STRUCTURED LOW-RANK MATRIX RECOVERY VIA OPTIMIZATION METHODS

by

Dehui Yang

© Copyright by Dehui Yang, 2018

All Rights Reserved

A thesis submitted to the Faculty and the Board of Trustees of the Colorado School of Mines in partial fulfillment of the requirements for the degree of Doctor of Philosophy (Electrical Engineering).

Golden, Colorado

Date _____

Signed: _____
Dehui Yang

Signed: _____
Dr. Michael B. Wakin
Thesis Advisor

Golden, Colorado

Date _____

Signed: _____
Dr. Atef Elsherbeni
Professor and Head
Department of Electrical Engineering

ABSTRACT

From single-molecule microscopy in biology, to collaborative filtering in recommendation systems, to quantum state tomography in physics, many scientific discoveries involve solving ill-posed inverse problems, where the number of parameters to be estimated far exceeds the number of available measurements. To make these daunting problems solvable, low-dimensional geometric structures are often exploited, and regularizations that promote underlying structures are used for various inference tasks. To date, one of the most effective and plausible low-dimensional models for matrix data is the low-rank structure, which assumes that columns of the data matrix are correlated and lie in a low-dimensional subspace. This helps make certain matrix inverse problems well-posed. However, in some cases, standard low-rank structure is not powerful enough for modeling the underlying data generating process, and additional modeling efforts are desired. This is the main focus of this research.

Motivated by applications from different disciplines in engineering and science, in this dissertation, we consider the recovery of three instances of structured matrices from limited measurement data, where additional structures naturally occur in the data matrices beyond simple low-rankness. The structured matrices that we consider include i) low-rank and spectrally sparse matrices in super-resolution imaging; ii) low-rank skew-symmetric matrices in pairwise comparisons; iii) and low-rank positive semidefinite matrices in physical and data sciences. Using optimization as a tool, we develop new regularizers and computationally efficient algorithmic frameworks to account for structured low-rankness in solving these ill-posed inverse problems. For some of the problems considered in this dissertation, theoretical analysis is also carried out for the proposed optimization programs. We show that, under mild conditions, the structured low-rank matrices can be recovered reliably from a minimal number of random measurements.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to thank my advisor Prof. Mike Wakin for his various forms of support, patience, encouragement, and sharing during the entire journey of my PhD study. I thank him for spending endless time discussing research problems with me and for giving me an opportunity to teach an undergraduate course, which helped improve my speaking and communication skills significantly. Throughout the years, I also learned very important life lessons from him. His diligence, humility, and optimistic attitude towards research and life will influence me for a long time.

I would also like to thank my committee members, Prof. Tyrone Vincent, Prof. Gongguo Tang, and Prof. DJ Yang, for their encouragement and invaluable input at various stages of my PhD study. My gratitude also extends to the fellow students in our research group, and I appreciate all of the help that I received from them over the years. Beyond Mines, I would like to thank Dr. Brian Eriksson for being a great mentor and for sharing many useful tips on making presentations with me while I was a research intern at Technicolor Research.

Outside of graduate study, I feel lucky being able to study and live in a quiet and beautiful foothill town, where the magnificent Rockies are in our backyard. One of the fond memories is during summer time, I took a long drive to the southwest of Colorado, hiked or backpacked in some remote wilderness of the San Juan Mountains, and was amazed by the spectacular mountain vistas and dynamic weather. I am grateful to have the opportunity to experience these great outdoor adventures.

Finally, I would like to thank my parents, Jianping Yang and Saifei Chen, for their unconditional love and support of my study in US. Although you did not know this research would take me five years to complete, you always believed in me and helped me build confidence when I lost perspective. I hope to have more time with you soon after graduation.

To my parents.

CHAPTER 1

INTRODUCTION

In this chapter, we start with a discussion of models in signal processing and machine learning. Then, we introduce a particular concise model, the low-rank model, which plays a foundational role in all of the topics developed throughout this dissertation. Finally, we outline the contributions of this dissertation as well as the organization of remaining chapters.

1.1 Models in Signal Processing and Machine Learning

Models, broadly speaking, can be considered as descriptions of the structure of the underlying objects. Powerful, accurate models will help practitioners better understand many real-world phenomena across the physical and data sciences. As a result, more efficient and reliable algorithmic frameworks can be designed by taking the corresponding models into consideration. The last decade has witnessed the revolutionary development of low-dimensional models and their roles in many signal processing and machine learning applications. By employing data-specific models, we can reveal redundancy as much as possible within data while keeping useful information intact. For instance, compressive sensing [1], a new paradigm for acquiring and compressing signals through random linear measurements, relies on the notion of sparsity, a typical low-dimensional structure that exists among many real-world signals, such as images, videos, etc. Sparsity not only provides a new metric for measuring information content contained within a signal and informs us how many compressive samples to take without loss of signal information, but also helps us distinguish signals from less structured ones, such as noise. Matrix completion is another example of using low-dimensional models [2], which finds applications in collaborative filtering. It poses the assumption that the incomplete data matrix under consideration is low-rank if the missing entries of the matrix are filled with reasonable and accurate numerical values. Other applications where low-dimensional geometric models play a crucial role include robust principal component

analysis (PCA) [3] where both sparsity and low-rankness are utilized, and spectral analysis with limited observations where a continuous analogy of sparsity is deployed, etc. For a more detailed review of these low-dimensional models and their practical applications, refer to the literature review sections in **Chapter 2**.

1.2 The Low-Rank Model

Rank is a core concept in linear algebra and is also one of the most important properties of a matrix [4]. Mathematically speaking, low-rankness means that the dimension of the vector space spanned by columns or rows of a matrix is relatively small compared to the ambient dimension of the columns or rows. This subspace structure has found numerous applications for regularizing and solving ill-posed inverse problems across signal processing and machine learning. Low-rankness implies that when data are collected as a matrix, a certain amount of redundancy is present, and the degree of the redundancy depends on the exact rank. Assume that $\boldsymbol{X} \in \mathbb{R}^{M\times N}$ is a rank-r matrix, i.e., $\operatorname{rank}(\boldsymbol{X}) = r$. Then, a simple calculation via a singular value decomposition (SVD) shows that the number of degrees of freedom in $\boldsymbol{X}$ is equal to $r(M + N - r)$. When M, N are much larger than r, we have $r(M+N-r) \approx r(M+N)$, which is much smaller than the ambient dimension of the matrix $\boldsymbol{X}$, $M \times N$. This simple calculation also shows that for a rank-r matrix, the redundancy is roughly proportional to $\frac{\min(M,N)-r}{\min(M,N)}$. This indicates that it might be possible to recover a low-rank matrix from partial or compressive measurements since correlation happens among columns of the matrix. It also implies that for a generic rank-r matrix $\boldsymbol{X} \in \mathbb{R}^{M\times N}$, when the number of observations is smaller than the number of degrees of freedom, i.e., $r(M+N-r)$, it is impossible to recover $\boldsymbol{X}$ from these observations.

Matrix sensing and matrix completion [2, 5] are two related problems that explicitly use the prior knowledge that the underlying matrix is low-rank. In matrix sensing, given a set of random linear measurements of a low-rank matrix, one would like to recover the matrix via a computationally feasible procedure, such as convex programming, or an iterative greedy algorithm, etc. Whereas in matrix completion observations are partial entries of a low-

rank matrix, the task is to infer the missing entries from available ones. It is possible that one could recover the low-rank matrix when the number of observations is on the order of the number of degrees of freedom in the matrix. In this dissertation, we develop efficient optimization schemes for solving such kinds of linear inverse problems.

1.3 Contributions and Organization

The main contributions of this dissertation include i) identifying important research problems across physical and data sciences, where low-dimensional geometric models (in particular, low-rank models) prevail and play an important role; ii) formulating appropriate inverse problems as the recovery of structured low-rank matrices from incomplete linear measurements with different sensing modalities; iii) utilizing convex optimization as a common tool for solving the corresponding linear inverse problems for recovery and useful information extraction.

A central theme that often arises throughout this dissertation is the sample complexity bound of solving structured low-rank matrix recovery problems from a set of random linear measurements using convex optimization. For some of the problems presented in this dissertation, under some mild conditions, we show that the sample complexity for exact recovery using convex programming is proportional to the number of degrees of freedom in the low-rank matrix under consideration. This implies that our proposed strategy is optimal, and there is little room for further improvement in terms of sample complexity.

Chapter 2 gives a detailed overview of some of the most useful and relevant low-dimensional geometric models. Then, a brief review of convex optimization and duality theory is also given as technical background. Lastly, we show that convex optimization can be combined with low-dimensional geometric models to yield an effective approach for solving linear inverse problems in signal processing and machine learning.

From **Chapter 3** to **Chapter 6**, we share three stories of recovering different structured low-rank matrices from a set of random linear measurements via optimization methods. This is also the main contribution of this dissertation.

In **Chapter 3** and **Chapter 4**, we present a structured matrix that is not only low-rank but also spectrally sparse. This matrix arises naturally in physical sciences and structural vibrations. Specifically, in **Chapter 3**, motivated by applications in optical microscopy and astronomy, we develop a mathematical framework of super-resolution of point sources from low-frequency measurements when the point spread functions associated with the underlying physical phenomena are also unknown. We show that this non-stationary blind super-resolution problem is solvable under some mild conditions on the point spread functions and it can be solved by recovering a spectrally sparse and low-rank matrix from a minimal number of linear measurements. This general framework also leads to new insights of modal analysis in structural health monitoring, which is the main topic of **Chapter 4**. In particular, in **Chapter 4**, we show that the problem of estimating mode shapes and natural frequencies from measurement data can be recast as the decomposition of a spectrally sparse and low-rank matrix studied in **Chapter 3**. Based on this observation, we propose a few random compression schemes, both in time and spatial domains, for reducing measurements while still being able to recover mode shapes and natural frequencies. This is achieved by solving linear inverse problems via computationally feasible convex programs.

The second story of the structured low-rank matrices discussed in this dissertation originates from the problem of ranking and pairwise comparisons in data science. It is the low-rank skew-symmetric matrices. In **Chapter 5**, we show that pairwise comparisons are naturally related to skew-symmetric matrices. Our contribution is the development of a new low-rank model that captures non-transitive behavior prevailed in pairwise comparisons. Then, we use matrix completion as a tool for dealing with missing data scenarios in practice. Along the way, we also develop new efficient alternating minimization schemes for pairwise comparison matrix completion.

Then, in **Chapter 6**, motivated by applications in statistics and data science, quantum state tomography in physics, we consider the recovery of a low-rank positive semidefinite matrix from its quadratic measurements. With the aid of a symmetric factorization, we

formulate the problem as an efficient second-order cone program given a valid anchor matrix. Furthermore, we propose a practical approach to construct a valid anchor matrix and provide insights into the exact recovery condition of the positive semidefinite matrix from measurement data. Numerical simulations are conducted to support the effectiveness of our proposed approach.

In **Chapter 7**, we present a few future directions of extending the topics presented in this dissertation.

CHAPTER 2

BACKGROUND - MODELS AND CONVEX OPTIMIZATION

In this chapter, we first give a short tour of a few common low-dimensional parsimonious models and their variants. Then, we review some concepts in convex optimization and duality theory, which will be frequently used in the following chapters. Lastly, we show how convex optimization, as a general tool, can be used for recovering low-dimensional geometric models from near optimal numbers of compressive measurements.

2.1 Notation

Throughout the dissertation, the following notation is adopted. We use boldface letters $\boldsymbol{X}, \boldsymbol{Y}$ and $\boldsymbol{x}, \boldsymbol{y}$ to denote matrices and vectors, respectively. For a vector $\boldsymbol{v}$, $\|\boldsymbol{v}\|_2$ is used to denote the ℓ_2 norm of $\boldsymbol{v}$. For a matrix $\boldsymbol{X}$, $\|\boldsymbol{X}\|$, $\|\boldsymbol{X}\|_F$ and $\sigma_{\min}(\boldsymbol{X})$ represent the operator norm, the Frobenius norm, and the minimum non-zero singular value of the matrix $\boldsymbol{X}$, respectively. An $M \times N$ zero matrix is denoted as $\mathbf{0}_{M\times N}$. We also use $\mathbf{0}_M$ and $\boldsymbol{I}_M$ to denote an $M \times M$ zero matrix and an $M \times M$ identity matrix, respectively. We use the matrix inequality notation $\boldsymbol{X} \preceq \boldsymbol{Y}$ to represent that $\boldsymbol{Y} - \boldsymbol{X}$ is positive semidefinite. Similarly, we use the vector inequality notation $\boldsymbol{x} \preceq \boldsymbol{y}$ to represent that $\boldsymbol{y}-\boldsymbol{x}$ is element-wise nonnegative. Conventional notations $\langle \cdot, \cdot \rangle$, trace$(\cdot)$, $(\cdot)^H$, $(\cdot)^T$ and $(\cdot)^*$ are used to denote the inner product, trace, Hermitian, transpose, and conjugation operations, respectively. For a set Ω, we use $|\Omega|$ to denote its cardinality. $\mathbb{E}$ and $\mathbb{P}\{\cdot\}$ denote expectation and probability of the underlying event. We use a calligraphic letter $\mathcal{B}$ to denote a linear operator.

2.2 Parsimonious Modeling

In many applications across science and engineering, we are often faced with solving ill-posed inverse problems, in which the number of measurements is much smaller than the number of unknowns that are of interest. For example, in biological microscopy, due to

the diffraction limit, the resolution of the light microscopy is not adequate as the high-frequency content of the molecules is missing. In order to achieve a higher resolution than the diffraction limit, we have to solve an ill-posed inverse problem from the low-frequency measurements. Additionally, in collaborative filtering, the task is to interpolate the unknown values of the missing entries in a data matrix for recommendations [6]. Again, one has to solve an inference problem merely based on available partial observations of the data matrix.

In general, it is impossible to solve these kinds of problems as the parameters to be inferred far exceed the available measurement constraints. To deal with the underdetermined nature of these inverse problems, one powerful strategy is to take advantage of the structures that the underlying parameters might have. These structures could help us constrain the search space where the plausible parameters lie, and thus make the daunting inverse problems solvable.

2.2.1 Sparsity

Sparsity is a parsimonious model that many real-world signals, such as natural images and videos, obey [7]. Given an N-dimensional signal $\boldsymbol{x}$, we say that $\boldsymbol{x}$ is K-sparse if and only if the number of nonzero entries in $\boldsymbol{x}$ is less than or equal to K. We can also write in a compact mathematical form $\|\boldsymbol{x}\|_0 = K$, where $\|\cdot\|_0$[1] counts the number of entrywise nonzero elements of the underlying argument. Let us take $N = 5$, $K = 2$, $\boldsymbol{x} = \begin{bmatrix} 2 & 0 & 5 & 0 & 0 \end{bmatrix}^T$ as an illustrative example. By observation, we can write

$$\boldsymbol{x} = \begin{bmatrix} 2 \\ 0 \\ 5 \\ 0 \\ 0 \end{bmatrix} = 2 \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 0 \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 5 \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 0 \cdot \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} + 0 \cdot \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \tag{2.1}$$

where $\boldsymbol{x}$ is a linear combination of a few canonical basis vectors in $\mathbb{R}^5$. Note that in many scenarios, we assume that $K \ll N$, which means that $\boldsymbol{x}$ can be written as a linear combination of at most K canonical basis vectors in $\mathbb{R}^N$. Furthermore, the notion of sparsity

[1]Here we have used the conventional norm notation $\|\cdot\|_0$, which is often referred as ℓ_0-norm in the literature. However, technically, $\|\cdot\|_0$ is not a valid norm.

can be easily related to the concept of a union of subspaces. When the locations of the K nonzero entries of $\boldsymbol{x}$ are fixed, $\boldsymbol{x} \in \mathbb{R}^N$ lies in a K-dimensional subspace. As there are only $\binom{N}{K}$ possible combinations of the locations of nonzero entries in $\boldsymbol{x}$, in general $\boldsymbol{x}$ lies in the union of subspaces with the number of subspaces being $\binom{N}{K}$. Thus, we can see that in contrast to general signals lying in N-dimensional space, K-sparse signals lie in a highly restricted union of subspaces embedded in $\mathbb{R}^N$. Therefore, when the high dimensional signal for recovery in an ill-posed inverse problem obeys a sparse model, we will see that certain regularization that helps promote sparsity can be used for efficient recovery from a minimal number of linear measurements.

Many variants and extensions beyond simple sparsity have been developed in the literature over the last decade. For instance, model-based sparsity has been developed by incorporating structural dependencies between the values and locations of the nonzero sparse coefficients [8]. In other cases when multiple similar signals are collected, certain types of joint sparse models can be utilized [9]. Capturing both the inner- and inter-signal structures will help improve the fidelity of signal processing tasks such as recovery and inference.

2.2.2 Low-rankness

Just as sparsity is a concise model for high dimensional vectors, low-rankness is a parsimonious model for matrices [5]. Denote $\boldsymbol{X} \in \mathbb{R}^{M \times N}$ as an $M \times N$ data matrix and assume that $\text{rank}(\boldsymbol{X}) = r$. Then, we see that the number of linearly independent columns or rows in $\boldsymbol{X}$ is r. Another convenient way to check the rank of $\boldsymbol{X}$ is by examining a singular value decomposition (SVD) of $\boldsymbol{X}$. It is well-known that the rank of $\boldsymbol{X}$ is equal to the number of nonzero singular values in the SVD of $\boldsymbol{X}$. In fact, the SVD of $\boldsymbol{X}$ also gives a rank-one decomposition of $\boldsymbol{X}$. If $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T$ is a reduced singular value decomposition of $\boldsymbol{X}$, then we have

$$
\begin{aligned}
\boldsymbol{X} &= \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T \\
&= \begin{bmatrix} \boldsymbol{u}_1 & \cdots & \boldsymbol{u}_r \end{bmatrix} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \begin{bmatrix} \boldsymbol{v}_1^T \\ \vdots \\ \boldsymbol{v}_r^T \end{bmatrix} \\
&= \sum_{i=1}^{r} \sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^T,
\end{aligned}
$$

which shows that $\boldsymbol{X}$ can be decomposed as a superposition of r rank-one matrices $\sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^T$. Note that the left singular vectors $\boldsymbol{u}_i, i = 1, \cdots, r$, are orthonormal with each other and the right singular vectors $\boldsymbol{v}_i, i = 1, \cdots, r$, are orthonormal with each other as well. In addition, we can count the number of degrees of freedom in $\boldsymbol{X}$ from the SVD by adding the number of degrees of freedom in $\boldsymbol{U}$, $\boldsymbol{\Sigma}$, and $\boldsymbol{V}$. A simple calculation reveals the number of degrees of freedom in the matrix $\boldsymbol{X}$ is $r(M+N-r)$. This implies that the number of degrees of freedom in a rank-r matrix $\boldsymbol{X}$ is much smaller than the actual matrix size when r is much smaller than $\min(M, N)$. This further suggests that it is possible to recover the rank-r matrix $\boldsymbol{X}$ only from a *sketch* of $\boldsymbol{X}$.

An extension of a low-rank model for a matrix is simultaneously low-rank and sparse [10]. Consider a matrix $\boldsymbol{X}$ of size $M \times N$. In the simultaneously low-rank and sparse model, the smallest submatrix that contains all of the nonzero entries of $\boldsymbol{X}$ is of size $s_1 \times s_2$ with $s_1 \ll M$ and $s_2 \ll N$. Furthermore, the $s_1 \times s_2$ submatrix is low-rank and $\text{rank}(\boldsymbol{X}) = r$. Hence, the number of degrees of freedom in $\boldsymbol{X}$ is $r(s_1 + s_2 - r)$, which is much smaller than the low-rank model.

2.2.3 Generalization to simple models

From our discussions of sparsity in Section 2.2.1 to those of low-rankness in Section 2.2.2, a common theme is to decompose the underlying object into a linear combination of a few elementary building blocks. For an arbitrary N-dimensional signal $\boldsymbol{x}$, we can always decompose it with building blocks $\mathcal{S} = \{\pm\boldsymbol{e}_1, \pm\boldsymbol{e}_2, \cdots, \pm\boldsymbol{e}_N\}$, where each $\boldsymbol{e}_i$ is a length-N canonical vector with its i-th position 1 and every other location 0. We call the set $\mathcal{S}$ an

atomic set and each element within $\mathcal{S}$ an atom. Then, a K-sparse signal can be decomposed as a superposition of at most K atoms from the building blocks in $\mathcal{S}$. Although this is not the only way to decompose a K-sparse vector, it is indeed the simplest way to do it. Similarly, for the matrix case, we can construct the atomic set $\mathcal{R}$ as

$$\mathcal{R} = \left\{ \boldsymbol{u}\boldsymbol{v}^T : \|\boldsymbol{u}\|_2 = 1, \|\boldsymbol{v}\|_2 = 1, \boldsymbol{u} \in \mathbb{R}^M, \boldsymbol{v} \in \mathbb{R}^N \right\}.$$

Note that the atomic set $\mathcal{R}$ is continuously parametrized, and there is an infinite number of elements in $\mathcal{R}$ and each element is a rank-one matrix. Then an arbitrary matrix $\boldsymbol{X}$ can be decomposed as a linear combination of elements from the set $\mathcal{R}$. In particular, for a rank-r matrix $\boldsymbol{X}$, we find the simplest way to decompose it via a reduced SVD.

Recent foundational work [11] generalizes sparsity and low-rankness to much broader simple models. It turns out that one can construct many other atomic sets beyond $\mathcal{S}$ for sparsity and $\mathcal{R}$ for low-rankness. For instance, in applications such as parameter estimation and radar imaging, one is often confronted with estimating the frequency content within a signal. In practice, the frequencies are often continuously parameterized and the building block [12]

$$\boldsymbol{e}_f = \begin{bmatrix} e^{i2\pi f0} \\ e^{i2\pi f1} \\ \vdots \\ e^{i2\pi f(N-1)} \end{bmatrix}$$

plays an important role. Assume that f is normalized, i.e., $0 \leq f < 1$. As we change f from 0 to 1, we obtain an atomic set

$$\mathcal{F} = \left\{ \boldsymbol{e}_f : \boldsymbol{e}_f = \begin{bmatrix} e^{i2\pi f0} \\ e^{i2\pi f1} \\ \vdots \\ e^{i2\pi f(N-1)} \end{bmatrix}, 0 \leq f < 1 \right\}.$$

In this case, there is an infinite number of atoms in the atomic set $\mathcal{F}$. We will see that this atomic set plays an important role in blind super-resolution in Chapter 3.

Another example of atomic set is

$$\mathcal{R}_s = \left\{ \boldsymbol{u}\boldsymbol{v}^T : \|\boldsymbol{u}\|_2 = 1, \|\boldsymbol{v}\|_2 = 1, \|\boldsymbol{v}\|_0 = s, \boldsymbol{u} \in \mathbb{R}^M, \boldsymbol{v} \in \mathbb{R}^N \right\},$$

which finds applications in sparse principle component analysis [13]. Note that the difference between the atomic set $\mathcal{R}$ and $\mathcal{R}_s$ lies in the fact that the factor $\boldsymbol{v}$ is sparse in $\mathcal{R}_s$.

2.3 Convex Optimization and Duality Theory

Convex optimization is a well-established subject in the optimization literature. We review some of the fundamental concepts that will be repeatedly used in this research. Most of the material presented here can be found in the textbook *Convex Optimization* by Stephen Boyd and Lieven Vandenberghe [14].

2.3.1 Basic concepts

Definition 2.1. *Convex Set*

A set $\mathcal{S}$ is convex if and only if for any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{S}$ and any $0 \leq \theta \leq 1$, we have

$$\theta \boldsymbol{x}_1 + (1-\theta)\boldsymbol{x}_2 \in \mathcal{S}.$$

Definition 2.2. *Convex Hull*

The convex hull of a set $\mathcal{S}$ is defined as the set given by

$$conv(\mathcal{S}) = \{\theta_1 \boldsymbol{x}_1 + \theta_2 \boldsymbol{x}_2 + \cdots + \theta_k \boldsymbol{x}_k \mid \boldsymbol{x}_i \in \mathcal{S}, i = 1, \ldots, k, \ \theta_1 + \theta_2 + \cdots + \theta_k = 1\}.$$

Definition 2.3. *Convex Cone*

A convex set $\mathcal{C}$ is a cone if and only if for arbitrary $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{C}$ and any $\theta_1 \geq 0, \theta_2 \geq 0$, we have

$$\theta_1 \boldsymbol{x}_1 + \theta_2 \boldsymbol{x}_2 \in \mathcal{C}.$$

Definition 2.4. *Polar Cone [11]*

$$\mathcal{C}^\star = \left\{ \boldsymbol{x} \in \mathbb{R}^N : \langle \boldsymbol{x}, \boldsymbol{z} \rangle \leq 0, \forall \boldsymbol{z} \in \mathcal{C} \right\}.$$

Definition 2.5. *Convex Function*

A function $f(\boldsymbol{x}) : \mathbb{R}^N \to \mathbb{R}$ *with* $\boldsymbol{x} \in \mathcal{S}$*, is convex if and only if*

- *the domain* $\mathcal{S}$ *itself is convex*
- *for any* $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{S}$ *and* $0 \leq \theta \leq 1$*, the following holds*

$$f(\theta \boldsymbol{x}_1 + (1-\theta)\boldsymbol{x}_2) \leq \theta f(\boldsymbol{x}_1) + (1-\theta) f(\boldsymbol{x}_2).$$

Note that every valid norm $\|\cdot\|$ is a convex function.

Definition 2.6. *Constrained Optimization, Lagrangian Multipliers, and Duality*

Constrained optimization (primal) problems typically have the form

$$\begin{aligned} &\textit{minimize} \quad f(\boldsymbol{x}) \\ &\textit{subject to} \quad f_i(\boldsymbol{x}) \leq 0, \ i = 1, \ldots, p \\ &\qquad\qquad\ \ g_i(\boldsymbol{x}) = 0, \ i = 1, \ldots, q \end{aligned} \tag{2.2}$$

and the Lagrangian $\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\mu}) : \mathbb{R}^N \times \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}$ *is defined as*

$$\mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\boldsymbol{x}) + \sum_{i=1}^{p} \boldsymbol{\lambda}(i) f_i(\boldsymbol{x}) + \sum_{i=1}^{q} \boldsymbol{\mu}(i) g_i(\boldsymbol{x})$$

where we have denoted $\boldsymbol{\lambda}(i)$*, the* i*th element of* $\boldsymbol{\lambda}$*, as the Lagrangian multiplier associated with the inequality constraint* $f_i(\boldsymbol{x})$*. Similarly,* $\boldsymbol{\mu}(i)$*, the* i*th component of* $\boldsymbol{\mu}$*, is the Lagrangian multiplier associated with the equality constraint* $g_i(\boldsymbol{x})$*. Then the Lagrangian dual function can be defined as*

$$\begin{aligned} d(\boldsymbol{\lambda}, \boldsymbol{\mu}) &= \inf_{\boldsymbol{x} \in \mathcal{S}} \mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \\ &= \inf_{\boldsymbol{x} \in \mathcal{S}} \left(f(\boldsymbol{x}) + \sum_{i=1}^{p} \boldsymbol{\lambda}(i) f_i(\boldsymbol{x}) + \sum_{i=1}^{q} \boldsymbol{\mu}(i) g_i(\boldsymbol{x}) \right) \end{aligned}$$

Denote $p^\star$ *the optimal value of the primal problem (2.2). When* $\boldsymbol{\lambda} \succeq \mathbf{0}$*, an important relationship between the dual function and* $p^\star$ *is*

$$d(\boldsymbol{\lambda}, \boldsymbol{\mu}) \leq p^\star,$$

which indicates that the dual function provides a lower bound on the optimum of the primal problem. This motivates us to find the tightest lower bound on $p^\star$ *by maximizing the dual*

function, i.e.,

$$\begin{array}{ll} \textit{maximize} & d(\boldsymbol{\lambda}, \boldsymbol{\mu}) \\ \textit{subject to} & \boldsymbol{\lambda} \succeq \mathbf{0}. \end{array} \tag{2.3}$$

Note that the primal problem is not necessarily convex as either $f(\boldsymbol{x})$ or $f_i(\boldsymbol{x}), i = 1, \ldots, p$, $g_i(\boldsymbol{x}), i = 1, \ldots, q$, can possibly be nonconvex. When $f(\boldsymbol{x}), f_i(\boldsymbol{x}), i = 1, \ldots, p$ are convex functions and $g_i(\boldsymbol{x})$ are affine functions, i.e., $g_i(\boldsymbol{x}) = \boldsymbol{a}_i^T \boldsymbol{x} - b_i$, the primal problem (2.2) becomes convex. In addition, if Slater's condition[2] is satisfied, we have

$$d^\star := d(\boldsymbol{\lambda}^\star, \boldsymbol{\mu}^\star) = p^\star.$$

This implies that one could possibly find the optimal value of the primal problem (2.2) by solving the dual problem (2.3) when program (2.2) is convex.

