting offers beneficial mathematical insights, and has led to a better understanding of recovery algorithms [Bredies and Pikkarainen, 2013]. In particular, these off-the-grid alternatives triggered major breakthroughs in super-resolution [Candès and Fernandez-Granda, 2014] or compressed sensing [Tang et al., 2013].

Let $\mu_0 = \sum_{k=1}^r a_{0,k} \delta_{x_{0,k}} \in \mathcal{M}(\mathbb{T}^d)$, where $a_{0,k} \in \mathbb{C}^*$ and $x_{0,k} \in \mathbb{T}^d$, be a (bounded) discrete Radon measure over the d -dimensional torus. We assume that we are given linear measurements y in a separable Hilbert space \mathcal{H} , of the form $y = \Phi \mu_0 + w \in \mathcal{H}$, where $\Phi : \mathcal{M}(\mathbb{T}^d) \rightarrow \mathcal{H}$ is a known linear operator and $w \in \mathcal{H}$ some unknown noise. The infinite-dimensional analogue of the LASSO, introduced in [Azaïs et al., 2015] as BLASSO, searches for sparse estimates by solving

$$\inf_{\mu \in \mathcal{M}(\mathbb{T}^d)} \frac{1}{2} \|\Phi \mu - y\|_2^2 + \lambda |\mu|(\mathbb{T}^d) \quad (2.3)$$

where the total variation norm $|\mu|(\mathbb{T}^d)$ of a measure $\mu \in \mathcal{M}(\mathbb{T}^d)$ is defined as

$$|\mu|(\mathbb{T}^d) \stackrel{\text{def.}}{=} \sup \left\{ \Re \left(\int_{\mathbb{T}^d} \bar{\eta} d\mu \right) ; \eta \in \mathcal{C}(\mathbb{T}^d), \quad \|\eta\|_\infty \leq 1 \right\} \quad (2.4)$$

where \Re denotes the real part. It extends the ℓ^1 -norm to the infinite-dimensional space of measures, favoring in particular the emergence of Dirac masses in the solution. The BLASSO has been extensively studied in the last years. Its recovery guarantees are well understood for uni-dimensional measures [Candès and Fernandez-Granda, 2014, Duval and Peyré, 2015, Denoyelle et al., 2017], and several breakthroughs have been made in understanding its properties in higher dimensions [Unser et al., 2017, Flinth and Weiss, 2017, Poon and Peyré, 2018, Boyer et al., 2019].

2.1.2 Off-the-grid approaches

While optimizing over Banach spaces allows for a sharper analysis, it is also more difficult to implement. Most of the time in practice, the search space is discretized beforehand, so that the problem reduces to finite-dimensional optimization. Nonetheless, this approach is subject to the aforementioned accuracy limitations, and many efforts are put in developing off-the-grid algorithms. These may be divided into two general categories.

Particle approaches. Instead of resorting to a fixed spatial discretization of the whole optimization domain in (2.3), one may rather discretize the sought measure and operate on its atoms directly over the continuous domain. Note that optimizing on the positions of the particles is a non-convex problem. It is shown in [Chizat and Bach, 2018, Chizat, 2019] that starting with many

particles and performing a non-convex gradient descent achieves global minimization nonetheless. On the other hand, several approaches rather focus on greedily updating the support at each iteration, using *e.g.* conditional gradient algorithms [Boyd et al., 2015, Denoyelle et al., 2020] or exchange algorithms [Flinth et al., 2019].

Moment-sos approaches. Another important line of work for solving (2.3) or its dual consist in lifting them to semidefinite programs. Semidefinite approaches for the total variation minimization problem were originally introduced in [Candès and Fernandez-Granda, 2014, Tang et al., 2013], but for unidimensional measures (*i.e.* $d = 1$) only. The multivariate case on the other hand raises a more challenging problem, which may be solved using Lasserre’s hierarchy [Lasserre, 2001], consisting in a sequence of increasingly better semidefinite approximations of (2.3) or its dual [de Castro et al., 2017, Dumitrescu, 2017].

Lasserre’s hierarchy is an approximation scheme that was originally introduced for finding global optima in polynomial optimization problems [Lasserre, 2001]. The approach has two aspects, that are related by Fenchel duality: the moment and the sum-of-squares (SOS) approach. Both rely on results from algebraic geometry.

Moment approach. The moment approach targets problems that can be put under the general form

$$\inf_{\mu \in \mathcal{M}_+(K)} \int_K f d\mu \quad \text{s.t.} \quad \int_K h_j d\mu = \gamma_j, \quad j \in \Gamma, \quad (2.5)$$

known as the generalized (polynomial) moment problem [Lasserre, 2010], where K may be a subset of \mathbb{T}^d , Γ is a (usually finite) set of indices, f, h_j are polynomials and γ_j are complex numbers. Many well-known problems fit into this framework, including the BLASSO. When K is a semi-algebraic set, that is a set defined by polynomial inequalities, it becomes natural to approximate the positive measure μ by a finite number of its moments. While the objective and constraints in (2.5) are then readily transformed, expressing which sequences are the moments of a positive measure is much more involved, as was already discussed in Chapter 1. Using characterizations of such moment sequences, that are notably due to Schmügdén [Schmügdén, 1991] or Curto and Fialkow [Curto and Fialkow, 1991, Curto and Fialkow, 1996], the moment hierarchy approximates this constraint by *necessary* semidefinite conditions on the sequences. Involving sequences of larger and larger size, one thus builds a hierarchy of semidefinite programs with higher and higher accuracy, in spaces of higher and higher dimensions. Each semidefinite program involved in the hierarchy is a relaxation of the initial problem (its feasible set is larger).

SOS approach. In mirror of the moment hierarchy, the SOS approach approximates the solutions of the dual to (2.5), given by

$$\sup \langle \gamma, \lambda \rangle \quad \text{s.t.} \quad \sum \lambda_j h_j \geq f,$$

by replacing the polynomial nonnegativity constraint with a *sufficient* certificate of positivity, *a.k.a.* Positivstellensatz, which involves sum-of-squares representations, see *e.g.* [Putinar, 1993, Jacobi and Prestel, 2001]. Again, these strengthened SOS conditions take the form of semidefinite constraints on the coefficients of the involved SOS polynomials, and the hierarchy consists in imposing degree bounds on these polynomials. The resulting SDP are strengthenings of the initial problem (their feasible sets are smaller), and can be shown to be the dual problem associated to the SDP relaxation (weakening) of the primal moment problem.

2.1.3 Contributions

The first contribution of this chapter, detailed in section 2.3, is a spectral approximation scheme of the forward operator. To any operator, we associate a spectral operator (Definition 8), which only acts on the first n Fourier coefficients of the measure. We prove that the corresponding minimizer for the BLASSO weakly converges toward the minimizer for the original operator. This new scheme allows to fit a broad range of sensing models into the semidefinite relaxation framework, and gives a common methodology to address a large variety of measurements. In particular, it provides a novel means to tackle non-translation-invariant observations.

In section 2.4 we develop a framework of Lasserre-inspired semidefinite liftings for the BLASSO, which generalizes the approach of [Tang et al., 2013] to multidimensional settings. Following classical results from the theory of Lasserre’s hierarchy, we prove in section 2.4.2 that our semidefinite relaxations indeed define a hierarchy of programs, that converges in energy towards the true solution (Proposition 6), in arbitrary dimension. We establish in section 2.4.3 precise connections between the solutions of these semidefinite programs and the polar decomposition of the measure solving the BLASSO (see Propositions 7 and 8). In doing so, we show that the semidefinite relaxations admit solutions of low rank, of which we take advantage in chapter 3.

2.2 Theoretical background on the Blasso

We recall in this section some useful and well-known notions about the Blasso and its dual problem.

2.2.1 Radon measures and total variation

Since \mathbb{T}^d is compact, the space of bounded Radon measures $\mathcal{M}(\mathbb{T}^d)$ may be seen as the topological dual of the space $\mathcal{C}(\mathbb{T}^d)$ of (complex-valued) continuous functions on \mathbb{T}^d , equipped with the uniform norm. It is a Banach space when endowed with the total variation norm (2.4). Note however that the dual space of $(\mathcal{M}(\mathbb{T}^d), |\cdot|(\mathbb{T}^d))$ is not $(\mathcal{C}(\mathbb{T}^d), \|\cdot\|_\infty)$, but a larger space.

The total variation (2.4) is convex and lower semi-continuous with respect to the weak-* topology. It is the dual norm of $\|\cdot\|_\infty$, and it naturally extends the ℓ^1 -norm to the space of Radon measures. For instance, if $\mu = \sum a_j \delta_{\mathbf{x}_j}$, then $|\mu|(\mathbb{T}^d) = \|a\|_1$. Similarly, if μ admits a density f with respect to the Lebesgue measure, then $|\mu|(\mathbb{T}^d) = \|f\|_1$. The subdifferential of the total variation norm reads

$$\partial|\mu|(\mathbb{T}^d) = \left\{ \eta \in \mathcal{C}(\mathbb{T}^d) ; \|\eta\|_\infty \leq 1 \quad \text{and} \quad \Re \left(\int_{\mathbb{T}^d} \bar{\eta} d\mu \right) = |\mu|(\mathbb{T}^d) \right\}. \quad (2.6)$$

In particular, for a discrete measure $\mu_{a,\mathbf{x}} = \sum a_j \delta_{\mathbf{x}_j}$, one has

$$\partial|\mu_{a,\mathbf{x}}|(\mathbb{T}^d) = \left\{ \eta \in \mathcal{C}(\mathbb{T}^d) ; \|\eta\|_\infty \leq 1 \quad \text{and} \quad \eta(\mathbf{x}_j) = \frac{a_j}{|a_j|} \quad \forall j \right\}.$$

2.2.2 Dual problem and certificates

A lot of information on the solutions of (2.3) is carried by its dual problem. In the case of ideal low-pass filter, the dual approach is also the key to the first off-the-grid numerical algorithms used in [Bhaskar and Recht, 2011, Candès and Fernandez-Granda, 2014, Azaïs et al., 2015]. The Fenchel dual problem to (2.3) is given by

$$\max_{p \in \mathcal{H}} \Re \langle y, p \rangle - \frac{\lambda}{2} \|p\|_2^2 \quad \text{s.t.} \quad \|\Phi^* p\|_\infty \leq 1 \quad (2.7)$$

Contrarily to the primal problem, the solution of (2.7) is always unique. We denote it by p_λ . Note that while the numerical difficulty in (2.3) comes from the infinite-dimensionality of the space of Radon measures, in (2.7) it is the infinite number of constraints that poses a numerical challenge.

Dual certificates Strong duality holds between (2.3) and (2.7) [Bredies and Pikkarainen, 2013]. In particular, any solution μ_λ of (2.3) is linked to the unique solution p_λ of (2.7) by the extremality relations

$$\begin{cases} \Phi^* p_\lambda \in \partial|\mu_\lambda|(\mathbb{T}^d) \\ \lambda p_\lambda = y - \Phi \mu_\lambda \end{cases}. \quad (2.8)$$

The vector $\eta_\lambda \stackrel{\text{def.}}{=} \Phi^* p_\lambda$ is called a *dual certificate* of μ_λ [Candès and Plan, 2010]. This terminology is justified by the fact that when $\lambda = 0$, finding a vector p satisfying $\Phi^* p \in \partial|\mu_0|(\mathbb{T}^d)$ – such a vector is not unique – gives a quick proof that μ_0 is a solution of the constrained problem

$$\inf_{\mu \in \mathcal{M}(\mathbb{T}^d)} |\mu|(\mathbb{T}^d) \quad \text{s.t.} \quad \Phi \mu = y_0. \quad (2.9)$$

More generally, η_λ informs significantly on the support of μ_λ , as shown in the next lemma.




Figure 2.1: Examples of dual certificates for the Dirichlet kernel, with cutoff frequency $f_c = 15$ in 1-D (left) and $f_c = 8$ in 2-D (right). With this kernel, the dual certificate η_λ is a trigonometric polynomial.

Lemma 12. *Let $S(p) = \{\mathbf{x} \in \mathbb{T}^d ; |\Phi^*p(\mathbf{x})| = 1\}$. Then $\text{Supp } \mu_\lambda \subset S(p_\lambda)$.*

Proof. The polar decomposition lemma for complex measures ensures that there exists a measurable function $\xi : \mathbb{T}^d \rightarrow \mathbb{C}$ such that $|\xi| = 1$ and $\mu = \xi|\mu|$. From (2.6) and (2.8) we get that

$$0 = |\mu_\lambda|(\mathbb{T}^d) - \Re \left(\int_{\mathbb{T}^d} \overline{\eta_\lambda} d\mu_\lambda \right) = \int_{\mathbb{T}^d} (1 - \Re(\overline{\eta_\lambda}\xi)) d|\mu_\lambda|$$

Since $\|\eta_\lambda\|_\infty \leq 1$ and $|\xi| = 1$, it follows that $1 \geq \Re(\overline{\eta_\lambda}\xi)$. Therefore $\Re(\overline{\eta_\lambda}\xi) = 1$ μ_λ -a.e., and thus $|\eta_\lambda| = 1$ μ_λ -a.e. since $1 \geq |\eta_\lambda|^2 = |\overline{\eta_\lambda}\xi|^2 \geq \Re(\overline{\eta_\lambda}\xi)^2$. \square

Robustness to noise. The Fenchel dual of problem (2.9) reads

$$\sup \Re\langle y_0, p \rangle \quad \text{s.t.} \quad \|\Phi^*p\|_\infty \leq 1. \quad (2.10)$$

A fundamental matter in the theory of sparse super-resolution is to understand how the solutions μ_λ of (2.3) behave with respect to the solutions μ_0 of (2.9) in regimes where the noise w and the parameter λ are small. The answer to this question is linked to the important notion of minimal norm certificate, defined as $\eta_0 \stackrel{\text{def.}}{=} \Phi^*p_0$ where

$$p_0 = \underset{p}{\text{argmin}} \{ \|p\|_2 ; p \text{ is a solution of (2.10)} \}, \quad (2.11)$$

It is known, see *e.g.* [Duval and Peyré, 2015, Denoyelle et al., 2017], that this particular certificate governs the stability of the solutions of (2.3), in the sense for instance that if it satisfies some non-degeneracy hypothesis – that we do not detail here as it goes beyond our purpose, then, in a regime of low noise and small λ , μ_λ is actually close to μ_0 and share many of its structural properties.

2.3 Spectral approximation of sensing operators

In this section we detail a spectral approximation method for the forward operator Φ , that is consistent with semidefinite approaches (see Section 2.4)

to solve (2.3). Our model encompasses deconvolution problems, as well as non-convolution ones – in particular, we focus on the case of subsampled convolution and spatially-varying filters (see Section 2.3.2), which are common in imaging problems. This approach is possible thanks to the use of a variational formulation, and gives a new way of addressing non-translation-invariant problems.

Sensing operators The variational problem (2.3) may be solved for any linear operator $\Phi : \mathcal{M}(\mathbb{T}^d) \rightarrow \mathcal{H}$ of the form

$$\Phi : \mu \in \mathcal{M}(\mathbb{T}^d) \mapsto \int_{\mathbb{T}^d} \varphi(x) d\mu(x), \quad (2.12)$$

where the kernel $\varphi : \mathbb{T}^d \rightarrow \mathcal{H}$ is assumed to be smooth. We denote by \mathcal{L} the set of such operators. This encompasses convolution operators, where $\varphi(x) = \tilde{\varphi}(\cdot - x)$, for $\tilde{\varphi} \in L^2(\mathbb{T}^d)$ and $\mathcal{H} = L^2(\mathbb{T}^d)$. When the Fourier transform of $\tilde{\varphi}$ is finitely supported on a set Ω , this is simply modeled in the Fourier domain by taking

$$\varphi(x) = (\hat{\tilde{\varphi}}(\omega) e^{-2i\pi\langle \omega, x \rangle})_{\omega \in \Omega}$$

with $\mathcal{H} = \mathbb{C}^{|\Omega|}$. For instance, ideal low-pass filtering with cutoff frequency $f_c \in \mathbb{N}$, as considered in [Candès and Fernandez-Granda, 2014], is modeled as $\varphi(x) = (e^{-2i\pi\langle k, x \rangle})_{k \in \Omega_c}$ where

$$\Omega_c \stackrel{\text{def.}}{=} \llbracket -f_c, f_c \rrbracket^d,$$

and $\mathcal{H} = \mathbb{C}^{|\Omega_c|}$, since it corresponds to a convolution with the Dirichlet kernel $\tilde{\varphi}_D(x) \stackrel{\text{def.}}{=} \sum_{k \in \Omega_c} e^{2i\pi\langle k, x \rangle}$. In that case, Φ is the Fourier operator \mathcal{F}_c , *i.e.*

$$\mathcal{F}_c : \mu \mapsto (c_k(\mu))_{k \in \Omega_c}, \quad (2.13)$$

where as before $c_k(\mu)$ denotes the k -th coefficient Fourier of μ .

Non-translation-invariant operators are also commonly encountered in imaging problems. Subsampled convolution for example, given a sampling grid $\mathcal{G} \subset \mathbb{T}^d$, may be modeled as $\varphi(x) = (\tilde{\varphi}(t - x))_{t \in \mathcal{G}}$, and $\mathcal{H} = \mathbb{C}^{|\mathcal{G}|}$. This is a good framework for the single molecule localization microscopy data that we consider in chapter 3. More generally, we may consider a kernel of the form $\varphi(x) = (\tilde{\varphi}(t, x))_{t \in \mathcal{G}}$. This allows for instance to deal with foveated measurements [Chang et al., 2000], see section 2.3.2.

2.3.1 Spectral approximation operator

Let $\Phi \in \mathcal{L}$ be an operator of the form (2.12), with kernel φ (in the following, we assume that φ is sufficiently smooth, namely $\varphi \in \mathcal{C}^j(\mathbb{T}^d; \mathcal{H})$, with $j \geq \lfloor \frac{d}{2} \rfloor + 1$). It is possible to define a Fourier series expansion of φ , using the coefficients

$$\forall k \in \mathbb{Z}^d, \quad c_k(\varphi) \stackrel{\text{def.}}{=} \int_{\mathbb{T}^d} e^{-2i\pi\langle k, x \rangle} \varphi(x) dx \in \mathcal{H}. \quad (2.14)$$

We refer to the paper of Kandil [Kandil, 1983] for the theory of the Fourier series of Hilbert-valued functions. In particular, as \mathcal{H} is separable, it is shown that the Parseval formula holds, *i.e.* for all $\psi \in L^2(\mathbb{T}^d; \mathcal{H})$,

$$\int_{\mathbb{T}^d} \|\psi(x)\|_{\mathcal{H}}^2 dx = \sum_{k \in \mathbb{Z}^d} \|c_k(\psi)\|_{\mathcal{H}}^2. \quad (2.15)$$

As a consequence, the following equality holds strongly in $L^2(\mathbb{T}^d; \mathcal{H})$,

$$\psi = \sum_{k \in \mathbb{Z}^d} c_k(\psi) e_k, \quad \text{where } e_k : x \mapsto e^{2i\pi \langle k, x \rangle}. \quad (2.16)$$

Now, given $f_c \in \mathbb{N}$, we are interested in approximating Φ with some operator Φ_c whose kernel φ_c has spectrum supported on Ω_c , *i.e.*

$$\Phi_c \in \mathcal{L}_c \stackrel{\text{def.}}{=} \{ \Psi \in \mathcal{L} ; c_k(\Psi) = 0 \quad \forall k \in \mathbb{Z}^d \setminus \Omega_c \}.$$

Definition 7 (Spectral approximation operator). *The spectral approximation operator of Φ is the operator $\Phi_c \in \mathcal{L}_c$ defined by*

$$c_k(\varphi_c) = \begin{cases} c_k(\varphi) & \forall k \in \Omega_c \\ 0 & \forall k \in \mathbb{Z}^d \setminus \Omega_c \end{cases}.$$

From Parseval's equality (2.15), we see that Φ_c is the best approximation of Φ in \mathcal{L}_c in terms of the Hilbert-Schmidt operator norm, $\Psi \mapsto \|\Psi\| \stackrel{\text{def.}}{=} \|\psi\|_{L^2(\mathbb{T}^d; \mathcal{H})}$, *i.e.*

$$\Phi_c = \operatorname{argmin}_{\Psi \in \mathcal{L}_c} \|\Phi - \Psi\|^2 = \operatorname{argmin}_{\Psi \in \mathcal{L}_c} \int_{\mathbb{T}^d} \|\varphi(x) - \psi(x)\|_{\mathcal{H}}^2 dx. \quad (2.17)$$

The next result shows that the solutions of the BLASSO when replacing Φ with Φ_c approximate the solutions of (2.3).

Proposition 5. *Let $\Phi \in \mathcal{L}$, with $\varphi \in \mathcal{C}^j(\mathbb{T}^d; \mathcal{H})$, $j \geq \lfloor \frac{d}{2} \rfloor + 1$. For each $f_c \in \mathbb{N}$, let $\Phi_c \in \mathcal{L}_c$ be its spectral approximation operator, and let μ_{f_c} be a minimizer of $\mathcal{E}_{f_c}(\mu) \stackrel{\text{def.}}{=}} \frac{1}{2\lambda} \|\Phi_c \mu - y\|^2 + |\mu|(\mathbb{T}^d)$ over $\mathcal{M}(\mathbb{T}^d)$.*

Then, the sequence $(\mu_{f_c})_{f_c \in \mathbb{N}}$ has accumulation points in the weak- topology and each of them is a solution to (2.3).*

Proof. By definition of \mathcal{E}_{f_c} and μ_{f_c} ,

$$|\mu_{f_c}|(\mathbb{T}^d) \leq \mathcal{E}_{f_c}(\mu_{f_c}) \leq \mathcal{E}_{f_c}(0) = \frac{1}{2\lambda} \|y\|_{\mathcal{H}}^2.$$

As the total variation ball (of radius $\frac{1}{2\lambda} \|y\|_{\mathcal{H}}^2$) is compact and metrizable for the weak-* topology, the sequence $(\mu_{f_c})_{f_c \in \mathbb{N}}$ has accumulation points. Let μ^* be any accumulation point and let us denote by $(\mu_n)_{n \in \mathbb{N}}$ any subsequence that converges towards μ^* in the weak-* sense. We denote by Φ_n (resp. φ_n) the corresponding operator (resp. kernel). Let $\mu \in \mathcal{M}(\mathbb{T}^d)$ be any Radon measure.

We note that, as $\varphi \in \mathcal{C}^j(\mathbb{T}^d; \mathcal{H})$ with $j \geq \lfloor \frac{d}{2} \rfloor + 1$,

$$\begin{aligned} & \sum_{k \in \mathbb{Z}^d \setminus \{0\}} \|c_k(\varphi)\|_{\mathcal{H}} \\ & \leq \frac{1}{2} \left(\sum_{k \in \mathbb{Z}^d \setminus \{0\}} \|c_k(\varphi)\|_{\mathcal{H}}^2 (k_1^j + \dots + k_d^j)^2 + \sum_{k \in \mathbb{Z}^d \setminus \{0\}} \frac{1}{(k_1^j + \dots + k_d^j)^2} \right) \\ & < +\infty, \end{aligned}$$

hence the series $\sum_{k \in \mathbb{Z}^d} c_k(\varphi) e_k$ converges uniformly on \mathbb{T}^d towards φ . As a result, φ_n converges uniformly towards φ , and

$$\lim_{n \rightarrow +\infty} \|\Phi_n \mu - y\|_{\mathcal{H}}^2 = \|\Phi \mu - y\|_{\mathcal{H}}^2.$$

Moreover, for any $h \in \mathcal{H}$,

$$\begin{aligned} |\langle h, (\Phi_n \mu_n - \Phi \mu^*) \rangle| & \leq \|h\|_{\mathcal{H}} \int_{\mathbb{T}^d} \|\varphi_n(x) - \varphi(x)\|_{\mathcal{H}} d|\mu_n|(x) \\ & \quad + \left| \int_{\mathbb{T}^d} \langle h, \varphi(x) \rangle d\mu_n(x) - \int_{\mathbb{T}^d} \langle h, \varphi(x) \rangle d\mu^*(x) \right|. \end{aligned}$$

The first term vanishes by uniform convergence of φ_n and boundedness of $(\mu_n)_{n \in \mathbb{N}}$, whereas the second one vanishes by the weak-* convergence of μ_n towards μ^* . Hence, $\Phi_n \mu_n \rightarrow \Phi \mu^*$ weakly in \mathcal{H} , and

$$\|\Phi \mu^* - y\|_{\mathcal{H}}^2 \leq \liminf_{n \rightarrow +\infty} \|\Phi_n \mu_n - y\|_{\mathcal{H}}^2. \quad (2.18)$$

To complete the proof, we note that by definition of μ_n , we have

$$\mathcal{E}_n(\mu_n) \leq \mathcal{E}_n(\mu), \quad (2.19)$$

and passing to the inferior limit, we get

$$\begin{aligned} \lambda |\mu^*|(\mathbb{T}^d) + \frac{1}{2} \|\Phi \mu^* - y\|_{\mathcal{H}}^2 & \leq \liminf_{n \rightarrow +\infty} \mathcal{E}_n(\mu_n) \\ & \leq \liminf_{n \rightarrow +\infty} \mathcal{E}_n(\mu) \\ & = \lambda |\mu|(\mathbb{T}^d) + \frac{1}{2} \|\Phi \mu - y\|_{\mathcal{H}}^2. \end{aligned}$$

As this is true for any $\mu \in \mathcal{M}(\mathbb{T}^d)$, we deduce that μ^* is a solution to (2.3). \square

By construction, $\Phi_c \mu$ only depends on the Fourier coefficients $c_k(\mu)$ for $k \in \Omega_c$. Indeed, for any $\mu \in \mathcal{M}(\mathbb{T}^d)$,

$$\Phi_c \mu = \int_{\mathbb{T}^d} \varphi_c(x) d\mu(x) = \sum_{k \in \Omega_c} c_k(\varphi) \int_{\mathbb{T}^d} e^{2i\pi \langle k, x \rangle} d\mu(x) = \sum_{k \in \Omega_c} c_{-k}(\varphi) (\mathcal{F}_c \mu)_k.$$

This leads to the following definition.




Figure 2.2: Examples of measurements. In the first two cases (left), the observations y lie in the Fourier domain, and we plot $\mathcal{F}_c^* y$. In the last two cases (right), y lies on a grid \mathcal{G} , and can be plotted directly.

Definition 8 (Spectral approximation factorization). *The spectral approximation operator may be factorized as*

$$\Phi_c = \mathcal{A}(\varphi)\mathcal{F}_c, \quad (2.20)$$

where $\mathcal{A}(\varphi) : \mathbb{C}^{|\Omega_c|} \rightarrow \mathcal{H}$ is defined by $\mathcal{A}(\varphi)_k \stackrel{\text{def}}{=} c_{-k}(\varphi)$. In particular, when $\mathcal{H} = \mathbb{C}^N$ and $\varphi = (\varphi_1, \dots, \varphi_N)$, the matrix of $\mathcal{A}(\varphi) \in \mathcal{M}_{N, |\Omega_c|}(\mathbb{C})$ is given by

$$\mathcal{A}(\varphi)_{j,k} = c_{-k}(\varphi_j), \quad \forall 1 \leq j \leq N, \quad \forall k \in \Omega_c.$$

We call this matrix the spectral approximation matrix of Φ .

In the rest of the paper, we therefore focus on solving the problem

$$\min_{\mu \in \mathcal{M}(\mathbb{T}^d)} \frac{1}{2\lambda} \|\mathcal{A}(\varphi)\mathcal{F}_c \mu - y\|_{\mathcal{H}}^2 + |\mu|(\mathbb{T}^d). \quad (\mathcal{P}_\lambda(y))$$

2.3.2 Examples

We give a few instances of problems for which the matrix $\mathcal{A}(\varphi)$ may be computed. Each case discussed below is illustrated in Fig. 2.2.

Deconvolution The convolution with some kernel $\tilde{\varphi} \in L^2(\mathbb{T}^d)$, obtained with $\varphi(x) = \tilde{\varphi}(\cdot - x)$, is equivalently obtained by multiplying the Fourier coefficients of the measure with those of $\tilde{\varphi}$. In other words, one might equivalently choose $\mathcal{H} = \ell^2(\mathbb{Z}^d)$, and

$$\varphi(x) = \left(c_k(\tilde{\varphi}) e^{-2i\pi\langle k, x \rangle} \right)_{k \in \mathbb{Z}^d}.$$

The spectral approximation matrix of a convolution operator is thus simply the diagonal matrix

$$\mathcal{A}(\varphi) = \text{Diag} (c_k(\tilde{\varphi}))_{k \in \Omega_c}.$$

Example 3 (Ideal low-pass filtering). This corresponds to convolving with the Dirichlet kernel, given by

$$\tilde{\varphi}_D(\mathbf{x}) = \sum_{\mathbf{k} \in \Omega_c} e^{2i\pi\langle \mathbf{k}, \mathbf{x} \rangle}.$$




Figure 2.3: Evolution of $\Phi_c \delta_{x_0}$, $x_0 \in \mathbb{T}^2$, for different values of f_c , in the Gaussian case. Left image is a true Gaussian convolution (over \mathbb{R}^2).

Thus, one simply has $\varphi_D(x) = (e^{-2i\pi\langle k, x \rangle})_{k \in \Omega_c}$ and $\mathcal{A}(\varphi_D) = \text{Id}$.

Example 4 (Gaussian filtering). Let

$$g(\mathbf{t}) \stackrel{\text{def.}}{=} e^{-\frac{1}{2}\langle \mathbf{t}, \Sigma^{-1} \mathbf{t} \rangle} \quad \forall \mathbf{t} \in \mathbb{R}^d$$

be the centered Gaussian function, with covariance Σ . The adequate definition of the Gaussian kernel over the torus is

$$\tilde{\varphi}_G(x) = \sum_{k \in \mathbb{Z}^d} g(x + k),$$

which also reads, by Poisson summation formula,

$$\tilde{\varphi}_G(x) = \sum_{k \in \mathbb{Z}^d} \hat{g}(k) e^{2i\pi\langle k, x \rangle}, \quad (2.21)$$

where \hat{g} denotes the continuous Fourier transform of g , *i.e.*

$$\hat{g}(\boldsymbol{\xi}) = (2\pi)^{\frac{d}{2}} (\det \Sigma)^{\frac{1}{2}} e^{-2\pi^2 \langle \boldsymbol{\xi}, \Sigma \boldsymbol{\xi} \rangle}, \quad \forall \boldsymbol{\xi} \in \mathbb{R}^d.$$

In that case the approximation matrix reads $\mathcal{A}(\varphi_G) = \text{Diag}(\hat{g}(\mathbf{k}))_{\mathbf{k} \in \Omega_c}$. The error induced by the spectral approximation in the Gaussian case is described in Fig. 2.3. It can be made negligible with f_c sufficiently large.

Subsampled convolution In practical cases, one often only has access to convolution measurements over some sampling grid \mathcal{G} . In fluorescence microscopy for instance [Group, 2013], the observations are accurately described as subsampled Gaussian measurements. In that case, given some convolution kernel $\tilde{\varphi} \in L^2(\mathbb{T}^d)$ (typically a Gaussian), φ may be defined as

$$\varphi(x) = (\tilde{\varphi}(t - x))_{t \in \mathcal{G}} \quad (2.22)$$

which leads to

$$\mathcal{A}(\varphi) = \left(c_k(\tilde{\varphi}) e^{2i\pi\langle k, t \rangle} \right)_{t \in \mathcal{G}, k \in \Omega_c}.$$

Example 5 (Microscopy). The acquisition process for microscopy may be modeled as a Gaussian convolution integrated over small areas [Denoyelle et al., 2020], i.e. at any point $\mathbf{t} \in \mathcal{G}$ of the observation grid, one has

$$(\Phi\mu)_{\mathbf{t}} = \frac{1}{2\varepsilon} \int_{\mathbf{t}+[-\varepsilon,\varepsilon]^d} \int_{\mathbf{x} \in \mathbb{T}^d} \tilde{\varphi}_G(\mathbf{t}' - \mathbf{x}) d\mu(\mathbf{x}) d\mathbf{t}'$$

for $\varepsilon > 0$. One may swap the two integrals, as well as the infinite sum (2.21), and obtain

$$(\Phi\mu)_{\mathbf{t}} = \int_{\mathbf{x} \in \mathbb{T}^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{g}(\mathbf{k}) e^{-2i\pi\langle \mathbf{k}, \mathbf{x} \rangle} \left(\frac{1}{2\varepsilon} \int_{\mathbf{t}+[-\varepsilon,\varepsilon]^d} e^{2i\pi\langle \mathbf{k}, \mathbf{t}' \rangle} d\mathbf{t}' \right) d\mu(\mathbf{x})$$

which leads to

$$(\Phi\mu)_{\mathbf{t}} = \int_{\mathbf{x} \in \mathbb{T}^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{g}(\mathbf{k}) e^{-2i\pi\langle \mathbf{k}, \mathbf{t}-\mathbf{x} \rangle} \prod_{j=1}^d \frac{\sin(2\pi k_j \varepsilon)}{2\pi k_j \varepsilon} d\mu(\mathbf{x})$$

Hence, noting $I_{\mathbf{k}}(\varepsilon) = \prod_j \text{sinc}(2\pi k_j \varepsilon)$, the corresponding approximation matrix reads

$$\mathcal{A}(\varphi) = \left(I_{\mathbf{k}}(\varepsilon) \hat{g}(\mathbf{k}) e^{2i\pi\langle \mathbf{k}, \mathbf{t} \rangle} \right)_{\mathbf{t} \in \mathcal{G}, \mathbf{k} \in \Omega_c}$$

Remark 8. Note that in the case where the grid \mathcal{G} is *regular* of dimension $L_1 \times \dots \times L_d$, and assuming that $2f_c < \min(L_1, \dots, L_d)$, the columns of $\mathcal{A}(\varphi)$ are orthogonal, and therefore the matrix $\mathcal{A}(\varphi)^* \mathcal{A}(\varphi)$ is diagonal. As we see in Section 3.3, this has an important numerical interest. If $2f_c \geq \min(L_1, \dots, L_d)$, $\mathcal{A}(\varphi)^* \mathcal{A}(\varphi)$ is not diagonal, but only a few of its diagonals are non-zeros.