2.3.2 SDP and SOCP

We review two classes of convex optimization problems that are of particular interest and find numerous applications in engineering. They are the semidefinite program (SDP) and second-order cone program (SOCP). Both SDP and SOCP can be solved in polynomial time with available solvers using an interior point method.

Denote $\mathcal{S}^N$ as the set of all possible $N \times N$ symmetric matrices. The standard form of an SDP is given as follows:

$$\begin{array}{ll} \text{minimize} & \text{trace}(\boldsymbol{C}\boldsymbol{X}) \\ \text{subject to} & \text{trace}(\boldsymbol{A}_i \boldsymbol{X}) = b_i, \ i = 1, \ldots, q, \\ & \boldsymbol{X} \succeq \mathbf{0} \end{array} \tag{2.4}$$

where both the variable $\boldsymbol{X}$ and the constant matrices $\boldsymbol{C}, \boldsymbol{A}_i, i = 1, \ldots, q$, belong to the set $\mathcal{S}^N$.

A particular subset of SDP is called SOCPs. They have the form

$$\begin{array}{ll} \text{minimize} & \boldsymbol{a}^T \boldsymbol{x} \\ \text{subject to} & \|\boldsymbol{A}_i \boldsymbol{x} + \boldsymbol{d}_i\|_2 \leq \boldsymbol{c}_i^T \boldsymbol{x} + b_i, \ i = 1, \ldots, p, \\ & \boldsymbol{F}\boldsymbol{x} = \boldsymbol{g} \end{array} \tag{2.5}$$

[2]Slater's condition is satisfied if there exists some $\boldsymbol{x}$ that belongs to the relative interior of $\mathcal{S}$ such that $f_i(\boldsymbol{x}) < 0, i = 1, \ldots, p$ and $\boldsymbol{a}_i^T \boldsymbol{x} = b_i, i = 1, \ldots, q$.

where $\boldsymbol{x} \in \mathbb{R}^N$ is the variable and $\boldsymbol{A}_i \in \mathbb{R}^{N_i \times N}$ and $\boldsymbol{F} \in \mathbb{R}^{q \times N}$. The constraints $\|\boldsymbol{A}_i\boldsymbol{x}+\boldsymbol{d}_i\|_2 \leq \boldsymbol{c}_i^T\boldsymbol{x} + b_i, i = 1, \ldots, p$, are called second-order cone constraints.

2.4 Solving Inverse Problems via Optimization

In this section, we introduce the atomic norm and explain how an atomic norm can be used for regularization and for recovering simple models from compressive measurements in various ill-posed inverse problem settings.

2.4.1 Recovery via an atomic norm minimization

Given an atomic set $\mathcal{A}$ with its elements being extreme points of $\text{conv}(\mathcal{A})$, we define the atomic norm as

$$\|\boldsymbol{x}\|_{\mathcal{A}} = \inf\{t > 0 : \boldsymbol{x} \in t \ \text{conv}(\mathcal{A})\} \tag{2.6}$$

Then, the dual norm of $\|\boldsymbol{x}\|_{\mathcal{A}}$ is

$$\|\boldsymbol{x}\|_{\mathcal{A}}^{\star} = \sup\{\langle \boldsymbol{x}, \boldsymbol{a}\rangle , \ \boldsymbol{a} \in \mathcal{A}\}.$$

Now, assume that $\boldsymbol{x}^\star$ is a linear combination of a few elements from the atomic set $\mathcal{A}$, i.e.,

$$\boldsymbol{x}^\star = \sum_{i=1}^{K} c_i \boldsymbol{a}_i, \ \boldsymbol{a}_i \in \mathcal{A}. \tag{2.7}$$

Then, we can try to find $\boldsymbol{x}^\star$ from linear measurements $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}^\star$ by solving the following program

$$\begin{array}{ll} \text{minimize} & \|\boldsymbol{x}\|_{\mathcal{A}} \\ \text{subject to} & \boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x} \end{array} \tag{2.8}$$

Note that the dual problem of (2.8) is given by

$$\begin{array}{ll} \text{maximize} & \boldsymbol{y}^T\boldsymbol{z} \\ \text{subject to} & \|\boldsymbol{\Phi}^*\boldsymbol{y}\|_{\mathcal{A}}^* \leq 1, \end{array}$$

where $\boldsymbol{\Phi}^*$ denotes the adjoint operator of the linear measurement operator $\boldsymbol{\Phi}$.

We illustrate two examples of using (2.7) as a way to construct a signal $\boldsymbol{x}^\star$ from an atomic set $\mathcal{A}$ and using the atomic norm minimization (2.8) to recover the signal from compressive measurements.

Example 1: Compressive Sensing

From Section 2.2.3, we see that for an arbitrary signal $\boldsymbol{x} \in \mathbb{R}^N$, the atomic set[3] is given by

$$\mathcal{A} = \{\pm \boldsymbol{e}_1, \pm \boldsymbol{e}_2, \cdots, \pm \boldsymbol{e}_N\}. \tag{2.9}$$

Now, assume that $\boldsymbol{x}^\star$ is K-sparse, i.e., $\boldsymbol{x}^\star = \sum_{i=1}^K c_i \boldsymbol{a}_i, \boldsymbol{a}_i \in \mathcal{A}$ with $\mathcal{A}$ defined in (2.9). Then (2.8) becomes a sparse recovery problem in compressive sensing:

$$\begin{array}{ll} \text{minimize} & \|\boldsymbol{x}\|_1 \\ \text{subject to} & \boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x} \end{array} \tag{2.10}$$

where $\|\boldsymbol{x}\|_1 = \sum_{i=1}^N |\boldsymbol{x}(i)|$ is the ℓ_1 norm of $\boldsymbol{x}$ and $\boldsymbol{\Phi} \in \mathbb{R}^{m \times N}$ is a random matrix with independent and identically distributed (i.i.d.) Gaussian entries with mean zero and variance $\frac{1}{m}$. Note that program (2.10) can be easily recast as a linear program.

Theorem 2.1. *[15] Assume that $\boldsymbol{x}^\star \in \mathbb{R}^N$ is a K-sparse signal and let $\beta > 1$. If $\boldsymbol{\Phi}$ is populated using i.i.d. Gaussian entries with mean zero and variance $\frac{1}{m}$ and $m \geq 2\beta K \log N + K$, then solving (2.10) recovers $\boldsymbol{x}^\star$ exactly with probability at least $1 - 2N^{-f(\beta,K)}$, where $f(\beta, K) = \left(\sqrt{\frac{\beta}{2K} + \beta - 1} - \sqrt{\frac{\beta}{2K}}\right)^2$.*

Theorem 2.1 indicates that when the linear measurement operator $\boldsymbol{\Phi}$ is Gaussian and m is proportional to the sparsity level K, up to a logarithmic factor, one can recover the K-sparse vector exactly with high probability. There is little room for further improvement as the information-theoretical lower bound for the number of measurements is K.

Example 2: Matrix Sensing and Matrix Completion

Recall that, in Section 2.2.3, we defined the atomic set for matrices as

$$\mathcal{A} = \left\{\boldsymbol{u}\boldsymbol{v}^T : \|\boldsymbol{u}\|_2 = 1, \|\boldsymbol{v}\|_2 = 1, \boldsymbol{u} \in \mathbb{R}^M, \boldsymbol{v} \in \mathbb{R}^N\right\}.$$

Then, one can retrieve a rank-r matrix $\boldsymbol{X}^\star = \boldsymbol{U}^\star \boldsymbol{\Sigma}^\star \boldsymbol{V}^{\star T} = \sum_{i=1}^r \sigma_i^\star \boldsymbol{u}_i^\star \boldsymbol{v}_i^{\star T}$ from a set of random Gaussian linear measurements via the following program

$$\begin{array}{ll} \text{minimize} & \|\boldsymbol{X}\|_* \\ \text{subject to} & \boldsymbol{y} = \boldsymbol{\Phi}(\boldsymbol{X}) \end{array} \tag{2.11}$$

[3]Instead of using $\mathcal{S}$ in the previous case, below we use the notation $\mathcal{A}$ for representing an atomic set.

where $\|\boldsymbol{X}\|_* = \sum_{i=1}^{\min(M,N)} \sigma_i$ is the nuclear norm of the matrix $\boldsymbol{X}$ and $\boldsymbol{\Phi}$ is a linear operator. For matrix sensing, each element of $\boldsymbol{y}$ is a (random) linear combination of all of the entries in $\boldsymbol{X}$. However, in the matrix completion problem, the linear operator $\boldsymbol{\Phi}$ is highly structured. Each entry of the measurements $\boldsymbol{\Phi}(\boldsymbol{X})$ is the inner product between $\boldsymbol{\Phi}_i$ and $\boldsymbol{X}$, i.e., $\langle \boldsymbol{\Phi}_i, \boldsymbol{X} \rangle$, where $\boldsymbol{\Phi}_i$ is a zero matrix except for one nonzero entry which is 1. The locations of the nonzero entries of the $\boldsymbol{\Phi}_i$s correspond to the row and column indexes of the observed entries in $\boldsymbol{X}$. Therefore, the measurements $\boldsymbol{y}$ are actually a partially observed version of $\boldsymbol{X}^\star$ with the unobserved entries filled with zeros. For notational convenience, for the matrix completion problem, we rewrite program (2.11) as

$$\begin{array}{ll} \text{minimize} & \|\boldsymbol{X}\|_* \\ \text{subject to} & \mathcal{P}_\Omega(\boldsymbol{X}^\star) = \mathcal{P}_\Omega(\boldsymbol{X}) \end{array} \tag{2.12}$$

where we have replaced the linear operator $\boldsymbol{\Phi}(\cdot)$ with a more amenable notation $\mathcal{P}_\Omega(\cdot)$ and Ω denotes the index set of observed entries.

Theorem 2.2. *[15] Let $\boldsymbol{X}^\star \in \mathbb{R}^{M \times N}$ be a rank-r matrix and $\boldsymbol{\Phi}(\cdot)$ be the linear measurement operator with entries of the corresponding Gaussian matrix being i.i.d. mean zero and variance $\frac{1}{m}$. Then, solving the convex program (2.11) recovers $\boldsymbol{X}^\star$ exactly with probability at least $1 - 2e^{(1-\beta)\max(M,N)/8}$ as long as the number of measurements $m \geq \beta r(3M + 3N - 5r)$ for some $\beta > 1$.*

In order to give a precise theoretical guarantee for the convex program (2.12), we also need the following definition of coherence of $\boldsymbol{X}^\star$, which quantifies the spikiness of $\boldsymbol{X}^\star$.

Definition 2.7. *Coherence [2, 16]*

The coherence of a rank-r matrix $\boldsymbol{X}^\star = \boldsymbol{U}^\star \boldsymbol{\Sigma}^\star \boldsymbol{V}^{\star T} = \sum_{i=1}^r \sigma_i^\star \boldsymbol{u}_i^\star \boldsymbol{v}_i^{\star T}$ is defined as

$$\mu(\boldsymbol{X}^\star) = \max\left\{ \frac{M}{r} \max_{1 \leq i \leq M} \|\boldsymbol{P}_{\boldsymbol{U}^\star} \boldsymbol{e}_i\|^2, \frac{N}{r} \max_{1 \leq i \leq N} \|\boldsymbol{P}_{\boldsymbol{V}^\star} \boldsymbol{e}_i\|^2 \right\}$$

where $\boldsymbol{P}_{\boldsymbol{U}^\star} := \boldsymbol{U}^\star \boldsymbol{U}^{\star T}$ is a projection matrix and $\boldsymbol{P}_{\boldsymbol{V}^\star}$ is defined similarly.

Theorem 2.3. *[17] Assume that each entry of a rank-r matrix $\boldsymbol{X}^\star$ is observed independently with probability p. Then, there exist numerical constants c_0, c_1, and c_2 such that when*

$$p \geq c_0 \frac{\mu(\boldsymbol{X}^\star) r \log^2(M+N)}{\min(M,N)},$$

solving the convex program (2.12) recovers $\boldsymbol{X}^\star$ *exactly with probability at least* $1 - c_1(M+N)^{-c_2}$.

Theorem 2.2 shows that it is possible to recover a rank-r matrix $\boldsymbol{X}^\star$ from a set of Gaussian linear measurements whose cardinality is proportional to the number of degrees of freedom in $\boldsymbol{X}^\star$, up to a constant factor. Theorem 2.3 implies that, for the matrix completion problem, one can recover an incoherent[4] rank-r matrix $\boldsymbol{X}^\star$ from a random selection of the entries whose expected cardinality is proportional to the degree of freedom in $\boldsymbol{X}^\star$, up to a polylogarithmic factor.

2.4.2 Sample complexity

We have shown in Section 2.4.1 that it is possible to recover a sparse signal or a low-rank matrix from a set of Gaussian linear measurements that is proportional to the degree of freedom therein. However, for a generic atomic set and an object obeying a simple model, one may wonder how many (possibly random) linear measurements are needed for exact recovery by solving program (2.8)?

It has been shown that it is possible to develop a generic theoretical guarantee for program (2.8) when the linear measurement operator is Gaussian with mean zero and variance $\frac{1}{m}$.

Definition 2.8. *Tangent Cone [11]*
Given an atomic set $\mathcal{A}$*, the tangent cone at a point* $\boldsymbol{x}$ *with respect to the scaled unit ball* $\|\boldsymbol{x}\|_{\mathcal{A}} conv(\mathcal{A})$ *is defined as*

$$\mathcal{T}_{\mathcal{A}}(\boldsymbol{x}) = cone\left\{\boldsymbol{z} - \boldsymbol{x} : \ \|\boldsymbol{z}\|_{\mathcal{A}} \leq \|\boldsymbol{x}\|_{\mathcal{A}}\right\}.$$

Theorem 2.4. *[11] Assume that* $\boldsymbol{\Phi} : \mathbb{R}^N \to \mathbb{R}^m$ *is a random linear map with i.i.d. Gaussian entries with mean zero and variance* $\frac{1}{m}$*. Let* $\Omega = \mathcal{T}_{\mathcal{A}}(\boldsymbol{X}^\star) \cap \mathbb{S}^{N-1}$ *be the tangent cone* $\mathcal{T}_{\mathcal{A}}(\boldsymbol{X}^\star)$ *constrained to the sphere in* $\mathbb{R}^N$*. Furthermore, let* $\lambda_n - w(\Omega)$ *be defined as follows:*

[4]Here we use *incoherent* to indicate that the coherence of the matrix can be bounded by a constant.

$$\mathbb{E}\left[\min_{\boldsymbol{z}\in\Omega}\|\boldsymbol{\Psi}\boldsymbol{z}\|_2\right] \geq \lambda_n - w(\Omega),$$

where $\boldsymbol{\Psi}$ *is a random map with entries zero mean and unit variance Gaussian. Then, solving convex program (2.8) recovers* $\boldsymbol{X}^\star$ *exactly with probability* $1 - e^{-\frac{1}{2}(\lambda_n - w(\Omega))^2}$*, as long as*

$$m \geq w(\Omega)^2 + 1, \tag{2.13}$$

where $w(\Omega) := \mathbb{E}_{\boldsymbol{g}}\left[\sup_{\boldsymbol{z}\in\Omega}\boldsymbol{g}^T\boldsymbol{z}\right]$ *is defined as the Gaussian width of the set* Ω[5]*.*

We give two examples where Gaussian widths can be explicitly upper bounded.

- Compressive sensing: $w(\mathcal{T}_{\mathcal{A}}(\boldsymbol{x}^\star)\cap\mathbb{S}^{N-1})^2 \leq 2K\log(\frac{N}{K}) + \frac{5}{4}K$;
- Matrix sensing: $w(\mathcal{T}_{\mathcal{A}}(\boldsymbol{X}^\star)\cap\mathbb{S}^{M\times N-1})^2 \leq 3r(M+N-r)$.

This implies that the measurement bound (2.13) in Theorem 2.4 is tight for compressive sensing and matrix sensing with a random Gaussian measurement scheme.

Although Theorem 2.4 gives a unified result on the number of random Gaussian linear measurements for exact recovery, it is not trivial to compute the Gaussian width $w(\Omega)$ given a generic atomic set $\mathcal{A}$ and the corresponding $\boldsymbol{X}^\star$. Furthermore, in many applications, the random linear map $\boldsymbol{\Phi}$ is not Gaussian and can be structured due to different sensing modalities. Therefore, new theorems and methodologies are needed to quantify the number of linear measurements when the atomic set $\mathcal{A}$ and the sensing operator $\boldsymbol{\Phi}$ are uniquely specified by particular applications.

[5]In this case, $\boldsymbol{g}\in\mathbb{R}^N$ is a random vector with entries obeying i.i.d. normal distributions.

CHAPTER 3

SUPER-RESOLUTION OF COMPLEX EXPONENTIALS FROM MODULATIONS WITH UNKNOWN WAVEFORMS

In this chapter[6], we introduce the non-stationary blind super-resolution problem that arises in physical and biological sciences. We show that this ill-posed inverse problem can be formulated as the recovery of a structured low-rank matrix from compressive measurements. Atomic norm minimization is then used to enforce the structured low-rankness, and is reformulated as a semidefinite program that is solvable in polynomial time. A theoretical analysis is carried out and we derive a near optimal sample complexity bound that is proportional to the number of degrees of freedom in the problem, up to a polylogarithmic factor. Numerical simulations support our theoretical findings, showing that non-stationary blind super-resolution using atomic norm minimization is possible.

3.1 Introduction

Super-resolution refers to techniques for enhancing the resolution of imaging systems. It finds applications in a variety of practical problems, including single-molecule microscopy, computational photography, astronomy, radar imaging. For example, in single-molecule imaging [20, 21], one is interested in studying the individual behavior of molecules from measurements of an ensemble of molecules. The measurements, however, only contain the average characteristics of the molecules with fine details smeared out by the point spread function of the imaging process. Super-resolution aims to recover these fine details by localizing individual molecules, and consequently, enhance the performance of the imaging system.

[6]This work is in collaboration with Prof. Mike Wakin and Prof. Gongguo Tang [18, 19].

3.1.1 Motivation

In this section, we consider super-resolution of unknown complex exponentials from their modulations with unknown waveforms. This extends super-resolution to the *blind* and *non-stationary*[7] scenario. More specifically, consider the observation model

$$\boldsymbol{y}(n) = \sum_{j=1}^{J} c_j e^{-i2\pi n\tau_j} \boldsymbol{g}_j(n), \tag{3.1}$$

where $\{\boldsymbol{y}(n) \in \mathbb{C}\}$ are samples of a continuous-time output, $\{c_j\} \subset \mathbb{C}$ and $\{\tau_j\} \subset [0, 1)$ are unknown coefficients and parameters associated with the complex exponentials, and $\{\boldsymbol{g}_j(n)\}$ are samples of unknown waveforms, whose forms vary with the index j. Our goal is to recover $\{\tau_j\}$ and $\{c_j\}$, as well as the samples of the unknown waveforms $\{\boldsymbol{g}_j(n)\}$. To make this otherwise ill-posed problem well-posed, we assume that the unknown waveforms $\{\boldsymbol{g}_j\}$ belong to a common and known low-dimensional subspace.

Model (5.1) encompasses a wide spectrum of applications. Here we list three stylized examples that can be modeled using our general mathematical framework.

Super-resolution with unknown point spread functions: In applications such as single-molecule microscopy, one is interested in super-resolving and localizing unknown point sources from their convolutions with point spread functions. Quite often, the point spread function, however, cannot be perfectly known. The point spread function may also depend on the locations of the point sources. This is the case in 3D single-molecule microscopy [23], where the point spread function depends on the depth (z-axis) of the target, demanding a super-resolution technique that handles unknown and space-varying system functions. Another example is the non-stationary blind deconvolution of seismic data [22]. Here the goal is to retrieve the time domain reflectivity of the earth from its convolution with (non-stationarily) attenuated seismic waves from samples of the seismic trace. Yet other non-stationary blind super-resolution applications include computational photography [24] and astronomy [25]. Finally, one further application involving simultaneous super-resolution and

[7]Our choice of the term "non-stationary" is inspired by its use in non-stationary deconvolution [22].

calibration of unknown waveforms is the blind multi-path channel identification problem in multi-user communication systems [26]. At the receiver, one must estimate the multi-path delays of unknown waveforms set by different users. For all of these applications, the goal is to determine the unknown delays $\{\tau_j\}$ and coefficients $\{c_j\}$ from observations of the form

$$y(t) = \sum_{j=1}^{J} c_j g_j(t - \tau_j) \tag{3.2}$$

with $\{g_j(t)\}$ being the unknown point spread functions. By taking the Fourier transform on both sides of (3.2), we obtain

$$\widehat{y}(f) = \sum_{j=1}^{J} c_j e^{-i2\pi f \tau_j} \widehat{g}_j(f), \tag{3.3}$$

which takes the form of (5.1) when sampled. The goal is to simultaneously recover $\{c_j\}, \{\tau_j\}$ and samples of the point spread functions $\{\widehat{g}_j(f)\}$.

Parameter estimation in radar imaging: In radar imaging [27], one is concerned with estimating the distances and velocities of the targets relative to the radar. These quantities can be inferred by estimating the unknown delay-Doppler parameters (μ_j, ν_j) from the following signal model:

$$y(t) = \sum_{j=1}^{J} c_j e^{i2\pi\nu_j t} x(t - \mu_j), \tag{3.4}$$

where both the transmitted waveform $x(t)$ and the received waveform $y(t)$ are known. We note that ν_j and μ_j can be arbitrary and do not necessarily lie on a grid. It is easy to see that sampling (3.4) also produces (5.1).

Frequency estimation with damping: In applications such as nuclear magnetic resonance spectroscopy [28], the signal is a superposition of complex exponentials with unknown frequencies $\{f_j\}$ and damping factors $\{\varsigma_j\}$:

$$y(t) = \sum_{j=1}^{J} c_j e^{i2\pi f_j t} e^{-\varsigma_j t}. \tag{3.5}$$

By sampling the continuous variable t in (3.5), we again obtain an instance of (5.1). Here the modulating waveforms $\boldsymbol{g}_j(n)$ are samples of the damping terms $e^{-\varsigma_j t}$.

In some cases, to help regularize the inverse problems above, it may be appropriate to assume, as we do, that the unknown waveforms $\{\boldsymbol{g}_j\}$ belong to a known low-dimensional subspace. In super-resolution imaging, for example, point spread functions can often be modeled as Gaussians; see [20, 23] and references therein. When the widths of the point spread functions are unknown, however, a dictionary can be constructed consisting of Gaussian functions with different variances. Applying principal component analysis (PCA) on the constructed dictionary reveals an approximate low-dimensional subspace structure that captures the unknown point spread functions. We demonstrate this in numerical experiments in Section 3.4. Further, in multi-user communication systems it may be reasonable to assume that the unknown waveforms transmitted by different users belong to a subspace; in addition, the multiple received copies of a single user's waveform will all be identical (save for a delay, which becomes part of the modulation term in (5.1)). On the other hand, in radar imaging, the subspace spanned by sampled, shifted copies of the transmitted waveform may not always have a low dimension. Related works such as [27] may give sharper guarantees in this case.

3.1.2 Related work

In the past few years, super-resolution via convex programming has become a popular approach since convex methods usually come with strong theoretical guarantees and robustness to noise and outliers. In [29], a general mathematical framework for super-resolution using total variation (TV) norm minimization is proposed. The goal there is to super-resolve the unknown locations in $[0, 1)$ of point sources from low-frequency samples of the spectrum. This TV norm minimization problem can be recast as a computationally efficient semidefinite program (SDP) [30]. It is shown that one can super-resolve J point sources from $O(J)$ samples under a minimum separation condition. We note that this approach, however, requires perfect information of the point spread function. Based on [29], [31, 32] study the

robustness of TV norm minimization for super-resolution by considering the noisy data case; [33] extends the super-resolution problem to the case when the point sources are positive; [34] examines the recovery property of sparse spikes using TV norm minimization through studying the non-degeneracy of the dual certificate. Recent work [35] studies the super-resolution problem without separation. In [36], the author considers super-resolution for demixing and super-resolution of multiple signals with a common support. In [12], driven by applications in line spectral estimation, an atomic norm minimization scheme is proposed for super-resolution of arbitrary unknown frequencies from random time samples of a superposition of complex exponentials. It has been shown that the sample complexity is proportional to the number of frequencies (up to a polylogarithmic factor) for exact frequency estimation. It is also worth mentioning subsequent work based on [12, 29]. In [37, 38], the authors study the problem of frequency estimation when multiple measurement vectors (MMV) are available. [39] proposes an enhanced matrix completion algorithm for frequency estimation from limited time samples by converting spectral sparsity in the model into a low-rank structure of the block Hankel matrix.

Another line of related work addresses the blind deconvolution problem. In [40], the bilinear blind deconvolution problem is reformulated as a rank-one matrix sensing problem. A nuclear norm minimization program is then utilized for rank-one matrix recovery. It is shown that by employing subspace models for both signals, one can recover two length-L vectors from their circular convolution when $L = O(Q + K)$, where Q and K are the dimensions of the two subspaces. Following this general idea of lifting for blind deconvolution using convex programming, [41] considers the problem of blind deconvolution when multiple unknown inputs belong to a known and diverse subspace; [42] extends the work in [40] from rank-one case to a general rank-r matrix sensing problem, achieving simultaneous blind deconvolution and demixing. In [43], the authors propose an alternating minimization scheme for blind deconvolution under a sparsity model for the underlying signals. In [44], the authors propose a nuclear norm minimization algorithm for blind deconvolution using

random masks. More recently, [45] generalizes the problem studied in [44] by considering the effect of subsampling in the measurement process. Other related works along this line include [46, 47], which study conditions for the uniqueness of blind deconvolution.

This work is most closely related to the recent works [48, 49]. In [49], a biconvex problem for simultaneous sparse recovery and unknown gain calibration is studied. In their work, a subspace model is employed for the unknown gains to make the problem well-posed. It is worth mentioning that they use ℓ_1 minimization as a convex program, which is different from ours. Then, a sample complexity bound that is suboptimal is derived for sparse recovery and self-calibration. Inspired by [49], [48] considers a super-resolution problem that has a similar setup to [29], except that the point spread function is assumed unknown. By employing a subspace model for the point spread function, an atomic norm minimization program is formulated for simultaneous super-resolution of point sources and recovery of the unknown point spread function. The atomic norm minimization problem therein is recast as an SDP. The sample complexity bound derived there, however, is suboptimal. As we explain in Section 3.1.3, this work further generalizes the model in [48] to the non-stationary case, where the point spread functions can vary with the point sources.

In [27], super-resolution radar is formulated as a convex optimization program. In particular, the signal is modeled as a superposition of delayed and Doppler shifted versions of the template waveform, which is the same as model (3.4). It should be pointed out that our model (5.1) can be utilized for this problem as well. Therefore, the proposed blind super-resolution method can conceivably be used for super-resolution radar.

Lastly, we would like to mention that the signal model in our work has both low-rank and spectrally sparse structures, and thus is simultaneously structured. Consistent with [10], we can achieve the information-theoretic limit on the measurement bound (up to a polylogarithmic factor) not by a combination of convex objectives but rather through a single convex objective—in this case via atomic norm minimization.

3.1.3 Main contributions

Our contributions are twofold. First, we propose a general model for non-stationary blind super-resolution, which arises in a variety of disciplines. Our non-stationary blind super-resolution problem is naturally non-convex. By utilizing a subspace model for the unknown waveforms and a lifting trick [40, 50], we relax the non-stationary blind super-resolution problem using atomic norm minimization, which can be further formulated as an SDP. Second, we derive a sample complexity bound that is near information-theoretically optimal under assumptions on the minimum separation of the τ_j's and on the randomness and incoherence properties of the subspace. Specifically, assuming that the subspace has dimension K, we show that when the number of measurements is proportional to the number of degrees of freedom in the problem, i.e., $O(JK)$ (up to a polylogarithmic factor), the non-stationary blind super-resolution problem is solvable by an SDP. Furthermore, we can faithfully recover $\{\tau_j\}$, $\{c_j\}$, and the samples of the unknown waveforms $\{\boldsymbol{g}_j(n)\}$.

It is also worth mentioning the recent work [48], which can be viewed as a special case of our general non-stationary blind super-resolution framework by assuming all of the unknown waveforms are the same. Our model is more realistic and powerful due to its generality. As illustrated by the examples in the introduction, our framework also covers a wider range of non-convex inverse problems beyond the super-resolution problem with unknown point spread functions, including blind multi-path channel identification in communication systems, parameter estimation in radar imaging and frequency estimation with damping. Additionally, on the theoretical side, we improve the sample complexity bound in [48] from $O(J^2K^2)$ to $O(JK)$, up to a polylogarithmic factor. We elaborate on comparisons with [48] in Section 3.3.2.