Spatially varying filtering In more general cases, φ may be defined as

$$\varphi(x) = (\tilde{\varphi}(t, x))_{t \in \mathcal{G}},$$

in which case the lines of $\mathcal{A}(\varphi)$ consist in the Fourier coefficients of $x \mapsto \tilde{\varphi}(t, x)$ at frequencies taken in Ω_c . Typical examples include foveation operators [Chang et al., 2000], for which $\tilde{\varphi}$ takes the form

$$\tilde{\varphi}(t, x) = g(\sigma^{-1}(x)(t - x)), \quad \forall t \in \mathcal{G}, \quad \forall x \in \mathbb{T}^d \quad (2.23)$$

for some smoothing function g , typically a Gaussian, and some (positive) covariance function σ . Spatial variations of filters are common in imaging science, due for instance to optical deformations in 2-D or temporal artifacts in 1-D. A typical example where filters are expected to be non-stationary is in astrophysical imaging, see for instance [Alard, 2000, Steinbrig et al., 2002, Gentile et al., 2013]. These variations are for instance important for the observation of the early universe, due to the impact of the lensing effect, see e.g. [Cropper et al., 2013, Ngolé and Starck, 2017]. While it is sometimes possible to account for these variations by deforming the observation, this is often non-trivial. Fig. 2.2 shows an example of Gaussian foveated measurements, for which the matrix $\mathcal{A}(\varphi)$ has no closed-form but may be approximated numerically using the discrete Fourier transform. The approximated foveation kernel is displayed in Fig. 2.4.




Figure 2.4: *Left*: display of the kernel matrix $(\tilde{\varphi}(t, x))_{t, x \in \mathcal{G}}$, for the 1D foveation kernel $\tilde{\varphi}$ of the form (2.23), with \mathcal{G} a regular grid over $[0, 1]$. *Right*: approximated kernel matrices $\mathcal{A}(\varphi)\mathcal{F}_c = (\sum_{k \in \Omega_c} c_{-k}(\tilde{\varphi}_t)e^{-2i\pi kx})_{t, x \in \mathcal{G}}$, for different values of f_c .

2.3.3 Discussion

Although the quality of the approximation proposed in this section depends on the chosen cutoff frequency f_c , and hence on the size of the approximation matrix $\mathcal{A}(\varphi)$, this matrix does not need to be fully stored in many situations. In the convolution case for instance, $\mathcal{A}(\varphi)$ is diagonal, and therefore poses no supplementary computational challenge. More interestingly, as noted above, if one considers subsampled convolution observations on a regular grid \mathcal{G} , then the matrix $\mathcal{A}(\varphi)^*\mathcal{A}(\varphi)$ is diagonal. Since the optimization problem $(\mathcal{P}_\lambda(y))$ may be expressed entirely in terms of $\mathcal{A}(\varphi)^*\mathcal{A}(\varphi)$ (this is further detailed in Section 3.3), an inexpensive implementation is possible, regardless of the size of $\mathcal{A}(\varphi)$.

However, for general non-translation-invariant operators (such as spatially varying filters), the matrix $\mathcal{A}(\varphi)^*\mathcal{A}(\varphi)$ is of size $(2f_c + 1)^d \times (2f_c + 1)^d$ (since $\mathcal{A}(\varphi)$ is of size $|\mathcal{G}| \times (2f_c + 1)^d$) and needs to be fully stored. Consequently, our approach is not efficient if the operator exhibits non-smooth variations, since it would require a wide frequency domain Ω_c , and therefore give rise to large approximation matrices – as it is the case for instance for a foveation kernel exhibiting large variations. Nonetheless, we believe the simplicity and versatility of this method (which allows to re-use an approach designed for convolution) makes it quite appealing in situations where the operator varies smoothly.

2.4 Lasserre’s hierarchy for the Blasso

In this section, we introduce semidefinite approximations of the Beurling LASSO and its dual in arbitrary dimension in the case where the forward operator is of the form (2.20). As in polynomial optimization [Lasserre, 2010, de Castro et al., 2017], the multidimensional case is much more involved than the unidimensional one, and needs the introduction of a hierarchy of semidefinite programs. Our formulations generalize the semidefinite programming reformulation of the atomic norm in [Tang et al., 2013] to the multidimensional case. They also differ from [de Castro et al., 2017] in that we recover a polar decomposition of the sought after measure, instead of a Jordan decomposition. We establish precise connections between the *sequences* solving the semidefinite

hierarchy and the *measure* solving the BLASSO.

2.4.1 Problem statement

Given an optimization problem over measures satisfying some mild assumptions (in particular, the variables should be defined over a semialgebraic domain, such as \mathbb{T}^d), moment approaches [Lasserre, 2001] consist in replacing measures with (infinite) moment sequences. Semidefinite hierarchies then result from truncating these moments, and invoking semidefinite characterizations of moment sequences [Curto and Fialkow, 1996], which take the general form of positivity condition on the sequence.

The semidefinite relaxations of the BLASSO (2.3) that we introduce in this section extend the atomic norm reformulation of [Tang et al., 2013] to the multi-dimensional setting. To make the connection between (2.3) and the atomic norm minimization problem of [Tang et al., 2013] explicit, one can see that (2.3) is actually equivalent to

$$\min_{c \in \mathbb{C}^{(2f_c+1)^d}} \frac{1}{2} \|y - \mathcal{A}(\varphi)c\|_{\mathcal{H}}^2 + \lambda \left(\min_{\mu \in \mathcal{M}(\mathbb{T}^d)} |\mu|(\mathbb{T}^d) \quad \text{s.t.} \quad (\mathcal{F}\mu)_k = c_k \quad \forall k \in \Omega_c \right). \quad (2.24)$$

Therefore, given $c \in \mathbb{C}^{(2f_c+1)^d}$, we focus in this section on the constrained problem

$$\min_{\mu \in \mathcal{M}(\mathbb{T}^d)} |\mu|(\mathbb{T}^d) \quad \text{s.t.} \quad (\mathcal{F}\mu)_k = c_k \quad \forall k \in \Omega_c. \quad (\mathcal{Q}_0(c))$$

Note that, by compactness and lower semi-continuity, the above problem has indeed a minimum. Moreover, its value is the so-called *atomic norm* of c introduced in [Chandrasekaran et al., 2012]. The purpose of the present section is to approximate $(\mathcal{Q}_0(c))$ with problems involving only a finite number of moments of an optimal measure μ and its absolute value $|\mu|$.

Multi-level Toeplitz matrices Let $\ell \geq f_c$, and $N \stackrel{\text{def.}}{=} (2\ell + 1)^d$. We assume that some ordering on multi-indices (*i.e.* elements of $\Omega_\ell \stackrel{\text{def.}}{=} \llbracket -\ell, \ell \rrbracket^d$) has been chosen (for instance the colexicographic order). We denote by \mathcal{H}_N^+ the set of Hermitian matrices of size $N \times N$ which are positive semidefinite.

Definition 9 (Multi-level Toeplitz matrix). *We say that $R \in \mathbb{C}^{N \times N}$ is a multi-level Toeplitz matrix, denoted by $R \in \mathfrak{T}_N$, if for every multi-indices $i, j, k \in \llbracket -\ell, \ell \rrbracket^d$ such that $\|i + k\|_\infty \leq \ell$ and $\|j + k\|_\infty \leq \ell$,*

$$R_{i+k, j+k} = R_{i, j}. \quad (2.25)$$

Remark 9. If R is the trigonometric moment matrix of some measure μ , *i.e.* such that $R_{i, j}$ is the $(i - j)$ -th trigonometric moment of μ , it obviously satisfies

$R \in \mathfrak{T}_N$, as

$$R_{i+k, j+k} = \int_{\mathbb{T}^d} e^{-2i\pi\langle i+k, x \rangle} e^{2i\pi\langle j+k, x \rangle} d\mu(x) = \int_{\mathbb{T}^d} e^{-2i\pi\langle i, x \rangle} e^{2i\pi\langle j, x \rangle} d\mu(x) = R_{i, j}.$$

If the ordering on the multi-indices is colexicographical, the generalized Toeplitz property rewrites

$$R = \sum_{k \in [-2\ell, 2\ell]^d} c_k \Theta_k \quad (2.26)$$

where $\Theta_k = \theta_{k_d} \otimes \dots \otimes \theta_{k_1} \in \mathbb{C}^{N \times N}$. Here θ_{k_j} denotes the $(2\ell + 1) \times (2\ell + 1)$ Toeplitz matrix with ones on its k_j -th diagonal and zeros everywhere else (with the convention that lower and upper diagonals have positive and negative indices respectively), and \otimes stands for the Kronecker product. For instance, with 2×2 matrices, for $d = 2$ and $k = (-1, 0)$, one has

$$\Theta_k = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Remark 10. It is important to note that unlike the moment matrices introduced in chapter 1, the matrices that we consider in this chapter are indexed by Ω_ℓ instead of Ω_ℓ^+ – in particular, note that with our notations $\sum_{k \in \Omega_{2\ell}} c_k \Theta_k = T_{4\ell}(c)$, where $T_{4\ell}(c)$ is defined in (1.7). This may be intuitively explained by the fact that these matrices actually encode the *absolute value* of a measure, as we are implicitly working with the polar decomposition of our variables. Alternatively, one could rather try to retrieve the Jordan decomposition of the solution, *i.e.* its positive and negative parts, but this would require to consider two semidefinite variables (one for each component), see e.g. [de Castro et al., 2017].

2.4.2 Moment relaxations for the Blasso

Now, we consider for $\ell \geq f_c$, and $N \stackrel{\text{def.}}{=} (2\ell + 1)^d$,

$$\min_{\substack{R \in \mathcal{H}_N^+, \\ \tilde{c} \in \mathbb{C}^N, \tau \in \mathbb{R}}} \frac{1}{2} \left(\frac{1}{N} \text{Tr}(R) + \tau \right) \quad \text{s.t.} \quad \begin{cases} (a) & \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^H & \tau \end{bmatrix} \succeq 0 \\ (b) & \tilde{c}_k = c_k, \quad \forall k \in \Omega_{f_c} \\ (c) & R \in \mathfrak{T}_N \end{cases} \quad (\mathcal{Q}_0^{(\ell)}(c))$$

In the following, we write

$$\mathcal{R} \stackrel{\text{def.}}{=} \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^H & \tau \end{bmatrix}. \quad (2.27)$$

Remark 11 (Alternative form). In fact, as noted in [Polisano et al., 2017], it is possible to show by an homogeneity argument that the term τ must be chosen equal to $\frac{1}{N} \text{Tr}(R)$. Indeed the positive semi-definiteness constraint in (a) is equivalent to the following three conditions (see for instance [Curto and Fialkow, 1996])

1. $R \succeq 0$,
2. there exists some $\alpha \in \mathbb{C}^N$ such that $R\alpha = \tilde{c}$,
3. $\tau \geq \langle \alpha, R\alpha \rangle$.

Therefore, at optimality, $\tau = \langle \alpha, R\alpha \rangle$, and replacing α, R with $t\alpha, 1/tR$ for $t > 0$ yields another feasible point with energy $\frac{1}{2} \left(\frac{1}{tN} \text{Tr}(R) + t\tau \right)$. Minimizing that quantity over t yields the equality of the two terms. Therefore, $(\mathcal{Q}_0^{(\ell)}(c))$ is equivalent to

$$\min_{\substack{R \in \mathcal{H}_N^+, \\ c \in \mathbb{C}^N}} \frac{1}{N} \text{Tr}(R) \quad \text{s.t.} \quad \begin{cases} (a') & \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^H & \frac{1}{N} \text{Tr}(R) \end{bmatrix} \succeq 0 \\ (b) & \tilde{c}_k = c_k, \quad \forall k \in \Omega_c \\ (c) & R \in \mathfrak{X}_N \end{cases} \quad (\tilde{\mathcal{Q}}_0^{(\ell)}(c))$$

Incidentally, notice that at optimality the rank of the large matrix in (a) and (a') is equal to $\text{rank}(R)$.

The following result explains that $(\mathcal{Q}_0^{(\ell)}(c))$ defines a relaxation of $(\mathcal{Q}_0(c))$.

Proposition 6. *Let $c \in \mathbb{C}^{(2f_c+1)^d}$. For any $\ell \geq f_c$,*

$$\min(\mathcal{Q}_0^{(\ell)}(c)) \leq \min(\mathcal{Q}_0^{(\ell+1)}(c)) \leq \min(\mathcal{Q}_0(c)). \quad (2.28)$$

Moreover, $\lim_{\ell \rightarrow +\infty} \min(\mathcal{Q}_0^{(\ell)}(c)) = \min(\mathcal{Q}_0(c))$.

The proof is an adaptation of the approach used in real polynomial optimization using Lasserre's hierarchies [Lasserre, 2001].

Proof. First, we note that if τ, R' and c' are feasible for $(\mathcal{Q}_0^{(\ell+1)}(c))$, then τ, R and \tilde{c} are feasible for $(\mathcal{Q}_0^{(\ell)}(c))$, where R and c respectively denote the restrictions of R' and c' to Ω_ℓ . Since R' is constant on its main diagonal, we get $\frac{1}{(2\ell+3)^d} \text{Tr}(R') = \frac{1}{N} \text{Tr}(R)$. That yields the first inequality.

Now, for the second inequality, let $\mu \in \mathcal{M}(\mathbb{T}^d)$ such that $(\mathcal{F}\mu)_k = c_k$, and $\xi(x) \stackrel{\text{def.}}{=} \frac{d\mu}{d|\mu|}(x)$ be its sign (defined $|\mu|$ -almost everywhere). For any $\ell \geq f_c$, consider $R \in \mathbb{C}^{N \times N}$, $c \in \mathbb{C}^N$ and τ defined by

$$R_{i,j} = \int_{\mathbb{T}^d} e^{-2i\pi(i-j, x)} d|\mu|(x), \quad c_j = \int_{\mathbb{T}^d} e^{-2i\pi(j, x)} d\mu(x) \quad \text{and} \quad \tau = |\mu|(\mathbb{T}^d). \quad (2.29)$$

It is immediate that (b) and (c) are satisfied. Moreover, for all $\mathbf{p} \in \mathbb{C}^N$, $q \in \mathbb{C}$,

$$\begin{aligned} \begin{bmatrix} p \\ q \end{bmatrix}^H \begin{bmatrix} R & c \\ c^h & \tau \end{bmatrix} \begin{bmatrix} p \\ q \end{bmatrix} &= \sum_{i,j} p_i^H R_{i,j} p_j + 2 \sum_j \Re(\bar{q} \langle c_j, p_j \rangle) + \tau |q|^2 \\ &= \int_{\mathbb{T}^d} \left(\left| \sum_j p_j e^{2i\pi \langle j, x \rangle} \right|^2 + 2 \Re(\bar{q} \xi^*(x) \sum_j p_j e^{2i\pi \langle j, x \rangle}) + |q|^2 \right) d|\mu|(x) \\ &= \int_{\mathbb{T}^d} \left| \sum_j p_j e^{2i\pi \langle j, x \rangle} + q \xi(x) \right|^2 d|\mu|(x) \geq 0, \end{aligned}$$

which yields (a). As a result, R, c, τ is admissible for $(\mathcal{Q}_0^{(\ell)}(c))$ with energy $\frac{1}{2} \left(\frac{1}{N} \text{Tr}(R) + \tau \right) = |\mu|(\mathbb{T}^d)$, hence $\inf(\mathcal{Q}_0^{(\ell)}(c)) \leq \inf(\mathcal{Q}_0(c))$.

To prove that the limit of the sequence is indeed $(\min(\mathcal{Q}_0(c)))$, let us consider the dual problem to $(\mathcal{Q}_0(c))$,

$$\sup_{p \in \mathbb{C}^{(2f_c+1)^d}} \Re \langle p, c \rangle \quad \text{s.t.} \quad \|\mathcal{F}_c^* p\|_\infty \leq 1 \quad (\mathcal{D}_0(c))$$

It is possible to check that $(\mathcal{D}_0(c))$ always has a solution [Duval and Peyré, 2015] and that strong duality holds (see for instance [Candès and Fernandez-Granda, 2014]), $\max(\mathcal{D}_0(c)) = \min(\mathcal{Q}_0(c))$.

On the other hand, one may show that a dual problem to $(\mathcal{Q}_0^{(\ell)}(c))$ is given by

$$\sup_{\substack{Q \in \mathcal{H}_N^+, \\ p \in \mathbb{C}^{2c}}} \Re \langle p, c \rangle \quad \text{s.t.} \quad \begin{cases} (a) & \begin{bmatrix} Q & \tilde{p} \\ \tilde{p}^H & 1 \end{bmatrix} \succeq 0, \\ (b) & \tilde{p}_k = \begin{cases} p_k & \text{if } k \in \Omega_c \\ 0 & \text{if } k \in \Omega_\ell \setminus \Omega_c \end{cases} \\ (c) & Q - \frac{1}{N} \mathbf{I}_N \in \mathfrak{T}_N^\perp \end{cases} \quad (\mathcal{D}_0^{(\ell)}(z))$$

where \mathfrak{T}_N^\perp is the orthogonal complement to \mathfrak{T}_N , *i.e.*

$$Q \in \mathfrak{T}_N^\perp \iff \forall k \in \llbracket -2\ell, 2\ell \rrbracket^d, \quad \sum_{\substack{i,j \in \llbracket -\ell, \ell \rrbracket^d \\ i+j=k}} Q_{i,j} = 0. \quad (2.30)$$

As before, there exists a solution to $(\mathcal{D}_0^{(\ell)}(z))$ and $\max(\mathcal{D}_0^{(\ell)}(z)) = \min(\mathcal{Q}_0^{(\ell)}(c))$.

Now, let $\varepsilon > 0$, let p be a solution to $(\mathcal{D}_0(c))$ and let $p_\varepsilon \stackrel{\text{def.}}{=} (1 - \varepsilon)p$. Since $\|\mathcal{F}_c^* p_\varepsilon\|_\infty < 1$, the bounded real lemma [Dumitrescu, 2017, Corollary 4.25] ensures that there exists $\ell \geq f_c$, a matrix $Q \in \mathcal{H}_N^+$ with $N = (2\ell + 1)^d$ such that $Q - \frac{1}{N} \mathbf{I}_N \in \mathfrak{T}_N^\perp$ and

$$\begin{bmatrix} Q & \tilde{p}_\varepsilon \\ \tilde{p}_\varepsilon^h & 1 \end{bmatrix} \succeq 0,$$

where \tilde{p}_ε extends p_ε in the sense that $\tilde{p}_{\varepsilon,k} = p_{\varepsilon,k}$ for all $k \in \Omega_c$, 0 otherwise. As a result, Q and \tilde{p}_ε are admissible for $(\mathcal{D}_0^{(\ell)}(z))$, hence

$$\min(\mathcal{Q}_0^{(\ell)}(c)) = \max(\mathcal{D}_0^{(\ell)}(z)) \geq (1 - \varepsilon) \max(\mathcal{D}_0(c)) = (1 - \varepsilon) \min(\mathcal{Q}_0(c)),$$

which yields the claimed convergence. \square

Back to the Blasso. In view of the formulation 2.24, concatenating the minimization in z and in (R, τ) , we are led to solve the following problem:

$$\min_{\substack{R \in \mathcal{H}_N^+, \\ \tilde{c} \in \mathbb{C}^N, \\ \tau \in \mathbb{R}}} \frac{1}{2} \left(\frac{\text{Tr}(R)}{N} + \tau \right) + \frac{1}{2\lambda} \|y - \mathcal{A}c\|_{\mathcal{H}}^2 \quad \text{s.t.} \quad \begin{cases} (a) & \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^H & \tau \end{bmatrix} \succeq 0 \\ (b) & c_k = \tilde{c}_k \quad \forall k \in \Omega_c \\ (c) & R \in \mathfrak{T}_N \end{cases}, \quad (\mathcal{P}_\lambda^{(\ell)}(y))$$

where $N = (2\ell + 1)^d$. The problem $(\mathcal{P}_\lambda^{(\ell)}(y))$ is the semidefinite relaxation of (2.3) at order ℓ .

Remark 12. If one considers the BLASSO with the additional constraint that the measure is positive, then the block formulation we have exposed, which corresponds to the polar decomposition of the signed measure (see Remark 10 above and Proposition 8 below), is not necessary, and the corresponding semidefinite relaxation can be more simply written as

$$\min_{\tilde{c} \in \mathbb{C}^{N(\ell)}} \frac{\text{Tr}(T_\ell(\tilde{c}))}{N(\ell)} + \frac{1}{2\lambda} \|y - \mathcal{A}c\|^2 \quad \text{s.t.} \quad \begin{cases} T_\ell(\tilde{c}) \succeq 0 \\ c_k = \tilde{c}_k \quad \forall k \in \Omega_c \end{cases}. \quad (2.31)$$

where $N(\ell) = (\ell + 1)^d$ and $T_\ell(\tilde{c})$ is defined in (1.7).

2.4.3 Finite convergence

The next proposition discusses the equality case between $(\mathcal{Q}_0^{(\ell)}(c))$ and $(\mathcal{Q}_0(c))$, which we sometimes refer to as *collapsing* of the hierarchy, by interpreting R as a moment matrix.

Proposition 7. *Let $\ell \geq f_c$. Then, $\min(\mathcal{Q}_0^{(\ell)}(c)) = \min(\mathcal{Q}_0(c))$ if and only if there exists (R, z, τ) solution to $(\mathcal{Q}_0^{(\ell)}(c))$ and μ solution to $(\mathcal{Q}_0(c))$ such that*

$$\tau = |\mu|(\mathbb{T}^d) \quad \text{and} \quad R_{i,j} = \int_{\mathbb{T}^d} e^{-2i\pi\langle i-j, x \rangle} d|\mu|(x) \quad (2.32)$$

for all $i, j \in \Omega_\ell$. In particular, if μ is a discrete measure with cardinal r , then $\text{rank } R \leq r$.

Proof. Assume that $\min(\mathcal{Q}_0^{(\ell)}(c)) = \min(\mathcal{Q}_0(c))$, and let μ be a solution to $(\mathcal{Q}_0(c))$. Define R, z and τ by (2.29). As in the proof of Proposition 6, we see that (R, z, τ) is admissible for $(\mathcal{Q}_0^{(\ell)}(c))$, with energy

$$\frac{1}{2} \left(\frac{1}{N} \text{Tr}(R) + \tau \right) = |\mu|(\mathbb{T}^d) = \min(\mathcal{Q}_0^{(\ell)}(c)).$$

Hence (R, z, τ) is a solution to $(\mathcal{Q}_0^{(\ell)}(c))$.

The converse implication is straightforward: if (R, z, τ) (resp. μ) is a solution to $(\mathcal{Q}_0^{(\ell)}(c))$ (resp. $(\mathcal{Q}_0(c))$) such that (2.32) holds, then $\frac{1}{2}(\frac{1}{N} \text{Tr}(R) + \tau) = |\mu|(\mathbb{T}^d)$ and we obtain $\min(\mathcal{Q}_0^{(\ell)}(c)) = \min(\mathcal{Q}_0(c))$.

If R satisfies (2.32) and μ has cardinal r , i.e. $\mu = \sum_{i=1}^r a_i \delta_{x_i}$ with $x_i \neq x_j$ for $i \neq j$, we note that this matrix R is of the form

$$R = \sum_{i=1}^r |a_i| v_\ell(x_i) v_\ell(x_i)^*, \quad \text{where } v_\ell(x) \stackrel{\text{def.}}{=} (e^{-2i\pi(k, x)})_{k \in \Omega_\ell}$$

Thus R is a sum of at most r rank one matrices, and $\text{rank } R \leq r$. \square

If R is the moment matrix of some positive measure ν , it be may retrieved using a matrix-pencil method, as described in chapter 1. The signed measure μ matching the moment c is then recovered by solving the usual Vandermonde system. The next lemma ensures that such a measure μ actually satisfies $|\mu| = \nu$.

Proposition 8. *Let $\mathcal{R} = \begin{pmatrix} R & \tilde{c} \\ \tilde{c}^H & \tau \end{pmatrix}$ be a solution of $(\mathcal{Q}_0^{(\ell)}(c))$. If R is the moment matrix of the sparse measure $\nu = \sum_{j=1}^s a_j \delta_{\mathbf{x}_j}$ ($\nu \geq 0$), then $\tilde{c} = \mathcal{F}\mu$ with $|\mu| = \nu$.*

Proof. R is a solution of $(\mathcal{Q}_0^{(\ell)}(c))$, hence as stated in remark 11, $R \succeq 0$, there exists $q \in \mathbb{C}^N$ such that $\tilde{c} = Rq$, and $\tau = q^H R q$. The relation $\tilde{c} = Rq$ may be written $\tilde{c} = V_\ell^\pm(\mathbf{x}) \text{Diag}(a) V_\ell^\pm(\mathbf{x})^H q$, where $V_\ell^\pm(\mathbf{x})$ is the Vandermonde matrix (1.6). Since $V_\ell(\mathbf{x})^H q = (Q^\mathbb{T}(\mathbf{x}_1), \dots, Q^\mathbb{T}(\mathbf{x}_s))$, where $Q^\mathbb{T}$ is the trigonometric polynomial associated to q (we write Q in the following for simplicity), we obtain

$$\tilde{c} = \mathcal{F}\mu \quad \text{with} \quad \mu = \sum_{j=1}^s (a_j Q(\mathbf{x}_j)) \delta_{\mathbf{x}_j}.$$

In the rest of the proof we show that $|Q(\mathbf{x}_j)| = 1$ for all j .

Consider the interpolating polynomials

$$L_j(\mathbf{x}) = \frac{1}{C_j} \prod_{k \neq j} \sum_{i=1}^d \sin^2(\pi(x_i - x_{k,i})),$$

where $C_j \stackrel{\text{def.}}{=} \prod_{k \neq j} \sum_{i=1}^d \sin^2(\pi(x_{j,i} - x_{k,i}))$, so that $L_j(\mathbf{x}_k) = \delta_{jk}$. Let

$$\ell' \stackrel{\text{def.}}{=} \max(\ell, \max_j \deg_\infty(L_j)). \quad (2.33)$$

Since R is the moment matrix of the positive measure ν , it can be flatly extended up to order ℓ' (theorem 6), if necessary. Let R' be the corresponding matrix. The vectors \tilde{c} and q can then be trivially extended into $\tilde{c}' = (\tilde{c}, 0, \dots, 0)$ and $q' = (q, 0, \dots, 0)$, so that $\tilde{c}' = R'q'$. Let $\tau' = q'^H R' q'$. Then, (R', \tilde{c}', τ') is solution of $(\mathcal{Q}_0^{(\ell')}(c))$. Indeed, it is feasible (cf remark 11), and we have that

$\tau' + \frac{1}{N'} \text{Tr}(R') = \tau + \frac{1}{N} \text{Tr}(R) = \min(\mathcal{Q}_0^{(\ell)}(\mathbf{c})) \leq \min(\mathcal{Q}_0^{(\ell')}(\mathbf{c}))$, which proves the optimality.

On the other hand, we have that $\tau = \frac{1}{N} \text{Tr}(R) = \mathbf{q}^H R \mathbf{q}$ (cf. remark 11), and hence that $\sum_j a_j = \sum_j a_j |Q(\mathbf{x}_j)|^2$. This implies by Cauchy-Schwartz that,

$$\sum_j a_j |Q(\mathbf{x}_j)| \leq \sum_j a_j \quad (2.34)$$

with equality if and only if $|Q(\mathbf{x}_j)| = 1$ for all j . Assume that there exists j_0 such that $|Q(\mathbf{x}_{j_0})| \neq 1$. Let

$$a_j'' \stackrel{\text{def.}}{=} \begin{cases} a_j |Q(\mathbf{x}_j)| & \text{if } j = j_0 \\ a_j & \text{otherwise} \end{cases},$$

and $R'' \stackrel{\text{def.}}{=} V_{\ell'}^{\pm} \text{Diag}(a'') V_{\ell'}^{\pm H}$. Let $\mathbf{q}'' \in \mathbb{C}^{N'}$ be such that

$$V_{\ell'}^{\pm H} \mathbf{q}'' = \left(\frac{Q(\mathbf{x}_1)}{|Q(\mathbf{x}_1)|}, \dots, \frac{Q(\mathbf{x}_s)}{|Q(\mathbf{x}_s)|} \right)^\top.$$

Such a vector exists, it is the coefficients of $\sum_j \frac{Q(\mathbf{x}_j)}{|Q(\mathbf{x}_j)|} L_j$, whose degree does not exceed ℓ' . Finally, let $\tau'' = \mathbf{q}''^H R'' \mathbf{q}''$. Then $(R'', R'' \mathbf{q}'', \tau'')$ is a feasible point of $(\mathcal{Q}_0^{(\ell')})$, and we deduce from (2.34) that

$$\frac{1}{N'} \text{Tr} R'' = \sum_j a_j |Q(\mathbf{x}_j)| < \sum_j a_j = \frac{1}{N'} \text{Tr} R'$$

which contradicts the optimality of (R, z, τ) . \square

In practice, convergence may be detected using the flatness property: if at some level of the hierarchy, the returned matrix is flat, then it admits a representing measure and from Proposition 7, we have equality between the minima.

Remark 13 (Finite convergence in 2-D). In dimension 2, a result of Scheiderer [Scheiderer, 2006] states that a positive polynomial admits a sum-of-squares representation. However, the result does not provide bounds on the degrees of the sos factors, and thus extends only partially Féjer-Riesz' theorem to bivariate polynomials. This suggests that finite convergence should be observed in 2-D, though the level at which it occurs may be arbitrarily large. Strikingly, in our numerical simulations for the BLASSO in 2-D, we always observe that the hierarchy converges from the first order of relaxation, without needing to reach higher levels, although no theoretical result is known that may validate this observation.

On the other hand, recent works on extensions of the Carathéodory-Toeplitz theorem to several dimensions and on the Vandermonde factorization of multi-level Toeplitz matrices, see [Yang et al., 2016, Andersson and Carlsson, 2017], indicates that if a multi-level Toeplitz matrix of size $(2f_c + 1)^d$ has a rank lower

than f_c , then it is the truncated moment matrix of a measure supported on at most f_c spikes. Therefore, if the rank assumption holds for a solution of the semidefinite relaxation of the BLASSO, then we should have that the hierarchy has converged from Propositions 8 and 7. Although there is no guarantee that there always exists such a low-rank solution, we believe that it does with high probability, especially in view of the presence of the trace norm in the objective of $(\mathcal{P}_\lambda^{(\ell)}(y))$, which favors low-rank solutions. This could explain our empirical observation that the hierarchy reaches finite convergence as soon as $\ell \geq f_c$.

Remark 14 (Low rank solutions). When finite convergence occurs, Proposition 7 ensures that the semidefinite relaxation $(\mathcal{P}_\lambda^{(\ell)}(y))$ admits a low-rank solution. Although this problem might not have a unique solution (it is not strictly convex with respect to variable R), the algorithm explained in chapter 3 focuses on capturing a low-rank solution to this problem. This structure cannot be exploited by usual interior point algorithms, which tend to return full-rank solutions.

2.4.4 Dual sum-of-squares strengthening

We conclude this chapter by formulating the semidefinite approximations of the dual (2.7) of the BLASSO. The dual aspect of the moment problem, on which is based the moment hierarchy, is the theory of positive polynomials. The link between the two is made by the Riesz-Haviland theorem [Haviland, 1935, Haviland, 1936], which relates the existence of a positive representing measure for an infinite (real) sequence to the positivity of a certain functional over all positive (real) polynomials. In the same way that it is hard to detect if a sequence admits a representing measure in several dimensions, checking whether a certain multivariate polynomial is positive is often intractable.

On the other hand, sum-of-squares (sos) polynomials are much simpler to identify. Indeed, a polynomial $p \in \mathbb{C}_n[\mathbf{z}, \mathbf{z}^{-1}]$ is sum-of-squares if and only if there exists a positive semidefinite matrix $Q \in \mathbb{C}^{N_n \times N_n}$ such that $p = \psi_n(\mathbf{z}^{-1})^H Q \psi_n(\mathbf{z})$, where $\psi_n(\mathbf{z}) = [\mathbf{z}^{\mathbf{k}}]_{\mathbf{k} \in \Omega_n}$ is the vector of monomials (see for instance [Dumitrescu, 2017, Theorem 3.15]). Unlike the univariate case, a *nonnegative* multivariate polynomial is not necessarily sos, but *positive* polynomials admit sos decompositions, with no bound however on the degrees of the sos factors in the generic case. Representative examples of such characterizations include Schmügdgen Positivstellensatz [Schmügdgen, 1991] or Putinar's Positivstellensatz [Putinar, 1993, Jacobi and Prestel, 2001].

Coming back to our super-resolution problem, one can see that the dual problem (2.7) asks for the polynomial $1 - |\Phi^* p|^2$ to be nonnegative. Relying on the aforementioned positivity characterizations for polynomials, the sos strengthening introduced by Lasserre [Lasserre, 2001] consist in reinforcing this nonnegativity constraint into a sum-of-squares constraint. Then, each level of the hierarchy imposes a bound on the degrees of the involved sos factors.

Thus, the dual of the BLASSO may be strengthened in sos-optimization prob-

lem which reads

$$\max \Re \langle y, q \rangle - \frac{\lambda}{2} \|q\|^2 \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} Q & \tilde{p} \\ \tilde{p}^* & 1 \end{bmatrix} \succeq 0 \\ \tilde{p}_{\mathbf{k}} = \begin{cases} (\mathcal{A}^* q)_{\mathbf{k}} & \text{if } \mathbf{k} \in \Omega_c \\ 0 & \text{otherwise} \end{cases}, \\ Q - \frac{1}{N} I \in \mathfrak{T}_N^\perp \end{cases}, \quad (\mathcal{D}_\lambda^{(\ell)}(y))$$

Note that $(\mathcal{D}_\lambda^{(\ell)}(y))$ is also the dual of the semidefinite relaxation $(\mathcal{P}_\lambda^{(\ell)}(y))$.

Chapter 3

The Fourier-based Frank Wolfe algorithm

Abstract

The direct resolution of the semidefinite relaxations introduced in Chapter 2 is intractable in large scale settings, since the problem size grows as f_c^{2d} (where f_c is the cutoff frequency and d the ambient dimension), while usual interior points solvers are limited to matrices of size a few hundreds. In this chapter, we introduce a Fourier-based Frank-Wolfe algorithm (FFW), that leverages both the low-rank and the Toeplitz structure of the matrices involved in Lasserre’s hierarchy, resulting in an $O(f_c^d \log f_c)$ complexity per iteration. Our algorithm combines Frank-Wolfe steps, corresponding to the extraction of an eigenvector, with non-convex updates, to build column-by-column a low-rank factorization of the sought after moment matrix. Positions and amplitudes of the target measure are finally recovered using the joint diagonalization algorithm described in Chapter 1. We evaluate the recovery performance of the algorithm on a fluorescent microscopy dataset.