3.2 Problem Formulation

In this section, we show that the non-stationary blind super-resolution problem can be solved efficiently by a convex program. In Section 3.2.1, we show that this nonconvex

problem can be recast as an atomic norm minimization program via a lifting trick. In Section 3.2.2, we demonstrate that the atomic norm minimization program can be solved via a computationally feasible SDP.

3.2.1 Problem setup via an atomic norm minimization

Consider the model

$$\boldsymbol{y}(n) = \sum_{j=1}^{J} c_j e^{-i2\pi n\tau_j} \boldsymbol{g}_j(n), \quad n = -2M, \ldots, 2M, \tag{3.6}$$

where $\boldsymbol{y}(n) \in \mathbb{C}$ are observations, (c_j, τ_j) are unknown parameters of complex exponentials $c_j e^{-i2\pi n\tau_j}$, and $\boldsymbol{g}_j(n)$ are samples of unknown waveforms. Without loss of generality, we assume that $\tau_j \in [0, 1), j = 1, \ldots, J$. Our goal is to recover τ_j, c_j, and $\boldsymbol{g}_j(n)$ from the samples $\boldsymbol{y}(n)$. It is apparent that one can only recover c_j and $\boldsymbol{g}_j(n)$ up to a scaling factor due to the multiplicative form in (3.6).

Unfortunately, this problem is severely ill-posed without any additional constraints on $\boldsymbol{g}_j$ since the number of samples in (3.6) is $N := 4M + 1$, while the number of unknowns in (3.6) is $O(JN)$, which is larger than N. To alleviate this, we solve our problem under the assumption that all $\boldsymbol{g}_j$ live in a common low-dimensional subspace spanned by the columns of a known $N \times K$ matrix $\boldsymbol{B}$ with $K \leq N$, which we denote as

$$\boldsymbol{B} = \begin{bmatrix} \boldsymbol{b}_{-2M} & \boldsymbol{b}_{-2M+1} & \cdots & \boldsymbol{b}_{2M-1} & \boldsymbol{b}_{2M} \end{bmatrix}^H$$

with $\boldsymbol{b}_n \in \mathbb{C}^{K\times 1}$. In other words, $\boldsymbol{g}_j = \boldsymbol{B}\boldsymbol{h}_j$ for some unknown $\boldsymbol{h}_j \in \mathbb{C}^{K\times 1}$. Henceforth, we assume that $\|\boldsymbol{h}_j\|_2 = 1$ without loss of generality. Under the subspace assumption, recovery of $\boldsymbol{g}_j$ is guaranteed if $\boldsymbol{h}_j$ can be recovered. Therefore, the number of degrees of freedom in (3.6) becomes $O(JK)$, which can possibly be smaller than the number of samples N when $J, K \ll N$.

Under the subspace assumption, we can rewrite (3.6) as

$$\boldsymbol{y}(n) = \sum_{j=1}^{J} c_j e^{-i2\pi n\tau_j} \boldsymbol{b}_n^H \boldsymbol{h}_j. \tag{3.7}$$

Defining

$$\boldsymbol{a}(\tau) = \begin{bmatrix} e^{i2\pi(-2M)\tau} & \cdots & e^{i2\pi(0)\tau} & \cdots & e^{i2\pi(2M)\tau} \end{bmatrix}^T,$$

we have

$$\begin{aligned} \boldsymbol{y}(n) &= \sum_{j=1}^{J} c_j \boldsymbol{a}(\tau_j)^H \boldsymbol{e}_n \boldsymbol{b}_n^H \boldsymbol{h}_j \\ &= \text{trace}\left(\boldsymbol{e}_n \boldsymbol{b}_n^H \sum_{j=1}^{J} c_j \boldsymbol{h}_j \boldsymbol{a}(\tau_j)^H \right) \\ &= \left\langle \sum_{j=1}^{J} c_j \boldsymbol{h}_j \boldsymbol{a}(\tau_j)^H, \boldsymbol{b}_n \boldsymbol{e}_n^H \right\rangle, \end{aligned} \tag{3.8}$$

where we have defined $\langle \boldsymbol{X}, \boldsymbol{Y} \rangle = \text{trace}(\boldsymbol{Y}^H \boldsymbol{X})$ and used $\boldsymbol{e}_n$, $-2M \leq n \leq 2M$, to denote the $(n + 2M + 1)$th column of the $N \times N$ identity matrix $\boldsymbol{I}_N$. We see that (3.8) leads to a parametrized rank-J matrix sensing problem, which we write as

$$\boldsymbol{y} = \mathcal{B}(\boldsymbol{X}_o),$$

where the linear operator $\mathcal{B} : \mathbb{C}^{K \times N} \to \mathbb{C}^N$ is defined as $[\mathcal{B}(\boldsymbol{X}_o)]_n = \left\langle \boldsymbol{X}_o, \boldsymbol{b}_n \boldsymbol{e}_n^H \right\rangle, n = -2M, \cdots, 2M$ with $\boldsymbol{X}_o = \sum_{j=1}^{J} c_j \boldsymbol{h}_j \boldsymbol{a}(\tau_j)^H$. Here we choose the number of measurements $N = 4M + 1$, which is purely for ease of theoretical analysis. We note that our result is not restricted to the symmetric case presented here and does not necessarily require that N should be an odd number. We refer the interested reader to Appendix A of [12] for a discussion of how to modify the argument for the general case.

In many scenarios, the number of complex exponentials J is small. Therefore, we use the atomic norm to promote sparsity. As in [48], define the atomic norm [11] associated with the following set of atoms

$$\mathcal{A} = \left\{ \boldsymbol{h}\boldsymbol{a}(\tau)^H : \ \tau \in [0, 1), \|\boldsymbol{h}\|_2 = 1, \boldsymbol{h} \in \mathbb{C}^{K \times 1} \right\}$$

as

$$\begin{aligned}\|\boldsymbol{X}\|_{\mathcal{A}} &= \inf\{t>0:\ \boldsymbol{X}\in t\mathbf{conv}(\mathcal{A})\}\\ &= \inf_{c_k,\tau_k,\|\boldsymbol{h}_k\|_2=1}\left\{\sum_k |c_k|:\ \boldsymbol{X}=\sum_k c_k\boldsymbol{h}_k\boldsymbol{a}(\tau_k)^H\right\}.\end{aligned}$$

To enforce the sparsity of the atomic representation, we solve

$$\begin{aligned}&\underset{\boldsymbol{X}}{\text{minimize}} \quad \|\boldsymbol{X}\|_{\mathcal{A}}\\ &\text{subject to} \quad \boldsymbol{y}(n)=\langle \boldsymbol{X}, \boldsymbol{b}_n\boldsymbol{e}_n^H\rangle,\quad n=-2M,\cdots,2M.\end{aligned} \tag{3.9}$$

Standard Lagrangian analysis shows that the dual of (3.9) is given by

$$\underset{\boldsymbol{\lambda}}{\text{maximize}} \quad \langle \boldsymbol{\lambda}, \boldsymbol{y}\rangle_{\mathbb{R}} \qquad \text{subject to} \quad \|\mathcal{B}^*(\boldsymbol{\lambda})\|_{\mathcal{A}}^* \le 1 \tag{3.10}$$

where $\langle \boldsymbol{\lambda}, \boldsymbol{y}\rangle_{\mathbb{R}} = \mathrm{Re}\left(\langle \boldsymbol{\lambda}, \boldsymbol{y}\rangle\right)$, $\mathcal{B}^*:\mathbb{C}^N\to\mathbb{C}^{K\times N}$ denotes the adjoint operator of $\mathcal{B}$ and $\mathcal{B}^*(\boldsymbol{\lambda})=\sum_n \boldsymbol{\lambda}(n)\boldsymbol{b}_n\boldsymbol{e}_n^H$, and $\|\cdot\|_{\mathcal{A}}^*$ is the dual norm of the atomic norm.

The following proposition characterizes the optimality condition of program (3.9) with the vector polynomial $\boldsymbol{q}(\tau)$ serving as a dual certificate to certify the optimality of $\boldsymbol{X}_o$ in the primal problem (3.9).

Proposition 3.1. *Suppose that the atomic set $\mathcal{A}$ is composed of atoms of the form $\boldsymbol{h}\boldsymbol{a}(\tau)^H$ with $\|\boldsymbol{h}\|_2=1, \tau\in[0,1)$. Define the set $\mathbb{D}=\{\tau_j, 1\le j\le J\}$. Let $\widehat{\boldsymbol{X}}$ be the optimal solution to (3.9). Then $\widehat{\boldsymbol{X}}=\boldsymbol{X}_o$ is the unique optimal solution if the following two conditions are satisfied:*

1) There exists a dual polynomial

$$\begin{aligned}\boldsymbol{q}(\tau) &= \mathcal{B}^*(\boldsymbol{\lambda})\boldsymbol{a}(\tau)\\ &= \sum_{n=-2M}^{2M}\boldsymbol{\lambda}(n)e^{i2\pi n\tau}\boldsymbol{b}_n\end{aligned} \tag{3.11}$$

satisfying

$$\boldsymbol{q}(\tau_j)=\mathrm{sign}(c_j)\boldsymbol{h}_j,\quad \forall\ \tau_j\in\mathbb{D}\ \ \textit{(interpolation condition)} \tag{3.12}$$

$$\|\boldsymbol{q}(\tau)\|_2<1,\quad \forall\ \tau\notin\mathbb{D}\ \ \textit{(boundedness condition).} \tag{3.13}$$

Here $\boldsymbol{\lambda}$ *is a dual optimizer and* $\operatorname{sign}(c_j) := \frac{c_j}{|c_j|}$.

2) $\left\{ \begin{bmatrix} \vdots \\ \boldsymbol{a}(\tau_j)^H \boldsymbol{e}_n \boldsymbol{b}_n^H \\ \vdots \end{bmatrix}, j = 1, \cdots, J \right\}$ *is a linearly independent set.*

The proof of Proposition 3.1 can be found in Appendix A.

3.2.2 SDP characterization

Since the convex hull of the set of atoms $\mathcal{A}$ can be characterized by a semidefinite program, $\|\boldsymbol{X}\|_{\mathcal{A}}$ admits an equivalent SDP representation.

Lemma 3.1. *[37, 38] For any* $\boldsymbol{X} \in \mathbb{C}^{K \times N}$,

$$\|\boldsymbol{X}\|_{\mathcal{A}} = \inf_{\boldsymbol{u}, \boldsymbol{T}} \left\{ \frac{1}{2N} \operatorname{trace}\left(\operatorname{Toep}(\boldsymbol{u})\right) + \frac{1}{2} \operatorname{trace}(\boldsymbol{T}) : \begin{bmatrix} \operatorname{Toep}(\boldsymbol{u}) & \boldsymbol{X}^H \\ \boldsymbol{X} & \boldsymbol{T} \end{bmatrix} \succeq \boldsymbol{0} \right\},$$

where $\boldsymbol{u}$ *is a complex vector whose first entry is real,* $\operatorname{Toep}(\boldsymbol{u})$ *denotes the* $N \times N$ *Hermitian Toeplitz matrix whose first column is* $\boldsymbol{u}$*, and* $\boldsymbol{T}$ *is a Hermitian* $K \times K$ *matrix.*

Hence, (3.9) can be solved efficiently using off-the-shelf SDP solvers such as CVX [51]. With many SDP solvers, one can also obtain a dual optimal solution $\boldsymbol{\lambda}$ to (3.10) for free by solving the primal program (3.9), and as we discuss below, this can be used to localize the supports of the point sources. We note that the dual optimal solution is not unique in general. As discussed in [12], the recovered support set from the dual solution must contain the true support set when the optimal primal solution is $\boldsymbol{X}_o$. Though it is possible that the recovered support set contains spurious parameters, solving the SDP with the interior point method will avoid this pathological situation and recover the true support exactly. See [12] for more technical discussions on this.

Now, given the dual optimal solution $\boldsymbol{\lambda}$, consider the trigonometric polynomial:

$$\begin{aligned} p(e^{i2\pi\tau}) &= 1 - \|\boldsymbol{q}(\tau)\|_2^2 \\ &= 1 - \boldsymbol{q}(\tau)^H \boldsymbol{q}(\tau) \\ &= 1 - \sum_{n=-4M}^{4M} u_n e^{i2\pi n\tau}, \end{aligned}$$

where $\boldsymbol{q}(\tau)$ is defined in terms of $\boldsymbol{\lambda}$ in (3.11), and u_n are some scalars that can be computed from $\boldsymbol{q}(\tau)$ explicitly. To localize the supports of the point sources, one can simply compute the roots of the polynomial $p(z)$ on the unit circle. This method allows for the recovery of point sources to very high precision as shown in [29].

Another way to recover the support is by discretizing $\tau \in [0, 1)$ on a fine grid up to a desired accuracy. Then, one can check the ℓ_2 norm of the dual polynomial $\boldsymbol{q}(\tau)$ and identify the $\{\tau_j\}$ by selecting the values of τ such that $\|\boldsymbol{q}(\tau)\|_2 = \|\sum_{n=-2M}^{2M} \boldsymbol{\lambda}(n) e^{i2\pi n\tau} \boldsymbol{b}_n\| \approx 1$. We use this heuristic in our numerical simulations.

Given an estimate of $\{\tau_j\}$, say $\{\widehat{\tau}_j\}$, one can plug these values back into (3.8) to form the following overdetermined linear system of equations:

$$\begin{bmatrix} \boldsymbol{a}(\widehat{\tau}_1)^H \boldsymbol{e}_{-2M} \boldsymbol{b}_{-2M}^H & \cdots & \boldsymbol{a}(\widehat{\tau}_J)^H \boldsymbol{e}_{-2M} \boldsymbol{b}_{-2M}^H \\ \vdots & \ddots & \vdots \\ \boldsymbol{a}(\widehat{\tau}_1)^H \boldsymbol{e}_{2M} \boldsymbol{b}_{2M}^H & \cdots & \boldsymbol{a}(\widehat{\tau}_J)^H \boldsymbol{e}_{2M} \boldsymbol{b}_{2M}^H \end{bmatrix} \begin{bmatrix} c_1 \boldsymbol{h}_1 \\ \vdots \\ c_J \boldsymbol{h}_J \end{bmatrix} = \begin{bmatrix} \boldsymbol{y}(-2M) \\ \vdots \\ \boldsymbol{y}(2M) \end{bmatrix}.$$

A unique solution for $\{c_j \boldsymbol{h}_j\}_{j=1}^J$ can be obtained by a simple least squares since the columns of the above matrix are linearly independent. However, we note that one cannot resolve the inherent scaling ambiguity between each c_j and the corresponding $\boldsymbol{h}_j$.

3.3 Main Results

This section presents the main result of this work. In particular, we show that nonstationary blind super-resolution is possible with the number of measurements proportional to the degree of freedom in the problem, up to a polylogarithmic factor.

3.3.1 Sample complexity bound for exact recovery

Given samples of the form $\boldsymbol{y}(n) = \langle \boldsymbol{X}_o, \boldsymbol{b}_n \boldsymbol{e}_n^H \rangle, n = -2M, \ldots, 2M$, we wish to quantify precisely how large the number of measurements $N = 4M + 1$ should be so that the atomic norm minimization (3.9) exactly recovers $\boldsymbol{X}_o$.

Before presenting the main result of this section, we discuss the assumptions that are used in the main theorem. The assumptions can be grouped into three categories: (a) randomness

and incoherence of the subspace spanned by the columns of $\boldsymbol{B}$, (b) minimum separation of the τ_j, and (c) uniform distribution of $\boldsymbol{h}_j$ on the complex unit sphere $\mathbb{CS}^{K-1}$.

We assume that the columns of the matrix $\boldsymbol{B}^H$, namely, $\boldsymbol{b}_n, -2M \leq n \leq 2M$, are independently sampled from a population $\mathcal{F}$ with the following properties [52]:

- **Isotropy property:** We assume that the distribution $\mathcal{F}$ obeys the isotropy property in that

$$\mathbb{E}\boldsymbol{b}\boldsymbol{b}^H = \boldsymbol{I}_K, \quad \boldsymbol{b} \sim \mathcal{F}. \tag{3.14}$$

- **Incoherence property:** We assume that $\mathcal{F}$ satisfies the incoherence property with coherence μ in that

$$\max_{1 \leq p \leq K} |\boldsymbol{b}(p)|^2 \leq \mu, \quad \boldsymbol{b} \in \mathcal{F} \tag{3.15}$$

  holds almost surely, where $\boldsymbol{b}(p)$ is the pth element of $\boldsymbol{b}$. For ease of theoretical analysis, hereafter we also assume that $\mu K \geq 1$, which can always be ensured by choosing μ sufficiently large. In particular, when the incoherence property (3.15) holds almost surely, the isotropy property (3.14) ensures that $\mu \geq 1$ and thus that $\mu K \geq 1$.

Furthermore, we require the following conditions on the parameters of the complex exponentials and the rotations of $\boldsymbol{g}_j$ in the subspace $\boldsymbol{B}$, namely, $\boldsymbol{h}_j$.

- **Minimum separation:** We assume that

$$\Delta_\tau = \min_{k \neq j} |\tau_k - \tau_j| \geq \frac{1}{M}$$

  where the distance $|\tau_k - \tau_j|$ is understood as the wrap-around distance on $[0, 1)$.

- **Random rotation:** We assume that the coefficient vectors $\boldsymbol{h}_j$ are drawn i.i.d. from the uniform distribution on the complex unit sphere $\mathbb{CS}^{K-1}$.

Theorem 3.5. *Assume that the minimum separation condition $\Delta_\tau \geq \frac{1}{M}$ is satisfied and that $M \geq 64$. Also, assume that $\boldsymbol{g}_j = \boldsymbol{B}\boldsymbol{h}_j$ with the columns of $\boldsymbol{B}^H$, namely, $\boldsymbol{b}_n$, being*

i.i.d. samples from a distribution $\mathcal{F}$ that satisfies the isotropy and incoherence properties with coherence parameter μ. Additionally, assume that $\boldsymbol{h}_j$ are drawn i.i.d. from the uniform distribution on the complex unit sphere $\mathbb{CS}^{K-1}$. Then, there exists a numerical constant C such that

$$M \geq C\mu JK \log\left(\frac{MJK}{\delta}\right) \log^2\left(\frac{MK}{\delta}\right) \tag{3.16}$$

is sufficient to guarantee that we can recover $\boldsymbol{X}_o$ via program (3.9) with probability at least $1-\delta$.

3.3.2 Discussion

Inspired by [48], we use the same assumptions on the random subspace model (the isotropy and incoherence properties) in order to prove our Theorem 3.5. The randomness assumption on the subspace does not appear to be critical in practice, as evidenced by our numerical experiments in Section 3.4; being able to replace this with a deterministic condition would increase the relevance of our theory to the applications discussed in Section 3.1.1.

Also, as noted in Section 3.1.3, our work generalizes the model in [48] to the non-stationary case. It may also be possible to extend the result developed in [48] to the non-stationary case; however, the sample complexity would still be $O(J^2K^2)$, up to a polylogarithmic factor. In contrast, we reduce the sample complexity to $O(JK)$, which is information theoretically optimal, up to a polylogarithmic factor. In order to do this, in the proof of Lemma 3.6 in Appendix 3.5, we apply a matrix Bernstein's inequality instead of Talagrand's concentration inequality which was used by [48]. Our theorem also relies on an additional assumption that was not present in [48], namely that the coefficient vectors $\boldsymbol{h}_j$ are drawn randomly. We do not believe that this randomness assumption is important in practice and suspect that it is merely an artifact of our proof.

Our bound on M suggests that when μ is a constant (e.g., when the rows of $\boldsymbol{B}$ are drawn from a sub-Gaussian distribution, μ can be bounded by a constant times $\log K$ with high

probability [52]), $M = O(JK)$ is sufficient for exact recovery and this matches the number of degrees of freedom in the problem, up to a polylogarithmic factor. Thus, our sample complexity bound is tight and there is little room for further improvement. When the dimension of the subspace is bounded by a constant, $M = O(J)$ (up to a polylogarithmic factor) is sufficient for exact recovery. This bound matches the one in the deterministic super-resolution framework [29], where $N = O(J)$ suffices to exactly localize the unknown spikes under the same minimum separation condition used here. We also see that our bound improves the one derived in [48] even when $\boldsymbol{g}_j = \boldsymbol{B}\boldsymbol{h}$, i.e., when $\boldsymbol{g}_j$ has no dependence on j. We note that the number of degrees of freedom in that problem is $O(J + K)$. It would be interesting to see if further improvement upon our bound is possible in the stationary scenario.

Finally, when the measurements are contaminated by noise, one can extend the observation model as

$$\boldsymbol{y} = \mathcal{B}(\boldsymbol{X}_o) + \boldsymbol{z},$$

where $\|\boldsymbol{z}\|_2 \leq \delta_{\text{noise}}$ and δ_{noise} is a parameter controlling the noise level. To recover an estimate of $\boldsymbol{X}_o$, we propose to solve the following inequality constrained program:

$$\underset{\boldsymbol{X}}{\text{minimize}} \ \ \|\boldsymbol{X}\|_{\mathcal{A}} \quad \text{subject to} \ \ \|\boldsymbol{y} - \mathcal{B}(\boldsymbol{X})\|_2 \leq \delta_{\text{noise}}. \tag{3.17}$$

We leave robust performance analysis of (3.17) for future research.

3.4 Numerical Simulations

In this section, we provide synthetic numerical simulations to support our theoretical findings. In all of the simulations, we solve the atomic norm minimization problem (3.9) using CVX. We start with a simple example to illustrate how to localize unknown spikes using the dual polynomial $\boldsymbol{q}(\tau)$ and recover samples of the unknown waveforms $\boldsymbol{g}_j(n)$. For the sake of illustration, we set $N = 64$, $J = 3$ and $K = 4$ and randomly generate the locations of J spikes on $[0, 1)$ from a uniform distribution. We regenerate the set of locations if a particular minimum separation condition is violated; in particular, we ensure that the

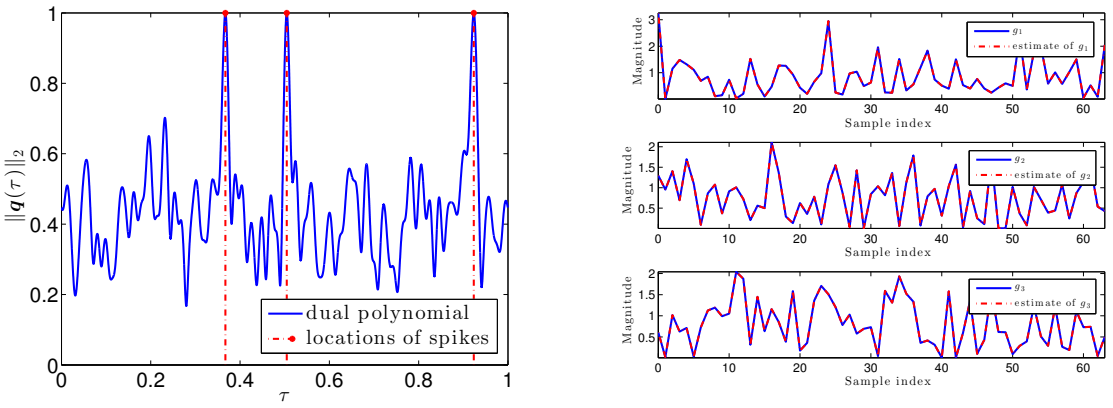

(a) (b)

Figure 3.1: (a) The ℓ_2-norm of the dual polynomial $\|\boldsymbol{q}(\tau)\|_2$ and the locations of the true spikes when the entries of $\boldsymbol{B}$ are built from the standard real Gaussian distribution. (b) The magnitude of samples of the waveforms $\boldsymbol{g}_1, \boldsymbol{g}_2, \boldsymbol{g}_3$ and their estimates $\widehat{\boldsymbol{g}}_1, \widehat{\boldsymbol{g}}_2, \widehat{\boldsymbol{g}}_3$ from least squares (best viewed in color).

minimum separation Δ_τ between spikes is not less than $\frac{1}{N}$, which is slightly smaller than $\frac{1}{M}$ required by our theorem. Each c_j is generated randomly with a dynamic range of 20dB and uniform phase. We build $\boldsymbol{B}$ with entries generated randomly from the standard Gaussian distribution. Each $\boldsymbol{h}_j$ is also generated using i.i.d. real Gaussian random variables and is then normalized. Figure 3.1(a) shows the ℓ_2-norm of the dual polynomial, namely, $\|\boldsymbol{q}(\tau)\|_2$ on a discretized set of points on [0,1] with discretization step size 10^{-4}. Once the τ_j's are identified, we solve a least squares problem using a pseudo-inverse to estimate the $\boldsymbol{h}_j$'s. Because we cannot exactly recover the $\boldsymbol{h}_j$'s, hence the $\boldsymbol{g}_j$'s, due to phase rotations, we plot the magnitude of the $\boldsymbol{g}_j$'s and the estimated $\widehat{\boldsymbol{g}}_j$'s in Figure 3.1(b). We also compute $|\langle \boldsymbol{h}_j, \widehat{\boldsymbol{h}}_j \rangle|, j = 1, 2, 3$, which equal $0.99999984, 0.99999973$ and 0.99999996, respectively. For $j = 1, 2, 3$, the agreements on the magnitude of $\boldsymbol{g}_j$ and $\widehat{\boldsymbol{g}}_j$ and on the absolute inner product between $\boldsymbol{h}_j$ and $\widehat{\boldsymbol{h}}_j$ confirm that we can faithfully recover $\boldsymbol{g}_j$.

Next, we characterize the phase transition of atomic norm minimization (3.9). We run 50 trials for each pair of the underlying changing variables. For each trial, we declare success if the relative reconstruction Frobenius norm error of X_o is less than 10^{-4}. In the

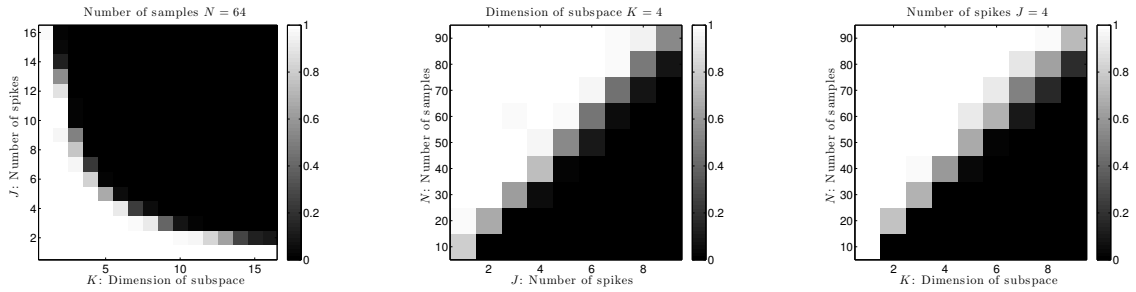

(a) (b) (c)

Figure 3.2: (a) The probability of success of non-stationary blind super-resolution using atomic norm minimization when $N = 64$ is fixed. (b) The phase transition of non-stationary blind super-resolution using atomic norm minimization when the dimension of subspace is fixed, $K = 4$. (c) The phase transition of non-stationary blind super-resolution using atomic norm minimization when the number of spikes is fixed, $J = 4$.

first experiment, we fix $N = 64$ and vary the values of J and K. Other specifications are the same as the warm-up experiment above. Figure 3.2(a) shows the phase transition in this situation. From Figure 3.2(a), we can roughly see that the success transition curve behaves like a hyperbola, which matches the bound appearing in Theorem 3.5. In the second experiment, we fix the dimension of the subspace $K = 4$ and study the phase transition between N and J. Figure 3.2(b) indicates a linear relationship between N and J when K is fixed. Finally, we test the phase transition between N and K when the number of spikes J is fixed. We set $J = 4$ and vary N and K in the experiment. Figure 3.2(c) shows the phase transition in this situation and also implies a linear relationship between N and K.

Finally, we test the atomic norm minimization (3.9) for localization of spikes in a more practical scenario where the matrix $\boldsymbol{B}$ is generated by extracting the principal components of a structured matrix, and thus has less randomness. To set up the experiment, we set $J = 3$ and generate the locations of $\{\tau_j\}$ uniformly at random between 0 and 1 under the minimum separation $\Delta_\tau = \frac{4}{N}$, which is roughly the same as what Theorem 3.5 requires. We assume that $\boldsymbol{g}_j(n)$ are samples of the zero mean Gaussian waveform $g_{\sigma^2}(t) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{t^2}{2\sigma^2}}$ with unknown variance $\sigma^2 \in [0.1, 1]$. We take $N = 100$ samples around the origin with sampling interval $\Delta t = 0.01$. To obtain the matrix $\boldsymbol{B}$, we first build a dictionary as follows:

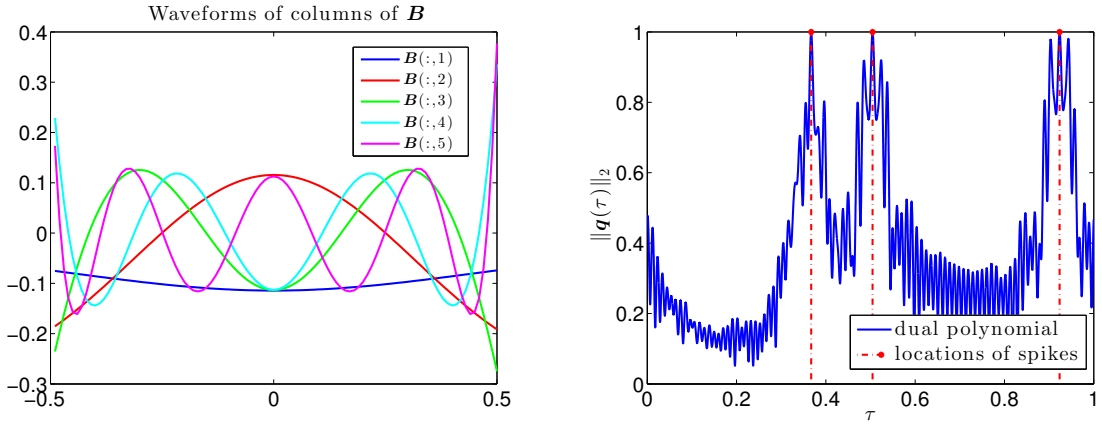

(a) (b)

Figure 3.3: (a) Plot of orthonormal columns of $\boldsymbol{B}$. (b) The ℓ_2-norm of the dual polynomial $\|\boldsymbol{q}(\tau)\|_2$ and the locations of the true spikes when $\boldsymbol{B}$ is built from the left singular vectors of the low-rank approximation of the dictionary $\boldsymbol{D_g}$ (best viewed in color).

$$\boldsymbol{D_g} = \left[\boldsymbol{g}_{\sigma^2=0.1} \quad \boldsymbol{g}_{\sigma^2=0.11} \quad \boldsymbol{g}_{\sigma^2=0.12} \quad \cdots \quad \boldsymbol{g}_{\sigma^2=1}\right],$$

where columns of $\boldsymbol{D_g}$ are from the samples of Gaussians with a discretized set of variances $\sigma^2 \in [0.1, 1]$. Then, we apply PCA on $\boldsymbol{D_g}$ and obtain the best rank-5 approximation, which we denote as $\boldsymbol{D}_{\boldsymbol{g},5}$. Finally, we construct $\boldsymbol{B}$ by taking its columns to be the left singular vectors of $\boldsymbol{D}_{\boldsymbol{g},5}$. Figure 3.3(a) depicts the waveforms of the columns of $\boldsymbol{B}$. We note that $\boldsymbol{D}_{\boldsymbol{g},5}$ gives a very good approximation to $\boldsymbol{D_g}$ due to the sharp decay of the singular values of $\boldsymbol{D_g}$. In particular, $\|\boldsymbol{D_g} - \boldsymbol{D}_{\boldsymbol{g},5}\|_F / \|\boldsymbol{D_g}\|_F = 1.9860 \times 10^{-8}$. We generate each $\boldsymbol{g}_j(n)$ by choosing its variance uniformly at random between 0.1 and 1, and then we build the measurements $\{\boldsymbol{y}(n)\}$. In particular, we note that $\{\boldsymbol{g}_j\}$ do not necessarily belong to the columns of $\boldsymbol{D_g}$. Finally, we solve (3.9) using CVX. Figure 3.3(b) shows the ℓ_2-norm of the dual polynomial, namely, $\|\boldsymbol{q}(\tau)\|_2$, on a discretized set of points on [0,1] with discretization step size 10^{-4}. We can see that one can still localize $\{\tau_j\}$ even when there is model mismatch between the subspace spanned by the columns of $\boldsymbol{B}$ and $\boldsymbol{g}_j(n)$.

3.5 Proof of the Main Theorem

In this part, we prove Theorem 3.5, the main result of Section 3.3. We divide the proof into three subsections. In Subsection 3.5.1, we construct the pre-certificate dual polynomial $\boldsymbol{q}(\tau)$. In Subsections 3.5.2 and 3.5.3, we show that our constructed dual polynomial $\boldsymbol{q}(\tau)$ satisfies conditions (3.12) and (3.13), respectively.

3.5.1 Construction of the dual polynomial

According to Proposition 3.1, our goal is to explicitly construct a dual polynomial $\boldsymbol{q}(\tau)$ such that conditions (3.12) and (3.13) hold. To proceed, we require that for all $\tau_j \in \mathbb{D}$

$$\boldsymbol{q}(\tau_j) = \operatorname{sign}(c_j)\boldsymbol{h}_j, \tag{3.18}$$

$$-\boldsymbol{q}'(\tau_j) = \mathbf{0}_{K\times 1}. \tag{3.19}$$

Note that (3.18) is exactly the same as (3.12), while (3.18) and (3.19) together form a necessary condition for (3.13) to hold. We construct the dual polynomial in (3.11) by finding a proper $\boldsymbol{\lambda}$ from solving the following weighted least energy minimization program with diagonal weighting matrix $\boldsymbol{W} = \operatorname{diag}\left(\begin{bmatrix} w_{-2M} & \cdots & w_{2M}\end{bmatrix}\right)$ to be determined later, $w_n > 0, -2M \le n \le 2M$:

$$\begin{aligned} \underset{\boldsymbol{\lambda}}{\text{minimize}} \quad & \|\boldsymbol{W}\boldsymbol{\lambda}\|_2^2 \\ \text{subject to} \quad & \boldsymbol{q}(\tau_j) = \operatorname{sign}(c_j)\boldsymbol{h}_j, \quad j = 1, \cdots, J, \\ & -\boldsymbol{q}'(\tau_j) = \mathbf{0}_{K\times 1}, \quad j = 1, \cdots, J. \end{aligned} \tag{3.20}$$

We can see that the equality constraints in (3.20) can be written as a linear system of equations (3.21).

$$\begin{bmatrix} \boldsymbol{b}_{-2M}\boldsymbol{e}_{-2M}^H\boldsymbol{a}(\tau_1) & \cdots & \boldsymbol{b}_{2M}\boldsymbol{e}_{2M}^H\boldsymbol{a}(\tau_1) \\ \vdots & \ddots & \vdots \\ \boldsymbol{b}_{-2M}\boldsymbol{e}_{-2M}^H\boldsymbol{a}(\tau_J) & \cdots & \boldsymbol{b}_{2M}\boldsymbol{e}_{2M}^H\boldsymbol{a}(\tau_J) \\ -i2\pi(-2M)\boldsymbol{b}_{-2M}\boldsymbol{e}_{-2M}^H\boldsymbol{a}(\tau_1) & \cdots & -i2\pi(2M)\boldsymbol{b}_{2M}\boldsymbol{e}_{2M}^H\boldsymbol{a}(\tau_1) \\ \vdots & \ddots & \vdots \\ -i2\pi(-2M)\boldsymbol{b}_{-2M}\boldsymbol{e}_{-2M}^H\boldsymbol{a}(\tau_J) & \cdots & -i2\pi(2M)\boldsymbol{b}_{2M}\boldsymbol{e}_{2M}^H\boldsymbol{a}(\tau_J) \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}(-2M) \\ \vdots \\ \boldsymbol{\lambda}(2M) \end{bmatrix} = \begin{bmatrix} \operatorname{sign}(c_1)\boldsymbol{h}_1 \\ \vdots \\ \operatorname{sign}(c_J)\boldsymbol{h}_J \\ \mathbf{0}_{K\times 1} \\ \vdots \\ \mathbf{0}_{K\times 1} \end{bmatrix}. \tag{3.21}$$

For notational simplicity, we denote equation (3.21) as

$$\boldsymbol{A}\boldsymbol{\lambda} = \boldsymbol{p}.$$

The optimality condition ensures that the optimal $\boldsymbol{\lambda}$ has the form

$$\begin{aligned}
\boldsymbol{\lambda} &= \left(\boldsymbol{W}^H\boldsymbol{W}\right)^{-1}\boldsymbol{A}^H\begin{bmatrix}\boldsymbol{\alpha}\\ \boldsymbol{\beta}\end{bmatrix}\\
&= \left(\boldsymbol{W}^H\boldsymbol{W}\right)^{-1}\left(\sum_{j=1}^{J}\begin{bmatrix}\boldsymbol{a}(\tau_j)^H\boldsymbol{e}_{-2M}\boldsymbol{b}_{-2M}^H\\ \vdots\\ \boldsymbol{a}(\tau_j)^H\boldsymbol{e}_{2M}\boldsymbol{b}_{2M}^H\end{bmatrix}\boldsymbol{\alpha}_j + \sum_{j=1}^{J}\begin{bmatrix}i2\pi(-2M)\boldsymbol{a}(\tau_j)^H\boldsymbol{e}_{-2M}\boldsymbol{b}_{-2M}^H\\ \vdots\\ i2\pi(2M)\boldsymbol{a}(\tau_j)^H\boldsymbol{e}_{2M}\boldsymbol{b}_{2M}^H\end{bmatrix}\boldsymbol{\beta}_j\right),
\end{aligned}$$

for some

$$\begin{bmatrix}\boldsymbol{\alpha}\\ \boldsymbol{\beta}\end{bmatrix} = \begin{bmatrix}\boldsymbol{\alpha}_1^H & \cdots & \boldsymbol{\alpha}_J^H & \boldsymbol{\beta}_1^H & \cdots & \boldsymbol{\beta}_J^H\end{bmatrix}^H, \boldsymbol{\alpha}_j, \boldsymbol{\beta}_j \in \mathbb{C}^{K\times 1}$$

satisfying the normal equation

$$\boldsymbol{A}\left(\boldsymbol{W}^H\boldsymbol{W}\right)^{-1}\boldsymbol{A}^H\begin{bmatrix}\boldsymbol{\alpha}\\ \boldsymbol{\beta}\end{bmatrix} = \boldsymbol{p}.$$

Inserting $\boldsymbol{\lambda}$ back into (3.11), we obtain the dual polynomial

$$\begin{aligned}
\boldsymbol{q}(\tau) &= \sum_{n=-2M}^{2M}\boldsymbol{\lambda}(n)e^{i2\pi n\tau}\boldsymbol{b}_n\\
&= \sum_{j=1}^{J}\left(\sum_{n=-2M}^{2M}\frac{1}{w_n^2}e^{-i2\pi n\tau_j}\boldsymbol{b}_n^H\boldsymbol{\alpha}_j e^{i2\pi n\tau}\boldsymbol{b}_n + \sum_{n=-2M}^{2M}\frac{i2\pi n}{w_n^2}e^{-i2\pi n\tau_j}\boldsymbol{b}_n^H\boldsymbol{\beta}_j e^{i2\pi n\tau}\boldsymbol{b}_n\right)\\
&= \sum_{j=1}^{J}\left(\sum_{n=-2M}^{2M}\frac{1}{w_n^2}e^{i2\pi n(\tau-\tau_j)}\boldsymbol{b}_n\boldsymbol{b}_n^H\boldsymbol{\alpha}_j + \sum_{n=-2M}^{2M}\frac{i2\pi n}{w_n^2}e^{i2\pi n(\tau-\tau_j)}\boldsymbol{b}_n\boldsymbol{b}_n^H\boldsymbol{\beta}_j\right)\\
&= \sum_{j=1}^{J}\boldsymbol{K}_M(\tau-\tau_j)\boldsymbol{\alpha}_j + \sum_{j=1}^{J}\boldsymbol{K}'_M(\tau-\tau_j)\boldsymbol{\beta}_j,
\end{aligned}$$

where

$$\boldsymbol{K}_M(\tau) = \sum_{n=-2M}^{2M}\frac{1}{w_n^2}e^{i2\pi n\tau}\boldsymbol{b}_n\boldsymbol{b}_n^H$$

and

$$\boldsymbol{K}'_M(\tau) = \sum_{n=-2M}^{2M} \frac{i2\pi n}{w_n^2} e^{i2\pi n\tau} \boldsymbol{b}_n \boldsymbol{b}_n^H,$$

which is the entry-wise derivative of $\boldsymbol{K}_M(\tau)$ with respect to τ. In the following, we choose $w_n = \sqrt{\frac{M}{g_M(n)}}$, where $g_M(n)$ is the convolution of two triangle functions given below:

$$g_M(n) = \frac{1}{M} \sum_{k=\max\{n-M,-M\}}^{\min\{n+M,M\}} \left(1 - \frac{|k|}{M}\right)\left(1 - \frac{|n-k|}{M}\right).$$

This particular choice of weights ensures that

$$\begin{aligned}\sum_{n=-2M}^{2M} \frac{1}{w_n^2} e^{i2\pi n\tau} &= \frac{1}{M} \sum_{n=-2M}^{2M} g_M(n) e^{i2\pi n\tau} \\ &= \left[\frac{\sin(\pi M\tau)}{M\sin(\pi\tau)}\right]^4 \\ &=: K_M(\tau),\end{aligned}$$

where $K_M(\tau)$ is known as the squared Fejér kernel. Consequently, $\boldsymbol{K}_M(\tau), \boldsymbol{K}'_M(\tau) \in \mathbb{C}^{K\times K}$ are the random matrix kernels produced from the squared Fejér kernel. They have the form

$$\boldsymbol{K}_M(\tau) = \frac{1}{M} \sum_{n=-2M}^{2M} g_M(n) e^{i2\pi n\tau} \boldsymbol{b}_n \boldsymbol{b}_n^H$$

and

$$\boldsymbol{K}'_M(\tau) = \frac{1}{M} \sum_{n=-2M}^{2M} (i2\pi n) g_M(n) e^{i2\pi n\tau} \boldsymbol{b}_n \boldsymbol{b}_n^H.$$

Denoting $\boldsymbol{K}^\ell_M(\tau)$ as the ℓth entry-wise derivative of $\boldsymbol{K}_M(\tau)$, we have

$$\begin{aligned}\mathbb{E}\boldsymbol{K}^\ell_M(\tau) &= \frac{1}{M} \sum_{n=-2M}^{2M} (i2\pi n)^\ell g_M(n) e^{i2\pi n\tau} \mathbb{E}\left[\boldsymbol{b}_n \boldsymbol{b}_n^H\right] \\ &= K^\ell_M(\tau) \boldsymbol{I}_K,\end{aligned}$$

where the second equation comes from the isotropy property of $\boldsymbol{b}_n$. We define

$$\begin{aligned}\overline{\boldsymbol{q}}^\ell(\tau) &= \sum_{j=1}^{J} \left[\mathbb{E}\boldsymbol{K}^\ell_M(\tau - \tau_j)\right] \overline{\boldsymbol{\alpha}}_j + \sum_{j=1}^{J} \left[\mathbb{E}\boldsymbol{K}^{\ell+1}_M(\tau - \tau_j)\right] \overline{\boldsymbol{\beta}}_j \\ &= \sum_{j=1}^{J} K^\ell_M(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j + \sum_{j=1}^{J} K^{\ell+1}_M(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j,\end{aligned} \tag{3.22}$$

where $\overline{\boldsymbol{\alpha}}_j$ and $\overline{\boldsymbol{\beta}}_j$ are solutions of the following linear system of equations:

$$\begin{aligned}\overline{\boldsymbol{q}}(\tau_s) &= \sum_{j=1}^{J} K_M(\tau_s-\tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{j=1}^{J} K'_M(\tau_s-\tau_j)\overline{\boldsymbol{\beta}}_j \\ &= \text{sign}(c_s)\boldsymbol{h}_s, \quad \tau_s \in \mathbb{D}, \\ -\overline{\boldsymbol{q}}'(\tau_s) &= -\left(\sum_{j=1}^{J} K'_M(\tau_s-\tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{j=1}^{J} K''_M(\tau_s-\tau_j)\overline{\boldsymbol{\beta}}_j\right) \\ &= \mathbf{0}_{K\times 1}, \quad \tau_s \in \mathbb{D}.\end{aligned}$$

We note that $\overline{\boldsymbol{q}}(\tau)$ is the vector dual polynomial constructed in [38] for the full data multiple measurement vectors problem. Thus, the vector dual polynomial $\boldsymbol{q}(\tau)$ constructed in our work is a random version of $\overline{\boldsymbol{q}}(\tau)$, with the randomness introduced by $\boldsymbol{b}_n$.

3.5.2 Showing the interpolation condition

To show (3.12), our effort becomes to explicitly find $\boldsymbol{\alpha}_j$ and $\boldsymbol{\beta}_j$ such that (3.12) holds. To accomplish this, for all $\tau_s \in \mathbb{D}$, we plug the form of $\boldsymbol{q}(\tau)$ back into (3.18) and (3.19). Thus, we can write

$$\begin{aligned}\boldsymbol{q}(\tau_s) &= \sum_{j=1}^{J} \boldsymbol{K}_M(\tau_s-\tau_j)\boldsymbol{\alpha}_j + \sum_{j=1}^{J} \boldsymbol{K}'_M(\tau_s-\tau_j)\boldsymbol{\beta}_j \\ &= \text{sign}(c_s)\boldsymbol{h}_s, \\ -\boldsymbol{q}'(\tau_s) &= -\left(\sum_{j=1}^{J} \boldsymbol{K}'_M(\tau_s-\tau_j)\boldsymbol{\alpha}_j + \sum_{j=1}^{J} \boldsymbol{K}''_M(\tau_s-\tau_j)\boldsymbol{\beta}_j\right) \\ &= \mathbf{0}_{K\times 1}.\end{aligned}$$

Equivalently, we have the following linear system of equations

$$\begin{bmatrix} \boldsymbol{D}_0 & \frac{1}{\sqrt{|K''_M(0)|}}\boldsymbol{D}_1 \\ -\frac{1}{\sqrt{|K''_M(0)|}}\boldsymbol{D}_1 & -\frac{1}{|K''_M(0)|}\boldsymbol{D}_2 \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \sqrt{|K''_M(0)|}\boldsymbol{\beta} \end{bmatrix} = \begin{bmatrix} \boldsymbol{h} \\ \mathbf{0}_{JK\times 1} \end{bmatrix},$$

where $\boldsymbol{h} = \begin{bmatrix} (\text{sign}(c_1)\boldsymbol{h}_1)^H & \dots & (\text{sign}(c_J)\boldsymbol{h}_J)^H \end{bmatrix}^H$, $[\boldsymbol{D}_\ell]_{sj} = \boldsymbol{K}^\ell_M(\tau_s-\tau_j)$, and $K''_M(0) = -\frac{4\pi^2(M^2-1)}{3}$.

Denote

$$\boldsymbol{D} = \begin{bmatrix} \boldsymbol{D}_0 & \frac{1}{\sqrt{|K_M''(0)|}}\boldsymbol{D}_1 \\ -\frac{1}{\sqrt{|K_M''(0)|}}\boldsymbol{D}_1 & -\frac{1}{|K_M''(0)|}\boldsymbol{D}_2 \end{bmatrix} \in \mathbb{C}^{2JK\times 2JK}.$$

Thus, as long as $\boldsymbol{D}$ is invertible, one can obtain $\boldsymbol{\alpha}_j$ and $\boldsymbol{\beta}_j$. To show the invertibility of $\boldsymbol{D}$, we first show that $\mathbb{E}\boldsymbol{D}$ is invertible under the minimum separation condition in Lemma 3.2. Then, in Lemma 3.4, we verify that $\boldsymbol{D}$ is invertible with high probability after arguing that $\boldsymbol{D}$ is close to $\mathbb{E}\boldsymbol{D}$ given enough measurements in Lemma 3.3. Defining

$$\boldsymbol{E}(n) = \begin{bmatrix} e^{i2\pi n\tau_1} \\ \vdots \\ e^{i2\pi n\tau_J} \\ \frac{-i2\pi n}{\sqrt{|K_M''(0)|}}e^{i2\pi n\tau_1} \\ \vdots \\ \frac{-i2\pi n}{\sqrt{|K_M''(0)|}}e^{i2\pi n\tau_J} \end{bmatrix},$$

we have

$$\boldsymbol{D} = \frac{1}{M}\sum_{n=-2M}^{2M}\left(\left(g_M(n)\boldsymbol{E}(n)\boldsymbol{E}(n)^H\right)\otimes\left(\boldsymbol{b}_n\boldsymbol{b}_n^H\right)\right)$$

and

$$\begin{aligned}\mathbb{E}\boldsymbol{D} &= \frac{1}{M}\sum_{n=-2M}^{2M}\left(\left(g_M(n)\boldsymbol{E}(n)\boldsymbol{E}(n)^H\right)\otimes\boldsymbol{I}_K\right)\\ &= \boldsymbol{D}'\otimes\boldsymbol{I}_K,\end{aligned}$$

where

$$\boldsymbol{D}' = \begin{bmatrix} \boldsymbol{D}_0' & \frac{1}{\sqrt{|K_M''(0)|}}\boldsymbol{D}_1' \\ -\frac{1}{\sqrt{|K_M''(0)|}}\boldsymbol{D}_1' & -\frac{1}{|K_M''(0)|}\boldsymbol{D}_2' \end{bmatrix} \in \mathbb{C}^{2J\times 2J}$$

with $[\boldsymbol{D}_\ell']_{sj} = K_M^\ell(\tau_s-\tau_j)$.

Lemma 3.2. *Suppose that $\Delta_\tau \geq \frac{1}{M}$. Then, $\mathbb{E}\boldsymbol{D}$ is invertible and*

$$\|\boldsymbol{I}_{2JK} - \mathbb{E}\boldsymbol{D}\| \leq 0.3623,$$

$$\|\mathbb{E}\boldsymbol{D}\| \leq 1.3623,$$

$$\left\|(\mathbb{E}\boldsymbol{D})^{-1}\right\| \leq 1.568.$$

The proof of Lemma 3.2 is given in Appendix B. We denote $\boldsymbol{D}'^{-1} = \begin{bmatrix} \boldsymbol{L}' & \boldsymbol{R}' \end{bmatrix}$, where $\boldsymbol{L}' \in \mathbb{C}^{2J\times J}$ and $\boldsymbol{R}' \in \mathbb{C}^{2J\times J}$. To have a concentration of measure result for $\boldsymbol{D}$ as given in the following lemma, we note that

$$\boldsymbol{D} - \mathbb{E}\boldsymbol{D} = \frac{1}{M}\sum_{n=-2M}^{2M} \left(\left(g_M(n)\boldsymbol{E}(n)\boldsymbol{E}(n)^H\right) \otimes \left(\boldsymbol{b}_n\boldsymbol{b}_n^H - \boldsymbol{I}_K\right)\right),$$

which is a sum of independent matrices of zero mean. It is derived in Lemma 2 of [48].

Lemma 3.3. *[48] For any $\varepsilon_1 \in (0, 0.6376), 0 < \delta_1 < 1$, if $M \geq \frac{80\mu JK}{\varepsilon_1^2}\log\left(\frac{4JK}{\delta_1}\right)$ and $\Delta_\tau \geq \frac{1}{M}$, then $\|\boldsymbol{D} - \mathbb{E}\boldsymbol{D}\| \leq \varepsilon_1$ with probability at least $1 - \delta_1$.*

Lemma 3.4. *Under the assumptions of Lemma 3.3, $\boldsymbol{D}$ is invertible with probability at least $1 - \delta_1$.*

Proof.

$$\begin{aligned} \|\boldsymbol{I}_{2JK} - \boldsymbol{D}\| &\leq \|\boldsymbol{I}_{2JK} - \mathbb{E}\boldsymbol{D}\| + \|\mathbb{E}\boldsymbol{D} - \boldsymbol{D}\| \\ &\leq 0.3623 + \varepsilon_1 \\ &< 1, \end{aligned}$$

where the first inequality uses the triangle inequality, and the second inequality follows from Lemmas 3.2 and 3.3. □

For $\varepsilon_1 > 0$, we define the event $\mathcal{E}_{1,\varepsilon_1} = \{\|\boldsymbol{D} - \mathbb{E}\boldsymbol{D}\| \leq \varepsilon_1\}$. The following lemma will also be useful later in our analysis.

Lemma 3.5. *[12, 48] Under the event $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in \left(0, \frac{1}{4}\right]$, we have*

$$\left\|\boldsymbol{D}^{-1} - (\mathbb{E}\boldsymbol{D})^{-1}\right\| \leq 2\left\|(\mathbb{E}\boldsymbol{D})^{-1}\right\|^2 \varepsilon_1,$$

and

$$\left\|\boldsymbol{D}^{-1}\right\| \leq 2\left\|(\mathbb{E}\boldsymbol{D})^{-1}\right\|.$$

Therefore, the construction of the dual polynomial ensures that condition (3.12) is satisfied automatically. It should also be pointed out that Lemma 3.4 essentially guarantees that the set formed in the second condition of Proposition 3.1 is linearly independent. In the following subsection, we focus on showing (3.13): $\|\boldsymbol{q}(\tau)\|_2 < 1$ for all $\tau \notin \mathbb{D}$. The proof is partitioned into the following three major steps:

- Showing that the random dual polynomial $\boldsymbol{q}(\tau)$ concentrates around $\overline{\boldsymbol{q}}(\tau)$ on a discrete set Ω_{Grid}.

- Then, showing that the random dual polynomial $\boldsymbol{q}(\tau)$ concentrates around $\overline{\boldsymbol{q}}(\tau)$ everywhere in $[0, 1)$.

- Eventually, showing that $\|\boldsymbol{q}(\tau)\| < 1, \tau \notin \mathbb{D}$.

This proof strategy was first developed in [12] for compressed sensing off the grid, and was later adopted by [37, 38] for line spectrum estimation with multiple measurement vectors and by [48] for blind sparse spikes deconvolution.