Contents

3.1 Introduction	89
3.1.1 Motivation	89
3.1.2 Efficient SDP solvers	90
3.1.3 Contributions	92
3.1.4 Notations	93
3.2 Toeplitz relaxation	94
3.2.1 The factorized problem	94
3.2.2 Toeplitz relaxation	95
3.2.3 Model stability	97
3.2.4 Connection with Augmented Lagrangian methods	100
3.3 Fast Fourier Transform-based Frank-Wolfe (FFW)	102
3.3.1 Frank-Wolfe over the PSD cone	103
3.3.2 The FFW algorithm	104
3.3.3 Fast-Fourier-Transform-based computations	105
3.3.4 Complexity	108
3.4 Numerics	108
3.4.1 Tests on synthetic data	109
3.4.2 Tests on SMLM data	113

In the purpose of solving the super-resolution problem, we propose an algorithm combining low-cost convex iterations with non-convex updates to solve the semidefinite relaxations of the BLASSO introduced in the previous chapter.

3.1 Introduction

After having seen in Chapter 1 how to recover a measure from its moment matrix, and studied in Chapter 2 a variational approach to retrieve a denoised moment matrix from a large variety of measurements, we now introduce a new solver for this variational formulation.

3.1.1 Motivation

Lasserre’s relaxation gives a systematic method to solve a large class of optimization problems over Radon measures, known as generalized moment problems [Lasserre, 2010]. This strong modeling power comes at the cost of solving a hierarchy of semidefinite programs, involving matrices of typical size n^d , where n is the polynomial degree and d the dimension. Such large-scale SDP are computationally demanding and often cannot be handled by usual algorithms such as interior-point (IP) solvers. Indeed, IP methods traditionally require computing exact Newton search directions, for which the computational effort may grow as fast as $O(N^6)$ for matrices of size N . This lack of efficient SDP solver is one of the reason why up to now, Lasserre’s hierarchy has essentially been successful

for combinatorial problems such as MaxCut [Goemans and Williamson, 1995], in which the number of variables (d) is high and the degree (n) is low – it is equal to 2 in the MaxCut case. Imaging problems on the contrary typically involve polynomials with few variables but high degrees, and thus require adapted algorithms to be solved using Lasserre’s approach.

3.1.2 Efficient SDP solvers

A remarkable property of semidefinite programs that interior-point algorithms fail to exploit is that they tend to admit low-rank minimizers – for instance for the BLASSO problem when the measure to recover is sparse. Indeed, IP methods necessarily output positive matrices, hence full-rank. Extensions have been proposed that enforce a partial low-rank structure on the IP iterates [Bellavia et al., 2020], but can therefore only approximate the true low-rank minimizer. Other attempts to leverage the particular structure of semidefinite programs are numerous and include improvements based on sparsity [Ahmadi et al., 2017], or symmetry [Gatermann and Parrilo, 2009]. We refer to [Majumdar et al., 2020] for a thorough overview of all recent developments of semidefinite programming.

In this section, we focus on two main lines of works that naturally take advantage of the low-rank properties of solutions to lessen the computational burden of SDP programs: Burer-Monteiro approaches and Frank-Wolfe methods.

Burer-Monteiro factorization. Assuming that the semidefinite problem admits a solution X_* of rank s , the Burer-Monteiro approach [Burer and Monteiro, 2005] works by factorizing the optimization variable as $X = UU^*$, where U is of size $n \times r$ with $r \geq s$. Optimizing over U then becomes tractable for algorithms, at the expense of having to solve a non-convex problem, hence with no guarantees of finding a global minimizer. Nonetheless, these algorithms often work extremely well in practice, and manage to globally solve the factorized problem as soon as r is slightly larger than rank X_* in many cases, see e.g. [Burer and Monteiro, 2003, Boumal, 2015, Journée et al., 2010]. This good behavior may be theoretically explained in specific cases [Bandeira et al., 2016], but it does not occur in the general case, and larger values for r are needed to guarantee global optimality [Boumal et al., 2019, Waldspurger and Waters, 2018].

One of the main hypothesis for the Burer-Monteiro approach to work is that the search space of the factorized problem be a smooth manifold [Boumal et al., 2016]. Beyond its theoretical purpose, this assumption further allows the use of Riemannian optimization algorithms. In the case of the Lasserre’s hierarchy (see chapter 2, this space typically is the set of matrices U such that UU^* is Toeplitz,

$$\{U \in \mathbb{C}^{N \times r} ; \Theta_{\perp}(UU^*) = 0\}.$$

where we denote by Θ the orthogonal projection on the set of generalized Toeplitz matrices \mathfrak{T}_N of size N (cf Definition 9), with $\Theta_{\perp} = I - \Theta$.

Frank-Wolfe. Frank-Wolfe algorithm [Frank and Wolfe, 1956], *a.k.a.* conditional gradient [Levitin and Polyak, 1966], builds its solution as a linear combination of a few simple atoms (*e.g.* spikes, rank-one matrices, etc.), greedily selected at each step. It has recently gained a lot of interest in the machine learning and optimization communities due to its simplicity and its relatively inexpensive iterations, which make it useable on large scale problems [Bredies et al., 2009, Jaggi, 2013, Locatello et al., 2017]. Another feature of this algorithm is that it does not require any Hilbertian structure to work, only convexity. This makes it a good candidate for optimizing over Banach spaces [Demyanov and Rubinov, 1967, Bredies and Pikkarainen, 2013], as it has already been exploited for super-resolution over Radon measures in [Boyd et al., 2015, Denoyelle et al., 2020].

The essence of the method is as follows: given a convex and continuously differentiable function f , Frank-Wolfe minimizes f over a compact convex subset C by first minimizing over C the linearization of f at the current point, which is always solved by at least one extreme point of the set C , and then moving towards this minimizer to build the next point (see Figure 3.1 for an illustration). This scheme ensures the “sparsity” of its iterates, since the solution after k iterations is a convex combination of at most k atoms. For instance, if these atoms consist in rank-one matrices, then the output after k iterations is a matrix of rank at most k . For some sets, *e.g.* the ℓ^1 – ball, the set of semidefinite matrices with bounded trace, etc., which are the convex hull of simple atomic sets, the extreme points can be computed efficiently, see [Jaggi, 2013] for an overview. This is an important asset of Frank-Wolfe compared to proximal or projected gradient algorithms which require projections that are often much more costly. Typically, for trace norm minimization, Frank-Wolfe only requires to compute the leading eigenvector of a matrix at each iteration, while proximal approaches need to compute full singular value decompositions [Jaggi and Sulovský, 2010]. On the other hand, conditional gradient methods converge slowly, in general at a rate of $O(1/k)$, making it often necessary to use corrective updates such as away steps [Wolfe, 1970, Guélat and Marcotte, 1986], or sometimes even more involved non-convex optimization steps [Boyd et al., 2015, Denoyelle et al., 2020].

Frank-Wolfe and the Blasso. The utility of Frank-Wolfe algorithms for inverse problems over Radon measures is highlighted in [Bredies and Pikkarainen, 2013], and the method is now at the heart of several state-of-the-art super-resolution algorithms, namely the alternating descent conditional gradient method (ADCG) [Boyd et al., 2015] and the sliding Frank-Wolfe algorithm (SFW) [Denoyelle et al., 2020]. Both these algorithms work by greedily adding at each step a new spike in the reconstructed measure. A non-convex corrective step is added to improve the slow convergence of Frank-Wolfe: at each iteration, a BFGS is used to slightly rectify either the positions of the current spikes in ADCG, or both positions and amplitudes in SFW. Remarkably, for SFW, this results in the algorithm converging after only a finite number of newly added spikes [Denoyelle et al., 2020], thus spectacularly improving the




Figure 3.1: Standard Frank-Wolfe iterations, with no line-search and a step size of $k/(k+2)$. Each iterate is obtained as a linear combination of the extremal points of the convex domain, pictured in cyan on the figure, and the initial point.

convergence rate of traditional Frank-Wolfe methods.

The main motivation of this chapter is to apply Frank-Wolfe algorithm to solve the semidefinite relaxations of the BLASSO introduced in the previous chapter (see section 2.4). In that case however, the optimization domain consists in the intersection of two sets: the semidefinite cone on the one hand, and the multi-level Toeplitz matrices on the other. Extensions of Frank-Wolfe to problems involving several intersecting constraints have been proposed in the literature, see *e.g.* [Harchaoui et al., 2015, Gidel et al., 2018, Silveti-Falls et al., 2019], but whose main improvements concern cases where each set is favorable to Frank-Wolfe iterations (for instance, positive and low-rank). In our case on the other hand, a more interesting decomposition would be to perform Frank-Wolfe iterations on the positive semidefinite cone, and proximal iterations to compute the Toeplitz projection. Such hybridizations of Frank-Wolfe have also been explored recently [Silveti-Falls et al., 2019, Yurtsever et al., 2018]. We discuss the approach of [Silveti-Falls et al., 2019] in section 3.2.4, in comparison with our method which consists in simply penalizing the Toeplitz constraint.

3.1.3 Contributions

We detail in this chapter the Fourier-based Frank-Wolfe (FFW) algorithm, that takes advantage of the low-rank, positive and Toeplitz structure of the variables to reduce the cost per iteration, The main motivation of this section is to provide an algorithm that fully takes advantage of both the low-rank and Toeplitz structures to reduce the cost per iteration, thus making addressable large-scale semidefinite programs. Our method is inspired by both Burer-Monteiro and Frank-Wolfe approaches. The code is available at <https://github.com/Paulcat>.

Section 3.2 introduces a penalized version of the semidefinite lifting ($\mathcal{P}_\lambda^{(\ell)}(y)$), whose solutions conserve the same low-rank properties as those of the original problem. The goal of this relaxation is to make the problem treatable with a Frank-Wolfe approach.

Section 3.3 details the FFW algorithm, a Fourier-based Frank-Wolfe scheme. This algorithm solves the penalized SDP while only handling low-rank factors of the iterates, in a Burer-Monteiro fashion. This factorization also allows to significantly reduce the numerical complexity, using mostly Fast Fourier Transforms. This is detailed in section 3.3.3. Overall, the approach reaches an $O(f_c^d \log f_c)$ complexity per iteration.

We present numerical results on both synthetic and real datasets in Section 3.4. We apply the algorithm on fluorescence microscopy data, using datasets from the Single-Molecule Localization Microscopy (SMLM) challenge [Group, 2013].

3.1.4 Notations

Toeplitz projection In the rest of this dissertation, we denote by $\Theta \in \mathcal{L}(\mathcal{H}_N)$ the projection on the set of multi-level Toeplitz matrices \mathfrak{T}_N , and by $\Theta_\perp \stackrel{\text{def.}}{=} \text{id}_N - \Theta$ the projection on the orthogonal, *i.e.*

$$R \in \mathfrak{T}_N \quad \Leftrightarrow \quad \Theta_\perp(R) = 0.$$

In particular, when working with the colexicographical order, Θ takes the form

$$\Theta : R \in \mathcal{H}_N \mapsto \sum_{\|\mathbf{k}\|_\infty \leq 2\ell} \frac{\langle R, \Theta_{\mathbf{k}} \rangle}{\|\Theta_{\mathbf{k}}\|^2} \Theta_{\mathbf{k}} \in \mathfrak{T}_N, \quad (3.1)$$

where $\Theta_{\mathbf{k}}$ is defined in (2.26). In dimension one for instance, this operator replaces each entry of R by the mean of the corresponding diagonal. We also denote by θ the operator returning the “diagonal” values of the closest (in the ℓ^2 sense) multi-level Toeplitz matrix, *i.e.*

$$\theta : R \in \mathcal{H}_N \mapsto \left(\frac{\langle R, \Theta_{\mathbf{k}} \rangle}{\|\Theta_{\mathbf{k}}\|^2} \right)_{\mathbf{k} \in \Omega_\ell} \in \mathbb{V}_\ell \quad (3.2)$$

where $\mathbb{V}_\ell \stackrel{\text{def.}}{=} \{(v_{\mathbf{k}})_{\mathbf{k} \in \Omega_\ell} ; v_{-\mathbf{k}} = \bar{v}_{\mathbf{k}}\}$. The adjoint of θ reads

$$\theta^* : c \mapsto \sum_{\mathbf{k} \in \Omega_\ell} \frac{c_{\mathbf{k}}}{\|\Theta_{\mathbf{k}}\|^2} \Theta_{\mathbf{k}}.$$

Problem of interest Given $R \in \mathcal{H}_N$, $\tilde{c} \in \mathbb{C}^N$ and $\tau \in \mathbb{R}$, we keep denoting by $\mathcal{R} \in \mathcal{H}_{N+1}$ the block matrix (2.27). To simplify the exposition, we often replace in our problems the triplet (R, \tilde{c}, τ) with simply \mathcal{R} .

Let

$$J_N \stackrel{\text{def.}}{=} \frac{1}{2} \begin{bmatrix} \frac{1}{N} I_N & 0 \\ 0 & 1 \end{bmatrix}, \quad (3.3)$$

so that for any $\mathcal{R} \in \mathcal{H}_{N+1}$ we have $\langle J_N, \mathcal{R} \rangle = \frac{1}{2N} \text{Tr}(\mathcal{R}) + \frac{1}{2}\tau$. With these notations, $(\mathcal{P}_\lambda^{(\ell)}(y))$ can be reformulated as

$$\min_{R \in \mathcal{H}_N, \tilde{c} \in \mathbb{C}^N, \tau \in \mathbb{R}} \langle J_N, \mathcal{R} \rangle + \frac{1}{2\lambda} \|y - \mathcal{A}c\|^2 \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^* & \tau \end{bmatrix} \succeq 0 \\ c_{\mathbf{k}} = \tilde{c}_{\mathbf{k}} \quad \forall \mathbf{k} \in \Omega_c, \\ \Theta_\perp(R) = 0, \end{cases} \quad (\mathcal{P}_\lambda^{(\ell)}(y))$$

Remark 15. The semidefinite program $(\mathcal{P}_\lambda^{(\ell)}(y))$ can be put in standard form, *i.e.* with a linear objective, as

$$\min_{\mathcal{R} \in \mathcal{H}_{N+1}, t \in \mathbb{R}} \langle J_N, \mathcal{R} \rangle + \frac{1}{2\lambda} t \quad \text{s.t.} \quad \begin{cases} \text{diag}(\mathcal{R}, B(t)) \succeq 0, \\ c_{\mathbf{k}} = \tilde{c}_{\mathbf{k}} \quad \forall \mathbf{k} \in \Omega_c, \\ \Theta_\perp(R) = 0 \end{cases}, \quad (\text{standard})$$

where

$$B(t) \stackrel{\text{def.}}{=} \begin{bmatrix} tI & y - \mathcal{A}c \\ (y - \mathcal{A}c)^* & t \end{bmatrix}.$$

3.2 Toeplitz relaxation

In this section, we highlight the difficulties that arise when trying to solve the semidefinite lifting of Blasso $(\mathcal{P}_\lambda^{(\ell)}(y))$ while maintaining low-rank iterates. To overcome these issues, we introduce a penalized alternative that is a simple SDP with no constraints but positivity. We show that the solutions of this new problem preserve many of the properties of the true minimizers, and still allow for good recovery performance.

3.2.1 The factorized problem

Given a positive Hermitian matrix $\mathcal{R} \in \mathcal{H}_{N+1}^+$ of rank r , we may factorize it as $\mathcal{R} = \mathcal{U}\mathcal{U}^*$, where $\mathcal{U} \in \mathcal{M}_{N+1, r}(\mathbb{C})$. Writing

$$\mathcal{U} \stackrel{\text{def.}}{=} \begin{bmatrix} U \\ \xi \end{bmatrix}, \quad (3.4)$$

the block decomposition (2.27) of \mathcal{R} may be obtained from (3.4) as $R = UU^*$, $\tilde{c} = U^*\xi$ and $\tau = \|\xi\|^2$. Given $q \geq r$, the Burer-Monteiro factorization of problem $(\mathcal{P}_\lambda^{(\ell)}(y))$ yields an optimization problem over the set

$$\{U \in \mathcal{M}_{N, q}(\mathbb{C}) ; \Theta_\perp(UU^*) = 0\}, \quad (3.5)$$

which is the set of Toeplitz matrices of size m and of rank at most q . Unfortunately, this set is not a smooth variety: for instance, the set of 3×3 rank-deficient symmetric Toeplitz matrices is the union of a plane and a cone, see Figure 3.2. Hence, our problem does not fall in the scope of [Boumal et al., 2016,




Figure 3.2: Representation of the set of 3×3 rank-deficient symmetric Toeplitz matrices, identified with elements of \mathbb{R}^3 . The two red lines correspond to matrices of rank one. The surface is defined by the implicit equation $(x-z)(x^2+xz-2y^2) = 0$, $(x, y, z) \in \mathbb{R}^3$, which correspond to symmetric Toeplitz matrices with diagonal elements (x, y, z) whose determinant is null. The result is thus the union of the plane $x = z$ and the Lorentz cone $x^2 + xz - 2y^2 = 0$.

[Boumal et al., 2019], and the Burer-Monteiro approach cannot be straightforwardly applied to problem $(\mathcal{P}_\lambda^{(\ell)}(y))$. Furthermore, optimizing over this domain is difficult, and should be tackled with specific approaches (see section 3.1.2).

3.2.2 Toeplitz relaxation

To overcome the difficulty induced by the Toeplitz constraint in $(\mathcal{P}_\lambda^{(\ell)}(y))$, we introduce a penalized version of this problem – which can also be seen as a perturbation of the atomic norm regularizer used in [Tang et al., 2013, Tang et al., 2015].

We consider the following program:

$$\begin{aligned} \min_{R, \tilde{c}, \tau} & \frac{1}{2} \left(\frac{1}{N} \text{Tr}(R) + \tau \right) + \frac{1}{2\lambda} \|y - \mathcal{A}c\|_{\mathcal{H}}^2 + \frac{1}{2\rho} \|R - \Theta(R)\|^2 \\ \text{s.t.} & \begin{cases} \begin{bmatrix} R & \tilde{c} \\ \tilde{c}^* & \tau \end{bmatrix} \succeq 0 \\ \tilde{c}_k = c_k, \quad \forall k \in \Omega_c \end{cases} \end{aligned} \quad (\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$$

where the parameter ρ controls the penalization of the Toeplitz constraint. We can derive its Fenchel dual.

Proposition 9. *The dual of $(\mathcal{P}_{\lambda,\rho}^{(\ell)}(y))$ reads*

$$\max \langle y, q \rangle - \frac{\lambda}{2} \|q\|^2 - \frac{\rho}{2} \|Q - \frac{1}{N}I\|^2 \quad \text{s.t.} \quad \begin{cases} \begin{pmatrix} Q & \tilde{p} \\ \tilde{p}^H & 1 \end{pmatrix} \succeq 0 \\ \tilde{p}_{\mathbf{k}} = \begin{cases} (\mathcal{A}^*q)_{\mathbf{k}} & \text{if } \mathbf{k} \in \Omega_c \\ 0 & \text{otherwise} \end{cases} \\ \Theta(Q) = \frac{1}{2N}I \end{cases}$$

and strong duality holds.

Proof. We may introduce the variable $z = \mathcal{A}c$ in $(\mathcal{P}_{\lambda,\rho}^{(\ell)}(y))$ and consider the Lagrangian

$$\begin{aligned} \mathcal{L}(R, Q, \tilde{c}, \tilde{p}, \tau, \alpha, z, q) \stackrel{\text{def.}}{=} & \frac{1}{2} \left(\frac{1}{N} \text{Tr}(R) + \tau \right) + \frac{1}{2\lambda} \|y - \mathcal{A}c\|_{\mathcal{H}}^2 + \frac{1}{2\rho} \|\Theta_{\perp}(R)\|^2 \\ & + \left\langle \begin{pmatrix} Q & \tilde{p} \\ \tilde{p}^H & \alpha \end{pmatrix}, \begin{pmatrix} R & \tilde{c} \\ \tilde{c}^H & \tau \end{pmatrix} \right\rangle + \langle q, \mathcal{A}c - z \rangle, \end{aligned}$$

and its associated Lagrangian problem

$$\min_{\mathcal{R}_{\succeq 0}, z} \max_{\mathcal{Q}_{\succeq 0}, q} \mathcal{L}(R, \tilde{c}, \tau, z, Q, \tilde{p}, \alpha, q). \quad (3.6)$$

Minimizing the above expression with respect to (\tilde{c}, z, τ) yields

$$\begin{cases} z = y + \lambda q \\ \tilde{p}_{\mathbf{k}} = \begin{cases} -\frac{1}{2}(\mathcal{A}^*p)_{\mathbf{k}} & \text{if } \mathbf{k} \in \Omega_c \\ 0 & \text{otherwise} \end{cases} \\ \alpha = -\frac{1}{2} \end{cases}.$$

Minimizing over R yields $\Theta_{\perp}(R) = -\rho(Q + \frac{1}{2}NI)$. Thus, exchanging the min and max in (3.6) and solving the inner minimization, one obtains,

$$\max_{Q, \alpha, q} -\langle y, q \rangle - \frac{\lambda}{2} \|q\|^2 - \frac{\rho}{2} \|Q + \frac{1}{2N}I\|^2 \quad \text{s.t.} \quad \begin{cases} \begin{pmatrix} Q & \tilde{p} \\ \tilde{p}^* & -\frac{1}{2} \end{pmatrix} \preceq 0 \\ \tilde{p}_{\mathbf{k}} = \begin{cases} -\frac{1}{2}(\mathcal{A}^*p)_{\mathbf{k}} & \text{if } \mathbf{k} \in \Omega_c \\ 0 & \text{otherwise} \end{cases} \\ Q + \frac{1}{2N}I \in \text{Im } \Theta_{\perp} \end{cases}$$

The solution is then obtained by making the change of variable $Q \leftarrow -Q$. \square

In the rest of the section, we prove that the solution of $(\mathcal{P}_{\lambda,\rho}^{(\ell)}(y))$ are sufficiently close to the ones of $(\mathcal{P}_{\lambda}^{(\ell)}(y))$, when ρ is small enough. We write $f_{\lambda,\rho}$ the objective of the penalized problem, and $\mathcal{R}_{\lambda,\rho}$ a solution. The next proposition shows that $\mathcal{R}_{\lambda,\rho}$ converges towards \mathcal{R}_{λ} as the parameter ρ goes to zero.

Proposition 10. *Let $(\mathcal{R}_{\lambda, \rho_n})$ be a sequence of solution of $(\mathcal{P}_{\lambda, \rho_n}^{(\ell)}(y))$, with $\rho_n \rightarrow 0$ as $n \rightarrow \infty$. Then any accumulation point of $(\mathcal{R}_{\lambda, \rho_n})$ is a solution of $(\mathcal{P}_{\lambda}^{(\ell)}(y))$.*

Proof. One has

$$0 \leq \frac{1}{N} \text{tr}(\mathcal{R}_{\lambda, \rho_n}) \leq f_{\lambda, \rho_n}(\mathcal{R}_{\lambda, \rho_n}) \leq f_{\lambda, \rho_n}(0) = \frac{1}{2\lambda} \|y\|_{\mathcal{H}}^2$$

hence there exists a subsequence $(\mathcal{R}_{\lambda, \rho_s})$ that converges. Let \mathcal{R}_{λ}^* be its limit, and let \mathcal{R}_{λ}^0 be a solution of $(\mathcal{P}_{\lambda}^{(\ell)}(y))$. Since $\frac{1}{2\rho_s} \|\mathcal{R}_{\lambda, \rho_s} - \Theta(\mathcal{R}_{\lambda, \rho_s})\|^2 \leq \frac{1}{2\lambda} \|y\|_{\mathcal{H}}^2$, one has $\|\mathcal{R}_{\lambda, \rho_s} - \Theta(\mathcal{R}_{\lambda, \rho_s})\| \rightarrow 0$ when $s \rightarrow \infty$, which ensures that $\mathcal{R}_{\lambda}^* \in \mathfrak{T}_N$. Furthermore, we have

$$f_{\lambda}(\mathcal{R}_{\lambda, \rho_s}) \leq f_{\lambda, \rho_s}(\mathcal{R}_{\lambda, \rho_s}) \leq f_{\lambda, \rho_s}(\mathcal{R}_{\lambda}^0) = f_{\lambda}(\mathcal{R}_{\lambda}^0).$$

Passing to the limit in these inequalities thus gives $f_{\lambda}(\mathcal{R}_{\lambda}^*) \leq f_{\lambda}(\mathcal{R}_{\lambda}^0)$. Since \mathcal{R}_{λ}^* is positive semidefinite (as the SDP cone is closed) and belongs to \mathfrak{T}_N , it is a solution of $(\mathcal{P}_{\lambda}^{(\ell)}(y))$. \square

3.2.3 Model stability

To ensure the validity of this regularized approach, it is crucial that the solutions $\mathcal{R}_{\lambda, \rho}$ share the main structural properties of the true solutions \mathcal{R}_{λ} , for reasonable values of the parameter ρ . In the rest of the section, we provide a numerical analysis of the sensibility of the penalized problem, with respect to ρ . For simplicity, our numerical illustrations are all in dimension one.

Specifically in this section, we introduce the functional

$$h(\mathcal{R}) \stackrel{\text{def.}}{=} \frac{1}{2} \left(\frac{\text{Tr}(\mathcal{R})}{N} + \tau \right) + \frac{1}{2\lambda} \|y - \mathcal{A}c\|^2 + i_{\mathcal{H}_N^+}(\mathcal{R}) = f_{\lambda}(\mathcal{R}) + i_{\mathcal{H}_N^+}(\mathcal{R}),$$

where \mathcal{R} has the block form (2.27), so that $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$ may be more simply written

$$\min_{\mathcal{R}} \frac{1}{2\rho} \|\Theta_{\perp}(\mathcal{R})\|^2 + h(\mathcal{R}).$$

Note that we make a slight abuse of notations here, since Θ_{\perp} actually only acts on the main sub-block of \mathcal{R} .

The stability of the solutions of the above problem as ρ goes to zero depends on whether the minimal norm certificate associated to $(\mathcal{P}_{\lambda}^{(\ell)}(y))$ satisfies some non-degeneracy condition. Such results are standard in inverse problems regularization theory, and we refer to [Vaïter et al., 2015] for a thorough overview. Recall that a dual certificate for a solution \mathcal{R}_{λ} of $(\mathcal{P}_{\lambda}^{(\ell)}(y))$ is an element η of the space

$$\text{Im } \Theta_{\perp}^* \cap \partial h(\mathcal{R}_{\lambda}),$$

and that among all these certificates, the minimal norm certificate η_0 is the one with minimal norm, that is

$$\eta_0 = \Theta_{\perp}^* (\text{argmin } \|\mathcal{Q}\|^2 \quad \text{s.t.} \quad \Theta_{\perp}^* \mathcal{Q} \in \partial h(\mathcal{R}_{\lambda})). \quad (3.7)$$

Definition 10 (Non-degenerate dual certificates). *A dual certificate $\eta = \Theta_{\perp}^* \mathcal{Q}$ is said to be non-degenerate if*

$$\eta \in \text{ri}(\partial h(\mathcal{R}_{\lambda})),$$

where ri denotes the relative interior of a set, i.e. its interior within its affine hull.

The following theorem from [Vaiteer et al., 2015] shows how to ensure rank stability.

Theorem 9 ([Vaiteer et al., 2015]). *If η_0 is non-degenerate, then there exists $\rho_0 > 0$ such that $\text{rank}(\mathcal{R}_{\lambda, \rho}) = \text{rank}(\mathcal{R}_{\lambda})$ for all $\rho \leq \rho_0$.*

Ensuring non-degeneracy of η_0 is a difficult problem, which we only study numerically. We give evidence of the non-degeneracy of η_0 , as well as empirical considerations on the regime of relaxation under which the solutions are well-behaving.

Non-degeneracy of η_0 The condition $\Theta_{\perp}^* \in \partial h(\mathcal{R}_{\lambda})$ appearing in (3.7) is equivalent in our case to

$$\Theta_{\perp}^* \mathcal{Q} - \nabla f_{\lambda}(\mathcal{R}_{\lambda}) \in \partial i_{\mathcal{H}_N^+}(\mathcal{R}_{\lambda}).$$

The subdifferential $\partial i_{\mathcal{H}_N^+}(\mathcal{R}_{\lambda})$ is the normal cone to the positive semidefinite cone at \mathcal{R}_{λ} (see e.g. [Rockafellar, 1970]), which is given by

$$\partial i_{\mathcal{H}_N^+}(\mathcal{R}_{\lambda}) = \{S \succeq 0; \text{Tr}(S\mathcal{R}_{\lambda}) = 0\}. \quad (3.8)$$

Therefore, the minimal norm certificate η_0 (3.7) may be computed as $\Theta_{\perp}^* \mathcal{Q}_0$ where

$$\mathcal{Q}_0 = \underset{\mathcal{Q}}{\text{argmin}} \|\mathcal{Q}\|^2 \quad \text{s.t.} \quad \begin{cases} \Theta_{\perp}^* \mathcal{Q} \preceq \nabla f_{\lambda}(\mathcal{R}_{\lambda}) \\ \langle \Theta_{\perp}^* \mathcal{Q} - \nabla f_{\lambda}(\mathcal{R}_{\lambda}), \mathcal{R}_{\lambda} \rangle = 0 \end{cases}.$$

Let v_1, \dots, v_s a basis of $\text{Im } \mathcal{R}_{\lambda}$ and v_{s+1}, \dots, v_m a basis of $\text{Ker } \mathcal{R}_{\lambda}$. The relative interior of $\partial i_{\mathcal{H}_N^+}(\mathcal{R}_{\lambda})$ is the set of matrices $S \succeq 0$ satisfying $Sv_i = 0$ for all $i = 1, \dots, s$ (coming from the constraint $\text{Tr}(S\mathcal{R}_{\lambda}) = 0$ in (3.8)) and $v_i^H S v_i < 0$ for all $i = s+1, \dots, m$, so that S is definite negative on $\text{Ker } \mathcal{R}_{\lambda}$. Numerical observations indicate that the non-degeneracy condition

$$\eta_0 - \nabla f_{\lambda}(\mathcal{R}_{\lambda}) \in \text{ri}(\partial i_{\mathcal{H}_N^+}(\mathcal{R}_{\lambda})) \quad (3.9)$$

is satisfied in our case, see Figure 3.3.




Figure 3.3: Display of the singular values of the primal solution \mathcal{R}_λ (top) and of the certificate $\eta_0 - \nabla f_\lambda(\mathcal{R}_\lambda)$ (bottom). The non-degeneracy condition of theorem 9 is satisfied for our Toeplitz relaxation. Indeed, results indicate that the certificate, whose image must contain $\text{Ker } \mathcal{R}_\lambda$ from (3.8), is indeed negative definite on the rest of the space, thus ensuring the non-degeneracy condition (3.9).




Figure 3.4: Rank drop (for Dirichlet measurements, with $f_c = 17$), with respect to ρ . Results are averaged over 100 random trials of positive 3-sparse initial measures; the minimal separation distance, *i.e.* the minimal distance between two consecutive spikes, is larger than $1/(10f_c)$ in all the cases.




Figure 3.5: Trajectories of the support $x_j = \arg(\zeta_j)$, $1 \leq j \leq 3$ (see Algorithm 2), with respect to ρ .

Empirical relaxation regime Numerical observations confirm that there is indeed a comfortable range of value of the parameter ρ for which the extraction remains robust. Figure 3.4 shows the evolution of $\text{rank } \mathcal{R}_{\lambda, \rho}$ with respect to ρ . We see that the rank of $\mathcal{R}_{\lambda, \rho}$ remains stable for low values of ρ , and equal to the sparsity of the initial measure.

Furthermore, we also observe numerical evidence of the robustness of the extraction procedure described in Section 1.3 in this penalized setting: although the solutions of $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$ do not exactly satisfy the generalized Toeplitz property, and therefore are not moment matrices, Algorithm 2 still yields a good estimation of the support of μ_λ . In particular, in a satisfyingly large regime of values of parameter ρ , the eigenvalues of the multiplication matrices (1.20) (see section 1.3.3) remain stable: in Fig. 3.5, we consider a 1D setting, and we plot $\arg z_j(\rho)$, where $z_j(\rho)$, $1 \leq j \leq r$ are the eigenvalues of one multiplication matrix, extracted from a solution of $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$. The width of the line is defined as $|\log(|1 - |z_j||)|$, so that the thicker the line is, the closer $z_j(\rho)$ is to the unit circle. We see that the extraction procedure is stable, up to a certain point, which coincides with the first rank drop of $\mathcal{R}_{\lambda, \rho}$.

3.2.4 Connection with Augmented Lagrangian methods

As we have seen, the necessity of considering a penalized problem comes from the fact that it is hard to maintain low-rank iterates that are positive and Toeplitz at the same time, so that Frank-Wolfe cannot be applied directly on $(\mathcal{P}_\lambda^{(\ell)}(y))$. Relaxing the Toeplitz constraint on the other hand, simplifies the constraint set and makes the Frank-Wolfe steps tractable.

The appeal in this relaxation procedure holds in that it yields good results in a wide regime of penalization (see numerical illustrations in section 3.2.3 above), while keeping the algorithm simple and efficient. Still, it is worth noticing that extensions of Frank-Wolfe have been proposed in the literature that are able to handle multiple constraints. These algorithms are more involved, and mostly rely on augmented Lagrangian (AL) techniques [Liu et al., 2015, Gidel et al., 2018, Silveti-Falls et al., 2019]. We briefly discuss in this section two possible alternative approaches to our problem.

Augmented Lagrangian over the Toeplitz constraint This approach actually share similarities with our penalization scheme. Writing the optimization problem as,

$$\min f_\lambda(\mathcal{R}) + i_{\mathcal{H}_N^+}(\mathcal{R}) \quad \text{s.t.} \quad \Theta_\perp(\mathcal{R}) = 0$$

it consists in primal-dual updates on the augmented Lagrangian functional

$$\mathcal{L}(\mathcal{R}, \alpha) = f_\lambda(\mathcal{R}) + \langle \alpha, \Theta_\perp(\mathcal{R}) \rangle + \frac{1}{2\rho} \|\Theta_\perp(\mathcal{R})\|^2 + i_{\mathcal{H}_N^+}(\mathcal{R})$$

At each iteration, the primal variable \mathcal{R} may be updated with a Frank-Wolfe step, followed by a gradient ascent on the dual variable α :

1. $\mathcal{R}_{t+1} = \mathcal{FW}(\mathcal{R}_t, \mathcal{L}(\cdot, \alpha_t))$
2. $\alpha_{t+1} = \alpha_t + \kappa_t \Theta_\perp(\mathcal{R}_{t+1})$

where $\kappa_t > 0$.

Remarkably, the Frank-Wolfe update is close to the update we perform in our penalized approach, the only difference being the presence of the linear term $\langle \alpha, \Theta_\perp(\mathcal{R}) \rangle$ in the objective. This supplementary term actually poses a challenge in terms of storage and computations. Indeed, due to the nature of Frank-Wolfe updates, the iterates \mathcal{R} may be kept under low-rank form $\mathcal{R} = \mathcal{U}\mathcal{U}^*$, which has obvious advantages in terms of memory space, but also of computational time as the projection $\Theta_\perp(\mathcal{U}\mathcal{U}^*)$ can be performed efficiently (see section 3.3.3 for further details). On the other hand, dual variables have no such low-rank structure. However, the form of the dual updates (2.) suggests that the dual variable may be handled as $\alpha = \alpha^L - \alpha^T$, where α^L is low-rank and α^T is Toeplitz. The main steps of the resulting algorithm are detailed in Algorithm 10. Theoretical guarantees and convergence rates for this algorithm are provided in [Silveti-Falls et al., 2019].

Algorithm 5: CGALP

- 1 Initialize $\mathcal{U}_0 = 0, \alpha_0^L = 0, \alpha_0^T = 0, (\kappa_t)_t$ non-decreasing
 - 2 **while** *not converge* **do**
 - 3 **end**
 - 4 AL gradient: $g_t = \nabla f_\lambda(\mathcal{U}_t \mathcal{U}_t^*) + \Theta_\perp(\alpha_t^L - \alpha_t^T) + \frac{1}{\rho} \Theta_\perp(\mathcal{U}_t \mathcal{U}_t^*)$
 - 5 LMO: $s_t \in \operatorname{argmin}_{s \geq 0} \langle g_t, s \rangle$, with $s_t \stackrel{\text{def.}}{=} v_t v_t^*$
 - 6 Line-search: $\mathcal{U}_{t+1} = V_t(a_t, b_t)$, where $V_t(a, b) = [a\mathcal{U}_t, bv_t]$, and
 - 7 $(a_t, b_t) = \operatorname{argmin}_{0 \leq a+b \leq 1} f_\lambda(V_t V_t^*) + \frac{1}{\rho} \Theta_\perp(V_t V_t^*)$
 - 8 Dual ascent
 - 9 $\alpha_{t+1}^L = \alpha_t^L + \kappa_t \mathcal{U}_{t+1} \mathcal{U}_{t+1}^*$
 - 10 $\alpha_{t+1}^T = \alpha_t^T + \kappa_t \Theta(\mathcal{U}_{t+1} \mathcal{U}_{t+1}^*)$
-

Augmented Lagrangian in the lifted space This approach actually is almost equivalent to the penalized scheme we propose. Algorithm 10 works by considering the augmented Lagrangian associated to the linear constraint $\Theta_{\perp}(\mathcal{R}) = 0$. Another way to handle simultaneously this constraint with the positivity constraint within a Frank-Wolfe scheme is to lift the optimization over the product space $\mathcal{H}_N \times \mathcal{H}_N$. Let M be the linear operator over $\mathcal{H}_N \times \mathcal{H}_N$ such that $M(\mathcal{R}_1, \mathcal{R}_2) = \mathcal{R}_1 - \mathcal{R}_2$ for any $(\mathcal{R}_1, \mathcal{R}_2) \in \mathcal{H}_N \times \mathcal{H}_N$. Then our optimization problem can actually be written as

$$\min \frac{1}{2}f_{\lambda}(\mathcal{R}_1) + \frac{1}{2}f_{\lambda}(\mathcal{R}_2) \quad \text{s.t.} \quad \mathcal{R}_1 \in \text{Im } \Theta, \mathcal{R}_2 \in \mathcal{H}_N^+, M(\mathcal{R}_1, \mathcal{R}_2) = 0.$$

One may then apply an augmented Lagrangian scheme relatively to the linear constraint $M(\mathcal{R}_1, \mathcal{R}_2) = 0$, that is perform primal-dual updates on the augmented Lagrangian $\mathcal{L}(\mathcal{R}_1, \mathcal{R}_2, \alpha)$ given by

$$\frac{1}{2}f_{\lambda}(\mathcal{R}_1) + i_{\text{Im } \Theta}(\mathcal{R}_1) + \frac{1}{2}f_{\lambda}(\mathcal{R}_2) + i_{\mathcal{H}_N^+}(\mathcal{R}_2) + \langle \alpha, \mathcal{R}_1 - \mathcal{R}_2 \rangle + \frac{1}{2\rho} \|\mathcal{R}_1 - \mathcal{R}_2\|^2.$$

This approach was first introduced in [Gidel et al., 2018], but only in the case where linear minimization oracles are easily accessible over each constraint set, so that the primal updates over \mathcal{R}_i ($i = 1, 2, \dots$) may be performed as Frank-Wolfe updates. In our case, the positive semidefinite cone is a good candidate for such Frank-Wolfe updates, while on the other hand one would rather use a proximal step over the Toeplitz constraint, as projection on the set of Toeplitz matrices can be performed easily. This combination of proximal and Frank-Wolfe updates is taken into account in the CGALP algorithm 10 [Silveti-Falls et al., 2019].

Due to the decomposition of the dual variable into low-rank and Toeplitz, the AL gradient in CGALP actually takes the form

$$\begin{aligned} g_t &= \nabla f_{\lambda}(\mathcal{U}_t \mathcal{U}_t^*) + \Theta_{\perp}(\alpha_t^L + \frac{1}{\rho} \mathcal{U}_t \mathcal{U}_t^*) \\ &= \nabla f_{\lambda}(\mathcal{U}_t \mathcal{U}_t^*) + \left(\frac{1}{\rho} + \kappa_{t-1} \right) \Theta_{\perp}(\mathcal{U}_t \mathcal{U}_t^*) \end{aligned}$$

assuming for instance $\kappa_0 = 0$. Thus, this approach is in the end equivalent to progressively decreasing the value of the parameter ρ relaxing the Toeplitz constraint during the iterations. However, as highlighted in Section 3.4, using small values of ρ tend to be numerically instable, so that we found that there is little gain in doing so.

3.3 Fast Fourier Transform-based Frank-Wolfe (FFW)

In this section, we propose an efficient numerical scheme for solving $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$, that takes advantage of the low-rank property of the solutions as well as of the

convolutive structure of the Toeplitz constraint. Given a matrix \mathcal{R} of the form (2.27), we consider in the rest of the paper the normalized objective function

$$f(\mathcal{R}) \stackrel{\text{def.}}{=} C_0 \left(\frac{1}{2} \left(\frac{\text{Tr}(\mathcal{R})}{N} + \tau \right) + \frac{1}{2\lambda} \|y - \mathcal{A}c\|_{\mathcal{H}}^2 + \frac{1}{2\rho} \|\mathcal{R} - \Theta(\mathcal{R})\|^2 \right), \quad (3.10)$$

where $C_0 = 2\lambda / \|y\|_{\mathcal{H}}^2$.

3.3.1 Frank-Wolfe over the PSD cone

The Frank-Wolfe algorithm (see section 3.1.2), *a.k.a.* conditional gradient, aims at minimizing a convex and continuously differentiable function f over a compact convex subset K of a vector space. The essence of the method is as follows: linearize f at the current position \mathcal{R}_t , solve the auxiliary linear problem of minimizing $\mathcal{S} \mapsto \langle \nabla f(\mathcal{R}_t), \mathcal{S} \rangle$ on K , and move towards the minimizer to obtain the next position. This scheme ensures the sparsity of its iterates, since the solution after k iterations is a convex combination of at most k atoms. We refer to [Jaggi, 2013] for a detailed overview of the method. Since no Hilbertian structure is required, it is a good candidate to solve problems in Banach space [Bredies and Pikkarainen, 2013]. Moreover, in many cases, the linear minimization oracle may be computed efficiently, as it amounts to extracting an extremal point of the set K .

Remark 16. In our case, K is the positive semidefinite cone, which is linearly spanned by unit-rank matrices. It is not bounded (hence not compact), but one can restrict $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$ over a bounded subset of the cone by noticing that for any solution \mathcal{R}^* of $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$, one has

$$\frac{1}{2} \left(\tau^* + \frac{\text{Tr}(\mathcal{R}^*)}{N} \right) \leq f(\mathcal{R}^*) \leq f(0),$$

suggesting a subset of the form $\left\{ \mathcal{R} \succeq 0 ; \langle \mathcal{R}, J_N \rangle \leq D_0 \stackrel{\text{def.}}{=} 2f(0) \right\}$, where

$$J_N \stackrel{\text{def.}}{=} \begin{bmatrix} \frac{1}{N} I_N & 0 \\ 0 & 1 \end{bmatrix},$$

so that $\langle \mathcal{R}, J_N \rangle = \tau + \frac{1}{N} \text{Tr}(\mathcal{R})$.

The linear minimization then consists in computing a minor eigenvector of $\nabla f(\mathcal{R})$.

Lemma 13. *Let $M \in \mathcal{M}_N(\mathbb{C})$ be a Hermitian matrix, and let $\{\lambda_1, \dots, \lambda_N\}$ be its eigenvalues, with $\lambda_1 \leq \dots \leq \lambda_N$. Then, for $D_0 \geq 0$,*

$$\underset{\substack{S \succeq 0 \\ \langle S, J_N \rangle \leq D_0}}{\text{argmin}} \langle M, S \rangle = \begin{cases} D_0 J_N^{-\frac{1}{2}} e_1 e_1^* J_N^{-\frac{1}{2}} & \text{if } \lambda_1 < 0 \\ 0 & \text{otherwise} \end{cases}$$

where $e_1 \in \text{Ker}(J_N^{-\frac{1}{2}} M J_N^{-\frac{1}{2}} - \lambda_1 I)$ such that $\|e_1\| = 1$.

Proof. By the change of variable $S' = J_N^{\frac{1}{2}} S J_N^{\frac{1}{2}}$, the linear program becomes

$$\operatorname{argmin} \langle M', S' \rangle \quad \text{s.t.} \quad \begin{cases} S' \succeq 0 \\ \operatorname{Tr}(S') \leq D_0 \end{cases}$$

with $M' = J_N^{-\frac{1}{2}} M J_N^{-\frac{1}{2}}$. Let $S' \in \{X \succeq 0; \operatorname{Tr}(X) \leq D_0\}$, and write $S' = \sum \alpha_i v_i v_i^*$, with $\alpha_i \geq 0$ and $\|v_i\| = 1$. Then

$$\langle M', S' \rangle = \sum \alpha_i v_i^* M' v_i \geq \sum \alpha_i e_1^* M' e_1 = \operatorname{Tr}(S') \lambda_1$$

with equality if and only if $S' = (\sum \alpha_i) e_1 e_1^*$, hence the desired result. \square

3.3.2 The FFW algorithm

We detail in this section the main features of the algorithm we introduce to solve $(\mathcal{P}_{\lambda, \rho}^{(\ell)}(y))$. It essentially relies on a Frank-Wolfe step, followed by a non-convex corrective step similar to [Denoyelle et al., 2020], see Algorithm 6. This last step compensates the slow convergence of Frank-Wolfe.

Linear minimization oracle As mentioned above, the linear minimization step of Frank-Wolfe in our case simply amounts to computing a minor eigenvector of the gradient of f (or more specifically of $J_N^{-\frac{1}{2}} \nabla f J_N^{-\frac{1}{2}}$, but we omit this scaling in the following for simplicity) at the current iterate. In practice, we perform this step efficiently with power iterations, see Section 3.3.3. In order to determine the lowest singular value of ∇f (simultaneously with a corresponding singular vector), we often need to run the algorithm two times consecutively: if after the first run, a negative eigenvalue $\lambda_1 < 0$ is returned, this is indeed the lowest one (since power iterations retrieve the eigenvalue whose magnitude is the greatest); if not, we re-run the algorithm on $\nabla f - \lambda_1 I$, and add the resulting value to λ_1 to obtain the lowest singular value of ∇f .

Low-rank storage To take advantage of the low-rank structure of the solutions, we store our iterates as $\mathcal{R} = \mathcal{U}\mathcal{U}^*$, and work only with the factor \mathcal{U} . This is a cornerstone of our approach, since in practice the matrix \mathcal{R} is too large to be stored entirely. Furthermore, this factorization allows an efficient implementation of several steps of FFW (see Section 3.3.3) that considerably lowers the complexity of the algorithm.

Consequently, at each step of the algorithm, the update consists in adding a column (namely a leading eigenvector of ∇f , as mentioned above) at the end of the matrix \mathcal{U} , thus increasing by one the rank of \mathcal{R} each time.

Non-convex corrective step The non-convex step that we add after each Frank-Wolfe update consists, as in [Boyd et al., 2015], in a gradient descent on $F : \mathcal{U} \mapsto f(\mathcal{U}\mathcal{U}^*)$. The idea is to continuously move the iterate \mathcal{U} in the manifold of fixed rank matrices to improve the value of the functional. This is similar to

the celebrated Burer-Monteiro non-convex method for low-rank minimization , which has proven to be very efficient in practice [Boumal et al., 2016]. We use a limited-memory BFGS descent in our implementation.

Stopping criterion It is known [Jaggi, 2013] that if \mathcal{S} is a solution of the linear minimization oracle, then it satisfies the inequality $\langle \mathcal{R} - \mathcal{S}, \nabla f(\mathcal{R}) \rangle \geq f(\mathcal{R}) - f(\mathcal{R}_*)$ for any \mathcal{R} . We use this property as a stopping criterion, ceasing the iterations if $\langle \mathcal{R} - \mathcal{S}, \nabla f(\mathcal{R}) \rangle$ goes below some tolerance ε .

Algorithm 6:

Input: $\mathcal{U}_0 = [0 \dots 0]^\top$, $D_0 = 2f(0)$
1 while $\langle \mathcal{U}_r \mathcal{U}_r^* - v_r v_r^*, \nabla f(v_r v_r^*) \rangle \geq \varepsilon f(x_0)$ **do**
2 linear minimization oracle:
 $v_r = D_0 J_N^{-\frac{1}{2}} \left(\arg \min_{\|v\| \leq 1} v^\top \cdot \left(J_N^{-\frac{1}{2}} \nabla f(\mathcal{U}_r \mathcal{U}_r^*) J_N^{-\frac{1}{2}} \right) \cdot v \right) J_N^{-\frac{1}{2}}$
3 update $\hat{\mathcal{U}}_{r+1} = [\sqrt{\alpha_r} \mathcal{U}_r, \sqrt{\beta_r} v_r]$, where
 $\alpha_r, \beta_r = \arg \min_{\alpha \geq 0, \beta \geq 0, \alpha + \beta \leq 1} f(\alpha \mathcal{U}_r \mathcal{U}_r^* + \beta v_r v_r^*)$
4 corrective step $\mathcal{U}_{r+1} =$
 bfgs $\left\{ \mathcal{U} \mapsto f(\mathcal{U} \mathcal{U}^*) ; \mathcal{U} \in \mathbb{C}^{(N+1) \times (r+1)}, \text{ starting from } \hat{\mathcal{U}}_{r+1} \right\}$
5 end
Output: $(\mathcal{U}_{i,j})_{1 \leq i \leq N, 1 \leq j \leq r+1}$

3.3.3 Fast-Fourier-Transform-based computations

We show in this section that all elementary operations in the FFW algorithm can be performed efficiently using Fast Fourier Transforms.

In what follows, we consider a variable \mathcal{R} of the form (2.27), of rank r , such that $\mathcal{R} = \mathcal{U} \mathcal{U}^*$ where

$$\mathcal{U} = \begin{bmatrix} U_1 \\ \zeta \end{bmatrix},$$

with $U_1 \in \mathbb{C}^{N \times r}$ and $u \in \mathbb{C}^r$. In particular, this yields $R = U_1 U_1^*$, $z = U_1^* \zeta$ and $\tau = \|\zeta\|^2$. For clarity, we keep using R , z and τ in our expressions, but in practice only U_1 and ζ are stored. Finally, we assume that the matrices are indexed following the colexicographic order, so that generalized Toeplitz matrices may be written in the form (2.26). In this section only, we use the same symbol \mathcal{F} to refer to the discrete Fourier transform (instead of the discrete-time Fourier transform).

Power Iterations Computing a minor eigenvector of ∇f can be done using power iterations, which consist in recursively applying ∇f to a vector. Given

the form (3.10) of the objective, its gradient at \mathcal{R} reads

$$\nabla f = C_0 \begin{bmatrix} \frac{1}{2N}I + \frac{1}{\rho}(R - \Theta(R)) & \frac{1}{2\lambda}\mathcal{A}^*(\mathcal{A}z - y) \\ \frac{1}{2\lambda}(\mathcal{A}^*(\mathcal{A}z - y))^* & \frac{1}{2} \end{bmatrix}, \quad (3.11)$$

so that multiplying it with a vector $w = \begin{bmatrix} w_1 \\ \omega \end{bmatrix}$ ($w_1 \in \mathbb{C}^N$, $\omega \in \mathbb{C}$) yields

$$(\nabla f)w = \frac{C_0}{2N} \begin{bmatrix} w_1 \\ 0 \end{bmatrix} + \frac{C_0}{\rho} \begin{bmatrix} R w_1 \\ 0 \end{bmatrix} + \frac{C_0}{2\lambda} \begin{bmatrix} \omega \mathcal{A}^*(\mathcal{A}z - y) \\ \langle \mathcal{A}^*(\mathcal{A}z - y), w_1 \rangle + \lambda \omega \end{bmatrix} - \frac{C_0}{\rho} \begin{bmatrix} \Theta(R)w_1 \\ 0 \end{bmatrix}.$$

While the first three terms in the sum above are quite straightforward to compute (remember that Rw_1 is computed as $U_1(U_1^*w_1)$), evaluating $\Theta(R)w_1$ on the other hand can be costly. The next two propositions show that it can actually be performed in $O(N \log N)$ operations using only fast Fourier transforms.

Proposition 11. *Let $U_1^{(1)}, \dots, U_1^{(r)}$ be the columns of U_1 . Then $\Theta(U_1 U_1^*) = \sum_{k \in \Omega_{2\ell}} u_k \Theta_k$, where*

$$u_k = \frac{1}{\text{card}\{(s, t) \in \Omega_\ell; s - t = k\}} \left[\sum_{j=1}^r \mathcal{F}^{-1} \left(\left| \mathcal{F}(\tilde{U}_1^{(j)}) \right|^2 \right) \right]_k,$$

where $\tilde{U}_1^{(j)} \in \mathbb{C}^{|\Omega_{2\ell}|}$ is defined as

$$\forall s \in \Omega_\ell, \quad \forall k \in \Omega_{2\ell}, \quad (\tilde{U}_1^{(j)})_{s-k} = \begin{cases} (U_1^{(j)})_{s-k} & \text{if } s - k \in \Omega_\ell \\ 0 & \text{otherwise} \end{cases},$$

Proof. Since $\|\Theta(U_1 U_1^*)\|_2^2 = \sum_{s, t \in \Omega_\ell} |u_{s-t}|^2$, we have:

$$\begin{aligned} \|\Theta(U_1 U_1^*) - U_1 U_1^*\|_2^2 &= \sum_{s, t \in \Omega_\ell} |u_{s-t} - \sum_j (U_1^{(j)})_s (U_1^{(j)})_t|^2 \\ &= \sum_{k \in \Omega_{2\ell}} \sum_{s-t=k} |u_k - \sum_j (U_1^{(j)})_s (U_1^{(j)})_{s-k}|^2 \end{aligned}$$

Minimizing this quantity with respect to u leads to

$$\forall k \in \Omega_{2\ell}, \quad u_k = \frac{1}{\text{card}\{(s, t) \in \Omega_\ell; s - t = k\}} \sum_{s \in \Omega_\ell} \sum_{j=1}^r (\tilde{U}_1^{(j)})_s (\tilde{U}_1^{(j)})_{s-k}.$$

Then

$$\sum_{s \in \Omega_\ell} (\tilde{U}_1^{(j)})_s (\tilde{U}_1^{(j)})_{s-k} = \left[\tilde{U}_1^{(j)} * \tilde{U}_1^{(j)-} \right]_k = \left[\mathcal{F}^{-1} \left(\mathcal{F}(\tilde{U}_1^{(j)}) \cdot \mathcal{F}(\tilde{U}_1^{(j)})^* \right) \right]_k,$$

which yields the desired result. \square

Remark 17. The quantity $\text{card}\{(s, t) \in \Omega_\ell; s - t = k\}$ can be obtained by computing $\mathbb{1}_{\Omega_\ell} * \mathbb{1}_{\Omega_\ell}$, using again fast Fourier transforms.

Proposition 12. *Let $T \in \mathfrak{T}_N$, and write $T = \sum_{k \in \Omega_{2\ell}} u_k \Theta_k$. Let $w_1 \in \mathbb{C}^N$. Then*

$$\forall k \in \Omega_\ell, \quad (Tw_1)_k = (\mathcal{F}^{-1}(\langle \mathcal{F}(u), \mathcal{F}(\tilde{w}_1) \rangle))_k,$$

where $\tilde{w}_1 \in \mathbb{C}^{|\Omega_{2\ell}|}$ is defined as

$$\forall k \in \Omega_{2\ell}, \quad (\tilde{w}_1)_k = \begin{cases} (w_1)_k & \text{if } k \in \Omega_\ell \\ 0 & \text{otherwise} \end{cases}$$

Proof. The product Tw_1 is the (aperiodic) convolution of u and w_1 , and may be formulated as a periodic convolution between u and a zero-padded version of w_1 , hence the result. \square

We conclude this section by giving the closed-form expression for the line-search coefficients α_r and β_r in step 2 of Algorithm 6, assuming in our notations that $\mathcal{U}_r = \mathcal{U}$ and $v_r = w$.

Proposition 13. *With the same notations as before, let*

$$\begin{cases} t_{11} \stackrel{\text{def.}}{=} C_0 \left(\frac{1}{2\lambda} \|\mathcal{A}z\|_{\mathcal{H}}^2 + \frac{1}{2\rho} (\|R\|^2 - \|\Theta(R)\|^2) \right) \\ t_{22} \stackrel{\text{def.}}{=} C_0 \left(\frac{\omega^2}{2\lambda} \|\mathcal{A}w_1\|_{\mathcal{H}}^2 + \frac{1}{2\rho} (\|w_1 w_1^*\|^2 - \|\Theta(w_1 w_1^*)\|^2) \right) \\ t_{12} \stackrel{\text{def.}}{=} C_0 \left(\frac{\omega}{\lambda} \Re \langle z, \mathcal{A}^* \mathcal{A}(w_1) \rangle_{\mathcal{H}} + \frac{1}{\rho} \Re \langle R, w_1 w_1^* \rangle - \langle \Theta(R), \Theta(w_1 w_1^*) \rangle \right) \\ t_1 \stackrel{\text{def.}}{=} C_0 \left(\frac{1}{2} \left(\tau + \frac{\text{Tr } R}{N} \right) - \frac{1}{\lambda} \Re \langle y, \mathcal{A}z \rangle_{\mathcal{H}} \right) \\ t_2 \stackrel{\text{def.}}{=} C_0 \left(\frac{1}{2} \left(\omega^2 + \frac{\|w_1\|^2}{N} \right) - \frac{\omega}{\lambda} \Re \langle y, \mathcal{A}w_1 \rangle_{\mathcal{H}} \right) \end{cases}$$

and let $\Delta = \{\alpha, \beta \in [0, 1]^2; \alpha + \beta \leq 1\}$. Then, the solutions α_r, β_r of the linesearch of step 2 in Algorithm 6 are given by

$$(\alpha_r, \beta_r) = \begin{cases} P_\Delta(0, -\frac{t_2}{2t_{22}}) & \text{if } t_{11} = 0 \text{ and } t_{22} \neq 0 \\ P_\Delta(-\frac{t_1}{2t_{11}}, 0) & \text{if } t_{22} = 0 \text{ and } t_{11} \neq 0 \\ P_\Delta\left(\frac{t_{12}t_2 - 2t_{22}t_1}{4t_{11}t_{22} - t_{12}^2}, \frac{t_{12}t_1 - 2t_{11}t_2}{4t_{11}t_{22} - t_{12}^2}\right) & \text{otherwise} \end{cases}$$

Proof. This is a straightforward consequence of minimizing the quadratic form $f(\alpha\mathcal{R} + \beta w w^*) = t_{11}\alpha^2 + t_{22}\beta^2 + t_{12}\alpha\beta + t_1\alpha + t_2\beta + \frac{1}{2\lambda}\|y\|_H^2$ over Δ . Furthermore, it is realistic to consider only the three cases above (in particular, one never has $t_{11} = t_{22} = 0$ or $4t_{11}t_{22} - t_{12}^2 = 0$ in practice). \square

Implementation of the approximation matrix \mathcal{A} Regarding the computations with the approximation matrix \mathcal{A} , evaluating the objective (3.10) or the gradient (3.11) both only require multiplications with $\mathcal{A}^* \mathcal{A}$: in the objective for instance, one has $\|y - \mathcal{A}z\|_{\mathcal{H}}^2 = \|y\|_{\mathcal{H}}^2 - 2\Re\langle \mathcal{A}^*y, z \rangle + \langle \mathcal{A}^* \mathcal{A}z, z \rangle$, and the vector \mathcal{A}^*y can be pre-computed. This may be well exploited in the implementation, since, as we mentioned in Section 2.3.2, in several cases like convolutions or subsampled convolutions on a regular grid, $\mathcal{A}^* \mathcal{A}$ is diagonal.

3.3.4 Complexity

Using the FFT implementations described in the previous section, we are able to decrease the computational cost of the two elementary operations in FFW: evaluating $f'(\mathcal{U}\mathcal{U}^*)w$ and evaluating $F'(\mathcal{U})$, for $\mathcal{U} \in \mathbb{C}^{(N+1) \times r}$ and $w \in \mathbb{C}^{N+1}$ (and $N = (2\ell + 1)^d$, $\ell \geq f_c$). Table 3.1 summarizes the costs of both these operations in the three settings we consider in this paper, *i.e.* convolution and subsampled convolution (for which $\mathcal{A}(\varphi)^* \mathcal{A}(\varphi)$ is diagonal) and spatially varying filtering (for which $\mathcal{A}(\varphi)^* \mathcal{A}(\varphi)$ is a dense matrix, of size $(2f_c + 1)^d \times (2f_c + 1)^d$).

	$f'(\mathcal{U}\mathcal{U}^*) w$	$F'(\mathcal{U})$
convolution	$O(r\ell^d \log \ell)$	$O(r^2\ell^d + r\ell^d \log \ell)$
subsampled convolution (with regular grid)	$O(r\ell^d \log \ell)$	$O(r^2\ell^d + r\ell^d \log \ell)$
spatially varying filtering	$O(r\ell^d \log \ell + f_c^{2d})$	$O(r^2\ell^d + r\ell^d \log \ell + f_c^{2d})$

Table 3.1: Computational costs.

3.4 Numerics

We study in this section the behavior of FFW with respect to parameters such as sparsity, minimal separation distance, λ , ρ or the number of BFGS iterations.

Scalings As already mentioned, in all our tests, we consider the objective multiplied by $C_0 = 2\lambda/\|y\|^2$ (and its gradient accordingly), so that $f(0) = 1$. If the observations lie on a grid \mathcal{G} , *i.e.* $\mathcal{H} = \mathbb{C}^{|\mathcal{G}|}$, we choose $\|\cdot\|_{\mathcal{H}} = \frac{1}{\sqrt{|\mathcal{G}|}} \|\cdot\|$. Finally, we scale the parameter λ of the BLASSO with $\|\Phi^*y\|_{\infty}$, *i.e.* we set $\lambda = \lambda_0 \|\Phi^*y\|_{\infty}$.

Power Iteration step The tolerance for the power iteration step is set to 10^{-8} , with a maximum of 2000 iterations. Iterations are stopped when the

angle between the eigenvectors returned by two consecutive steps goes below the tolerance.

BFGS step We use Mark Schmidt’s code for the BFGS solver [Schmidt, 2005]. The tolerance for this step is set to 10^{-11} (in terms of functions or parameters changes), with typically 500 as the maximum number of iterations. This step is crucial to ensure the finite convergence of the algorithm. When ρ tends to zero, plain Frank-Wolfe steps become significantly insufficient, and the number of BFGS iterations necessary to converge at each step increases, see Fig. 3.8.

Stopping criterion As explained in Section 3.3.2, the linear minimization oracle in Frank-Wolfe gives access to a bound on the current duality gap, which can then be used to decide when to stop the algorithm. However, it is difficult to define a stopping criterion based on this property that remains stable from one kernel to another. Therefore, in our implementation, we rather use the objective decrease as stopping criterion, and stop the iterations when $|f(\mathcal{R}_{t+1}) - f(\mathcal{R}_t)|$ goes below some tolerance ε (where f is the normalized objective). In the tests presented in this section, ε is set to 10^{-8} .

Support extraction Once a matrix is returned by the FFW algorithm, the recovery of the measure is performed using the extraction procedure detailed in Chapter 1. Since this chapter considers a discrete setting, we use in the tests of this section a random linear combination of the multiplication matrices (see Section 1.3.3) for the joint diagonalization step (see Section 1.3.3 and 1.6 for more details on these steps).

3.4.1 Tests on synthetic data

We focus in this section on synthetic experiments. We generate ground-truth measures randomly by drawing random positions uniformly over \mathbb{T}^d , and random amplitudes uniformly over $[-1, 1]$. Figure 3.6 gives instances of reconstruction in each setting described in Section 2.3.2, with reconstruction errors with respect to μ_0 .

Finite convergence Remarkably, FFW converges in few steps, usually as many as the number of spikes composing the solutions. Fig. 3.7 shows the number of FFW iterations with respect to the sparsity of the initial measure, averaged over 200 random trials, *i.e.* random positions and random amplitudes. The red curve correspond to a subset of these trials for which the minimal separation distance is greater than $1/f_c$. We see that in these simpler cases, FFW converges exactly in r -steps, r being the number of spikes composing the solution.

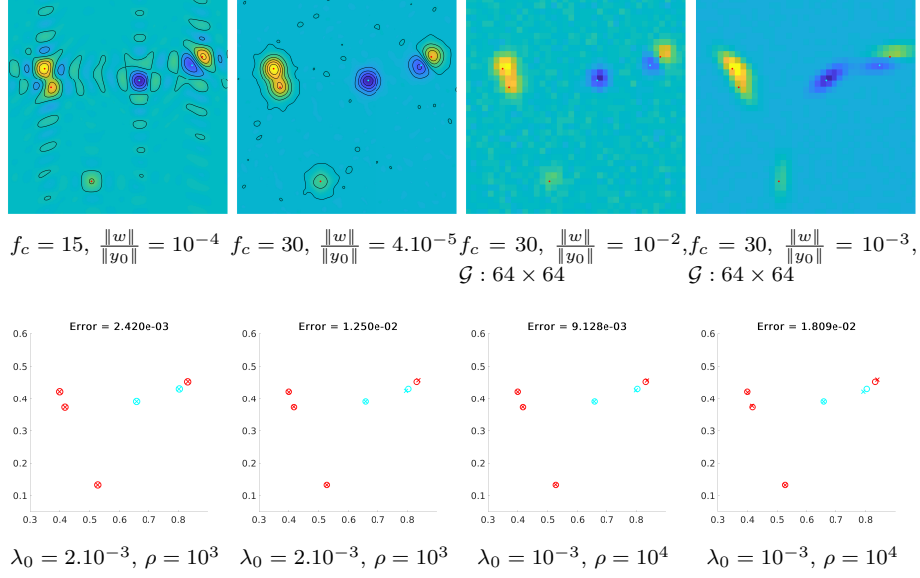


Figure 3.6: From left to right: Measurements $y = \Phi\mu_0 + w$ (we plot \mathcal{F}_c^*y in the first two figures) in the case of Dirichlet convolution, Gaussian convolution, Subsampled Gaussian convolution, and (Subsampled) Gaussian foveation. The support of μ_0 is represented by red (positive spikes) and blue (negative spikes) dots. In each case, the SDP problem is solved using the spectral approximation described in Section 2.3.2. On the bottom line, the indicated errors are defined as $\|x_0 - x_r\|/\|x_0\|$, x_0 and x_r being respectively the ground-truth and the reconstructed supports.




Figure 3.7: Number of FFW iterations with respect to sparsity of the initial measure.

Performance Several metrics can be used to measure the recovery performance of FFW. We use two of them in our tests: the Jaccard index [Jaccard, 1901], and the flat norm [Federer, 1969], also called dual bounded Lipschitz norm.

The flat metric of parameter τ between μ and ν is defined as

$$F_\tau(\mu, \nu) \stackrel{\text{def.}}{=} \sup_{f \in \text{BL}(\mathbb{T}^d, \mathbb{R})} \left\{ \left| \int_{\mathbb{T}^d} f d(\mu - \nu) \right| ; \|f\|_\infty \leq \tau \quad \text{and} \quad \text{Lip}(f) \leq 1 \right\}$$

where $\text{BL}(\mathbb{T}^d, \mathbb{R})$ is the set of bounded Lipschitz functions on \mathbb{T}^d , and $\text{Lip}(f)$ denotes the minimal Lipschitz constant for f . It can be computed using linear programming. The flat metric plays an important role in unbalanced optimal transport theory [Chizat et al., 2018], as it measures the cost of moving mass from μ to ν while allowing for creation and destruction of mass. The trade-off between the two behaviors is controlled by τ .

The Jaccard index measures the similarity between the initial (finite) support $S_0 \subset \mathbb{T}^d$ and the (finite) reconstructed support $S_r \subset \mathbb{T}^d$. It is defined as

$$J \stackrel{\text{def.}}{=} \frac{|S_0 \cap S_r|}{|S_0 \cup S_r|} = \frac{|S_0 \cap S_r|}{|S_0| + |S_r| - |S_0 \cap S_r|}. \quad (3.12)$$

Given an initial set S_0 and its estimate S_r , we may determine the true positive, false positive and false negative using Algorithm 7.