3.5.3 Showing the boundedness condition

We first show that $\boldsymbol{q}^\ell(\tau)$ concentrates around $\overline{\boldsymbol{q}}^\ell(\tau)$ on a finite discrete set Ω_{Grid}, whose size $|\Omega_{\text{Grid}}|$ will be determined later on. For this purpose, we introduce

$$\begin{aligned}\boldsymbol{V}_\ell(\tau) &= \frac{1}{\sqrt{|K_M''(0)|}^\ell} \begin{bmatrix} \boldsymbol{K}_M^\ell(\tau-\tau_1)^H \\ \vdots \\ \boldsymbol{K}_M^\ell(\tau-\tau_J)^H \\ \frac{1}{\sqrt{|K_M''(0)|}} \boldsymbol{K}_M^{\ell+1}(\tau-\tau_1)^H \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}} \boldsymbol{K}_M^{\ell+1}(\tau-\tau_J)^H \end{bmatrix} \\ &= \frac{1}{M} \sum_{n=-2M}^{2M} g_M(n) \left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}} \right)^\ell e^{-i2\pi n\tau} \boldsymbol{E}(n) \otimes \boldsymbol{b}_n \boldsymbol{b}_n^H \in \mathbb{C}^{2JK \times K}.\end{aligned}$$

Taking the expectation of the above leads to

$$\begin{aligned}
\mathbb{E}\boldsymbol{V}_\ell(\tau) &= \frac{1}{M}\sum_{n=-2M}^{2M} g_M(n)\left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}\right)^\ell e^{-i2\pi n\tau}\boldsymbol{E}(n)\otimes\boldsymbol{I}_K \\
&= \frac{1}{\sqrt{|K_M''(0)|}^\ell}\begin{bmatrix} K_M^\ell(\tau-\tau_1)^* \\ \vdots \\ K_M^\ell(\tau-\tau_J)^* \\ \frac{1}{\sqrt{|K_M''(0)|}}K_M^{\ell+1}(\tau-\tau_1)^* \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}}K_M^{\ell+1}(\tau-\tau_J)^* \end{bmatrix}\otimes\boldsymbol{I}_K \\
&=: \boldsymbol{v}_\ell(\tau)\otimes\boldsymbol{I}_K.
\end{aligned}$$

Setting $\boldsymbol{D}^{-1} = \begin{bmatrix}\boldsymbol{L} & \boldsymbol{R}\end{bmatrix}$, with $\boldsymbol{L},\boldsymbol{R}\in\mathbb{C}^{2JK\times JK}$ and using the fact that

$$\boldsymbol{D}\begin{bmatrix}\boldsymbol{\alpha}\\ \sqrt{|K_M''(0)|}\boldsymbol{\beta}\end{bmatrix} = \begin{bmatrix}\boldsymbol{h}\\ \boldsymbol{0}_{JK\times 1}\end{bmatrix},$$

we have

$$\begin{aligned}
\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau) &= \sum_{j=1}^J \frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{K}_M^\ell(\tau-\tau_j)\boldsymbol{\alpha}_j + \sum_{j=1}^J\sqrt{|K_M''(0)|}\frac{1}{\sqrt{|K_M''(0)|}^{\ell+1}}\boldsymbol{K}_M^{\ell+1}(\tau-\tau_j)\boldsymbol{\beta}_j \\
&= \boldsymbol{V}_\ell(\tau)^H\boldsymbol{D}^{-1}\begin{bmatrix}\boldsymbol{h}\\ \boldsymbol{0}_{JK\times 1}\end{bmatrix} \\
&= \boldsymbol{V}_\ell(\tau)^H\boldsymbol{L}\boldsymbol{h}.
\end{aligned}$$

We decompose $\boldsymbol{V}_\ell(\tau)^H\boldsymbol{L}\boldsymbol{h}$ into three parts as follows:

$$\begin{aligned}
\boldsymbol{V}_\ell(\tau)^H\boldsymbol{L}\boldsymbol{h} &= \left(\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)+\mathbb{E}\boldsymbol{V}_\ell(\tau)\right)^H\left(\boldsymbol{L}-\boldsymbol{L}'\otimes\boldsymbol{I}_K+\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)\boldsymbol{h} \\
&= \left[\mathbb{E}\boldsymbol{V}_\ell(\tau)\right]^H\left(\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)\boldsymbol{h} + \left(\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\right)^H\boldsymbol{L}\boldsymbol{h} + \left[\mathbb{E}\boldsymbol{V}_\ell(\tau)\right]^H\left(\boldsymbol{L}-\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)\boldsymbol{h} \\
&= \left[\mathbb{E}\boldsymbol{V}_\ell(\tau)\right]^H\left(\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)\boldsymbol{h} + \boldsymbol{I}_1^\ell(\tau)+\boldsymbol{I}_2^\ell(\tau),
\end{aligned}$$

where we have defined

$$\boldsymbol{I}_1^\ell(\tau) = \left(\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\right)^H\boldsymbol{L}\boldsymbol{h},$$

$$\boldsymbol{I}_2^\ell(\tau) = \left[\mathbb{E}\boldsymbol{V}_\ell(\tau)\right]^H\left(\boldsymbol{L}-\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)\boldsymbol{h}.$$

For the first term $[\mathbb{E}\boldsymbol{V}_\ell(\tau)]^H(\boldsymbol{L}'\otimes\boldsymbol{I}_K)\boldsymbol{h}$ above, we notice that

$$\begin{aligned}[\mathbb{E}\boldsymbol{V}_\ell(\tau)]^H(\boldsymbol{L}'\otimes\boldsymbol{I}_K)\boldsymbol{h} &= (\boldsymbol{v}_\ell(\tau)\otimes\boldsymbol{I}_K)^H(\boldsymbol{L}'\otimes\boldsymbol{I}_K)\boldsymbol{h}\\ &= (\boldsymbol{v}_\ell(\tau)\otimes\boldsymbol{I}_K)^H\begin{bmatrix}\boldsymbol{L}'\otimes\boldsymbol{I}_K & \boldsymbol{R}'\otimes\boldsymbol{I}_K\end{bmatrix}\begin{bmatrix}\boldsymbol{h}\\ \boldsymbol{0}_{JK\times 1}\end{bmatrix}\\ &= (\boldsymbol{v}_\ell(\tau)\otimes\boldsymbol{I}_K)^H(\boldsymbol{D}'\otimes\boldsymbol{I}_K)^{-1}\begin{bmatrix}\boldsymbol{h}\\ \boldsymbol{0}_{JK\times 1}\end{bmatrix}\\ &= (\boldsymbol{v}_\ell(\tau)\otimes\boldsymbol{I}_K)^H\begin{bmatrix}\overline{\boldsymbol{\alpha}}\\ \sqrt{|K_M''(0)|}\overline{\boldsymbol{\beta}}\end{bmatrix},\end{aligned} \tag{3.23}$$

where $\overline{\boldsymbol{\alpha}}=\begin{bmatrix}\overline{\boldsymbol{\alpha}}_1^H & \cdots & \overline{\boldsymbol{\alpha}}_J^H\end{bmatrix}^H$ and $\overline{\boldsymbol{\beta}}=\begin{bmatrix}\overline{\boldsymbol{\beta}}_1^H & \cdots & \overline{\boldsymbol{\beta}}_J^H\end{bmatrix}^H$. Matching (3.23) to (3.22), we have

$$[\mathbb{E}\boldsymbol{V}_\ell(\tau)]^H(\boldsymbol{L}'\otimes\boldsymbol{I}_K)\boldsymbol{h}=\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau).$$

As a consequence, we obtain the following decomposition for $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau)$:

$$\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau)=\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau)+\boldsymbol{I}_1^\ell(\tau)+\boldsymbol{I}_2^\ell(\tau). \tag{3.24}$$

We now present the roadmap for the rest of the proof. First, we demonstrate that $\left\|\boldsymbol{I}_1^\ell(\tau)\right\|_2$ and $\left\|\boldsymbol{I}_2^\ell(\tau)\right\|_2$ are small on the set of grid points Ω_{Grid}, respectively. Then, we combine these facts to establish that $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau)$ is close to $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau)$ on Ω_{Grid}. Finally, we extend the argument to the whole continuous domain $[0,1)$ and show that $\|\boldsymbol{q}(\tau)\|_2<1, \tau\in[0,1)\backslash\mathbb{D}$.

- **Bound $\left\|\boldsymbol{I}_1^\ell(\tau)\right\|_2$ on the set of points Ω_{Grid}**

In order to bound $\left\|\boldsymbol{I}_1^\ell(\tau)\right\|_2$, we will apply the matrix Bernstein inequality by exploiting the randomness of $\boldsymbol{h}_j$. For this purpose, we first need to control $\left\|(\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau))^H\boldsymbol{L}\right\|$, which further requires estimating $\|\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\|$. The latter is established in Lemma 3.6, whose proof uses the matrix Bernstein inequality and can be found in Appendix C.

Lemma 3.6. *Fix $\tau\in[0,1)$, and let $0<\varepsilon_2<1$. Then, for $\ell=0,1,2,3$, $\|\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\|\le\varepsilon_2$ holds with probability at least $1-4\delta_2$ provided that $M\ge\frac{640\cdot 4^{2\ell}\mu JK}{3\varepsilon_2^2}\log\left(\frac{2JK+K}{\delta_2}\right)$.*

We define the event $\mathcal{E}_{2,\varepsilon_2}=\{\|\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\|\le\varepsilon_2,\ell=0,1,2,3\}$. The following two lemmas control $\left\|(\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau))^H\boldsymbol{L}\right\|$ and $\|\boldsymbol{V}_\ell(\tau)-\mathbb{E}\boldsymbol{V}_\ell(\tau)\|_F$, respectively.

Lemma 3.7. *Assume that $\varepsilon_1 \in (0, \frac{1}{4}]$. Consider the finite set of grid points $\Omega_{\text{Grid}} = \{\tau_d\}$, whose cardinality $|\Omega_{\text{Grid}}|$ will be determined later. We have*

$$\mathbb{P}\left\{\sup_{\tau_d \in \Omega_{\text{Grid}}} \left\| \left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L} \right\| \geq 4\varepsilon_2, \ell = 0,1,2,3 \right\} \leq |\Omega_{\text{Grid}}| 4\delta_2 + \mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right)$$

provided that $M \geq \frac{640 \cdot 4^{2\ell} \mu JK}{3\varepsilon_2^2} \log\left(\frac{2JK+K}{\delta_2}\right)$.

Lemma 3.7 is a consequence of Lemmas 3.5 and 3.6. Its proof is given in Appendix D.

Lemma 3.8. *Conditioned on the event $\mathcal{E}_{2,\varepsilon_2}$, we have*

$$\left\|\boldsymbol{V}_\ell(\tau) - \mathbb{E}\boldsymbol{V}_\ell(\tau)\right\|_F \leq \sqrt{K} \left\|\boldsymbol{V}_\ell(\tau) - \mathbb{E}\boldsymbol{V}_\ell(\tau)\right\| \leq \sqrt{K}\varepsilon_2.$$

Based on Lemmas 3.7 and 3.8, the following lemma shows that $\left\|\boldsymbol{I}_1^\ell(\tau)\right\|_2$ can be well-controlled on the set of points Ω_{Grid} with high probability.

Lemma 3.9. *Assume that $\boldsymbol{h}_j \in \mathbb{C}^K, j = 1, \ldots, J$, are i.i.d. symmetric random samples from the complex unit sphere $\mathbb{CS}^{K-1}$. There exists a numerical constant C such that if*

$$M \geq C\mu JK \max\left\{\frac{1}{\varepsilon_4^2} \log\left(\frac{|\Omega_{\text{Grid}}|JK}{\delta}\right) \log^2\left(\frac{|\Omega_{\text{Grid}}|K}{\delta}\right), \log\left(\frac{JK}{\delta}\right)\right\},$$

then we have

$$\mathbb{P}\left\{\sup_{\tau_d \in \Omega_{\text{Grid}}} \left\|\boldsymbol{I}_1^\ell(\tau_d)\right\|_2 \leq \varepsilon_4, \ell = 0,1,2,3\right\} \geq 1 - 12\delta.$$

The detailed proof can be found in Appendix E, which is based on an application of the matrix Bernstein inequality for the vector case.

- **Bound $\left\|\boldsymbol{I}_2^\ell(\tau)\right\|_2$ on the set of points Ω_{Grid}**

The strategy for bounding $\left\|\boldsymbol{I}_2^\ell(\tau)\right\|_2$ is similar to the one for bounding $\left\|\boldsymbol{I}_1^\ell(\tau)\right\|_2$. First of all, we use the following lemma to bound $\left\|\left(\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\right)^H \mathbb{E}\boldsymbol{V}_\ell(\tau)\right\|_F^2$. Its proof is given in Appendix F.

Lemma 3.10. *Conditioned on the event $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$, we have*

$$\left\|\left(\boldsymbol{L}-\boldsymbol{L}'\otimes\boldsymbol{I}_K\right)^H\mathbb{E}\boldsymbol{V}_\ell(\tau)\right\|_F^2\leq CK\varepsilon_1^2$$

for some constant C.

Lemma 3.10 together with the matrix Bernstein inequality ensure that $\left\|\boldsymbol{I}_2^\ell(\tau)\right\|_2$ is small on the set of points Ω_{Grid} with high probability.

Lemma 3.11. *Assume that each $\boldsymbol{h}_j\in\mathbb{C}^K$ is an i.i.d. symmetric random sample from the complex unit sphere $\mathbb{CS}^{K-1}$. There exists a numerical constant C such that if*

$$M\geq C\frac{\mu JK}{\varepsilon_5^2}\log\left(\frac{JK}{\delta}\right)\log^2\left(\frac{|\Omega_{\text{Grid}}|K}{\delta}\right),$$

we have

$$\mathbb{P}\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\boldsymbol{I}_2^\ell(\tau_d)\right\|_2\leq\varepsilon_5,\ell=0,1,2,3\right\}\geq 1-8\delta.$$

The proof of Lemma 3.11 is provided in Appendix G.

- **Show that $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau)$ is close to $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau)$ on the set of points Ω_{Grid}**

Define the event

$$\mathcal{E}=\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau_d)-\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau_d)\right\|_2\leq\frac{\varepsilon}{3},\ell=0,1,2,3\right\}.$$

Lemmas 3.9 and 3.11 along with the decomposition (3.24) immediately result in the following proposition.

Proposition 3.2. *Assume that $\Omega_{\text{Grid}}\subset[0,1)$ is a finite set of points. Let $\delta\in(0,1)$ be the failure probability. Then, there exists a numerical constant C such that when*

$$M\geq C\mu JK\max\left\{\frac{1}{\varepsilon^2}\log\left(\frac{|\Omega_{\text{Grid}}|JK}{\delta}\right)\log^2\left(\frac{|\Omega_{\text{Grid}}|K}{\delta}\right),\log\left(\frac{JK}{\delta}\right)\right\},$$

we have

$$\mathbb{P}(\mathcal{E})\geq 1-\delta.$$

- **Show that $\|\boldsymbol{q}(\tau)\|_2 < 1$ everywhere, $\tau \in [0,1) \setminus \mathbb{D}$**

The following lemma shows that $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau)$ is close to $\frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau)$ everywhere in $[0,1)$. The proof, which involves Bernstein's polynomial inequality, is given in Appendix H.

Lemma 3.12. *Assume that $\Delta_\tau \geq \frac{1}{M}$. Then for all $\tau \in [0,1)$ and $\ell = 0,1,2,3$, the following holds with probability at least $1-\delta$*

$$\left\| \frac{1}{\sqrt{|K_M''(0)|}^\ell}\boldsymbol{q}^\ell(\tau) - \frac{1}{\sqrt{|K_M''(0)|}^\ell}\overline{\boldsymbol{q}}^\ell(\tau) \right\|_2 \leq \varepsilon$$

provided that $M \geq C\mu JK \max\left\{\frac{1}{\varepsilon^2}\log\left(\frac{MJK}{\varepsilon\delta}\right)\log^2\left(\frac{MK}{\varepsilon\delta}\right), \log\left(\frac{JK}{\delta}\right)\right\}$ for some numerical constant C.

Next, we show that $\|\boldsymbol{q}(\tau)\|_2 < 1$ everywhere, $\tau \in [0,1) \setminus \mathbb{D}$. To do this, define

$$\Omega_{\text{near}} = \bigcup_{j=1}^{J} [\tau_j - \tau_{b,1}, \tau_j + \tau_{b,1}],$$

$$\Omega_{\text{far}} = [0,1) \setminus \Omega_{\text{near}}$$

with $\tau_{b,1} = 8.245 \times 10^{-2}\frac{1}{M}$. An immediate consequence of Lemma 3.12 is the following result, which verifies that $\|\boldsymbol{q}(\tau)\|_2 < 1, \forall\ \tau \in \Omega_{\text{far}}$.

Lemma 3.13. *Assume that $\Delta_\tau \geq \frac{1}{M}$ and that*

$$M \geq C\mu JK \log\left(\frac{MJK}{\delta}\right)\log^2\left(\frac{MK}{\delta}\right),$$

for some positive numerical constant C. Then

$$\|\boldsymbol{q}(\tau)\|_2 < 1, \qquad \forall\ \tau \in \Omega_{\text{far}}$$

with probability at least $1-\delta$.

Proof. Taking $\varepsilon = 10^{-5}$ in Lemma 3.12, we obtain that

$$\begin{aligned} \|\boldsymbol{q}(\tau)\|_2 &\leq \|\boldsymbol{q}(\tau) - \overline{\boldsymbol{q}}(\tau)\|_2 + \|\overline{\boldsymbol{q}}(\tau)\|_2 \\ &\leq 10^{-5} + \|\overline{\boldsymbol{q}}(\tau)\|_2 . \end{aligned}$$

In order to bound $\|\boldsymbol{q}(\tau)\|_2$, it remains to control $\|\overline{\boldsymbol{q}}(\tau)\|_2$ for $\tau \in \Omega_{\text{far}}$. From (3.23), note that

$$
\begin{aligned}
\|\overline{\boldsymbol{q}}(\tau)\|_2 &= \sup_{\boldsymbol{u}:\|\boldsymbol{u}\|_2=1} \boldsymbol{u}^H \left([\mathbb{E}\boldsymbol{V}(\tau)]^H (\boldsymbol{L}' \otimes \boldsymbol{I}_K)\boldsymbol{h} \right) \\
&= \sup_{\boldsymbol{u}:\|\boldsymbol{u}\|_2=1} \boldsymbol{u}^H \left((\boldsymbol{v}(\tau) \otimes \boldsymbol{I}_K)^H (\boldsymbol{L}' \otimes \boldsymbol{I}_K)\boldsymbol{h} \right) \\
&= \sup_{\boldsymbol{u}:\|\boldsymbol{u}\|_2=1} \sum_{j=1}^{J} \left(\boldsymbol{v}(\tau)^H \boldsymbol{L}' \right)(j) (\boldsymbol{u}^H \mathrm{sign}(c_j)\boldsymbol{h}_j) \\
&\leq 0.99992,
\end{aligned}
$$

where for the third line above we have, with some abuse of notation, denoted $\left(\boldsymbol{v}(\tau)^H \boldsymbol{L}' \right)(j)$ as the jth entry of the row vector $\boldsymbol{v}(\tau)^H \boldsymbol{L}'$. The fourth line follows from the fact that $|\boldsymbol{u}^H \mathrm{sign}(c_j)\boldsymbol{h}_j| \leq 1$ and from the proof of Lemma 2.4 in [29] for $\tau \in \Omega_{\text{far}}$. Thus, we have shown

$$
\|\boldsymbol{q}(\tau)\|_2 < 1, \quad \forall\ \tau \in \Omega_{\text{far}}.
$$

□

The next lemma shows that $\|\boldsymbol{q}(\tau)\|_2 < 1$ when $\tau \in \Omega_{\text{near}}$.

Lemma 3.14. *Assume that $\tau \in \Omega_{\text{near}}$. Then as long as*

$$
M \geq C\mu JK \log\left(\frac{MJK}{\delta}\right) \log^2\left(\frac{MK}{\delta}\right),
$$

we have

$$
\|\boldsymbol{q}(\tau)\|_2 < 1, \quad \forall\ \tau \in \Omega_{\text{near}}
$$

with probability at least $1 - \delta$.

Proof. First of all, for $\tau_j \in \mathbb{D}$, we have $\frac{d\|\boldsymbol{q}(\tau)\|_2^2}{d\tau}|_{\tau=\tau_j} = 2\langle \boldsymbol{q}'(\tau_j), \boldsymbol{q}(\tau_j)\rangle_{\mathbb{R}} = 0$. The Taylor's expansion of $\|\boldsymbol{q}(\tau)\|_2^2$ in the interval $[\tau_j - \tau_{b,1}, \tau_j + \tau_{b,1}]$ is given by

$$
\|\boldsymbol{q}(\tau)\|_2^2 = \|\boldsymbol{q}(\tau_j)\|_2^2 + \frac{d\,\|\boldsymbol{q}(\tau)\|_2^2}{d\tau}|_{\tau=\tau_j}(\tau - \tau_j) + \frac{1}{2}\frac{d^2\,\|\boldsymbol{q}(\tau)\|_2^2}{d\tau^2}|_{\tau=\xi}(\tau - \tau_j)^2,
$$

for some $\xi \in [\tau_j - \tau_{b,1}, \tau_j + \tau_{b,1}]$. This implies that in order to show $\|\boldsymbol{q}(\tau)\|_2 < 1, \tau \in \Omega_{\text{near}}$, it is sufficient to verify that

$$
\frac{1}{2}\frac{d^2\,\|\boldsymbol{q}(\tau)\|_2^2}{d\tau^2} = \|\boldsymbol{q}'(\tau)\|_2^2 + \mathrm{Re}\left\{ \boldsymbol{q}''(\tau)^H \boldsymbol{q}(\tau) \right\} < 0
$$

for $\tau \in \Omega_{\text{near}}$. Note that

$$\begin{aligned}\frac{1}{|K_M''(0)|}\left\|\boldsymbol{q}'(\tau)\right\|_2^2 &= \left\|\frac{1}{\sqrt{|K_M''(0)|}}\boldsymbol{q}'(\tau)\right\|_2^2 \\ &= \left\|\frac{1}{\sqrt{|K_M''(0)|}}\left(\boldsymbol{q}'(\tau)-\overline{\boldsymbol{q}}'(\tau)+\overline{\boldsymbol{q}}'(\tau)\right)\right\|_2^2 \\ &\leq \varepsilon^2 + 2\varepsilon\left\|\frac{1}{\sqrt{|K_M''(0)|}}\overline{\boldsymbol{q}}'(\tau)\right\|_2 + \left\|\frac{1}{\sqrt{|K_M''(0)|}}\overline{\boldsymbol{q}}'(\tau)\right\|_2^2,\end{aligned}$$

where the inequality above follows from Lemma 3.12. We also need the following estimate on $\left\|\overline{\boldsymbol{q}}'(\tau)\right\|_2$:

$$\left\|\overline{\boldsymbol{q}}'(\tau)\right\|_2 \leq 1.5765M$$

as given in Appendix I. Therefore, we have

$$\begin{aligned}\frac{1}{|K_M''(0)|}\left\|\boldsymbol{q}'(\tau)\right\|_2^2 &< \varepsilon^2 + 1.7383\varepsilon + \left\|\frac{1}{\sqrt{|K_M''(0)|}}\overline{\boldsymbol{q}}'(\tau)\right\|_2^2 \\ &= \varepsilon^2 + 1.7383\varepsilon + \frac{1}{|K_M''(0)|}\left\|\overline{\boldsymbol{q}}'(\tau)\right\|_2^2,\end{aligned}$$

where the inequality above follows from the fact that $\frac{1}{\sqrt{|K_M''(0)|}} < \frac{1}{M}\sqrt{\frac{3}{\pi^2}}$ for $M \geq 2$.

Next, we bound $\frac{1}{|K_M''(0)|}\text{Re}\left\{\boldsymbol{q}''(\tau)^H\boldsymbol{q}(\tau)\right\}$:

$$\begin{aligned}&\frac{1}{|K_M''(0)|}\text{Re}\left\{\boldsymbol{q}''(\tau)^H\boldsymbol{q}(\tau)\right\} \\ &= \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\left(\boldsymbol{q}''(\tau)-\overline{\boldsymbol{q}}''(\tau)\right)+\frac{1}{|K_M''(0)|}\overline{\boldsymbol{q}}''(\tau)\right)^H\left(\boldsymbol{q}(\tau)-\overline{\boldsymbol{q}}(\tau)+\overline{\boldsymbol{q}}(\tau)\right)\right\} \\ &= \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\left(\boldsymbol{q}''(\tau)-\overline{\boldsymbol{q}}''(\tau)\right)\right)^H\left(\boldsymbol{q}(\tau)-\overline{\boldsymbol{q}}(\tau)\right)\right\} + \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\overline{\boldsymbol{q}}''(\tau)\right)^H\overline{\boldsymbol{q}}(\tau)\right\} \\ &\quad + \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\left(\boldsymbol{q}''(\tau)-\overline{\boldsymbol{q}}''(\tau)\right)\right)^H\overline{\boldsymbol{q}}(\tau)\right\} + \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\overline{\boldsymbol{q}}''(\tau)\right)^H\left(\boldsymbol{q}(\tau)-\overline{\boldsymbol{q}}(\tau)\right)\right\} \\ &\leq \varepsilon^2 + 4.2498\varepsilon + \text{Re}\left\{\left(\frac{1}{|K_M''(0)|}\overline{\boldsymbol{q}}''(\tau)\right)^H\overline{\boldsymbol{q}}(\tau)\right\},\end{aligned}$$

where in the last inequality we have used Lemma 3.12, the fact that $\left\|\overline{\boldsymbol{q}}(\tau)\right\|_2 \leq 1.0361$ for $\tau \in \Omega_{\text{near}}$, and the fact that $\left\|\overline{\boldsymbol{q}}''(\tau)\right\|_2 \leq 21.1451M^2$ for $\tau \in \Omega_{\text{near}}$ from Appendix I. Thus, we

have

$$\begin{aligned}\frac{1}{|K_M''(0)|}\frac{1}{2}\frac{d^2\left\|\boldsymbol{q}(\tau)\right\|_2^2}{d\tau^2} &= \frac{1}{|K_M''(0)|}\left\{\left\|\boldsymbol{q}'(\tau)\right\|_2^2 + \mathrm{Re}\left\{\boldsymbol{q}''(\tau)^H\boldsymbol{q}(\tau)\right\}\right\} \\ &< 2\varepsilon^2 + 5.9881\varepsilon + \frac{1}{|K_M''(0)|}\left\{\left\|\overline{\boldsymbol{q}}'(\tau)\right\|_2^2 + \mathrm{Re}\left\{\overline{\boldsymbol{q}}''(\tau)^H\overline{\boldsymbol{q}}(\tau)\right\}\right\} \\ &\leq 2\varepsilon^2 + 5.9881\varepsilon + \frac{1}{|K_M''(0)|}(-0.3756M^2) \\ &< 2\varepsilon^2 + 5.9881\varepsilon - 0.0285 \\ &< 0,\end{aligned}$$

where in the fourth line above we have used the fact that $\left\|\overline{\boldsymbol{q}}'(\tau)\right\|_2^2 + \mathrm{Re}\left\{\overline{\boldsymbol{q}}''(\tau)^H\overline{\boldsymbol{q}}(\tau)\right\} \leq -0.3756M^2$ from Appendix I; see also [38] for a similar argument. The fifth line follows because $|K_M''(0)| < \frac{4\pi^2M^2}{3}$ and the last line follows by choosing ε small enough, say $\varepsilon = 10^{-5}$. This completes the proof of Lemma 3.14. □

Combining Lemmas 3.13 and 3.14 immediately shows that $\left\|\boldsymbol{q}(\tau)\right\|_2 < 1$ everywhere, $\tau \in [0,1) \setminus \mathbb{D}$.

CHAPTER 4

MODAL ANALYSIS FROM COMPRESSED MEASUREMENTS

The use of wireless sensor networks for structural health monitoring (SHM) has gained considerable popularity over the past years. In SHM, among many other tasks, performing modal analysis is a critical component in assessing the quality of infrastructure. In this chapter[8], we consider the problem of modal analysis using collected vibration data from wireless sensor networks. In particular, we propose the use of computationally feasible atomic norm minimizations for extraction of mode shapes and natural frequencies from spatial and temporal compressed measurements, respectively. Theoretical analysis and numerical simulations are carried out, showing that modal analysis can be accomplished from a minimal number of spatial and temporal compressed measurements.

4.1 Introduction

Structural health monitoring (SHM) refers to the process of evaluating and detecting damage in engineering structures, such as buildings or bridges. It has become an essential part of maintaining the quality and safety of modern infrastructure [56–58]. Traditionally, this inspection process is accomplished by domain experts, who examine some key aspects of the structures with manual procedures. However, these manual processes are time consuming, expensive, and labor intensive. Recently, there is a surge of interest in using wireless sensor networks for collecting vibration data from structures, and subsequently performing data analysis for various inspections [59]. There are several reasons for doing so. First of all, developing a wireless sensor network system for SHM is becoming more and more affordable in recent years and it has already been in usage for many real-world applications. Furthermore, wireless sensor networks provide data collection and transmission for analysis in real-time, thus enabling damage inspection in real-time as well. Thirdly, once the data

[8]This work is in collaboration with Prof. Mike Wakin, Prof. Gongguo Tang, and Miss Shuang Li [53–55].

are transmitted by the wireless sensors to the data center, sophisticated algorithms can be developed and tested for modal analysis and damage detection.

However, utilizing wireless sensor networks for SHM has its own limitations and challenges. For example, the battery lifecycle of the sensors can be short and the batteries need to be changed periodically. Furthermore, sensor failures and replacements are other important factors that need to be taken into consideration in wireless sensor networks. Therefore, new sensing schemes that are energy efficient and modal analysis techniques that are robust to sensor failures and faulty data are desirable for SHM. In this chapter, we propose novel schemes for vibration data collection using compressive sensing techniques. Then, with the collected compressed measurements, atomic norm minimization programs are used for estimation of the natural frequencies and modal shapes. We show that our approach is in fact a better way to perform modal analysis than related ones available in the literature. Under some mild conditions, theoretical analysis is also carried out, showing that one can recover mode shapes and natural frequencies from a near optimal number of compressive measurements.

Consider the model

$$\boldsymbol{x}^\star(t) = \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star A_j^\star e^{i2\pi f_j^\star t}, \tag{4.1}$$

where the vectors $\boldsymbol{\psi}_j^\star \in \mathbb{R}^N, j = 1, \ldots, J$, are the mode shapes, $f_j^\star \in [0, 1), j = 1, \ldots, J$, are natural frequencies[9], and $A_j^\star \in \mathbb{R}, j = 1, \ldots, J$, are the amplitudes. We have used the bold notation $\boldsymbol{x}^\star(t)$ to indicate that $\boldsymbol{x}^\star(t) \in \mathbb{R}^N$ is a vector for a given time index t because the mode shapes $\{\boldsymbol{\psi}_j^\star\}$ are vectors. As we will show in Section 4.1.1, model (4.1) is ultimately related to a multiple-degree-of-freedom system in SHM.