Algorithm 7: True positive, False positive, False negative

Input: $S_0 = (x_{0,i})_i$, $S_r = (x_j)_j$, a tolerance δ

- 1 initialize:
- 2 - $\mathcal{D} = \{\}$ (detected)
- 3 - $\mathcal{F} = S_r$ (falsely detected)
- 4 - $\mathcal{U} = S_0$ (undetected)
- 5 **while** $|\mathcal{F}| > 0$ and $|\mathcal{U}| > 0$ **do**
- 6 compute distances $d_{ij} = d(x_{0,i}, x_{e,j})$ between all points of S_0 and S_r
- 7 find (i, j) corresponding to the minimal distance d_m
- 8 **if** $d_m \leq \delta$ **then**
- 9 | add $x_{0,i}$ to TP
- 10 **else**
- 11 | **return**
- 12 **end**
- 13 remove $x_{0,i}$ from \mathcal{F}
- 14 remove x_j from \mathcal{U}
- 15 **end**

Output: $\mathcal{D}, \mathcal{F}, \mathcal{U}$

From the lists $\mathcal{D}, \mathcal{F}, \mathcal{U}$ returned by Algorithm 7, the number of true positive, false positive and false negative, with respect to the tolerance δ , are defined as $\text{TP} \stackrel{\text{def.}}{=} |\mathcal{D}|$, $\text{FP} \stackrel{\text{def.}}{=} |\mathcal{F}|$ and $\text{FN} \stackrel{\text{def.}}{=} |\mathcal{U}|$, respectively. The Jaccards index then




Figure 3.8: Performance (left) and computational cost (right).

corresponds to the quantity

$$J = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}}.$$

We may also define from these numbers two other useful indicator of success. The recall is defined as

$$R \stackrel{\text{def.}}{=} \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (3.13)$$

and measures the performance of detection with respect to the number of undetected spikes (values close to 1 indicate that a high ability to detect spikes). Finally, the precision is defined as

$$P \stackrel{\text{def.}}{=} \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (3.14)$$

and measures the performance with respect to the number of falsely detected spikes (values close to 1 indicate that the spikes are reconstructed in close neighborhoods of the spikes of the initial measures).

The first plot in Fig. 3.8 shows the flat distance between the measure $\mu_{\lambda, \rho}$ reconstructed by FFW, and a solution μ_{λ} of (2.3) computed using MOSEK. The results are averaged over 680 trials (random positions, random amplitudes and random sparsities in $\llbracket 2, 8 \rrbracket$), and sorted with respect to the minimal separation distance, either lower than $1/f_c$ (dashed line, 417 cases) or greater (solid line, 263 cases). As expected, the quality of the reconstruction decreases as the relaxation parameter ρ increases. On the other hand, the computational cost, represented in the second figure in terms of total number of FFT performed, decreases as ρ increases. For this experiment, the maximum number of BFGS iterations was set to 1000. This cost comes essentially from the BFGS iterations. In all our 1D tests, setting ρ between 1 and 10 gave good performance. The sweet spot depends on the dimension d : for $d = 2$, better performance is achieved with ρ of the order of 10^3 or 10^4 , see Fig. 3.12.




Figure 3.9: Example of reconstruction on data from the smlm challenge. Relative error is $\|x_{\text{rec}} - x_0\|/\|x_0\| = 1.57 \times 10^{-2}$

3.4.2 Tests on SMLM data

We present some results of FFW applied to data taken from the SMLM challenge [Group, 2013]. Fig. 3.9 shows an example of reconstruction for one image of the challenge. On these data, the performance is measured by the Jaccard index, since the challenge gives information only about the locations (and not the amplitudes) of the Dirac masses. Figure 3.10 shows an example of the overall reconstruction from the **bundled tubes long sequence** dataset of SMLM challenge, containing 12000 frames.

In FFW, the most costly step is the BFGS step. Fig. 3.11 shows the impact of diminishing the maximum number of BFGS iterations on the quality of the reconstruction (measured in terms of Jaccard index). The red solid line represents the time taken by FFW, in seconds, and the dashed line the time spent in the BFGS iterations. Results are averaged over 20 random images taken from the challenge. We see that a low bound on the number of BFGS iterations deteriorates the performance. On the other hand, we do not gain much by setting this bound higher than 250.

Finally, Figure 3.12 shows the Jaccard index with respect to parameters ρ and λ_0 . Each pixel is obtained by averaging over 20 random images taken from the challenge. This gives an idea on the range of choices for ρ and λ_0 in which FFW performs well. Although the choices for λ does not change much following the different settings (kernel, dimension), the best values for ρ depends on d , as mentioned above. When $d = 2$, the approach works better with $\rho \approx 10^3$.




Figure 3.10: Recovery from a full dataset of SMLM challenge. This result is obtained by combining the super-resolved output 12000 individual frames of the same type than the one displayed in Figure 3.9




Figure 3.11: Performance versus maximum number of BFGS iterations




Figure 3.12: We measure the performance (in terms of Jaccard index (3.12), recall (3.13) and precision (3.14)) of FFW with respect to parameters λ and ρ . Each pixel is obtained by averaging over 20 images, taken each time randomly within the dataset.

Chapter 4

Semidefinite Approaches for Optimal transport

Abstract

The algorithm developed in Chapter 3 overcomes the computational challenge of Lasserre’s hierarchy when applied to the BLASSO with Fourier type measurements, making the approach tractable on 2-D super-resolution problems. Encouraged by this improvement, we propose in this chapter to broaden the scope of our algorithm, and consider applications in optimal transport. The idea to apply Lasserre’s hierarchy to mass transportation problems was already mentioned in early works of Lasserre [Lasserre, 2008] but was never further investigated. We provide a theoretical framework for this original approach, as well as a numerical solver relying on the algorithm of Chapter 3 and on the multivariate extraction procedure of Chapter 1. In particular, we apply our method on a relaxation of the group-Lasso problem, replacing the constraint of equality between the support of the multiple signals to recover with a Wasserstein penalization.

Contents

4.1	Introduction	117
4.1.1	Developments of optimal transport	118
4.1.2	Contributions	119
4.2	Semidefinite formulations for OT	120
4.2.1	\mathcal{W}_2 and \mathcal{W}_1 costs	120
4.2.2	Semidefinite approximations	122
4.2.3	Numerical illustrations	125
4.2.4	An alternative formulation for \mathcal{W}_1	127
4.3	MMD unbalanced optimal transport	128
4.3.1	MMD penalization	129
4.3.2	Stability analysis	131
4.3.3	Semidefinite formulations for MMD transport	133
4.3.4	Numerical illustrations	133
4.4	SDP transport in higher dimensions	135
4.4.1	Problem statement	135
4.4.2	FFW algorithm for optimal transport	135
4.4.3	Numerical illustrations	137
4.5	Group-Blasso	138
4.6	Wasserstein-Blasso	141
4.6.1	A Wasserstein model for the 2-task BLASSO	141
4.6.2	Multi-marginals extension	143
4.6.3	The particular case of \mathcal{W}_1	144
4.6.4	Numerical illustrations	144

We study in this chapter the use of Lasserre’s hierarchy to solve the optimal transport problem. Since its solution is not sparse in the generic case, we resort to the extraction step described in Section 1.6 to recover the measures from the moments provided by the hierarchy.

4.1 Introduction

The optimal transport (OT) problem was originally formulated by Monge [Monge, 1781] as the matter of moving a collection of particles distributed according to a probability distribution μ_1 (over \mathbb{T}^d in our case), so that their new distribution is described by a second density μ_2 , and the transport is done with minimal effort, given that a local displacement from \mathbf{x} to \mathbf{y} costs $h(\mathbf{x}, \mathbf{y})$, for any $\mathbf{x}, \mathbf{y} \in \mathbb{T}^d$. The optimal transport is then described by a map T , that gives the destination of each particle. The Monge problem thus reads

$$\inf_{T_{\#}\mu_1=\mu_2} \int h(\mathbf{x}, T(\mathbf{x}))d\mu_1(\mathbf{x})$$

where $T_{\#}\mu_1$ is the image measure of μ_1 , such that $(T_{\#}\mu_1)(A) = \mu_1(T^{-1}(A))$ for every measurable set A .

4.1.1 Developments of optimal transport

Kantorovitch’s formulation Monge’s problem over transport maps might fail to have a solution, and the first breakthrough on this question was the reformulation of OT by Kantorovitch as a linear program over joint probability distributions [Kantorovich, 1942],

$$\mathcal{W}(\mu_1, \mu_2) = \inf_{\gamma \in \mathcal{M}_+(\mathbb{T}^d \times \mathbb{T}^d)} \int_{\mathbb{T}^d \times \mathbb{T}^d} h(x, y) d\gamma \quad \text{s.t.} \quad \begin{cases} \pi_1(\gamma) = \mu_1 \\ \pi_2(\gamma) = \mu_2 \end{cases} \quad (\text{OT})$$

which always have solutions. Here $\pi_1(\gamma)$ and $\pi_2(\gamma)$ are the two marginals of the coupling γ . Figure 4.1 shows examples of such transport plans. When h is a distance, \mathcal{W} is known as a Wasserstein distance. This convex reformulation is now the one which is most often considered when it comes to numerical computations. A second major breakthrough is Brenier’s theorem, which states the equivalence between the two formulations of Monge and Kantorovitch, in the cases where the source measures have a density with respect to Lebesgue measure. Optimal transport is now a well established field at the crossroads between convex analysis, PDE’s and probability. Among the many spectacular outcomes of this line of works, let us cite the Benamou-Brenier convex reformulation of OT geodesics [Benamou and Brenier, 2000], the formulation of curvature [Figalli and Villani, 2011] and the study of some class of PDEs as OT gradient flows [Jordan et al., 1998]. For more details on the theory of OT, we refer to the recent monograph [Santambrogio, 2015].




Figure 4.1: Examples of optimal transports over the torus, computed on a grid via linear programming, between sparse measures, uniform measures, and Gaussian measures. The displacement cost is taken as $h(\mathbf{x}, \mathbf{y}) = \|\sin(\pi(\mathbf{x} - \mathbf{y}))\|_2^2$ to approximate the squared Euclidean distance (see Section 4.2 for more details).

Variational OT problems. Besides these theoretical contributions, OT has progressively become a standard tool to define geometric loss functions to solve many problems in data sciences. In imaging sciences, it serves as a discrepancy term for inverse problems such as seismic imaging [Métivier et al., 2016]. In medical imaging [Feydy et al., 2019] and shape registration [Chui and Rangarajan, 2003], it is common to use optimal matching energy to compute deformation between images or surfaces. In unsupervised learning, OT models optimal quantization and clustering [Canas and Rosasco, 2012]. For multi-class supervised learning, Wasserstein

loss can improve over more classical loss such as the Kullback-Leibler divergence, by taking advantage of some metric on the space of labels [Frogner et al., 2015]. In computational biology, it may be used to extract evolving densities of cells to model genetic differentiation [Schiebinger et al., 2019]. OT is also a natural energy to consider when doing density fitting for generative models, such as for instance in Wasserstein Generative Adversarial Networks (WGANs) [Arjovsky et al., 2017, Genevay et al., 2018]

Lastly, beyond its use to refine loss functions, OT may also serve as a regularizer in inverse problems. This was recently applied in the context of multi-task sparse optimization in [Janati et al., 2019b, Janati et al., 2019a], for the simultaneous recovery of pointwise sources from multiple data. While many typical multi-task solvers such as group-Lasso [Yuan and Lin, 2006] impose a perfect overlap between the recovered points, the use of a Wasserstein prior allows to relax this unrealistic assumption and improve the performance of the recovery. As we highlight in section 4.3, this idea of integrating OT in LASSO/BLASSO is similar to the computation of unbalanced OT, which extends OT to positive measures of arbitrary mass (*i.e.* not necessarily probability distributions), as done initially in [Liero et al., 2017, Chizat et al., 2018].

Computational methods A straightforward way to estimate OT between arbitrary input probability measures (possibly being continuous densities) is to discretize them on some fixed grid. This can be too costly for distributions in 2 (or higher) dimensions. This numerical complexity can be mitigated by using an entropic regularization, which in turn can be solved using Sinkhorn’s algorithm [Sinkhorn, 1964]. This method has been revitalized in the imaging and ML communities thanks to Cuturi [Cuturi, 2013] who highlights the smoothness properties of the corresponding loss function, and its efficient computation on GPU architectures. For quadratic Euclidean cost, an efficient approach consists in discretizing only one of the two distributions, and use a semi-discrete solver which relies on the efficient computation of Voronoi diagram in 3D, as advocated in [Mérigot, 2011]. For higher dimensional problem, it is possible to use a similar approach in conjunction with stochastic gradient method [Genevay et al., 2016]. We refer to the monograph [Peyré and Cuturi, 2019] for a review of numerical schemes for OT.

4.1.2 Contributions

In this dissertation, we explore a radically different approach, where the spatial discretization is replaced by a spectral discretization, and the measure is represented through its moments in some basis (here in the Fourier basis). This approach was initially mentioned in [Lasserre, 2008], and was also explored recently in [Condat, 2020] for computing Wasserstein distances between 1-D distributions. We propose here to apply it to a class of problem involving OT (such as extensions of the BLASSO) and to study this approach numerically for both discrete and continuous input measures. To the best of our knowledge,

this is the first time the approach of moments is used to compute OT between distributions in dimensions greater than one.

Section 4.2 introduces semidefinite formulations for optimal transport in arbitrary dimension. The primal approximation was originally mentioned in [Lasserre, 2008]. We prove the convergence of the hierarchy in this optimal transport setting. We also detail an alternative semidefinite formulation for the Wasserstein-1 case.

Section 4.3 introduces the maximum mean discrepancy (MMD) optimal transport, which allows to handle unbalanced marginals. We detail the semidefinite approximations corresponding to this unbalanced transport. We also highlight the resemblance between the MMD formulation and the BLASSO, and provide a numerical analysis of the sensitivity of the penalized solutions.

In section 4.4, we perform a numerical investigation of our solver from chapter 3 for solving the optimal transport problem in dimension higher than one. Note that computing a transport between two 2-D measures requires to retrieve a 4-dimensional coupling measure, making the problem computationally hard.

Finally, section 4.6 focuses on the multi-task sparse recovery problem. We introduce a Wasserstein-penalized version of the Blasso that extends the work of [Janati et al., 2019b] to an off-the-grid setting. Although it enables the use of simpler solvers, we show that a Wasserstein-1 penalization leads to degenerate solutions, which motivates the use of our approach. Numerical illustrations highlight the usefulness of the proposed approach.

4.2 Semidefinite formulations for OT

Supposing that the cost function h is a trigonometric polynomial in problem (OT), the latter becomes a good candidate for Lasserre’s approximation, giving in turn access to new numerical solvers such as the one detailed in Chapter 3. This section aims at introducing the semidefinite relaxation (or strengthening in the dual case) of problem (OT) and its dual.

4.2.1 \mathcal{W}_2 and \mathcal{W}_1 costs

The type of costs that can be used for optimal transport over the torus are studied in [Delon et al., 2009]. In our framework, in order to be able to use Lasserre’s framework, we assume that h is a trigonometric polynomial, *i.e.*

$$h(\mathbf{x}, \mathbf{y}) = \sum \hat{h}_{\mathbf{k}, \mathbf{l}} e^{2i\pi\langle \mathbf{k}, \mathbf{x} \rangle} e^{2i\pi\langle \mathbf{l}, \mathbf{y} \rangle}, \quad (4.1)$$

so that the optimal transport may be expressed in terms of trigonometric moments of measures only. We detail in the following the two polynomials that we consider in the rest of this thesis, which correspond to trigonometric approximations of the usual Wasserstein-2 and Wasserstein-1 distances over the torus.




Figure 4.2: Display of the trigonometric approximation (4.2) for the Wasserstein-2 cost, in 1-D (3 Fourier coefficients). The true cost is shown in blue. Up to a rescaling which does not change the optimal transport problem, the two costs are identical around zero.

- *Wasserstein-2*. The squared Wasserstein-2 distance is commonly obtained by defining the cost as $\min_{\mathbf{k} \in \mathbb{Z}^d} \|\mathbf{x} - \mathbf{y} + \mathbf{k}\|_2^2$, for which a reasonable trigonometric approximation is given by

$$h(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^d \sin^2(\pi(x_i - y_i))^2, \quad (4.2)$$

with corresponding coefficients

$$\hat{h}_{\mathbf{k}, \mathbf{l}} = \begin{cases} d/2 & \text{if } (\mathbf{k}, \mathbf{l}) = (\mathbf{0}, \mathbf{0}) \\ -1/4 & \text{if } (\mathbf{k}, \mathbf{l}) \in \cup_{i=1}^d \{(-\mathbf{e}_i, \mathbf{e}_i), (\mathbf{e}_i, -\mathbf{e}_i)\} \\ 0 & \text{otherwise} \end{cases}. \quad (4.3)$$

This approximation is simple – only $2d + 1$ coefficients are required – and, up to a rescaling, which does not change the solution, it fits up to second order the Euclidean distance, see Figure 4.2. We use it in most of our numerical simulations, in particular in Section 4.6 when using optimal transport as a prior for the BLASSO. For this application, this approximation is legit because one aims at comparing close distributions.

- *Wasserstein-1*. The usual cost is $\min_{\mathbf{k} \in \mathbb{Z}^d} \|\mathbf{x} - \mathbf{y} + \mathbf{k}\|$, which may be replaced with its Fourier expansion, computed numerically. Compared to the Wasserstein-2 approximation (4.2), more coefficients from the Fourier expansion are needed to deal with the singularity of the Wasserstein-1 cost at zero, see Figure 4.3. As discussed below, this approximation of \mathcal{W}_1 may be circumvented by using the dual formulation.




Figure 4.3: Display of the trigonometric approximation of the Wasserstein-1 cost, in 1-D, with 7 Fourier coefficients

4.2.2 Semidefinite approximations

In order to define semidefinite relaxations of the optimal transport problem, it is necessary to truncate at some order the moments constraint on the marginals, that are in infinite number in the original problem. We derive in this section the semidefinite relaxations of OT and its dual, and prove the convergence of the hierarchy in this specific setting.

Primal relaxation Since it involves a number of constraints that is a priori infinite (the constraints over marginals in (OT) are equivalent to constraints over infinite sequences of moments), the optimal transport problem (OT) falls beyond the scope of the original Lasserre’s framework. However, as already noted by Lasserre [Lasserre, 2008, Section 4.4], the usual SDP-relaxation has the natural extension

$$\mathcal{W}^{(\ell)}(\mu_1, \mu_2) = \inf_{\mathbf{c} \in \mathbb{C}^{N(\ell) \times N(\ell)}} \langle \hat{h}, \mathbf{c} \rangle \quad \text{s.t.} \quad \begin{cases} \Pi_1(\mathbf{c}) = \mathcal{F}_\ell \mu_1 \\ \Pi_2(\mathbf{c}) = \mathcal{F}_\ell \mu_2, \\ T_\ell(\mathbf{c}) \succeq 0 \end{cases} \quad (\text{OT}^{(\ell)})$$

where \mathbf{c} intuitively encodes the moments of the sought coupling γ , and Π_1 and Π_2 denote the operators such that, given a sequence $(c_{\mathbf{k}, \mathbf{l}})$, $\Pi_1(\mathbf{c}) = (c_{\mathbf{k}, \mathbf{0}})_{\mathbf{k}} \in \mathbb{C}^{N(\ell)}$ and $\Pi_2(\mathbf{c}) = (c_{\mathbf{0}, \mathbf{l}})_{\mathbf{l}} \in \mathbb{C}^{N(\ell)}$. Note that here the matrix $T_\ell(\mathbf{c})$ has size $(N_+(\ell))^2 \times (N_+(\ell))^2$, where we recall that $N_+(\ell) = (\ell + 1)^d$.

The main difference between problem (OT^(ℓ)) and the conventional semidefinite relaxations is that here the number of linear constraints increases with the order, as the relaxation (OT^(ℓ)) constrains all “moments” up to order ℓ . By comparison, the semidefinite relaxation of the BLASSO ($\mathcal{P}_\lambda^{(\ell)}(y)$) always involves $|\Omega_{f_c}|$ linear constraints, independently of the order ℓ (assuming that $\ell \geq f_c$). Nonetheless, (OT^(ℓ)) still defines a relaxation of (OT), as detailed in the Proposition 14.

Dual strengthening Semidefinite approximations can also be applied to the dual of (OT), given by (see *e.g.* [Santambrogio, 2015])

$$\sup_{\varphi, \psi \in \mathcal{C}(\mathbb{T}^d)} \int_{\mathbb{T}^d} \varphi d\mu_1 + \int_{\mathbb{T}^d} \psi d\mu_2 \quad \text{s.t.} \quad \varphi \oplus \psi \leq h, \quad (\text{D-OT})$$

where the operation \oplus is defined as

$$(\varphi \oplus \psi)(\mathbf{x}, \mathbf{y}) \stackrel{\text{def.}}{=} \varphi(\mathbf{x}) + \psi(\mathbf{y}).$$

The two functions φ and ψ solving (D-OT) are usually called dual potentials. Let $\eta \stackrel{\text{def.}}{=} h - (\varphi \oplus \psi)$. It is nonnegative, and the primal-dual optimality relations also state that $\eta(x, y) = 0$ for any (x, y) supporting the transport plan solving (OT). Such a function that is nonnegative and vanishes on the support of the transport plan is called a dual certificate.

In mirror of the primal relaxation, the strengthening of the dual problem (D-OT) is obtained by replacing the positivity constraint $\varphi \oplus \psi \leq h$ with a more restrictive sum-of-squares (sos) constraint, which is SDP-representable. Again, the dual problem (D-OT) does not exactly fit Lasserre's framework, since φ and ψ have no reason to be polynomials. However, using the (partial) Fourier expansions of these potentials gives a natural extension of the usual sos-strengthening:

$$\sup_{\substack{\mathbf{a}, \mathbf{b} \in \mathbb{C}^{N(\ell)} \\ Q \in \mathbb{C}^{N_+(\ell)^2 \times N_+(\ell)^2}}} \langle \mathbf{a}, \mathcal{F}_\ell \mu_1 \rangle + \langle \mathbf{b}, \mathcal{F}_\ell \mu_2 \rangle \quad \text{s.t.} \quad \begin{cases} Q \succeq 0 \\ \hat{h} - (\mathbf{a} \oplus \mathbf{b}) = T_\ell^*(Q) \end{cases} \quad (\text{D-OT}^{(\ell)})$$

where the adjoint of the Toeplitz operator T_n^* is defined by

$$T_n^* : \begin{array}{l} \mathbb{C}^{N_+(n) \times N_+(n)} \rightarrow \mathbb{C}^{N(n)} \\ Q \mapsto ((Q, \Theta_{\mathbf{k}}))_{\mathbf{k} \in \Omega_n} \end{array}. \quad (4.4)$$

Here, unlike the traditional sos hierarchies, the total degree of the sos polynomials involved in (D-OT^(ℓ)) increases with the order.

Convergence of the hierarchy The next proposition adapts the usual proof of convergence of the Lasserre's hierarchy to the specific case of (OT^(ℓ)).

Proposition 14. *For any $\ell \geq \deg h$, one has*

$$\mathcal{W}^{(\ell)}(\mu_1, \mu_2) \leq \mathcal{W}^{(\ell+1)}(\mu_1, \mu_2) \leq \mathcal{W}(\mu_1, \mu_2).$$

Furthermore

$$\lim_{\ell \rightarrow \infty} \mathcal{W}^{(\ell)}(\mu_1, \mu_2) = \mathcal{W}(\mu_1, \mu_2).$$

Proof. The proof for the double inequality can be readily adapted from the same proof for the BLASSO, see Proposition 6. We only detail the convergence result. Let (φ, ψ) be a solution of the dual problem (D-OT), and let $\varphi_\varepsilon = \varphi - \varepsilon$ and $\psi_\varepsilon = \psi - \varepsilon$ for $\varepsilon > 0$, so that

$$h(\mathbf{x}, \mathbf{y}) - \varphi_\varepsilon(\mathbf{x}) - \psi_\varepsilon(\mathbf{y}) \geq 2\varepsilon > 0, \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{T}^d.$$

The functions φ_ε and ψ_ε may be approximated by trigonometric polynomials: let $\ell \geq \deg(h)$ and $a_\varepsilon, b_\varepsilon \in \mathbb{C}^{N_\ell}$ be such that

$$|\varphi_\varepsilon - \mathcal{F}_\ell^* a_\varepsilon| < \frac{\varepsilon}{2} \quad \text{and} \quad |\psi_\varepsilon - \mathcal{F}_\ell^* b_\varepsilon| < \frac{\varepsilon}{2}.$$

Then, for all $\mathbf{x}, \mathbf{y} \in \mathbb{T}^d$,

$$h(\mathbf{x}, \mathbf{y}) - \mathcal{F}_\ell^* a_\varepsilon - \mathcal{F}_\ell^* b_\varepsilon \geq h(\mathbf{x}, \mathbf{y}) - \varphi_\varepsilon(\mathbf{x}) - \psi_\varepsilon(\mathbf{y}) - \varepsilon \geq \varepsilon > 0.$$

Therefore, from [Dumitrescu, 2017, Theorem 4.15], there exists a matrix $Q \succeq 0$ such that $\hat{h} - (a_\varepsilon \oplus b_\varepsilon) = T_n^*(Q)$, so that a_ε and b_ε are feasible for (D-OT^(ℓ)), and one has

$$\begin{aligned} \mathcal{D}^{(\ell)}(\mu_1, \mu_2) &\geq \langle a_\varepsilon, \mathcal{F}_\ell \mu_1 \rangle + \langle b_\varepsilon, \mathcal{F}_\ell \mu_2 \rangle \\ &\geq \int_{\mathbb{T}^d} (\varphi - \varepsilon - \frac{\varepsilon}{2}) d\mu_1 + \int_{\mathbb{T}^d} (\psi - \varepsilon - \frac{\varepsilon}{2}) d\mu_2 \\ &= \mathcal{D}(\mu_1, \mu_2) - \frac{3\varepsilon}{2}(\mu_1(\mathbb{T}^d) + \mu_2(\mathbb{T}^d)). \end{aligned}$$

Since strong duality holds between (OT) and (D-OT), and (OT^(ℓ)) and (D-OT^(ℓ)), we obtain

$$\begin{aligned} \mathcal{W}^{(\ell)}(\mu_1, \mu_2) &= \mathcal{D}^{(\ell)}(\mu_1, \mu_2) \geq \mathcal{D}(\mu_1, \mu_2) - \frac{3\varepsilon}{2}(\mu_1 + \mu_2)(\mathbb{T}^d) \\ &= \mathcal{W}(\mu_1, \mu_2) - \frac{3\varepsilon}{2}(\mu_1 + \mu_2)(\mathbb{T}^d) \end{aligned}$$

which yields the desired convergence. \square

Remark 18 (Spectral discretization). At a given order ℓ , problem (OT^(ℓ)) may be understood as a “spectral” discretization of the actual optimal transport problem (OT), where one only needs to know the low-frequency moments of the marginals. This contrasts with the spatial discretization that is traditionally used in numerical solvers for optimal transport. As a result, problem (OT^(ℓ)) also shares similarities with the Beurling Lasso problem, since only the first moments of the marginals are matched. Note that such partially constrained optimal transport problems have also recently been studied in [Alfonsi et al., 2019].

Finite convergence The cases of equality between (OT) and (OT^(ℓ)) may be studied in the case where (OT) admits a sparse solution, which happens when the marginals are sparse.

Proposition 15 ([Santambrogio, 2015]). *If μ_1 and μ_2 are k_1 -sparse and k_2 -sparse respectively, then there exists a transport plan solving (OT) which is at most $(k_1 + k_2 - 1)$ -sparse.*

In that case, finite convergence of the hierarchy may be detected via a flatness criterion as before. Note however that flatness of a sequence $\mathbf{c} \in \mathbb{C}^{N \times N}$ does not imply that the subsequences $\Pi_1(c)$ or $\Pi_2(c)$ be flat, see *e.g.* Figure 4.5 for an illustration. Hence, unlike the BLASSO, flatness of the main variable \mathbf{c} in $(\text{OT}^{(\ell)})$ may not suffice to guarantee that the hierarchy has converged.

Proposition 16. *Let c be the solution of $(\text{OT}^{(\ell)})$, at some order $\ell \geq \deg(h)$. If c , $\Pi_1(c)$ and $\Pi_2(c)$ are flat, then $\mathcal{W}^{(\ell)}(\mu_1, \mu_2) = \mathcal{W}(\mu_1, \mu_2)$.*

Proof. If flatness of the three sequences, then c admits a unique positive representing measure γ , that is discrete and whose marginals are μ_1 and μ_2 . Hence γ is feasible for (OT), and $\mathcal{W}(\mu_1, \mu_2) \leq \mathcal{W}^{(\ell)}(\mu_1, \mu_2)$. \square

Remark 19. Although we may determine a posteriori if problems (OT) or (D-OT) are equivalent to their relaxation or strengthening, using the flatness criterion on the primal matrix, there are however no theoretical guarantees about whether finite convergence occurs a priori, regardless of the dimension of the underlying signal. This contrasts for instance with the BLASSO, for which it is known that the hierarchy converges at a known finite order in 1d, due to Féjer-Riesz theorem, and at an unknown (and maybe arbitrarily large) finite order in 2d, due to similar (but weaker) arguments, see *e.g.* [Schmügden, 2017, page 385]. For (D-OT) on the other hand, the optimization variables are not polynomials, so that we cannot fully bridge the gap the variables φ, ψ , which are continuous functions, and a, b , which are SOS polynomials.

4.2.3 Numerical illustrations

We present some results obtained by solving the semidefinite hierarchies introduced above for the optimal transport problem, using the interior point solver SDPT3 [Toh et al., 1999]. Note that this only allows to consider transport between 1d measures, as otherwise the size of the involved semidefinite matrices, of the order of ℓ^{2d} , would be too large. We discuss the extension to higher dimensions in Section 4.4, where we develop an approximation which can be tackled using the FFW algorithm of Chapter 3.

After having recovered the matrix solving $(\text{OT}^{(\ell)})$, a support is recovered by applying the approximate joint diagonalization scheme described in Section 1.6, thus allowing to consider non-sparse measures. Figure 4.4 displays the resulting recovery for several order of relaxation ℓ , as well as the corresponding dual polynomial, solving $(\text{D-OT}^{(\ell)})$.

It is important to keep in mind that, as discussed above, finite convergence of the hierarchy almost never occurs for optimal transport problems, except maybe in cases of sparse transport plans. And even then, flatness of the sequence of moments is not a sufficient guarantee any more, due to the additional marginal

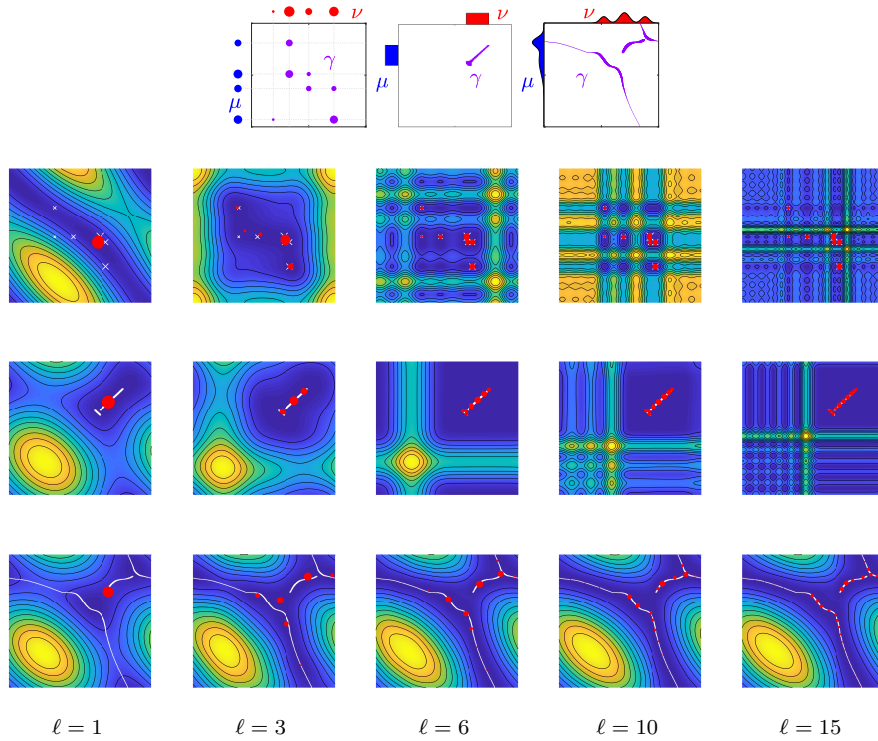


Figure 4.4: Recovery results for transport between sparse (top row), uniform (middle row) and Gaussian measures (bottom row), for increasing orders of relaxation. The positions (and amplitude) of the red dots are recovered by applying the approximate joint diagonalization scheme of Section 1.6.3 to the matrix $T_n(\mathbf{c})$ solving $(\text{OT}^{(\ell)})$. The white crosses indicate the support of the true transport plan, computed by linear programming (over a predefined grid). The surfaces shown are the dual polynomials $\eta^{(\ell)} = \mathcal{F}_\ell^*(\hat{h} - (a \oplus b))$ associated to the dual semidefinite problem $(\text{D-OT}^{(\ell)})$. Both semidefinite programs are solved using SDPT3.




Figure 4.5: We consider here the same sparse and Gaussian transport problems that are illustrated in Figure 4.4. The red continuous (resp. dashed) lines give the rank of the solutions $T_n(\mathbf{c})$ of $(\text{OT}^{(\ell)})$ (resp. of their principal submatrix $T_{n-1}(\mathbf{c})$). We also represent in blue the rank of $T_n(\pi_1(\mathbf{c}))$ (resp. $T_{n-1}(\pi_1(\mathbf{c}))$). Flatness correspond to the points where the dashed line matches the continuous one. We see clearly that the moment sequence of the coupling may be flat while the one of its marginal is not.

constraints, that may not be fully encoded in the relaxation. We illustrate in Figure 4.5 how flatness does or does not occur for both the sparse and the Gaussian examples given in Figure 4.4.