In order to perform data analysis, we collect measurements from model (4.1) by sampling the time index t. Consider the uniform sampling regime and without loss of generality, assume that indices $t_1 = 1, t_2 = 2, \ldots, t_M = M$ are the sampling time indices[10]. Then,

[9]Without loss of generality, we have assumed that the natural frequencies are normalized.
[10]We can always normalize the consecutive time intervals such that the uniform sampling interval is 1.

we stack all of the data vectors from different time indices together, yielding the following $N \times M$ data matrix:

$$\begin{aligned}\boldsymbol{X}^\star &= \begin{bmatrix}\boldsymbol{x}^\star(t_1) & \boldsymbol{x}^\star(t_2) & \cdots & \boldsymbol{x}^\star(t_M)\end{bmatrix} \\ &= \begin{bmatrix}\boldsymbol{x}^\star(1,1) & \boldsymbol{x}^\star(1,2) & \cdots & \boldsymbol{x}^\star(1,M) \\ \boldsymbol{x}^\star(2,1) & \boldsymbol{x}^\star(2,2) & \cdots & \boldsymbol{x}^\star(2,M) \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{x}^\star(N,1) & \boldsymbol{x}^\star(N,2) & \cdots & \boldsymbol{x}^\star(N,M)\end{bmatrix},\end{aligned}$$

where $\boldsymbol{X}^\star(p,q) = \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(p) A_j^\star e^{i2\pi f_j^\star q}, 1 \le p \le N, 1 \le q \le M$. The purpose of modal analysis is to extract mode shapes $\boldsymbol{\psi}_j^\star, j = 1, \ldots, J$, and frequencies $f_j^\star, j = 1, \ldots, J$, from the data matrix $\boldsymbol{X}^\star$.

As mentioned previously, it is desirable to reduce the consumption of energy in data collection of wireless sensor networks as this will prolong the life of battery embedded in the sensors and eventually will reduce the cost of maintaining a wireless sensor network. One way to increase the battery lifecycle of wireless sensor networks is to sample less data while still being able to gather enough information of vibration signals for estimating mode shapes and natural frequencies. Note that the number of samples within the data matrix is MN while the number of degrees of freedom in model (4.1) is proportional to JN. Therefore, when $J \ll M$ (in practice, quite often, there are only a few dominant natural frequencies), certain amounts of redundancy exist in the data matrix $\boldsymbol{X}^\star$. This further suggests that one could possibly compress the data matrix $\boldsymbol{X}^\star$ significantly. In this chapter, we consider both time and spatial compression schemes for reducing the number of measurements for the estimation of mode shapes and frequencies.

4.1.1 Background

In this section, we show that model (4.1) represents a class of systems that often arises in many real-world vibrational structures.

Consider a multiple-degree-of-freedom system that governs the motion of structures:

$$\boldsymbol{M}\ddot{\boldsymbol{x}}^\star(t) + \boldsymbol{C}\dot{\boldsymbol{x}}^\star(t) + \boldsymbol{K}\boldsymbol{x}^\star(t) = \boldsymbol{0}(t), \tag{4.2}$$

where the diagonal matrix $\boldsymbol{M}$ is an $N \times N$ mass matrix, $\boldsymbol{C} \in \mathbb{R}^{N\times N}$ is the symmetric damping matrix, and $\boldsymbol{K} \in \mathbb{R}^{N\times N}$ represents the symmetric stiffness matrix. For notational consistency, we have also denoted $\boldsymbol{x}^\star(t) \in \mathbb{R}^{N\times 1}$ as the displacement signal. Also, for simplicity, we have assumed that the system does not have external force, i.e., the right side of equation (4.2) is $\mathbf{0}(t)$. While damping is an important factor that cannot be ignored in reality, in order to simplify the analysis, in the following, we assume an undamped system and set $\boldsymbol{C} = \mathbf{0}$. Thus, the multiple-degree-of-freedom system simplifies to

$$\boldsymbol{M}\ddot{\boldsymbol{x}}^\star(t) + \boldsymbol{K}\boldsymbol{x}^\star(t) = \mathbf{0}(t). \tag{4.3}$$

It is shown that the solution to the above second-order differential equation is characterized by the following analytical formula

$$\boldsymbol{x}^\star(t) = \sum_{j=1}^{N} \boldsymbol{\psi}_j^\star \rho_j^\star \sin(2\pi f_j^\star t + \theta_j^\star), \tag{4.4}$$

where $\boldsymbol{\psi}_j^\star, j = 1, \ldots, N$, are the mode shapes and $\rho_j^\star, j = 1, \ldots, N$, are the amplitudes associated with the mode shapes $\left\{\boldsymbol{\psi}_j^\star\right\}$. At first glance, the solution equation (4.4) to model (4.3) is different from (4.1). In particular, there are two distinct differences. First of all, the summation in (4.1) ranges from 1 to J instead of from 1 to N. Secondly, the complex exponential term $e^{i2\pi f_j^\star t}$ in (4.1) is replaced by $\sin(2\pi f_j^\star t + \theta_j^\star)$ in (4.4).

However, if we assume that only J frequencies are excited, there will be only J nonzero terms appearing in model (4.1). Without loss of generality, in model (4.4), we assume that $\rho_{J+1}^\star = \rho_{J+2}^\star = \cdots = \rho_N^\star = 0$ and these corresponding terms disappear. Thus, equation (4.4) becomes

$$\boldsymbol{x}^\star(t) = \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star \rho_j^\star \sin(2\pi f_j^\star t + \theta_j^\star). \tag{4.5}$$

In addition, it can be shown that the analytical signal of $\boldsymbol{x}^\star(t)$ in (4.5) is exactly equivalent to model (4.1) considered in this chapter. Therefore, in the following sections, we will only consider model (4.1).

4.1.2 Contributions

Our contributions in this chapter are threefold. First of all, we show that the data matrix $\boldsymbol{X}^\star$ obeying model (4.1) enjoys an atomic decomposition, which can be found via a computationally feasible atomic norm minimization. Based on this, we can extract the mode shapes and frequencies from the atomic decomposition. Secondly, we apply a spatial compression scheme to the data matrix $\boldsymbol{X}^\star$ and show that exact recovery of mode shapes and frequencies is still possible as long as the number of spatial measurements is proportional to the number of degrees of freedom in $\boldsymbol{X}^\star$, up to a polylogarithmic factor. Thirdly, we investigate the possibility of applying a temporal compression scheme to the data matrix $\boldsymbol{X}^\star$ and show numerically that exact recovery of mode shapes and natural frequencies is possible with a minimal number of temporal measurements.

4.1.3 Related work

Our work is mostly related to the work by J. Park et al [60]. In [60], the authors propose to use an SVD approach for estimating the mode shapes and natural frequencies. Specifically, they directly apply an SVD to the data matrix $\boldsymbol{X}^\star$ and take the left singular vectors of $\boldsymbol{X}^\star$ as the estimates of the mode shapes. Theoretical analysis is also carried out, showing that the number of measurements is inversely proportional to the minimum separation condition of an arbitrary pair of frequencies, a similar condition that we will use as well. However, their approach has obvious drawbacks. It is well-known that the singular vectors in any SVD are orthogonal with each other. However, in practice, the mode shapes are often not orthogonal and can be highly correlated. When the mode shapes are not orthogonal, the SVD based approach gives poor estimates of the mode shapes. In [60], a random temporal compression scheme is also proposed to reduce the number of measurements and theoretical analysis is carried out using techniques developed in the compressive sensing literature. Again, their random temporal compression scheme is limited by the drawbacks of SVD based approaches.

The main reason that SVD baed approaches are not the best candidates for modal analysis is that they only take advantages of the fact that the data matrix $\boldsymbol{X}^\star$ is low-rank and do not exploit other rich structures that $\boldsymbol{X}^\star$ obeys. In fact, certain additional structure occurs in model (4.1) and the atomic norm minimization based approaches developed in this chapter will take a full advantage of this particular structure.

4.2 Modal Analysis via an Atomic Norm Minimization

This section is outlined as follows. In Section 4.2.1, we show that the data matrix $\boldsymbol{X}^\star$ admits an atomic decomposition, which enables us to extract mode shapes and frequencies. Based on this, in Section 4.2.2, we propose a spatial compression scheme for the data matrix $\boldsymbol{X}^\star$, which is a direct application of the main theorem in Chapter 3. Finally, in Section 4.2.3, we show that it is also possible to deploy a temporal compression scheme for the data matrix $\boldsymbol{X}^\star$.

4.2.1 Modal analysis via an atomic decomposition

Recall that the data matrix

$$\boldsymbol{X}^\star = \begin{bmatrix} \boldsymbol{x}^\star(1,1) & \boldsymbol{x}^\star(1,2) & \cdots & \boldsymbol{x}^\star(1,M) \\ \boldsymbol{x}^\star(2,1) & \boldsymbol{x}^\star(2,2) & \cdots & \boldsymbol{x}^\star(2,M) \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{x}^\star(N,1) & \boldsymbol{x}^\star(N,2) & \cdots & \boldsymbol{x}^\star(N,M) \end{bmatrix},$$

where $\boldsymbol{X}^\star(p,q)$ is the pth component of the signal vector

$$\boldsymbol{x}^\star(q) = \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star A_j^\star e^{i2\pi f_j^\star q}.$$

Thus, we can write the data matrix $\boldsymbol{X}^\star$ in a more explicit form:

$$\boldsymbol{X}^\star = \begin{bmatrix} \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star M} \\ \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star M} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^{J} \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star M} \end{bmatrix}.$$

Note that we can interpret $\boldsymbol{X}^\star(p,q)$ as sensor p's measurement at time index q. One can imagine that N sensors are displaced on a structure and they measure vibrational signal simultaneously at time indices $t = 1, 2, \ldots, M$.

Define the atomic set (see Section 3.2.1 in Chapter 3 for a detailed definition)

$$\mathcal{A} = \left\{ \boldsymbol{h}\boldsymbol{a}(f)^H : \ f \in [0,1), \|\boldsymbol{h}\|_2 = 1, \boldsymbol{h} \in \mathbb{C}^{N\times 1} \right\}$$

with $\boldsymbol{a}(f)$ defined as

$$\boldsymbol{a}(f) = \begin{bmatrix} e^{-i2\pi f1} & e^{-i2\pi f2} & \cdots & e^{-i2\pi fM} \end{bmatrix}^T.$$

Then we have

$$\boldsymbol{X}^\star = \sum_{j=1}^{J} A_j^\star \|\boldsymbol{\psi}_j^\star\|_2 \begin{bmatrix} \frac{\boldsymbol{\psi}_j^\star(1)}{\|\boldsymbol{\psi}_j^\star\|_2} \\ \frac{\boldsymbol{\psi}_j^\star(2)}{\|\boldsymbol{\psi}_j^\star\|_2} \\ \vdots \\ \frac{\boldsymbol{\psi}_j^\star(N)}{\|\boldsymbol{\psi}_j^\star\|_2} \end{bmatrix} \boldsymbol{a}(f_j^\star)^H. \tag{4.6}$$

From (4.6), we see clearly that $\boldsymbol{X}^\star$ is essentially a linear combination of J atoms from the atomic set $\mathcal{A}$. Motivated by this observation, we propose to use the following program to find the rank-one decomposition of $\boldsymbol{X}^\star$:

$$\begin{aligned} \underset{\boldsymbol{X}}{\text{minimize}} \quad & \|\boldsymbol{X}\|_{\mathcal{A}} \\ \text{subject to} \quad & \boldsymbol{X} = \boldsymbol{X}^\star, \end{aligned} \tag{4.7}$$

where the atomic norm is defined as

$$\begin{aligned} \|\boldsymbol{X}\|_{\mathcal{A}} &= \inf \left\{ t > 0 : \ \boldsymbol{X} \in t\mathbf{conv}(\mathcal{A}) \right\} \\ &= \inf_{c_k, \tau_k, \|\boldsymbol{h}_k\|_2 = 1} \left\{ \sum_k |c_k| : \ \boldsymbol{X} = \sum_k c_k \boldsymbol{h}_k \boldsymbol{a}(\tau_k)^H \right\}. \end{aligned}$$

We would like to give a few comments on investigating the atomic norm minimization program (4.7):

- Given the data matrix $\boldsymbol{X}^\star$ collected from N sensors, our aim in modal analysis is to identify the correct natural frequencies $\left\{ f_j^\star \right\}$ and mode shapes $\left\{ \boldsymbol{\psi}_j^\star \right\}$. As we will show in a moment, if there is a unique solution to (4.7) (i.e., the unique atomic decomposition (4.6)), one can find natural frequencies $\left\{ f_j^\star \right\}$ and mode shapes $\left\{ \boldsymbol{\psi}_j^\star \right\}$ exactly via a dual

polynomial.

- Furthermore, program (4.7) is convex and can be transformed into an equivalent semidefinite program

$$\begin{aligned} \text{minimize}_{\boldsymbol{u},\boldsymbol{T},\boldsymbol{X}} \quad & \frac{1}{2N}\text{trace}\left(\text{Toep}(\boldsymbol{u})\right) + \frac{1}{2}\text{trace}(\boldsymbol{T}) \\ \text{subject to} \quad & \begin{bmatrix} \text{Toep}(\boldsymbol{u}) & \boldsymbol{X}^H \\ \boldsymbol{X} & \boldsymbol{T} \end{bmatrix} \succeq \boldsymbol{0} \\ & \boldsymbol{X} = \boldsymbol{X}^\star, \end{aligned}$$

  which can be solved efficiently using off-the-shelf solvers.

- The dual problem of (4.7) is given by

$$\begin{aligned} \underset{\boldsymbol{\Lambda}}{\text{maximize}} \quad & \langle \boldsymbol{\Lambda}, \boldsymbol{X} \rangle \\ \text{subject to} \quad & \|\boldsymbol{\Lambda}\|_{\mathcal{A}}^* \leq 1, \end{aligned} \tag{4.8}$$

  where $\|\cdot\|_{\mathcal{A}}^*$ is the dual norm of the atomic norm $\|\cdot\|_{\mathcal{A}}$. Let $\widetilde{\boldsymbol{\Lambda}}$ be a dual solution to (4.8). Then, we can localize the estimate of natural frequencies $\left\{f_j^\star\right\}$ via the dual polynomial

$$\boldsymbol{q}(f) = \widetilde{\boldsymbol{\Lambda}}\boldsymbol{a}(f)$$

  by selecting out those f's such that $\|\boldsymbol{q}(f)\|_2 = 1$.

- Once the estimates of the frequencies $\left\{f_j^\star\right\}$ are localized, we can plug those estimates back into (4.6) and get the estimates of $\left\{\boldsymbol{\psi}_j^\star\right\}$ via a simple least squares.

- In general, it cannot be guaranteed that the decomposition obtained from solving (4.7) coincides with that by (4.6). However, under certain conditions, solving (4.7) returns the decomposition (4.6) exactly.

Denote $\Delta_f = \min_{i\neq j} |f_i^\star - f_j^\star|, i, j \in \{1, 2, \ldots, J\}$, which is understood as the wrap-around distance on the unit circle. Then, we have the following theorem, which is first established in [38].

Theorem 4.6. *[38] Assume that the data matrix $\boldsymbol{X}^\star$ obeys the atomic decomposition by (4.6). Then, solving the convex program (4.7) recovers the atomic decomposition exactly provided that 1) $\Delta_f \geq \frac{4}{M-1}$; 2) $M \geq 257$.*

Note that Theorem 4.6 is deterministic and requires a minimum separation between an arbitrary pair of frequencies from the set $\{f_1^\star, f_2^\star, \ldots, f_J^\star\}$. It can be clearly seen from the theorem that the closer any pair of frequencies gets, the larger M (the more time-domain measurements) needs to be for exact decomposition and subsequently, exact recovery of $\left\{f_j^\star\right\}$. Finally, we note that program (4.7) is also called frequency estimation from multiple-measurement vectors (MMV) in the signal processing community [37, 38].

4.2.2 Modal analysis from spatially compressed measurements

In this section, we apply a novel compression scheme to the data matrix $\boldsymbol{X}^\star$, leading to a potentially practical sensing strategy for wireless sensor networks in SHM. Recall that

$$\boldsymbol{X}^\star = \begin{bmatrix} \sum_{j=1}^J \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(1) A_j^\star e^{i2\pi f_j^\star M} \\ \sum_{j=1}^J \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(2) A_j^\star e^{i2\pi f_j^\star M} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{j=1}^J \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star 1} & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star 2} & \cdots & \sum_{j=1}^J \boldsymbol{\psi}_j^\star(N) A_j^\star e^{i2\pi f_j^\star M} \end{bmatrix}$$

and each column of $\boldsymbol{X}^\star$ is a collection of measurements from N sensors across the structure at a given synchronized time index. Our strategy for the spatial compression scheme tries to compress each column of the matrix with one measurement, which is a random weighted sum of the entries of the column. Specifically, denoting the qth column of $\boldsymbol{X}^\star$ as $\boldsymbol{X}^\star(:,q)$[11], then the compressed measurement is given by

$$\boldsymbol{y}(q) = \boldsymbol{b}_q^H \boldsymbol{X}^\star(:,q), \tag{4.9}$$

[11] We have used the Matlab notation $\boldsymbol{X}(:,q)$ to represent a column of the matrix $\boldsymbol{X}$.

where $\boldsymbol{b}_q$ is a length-N sensing vector, which has some properties of randomness as we will see later on. Note that we can rewrite (4.9) into

$$\begin{aligned} \boldsymbol{y}(q) &= \boldsymbol{b}_q^H \boldsymbol{X}^\star(:,q) \\ &= \boldsymbol{b}_q^H \boldsymbol{X}^\star \boldsymbol{e}_q \\ &= \langle \boldsymbol{X}^\star, \boldsymbol{b}_q \boldsymbol{e}_q^H \rangle, \end{aligned} \tag{4.10}$$

where we have denoted $\boldsymbol{e}_q, q = 1, \ldots, M$, as the qth column of an $M \times M$ identity matrix. As explained before, our goal is to recover $\boldsymbol{X}^\star$ with the number of measurements $\boldsymbol{y}(q)$ as small as possible.

The sensing mechanism (4.10) appears to be the same as that in the blind non-stationary super-resolution problem discussed in Chapter 3 (see equation (3.8) for a comparison). Therefore, we can apply the atomic norm minimization program utilized in Chapter 3 to the modal analysis with random spatial compression problem considered here. Again, we solve the following convex program

$$\begin{aligned} \underset{\boldsymbol{X}}{\text{minimize}} \quad & \|\boldsymbol{X}\|_{\mathcal{A}} \\ \text{subject to} \quad & \boldsymbol{y}(q) = \langle \boldsymbol{X}, \boldsymbol{b}_q \boldsymbol{e}_q^H \rangle, \quad q = 1, \cdots, M. \end{aligned} \tag{4.11}$$

Furthermore, we obtain a theoretical guarantee for the above program by borrowing Theorem 3.5 developed in Chapter 3 directly. In words, we have the following theorem.

Theorem 4.7. *Assume that the data matrix $\boldsymbol{X}^\star$ is measured through a random spatial compression scheme via (4.9). Let the random vectors $\boldsymbol{b}_q$ be i.i.d. samples from an isotropic and μ-incoherent distribution (More detailed has been developed in Chapter 3). Also, assume that the normalized mode shapes $\frac{\boldsymbol{\psi}_j^\star}{\|\boldsymbol{\psi}_j^\star\|_2}$ are drawn i.i.d. from the uniform distribution on the complex unit sphere and the minimum separation condition $\Delta_f \geq \frac{4}{M-1}$ is met. Then there exists a numerical constant C such that*

$$M \geq C\mu JN \log\left(\frac{MJN}{\delta}\right) \log^2\left(\frac{MN}{\delta}\right) \tag{4.12}$$

is sufficient to guarantee that we can recover $\boldsymbol{X}^\star$ via program (4.11) with probability at least $1 - \delta$.

Before moving on to the random temporal compression scheme, we would like to briefly interpret Theorem 4.7 in the scope of modal analysis. The theoretical bound (4.12) informs us how many spatial compressive measurements to take in order to be able to retrieve natural frequencies and mode shapes exactly. When we deploy incoherent sensing vectors $\boldsymbol{b}_q$ (i.e., when μ can be bounded by a constant), (4.12) simply says that as long as the number of spatial measurements is proportional to the number of degrees of freedom in those mode shape (i.e., JN), up to a polylogarithmic factor, it is sufficient for exact recovery. We also would like to point out that the frequencies can be localized with the aid of the dual polynomial and mode shapes can be estimated via least squares once the frequencies are identified (details can be found in Section 3.2.2).

4.2.3 Modal analysis from temporally compressed measurements

In this section, we investigate the possibility of compressing the data matrix $\boldsymbol{X}^\star$ with a temporal compression scheme. In the spatial compression scheme, we get a sketch of $\boldsymbol{X}^\star$ by replacing each column in $\boldsymbol{X}^\star$ with a numeric value. Similarly, we can obtain a sketch of $\boldsymbol{X}^\star$ by compressing each row of $\boldsymbol{X}^\star$ with a reduced number of measurements that fuse entries of each row randomly, which is the random temporal compression scheme that we will present in the following.

Assume that each row of $\boldsymbol{X}^\star$ is measured as follows

$$\boldsymbol{y}_p^H = \boldsymbol{X}^\star(p,:)\boldsymbol{\Phi}_p, \quad p = 1, \ldots, N. \tag{4.13}$$

Unlike getting one compressed measurement per column in the random spatial compression case, we cannot merely get one fused measurement per row for the temporal compression scheme. Otherwise, we will only collect N measurements, which is much smaller than $O(JN)$, the number of degrees of freedom in the problem. Then, exact recovery of natural frequencies and mode shapes is impossible. Therefore, for each row of $\boldsymbol{X}^\star$, we right-multiply it with a (random) $M \times m$ matrix and collect m measurements with $m < M$. The ratio $\frac{m}{M}$ reflects the compression ratio for the temporal compression scheme. Depending on whether we use the

same $\boldsymbol{\Phi}_p$ or not for different rows of $\boldsymbol{X}^\star$, we have the following two different sensing schemes:

1) The same sensing matrix across rows of $\boldsymbol{X}^\star$

We use $\boldsymbol{\Phi} \in \mathbb{R}^{M\times m}$ to represent the same sensing matrix. In this case, we can rewrite (4.13) as

$$\begin{bmatrix} \boldsymbol{y}_1^H \\ \boldsymbol{y}_2^H \\ \vdots \\ \boldsymbol{y}_N^H \end{bmatrix} = \boldsymbol{X}^\star \boldsymbol{\Phi}.$$

Then, to recover $\boldsymbol{X}^\star$, identify natural frequencies $\left\{f_j^\star\right\}$, and estimate mode shapes $\left\{\boldsymbol{\psi}_j^\star\right\}$, we propose to use the following atomic norm minimization program

$$\begin{aligned} \underset{\boldsymbol{X}}{\text{minimize}} \quad & \|\boldsymbol{X}\|_{\mathcal{A}} \\ \text{subject to} \quad & \begin{bmatrix} \boldsymbol{y}_1^H \\ \boldsymbol{y}_2^H \\ \vdots \\ \boldsymbol{y}_N^H \end{bmatrix} = \boldsymbol{X}\boldsymbol{\Phi}. \end{aligned} \tag{4.14}$$

Program (4.14) can be recast as the following SDP

$$\begin{aligned} &\text{minimize}_{\boldsymbol{u},\boldsymbol{T},\boldsymbol{X}} \quad \frac{1}{2N}\text{trace}\left(\text{Toep}(\boldsymbol{u})\right) + \frac{1}{2}\text{trace}(\boldsymbol{T}) \\ &\text{subject to} \quad \begin{bmatrix} \text{Toep}(\boldsymbol{u}) & \boldsymbol{X}^H \\ \boldsymbol{X} & \boldsymbol{T} \end{bmatrix} \succeq \boldsymbol{0} \\ &\qquad\qquad \begin{bmatrix} \boldsymbol{y}_1^H \\ \boldsymbol{y}_2^H \\ \vdots \\ \boldsymbol{y}_N^H \end{bmatrix} = \boldsymbol{X}\boldsymbol{\Phi}. \end{aligned}$$

2) Different sensing matrices across rows of $\boldsymbol{X}^\star$

In this case, we use $\boldsymbol{\Phi}_p$ to represent the sensing matrix corresponding to the pth row of $\boldsymbol{X}^\star$. Then, we solve the following slightly different atomic norm minimization program

$$\begin{aligned} \underset{\boldsymbol{X}}{\text{minimize}} \quad & \|\boldsymbol{X}\|_{\mathcal{A}} \\ \text{subject to} \quad & \boldsymbol{y}_p = \boldsymbol{\Phi}_p^H \boldsymbol{x}_p, \quad p = 1, \ldots, N \\ & \boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1^H \\ \boldsymbol{x}_2^H \\ \vdots \\ \boldsymbol{x}_N^H \end{bmatrix}. \end{aligned} \tag{4.15}$$

Similarly, (4.15) has the following equivalent SDP

$$\begin{aligned}
&\text{minimize}_{\boldsymbol{u},\boldsymbol{T},\boldsymbol{X}} \quad \frac{1}{2N}\text{trace}\left(\text{Toep}(\boldsymbol{u})\right) + \frac{1}{2}\text{trace}(\boldsymbol{T}) \\
&\text{subject to} \quad \begin{bmatrix} \text{Toep}(\boldsymbol{u}) & \boldsymbol{X}^H \\ \boldsymbol{X} & \boldsymbol{T} \end{bmatrix} \succeq \boldsymbol{0} \\
&\qquad\qquad \boldsymbol{y}_p = \boldsymbol{\Phi}_p^H \boldsymbol{x}_p, \quad p = 1, \ldots, N \\
&\qquad\qquad \boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1^H \\ \boldsymbol{x}_2^H \\ \vdots \\ \boldsymbol{x}_N^H \end{bmatrix}.
\end{aligned}$$

Discussions

- We can interpret (4.15) as a joint recovery approach since program (4.15) recovers $\boldsymbol{X}^\star$ at once instead of row by row.

- In contrast, one could also solve a series of the following SDP programs and recover each row of $\boldsymbol{X}^\star$ separately

$$\begin{aligned}
&\underset{\boldsymbol{x}}{\text{minimize}} \quad \|\boldsymbol{x}_p\|_{\mathcal{A}} \\
&\text{subject to} \quad \boldsymbol{y}_p = \boldsymbol{\Phi}_p^H \boldsymbol{x}_p,
\end{aligned} \tag{4.16}$$

  where p ranges from 1 to N. In other words, (4.16) recovers signals sensor by sensor as each row of the data matrix $\boldsymbol{X}^\star$ represents a signal measured by one sensor.

- We make numerical comparative studies between the joint recovery scheme (4.15) and the separate recovery scheme (4.16).

4.3 Numerical Simulations

In this section, we perform numerical simulations to show the effectiveness of random spatial and temporal compression schemes for modal analysis. For the sake of simplicity, we only consider the temporal compression case; we skip the random spatial compression scheme as we have already observed the numerical behavior in Chapter 3. To set up the experiment, we set the number of uniform samples $M = 80$ and consider various compression ratios by varying M' from 3 to 30. The natural frequencies are chosen as $\{f_1^\star = 0.10, f_2^\star = 0.15, f_3^\star = 0.50\}$.

The corresponding mode shapes are constructed with i.i.d. Gaussian entries and are then normalized to unit norm. We compare the performance of joint recovery (i.e., solving program (4.15)) with separate recovery by solving (4.16) for each sensor. Note that for different N, i.e., different number of sensors, we will have a different ground truth $\boldsymbol{X}^\star$. We declare the recovery a success if the normalized recovery error $\frac{\|\boldsymbol{X}^\star - \widetilde{\boldsymbol{X}}\|_F}{\|\boldsymbol{X}^\star\|_F} < 10^{-5}$, where $\widetilde{\boldsymbol{X}}$ is denoted as an estimate of $\boldsymbol{X}^\star$. The performance is demonstrated in Figure 4.1. As we can see from Figure 4.1 (a) and Figure 4.1 (b), the joint recovery scheme has considerably better performance over the separate recovery approach across different values of M'.

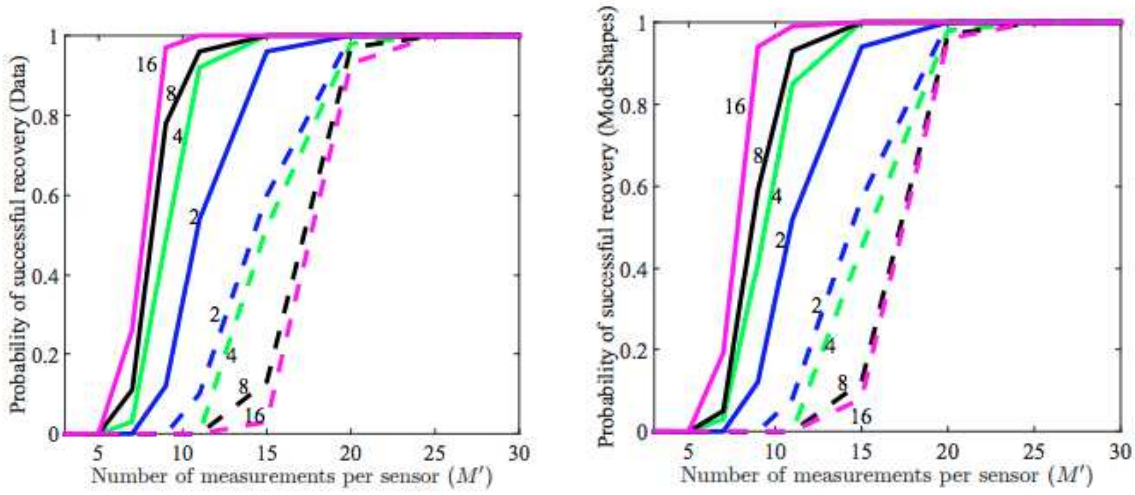

(a) (b)

Figure 4.1: (a) Comparative study between joint recovery and separate recovery of $\boldsymbol{X}^\star$. (b) Comparative study between joint recovery and separate recovery of mode shapes (best viewed in color).

4.4 Conclusion

In this chapter, we have developed a new computationally efficient framework for modal analysis from measurement data in wireless sensor networks. First of all, we have showed that the problem of estimating natural frequencies and mode shapes for modal analysis is essentially an MMV problem, and thus can be solved efficiently via an atomic norm minimization. Based on this fact, we have developed novel spatial and temporal compression schemes for

modal analysis from compressed measurements. Numerical simulations on synthetic data have also been conducted to support the effectiveness of our proposed approaches. Future works include testing the proposed atomic norm minimization programs on real vibration datasets and developing theoretical guarantees for noisy data for both spatial and temporal compression schemes.

CHAPTER 5

MATRIX COMPLETION FOR PAIRWISE COMPARISON MATRICES

In this chapter[12], we introduce a class of skew-symmetric matrices, the pairwise comparison matrices, that naturally arise in many domains. Then, we propose a new low-rank model for pairwise comparison matrices that accommodates non-transitive pairwise comparisons. Based on this model, we consider the regime where one has limited observations of a pairwise comparison matrix. To infer the missing entries, we apply matrix completion techniques to incomplete pairwise comparison matrices.

5.1 Introduction

Quantitative pairwise comparisons arise in a number of areas ranging from recommendation systems, economics, and psychology to competitions involving elections, sports, etc. For example, in recommendation systems, pairwise comparisons can be formed by using available user movie ratings. A pairwise comparison matrix $\boldsymbol{Y}$ encodes all possible comparisons between pairs of items in a set: for a set of N items, $\boldsymbol{Y}$ will have size $N \times N$, and $\boldsymbol{Y}(i, j)$ will denote the strength of the preference of item i over item j.

5.1.1 Motivation

In many applications, due to the increasingly large volume of data sets, pairwise comparison matrices often have some number of missing entries. Given the available pairwise comparisons, inferring the missing entries of the matrix can be valuable for making better decisions and recommendations.

Pairwise comparison matrices are also relevant to the rank aggregation problem, where one seeks to rank a collection of items based on pairwise comparisons of the items. In [62], the

[12]This work is in collaboration with Prof. Mike Wakin [61].

authors assume that each item i, $1 \leq i \leq N$, has an intrinsic value $\boldsymbol{s}(i)$ and describe a rank-two model for the resulting pairwise comparison matrix. In their model, $\boldsymbol{Y}(i,j)$ takes the value $\boldsymbol{s}(i)-\boldsymbol{s}(j)$; this dictates that the full matrix $\boldsymbol{Y}$ has rank two. Given partial observations of $\boldsymbol{Y}$, they propose to fill in the missing entries of $\boldsymbol{Y}$ using nuclear norm minimization, extract the intrinsic score vector $\boldsymbol{s} = \begin{bmatrix} \boldsymbol{s}(1) & \boldsymbol{s}(2) & \cdots & \boldsymbol{s}(N) \end{bmatrix}^T \in \mathbb{R}^N$, and then rank the items based on the values of the scores. However, there are inherent limitations to this approach. For example, the rank-two model above is only meaningful when the pairwise comparisons are transitive, i.e., when $\boldsymbol{Y}(i,j) + \boldsymbol{Y}(j,k) + \boldsymbol{Y}(k,i) = 0$ for all i, j, k. In fact, this transitive property requires the underlying pairwise comparison matrix $\boldsymbol{Y}$ to be rank-two, to be skew-symmetric (i.e., to satisfy $\boldsymbol{Y}(i,j) = -\boldsymbol{Y}(j,i)$ for all i, j), and to have the form of $\boldsymbol{Y}(i,j) = \boldsymbol{s}(i) - \boldsymbol{s}(j)$ for some $\boldsymbol{s} \in \mathbb{R}^N$ [62].

However, social choice theory indicates that real-world data provided by human beings is far more complicated and often exhibits non-transitive behavior [63]. For instance, while people who like item A more than item B, and like item B more than item C, will generally like item A more than item C, this is not always the case. The comparison of two items may rely on multiple factors nonlinearly rather than simply on the difference of two individual intrinsic values. For such scenarios, it is desirable to have a new model that maintains a low complexity but has the ability to capture these complicated pairwise interactions and non-transitive behavior. An efficient algorithm that can infer the missing pairwise comparisons from available ones is also required. In these situations, where comparisons can rely on multiple intrinsic (and typically hidden) factors and many of the pairwise comparisons are non-transitive, one can also argue that it is more meaningful to focus on recovering the pairwise comparison matrix $\boldsymbol{Y}$ itself rather than on recovering a scalar ranking of the items.

5.1.2 Contributions

In this chapter, we develop a new low-rank model for non-transitive pairwise comparison matrices. Our model is based on the fact that any skew-symmetric matrix $\boldsymbol{Y}$ with rank at

most $2r$ can be decomposed in a particular way as a sum of low-rank components:

$$\boldsymbol{Y} = \sum_{k=1}^{r} \boldsymbol{s}_k \boldsymbol{a}_k^T - \boldsymbol{a}_k \boldsymbol{s}_k^T,$$

for some vectors $\boldsymbol{s}_1, \boldsymbol{s}_2, \ldots, \boldsymbol{s}_r \in \mathbb{R}^N$ and $\boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_r \in \mathbb{R}^N$. Based on this model, we adopt the alternating minimization algorithm [64] for low-rank matrix completion [2], allowing reconstruction of missing pairwise comparisons.

5.2 Models for Pairwise Comparison Matrices

In this section, we present our model and discuss its connections with the transitive pairwise comparisons and ranking.

5.2.1 A new model for pairwise comparison matrices

Our model for pairwise comparison matrices is based on the following observation. Suppose that there are r latent factors on which pairwise comparisons among N items are based. For each k, $1 \leq k \leq r$, the comparison between item i and item j based on the k^{th} factor is determined by the difference

$$\boldsymbol{s}_k(i)\boldsymbol{a}_k(j) - \boldsymbol{s}_k(j)\boldsymbol{a}_k(i),$$

where $\boldsymbol{s}_k = \begin{bmatrix} \boldsymbol{s}_k(1) & \boldsymbol{s}_k(2) & \cdots & \boldsymbol{s}_k(N) \end{bmatrix}^T \in \mathbb{R}^N$ is a vector of values for the N items associated with feature k, and $\boldsymbol{a}_k = \begin{bmatrix} \boldsymbol{a}_k(1) & \boldsymbol{a}_k(2) & \cdots & \boldsymbol{a}_k(N) \end{bmatrix}^T \in \mathbb{R}^N$ is a vector of weights for the N items associated with feature k. (In many practical scenarios, $\boldsymbol{s}_k$ and $\boldsymbol{a}_k$ should have nonnegative elements to make the factors interpretable.) The interactions of the "cross terms" in the expression $\boldsymbol{s}_k(i)\boldsymbol{a}_k(j) - \boldsymbol{s}_k(j)\boldsymbol{a}_k(i)$ allow $\boldsymbol{a}_k(j)$ to inhibit $\boldsymbol{s}_k(i)$ or $\boldsymbol{a}_k(i)$ to inhibit $\boldsymbol{s}_k(j)$. For the final pairwise comparison between item i and item j, we sum the comparisons based on the r factors:

$$\boldsymbol{Y}(i,j) = \sum_{k=1}^{r} \boldsymbol{s}_k(i)\boldsymbol{a}_k(j) - \boldsymbol{s}_k(j)\boldsymbol{a}_k(i).$$

This model has a simple expression in matrix form. Defining $\boldsymbol{Y}_k := \boldsymbol{s}_k\boldsymbol{a}_k^T - \boldsymbol{a}_k\boldsymbol{s}_k^T$, we have

$$\boldsymbol{Y} = \sum_{k=1}^{r} \boldsymbol{Y}_k = \sum_{k=1}^{r} \boldsymbol{s}_k\boldsymbol{a}_k^T - \boldsymbol{a}_k\boldsymbol{s}_k^T. \tag{5.1}$$

We give an example with $r = 1$, where non-transitive behavior occurs. The rank two skew-symmetric matrix is constructed by the following factors:

$$\begin{aligned}\boldsymbol{s} &= \begin{bmatrix}-0.5749 & 0.7154 & 1.8577 & 0.0780\end{bmatrix}^T \\ \boldsymbol{a} &= \begin{bmatrix}0.1 & 0.5 & 0.5 & 1\end{bmatrix}^T\end{aligned}$$

Then, we have

$$\begin{aligned}\boldsymbol{Y} &= \boldsymbol{s}\boldsymbol{a}^T - \boldsymbol{a}\boldsymbol{s}^T \\ &= \begin{bmatrix} 0 & -0.3590 & -0.4732 & -0.5827 \\ 0.3590 & 0 & -0.5712 & 0.6764 \\ 0.4732 & 0.5712 & 0 & 1.8187 \\ 0.5827 & -0.6764 & -1.8187 & 0 \end{bmatrix}\end{aligned}$$

From $\boldsymbol{Y}$, we have the following:

- $\boldsymbol{Y}(1,2) + \boldsymbol{Y}(2,3) = -0.9302 < -0.4732 = \boldsymbol{Y}(1,3)$
- $\boldsymbol{Y}(1,2) + \boldsymbol{Y}(2,4) = 0.3174 > -0.5827 = \boldsymbol{Y}(1,4)$
- $\boldsymbol{Y}(1,3) + \boldsymbol{Y}(3,4) = 1.3455 > -0.5827 = \boldsymbol{Y}(1,4)$
- $\boldsymbol{Y}(2,3) + \boldsymbol{Y}(3,4) = 1.2475 > 0.6764 = \boldsymbol{Y}(2,4)$

We can see that $\boldsymbol{Y}(1,2) + \boldsymbol{Y}(2,4)$ is positive while $\boldsymbol{Y}(1,4)$ is negative, implying that a certain type of non-transitivity happens. This is also the case for $\boldsymbol{Y}(1,3) + \boldsymbol{Y}(3,4)$ and $\boldsymbol{Y}(1,4)$. This toy example implies that the proposed model (5.1) can capture non-transitive behavior.

Our proposed model also has a number of possible applications in practical scenarios. Consider an example involving competitions between football teams. Suppose $\boldsymbol{s}_1 \in \mathbb{R}^N$ represents the strength of each team's offense (with high numbers being better), and suppose $\boldsymbol{a}_1 \in \mathbb{R}^N$ represents the strength of each team's defense (with low numbers being better). Then $\boldsymbol{s}_1(i)\boldsymbol{a}_1(j)$ could be a reasonable model for the number of points that team i is expected to score against team j, $\boldsymbol{s}_1(j)\boldsymbol{a}_1(i)$ could model the number of points that team j is

expected to score against team i, and the difference $\boldsymbol{s}_1(i)\boldsymbol{a}_1(j) - \boldsymbol{s}_1(j)\boldsymbol{a}_1(i)$ could model the expected margin of victory (or loss, if negative) of team i over team j. The incorporation of additional offensive factors $\boldsymbol{s}_2, \ldots, \boldsymbol{s}_r$ and defensive factors $\boldsymbol{a}_2, \ldots, \boldsymbol{a}_r$ could model the use of different player configurations, special teams (field goals and kickoffs), etc. Such models could be useful for modeling other scenarios where comparisons involve multiple features and interactions where features are inhibited or accentuated.

By inspection, one can deduce that the matrix $\boldsymbol{Y}$ defined in (5.1) is skew-symmetric and has rank at most $2r$. In fact, the property of skew-symmetry is intimately related to decompositions of the form appearing in (5.1). This is formalized in Lemma 5.15.

Lemma 5.15. *[65] Any skew-symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{N\times N}$ with rank at most $2r$ can be decomposed as $\boldsymbol{A} = \sum_{k=1}^{r} \boldsymbol{x}_k\boldsymbol{y}_k^T - \boldsymbol{y}_k\boldsymbol{x}_k^T$ for some vectors $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_r \in \mathbb{R}^N$ and $\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_r \in \mathbb{R}^N$.*

Furthermore, any skew-symmetric matrix $\boldsymbol{A}$ has the following decomposition.

Lemma 5.16. *[62] Suppose that $\boldsymbol{A} = -\boldsymbol{A}^T$ is an $N \times N$ real skew-symmetric matrix with rank $2r$. Then $\boldsymbol{A}$ can be factorized into the form*

$$\boldsymbol{A} = \boldsymbol{X} \underbrace{\begin{bmatrix} 0 & \lambda_1 & & & \\ -\lambda_1 & 0 & & & \\ & & 0 & \lambda_2 & \\ & & -\lambda_2 & 0 & \\ & & & & \ddots \end{bmatrix}}_{\boldsymbol{T}} \boldsymbol{X}^T,$$

where $\boldsymbol{X}$ is an orthonormal matrix that depends on $\boldsymbol{A}$, and $\boldsymbol{T}$ is a $2r \times 2r$ block diagonal matrix determined by $\lambda_1, \lambda_2, \ldots, \lambda_r$. Since $\boldsymbol{A}$ is a real skew-symmetric matrix, the eigenvalues of $\boldsymbol{A}$ are purely imaginary and are given by $\pm i\lambda_1, \pm i\lambda_2, \ldots, \pm i\lambda_r$, where i is the imaginary unit.

5.2.2 Connection to the transitive rank-two skew-symmetric model

Let $\boldsymbol{e} \in \mathbb{R}^N$ denote the vector with all entries equal to 1. We note that when all $\boldsymbol{a}_k = \boldsymbol{e}$ for all k, the pairwise comparison matrix in (5.1) takes the form $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$, where

$\boldsymbol{s} = \sum_{k=1}^{r} \boldsymbol{s}_k$. In this case, our model reduces to the rank-two skew-symmetric case studied in [62]. In the following, we build a connection between transitivity and ranking, showing that transitivity is necessary and sufficient for obtaining a consistent ranking.

Theorem 5.8. *The skew-symmetric matrix $\boldsymbol{Y} \in \mathbb{R}^{N\times N}$ is transitive, i.e., for an arbitrary triple (i,j,k), $\boldsymbol{Y}(i,j) + \boldsymbol{Y}(j,k) = \boldsymbol{Y}(i,k)$, if and only if $\boldsymbol{Y}$ is of the form $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$ for some score vector $\boldsymbol{s}$.*

Proof. Observe that when $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$, it is trivial to show that for an arbitrary triple (i,j,k), $\boldsymbol{Y}(i,j) + \boldsymbol{Y}(j,k) = \boldsymbol{Y}(i,k)$ always holds. So, we only show the other direction, i.e., when $\boldsymbol{Y}(i,j) + \boldsymbol{Y}(j,k) = \boldsymbol{Y}(i,k)$ for an arbitrary triple, $\boldsymbol{Y}$ must be of the form $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$. The proof can be accomplished by mathematical induction. To begin with, let us only consider three items i, j, k.

Step one: Show the 3×3 matrix formed by comparisons among items i, j, k can be written as $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$ for some $\boldsymbol{s}$. Assume that $\boldsymbol{Y}(i,j) = \boldsymbol{s}(i) - \boldsymbol{s}(j)$ and $\boldsymbol{Y}(j,k) = \boldsymbol{s}(j) - \boldsymbol{s}(k)$ for some $\boldsymbol{s}(i), \boldsymbol{s}(j)$, and $\boldsymbol{s}(k)$. This is possible because given the values of $\boldsymbol{Y}(i,j)$ and $\boldsymbol{Y}(j,k)$, we can always initialize a triple $\boldsymbol{s}(i), \boldsymbol{s}(j)$, and $\boldsymbol{s}(k)$. Since $\boldsymbol{Y}(i,k) = \boldsymbol{Y}(i,j) + \boldsymbol{Y}(j,k)$, we have

$$\boldsymbol{Y}(i,k) = \boldsymbol{s}(i) - \boldsymbol{s}(j) + \boldsymbol{s}(j) - \boldsymbol{s}(k) = \boldsymbol{s}(i) - \boldsymbol{s}(k).$$

Thus, this 3×3 skew-symmetric matrix can be written into the form $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$ with $\boldsymbol{s} = \begin{bmatrix} \boldsymbol{s}(i) & \boldsymbol{s}(j) & \boldsymbol{s}(k) \end{bmatrix}^T$.

Step two: Let ℓ be a new item that comes in the comparison. Suppose that we can find an $\boldsymbol{s}(\ell)$ such that $\boldsymbol{Y}(i,\ell) = \boldsymbol{s}(i) - \boldsymbol{s}(\ell)$. As $\boldsymbol{Y}(\ell,j) = \boldsymbol{Y}(\ell,i) + \boldsymbol{Y}(i,j)$, we have

$$\begin{aligned} \boldsymbol{Y}(\ell,j) &= -(\boldsymbol{s}(i) - \boldsymbol{s}(\ell)) + \boldsymbol{s}(i) - \boldsymbol{s}(j) \\ &= \boldsymbol{s}(\ell) - \boldsymbol{s}(j). \end{aligned}$$

Next, we can decide the value of $\boldsymbol{Y}(\ell,k)$ using transitivity

$$\begin{aligned} \boldsymbol{Y}(\ell,k) &= \boldsymbol{Y}(\ell,j) + \boldsymbol{Y}(j,k) \\ &= \boldsymbol{s}(\ell) - \boldsymbol{s}(j) + \boldsymbol{s}(j) - \boldsymbol{s}(k) \\ &= \boldsymbol{s}(\ell) - \boldsymbol{s}(k). \end{aligned}$$

Therefore, we conclude that $\boldsymbol{Y} = \boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$ for $\boldsymbol{s} = \begin{bmatrix} \boldsymbol{s}(i) & \boldsymbol{s}(j) & \boldsymbol{s}(k) & \boldsymbol{s}(\ell) \end{bmatrix}^T$.

Step three: we incorporate one more item sequentially until N items and show that, in each step, the enlarged skew-symmetric matrix can always be written into the form of $\boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T$ for some score vector $\boldsymbol{s}$.

This completes the proof. □

By introducing more degrees of freedom (via the $\boldsymbol{a}_k$) into the model, inhibitions and non-transitive behavior can be captured. In fact, this model allows for non-transitive matrices $\boldsymbol{Y}$ even in the case where $r = 1$ (by choosing $\boldsymbol{a}_1 \neq \boldsymbol{e}$). Therefore, our model generalizes the rank-two skew-symmetric case [62] both by allowing for non-transitivity and by allowing for higher rank.

The combinatorial Hodge theory developed in [63] allows us to quantify the degree of non-transitivity expressed in our model.

Lemma 5.17. *[63] Given an arbitrary skew-symmetric matrix $\boldsymbol{Y}$, denote*

$$R(\boldsymbol{s}) = \|\boldsymbol{Y} - (\boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T)\|_F$$

as the Frobenius norm of the difference between $\boldsymbol{Y}$ and a transitive rank-two skew-symmetric matrix induced by a score vector $\boldsymbol{s} \in \mathbb{R}^N$. Then, the minimum norm solution of the following optimization problem

$$\underset{\boldsymbol{s}}{\text{minimize}} \quad R(\boldsymbol{s}) \tag{5.2}$$

is given by $\widehat{\boldsymbol{s}} = \frac{1}{n}\boldsymbol{Y}\boldsymbol{e}$.

The proof of the lemma above is quite straightforward and can be found in [66].

To appreciate the implications of Lemma 5.17, consider a scenario where many triples (i, j, k) satisfy $|\boldsymbol{Y}(i, j) + \boldsymbol{Y}(j, k) + \boldsymbol{Y}(k, i)| > \gamma$ with γ being a positive number. In this case, $\boldsymbol{Y}$ is a pairwise comparison matrix with many non-transitive comparisons. In such a case, we might suspect that it is not appropriate to look for a score vector $\boldsymbol{s}$ that provides a

ranking consistent with the pairwise comparisons. In [63], the degree of non-transitivity in $\boldsymbol{Y}$ is characterized by the residual $R(\widehat{\boldsymbol{s}})$ of the least squares problem (5.2). Hence, examining $R(\widehat{\boldsymbol{s}})$ can reveal whether a score vector exists that will provide a consistent ranking. If $R(\widehat{\boldsymbol{s}})$ is small, the pairwise comparison matrix $\boldsymbol{Y}$ is less non-transitive and the extracted score vector $\widehat{\boldsymbol{s}}$ is plausible. In the extreme case where $R(\widehat{\boldsymbol{s}}) = 0$, the extracted score vector $\widehat{\boldsymbol{s}}$ perfectly accounts for the pairwise comparisons, in that $\boldsymbol{Y}(i,j) = \widehat{\boldsymbol{s}}(i) - \widehat{\boldsymbol{s}}(j)$. On the other hand, if $\|\boldsymbol{R}(\widehat{\boldsymbol{s}})\|_F$ is large, $\boldsymbol{Y}$ exhibits a high degree of non-transitivity, no score vector $\boldsymbol{s}$ exists that accurately explains the pairwise comparisons, and the transitive rank-two skew-symmetric model is not appropriate. Furthermore, we can quantify the amount of non-transitivity in the model using a metric discussed in [63]. The following theorem gives an upper bound for the non-transitive residual when $r = 1$ (the rank-two skew-symmetric case).

Theorem 5.9. *Given a rank-two skew-symmetric matrix $\boldsymbol{Y} = \boldsymbol{s}_1\boldsymbol{a}_1^T - \boldsymbol{a}_1\boldsymbol{s}_1^T$, denote $R(\boldsymbol{s}) = \|\boldsymbol{Y} - (\boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T)\|_F$ as the residual between $\boldsymbol{Y}$ and a transitive rank-two skew-symmetric matrix induced by a score vector $\boldsymbol{s} \in \mathbb{R}^N$. Let $\widehat{\boldsymbol{s}}$ be a minimizer of $R(\boldsymbol{s})$. Then*

$$R(\widehat{\boldsymbol{s}}) \leq 2 \left\| \boldsymbol{s}_1 \left(\mathcal{P}^{\perp}_{[\boldsymbol{s}_1 \ \boldsymbol{e}]} \boldsymbol{a}_1 \right)^T \right\|_F,$$

where $\mathcal{P}^{\perp}_{[\boldsymbol{s}_1 \ \boldsymbol{e}]}$ is the projection orthogonal to span$([\boldsymbol{s}_1 \ \boldsymbol{e}])$.

Proof. Recall that the non-transitive residual is defined as

$$R(\widetilde{\boldsymbol{s}}) = \|\boldsymbol{s}_1\boldsymbol{a}_1^T - \boldsymbol{a}_1\boldsymbol{s}_1^T - (\widetilde{\boldsymbol{s}}\boldsymbol{e}^T - \boldsymbol{e}\widetilde{\boldsymbol{s}}^T)\|_F. \tag{5.3}$$

Denote $\widehat{\boldsymbol{s}}$ as the optimal vector for minimizing $R(\widetilde{\boldsymbol{s}})$. According to Lemma 5.17, the optimal $\widehat{\boldsymbol{s}}$ is given by

$$\widehat{\boldsymbol{s}} = \frac{1}{N} \left(\boldsymbol{s}_1\boldsymbol{a}_1^T - \boldsymbol{a}_1\boldsymbol{s}_1^T \right) \boldsymbol{e} \tag{5.4}$$

Denote $\boldsymbol{U} = \mathbf{span}\left([\boldsymbol{s}_1, \boldsymbol{e}]\right)$ and $\mathcal{P}_{\boldsymbol{U}}, \mathcal{P}^{\perp}_{\boldsymbol{U}}$ as the orthogonal projection and orthogonal complement onto $\boldsymbol{U}$. Then, we have

$$\begin{aligned} \boldsymbol{a}_1 &= \mathcal{P}_{\boldsymbol{U}}\boldsymbol{a}_1 + \mathcal{P}^{\perp}_{\boldsymbol{U}}\boldsymbol{a}_1 \\ &= (\alpha\boldsymbol{s}_1 + \beta\boldsymbol{e}) + \mathcal{P}^{\perp}_{\boldsymbol{U}}\boldsymbol{a}_1 \end{aligned}$$

for some $\alpha, \beta \in \mathbb{R}$. Thus we have

$$\begin{aligned} R(\widehat{\boldsymbol{s}}) &\leq R(\beta \boldsymbol{s}_1) \\ &= \left\| \boldsymbol{s}_1 \left(\alpha \boldsymbol{s}_1 + \beta \boldsymbol{e} + \mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right)^T - \left(\alpha \boldsymbol{s}_1 + \beta \boldsymbol{e} + \mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right) \boldsymbol{s}_1^T - \left(\beta \boldsymbol{s}_1 \boldsymbol{e}^T - \boldsymbol{e} (\beta \boldsymbol{s}_1)^T \right) \right\|_F \\ &= \left\| \boldsymbol{s}_1 \left(\mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right)^T - \left(\mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right) \boldsymbol{s}_1^T \right\|_F \\ &\leq 2 \left\| \boldsymbol{s}_1 \left(\mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right)^T \right\|_F . \end{aligned}$$

This completes the proof. □

This bound shows that as long as $\left\| \boldsymbol{s}_1 \left(\mathcal{P}_{\boldsymbol{U}}^{\perp} \boldsymbol{a}_1 \right)^T \right\|_F$ is small, the residual $R(\widehat{\boldsymbol{s}})$ will be small.

Theorem 5.9 characterizes the non-transitive residual for the rank-two case ($r = 1$). The following theorem generalizes the argument to the more general rank-$2r$ case.

Theorem 5.10. *Let $\boldsymbol{Y} = \sum_{i=1}^{r} \boldsymbol{s}_i \boldsymbol{a}_i^T - \boldsymbol{a}_i \boldsymbol{s}_i^T$ be a rank-$2r$ skew-symmetric matrix. Denote $R(\boldsymbol{s}) = \| \boldsymbol{Y} - (\boldsymbol{s}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{s}^T) \|_F$ as the residual between $\boldsymbol{Y}$ and a transitive rank-two skew-symmetric matrix induced by a score vector $\boldsymbol{s} \in \mathbb{R}^N$. Let $\widehat{\boldsymbol{s}}$ be a minimizer of $R(\boldsymbol{s})$. Then*

$$R(\widehat{\boldsymbol{s}}) \leq 2 \sum_{i=1}^{r} \left\| \boldsymbol{s}_i \left(\mathcal{P}_{[\boldsymbol{s}_i \ \boldsymbol{e}]}^{\perp} \boldsymbol{a}_i \right)^T \right\|_F ,$$

where $\mathcal{P}_{[\boldsymbol{s}_i \ \boldsymbol{e}]}^{\perp}$ is the projection orthogonal to span($[\boldsymbol{s}_i \ \boldsymbol{e}]$).

For simplicity, we omit the proof of Theorem 5.10 as it can be easily adapted from the rank-two case.

5.3 Pairwise Comparison Matrix Completion

In this section, we propose two algorithms for pairwise comparison matrix completion based on alternating minimization. Both algorithms explicitly use the idea of matrix factorization techniques for efficient optimization. Numerical simulations are carried out to support the effectiveness of the proposed algorithms. Empirically, it is shown that pairwise comparison matrix completion is possible with a minimal number of observations.

Algorithm 1 Alternating Minimization for Pairwise Comparison Matrix Completion

input: partial pairwise comparison matrix $\mathcal{P}_\Omega(\boldsymbol{Y})$, number of latent factors r
initialize: $\boldsymbol{U}^0$: the top $2r$ left singular vectors of $\frac{1}{p}\mathcal{P}_\Omega(\boldsymbol{Y})$, where p is the ratio between the number of measurements and N^2
while: stopping criterion not met **do**

$$\textbf{solve: } \boldsymbol{V}^{k+1} = \underset{\boldsymbol{V}\in\mathbb{R}^{N\times 2r}}{\arg\min}\ ||\mathcal{P}_\Omega(\boldsymbol{Y} - \boldsymbol{U}^k\boldsymbol{V}^T)||_F^2$$

$$\boldsymbol{U}^{k+1} = \underset{\boldsymbol{U}\in\mathbb{R}^{N\times 2r}}{\arg\min}\ ||\mathcal{P}_\Omega(\boldsymbol{Y} - \boldsymbol{U}\boldsymbol{V}^{k+1^T})||_F^2$$

output: $\widehat{\boldsymbol{U}} = \boldsymbol{U}^k$, $\widehat{\boldsymbol{V}} = \boldsymbol{V}^k$

5.3.1 Alternating minimization

In many applications, pairwise comparison matrices may be known only partially. In order to infer the missing entries, we consider an approach involving alternating minimization [64]. Alternating minimization emerged as a useful tool for matrix completion during the Netflix Prize competition. However, only recently has it been shown that alternating minimization can have global convergence as long as certain conditions (e.g., a satisfactory initialization, a sufficient number of measurements, an incoherence condition) are met. Our aim is to solve

$$\min_{\boldsymbol{U},\boldsymbol{V}} ||\mathcal{P}_\Omega(\boldsymbol{Y} - \boldsymbol{U}\boldsymbol{V}^T)||_F^2, \tag{5.5}$$

where Ω is the observed index set, $\mathcal{P}_\Omega$ zeros out all entries of a matrix outside Ω, and $\boldsymbol{Y}$ is the true pairwise comparison matrix. Under the assumption that $\boldsymbol{Y}$ obeys the model (5.1) (and thus has rank at most $2r$), we consider $\boldsymbol{U} \in \mathbb{R}^{N\times 2r}$ and $\boldsymbol{V} \in \mathbb{R}^{N\times 2r}$. The alternating minimization algorithm for solving (5.5) is shown in Algorithm 1; see also [64]. The completed pairwise comparison matrix can be calculated via $\widehat{\boldsymbol{Y}} = \widehat{\boldsymbol{U}}\widehat{\boldsymbol{V}}^T$.