4.2.4 An alternative formulation for \mathcal{W}_1

We conclude this section by detailing an alternative approach for computing the Wasserstein-1 distance

$$\mathcal{W}_1(\mu_1, \mu_2) = \inf_{\gamma \in \mathcal{M}_+(\mathbb{T}^d \times \mathbb{T}^d)} \int_{\mathbb{T}^d \times \mathbb{T}^d} \|\mathbf{x} - \mathbf{y}\|_2 d\gamma(\mathbf{x}, \mathbf{y}) \quad \text{s.t.} \quad \begin{cases} \pi_1(\gamma) = \mu_1 \\ \pi_2(\gamma) = \mu_2 \end{cases}.$$

In the case of \mathcal{W}_1 , it is possible to avoid resorting to the cost approximation described in Section 4.2.1 by rather considering the alternative (dual) formulation (see *e.g.* [Santambrogio, 2015])

$$\mathcal{W}_1(\mu_1, \mu_2) = \max_{f \in \mathcal{C}(\mathbb{T}^d)} \int_{\mathbb{T}^d} f d(\mu_1 - \mu_2) \quad \text{s.t.} \quad \text{Lip}(f) \leq 1. \quad (4.5)$$

This formulation has also the advantage of involving no coupling, thus potentially relieving some of the computational burden as the dual variable is a function over \mathbb{T}^d instead of $\mathbb{T}^d \times \mathbb{T}^d$. By Rademacher's theorem, the 1-Lipschitz ball $\{f \in \mathcal{C}(\mathbb{T}^d); \text{Lip}(f) \leq 1\}$ is equal to the set $\{f \in W^{1,\infty}(\mathbb{T}^d); \|\nabla f\|_\infty \leq 1\}$, where ∇f is the weak derivative of f . If f is a trigonometric polynomial of degree ℓ , with coefficients $(p_{\mathbf{k}}) \in \mathbb{C}^{N(\ell)}$, the existence of a matrix $Q \in \mathcal{M}_{N(\ell)}(\mathbb{C})$ such that

$$\begin{bmatrix} Q & Dp \\ (Dp)^H & 1 \end{bmatrix} \succeq 0 \quad \text{and} \quad T_\ell^*(Q) = \delta_0, \quad (4.6)$$

where $D = \text{Diag}(0, 1, 2, \dots)$, implies that $\|\nabla f\|_\infty \leq 1$ [Dumitrescu, 2017]. Thus, operating the same partial expansion as before, one may build a semidefinite strengthening of (4.5) as

$$p \in \mathbb{C}^{N(\ell)}, Q \in \mathbb{C}^{N(\ell) \times N(\ell)} \quad \langle p, \mathcal{F}_\ell \mu_1 - \mathcal{F}_\ell \mu_2 \rangle \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} Q & Dp \\ (Dp)^H & 1 \end{bmatrix} \succeq 0 \\ Q - \frac{1}{N(\ell)} I_{N(\ell)} \in \mathfrak{T}_{N(\ell)}^\perp \end{cases}. \quad (4.7)$$

Computing the dual of (4.7) then gives an alternative formulation for the primal relaxation (OT^(ℓ)) in the Wasserstein-1 case. It is given by

$$R \in \mathbb{C}^{N(\ell) \times N(\ell)}, z \in \mathbb{C}^{N(\ell)}, \tau \in \mathbb{C} \quad \frac{1}{2} \left(\tau + \frac{\text{Tr } R}{N} \right) \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} R & z \\ z^H & \tau \end{bmatrix} \succeq 0 \\ 2Dz = \mathcal{F}_\ell \mu_1 - \mathcal{F}_\ell \mu_2 \\ R \in \mathfrak{T}_N \end{cases}. \quad (4.8)$$

This formulation is simpler than (OT^(ℓ)) mostly because the matrices involved are much smaller ($N(\ell) \times N(\ell)$) instead of $N_+(\ell)^2 \times N_+(\ell)^2$. However, we do not use \mathcal{W}_1 as a prior for BLASSO in Section 4.6 because it unfortunately leads to degenerate solutions, see Section 4.6.3.

4.3 MMD unbalanced optimal transport

A major constraint of usual transportation metrics is that they are restricted to measures of equal total mass (typically probability distributions), which considerably reduces their range of applications. Unbalanced optimal transport [Benamou, 2003] on the other hand is able to handle un-normalized measures, allowing for mass creation or destruction, which is more consistent in many real applications, including multi-task optimization. In unbalanced formulations, the hard marginal constraints are replaced by a soft penalization, usually using Csiszár f -divergences [Liero et al., 2017]. Given a divergence $D(\cdot|\cdot)$ over $\mathcal{M}(\mathbb{T}^d) \times \mathcal{M}(\mathbb{T}^d)$, the unbalanced transport problem reads

$$\mathcal{W}_D(\mu_1, \mu_2) = \min_{\gamma \in \mathcal{M}(\mathbb{T}^d \times \mathbb{T}^d)} \int_{\mathbb{T}^d \times \mathbb{T}^d} h d\gamma + D(\pi_1(\gamma)|\mu_1) + D(\pi_2(\gamma)|\mu_2). \quad (4.9)$$

Typically, D can be chosen to be a Kullback-Leibler divergence or a χ -square divergence (see Figure 4.6 for an example of the χ -square penalized problem). However, these choices are not suited for semidefinite relaxations, as these divergences cannot be expressed in terms of moments of measures. To this end, we introduce a maximum mean discrepancy (MMD) penalization, controlling the discrepancy between moments of the coupling and its marginals. We also study numerically the stability of the solutions of the resulting problem with respect to the penalization.




Figure 4.6: Illustration of unbalanced optimal transport, with χ -square divergence (computed over a grid). The transport plan $(P_{i,j})$ is computed via linear programming, and the marginals constraint $\sum_j P_{i,j} = a_\mu$ and $\sum_i P_{i,j} = a_\nu$ are penalized in ℓ^2 -norm.

4.3.1 MMD penalization.

Maximum mean discrepancy norms consist in weighted ℓ^2 -norms of sequence of moments. Translation-invariant MMD are defined using Fourier moments. Given measures $\alpha, \beta \in \mathcal{M}(\mathbb{T}^d)$, their respective sequence of trigonometric moments $\mathcal{F}\alpha, \mathcal{F}\beta \in \mathbb{C}^{\mathbb{Z}^d}$ and a diagonal matrix of weights Γ , the MMD discrepancy between α and β reads

$$\text{MMD}(\alpha, \beta) \stackrel{\text{def.}}{=} \|\Gamma(\mathcal{F}\alpha - \mathcal{F}\beta)\|_2^2 \quad (4.10)$$

Typically, the weights Γ set to zero the coefficients beyond some order. We may then introduce the MMD-penalized unbalanced transport problem:

$$\mathcal{W}_\lambda(\mu_1, \mu_2) \stackrel{\text{def.}}{=} \inf_{\gamma \geq 0} \int h d\gamma + \frac{1}{2\lambda} \text{MMD}(\pi_1(\gamma), \mu_1) + \frac{1}{2\lambda} \text{MMD}(\pi_2(\gamma), \mu_2) \quad (4.11)$$

In the following, let

$$\mathcal{F}^\oplus \gamma \stackrel{\text{def.}}{=} [\mathcal{F}\pi_1(\gamma) \quad \mathcal{F}\pi_2(\gamma)] \quad (4.12)$$

be the moments of a coupling along the centered cross, and

$$u \stackrel{\text{def.}}{=} [\mathcal{F}\mu_1 \quad \mathcal{F}\mu_2] \quad (4.13)$$

be the moments of the two marginals. We also denote by \mathcal{F}_ℓ^\oplus and u_ℓ the moments truncated at order ℓ . Problem (4.11) is then equivalent to

$$\mathcal{W}_\lambda(\mu_1, \mu_2) = \inf_{\gamma \geq 0} \int h d\gamma + \frac{1}{2\lambda} \|\Gamma(\mathcal{F}^\oplus \gamma - u)\|_2^2. \quad (\text{OT-MMD}_\lambda)$$

Put under this form, the unbalanced problem shares obvious similarities with the Blasso under positivity constraint, the only differences being the cost function (which is 1 in the positive Blasso case) and the sensing operator, which here returns moments sampled along a cross instead of a full square for the Blasso. These resemblance suffice however to motivate an analysis of the stability of (OT-MMD_λ) , *i.e.* in particular whether the support of the solution of (OT-MMD_λ) remains stable with respect to the solution of (OT-MMD_0) as λ goes to zero [Duval and Peyré, 2015].

The limit $\lambda = 0$. Unlike the usual f -divergences such as Kullback-Leibler, χ -square or even total variation, having $\text{MMD}(\alpha, \beta) = 0$ does not necessarily imply that $\alpha = \beta$. Hence, when setting $\lambda = 0$ in (OT-MMD_λ) , we may not fall back on the classical transport problem (OT) . The problem we obtain instead reads

$$\mathcal{W}_0(\mu_1, \mu_2) = \inf_{\gamma \geq 0} \int h d\gamma \quad \text{s.t.} \quad \Gamma \mathcal{F}^\oplus \gamma = \Gamma u, \quad (\text{OT-MMD}_0)$$

with dual

$$\sup_{\mathbf{a}, \mathbf{b}} \int A d\mu_1 + \int B d\mu_2 \quad \text{s.t.} \quad \begin{cases} A = \mathcal{F}^* \Gamma \mathbf{a} \\ B = \mathcal{F}^* \Gamma \mathbf{b} \\ A \oplus B \leq h \end{cases}$$

and we may derive the optimality relations between the two problems.

Proposition 17. *A transport plan $\gamma_0 = \sum a_i \delta_{(\mathbf{x}_i, \mathbf{y}_i)}$ is solution of (OT-MMD_0) if and only if $\Gamma \mathcal{F}^\oplus \gamma = \Gamma u$ and there exist $A, B \in \text{Im } \mathcal{F}^* \Gamma$ such that*

$$A \oplus B \leq h \quad \text{and} \quad A(\mathbf{x}_i) + B(\mathbf{y}_i) = h(\mathbf{x}_i, \mathbf{y}_i). \quad (4.14)$$

At this point, a natural question is whether the optimal transport γ_* solving the original problem (OT) is also a solution of (OT-MMD_0) . This is obviously true for instance if Γ is strictly positive, *e.g.* Gaussian weights, in which case (OT-MMD_0) is rigorously equivalent to (OT) . Let us assume now that Γ only targets a finite number of moments – in fact, we assume in the following that $(\Gamma_{\mathbf{k}, \mathbf{k}})$ is equal to 1 if $\mathbf{k} \in \Omega_\ell$, and to 0 otherwise, so that $\Gamma \mathcal{F}^\oplus = \mathcal{F}_\ell^\oplus$. The gap between (OT) and (OT-MMD_0) may be bridged when considering sparse transports via the theory of dual certificates. Indeed, checking if γ_* solves (OT-MMD_0) amounts to finding polynomials satisfying (4.14). The existence of such polynomials may be discussed in some simple situations. We always assume that the cost h is of the form $h(\mathbf{x}, \mathbf{y}) = \tilde{h}(\mathbf{x} - \mathbf{y})$ (typically, $h(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^d \sin^2(\pi(x_i - y_i))$).

Proposition 18. *If $\mu_1 = \mu_2 \stackrel{\text{def.}}{=} \sum_{j=1}^s a_j \delta_{\mathbf{x}_j}$, in which case $\gamma_* = \sum_j a_j \delta_{(\mathbf{x}_j, \mathbf{x}_j)}$, then the polynomial $\eta_I \stackrel{\text{def.}}{=} h - (P \oplus P)$ where*

$$P(\mathbf{x}) \stackrel{\text{def.}}{=} - \prod_j \left(\sum_{i=1}^d \sin^2(\pi(x_i - x_{j,i})) \right) \quad \forall \mathbf{x} \in \mathbb{T}^d.$$

is a valid certificate for γ_ .*

Proof. P is nonpositive, and $P(\mathbf{x}_j) = 0$ for all $j = 1, \dots, s$. Thus, (P, P) satisfies the optimality relations (4.14), which suffices to conclude. \square

The certificate η_I is displayed in Figure 4.7. For less trivial situations than $\mu_1 = \mu_2$, it might be more difficult to give an explicit construction of a certificate. However, if μ_1 and μ_2 are sufficiently close, *e.g.* $\mu_1 = \sum_{j=1}^s a_j \delta_{\mathbf{x}_j}$ and $\mu_2 = \sum_{j=1}^s a_j \delta_{\mathbf{x}_j + \varepsilon \mathbf{v}_j}$ for some $\varepsilon \ll 1$, then we expect that a valid certificate could be obtained by slightly deforming η_I .




Figure 4.7: A certificate for the identity map, with Wasserstein-2 cost.

4.3.2 Stability analysis

We now discuss the stability of the solutions of the unbalanced problem (OT-MMD_λ) as λ goes to zero. This analysis involves the solutions of the dual problem, given by

$$\max_{\mathbf{a}, \mathbf{b}} \int A d\mu_1 + \int B d\mu_2 - \frac{\lambda}{2} (\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2) \quad \text{s.t.} \quad \begin{cases} A = \mathcal{F}^* \Gamma \mathbf{a} \\ B = \mathcal{F}^* \Gamma \mathbf{b} \\ A \oplus B \leq h \end{cases} \quad (\text{D-OT-MMD}_\lambda)$$

When $\lambda > 0$, the optimality relations between the above dual problem and its primal (OT-MMD_λ) give an explicit formulation of the dual potentials (P_λ, Q_λ) with respect to a transport plan solving (OT-MMD_λ).

Proposition 19. *For $\lambda > 0$, let $\gamma_\lambda \in \mathcal{M}_+(\mathbb{T}^d \times \mathbb{T}^d)$, let*

$$(\mathbf{a}_\lambda, \mathbf{b}_\lambda) \stackrel{\text{def.}}{=} \frac{1}{\lambda} \mathcal{F}_\ell^{\oplus*} (u_\ell - \mathcal{F}_\ell^{\oplus} \gamma_\lambda). \quad (4.15)$$

and let (A_λ, B_λ) be the corresponding polynomials. Then γ_λ is a solution of (OT-MMD_λ) if and only if $A_\lambda \oplus B_\lambda \leq h$ and $A_\lambda \oplus B_\lambda = h$ on $\text{Supp } \gamma_\lambda$.

Let $\eta_\lambda \stackrel{\text{def.}}{=} h - (A_\lambda \oplus B_\lambda)$ be the dual certificate associated to (D-OT-MMD_λ). It is known that the stability of γ_λ with respect to λ is governed by the minimal norm certificate η_0 (see *e.g.* (2.11) for a definition of the minimal norm certificate in the case of the BLASSO), which is the uniform limit of η_λ as λ goes to zero. In the case where η_0 is tight, *i.e.* does not saturate the constraint outside of the support of the transport plan, then the solutions of (OT-MMD_λ) are stable with respect to λ , which may give important information on their structure, such as the number of Diracs composing them.

However, the minimal norm certificate is difficult to compute in general. To remedy the problem, one may rather consider the vanishing derivatives pre-




Figure 4.8: Pre-certificates for the MMD-penalized OT problem, with uniform weights on the moment ($\Gamma = I_{\Omega_{f_c}}$). Results are shown for translation ($\mathbf{x}_{\mu_2} = \mathbf{x}_{\mu_1} + \varepsilon$), splitting, and warping ($\mathbf{x}_{\mu_2} = (1 - \varepsilon)\mathbf{x}_{\mu_1} + \varepsilon\mathbf{x}_{\mu_1}^3$), with $f_c = 15$ (top row) or $f_c = 50$ (bottom row). The exact optimal transport is displayed in red.

certificates [Duval and Peyré, 2015]

$$\begin{bmatrix} A_V \\ B_V \end{bmatrix} = \underset{\mathbf{a}, \mathbf{b}}{\operatorname{argmin}} \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 \quad \text{s.t.} \quad \begin{cases} A = \mathcal{F}^* \Gamma \mathbf{a} \\ B = \mathcal{F}^* \Gamma \mathbf{b} \\ A(\mathbf{x}_i) + B(\mathbf{y}_i) = h(\mathbf{x}_i, \mathbf{y}_i) \\ A'(\mathbf{x}_i) = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}_i, \mathbf{y}_i) \\ B'(\mathbf{x}_i) = \frac{\partial h}{\partial \mathbf{y}}(\mathbf{x}_i, \mathbf{y}_i) \end{cases} \quad (4.16)$$

which replaces the constraint $A \oplus B \leq h$ by simpler affine constraints corresponding to first order interpolation with the constraints, and may be computed simply by solving a linear system. The polynomial $\eta_V = h - (A_V \oplus B_V)$ is a good candidate to be a minimal norm certificate as it shares many of its properties.

Proposition 20 ([Duval and Peyré, 2015]). *One has $\eta_V \geq 0$ if and only if $\eta_V = \eta_0$. Furthermore, if $\eta_V(\mathbf{x}, \mathbf{y}) < 0$ for some (\mathbf{x}, \mathbf{y}) , then the minimal norm certificate is degenerate.*

Figure 4.8 shows examples of such pre-certificates, which appear to be degenerate in many optimal transport settings. As a consequence, we expect the support of the transport to be less stable as the latter deviate from identity,

and hence to not be exactly retrievable, see Section 4.3.4 for further details. In any case, in the framework of Section 4.6, where Wasserstein distances are used as priors for the BLASSO, this should not matter too much as the goal is to minimize the transport, therefore leading to couplings close to identity.

4.3.3 Semidefinite formulations for MMD transport

In order to solve problems (OT-MMD $_{\lambda}$) or (D-OT-MMD $_{\lambda}$), we resort to their semidefinite approximations. These are straightforward extensions of (OT $^{(\ell)}$) and (D-OT $^{(\ell)}$). Given $c \in \mathbb{C}^{N \times N}$, let

$$\Pi^{\oplus} c \stackrel{\text{def.}}{=} [\Pi_1 c \quad \Pi_2 c]. \quad (4.17)$$

The semidefinite relaxation of (OT-MMD $_{\lambda}$) reads

$$\begin{aligned} \mathcal{W}_{\lambda}^{(\ell)}(\mu_1, \mu_2) = \min_{c \in \mathbb{C}^{N^{(\ell)} \times N^{(\ell)}}} & \langle \hat{h}, c \rangle + \frac{1}{2\lambda} \|\Pi^{\oplus} c - u_{\ell}\|^2 \\ \text{s.t.} & \quad T_n(c) \succeq 0 \end{aligned} \quad (\text{OT-MMD}_{\lambda}^{(\ell)})$$

On the other hand, the dual strengthening of (D-OT-MMD $_{\lambda}$), which is also the dual problem of (OT-MMD $_{\lambda}^{(\ell)}$), is given by

$$\begin{aligned} \sup_{\substack{\mathbf{a}, \mathbf{b} \in \mathbb{C}^{N^{(\ell)}} \\ Q \in \mathbb{C}^{N_+^{(\ell)} \times N_+^{(\ell)}}}} & \langle \mathbf{a}, \mathcal{F}_{\ell} \mu_1 \rangle + \langle \mathbf{b}, \mathcal{F}_{\ell} \mu_2 \rangle - \lambda(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2) \\ \text{s.t.} & \quad \begin{cases} Q \succeq 0 \\ \hat{h} - (\mathbf{a} \oplus \mathbf{b}) = T_n^*(Q) \end{cases} \end{aligned} \quad (\text{D-OT-MMD}_{\lambda}^{(\ell)})$$

Similarly to (4.15), the primal-dual optimality relations give a simple link between the dual potentials \mathbf{a}, \mathbf{b} and the sequence c , when $\lambda > 0$.

Proposition 21. *The primal-dual optimality relations between solutions of (OT-MMD $_{\lambda}^{(\ell)}$) and (D-OT-MMD $_{\lambda}^{(\ell)}$) reads*

$$\begin{cases} \mathbf{a} = \frac{1}{2\lambda}(u - \Pi_1 c) \\ \mathbf{b} = \frac{1}{2\lambda}(v - \Pi_2 c) \end{cases}.$$

4.3.4 Numerical illustrations

We solve both problems (OT-MMD $_{\lambda}^{(\ell)}$) and (D-OT-MMD $_{\lambda}^{(\ell)}$) using SDPT3 [Toh et al., 1999]. All the tests of this section concern transports between 1-D measures, with a hierarchy order set to $\ell = 13$.




Figure 4.9: Illustration of the MMD-penalized unbalanced transport problem ($\text{OT-MMD}_\lambda^{(\ell)}$) when λ goes to zero, in the splitting example of Figure 4.6. The semidefinite optimization is solved using `cvx`, and the measure (red dots) is recovered using the multivariate Prony’s algorithm 2. We display the dual polynomial $\hat{h} - (\mathbf{a} \oplus \mathbf{b})$. We observe that when λ goes to zero, the correct transport is recovered, while the dual polynomial is degenerate.




Figure 4.10: Evolution of dual polynomial and recovered measure as λ increases. For high values of λ , the dual polynomial saturates on the diagonal.

Regime $\lambda \gg 1$. As λ goes to infinity, the dual potentials \mathbf{a} and \mathbf{b} tend to zero, so that the limit of the dual polynomial $h - (A \oplus B)$ (where $A \stackrel{\text{def.}}{=} \mathcal{F}_\ell^* \mathbf{a}$ and $B \stackrel{\text{def.}}{=} \mathcal{F}_\ell^* \mathbf{b}$) is h . Consequently, as our transport cost is h zero on the diagonal $\mathbf{x} = \mathbf{y}$, one should expect that the measure recovered from the primal variable \mathbf{c} in this regime be located on this diagonal. This is what we observe, see *e.g.* Figure 4.10. This behavior contrasts with Csiszár divergences penalizations, which are harder and impose more constraints on the transport plan. For instance, in the case of sparse transport, the plan remains located on the grid defined by the marginals, leading to a staircase structure, see Figure 4.6.

Regime $\lambda \ll 1$ and degeneracy For the MMD-penalized optimal transport problem (OT-MMD_λ), numerical simulations indicate that there is no stability of the solutions in most cases. Figure 4.8 shows the pre-certificates η_V for various configurations of transport plans. In all cases shown, η_V is degenerate, indicating that the minimal norm certificate should be too. This is coherent with what we observe in Figure 4.9, for the splitting example, where the minimal norm certificate saturates on a whole line.

4.4 SDP transport in higher dimensions

In this section we detail the FFW algorithm for solving the unbalanced optimal transport problem (OT-MMD $_{\lambda}^{(\ell)}$), and provide some numerical results of transport between 2-dimensional measures. Note that in this case, the plan to retrieve is 4-dimensional, making the size of the SDPs prohibitive for standard algorithms.

4.4.1 Problem statement

We make use of the same notations as in Chapter 3 for the Toeplitz projector Θ (3.1) and the Toeplitz values operator θ (3.2). The FFW algorithm introduced in Chapter 3 works by penalizing the Toeplitz constraint in (OT-MMD $_{\lambda}^{(\ell)}$). Therefore, we consider the following energy to minimize:

$$f_{\lambda,\rho}(R) \stackrel{\text{def.}}{=} \langle \hat{h}, \theta(R) \rangle + \frac{1}{4\lambda} \|\Pi^{\oplus} \theta(R) - u_{\ell}\|^2 + \frac{1}{2\rho} \|\Theta_{\perp}(R)\|^2. \quad (4.18)$$

where $R \in \mathcal{H}_{N_+(\ell)}$. To simplify notations, let $\Lambda \stackrel{\text{def.}}{=} \Pi^{\oplus} T$. For a given order ℓ , we want to solve the optimization problem

$$\min_{R \in \mathbb{C}^{N_+(\ell) \times N_+(\ell)}} f_{\lambda,\rho}(R) \quad \text{s.t.} \quad R \succeq 0. \quad (4.19)$$

We may derive the dual problem and optimality relations for (4.19).

Proposition 22. *The dual problem of (4.19) reads*

$$\begin{aligned} \max_{\mathbf{a}, \mathbf{b}} \quad & \langle \mathcal{F}_{\ell} \mu_1, \mathbf{a} \rangle + \langle \mathcal{F}_{\ell} \mu_2, \mathbf{b} \rangle - \frac{\lambda}{2} \|\mathbf{a}\|^2 - \frac{\lambda}{2} \|\mathbf{b}\|^2 - \frac{\rho}{2} \|Q - T_n(\hat{h} - \mathbf{a} \oplus \mathbf{b})\|^2 \\ \text{s.t.} \quad & \begin{cases} Q \succeq 0 \\ \hat{h} - \mathbf{a} \oplus \mathbf{b} = T_n^*(Q) \end{cases} \end{aligned} \quad (4.20)$$

At optimality, the solutions R and (\mathbf{a}, \mathbf{b}) satisfies

$$\begin{cases} \frac{1}{2\lambda} (\Pi_1 \theta(R) - \mathcal{F}_{\ell} \mu_1) + \mathbf{a} = 0 \\ \frac{1}{2\lambda} (\Pi_2 \theta(R) - \mathcal{F}_{\ell} \mu_2) + \mathbf{b} = 0 \\ Q + \theta^*(\hat{h} - \mathbf{a} \oplus \mathbf{b}) = -\frac{1}{\rho} \Theta_{\perp}(R) \end{cases} .$$

4.4.2 FFW algorithm for optimal transport

The gradient of the objective $f_{\lambda,\rho}$ can be straightforwardly derived. It reads

$$\nabla f_{\lambda,\rho}(R) = \theta^*(\hat{h}) + \frac{1}{2\lambda} \Lambda^*(\Lambda(R) - u_{\ell}) + \frac{1}{\rho} \Theta_{\perp}(R). \quad (4.21)$$




Figure 4.11: Effect of the extended line-search on a toy example. We minimize $\min_{X \succeq 0} \|X - X_0\|^2$ with $X_0 \succeq 0$, and where X and X_0 are real symmetric of size 2×2 . The displays on the left give a visual representation of the PSD cone with trace bounded by 6, and of the Frank-Wolfe iterates, in the case of a bounded line-search (top) and extended line-search (bottom). The right plot shows the corresponding decrease in energy. On this simple example, we see that the extended line-search improves significantly the convergence rate of the algorithm.

As before, the variable $R \in \mathcal{H}_{N_+(\ell)}$ may be stored under the form $R = UU^H$ throughout the iterations, where $U \in \mathbb{C}^{N_+(\ell) \times k}$, thus enabling the Fast Fourier Transform machinery detailed in Chapter 3. In the following, let $U \in \mathbb{C}^{N_+(\ell) \times k}$ be the k -th iterate, $v \in \mathbb{C}^{N_+(\ell)}$ be the vector selected by the linear oracle at the k -iteration, and $R = UU^H$, $S = vv^H$ be the corresponding semidefinite iterates.

Extended line-search Contrarily to the BLASSO case (see Remark 16), we do not have a simple explicit bound on the norm of the solutions of (4.19). Nonetheless, Frank-Wolfe algorithm readily extends to an “unbounded” setting. Seeing the linear oracle of Frank-Wolfe as the extraction of a *direction* rather than an atom, thus independent of the bounds on the domain, the linesearch over (α, β) to determine a couple minimizing $f_{\lambda, \rho}(\alpha R + \beta S)$ (cf. Section 3.3 and Algorithm 6) may then be performed without imposing that $\alpha + \beta \leq 1$, as it is done in usual implementations of Frank-Wolfe. This extended line-search may actually improve the rate of convergence of the algorithm compared to the bounded version, as illustrated in Figure 4.11 on a toy example.

The next proposition gives the closed-form formula for the extended line-

search step when applied on problem (4.19), in analogy with Proposition 13.

Proposition 23. *Let $c_1(v) \stackrel{\text{def.}}{=} \Pi_1 \boldsymbol{\theta}(vv^H)$ and $c_2(v) \stackrel{\text{def.}}{=} \Pi_2 \boldsymbol{\theta}(vv^H)$. Let*

$$\begin{cases} t_{11} \stackrel{\text{def.}}{=} \frac{1}{4\lambda} (\|c_1(U)\|^2 + \|c_2(U)\|^2) + \frac{1}{2\rho} \|\boldsymbol{\Theta}_\perp(R)\|^2 \\ t_{22} \stackrel{\text{def.}}{=} \frac{1}{4\lambda} (\|c_1(v)\|^2 + \|c_2(v)\|^2) + \frac{1}{2\rho} \|\boldsymbol{\Theta}_\perp(S)\|^2 \\ t_{12} \stackrel{\text{def.}}{=} \frac{1}{2\lambda} (\langle c_1(U), c_1(v) \rangle + \langle c_2(U), c_2(v) \rangle) + \frac{1}{\rho} (\langle \boldsymbol{\Theta}_\perp(R), \boldsymbol{\Theta}_\perp(S) \rangle) \\ t_1 \stackrel{\text{def.}}{=} \langle \hat{h}, \boldsymbol{\theta}(R) \rangle - \frac{1}{2\lambda} \langle c_1(U), \mathcal{F}_\ell \mu_1 \rangle - \frac{1}{2\lambda} \langle c_2(U), \mathcal{F}_\ell \mu_2 \rangle \\ t_2 \stackrel{\text{def.}}{=} \langle \hat{h}, \boldsymbol{\theta}(S) \rangle - \frac{1}{2\lambda} \langle c_1(v), \mathcal{F}_\ell \mu_1 \rangle - \frac{1}{2\lambda} \langle c_2(v), \mathcal{F}_\ell \mu_2 \rangle \end{cases}$$

and let

$$\begin{aligned} \tau_1 &\stackrel{\text{def.}}{=} \max(0, -\frac{t_1}{2t_{22}}), & \tau_2 &\stackrel{\text{def.}}{=} \max(0, -\frac{t_2}{2t_{11}}) \\ \kappa_1 &\stackrel{\text{def.}}{=} \max(0, \frac{t_{12}t_2 - 2t_{22}t_1}{4t_{11}t_{22} - t_{12}^2}), & \kappa_2 &\stackrel{\text{def.}}{=} \max(0, \frac{t_{12}t_1 - 2t_{11}t_2}{4t_{11}t_{22} - t_{12}^2}). \end{aligned}$$

The solution (α_k, β_k) of the extended line-search

$$\min_{\alpha, \beta} f_{\lambda, \rho}(\alpha R + \beta S) \quad \text{s.t.} \quad \alpha, \beta \geq 0$$

is given by

$$(\alpha_k, \beta_k) = \begin{cases} (0, \tau_2) & \text{if } t_{11} = 0 \quad \text{and} \quad t_{22} \neq 0 \\ (\tau_1, 0) & \text{if } t_{22} = 0 \quad \text{and} \quad t_{11} \neq 0 \\ (\kappa_1, \kappa_2) & \text{otherwise} \end{cases}.$$

4.4.3 Numerical illustrations

Numerical setting. We consider transports between several types of distributions: discrete, Gaussian, uniform on the circle and uniform on the disc. For the Gaussian and uniform densities, moments which serve as input to the algorithm are generated from discrete approximations, with high number of Diracs (10000). The spikes are positive, have all the same mass, and are in equal numbers in both input marginals (the transport is balanced).

Scaling of the parameters. The parameters λ and ρ involved in FFW are fixed in our tests. As before, we choose specific scalings for these parameters (see Section 3.4). In mirror of the BLASSO, we scale λ proportionally to $\|\mathcal{F}_\ell^{\oplus*} u_\ell\|_\infty$. Furthermore, we scale ρ with respect to $N_+(\ell)$, in order to limit the weight of the Toeplitz penalization term when ℓ increases (as the size of the involved matrices grows as ℓ^4 in that case). In the tests of this section, we set $\lambda = 10^{-3} \|\mathcal{F}_\ell^{\oplus*} u_\ell\|_\infty$ and $\rho = 10^{-1} N_+(\ell)$.

Support recovery. Extraction of a support from the matrix recovered by Frank-Wolfe is performed using the Prony’s approach described in Chapter 1. Since we consider transport between 2-D measures, the transport plans are 4-D. Thus, each 4-D point in the reconstructed measure corresponds to two connected (*i.e.* mass is transported from one to the other) 2-D points, one in each marginal. We also make use of our optimization-based joint diagonalization step (see Section 1.6.3) as we consider optimal transport between non-sparse measures.

Sparse transport. The FFW algorithm behaves similarly on the optimal transport problem between sparse measures than it does on the BLASSO (see Section 3.4). We observe a fast convergence after only a number of iterations corresponding roughly to the number of points in the coupling (see Figure 4.12).

Transport between non-sparse measures. For non-sparse measures, the FFW algorithm usually reconstructs in the end as many points as there has been iterations (20 in our examples), as a discretization of the marginals. Outliers may appear, but usually correspond to points with low amplitudes (see Figure 4.13). Although we do not impose the amplitudes to be positive when solving the least-square Vandermonde system (1.9) after having retrieved the support, in our examples the points we recover have all positive amplitudes.

Limitations. A main limitation of our approach remains the size of the involved matrices, which limits the cutoff frequency ℓ which can be used when optimizing over 4-D domains. As a result, and as we can see in Figure 4.13, although the marginals are correctly recovered, complicated transport maps suffer from a crude approximation. These poor performances on more complicated examples are due to the use of a too low value of ℓ . In spite of our efficient storage and computations (see Section 3.3.3), the algorithm still becomes slow when considering values of ℓ moderately large, constraining us to relatively small values (in the example of Figure 4.13 for instance, we use moments up to order $\ell = 10$). Compared to the proximity of the measures with each other, these values might not be sufficient to capture the correct mappings. The possibility of improving further the efficiency of our computations is still an open question, which could allow to solve more robustly problems with large values of ℓ .

4.5 Group-Blasso

In the rest of this thesis, we consider an extension of the BLASSO super-resolution problem studied in Chapter 2, which we call the “multi-task” sparse super-resolution problem. The goal is to simultaneously recover pointwise sources across several similar channels, given some low-pass measurements. Indeed, in practice, the sources to recover may be linked to physical or biological phenomena, and one may wish to constrain their positions across the observable data. For example, in brain imaging, they may be linked to a location in the brain




Figure 4.12: Example of sparse transport reconstruction with the FFW algorithm. The first plot shows the evolution of the energy. The second plot displays the error $\|R - R_0\|/\|R_0\|$ along the iterations, where $R = UU^H$ is the solution of FFW and R_0 is the (truncated) moment matrix of the true transport. Since R_0 is not a solution of the problem solved by FFW (due to the Toeplitz penalization), this quantity does not tend to zero.

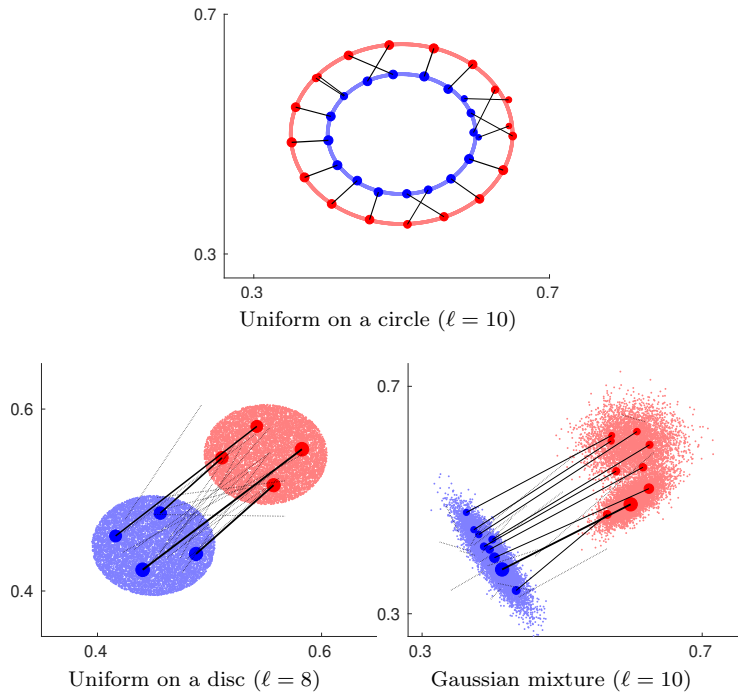


Figure 4.13: Examples of “continuous” transports recovery with the FFW algorithm. Links between points in dashed lines correspond to points with relative amplitudes (with respect to the maximum amplitude) less than $2 \cdot 10^{-1}$. These results are obtained after 20 FFW iterations.

[Owen et al., 2009], or to a position in the genome in genomics [Laurent, 2010]. Taking advantage of the cross-information coming from several measurements may allow to statistically enhance the performance of the recovery. This is important when dealing with high noise levels.

Group Blasso When the sought after measures have the same support, a popular approach to tackle such problems is the Group-LASSO [Yuan and Lin, 2006], which promotes group-sparsity by using a ℓ^1/ℓ^2 penalization, defined for a matrix $A \stackrel{\text{def.}}{=} [a_1 \ \dots \ a_t]$ as

$$\|A\|_{\ell^1/\ell^2} = \sum_j \|a_j\|_2 \quad (4.22)$$

where a_j is the j -th column of A . The group-LASSO favors solutions with only a few non-zero column vectors, and may thus recover simultaneously several sparse signals sharing a common support.

The off-the-grid generalization of this ℓ^1/ℓ^2 -norm is the well-known vectorial total variation norm [Ambrosio et al., 2000], defined for a vector-valued measure $\boldsymbol{\mu} = (\mu_1, \dots, \mu_t)$ as

$$|\boldsymbol{\mu}|((\mathbb{T}^d)^t) \stackrel{\text{def.}}{=} \sup_{\eta_1, \dots, \eta_t \in \mathcal{C}(\mathbb{T}^d)} \left\{ \Re\left(\sum_i \int_{\mathbb{T}^d} \bar{\eta}_i d\mu_i\right); \forall \mathbf{x} \in \mathbb{T}^d, \sqrt{\sum_i \eta_i(\mathbf{x})^2} \leq 1 \right\} \quad (4.23)$$

As TV compared to ℓ^1 , the vectorial TV norm is the natural extension of the ℓ^1/ℓ^2 norm to infinite-dimensional spaces. In particular, if (μ_1, \dots, μ_t) are discrete with extended amplitude vectors (a_1, \dots, a_t) , that is, the amplitudes of the measures over their joint support ($\text{Supp } \mu_1 \cup \dots \cup \text{Supp } \mu_t$), then $|\boldsymbol{\mu}|((\mathbb{T}^d)^t) = \|A\|_{\ell^1/\ell^2}$, where $A = [a_1 \ \dots \ a_t]$.

To perform the joint super-resolution of several similar signals, one may then consider the following natural extension of the BLASSO [Fernandez-Granda, 2015, Yang and Xie, 2016, Li and Chi, 2016], which we refer to as “group-BLASSO”,

$$\inf_{\mu_1, \dots, \mu_t \in \mathcal{M}(\mathbb{T}^d)} \frac{1}{2\lambda t} \sum_{i=1}^t \|y_i - \Phi \mu_i\|^2 + |\boldsymbol{\mu}|((\mathbb{T}^d)^t). \quad (\text{G-Blasso})$$

Figure 4.14 shows a toy example of recovery by group-BLASSO.

Recovering measures with disjoint support As the BLASSO, the group-BLASSO may be solved via a moment, enabling the algorithms studied in Chapter 2 and Chapter 3. However, a major limitation of the group methods is that they demand a perfect overlap between the support of the signals, see Figure 4.15. This hypothesis is unrealistic in practice: biological measurements for instance vary from a subject to another. To remedy this problem, the ℓ^1/ℓ^2 penalty in the group Lasso may be replaced by a penalization of an optimal transport distance between the different tasks [Janati et al., 2019b,




Figure 4.14: Recovery of two measures with same support from noisy observations, independently by two independent BLASSO (up) and jointly by group-BLASSO (bottom). The joint method recovers a spike of low amplitude while BLASSO fails to do so. The noise ratio is $\|w\| = 0.15\|\mathcal{F}\mu_0\|$, and $\lambda = 1.5$.




Figure 4.15: Joint recovery of two measures with distinct support, using group Blasso. The method tends to recover the full extended support, *i.e.* $\text{Supp } \mu_1 \cup \text{Supp } \mu_2$. Furthermore, in this particular case, it fails to retrieve the low amplitude spike.

[Janati et al., 2019a]. We discuss in the next section an off-the-grid extension of this idea.

4.6 Wasserstein-Blasso

Optimal transport distances are a natural tool to encode geometric similarities between measures. Inspired by an idea introduced in [Janati et al., 2019b] for the LASSO (hence over a grid), we investigate in this section the use of a Wasserstein coupling between several BLASSO.

4.6.1 A Wasserstein model for the 2-task Blasso

To clarify the exposition, we assume in the following that we are looking for two (close) discrete positive Radon measures μ_1 and μ_2 given two observations $y_1 = \Phi\mu_1 + w$ and $y_2 = \Phi\mu_2 + \varepsilon$. The extension to more sources is straightforward and discussed in the next section. As in Chapter 2, we consider a forward operator Φ of the form $\mathcal{A}\mathcal{F}_{f_c}$, for some cutoff frequency $f_c \in \mathbb{N}^*$.

As before, we denote by Φ^\oplus the operator on $\mathcal{M}(\mathbb{T}^d \times \mathbb{T}^d)$ such that $\Phi^\oplus \gamma \stackrel{\text{def.}}{=} [\Phi\pi_1(\gamma) \ \Phi\pi_2(\gamma)]$. Let also $y \stackrel{\text{def.}}{=} [y_1 \ y_2]$. We introduce in this section the Wasserstein BLASSO

$$\inf_{\mu_1, \mu_2 \in \mathcal{M}_+(\mathbb{T}^d)} \frac{1}{2} \sum_{i=1,2} \left[\frac{1}{2\lambda} \|\Phi(\mu_i) - y_i\|^2 + \mu_i(\mathbb{T}^d) \right] + \tau \mathcal{W}(\mu_1, \mu_2). \quad (\text{W-BLASSO}_{\lambda, \tau})$$

Considering the couplings $\gamma \in \mathcal{M}(\mathbb{T}^d \times \mathbb{T}^d)$ whose marginals are μ_1 and μ_2 , we may formulate $(\text{W-BLASSO}_{\lambda, \tau})$ as an optimization problem over product measures

$$\inf_{\gamma \in \mathcal{M}_+(\mathbb{T}^d \times \mathbb{T}^d)} \frac{1}{4\lambda} \|\Phi^\oplus \gamma - y\|^2 + \int (\tau h + 1) d\gamma, \quad (4.24)$$

where as before we assume that $h : \mathbb{T}^d \times \mathbb{T}^d \rightarrow \mathbb{R}_+$ is a trigonometric polynomial. By compactness and lower semi-continuity, problem (4.24) has indeed a minimum. Note that the formulation (4.24) highlights the resemblance of $(\text{W-BLASSO}_{\lambda, \tau})$ with MMD-penalized OT (OT-MMD_λ), and more generally with the positive BLASSO. We may derive the dual problem.

Proposition 24. *The dual of $(\text{W-BLASSO}_{\lambda, \tau})$ reads*

$$\max_{p_1, p_2} \sum_{i=1,2} [\langle y_i - 2\lambda\delta_0, p_i \rangle - \lambda\tau \|p_i\|^2] \quad \text{s.t.} \quad \begin{cases} \varphi_1 = \Phi^* p_1 \\ \varphi_2 = \Phi^* p_2 \\ \varphi_1 \oplus \varphi_2 \leq h \end{cases} \quad (\text{D-W-BLASSO})$$

where δ_0 is the vector such that $\Phi^* \delta_0 = 1$.

Proof. Applying Fenchel duality to problem (4.24) straightforwardly yields the dual problem

$$\sup_p \langle y, p \rangle - \lambda \|p\|^2 \quad \text{s.t.} \quad (\Phi^\oplus)^* p \leq \tau h + 1,$$

where $p \stackrel{\text{def.}}{=} [p_1 \ p_2] \in$ and $(\Phi^\oplus)^* p = \Phi^* p_1 \oplus \Phi^* p_2$. Applying the change of variable $p_i \leftarrow \tau p_i + \frac{\delta_0}{2}$ then yields the desired result. \square

Unlike the generic optimal transport problem, the presence in $(\text{W-BLASSO}_{\lambda, \tau})$ of the bandlimited operator Φ (or Φ^* for the dual) make these problems natural candidates for semidefinite hierarchies. Indeed, the primal $(\text{W-BLASSO}_{\lambda, \tau})$ only involves a finite number of moments of its variables, and the dual potentials φ_1 and φ_2 are truly trigonometric polynomials in that case.

Semidefinite hierarchies for Wasserstein Blasso The semidefinite relaxations of $(\text{W-BLASSO}_{\lambda, \tau})$ (or strengthenings for the dual) are built from the formulation (4.24) over product measures. In case of finite convergence, the moments returned by the hierarchy are those of the transport plan γ from μ_1 to μ_2 , the sought measures. In the following, let $\ell \geq \max(f_c, \deg(h))$. This assumption ensures that the number of moments involved in the ℓ -th level of the

hierarchy is at least equal to the number of moments involved in the objective to minimize.

Let Π^\oplus be the projection operator (4.17). The semidefinite hierarchy for (W-BLASSO $_{\lambda,\tau}$) at order ℓ can be straightforwardly deduced from the relaxations of the positive BLASSO (2.31). It reads

$$\inf_{\mathbf{c} \in \mathbb{C}^{N^{(\ell)} \times N^{(\ell)}}} \frac{1}{4\lambda} \|\Pi^\oplus \tilde{\mathbf{c}} - \mathbf{y}\|_2^2 + c_{0,0} + \tau \langle \hat{h}, \mathbf{c} \rangle \quad \text{s.t.} \quad \begin{cases} T_\ell(\mathbf{c}) \succeq 0 \\ \tilde{c}_{\mathbf{k},\mathbf{l}} = c_{\mathbf{k},\mathbf{l}} \quad \forall \mathbf{k}, \mathbf{l} \in \Omega_{f_c} \end{cases} \quad (\text{W-BLASSO}_{\lambda,\tau}^{(\ell)})$$

We state without proof the traditional convergence result for these relaxations, as the methodology is the same as for the BLASSO.

Proposition 25. *For $\ell \geq \max(f_c, \deg(h))$, one has*

$$(\text{W-BLASSO}_{\lambda,\tau}^{(\ell)}) \leq (\text{W-BLASSO}_{\lambda,\tau}^{(\ell+1)}) \leq (\text{W-BLASSO}_{\lambda,\tau}).$$

Furthermore, $\lim_{\ell \rightarrow \infty} (\text{W-BLASSO}_{\lambda,\tau}^{(\ell)}) = (\text{W-BLASSO}_{\lambda,\tau})$.

Similarly, the sos strengthening of the dual problem (D-W-BLASSO) is

$$\begin{aligned} \max_{\mathbf{p}_1, \mathbf{p}_2 \in \mathbb{C}^{N(f_c)}} \quad & \sum_{i=1,2} \langle y_i - \lambda \delta_0, \mathbf{p}_i \rangle - \lambda \tau \|\mathbf{p}_i\|^2 \\ \text{s.t.} \quad & \begin{cases} Q \succeq 0 \\ (\tilde{p}_{1,2})_{\mathbf{k}} = \begin{cases} (p_{1,2})_{\mathbf{k}} & \text{if } \mathbf{k} \in \Omega_{f_c} \\ 0 & \text{otherwise} \end{cases} \\ \hat{h} - (\tilde{\mathbf{p}}_1 \oplus \tilde{\mathbf{p}}_2) = T_\ell^*(Q) \end{cases} \end{aligned} \quad (\text{D-W-BLASSO}_{\lambda,\tau}^{(\ell)})$$

4.6.2 Multi-marginals extension

One may straightforwardly adapt (W-BLASSO $_{\lambda,\tau}$) to the case where there are more than two unknown measures by introducing a barycenter variable $\bar{\mu}$. Assuming that we are given M observations y_1, \dots, y_M coming from M sources μ_1, \dots, μ_M , the multi-marginals Wasserstein BLASSO reads

$$\inf_{\{\mu_i\}, \bar{\mu} \in \mathcal{M}_+(\mathbb{T}^d)} \frac{1}{M} \sum_{i=1}^M \left[\frac{1}{2\lambda} \|\Phi \mu_i - y_i\|^2 + \mu_i(\mathbb{T}^d) + \tau \mathcal{W}(\mu_i, \bar{\mu}) \right],$$

and can be formulated more simply over coupling measures as

$$\begin{aligned} \inf_{\{\gamma_i\} \in \mathcal{M}_+(\mathbb{T}^d \times \mathbb{T}^d)} \quad & \frac{1}{M} \sum_{i=1}^M \left[\frac{1}{2\lambda} \|\Phi \pi_1(\gamma_i) - y_i\|^2 + \tau \int h d\gamma_i \right] + \gamma_1(\mathbb{T}^d \times \mathbb{T}^d) \\ \text{s.t.} \quad & \pi_2(\gamma_i) = \pi_2(\gamma_j) \quad \forall i, j \end{aligned}$$

4.6.3 The particular case of \mathcal{W}_1

We illustrate in this section the particular case of using the Wasserstein-1 distance as a prior for the BLASSO. We show that in this case, the solution is always degenerate, *i.e.* that the marginals must satisfy $\mu_1 = \mu_2$.

Recall the dual formulation (4.5) of \mathcal{W}_1 over Lipschitz functions. Then, the \mathcal{W}_1 -penalized BLASSO may be put under the saddle-point form

$$\min_{\mu_1, \mu_2 \in \mathcal{M}_+(\mathbb{T}^d)} \max_{f \in \mathcal{C}(\mathbb{T}^d)} \frac{1}{2} \sum_{i=1,2} \frac{1}{2\lambda} \|\Phi \mu_i - y_i\|^2 + \int f d(\mu_1 - \mu_2) + i_{B_{\text{Lip}1}}(f), \quad (4.25)$$

where $i_{B_{\text{Lip}1}}$ is the indicator of the ball

$$\{f \in W^{1,\infty}(\mathbb{T}^d); \|\nabla f\|_\infty \leq 1\}.$$

Introducing the Lagrange variable $\mathbf{m} = (m_1, \dots, m_d) \in (\mathcal{M}_+(\mathbb{T}^d))^d$ associated to the indicator constraint, problem (4.25) becomes

$$\min_{\mu_1, \mu_2 \in \mathcal{M}_+} \max_{\substack{f \in \mathcal{C}(\mathbb{T}^d) \\ \mathbf{m} \in \mathcal{M}_+^d}} \frac{1}{2} \sum_{i=1,2} \frac{1}{2\lambda} \|\Phi \mu_i - y_i\|^2 + \int_{\mathbb{T}^d} f d(\mu_1 - \mu_2) - \langle \nabla f, \mathbf{m} \rangle + |\mathbf{m}|((\mathbb{T}^d)^d),$$

where $|\mathbf{m}|((\mathbb{T}^d)^d)$ is the vectorial total variation norm of \mathbf{m} , defined in (4.23). Differentiating with respect to μ_1, μ_2, f and \mathbf{m} respectively leads to the following optimality relations

$$\begin{cases} \frac{1}{2\lambda} \Phi^*(\Phi \mu_1 - y_1) + f = 0 \\ \frac{1}{2\lambda} \Phi^*(\Phi \mu_2 - y_2) - f = 0 \\ \mu_1 - \mu_2 + \text{div}(\mathbf{m}) = 0 \\ -\nabla f \in \partial |\mathbf{m}|(X) \end{cases}. \quad (4.26)$$

Thus, at optimality, $f \in \text{Im } \Phi^*$ is a trigonometric polynomial, and so is ∇f . The gradient ∇f cannot be constant and non null since this would contradict the fact that f is a trigonometric polynomial. Excluding the degenerate case where $\nabla f = 0$ (*i.e.* f is constant), \mathbf{m} is therefore a discrete measure, since from the last optimality relation it must be supported within the set of roots of ∇f , which is finite. On the other hand, the measure $\mu_1 - \mu_2$ is the divergence of \mathbf{m} . This necessarily implies that $\mathbf{m} = 0$, so that, $\mu_1 = \mu_2$.

4.6.4 Numerical illustrations

We illustrate the use of the Wasserstein-BLASSO, with \mathcal{W}_2 as a prior, on simple examples in 1-D. The goal of these experiments is to showcase the qualitative behavior of the method, and in particular its ability to leverage some localization property without imposing the strong support equality of the group

BLASSO. Benchmarking this method on real data sets deserves further study for future works.

We use SDPT3 [Toh et al., 1999] to solve the semidefinite program (W-BLASSO $_{\lambda, \tau}^{(\ell)}$), in which we use the approximation (4.2) for the \mathcal{W}_2 cost. We set the cutoff frequency of the operator Φ to be $f_c = 11$, and the level of the hierarchy at the same value ($\ell = f_c$).

The main feature of the Wasserstein-BLASSO is that it interpolates in some sense between the behavior of two independent BLASSO and the one of the group-BLASSO. Increasing the value of τ progressively forces the two measures to have the same support (which is the behavior of the group BLASSO, obtained for $\tau = \text{inf}$), see Figure 4.16.

The Wasserstein-BLASSO improves over two independent BLASSO (corresponding to $\tau = 0$) in noisy regimes, or for recovering spikes with low amplitudes. Taking advantage of the cross information between the multiple inputs, Wasserstein-BLASSO is able to outperform the BLASSO on these difficult tasks, see *e.g.* Figure 4.17. This requires the tuning of the parameter τ , to find the correct trade-off between transport and deviation from the input measures. On the example displayed on Figure 4.17, the optimal parameter (according to the flat norm error to the input measure) is $\tau \approx 10^2$.




Figure 4.16: Illustration of the 1-D Wasserstein-BLASSO ($\text{W-BLASSO}_{\lambda, \tau}$) on a toy example: only one spike in each measure, with same amplitude, and noiseless measurements. On top row, we use uniform weights on the moments ($\Phi = \mathcal{F}_c$). On bottom row, we use Gaussian weights on the moments ($\Phi = \mathcal{A}(\varphi_G)\mathcal{F}_c$, see Section 2.3.2 for further details). The behavior of the solution interpolates between two separate BLASSO (left) and the group-BLASSO (right), where the two reconstructed measures have the same (extended) support. When the spacings between the two input measures is too large, or the filter not smooth enough, spurious additional spikes might appear, as shown on the top row for $\tau \geq 10^5$ and the bottom row for $\tau = 10^7$. The bottom row, corresponding to a Gaussian-like filtering, shows that the value of $\tau = 10^5$ reaches a sweet spot, where the support of the two measures almost coincide, and no spurious spike appears. This shows the advantage of the Wasserstein-BLASSO over the group-BLASSO.




Figure 4.17: Results of Wasserstein-BLASSO compared to two independent BLASSO on an example with noise (top figure). The observations y_i , $i = 1, 2$ consists in Fourier moments of the inputs between $-f_c$ and f_c . The noise ratio is set to $\|w_i\| = 0.5\|y_{0,i}\|$ where $y_i = y_{0,i} + w_i$. The parameter λ controlling the marginal constraints is set to $2 \cdot 10^{-2} (\|\mathcal{F}_c^* y_1\|_\infty + \|\mathcal{F}_c^* y_2\|_\infty)$. We see that the Wasserstein-BLASSO performs a better recovery of the spikes of small amplitude, and has less spurious spikes than the BLASSO.

Conclusion

Providing highly efficient algorithms to perform tasks as common and useful as recovering deteriorated or incomplete signals, or comparing complex objects such as probability distributions or shapes, has become an ubiquitous matter in the recent context of data science. Due to recent promising theoretical advances, off-the-grid methods are becoming mainstream in several applications, but they still face computational bottlenecks. The work presented in this thesis lifts some of these bottlenecks while maintaining theoretical guarantees.

Take-home messages

The contributions of this thesis put together several key ideas which work well hand in hand to define an integrated workflow. Let us recap what we think are these main ingredients.

Lasserre’s hierarchy for convex optimization. While Lasserre’s hierarchies have been successful to attack non-convex control problems or NP-hard combinatorial problems, it might seem surprising that they are also useful to tackle convex problems (such as super-resolution or optimal transport), which may appear simpler at first sight. Our work highlights the fact that these hierarchies are also a good fit for these imaging problems, where the degree of the polynomials involved are large but where one can leverage the convolutional structure of the problem.

Breaking the limits of Prony. A simple improvement of Prony’s method performs very well in a regime where no theoretical guarantee exist. We found that the crucial step in Prony’s type methods is the joint-diagonalization of non-symmetric matrices. This step requires some special care, and developing dedicated non-convex optimization solvers is the key to be able to use safely these methods outside of the comfort zone of moment relaxations. This broadens the scope of these approaches, making them usable to retrieve non-sparse measures (for instance supported on curves or surfaces), which also motivates our investigation of optimal transport applications.

Interplay between convex and non-convex methods. Hybridizing convex algorithms with non-convex updates may offer the best of both worlds, simple and low-cost computations on the one hand and faster convergence on the other. Alternating Frank-Wolfe updates with descent steps over the manifold of fixed-rank matrices in the FFW algorithm proves remarkably effective for super-resolution, reaching convergence after only a finite number of steps.

Perspectives

The work achieved in this thesis opens several interesting directions for future research, both from a theoretical and an applied perspective. Let us highlight a few representative ones.

Theoretical open problems

- Our promising observations on the robustness of Prony's approach when applied to measure supported on curves and surfaces paves the way for a further theoretical investigation of this regime. Analyzing the convergence of the discrete measure computed by Prony's method towards the continuous one, or any other relations between the two appear to be fertile questions, which could lead to new algorithms with stronger recovery guarantees and broader applicability.
- Many theoretical questions have still to be solved in regards of the identifiability and stability with respect to noise of variational models for super-resolution in dimension higher than one. In the multivariate case, such results depend on many additional factors, such as the geometrical configuration of the points to super-resolve. On the other hand, moment hierarchies provide a systematic approach to solve these problems, independently of their success. Investigating the interplay between scenarios where super-resolution is indeed achievable via convex methods, and cases of finite convergence of the hierarchy, could be a good starting point towards a better understanding of these matters.

Applicative open problems

- In our work, we have shown that Fourier-based methods can still be used to solve problems which are non-translation-invariant via a projection scheme. Using other type of moments might also be useful, and it is not yet clear what is the best option for challenging 3-D imaging problems with spatially-varying operators.
- On another note, our work also brings new promising results concerning the use of moment approaches in machine learning applications. Moment-sos approaches have already proven to be useful for data science, in particular via Christoffel polynomials. The design of new moment-based or

sos-based algorithms to perform tasks more challenging than the recovery of sparse measures is an exciting topic, offering many challenges yet to be solved.

Bibliography

- [Abergel and Moisan, 2017] Abergel, R. and Moisan, L. (2017). The Shannon total variation. *J. Math. Imaging Vis.*, 59:341–370.
- [Absil and Gallivan, 2006] Absil, P.-A. and Gallivan, K. (2006). Joint diagonalization on the oblique manifold for independent component analysis. In *ICASSP Proceedings*, pages 945–948.
- [Ahmadi et al., 2017] Ahmadi, A., Hall, G., Papachristodoulou, A., Saunderson, J., and Zheng, Y. (2017). Improving efficiency and scalability of sums of squares optimization: recent advances and limitations. In *Proceedings of the Conference on Decision and Control*.
- [Akhiezer and Krein, 1938] Akhiezer, N. and Krein, M. (1938). *Some Questions in the Theory of Moments*, volume 2 of *Transl. Math. Monographs*. AMS.
- [Akuzawa and Murata, 2001] Akuzawa, T. and Murata, N. (2001). Multiplicative nonholonomic Newton-like algorithm. *Chaos, Solitons and Fractals*, 12(4):785–793.
- [Alard, 2000] Alard, C. (2000). Image subtraction using a space-varying kernel. *Astronomy and Astrophysics Supplement*, 144:363–370.
- [Alfonsi et al., 2019] Alfonsi, A., Coyaud, R., Ehrlicher, V., and Lombardi, D. (2019). Approximation of optimal transport problems with marginal moments constraints.
- [Amari, 1998] Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Comput.*, 10:251–276.
- [Amari et al., 2000] Amari, S.-I., Chen, T., and Cichocki, A. (2000). Nonholonomic orthogonal learning algorithms for blind source separation. *Neural Comput.*, 12(6):1563–1484.
- [Ambrosio et al., 2000] Ambrosio, L., Fusco, N., and Pallara, D. (2000). *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford Mathematical Monographs.

- [Andersson and Carlsson, 2017] Andersson, F. and Carlsson, M. (2017). On the structure of positive semi-definite finite rank general domain Hankel and Toeplitz operators in several variables. *Complex Anal. Oper. Theory*, 11:755–784.
- [Arjovsky et al., 2017] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial network. *Proc ICML’17*, 60(4):941–965.
- [Auzinger and Stetter, 1988] Auzinger, W. and Stetter, H. (1988). An elimination algorithm for the computation of all zeros of a system of multivariate polynomial equations. In *International Series in Numerical Mathematics*, volume 86, pages 11–30. Birkhäuser.
- [Azaïs et al., 2015] Azaïs, J. M., de Castro, Y., and Gamboa, F. (2015). Spike detection from inaccurate sampling. *Appl. Comput. Harmon. Anal.*, 38(2):177–195.
- [Bandeira et al., 2016] Bandeira, A., Boumal, N., and Voroninski, V. (2016). On the low-rank approach for semidefinite programs arising in synchronization and community detection. In *COLT*.
- [Batenkov et al., 2020] Batenkov, D., Demanet, L., Goldman, G., and Yomdin, Y. (2020). Conditioning of partial nonuniform Fourier matrices with clustered nodes. *SIAM J. Matrix Anal. Appl.*, 41(1):199–220.
- [Beck and Teboulle, 2009] Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202.
- [Beck and Teboulle, 2010] Beck, A. and Teboulle, M. (2010). *Gradient-Based Algorithms with Applications to Signal-Recovery Problems*, pages 42–88. Cambridge Univ. Press.
- [Bellavia et al., 2020] Bellavia, S., Gondzio, J., and Porcelli, M. (2020). A relaxed interior point method for low-rank semidefinite programming problems with applications to matrix completion.
- [Benamou, 2003] Benamou, J. (2003). Numerical resolution of an “unbalanced” mass transport problem. *ESAIM Math. Model. Numer. Anal.*, 37(5):851–868.
- [Benamou and Brenier, 2000] Benamou, J.-D. and Brenier, Y. (2000). A computational fluid mechanics solution of the Monge-Kantorovitch mass transfer problem. *Numer. Math.*, 84:375–393.
- [Benamou et al., 2016] Benamou, J.-D., Collino, F., and Mirebeau, J.-M. (2016). Monotone and consistent discretization of the Monge-Ampère operator. *Mathematics of Computation*, 85(302):2743–2775.
- [Benamou et al., 2015] Benamou, J.-D., Froese, B., and Oberman, A. (2015). Numerical solution of the optimal transportation problem using the Monge-Ampère equation. *Journal of Computational Physics*, 260:107–126.

- [Betzig et al., 2006] Betzig, E., Patterson, G., Sougrat, R., Lindwasser, O., Olenych, S., Bonifacino, J., Davidson, M., Lippincott-Schwartz, J., and Hess, H. (2006). Imaging intracellular fluorescent proteins at nanometer resolution. *Science*.
- [Bhaskar and Recht, 2011] Bhaskar, B. and Recht, B. (2011). Atomic norm denoising with applications to line spectral estimation. In *49th Annual Allerton Conference on Communication, Control, and Computing*, pages 261–268.
- [Blu et al., 2008] Blu, T., Dragotti, P., Vetterli, M., Marziliano, P., and Coulot, L. (2008). Sparse sampling of signal innovations. *IEEE Signal Proc. Mag.*, 25:31–40.
- [Blumensath and Davies, 2008] Blumensath, T. and Davies, M. (2008). Iterative thresholding for sparse approximations. *J. Fourier Anal. Appl.*, 14(5-6):629–654.
- [Bonneel et al., 2016] Bonneel, N., Peyré, G., and Cuturi, M. (2016). Wasserstein barycentric coordinates: Histogram regression using optimal transport. In *SIGGRAPH*.
- [Boumal, 2015] Boumal, N. (2015). A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. Technical report.
- [Boumal et al., 2019] Boumal, N., Voroninski, V., and Bandeira, A. (2019). Deterministic guarantees for Burer-Monteiro factorizations of smooth semidefinite programs. *Comm. Pure App. Math.*, 73(3):581–608.
- [Boumal et al., 2016] Boumal, N., Voroninski, V., and Bandeira, A. S. (2016). The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *NIPS*, pages 2765–2773.
- [Boyd et al., 2015] Boyd, N., Schiebinger, G., and Recht, B. (2015). The alternating descent conditional gradient method for sparse inverse problems. In *CAMSAP*, pages 57–60.
- [Boyd et al., 2011] Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.
- [Boyer et al., 2019] Boyer, C., Chambolle, A., de Castro, Y., Duval, V., de Gournay, F., and Weiss, P. (2019). On representer theorems and convex regularization. *SIAM J. Optim.*, 29(2):1260–1281.
- [Bredies et al., 2009] Bredies, K., Lorenz, D. A., and Maass, P. (2009). A generalized conditional gradient method and its connection to an iterative shrinkage method. *Computational Optimization and Applications*, 42(2):173–193.

- [Bredies and Pikkarainen, 2013] Bredies, K. and Pikkarainen, H. K. (2013). Inverse problems in spaces of measures. *ESAIM: Control, Optimization and Calculus of Variations*, 19(1):190–218.
- [Burer and Monteiro, 2003] Burer, S. and Monteiro, R. (2003). A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.*, 95(2):329–357.
- [Burer and Monteiro, 2005] Burer, S. and Monteiro, R. (2005). Local minima and convergence in low-rank semidefinite programming. *Math. Program.*, 103(3):427–444.
- [Canas and Rosasco, 2012] Canas, G. and Rosasco, L. (2012). Learning probability measures with respect to optimal transport metrics. In *Advances in Neural Information Processing Systems*.
- [Candès and Fernandez-Granda, 2013] Candès, E. and Fernandez-Granda, C. (2013). Super-resolution from noisy data. *J. Fourier Anal. Appl.*, 19:1229–1254.
- [Candès and Fernandez-Granda, 2014] Candès, E. and Fernandez-Granda, C. (2014). Towards a mathematical theory of super-resolution. *Comm. Pure Appl. Math.*, 67(6):906–956.
- [Candès et al., 2006] Candès, E., Romberg, J., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate data. *Comm. Pure Appl. Math.*, 59:1207–1223.
- [Candès and Plan, 2010] Candès, E. J. and Plan, Y. (2010). A probabilistic and ripples theory of compressed sensing. *IEEE Trans. Inform. Theory*, 57(11):7235–7524.
- [Carathéodory, 1907] Carathéodory, C. (1907). über den Variabilitätsbereich der Koeffizienten von Potenzreihen, die gegebene Werte nicht annehmen. *Math. Ann.*, 64:95–115.
- [Carathéodory, 1911] Carathéodory, C. (1911). über den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen. *Rend. Circ. Mat. Palermo*, 32:193–217.
- [Cardoso and Souloumiac, 1993] Cardoso, J. and Souloumiac, A. (1993). Blind beamforming for non Gaussian signals. *IEEE Proc.-F*, 140(6):362–370.
- [Cardoso and Souloumiac, 1996] Cardoso, J. and Souloumiac, A. (1996). Jacobi angles for simultaneous diagonalization. *SIAM J. Mat. Anal. Appl.*, 17(1).
- [Cardoso and Laheld, 1996] Cardoso, J.-F. and Laheld, B. (1996). Equivariant adaptive source separation. *IEEE Trans. Signal Process.*, 44(12).
- [Chambolle, 2004] Chambolle, A. (2004). An algorithm for total variation minimization and applications. *J. Math. Imaging Vis.*, 20(1-2):89–97.

- [Chambolle et al., 2016] Chambolle, A., Duval, V., Peyré, G., and Poon, C. (2016). Geometric properties of solutions to the total variation denoising problem. *Inverse Problems*, 33(1):015002.
- [Chambolle and Pock, 2011] Chambolle, A. and Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.*, 40(1):120–145.
- [Chandrasekaran et al., 2012] Chandrasekaran, V., Recht, B., Parrilo, P., and Willsky, A. (2012). The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849.
- [Chang et al., 2000] Chang, E., Mallat, S., and Yap, C. (2000). Wavelet foveation. *J. Applied and Computational Harmonic Analysis*, 9(3):312–335.
- [Chen et al., 2012] Chen, C., Marziliano, P., and Kot, A. (2012). 2d finite rate of innovation reconstruction method for step edge and polygon signals in the presence of noise. *IEEE Trans. Signal Process.*, 60:2851–2859.
- [Chen et al., 2001] Chen, S. S., Donoho, D. L., and Saunders, M. A. (2001). Atomic decomposition by basis pursuit.
- [Chi et al., 2011] Chi, Y., Scharf, L., Pzeshki, A., and Calderbank, A. (2011). Sensitivity to basis mismatch in compressed sensing. *IEEE Trans. Signal Process.*, 59(5):2182–2195.
- [Chizat, 2019] Chizat, L. (2019). Sparse optimization on measures with over-parameterized gradient descent.
- [Chizat and Bach, 2018] Chizat, L. and Bach, F. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. In *NeurIPS*.
- [Chizat et al., 2018] Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2018). Unbalanced optimal transport: Dynamic and Kantorovitch formulations. *J. Fun. Anal.*, 274:3090–3123.
- [Chui and Rangarajan, 2003] Chui, H. and Rangarajan, A. (2003). A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89:114–141.
- [Claerbout and Muir, 1973] Claerbout, J. and Muir, F. (1973). Robust modeling with erratic data. *Geophysics*, 38(5):826–844.
- [Combettes and Pesquet, 2011] Combettes, P. and Pesquet, J.-C. (2011). Proximal splitting methods in signal processing. In *Springer Optim. Appl.*, volume 49, pages 185–212. Springer, New York.
- [Condat, 2017] Condat, L. (2017). Discrete total variation: New definition and minimization. *SIAM J. Imaging Sc.*, 10(3):1258–1290.

- [Condat, 2020] Condat, L. (2020). Atomic norm minimization for decomposition into complex exponentials and optimal transport in fourier domain. *J. Approx. Th.*, 258.
- [Corless, 1996] Corless, R. (1996). Editor’s corner: Groöbner bases and matrix eigenproblems. *ACM Sigsam Bulletin: Comm. Comput. Algebra*, 30.
- [Corless et al., 1997] Corless, R., Gianni, P., and Trager, B. (1997). A reordered scur factorization method for zero-dimensional polynomial systems with multiple roots. In *Proceedings of 1997 ISSAC*, pages 133–140. ACM.
- [Cropper et al., 2013] Cropper, M., Hoekstra, H., Kitching, T., Massey, R., Amiaux, J., Miller, L., Mellier, Y., Rhodes, J., Rowe, B., Pires, S., Saxton, C., and Scaramella, R. (2013). Defining a weak lensing experiment in space. *Monthly Notices on the Royal Astronomical Society*, 431(4):3103–3126.
- [Curto and Fialkow, 1991] Curto, R. and Fialkow, L. (1991). Recursiveness, positivity and truncated moment problems. *Houston J. Math.*, 17(4):603–635.
- [Curto and Fialkow, 1996] Curto, R. and Fialkow, L. (1996). Solution of the truncated complex moment problem for flat data. *Memoirs of the AMS*, (568).
- [Curto and Fialkow, 1998] Curto, R. and Fialkow, L. (1998). Flat extensions of positive moment matrices: Recursively generated relations. *Memoirs of the AMS*, 648.
- [Curto and Fialkow, 2002] Curto, R. and Fialkow, L. (2002). Solution of the singular quartic moment problem. *J. Operator Theory*, 42:315–354.
- [Cuturi, 2013] Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*.
- [Cuturi and Avis, 2014] Cuturi, M. and Avis, D. (2014). Ground metric learning. *J. Machine Learning Research*, 15:533–564.
- [Cuyt and Lee, 2018] Cuyt, A. and Lee, W. (2018). Multivariate exponential analysis from the minimal number of samples. *Adv. Comput. Math.*, 44:987–1002.
- [Daubechies et al., 2004] Daubechies, I., Defrise, M., and De Mol, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57(11):1413–1457.
- [de Castro and Gamboa, 2012] de Castro, Y. and Gamboa, F. (2012). Exact reconstruction using Beurling minimal extrapolation. *J. Math. Anal. Appl.*, 395(1):336–354.

- [de Castro et al., 2017] de Castro, Y., Gamboa, F., Henrion, D., and Lasserre, J. (2017). Exact solutions to super resolution on semi-algebraic domains in higher dimensions. *IEEE Trans. Inform. Theory*, 63(1):621–630.
- [De Goes et al., 2012] De Goes, F., Breeden, K., Ostromoukhov, V., and Desbrun, M. (2012). Blue noise through optimal transport. *ACM Transactions on Graphics*, 31(6):171.
- [Delon et al., 2009] Delon, J., Rabin, J., and Gousseau, Y. (2009). Transportation distances on the circle and applications.
- [Demagnet and Nguyen, 2015] Demagnet, L. and Nguyen, N. (2015). The recoverability limit for superresolution via sparsity.
- [Demyanov and Rubinov, 1967] Demyanov, V. F. and Rubinov, A. M. (1967). The minimization of a smooth convex functional on a convex set. *SIAM Journal on Control*, 5(2).
- [Demyanov and Rubinov, 1970] Demyanov, V. F. and Rubinov, A. M. (1970). *Approximate methods in optimization problems*. American Elsevier Pub. Co New York.
- [Denoyelle et al., 2017] Denoyelle, Q., Duval, V., and Peyré, G. (2017). Support recovery for sparse super-resolution of positive measures. *J. Fourier Anal. Appl.*, 23:1153–1194.
- [Denoyelle et al., 2020] Denoyelle, Q., Duval, V., Peyré, G., and Soubies, E. (2020). The sliding frank-wolfe algorithm and its application to super-resolution microscopy. *Inverse Problems*, 36(1):014001.
- [Diederichs and Iske, 2015] Diederichs, B. and Iske, A. (2015). Parameter estimation for bivariate exponential sums. In *International Conference on Sampling Theory and Applications*, pages 493–497. IEEE.
- [Donoho, 1992] Donoho, D. L. (1992). Superresolution via sparsity constraints. *SIAM J. Math. Anal.*, 23(5):1309–1331.
- [Dragotti et al., 2005] Dragotti, P., Vetterli, M., and Blu, T. (2005). Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets strang-fix. *IEEE Trans. Signal Process.*, 55:1741–1757.
- [Dumitrescu, 2017] Dumitrescu, B. A. (2017). *Positive trigonometric Polynomials and Signal Processing Applications*. Signals and Communication Technology. Springer International Publishing.
- [Dunkl and Xu, 2001] Dunkl, C. and Xu, Y. (2001). Cambridge Univ. Press.
- [Duval and Peyré, 2015] Duval, V. and Peyré, G. (2015). Exact support recovery for sparse spikes deconvolution. *Foundations of Computational Mathematics*, 15(5):1315–1355.

- [Duval and Peyré, 2017] Duval, V. and Peyré, G. (2017). Sparse regularization on thin grids I: the Lasso. *Inverse Problems*, 33(5).
- [Ehler et al., 2019] Ehler, M., Kunis, S., Peter, T., and Richter, C. (2019). A randomized multivariate matrix pencil method for superresolution microscopy. *Elec. Trans. Numer. Anal.*, 51:63–74.
- [Elkadi and Mourrain, 2007] Elkadi, M. and Mourrain, B. (2007). *Introduction à la Résolution des Systèmes Polynomiaux*. Mathématiques et Applications. Springer.
- [Fannjiang, 2016] Fannjiang, A. (2016). Compressive spectral estimation with single-snapshot ESPRIT: Stability and resolution.
- [Faugère et al., 1993] Faugère, J., Gianni, P., Lazard, D., and Mora, T. (1993). Efficient computation of zero-dimensional Gröbner bases by change of ordering. *J. Symb. Comput.*, 16:329–344.
- [Federer, 1969] Federer, H. (1969). *Geometric Measure Theory*. Classics in Mathematics. Springer Berlin Heidelberg.
- [Fejér, 1916] Fejér, L. (1916). über trionometrische polynome. *J. Reine Angew. Math.*, 146:53–82.
- [Fernandez-Granda, 2015] Fernandez-Granda, C. (2015). Super-resolution of point sources via convex programming. *2015 IEEE 6th International Workshop CAMSAP*, pages 41–44.
- [Ferreira Da Costa and Dai, 2018] Ferreira Da Costa, M. and Dai, W. (2018). A tight converse to the spectral resolution limit via convex programming. In *ISIT*. IEEE.
- [Feydy et al., 2017] Feydy, J., Charlier, B., Vialard, F.-X., and Peyré, G. (2017). Optimal transport for diffeomorphic registration. In *MICCAI 2017*, pages 291–299.
- [Feydy et al., 2019] Feydy, J., Roussillon, P., Trounev, A., and Gori, P. (2019). Fast and scalable optimal transport for brain tractograms. In *MICCAI*.
- [Figalli and Villani, 2011] Figalli, A. and Villani, C. (2011). *Optimal Transport and Curvature*, volume 2028 of *Lecture Notes in Mathematics*, pages 171–217. Springer.
- [Flinth et al., 2019] Flinth, A., de Gournay, F., and Weiss, P. (2019). On the linear convergence rates of exchange and continuous methods for total variation minimization.
- [Flinth and Weiss, 2017] Flinth, A. and Weiss, P. (2017). Exact solutions of infinite dimensional total-variation regularized problems.

- [Flury and Gautschi, 1986] Flury, B. and Gautschi, W. (1986). An algorithm for simultaneous orthogonal transformation of several positive definite symmetric matrices to nearly diagonal form. *SIAM J. Sci. Stat. Comput.*, 7(1):169–184.
- [Frank and Wolfe, 1956] Frank, M. and Wolfe, P. (1956). An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110.
- [Friedman et al., 2010] Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.*, 33(1).
- [Frogner et al., 2015] Frogner, C., Zhang, C., Mobahi, H., Araya, M., and Poggio, T. A. (2015). Learning with a Wasserstein loss. In *Advances in Neural Information Processing Systems*.
- [Fu, 1998] Fu, W. (1998). Penalized regressions: the bridge versus the lasso. *J. Comput. Graph. Statist.*, 7(3):397–416.
- [Gatermann and Parrilo, 2009] Gatermann, K. and Parrilo, P. (2009). Symmetry groups, semidefinite programs, and sums of squares. *J. Pure Appl. Alg.*, 192(1–3):95–128.
- [Genevay et al., 2016] Genevay, A., Cuturi, M., Peyré, G., and Bach, F. (2016). Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, pages 3432–3440.
- [Genevay et al., 2018] Genevay, A., Peyré, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617.
- [Gentile et al., 2013] Gentile, M., Courbin, F., and Meylan, G. (2013). Interpolating point spread function anisotropy. *Astronomy and Astrophysics*, 549:A1.
- [Gidel et al., 2018] Gidel, G., Pedregosa, F., and Lacoste-Julien, S. (2018). Frank-Wolfe splitting via augmented Lagrangian method.
- [Goemans and Williamson, 1995] Goemans, M. and Williamson, D. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145.
- [Golub et al., 1999] Golub, G., Milanfar, P., and Varah, J. (1999). A stable numerical method for inverting shape from moments. *SIAM J. Sci. Comput.*, 21(4):1222–1243.
- [Gribonval et al., 2017] Gribonval, R., Blanchard, G., Keriven, N., and Traonmilin, Y. (2017). Compressive statistical learning with random feature moments. *arXiv:1706.07180*.
- [Group, 2013] Group, B. I. (2013). Benchmarking of single-molecule localization microscopy software.

- [Guélat and Marcotte, 1986] Guélat, J. and Marcotte, P. (1986). Some comments on Wolfe’s “away step”. *Math. Program.*, 35(1).
- [Gustafsson et al., 2009] Gustafsson, B., Putinar, M., Saf, E., and Stylianopoulos, N. (2009). Bergman polynomials on an archipelago: Estimates, zeros and shape reconstruction. *Adv. Math.*, 222(4):1405–1460.
- [Haardt and Nossek, 1998] Haardt, M. and Nossek, J. (1998). Simultaneous schur decomposition of several nonsymmetric matrices to achieve automatic pairing in multidimensional harmonic retrieval problems. *IEEE Trans. Signal Process.*, 46(1).
- [Harchaoui et al., 2015] Harchaoui, Z., Juditsky, A., and Nemirovski, A. (2015). Conditional gradient algorithms for norm-regularized smooth convex optimization. *Math. Program.*, 152(1-2):75–112.
- [Harmouch et al., 2017] Harmouch, J., Khalil, H., and Mourrain, B. (2017). Structured low rank decomposition of multivariate Hankel matrices. *Linear Algebra and its Applications*, 542:162 – 185.
- [Haviland, 1935] Haviland, E. (1935). On the momentum problem for distributions in more than one dimension, i. *Amer. J. Math.*, 57:562–568.
- [Haviland, 1936] Haviland, E. (1936). On the momentum problem for distributions in more than one dimension, ii. *Amer. J. Math.*, 58:164–168.
- [Henrion and Lasserre, 2005] Henrion, D. and Lasserre, J. (2005). *Detecting Global Optimality and Extracting Solutions in Gloptipoly*, pages 293–310. Springer.
- [Hess et al., 2006] Hess, S., Girirajan, T., and Mason, M. (2006). Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophysical Journal*, 91(11):4258–4272.
- [Hua, 1992] Hua, Y. (1992). Estimating two-dimensional frequencies by matrix enhancement and matrix pencil. *IEEE Trans. Signal Process.*, 40(9):2267–2280.
- [Hua and Sarkar, 1989] Hua, Y. and Sarkar, T. (1989). Generalized pencil-of-function method for extracting poles of an em system from its transient response. *IEEE Trans. Antennas Propagation*, 37(2).
- [Jaccard, 1901] Jaccard, P. (1901). Etude de la distribution florale dans une portion des Alpes et du Jura. *Bulletin de la Société Vandoise des Sciences Naturelles*, 37:547–579.
- [Jacobi and Prestel, 2001] Jacobi, T. and Prestel, A. (2001). Distinguished representations of strictly positive polynomials. *J. Reine Angew. Math.*, 532:223–235.

- [Jaggi, 2013] Jaggi, M. (2013). Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *ICML*, volume 28.
- [Jaggi and Sulovský, 2010] Jaggi, M. and Sulovský, M. (2010). A simple algorithm for nuclear norm regularized problems. In *ICML*, pages 471–478.
- [Janati et al., 2019a] Janati, H., Bazeille, T., Thirion, B., Cuturi, M., and Gramfort, A. (2019a). Group level MEG/EEG source imaging via optimal transport: Minimum Wasserstein estimates. In *IPMI*, volume Lecture Notes in Computer Science.
- [Janati et al., 2019b] Janati, H., Cuturi, M., and Gramfort, A. (2019b). Wasserstein regularization for sparse multi-task regression. In *AISTATS*, volume 89.
- [Jordan et al., 1998] Jordan, R., Kinderlehrer, D., and Otto, F. (1998). The variational formulation of the Fokker-Planck equation. *SIAM J. Math. Anal.*, 29(1):1–17.
- [Journée et al., 2010] Journée, M., Bach, F., Absil, P.-A., and Sepulchre, R. (2010). Low-rank optimization on the cone of positive semidefinite matrices. *SIAM J. Optim.*, 20(5):2327–2351.
- [Kandil, 1983] Kandil, A. (1983). Parseval’s identity on Banach space. *Kyungpook Mathematical J.*, 23:99–104.
- [Kantorovich, 1942] Kantorovich, L. (1942). On the transfer of masses (in russian). *Doklady Akademii Nauk*, 37(2):227–229.
- [Klep et al., 2018] Klep, I., Povh, J., and Volčič, J. (2018). Minimizer extraction in polynomial optimization is robust. *SIAM J. Optim.*, 28(4):3177–3207.
- [Krim and Viberg, 1996] Krim, H. and Viberg, M. (1996). Two decades of array signal processing research: the parametric approach. *IEEE Signal Processing Magazine*, 13(4):67–94.
- [Kroó and Lubinsky, 2012] Kroó, A. and Lubinsky, D. (2012). Christoffel functions and universality in the bulk for multivariate orthogonal polynomials. *Can. J. Math.*, 65(3):600–620.
- [Kunis et al., 2018] Kunis, S., H.M., M., Peter, T., and U., V. D. O. (2018). Prony’s method under an almost sharp multivariate ingham inequality. *J. Fourier Anal. Appl.*, 24:1306–1318.
- [Kunis and Nagel, 2018] Kunis, S. and Nagel, D. (2018). On the condition number of Vandermonde matrices with pairs of nearly-colliding nodes.
- [Kunis et al., 2016] Kunis, S., Peter, T., Römer, T., and Von der Ohe, U. (2016). A multivariate generalization of Prony’s method. *Linear Algebra Appl.*, 490:31–47.

- [Lasserre and Pauwels, 2016] Lasserre, J. and Pauwels, E. (2016). Sorting out typicality with the inverse moment matrix SOS polynomial. In *NIPS*.
- [Lasserre and Pauwels, 2019] Lasserre, J. and Pauwels, E. (2019). The empirical Christoffel function with applications in data analysis. *Advances in Computational Mathematics*, 45(3):1439–1468.
- [Lasserre, 2001] Lasserre, J. B. (2001). Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817.
- [Lasserre, 2008] Lasserre, J.-B. (2008). A semidefinite approach to the generalized problem of moments. *Math. Program.*, 112:65–92.
- [Lasserre, 2010] Lasserre, J. B. (2010). *Moments, positive polynomials and their applications*. Imperial College Press Optimization Series. Imperial College Press, London.
- [Laurent et al., 2009] Laurent, J., Obozinski, G., and Vert, J.-P. (2009). Group-lasso with overlap and graph lasso. *ICML*, pages 43–440.
- [Laurent, 2005] Laurent, M. (2005). Revisiting two theorems of curto and alikow on moment matrices. *Proceedings of the AMS*, 133(10):2965–2976.
- [Laurent, 2010] Laurent, M. (2010). Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, volume 149. Springer new York.
- [Levitin and Polyak, 1966] Levitin, E. and Polyak, B. (1966). Constrained minimization methods. *USSR Computational Mathematics and Mathematical Physics*, 6(5):1–50.
- [Lévy, 2015] Lévy, B. (2015). A numerical algorithm for ℓ^2 semi-discrete optimal transport in 3d. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1693–1715.
- [Levy and Fullagar, 1981] Levy, S. and Fullagar, P. (1981). Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution. *Geophysics*, 46(9):1235–1243.
- [Li and Liao, 2019] Li, W. and Liao, W. (2019). Conditioning of restricted fourier matrices and super-resolution of music. In *SampTA*.
- [Li et al., 2019] Li, W., Liao, W., and Fannjiang, A. (2019). Super-resolution limit of the ESPRIT algorithm.
- [Li and Chi, 2016] Li, Y. and Chi, Y. (2016). Off-the-grid line spectrum denoising and estimation with multiple measurement vectors. *IEEE Trans. Signal Process.*, 64(5):1257–1269.
- [Liao and Fannjiang, 2014] Liao, W. and Fannjiang, A. (2014). MUSIC for single-snapshot spectral estimation: Stability and super-resolution. *CoRR*.

- [Liero et al., 2017] Liero, M., Mielke, A., and Savaré, G. (2017). Optimal entropy-transport problems and a new Hellinger-Kantorovitch distance between positive measures. *Invent. Math.*, 211:969–1117.
- [Liu et al., 2015] Liu, Y.-F., Wang, X., Liu, X., and Ma, S. (2015). A scalable Frank-Wolfe based augmented Lagrangian method for linearly constrained composite convex programming.
- [Locatello et al., 2017] Locatello, F., Khanna, R., Tschannen, M., and Jaggi, M. (2017). A unified optimization view on generalized matching pursuit and Frank-Wolfe. In *AISTATS*, pages 860–868.
- [Majumdar et al., 2020] Majumdar, A., Hall, G., and Ahmadi, A. (2020). Recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:331–360.
- [Maravic and Vetterli, 2005] Maravic, I. and Vetterli, M. (2005). Sampling and reconstruction of signals with finite rate of innovation in the presence of noise. *IEEE Trans. Signal Process.*, 53:2788–2805.
- [Marx et al., 2019] Marx, S., Pauwels, E., Weisser, T., Henrion, D., and Lasserre, J. (2019). Tractable semi-algebraic approximation using Christoffel-Darboux kernel.
- [Matsuoka et al., 1995] Matsuoka, K., Ohya, M., and Kawamoto, M. (1995). A neural net for blind separation of nonstationary signals. *Neural Networks*, 8:411–419.
- [Mérigot, 2011] Mérigot, Q. (2011). A multiscale approach to optimal transport. In *Computer Graphics Forum*, volume 30, pages 1583–1592. Wiley Online Library.
- [Métivier et al., 2016] Métivier, L., Brossier, R., Mérigot, Q., Oudet, E., and Virieux, J. (2016). An optimal transport approach for seismic tomography: Application to 3d full waveform inversion. *Inverse Problems*, 32(11).
- [Milanfar et al., 1995] Milanfar, P., Verghese, G., Karl, W., and Willsky, A. (1995). Reconstructing polygons from moments with connections to array processing. *IEEE Trans. Signal Proc.*, 43:432–443.
- [Moerner and Fromm, 2003] Moerner, W. and Fromm, D. (2003). Methods of single-molecule fluorescence spectroscopy and microscopy. *Rev. Sci. Instr.*, 74(8):3597–3619.
- [Moitra, 2015] Moitra, A. (2015). Super-resolution, extremal functions and the condition number of Vandermonde matrices. In *STOC*.
- [Möller, 1993] Möller, H. (1993). Systems of algebraic equations solved by means of endomorphisms. In *”App. Algebra, Algebraic Algo. and Error-Correcting Codes*, pages 43–56. Berlin.

- [Möller and Stetter, 1995] Möller, H. and Stetter, H. (1995). Multivariate polynomial equations with multiple zeros solved by matrix eigenproblems. *Numer. Math.*, 70:311–329.
- [Monge, 1781] Monge, G. (1781). Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris*, pages 666–704.
- [Morgenshtern, 2020] Morgenshtern, V. (2020). Super-resolution of positive sources on an arbitrarily fine grid.
- [Mosher et al., 1993] Mosher, J.-C., Leahy, R., and Lewis, P. (1993). Biomagnetic localization from transient quasi-static events. In *IEEE Int. Conf. Acous. Speech Signal Process.*, volume 1, pages 91–94.
- [Ngolé and Starck, 2017] Ngolé, F. and Starck, J.-L. (2017). Point spread function field learning based on optimal transport distances. *SIAM Journal on Imaging Sciences*, 10(3):1549–1578.
- [Ongie and Jacob, 2016] Ongie, G. and Jacob, M. (2016). Off-the-grid recovery of piecewise constant images from few Fourier samples. *SIAM J. Imaging Sci.*, 9(3):1004–1041.
- [Owen et al., 2009] Owen, J., Attias, H., Sekihara, K., Nagarajan, S., and Wipf, D. (2009). Estimating the location and orientation of complex, correlated neural activity using MEG. *NIPS*, pages 1777–1784.
- [Pan et al., 2014] Pan, H., Blu, T., and Dragotti, P. (2014). Sampling curves with finite rate of innovation. *IEEE Trans. Signal Process.*, 62(2):458–471.
- [Pauwels et al., 2020] Pauwels, E., Putinar, M., and Lasserre, J. (2020). Data analysis from empirical moments and the Christoffel function. *F. Comp. Math.*
- [Peter et al., 2015] Peter, T., Plonka, G., and Schaback, R. (2015). Prony’s method for multivariate signals. *Proc. Appl. Math. Mech.*, 15:665–666.
- [Peyré and Cuturi, 2019] Peyré, G. and Cuturi, M. (2019). Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- [Pham, 2001] Pham, D.-T. (2001). Joint approximate diagonalization of positive definite Hermitian matrices. *SIAM J. Matrix Anal. Appl.*, 22(4):1136–1152.
- [Pham and Congedo, 2009] Pham, D.-T. and Congedo, M. (2009). Least square joint diagonalization of matrices under an intrinsic scale constraint. In *Independent Component Analysis and Signal Separation*, pages 298–305. Springer.
- [Pisarenko, 1973] Pisarenko, V. (1973). The retrieval of harmonics from a covariance function. *Geophysical J. Royal Astro. Soc.*, 33:347–366.

- [Polisano et al., 2017] Polisano, K., Condat, L., Clausel, M., and Perrier, V. (2017). A convex approach to super-resolution and regularization of lines in images.
- [Poon et al., 2020] Poon, C., Keriven, N., and Peyré, G. (2020). The geometry of off-the-grid compressed sensing.
- [Poon and Peyré, 2018] Poon, C. and Peyré, G. (2018). Multi-dimensional sparse super-resolution. *SIAM J. Math. Anal.*, 51(1):1–44.
- [Potts and Tasche, 2011] Potts, D. and Tasche, M. (2011). Nonlinear approximation by sums of nonincreasing exponentials. *Applicable Analysis*, 90(3–4):609–626.
- [Potts and Tasche, 2013] Potts, D. and Tasche, M. (2013). Parameter estimation for multivariate exponential sums. *Electron. Trans. Numer. Anal.*, 40:204–224.
- [Potts and Tasche, 2017] Potts, D. and Tasche, M. (2017). *Error Estimates for the ESPRIT Algorithm*, pages 621–648. Oper. Theory Adv. Appl. Springer, Cham.
- [Putinar, 1993] Putinar, M. (1993). Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984.
- [R. de Prony, 1795] R. de Prony, G. (1795). Essai expérimental et analytique: Sur les lois de la dilatabilité des fluides élastiques et sur celles de la force expansive de la vapeur de l’eau et de la vapeur de l’alkool, à différentes températures. *Journal de l’École Polytechnique Floréal et Plairial*, 1(cahier 22):24–76.
- [Riesz, 1916] Riesz, F. (1916). über ein Problem des Herrn Carathéodory. *J. Reine Angew. Math.*, 146:83–87.
- [Rockafellar, 1970] Rockafellar, R. (1970). *Convex Analysis*. Princeton Math. Series. Princeton Univ. Press.
- [Rouquette and Najim, 2001] Rouquette, S. and Najim, M. (2001). Estimation of frequencies and damping factors by two-dimensional esprit type methods. *IEE Trans. Signal Process.*, 49(1).
- [Roy and Kailath, 1989] Roy, R. and Kailath, T. (1989). ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Trans. Acoustics Speech Signal Process.*, 37(7):984–995.
- [Roy et al., 1986] Roy, R., Paulraj, A., and Kailath, T. (1986). Esprit – a subspace rotation approach to estimation of parameters of cisoids in noise. *IEEE Trans. Acoustics Speech Signal Process.*, 34(5):1340–1342.

- [Rust et al., 2006] Rust, M., Bates, M., and Zhuang, X. (2006). Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM). *Nature Methods*, 3:793–796.
- [Sacchini et al., 1993] Sacchini, J., Steedly, W., and Moses, R. (1993). Two-dimensional Prony modeling and parameter estimation. *IEEE Trans. Signal Process.*, 41(11).
- [Santambrogio, 2015] Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94.
- [Santosa and Symes, 1986] Santosa, F. and Symes, W. (1986). Linear inversion of band-limited reflection seismograms. *SIAM J. Sci. Stat. Comput.*, 7(4):1307–1330.
- [Sauer, 2017] Sauer, T. (2017). Prony’s method in several variables. *Numer. Math.*, 136:411–438.
- [Sauer, 2018] Sauer, T. (2018). Prony’s method in several variables: Symbolic solutions by universal interpolation. *J. Symbolic Comput.*, 84:95–112.
- [Scheiderer, 2006] Scheiderer, C. (2006). Sums of squares on real algebraic surfaces. *Manuscripta Math.*, 119:395–410.
- [Schiebinger et al., 2019] Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Liu, S., Lin, S., Berube, P., Lee, L., et al. (2019). Reconstruction of developmental landscapes by optimal-transport analysis of single-cell gene expression sheds light on cellular reprogramming. *Cell*, (176):928–943.
- [Schmidt, 2005] Schmidt, M. (2005). minFunc: Unconstrained differentiable multivariate optimization in Matlab. <http://www.cs.ubc.ca/~schmidt/Software/minFunc.html>.
- [Schmidt, 1986] Schmidt, R. (1986). Multiple emitter location and signal parameter estimation. *IEEE Trans. Antennas Propagation*, 34(3):276–280.
- [Schmügdén, 1991] Schmügdén, K. (1991). The k-moment problem for compact semi-algebraic sets. *Math. Ann.*, 289:203–206.
- [Schmügdén, 2017] Schmügdén, K. (2017). volume 277 of *Graduate Texts in Mathematics*. Springer.
- [Shukla and Dragotti, 2007] Shukla, P. and Dragotti, P. (2007). Sampling schemes for multidimensional signals with finite rate of innovation. *IEEE Trans. Signal Process.*, 55:3670–3686.
- [Silveti-Falls et al., 2019] Silveti-Falls, A., Molinari, C., and Fadili, J. (2019). Generalized conditional gradient with augmented Lagrangian for composite minimization. *arXiv:1901.01287*.

- [Sinkhorn, 1964] Sinkhorn, R. (1964). A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35:876–879.
- [Steinbrig et al., 2002] Steinbrig, E. Faber, S., Hinkley, S., Macintosh, B., Gavel, D., Gates, E., Christou, J., Le Louarn, M., Raschke, L., Severson, S., Rigaut, F., Crampton, D., J.P., L., and Graham, J. (2002). Characterizing the adaptive optics off-axis point-spread function - i: A semi-empirical method for use in natural-guide-star observations. *Publications of the Astronomical Society of the Pacific*, 114(801):1267–1280.
- [Stetter, 1996] Stetter, H. (1996). Matrix eigenproblems are at the heart of polynomial system solving. *ACM Sigsam Bulletin*, 30(4):22–25.
- [Stoica and Babu, 2012] Stoica, P. and Babu, P. (2012). Sparse estimation of spectral lines: Grid selection problems and their solutions. *IEEE Trans. Signal Process.*, 60(2):962–967.
- [Stone, 1932] Stone, M. (1932). *Linear Transformations in Hilbert Space and Their Applications to Analysis*, volume 15 of *Coll. Publ.* AMS.
- [Ström et al., 1994] Ström, E., Parkvall, S., and Ottersten, B. (1994). Near-fast resistant propagation delay estimators for asynchronous direct-sequence code division multiple access systems. In *Mobile Communications Advanced Systems and Components*, pages 251–260. Springer.
- [Sturmfels, 2002] Sturmfels, B. (2002). *Solving Systems of Polynomial Equations*, volume 97 of *CBMS Regional Conferences Series*. AMS.
- [Szego, 1954] Szego, G. (1954). *On a theorem of C. Carathéodory*.
- [Tang, 2015] Tang, G. (2015). Resolution limits for atomic decompositions via Markov-Bernstein type inequalities. *2015 International Conference on Sampling Theory and Applications (SampTA)*, pages 548–552.
- [Tang et al., 2015] Tang, G., Bhaskar, B., and Recht, B. (2015). Near minimax line spectral estimation. *IEEE Trans. Inform. Theory*, 61(1):499–512.
- [Tang et al., 2013] Tang, G., Bhaskar, B. N., Shah, P., and Recht, B. (2013). Compressed sensing off the grid. *IEEE Trans. Inform. Theory*, 59(11):7465–7490.
- [Tibshirani, 1996] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Statist. Soc.*, 58(1):267–288.
- [Tichavský and Yeredor, 2009] Tichavský, P. and Yeredor, A. (2009). Fast approximate joint diagonalization incorporating weight matrices. *IEEE Trans. on Signal Process.*, 57(3):878–891.
- [Toeplitz, 1911] Toeplitz, O. (1911). über die Fouriersche Entwicklung positiver Funktionen. *Rend. Circ. Mat. Palermo*, 32:191–192.

- [Toh et al., 1999] Toh, K., Todd, M., and Tutuncu, R. (1999). SDPT3 - a Matlab software package for semidefinite programming. *Optimization Methods and Software*, 11:545 – 581.
- [Unser et al., 2017] Unser, M., Fageot, J., and Ward, J. (2017). Splines are universal solutions of linear inverse problems with generalized tv regularization. *SIAM Review*, 59(4):769–793.
- [Vaiter et al., 2015] Vaiter, S., Peyré, G., and Fadili, J. (2015). *Low Complexity Regularization of Linear Inverse Problems*, pages 103–153. App. Numer. Harm. Anal. Springer.
- [Vetterli et al., 2002] Vetterli, M., Marziliano, P., and Blu, T. (2002). Sampling signals with finite rate of innovation. *IEEE Trans. Signal Process.*, 50(6).
- [Waldspurger and Waters, 2018] Waldspurger, I. and Waters, A. (2018). Rank optimality for the Burer-Monteiro factorization.
- [Wolfe, 1970] Wolfe, P. (1970). Convergence theory in nonlinear programming. *Integer and Nonlinear Program*.
- [Yang and Xie, 2016] Yang, Z. and Xie, L. (2016). Exact joint sparse frequency recovery via optimization methods. *IEEE Trans. Signal Process.*, 64(19):5145–5157.
- [Yang et al., 2016] Yang, Z., Xie, L., and Stoica, P. (2016). Vandermonde decomposition of multilevel toeplitz matrices with applications to multidimensional super-resolution. *IEEE Trans. Inform. Theory*, 62(6):3685–3701.
- [Yeredor, 2002] Yeredor, A. (2002). Non-orthogonal joint diagonalization in the least-squares sense with application in blind source separation. *IEEE Trans. Signal Process.*, 50(7):1545–1553.
- [Yuan and Lin, 2006] Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Statist. Soc.*, 68(1):49–67.
- [Yurtsever et al., 2018] Yurtsever, A., Fercoq, O., Locatello, F., and Cevher, V. (2018). A conditional gradient framework for composite convex minimization with applications to semidefinite programming. In *ICML*, volume 80, pages 5713–5722.

RÉSUMÉ

Cette thèse propose des avancées théoriques et algorithmiques pour les relaxations semi-définies positives et leurs applications en science des données. Ces relaxations, dites de Lasserre, fondées sur la substitution aux mesures boréliennes de leurs moments trigonométriques, permettent de résoudre des problèmes de super-résolution sans discrétisation spatiale, mais nécessitent en contrepartie la résolution de problèmes d'optimisation convexe de grande taille. Les contributions de cette thèse montrent comment faire passer ces méthodes à l'échelle en exploitant certaines propriétés et invariances des problèmes d'imagerie. Nous proposons dans un premier temps une nouvelle méthode pour la reconstruction de mesures continues à partir d'un nombre fini de leurs moments, reposant sur un algorithme de co-diagonalisation approchée. Nous étudions ensuite le problème de super-résolution, sous sa forme variationnelle appelée BLASSO, et son approximation par la hiérarchie de Lasserre. Une étape préalable de projection spectrale de l'opérateur d'acquisition rend possible cette approximation et permet également son implémentation efficace via un nouvel algorithme, le Fourier-based Frank-Wolfe (FFW), tirant profit de la structure convolutive et de faible rang des matrices impliquées. Nous appliquons notre méthode sur des données de microscopie par fluorescence. Enfin, combinant la reconstruction de mesures continues avec l'implémentation rapide de FFW, nous employons les hiérarchies de Lasserre afin d'approcher le transport optimal entre deux mesures, qui peut également être couplé avec le Blasso.

MOTS CLÉS

Hiérarchies de Lasserre, programmation SDP, méthode de Prony, moments trigonométriques, matrices de moments, super-résolution, parcimonie, BLASSO, Frank-Wolfe, Burer-Monteiro, transport optimal, distances de Wasserstein

ABSTRACT

This thesis proposes theoretical and algorithmic advances for positive semi-definite relaxations and their applications in data science. These so-called Lasserre's hierarchies allow one to solve super-resolution problems without resorting to spatial discretization, by replacing measures with their trigonometric moments. However, they require the resolution of large convex optimization problems. The contributions of this thesis show how to scale these methods by exploiting certain properties and invariances of imaging problems. We first propose a new approach for the recovery of continuous measures from a finite number of moments, based on the approximate joint-diagonalization of a few non-Hermitian matrices. Then, we study the super-resolution problem, and its approximation by Lasserre's hierarchy. A preliminary step consisting in the spectral projection of the forward operator makes this approximation possible, and further allows a fast implementation, the Fourier-based Frank Wolfe (FFW), taking advantage of the convolutional and low-rank structure of the involved matrices. We apply our method on fluorescence microscopy data. Finally, combining the recovery of continuous measures with the fast implementation of FFW, we employ Lasserre's hierarchies to approximate the optimal transport between measures, which can also be coupled with the Blasso.

KEYWORDS

Lasserre's hierarchy, semidefinite programming, Prony's method, trigonometric moments, moment matrices, super-resolution, sparsity, BLASSO, Frank-Wolfe, Burer-Monteiro, optimal transport, Wasserstein distances.