To guarantee the performance of Algorithm 1, we need to know the coherence of skew-symmetric matrices $\boldsymbol{Y}$ arising in our model (5.1). The coherence of a subspace $\boldsymbol{U}$ with dimension r is defined by

$$\mu(\boldsymbol{U}) = \frac{N}{r}\max_{1\le i\le N} ||\boldsymbol{P}_{\boldsymbol{U}}\boldsymbol{e}_i||^2, \tag{5.6}$$

where $\boldsymbol{P_U}$ denotes the orthogonal projection onto $\boldsymbol{U}$ and $\boldsymbol{e}_i$ is a vector of all zeros with a single 1 in position i. Suppose that $\boldsymbol{Y}$ has rank $2r$ exactly. Then $[\boldsymbol{s}_1, \boldsymbol{s}_2, \ldots, \boldsymbol{s}_r, \boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_r]$ must be linearly independent. Therefore, according to Lemma 5.16, there exists a matrix $\boldsymbol{X} = [\widetilde{\boldsymbol{s}}_1, \widetilde{\boldsymbol{s}}_2, \ldots, \widetilde{\boldsymbol{s}}_r, \widetilde{\boldsymbol{a}}_1, \widetilde{\boldsymbol{a}}_2, \ldots, \widetilde{\boldsymbol{a}}_r]$ with orthonormal columns such that $\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{T}\boldsymbol{X}^T$, and this $\boldsymbol{X}$ can be obtained by taking the SVD of $\boldsymbol{Y}$. The coherence μ of an arbitrary matrix is given by the maximum of the coherences of its row and column spaces; therefore, in our model $\boldsymbol{X}$ has the same coherence as the column spans of both $\boldsymbol{U_Y}$ and $\boldsymbol{V_Y}$. That is, $\boldsymbol{U_Y}$ and $\boldsymbol{V_Y}$ have coherence μ iff $\boldsymbol{X}$ has coherence μ. Based on [64], we provide a theoretical guarantee for the case $r = 1$; this extends a result appearing in [62].

Theorem 5.11. *Suppose that* $\boldsymbol{Y} = \lambda_1 \boldsymbol{s}_1 \boldsymbol{a}_1^T - \lambda_1 \boldsymbol{a}_1 \boldsymbol{s}_1^T$ *for some orthonormal vectors* $\boldsymbol{s}_1$ *and* $\boldsymbol{a}_1$ *with the matrix* $[\boldsymbol{s}_1, \boldsymbol{a}_1]$ *having coherence* μ*. Suppose that we sample* $m = O(\mu N \log^2 N)$ *observations of* $\boldsymbol{Y}$ *uniformly at random. Then Algorithm 1 converges geometrically to the true* $\boldsymbol{Y}$ *with high probability.*[13]

To accommodate the skew-symmetric nature of pairwise comparison matrices, we also propose a skew-symmetric variant of the alternating minimization algorithm:

$$\text{minimize}_{\boldsymbol{P},\boldsymbol{Q} \in \mathbb{R}^{N \times r}} \quad \|\mathcal{P}_\Omega(\boldsymbol{Y} - (\boldsymbol{P}\boldsymbol{Q}^T - \boldsymbol{Q}\boldsymbol{P}^T))\|_F^2. \tag{5.7}$$

Note that (5.7) can also be solved in an alternating fashion; note also that $\boldsymbol{P}$ and $\boldsymbol{Q}$ each have only r columns.

5.3.2 Numerical simulations

We provide simulations on synthetic data to demonstrate the effectiveness of the two proposed alternating minimization algorithms. We set $N = 100$ and consider the cases $r = 1$ and $r = 2$. We generate elements of $\boldsymbol{s}_k$ and $\boldsymbol{a}_k$ uniformly at random between 0 and 1. Then, we construct the pairwise comparison matrix according to (5.1). We reconstruct $\boldsymbol{Y}$ from various numbers of random observations. For each trial, we declare success if the recon-

[13] Technically, for this theorem to hold one must partition the observation set Ω and use a unique set of measurements during each iteration of Algorithm 1. See [64] for further details.

struction relative error is less than 10^{-3}. Figure 5.1 and Figure 5.2 show the reconstruction performance of the alternating minimization algorithms and the singular value projection (SVP) algorithm for $r = 1$ (rank 2) and $r = 2$ (rank 4), respectively. SVP is a matrix completion algorithm first proposed in [67] and adopted in [62] for pairwise comparison matrix completion. We note that theoretical guarantees do not exist for the performance of SVP.

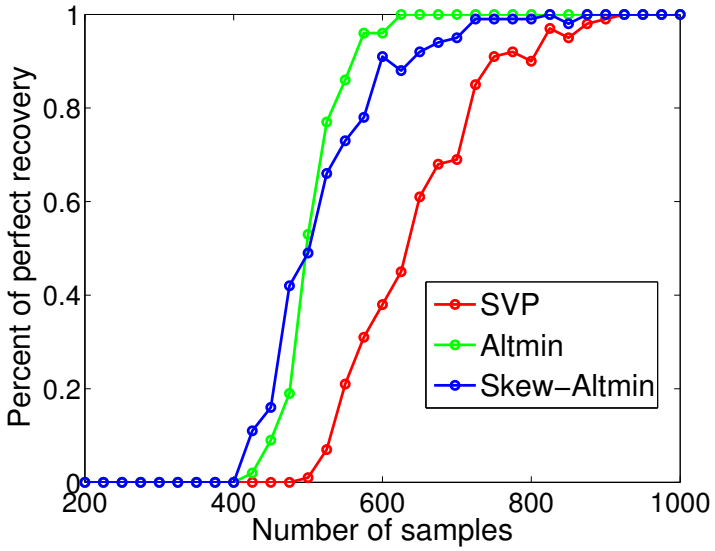

Figure 5.1: Recovery performance for $r = 1$. We generate elements of $\boldsymbol{s}_1 \in \mathbb{R}^{100\times 1}$ and $\boldsymbol{a}_1 \in \mathbb{R}^{100\times 1}$ uniformly at random between 0 and 1. Then, we construct the pairwise comparison matrix according to (5.1) and reconstruct $\boldsymbol{Y}$ from various numbers of random observations. For each trial, we declare success if the relative reconstruction error is less than 10^{-3}.

5.4 Concluding Remarks

In this chapter, we have developed a new low-rank model for pairwise comparison matrices. The proposed model generalized the simple rank-two transitive pairwise comparison matrices to non-transitive ones. Thus, it allows to capture inherent non-transitive behavior appeared in many real-world pairwise comparison datasets. Furthermore, we considered the scenario where only partial observations of the pairwise comparison are available and applied matrix completion techniques to infer the missing entries of the matrix. Fast alternating minimization algorithms are proposed to solve the problem and the performance of

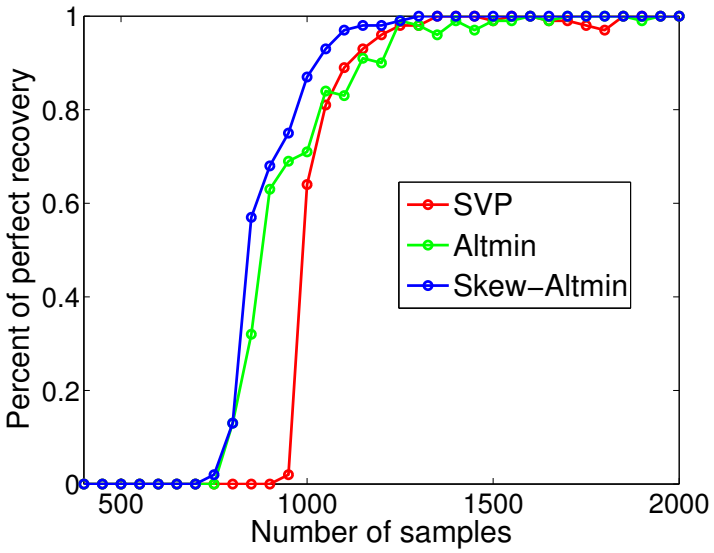

Figure 5.2: Recovery performance for $r = 2$. We generate elements of $\boldsymbol{s}_1, \boldsymbol{s}_2 \in \mathbb{R}^{100\times 1}$ and $\boldsymbol{a}_1, \boldsymbol{a}_2 \in \mathbb{R}^{100\times 1}$ uniformly at random between 0 and 1.

the algorithms is validated on simulated synthetic datasets. Future works include testing our proposed model in real-world datasets, and providing a theoretical guarantee for the alternating minimization algorithm that preserves skew-symmetry.

CHAPTER 6

RECOVERING POSITIVE SEMIDEFINITE MATRICES FROM QUADRATIC MEASUREMENTS

Many problems across different disciplines of engineering and science can be formulated as the recovery of a low-rank positive semidefinite (PSD) matrix from a set of linear measurements. Examples include phase retrieval in physical sciences, covariance estimation in statistics and data science, and quantum state tomography in computational physics. In this chapter[14], we consider the problem of the low-rank PSD matrix recovery from a set of structured linear measurements, the rank-one PSD measurements. We show that this problem can be reformulated as a flexible second-order cone program (SOCP) via a symmetric low-rank matrix factorization and an anchor matrix. We construct a valid anchor matrix from measurement data and sensing vectors. Theoretical insights for exact recovery of symmetric factors of the PSD matrix are carried out and numerical simulations are also conducted to support the effectiveness of our proposed algorithm.

6.1 Introduction

Consider the problem of estimating a rank-r positive semidefinite (PSD) matrix $\boldsymbol{X}^\star \in \mathbb{R}^{N\times N}$ from a set of quadratic measurements[15]

$$\boldsymbol{y}(i) = \boldsymbol{a}_i^T \boldsymbol{X}^\star \boldsymbol{a}_i, \quad i = 1, \ldots, M, \tag{6.1}$$

where $\left\{\boldsymbol{a}_i \in \mathbb{R}^N\right\}$ are sensing vectors. This problem arises in a range of applications across different disciplines of engineering and science. To motivate the problem, we list two stylized applications.

High dimensional covariance matrix recovery from sketches

[14]This ongoing work is in collaboration with Prof. Mike Wakin, and the manuscript is under preparation.

[15]For simplicity, in this chapter we only consider the real-valued scenario; the complex-valued case can be extended in a straightforward way.

Covariance matrices are most important and frequently encountered matrices in statistics and data science [68], and they are naturally positive semidefinite. Due to the curse of dimensionality in this big data era, estimating PSD covariance matrices is becoming a computationally expensive challenge. In many applications, however, the effective dimension of the PSD matrices is often much smaller than the ambient dimension of the matrices as the PSD matrices are usually low-rank. Therefore, equation (6.1) could be a handy design to obtain the sensing measurements of $\boldsymbol{X}^\star$ as each sample $\boldsymbol{y}(i)$ is simply a scalar [69, 70]. Then, the covariance matrix estimation problem becomes the inference of the low-rank PSD matrix $\boldsymbol{X}^\star$ from quadratic measurements $\boldsymbol{y}(i)$.

Phase retrieval

Phase retrieval refers to the process of inverting the phase of a signal $\boldsymbol{x}^\star$ from intensity only measurements

$$\boldsymbol{y}(i) = |\boldsymbol{a}_i^T \boldsymbol{x}^\star|^2. \tag{6.2}$$

It is a long-standing daunting inverse problem in physical sciences such as X-ray crystallography due to the nonlinearity involved in the measurement process [50, 71]. It is known that this nonlinear problem can be lifted to a linear inverse problem in a higher dimensional space [50, 72]. Specifically, with some mathematical manipulations, we can write

$$\begin{aligned} \boldsymbol{y}(i) &= |\boldsymbol{a}_i^T \boldsymbol{x}^\star|^2 \\ &= \boldsymbol{a}_i^T \boldsymbol{x}^\star \boldsymbol{x}^{\star T} \boldsymbol{a}_i. \end{aligned} \tag{6.3}$$

Then, equation (6.3) becomes a special case of (6.1) with $\boldsymbol{X}^\star = \boldsymbol{x}^\star \boldsymbol{x}^{\star T}$, a rank-one PSD matrix. This interpretation also provides us another way to view the sensing mechanism (6.3) because

$$\begin{aligned} \boldsymbol{y}(i) &= \mathrm{trace}(\boldsymbol{a}_i \boldsymbol{a}_i^T \boldsymbol{X}^\star) \\ &= \langle \boldsymbol{X}^\star, \boldsymbol{a}_i \boldsymbol{a}_i^T \rangle. \end{aligned} \tag{6.4}$$

This implies that each $\boldsymbol{y}(i)$ is a linear measurement of $\boldsymbol{X}^\star$ using the rank-one PSD sensing matrix $\boldsymbol{a}_i \boldsymbol{a}_i^T$. Once we get an estimate of $\boldsymbol{X}^\star$ from the measurements $\boldsymbol{y}(i)$, we obtain an estimate of $\boldsymbol{x}^\star$ via a simple matrix factorization, up to an inherent global phase ambiguity.

6.1.1 Related work

In recent years, many optimization-based works have been proposed in the research community to tackle the inverse problem (6.1). We review some of the works that are closely related to ours. Recent works in the literature can roughly be summarized into two categories: SDP-based convex relaxations and non-convex optimization algorithms.

In [50, 72], an SDP relaxation based approach is developed in the rank-one case for phase retrieval. Specifically, the following program is used for recovering $\boldsymbol{X}^\star = \boldsymbol{x}^\star \boldsymbol{x}^{\star T}$

$$\begin{aligned} \underset{\boldsymbol{X}}{\text{minimize}} \quad & \text{trace}(\boldsymbol{X}) \\ \text{subject to} \quad & \langle \boldsymbol{X}, \boldsymbol{a}_i \boldsymbol{a}_i^T \rangle = \boldsymbol{y}(i) \quad i = 1, \ldots, M, \\ & \boldsymbol{X} \succeq \boldsymbol{0}. \end{aligned} \tag{6.5}$$

It is shown that $M = O(N \log N)$ quadratic random Gaussian measurements suffice for exact recovery of the $N \times N$ matrix $\boldsymbol{X}^\star$ with high probability. Furthermore, a feasibility program is also proposed in the literature [73] and it is shown that, with high probability, solving

$$\begin{aligned} \text{find} \quad & \boldsymbol{X} \\ \text{subject to} \quad & \langle \boldsymbol{X}, \boldsymbol{a}_i \boldsymbol{a}_i^T \rangle = \boldsymbol{y}(i) \quad i = 1, \ldots, M, \\ & \boldsymbol{X} \succeq \boldsymbol{0}. \end{aligned} \tag{6.6}$$

is sufficient to recover $\boldsymbol{X}^\star$ exactly with $O(N)$ random Gaussian measurements [74], which is also the information-theoretic lower bound. This rank-one phase retrieval problem is further extended to rank-r case, allowing for the recovery of a low-rank PSD matrix from a minimal number of quadratic measurements [69, 70]. In particular, given a rank-r PSD matrix $\boldsymbol{X}^\star = \boldsymbol{U}^\star \boldsymbol{U}^{\star T}$, one can also deploy (6.5) directly for recovery and it is shown that with a random Gaussian measurement scheme, $O(rN)$ is sufficient for exact recovery. In addition, robust variants of program (6.5) are also proposed to deal with random Gaussian noise and arbitrary outliers in phase retrieval and low-rank PSD matrix recovery problems [75, 76].

While the SDP-based approaches are elegant and provide sample complexity that is proportional to the information-theoretical lower bound when the sensing vectors are random Gaussian, solving programs like (6.5) involves doing iterations over $N \times N$ matrices when

SDP-based numerical solvers are invoked. Thus, program (6.5) is limited to small-scale problems and is not efficient for real-world problems, such as high dimensional images, where N itself can be very large. This motivates the research for solving (6.1) in the factored domain directly, i.e., searching for $\boldsymbol{x}^\star$ in the rank-one case and $\boldsymbol{U}^\star$ in the rank-r case. In [50, 77], the authors propose a non-convex optimization scheme for phase retrieval and solve

$$\underset{\boldsymbol{x}}{\text{minimize}} \quad \frac{1}{2M}\sum_{i=1}^{2M}(\boldsymbol{y}(i) - |\boldsymbol{a}_i^T\boldsymbol{x}|^2)^2 \tag{6.7}$$

with a carefully chosen initialization and then via a gradient descent starting from the initialization. It is shown that even with a non-convex objective, $O(N)$ random Gaussian measurements are sufficient for exact recovery, up to a ploylogarithmic factor. Again, this rank-one non-convex scheme is extended to the rank-r case nontrivially, allowing for the recovery of a low-rank PSD matrix via a two-step non-convex optimization [78, 79]. Later on, it is shown [80] that (6.7) enjoys a benign geometry; the landscape of the object function in (6.7) does not involve any spurious local minima and all of the saddle points are ridable[16]. Therefore, careful initialization is not needed and gradient descent with random initialization is sufficient for converging to a global minimum with high probability [80, 81].

Recently, there is another line of research on using convex optimization for recovering the factor $\boldsymbol{x}^\star$ directly in phase retrieval [82, 83]. It solves an efficient linear program provided that there is an anchor vector $\boldsymbol{x}_0$ that is sufficiently positively correlated with the ground truth $\boldsymbol{x}^\star$. It is also shown that $O(N)$ random Gaussian measurements are sufficient for exact recovery, up to a polylogarithmic factor. Simplified analysis and a robust version of this simple approach are also proposed in the literature [84, 85]. We would like to note that this work is most closely related to our work presented in this chapter.

[16]The saddle points are called ridable if the Hessian matrices at the saddle points have at least one strictly negative eigenvalue.

6.1.2 Contributions

The main contribution of this chapter is the design of new efficient convex optimization algorithms for low-rank PSD matrix recovery using symmetric factors induced by a matrix factorization. We show that the factors of the low-rank PSD matrix can be recovered from quadratic measurements by solving a second-order cone program with a low computational complexity. The key to the success of the proposed method is the design of an anchor matrix that is positively correlated with the ground truth and the optimization objective that maximizes the correlation between the symmetric factor variable and the chosen anchor matrix. We show that such an anchor matrix can be constructed from measurement data and sensing vectors. Theoretical insights on conditions for exact recovery of the PSD matrices are also developed.

6.2 An SOCP Approach

Motivated by the seminal paper [82], we extend the PSD low-rank matrix recovery using an anchor vector from the rank-one case for phase retrieval to the general rank-r case. To begin with, our algorithm exploits the fact that any rank-r PSD matrix $\boldsymbol{X}^\star$ admits a symmetric low-rank matrix factorization

$$\boldsymbol{X}^\star = \boldsymbol{U}^\star \boldsymbol{U}^{\star T},$$

where $\boldsymbol{U}^\star \in \mathbb{R}^{N\times r}$ is a low-rank factor of $\boldsymbol{X}^\star$. Note that such a factorization is not unique and a global ambiguity exists. In particular, for any given $r \times r$ orthonormal matrix $\boldsymbol{O}$, $\boldsymbol{U}^\star\boldsymbol{O}$ is also a low-rank factor of $\boldsymbol{X}^\star$ as $\boldsymbol{X}^\star = \boldsymbol{U}^\star\boldsymbol{O}\boldsymbol{O}^T\boldsymbol{U}^{\star T} = \boldsymbol{U}^\star\boldsymbol{U}^{\star T}$. Now, assume that we have an initial estimate of $\boldsymbol{U}^\star$, denoted as $\boldsymbol{U}_0$, which is positively correlated with $\boldsymbol{U}^\star$, i.e., $\langle \boldsymbol{U}_0, \boldsymbol{U}^\star\rangle > 0$. Then, we propose the following program[17] to find $\boldsymbol{U}^\star$

$$\begin{aligned} \underset{\boldsymbol{U}}{\text{maximize}} \quad & \langle \boldsymbol{U}_0, \boldsymbol{U}\rangle \\ \text{subject to} \quad & \boldsymbol{a}_i^T\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{a}_i \le \boldsymbol{y}(i) \quad i = 1,\dots,M. \end{aligned} \tag{6.8}$$

[17]Note that we can only recover $\boldsymbol{U}^\star$, up to a $r \times r$ rotational ambiguity.

We can write (6.8) equivalently as

$$\begin{aligned} &\underset{\boldsymbol{U}}{\text{maximize}} \quad \langle \boldsymbol{U}_0, \boldsymbol{U} \rangle \\ &\text{subject to} \quad \|\boldsymbol{U}^T \boldsymbol{a}_i\|_2 \leq \sqrt{\boldsymbol{y}(i)} \quad i = 1, \ldots, M. \end{aligned} \tag{6.9}$$

The optimization program (6.9) is a convex program, in particular, a second-order cone program with the cone constraints of dimension $r+1$. The cone constraints in (6.9) can be related to a linear matrix inequality [86] and (6.9) is equivalent to

$$\begin{aligned} &\underset{\boldsymbol{U}}{\text{maximize}} \quad \langle \boldsymbol{U}_0, \boldsymbol{U} \rangle \\ &\begin{bmatrix} \sqrt{\boldsymbol{y}(i)} \boldsymbol{I}_{r\times r} & \boldsymbol{U}^T \boldsymbol{a}_i \\ \boldsymbol{a}_i^T \boldsymbol{U} & \sqrt{\boldsymbol{y}(i)} \end{bmatrix} \succeq \boldsymbol{0}, \quad i = 1, \ldots, M. \end{aligned} \tag{6.10}$$

Therefore, off-the-self solvers such as CVX [51] can be used to solve (6.10) efficiently. Note that the PSD matrices in the constraints above are of size $(r+1) \times (r+1)$, which is much smaller than that of $N \times N$ in the lifting-based SDP approaches. Therefore, program (6.10) is more computationally friendly compared to program (6.5).

6.3 Analysis

In this section, we provide strategies for obtaining a valid anchor matrix $\boldsymbol{U}_0$ from measurements $\boldsymbol{y}(i)$ and propose conditions under which solving (6.9) returns $\boldsymbol{U}^\star$ exactly, up to a rotational ambiguity. The proposed anchor matrix $\boldsymbol{U}_0$ is related to the initial estimate of $\boldsymbol{U}^\star$ in the non-convex optimization literature.

6.3.1 Anchor matrix

In this section, we propose a method to obtain an anchor matrix that is positively correlated with the ground truth $\boldsymbol{U}^\star$, i.e., $\langle \boldsymbol{U}_0, \boldsymbol{U}^\star \rangle > 0$. Consider the following matrix formed by measurement data and sensing vectors

$$\boldsymbol{I}_0 = \frac{1}{2M} \sum_{i=1}^{M} \boldsymbol{y}(i) \boldsymbol{a}_i \boldsymbol{a}_i^T.$$

From the construction, we see that $\boldsymbol{I}_0$ is a PSD matrix. Our construction of the anchor matrix is related to the eigen-decomposition of $\boldsymbol{I}_0$. Note that, in general, $\boldsymbol{I}_0$ can be full

rank. Assume that

$$\begin{aligned}\boldsymbol{I}_0 &= \boldsymbol{P}\boldsymbol{\Sigma}\boldsymbol{P}^T \\ &= \begin{bmatrix}\boldsymbol{p}_1 & \boldsymbol{p}_2 & \cdots & \boldsymbol{p}_N\end{bmatrix} \begin{bmatrix}\lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_N\end{bmatrix} \begin{bmatrix}\boldsymbol{p}_1^T \\ \boldsymbol{p}_2^T \\ \vdots \\ \boldsymbol{p}_N^T\end{bmatrix}\end{aligned} \tag{6.11}$$

is an eigen-decomposition of $\boldsymbol{I}_0$ with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N \geq 0$. We construct $\boldsymbol{U}_0$ as

$$\boldsymbol{U}_0 = \begin{bmatrix}\boldsymbol{p}_1 & \boldsymbol{p}_2 & \cdots & \boldsymbol{p}_r\end{bmatrix} \begin{bmatrix}\sqrt{\lambda_1 - \lambda_{r+1}} & & & \\ & \sqrt{\lambda_2 - \lambda_{r+1}} & & \\ & & \ddots & \\ & & & \sqrt{\lambda_r - \lambda_{r+1}}\end{bmatrix} := \boldsymbol{P}_r \boldsymbol{\Sigma}_r^{1/2}. \tag{6.12}$$

Note that this construction is first established in [79] and it is used as an initialization for a two-stage non-convex scheme for the recovery of the low-rank PSD matrix. When $\{\boldsymbol{a}_i\}$ are random vectors with i.i.d. Gaussian entries, we have the following theorem characterizing the relationship between $\boldsymbol{U}_0$ and $\boldsymbol{U}^\star$.

Lemma 6.18. *[79] Assume that the number of samples M satisfies $M \geq C\beta\lambda_r^{-4}\|\boldsymbol{U}^\star\|_F^8 r^2 N \log^2 N$, where $\lambda_1, \cdots, \lambda_r$ are defined in (6.11). Then, the initialization constructed using (6.12) satisfies*

$$d(\boldsymbol{U}_0) < \frac{9}{100\|\boldsymbol{U}^\star\|_F^2}\lambda_r^2, \tag{6.13}$$

with probability at least $1 - 3e^{-\beta r N} - \frac{7}{M^2}$. The term $d(\boldsymbol{U}_0)$ in (6.13) is defined as

$$d(\boldsymbol{U}_0) = \min_{\boldsymbol{O} \in \mathcal{O}(r)} \|\boldsymbol{U}^\star \boldsymbol{O} - \boldsymbol{U}_0\|_F^2 \tag{6.14}$$

with $\mathcal{O}(r)$ the set of all $r \times r$ orthonormal matrices.

We use the initialization $\boldsymbol{U}_0$ above for the design of the anchor matrix. Note that the sample complexity for obtaining such a anchor matrix is $O(r^2 N \log^2 N)$ when the Frobenius norm of the matrix is bounded. Based on (6.13), it is not hard to show that $\boldsymbol{U}_0$ is positively correlated with $\boldsymbol{U}^\star$.

6.3.2 Unique recovery condition

Assume that an anchor matrix $\boldsymbol{U}_0$ that satisfies the condition $\langle \boldsymbol{U}_0, \boldsymbol{U}^\star \rangle > 0$ is already given. We would like to derive conditions under which solving the computationally efficient SOCP program (6.10) recovers $\boldsymbol{U}^\star$ exactly. Our analysis strategy follows from the recent work [84] and we have the following theorem.

Theorem 6.12. *A sufficient condition for perfect recovery of $\boldsymbol{U}^\star$ using (6.10) (up to a global rotation) is*

$$\langle \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T + \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T}, \boldsymbol{a}_i \boldsymbol{a}_i^T \rangle \leq 0, \quad \Rightarrow \quad \langle \boldsymbol{U}_0, \boldsymbol{\Delta} \rangle < 0,$$

where $\Delta \in \mathbb{R}^{N \times r}$ is an arbitrary matrix and $\boldsymbol{O} \in \mathcal{O}(r)$.

Proof. Assume that $\widetilde{\boldsymbol{U}}$ is a solution to program (6.10), which is different from $\boldsymbol{U}^\star$ within a rotational ambiguity. Let us denote $\boldsymbol{\Delta} = \widetilde{\boldsymbol{U}} - \boldsymbol{U}^\star \boldsymbol{O}$, where $\boldsymbol{O}$ is defined such that the distance metric (6.14) is minimized. Due to the optimality of $\widetilde{\boldsymbol{U}}$, it must hold that

$$\langle \boldsymbol{U}_0, \widetilde{\boldsymbol{U}} \rangle \geq \langle \boldsymbol{U}_0, \boldsymbol{U}^\star \boldsymbol{O} \rangle,$$

which is equivalent to $\langle \boldsymbol{U}_0, \boldsymbol{\Delta} \rangle \geq 0$. On the other hand, since $\widetilde{\boldsymbol{U}}$ is a feasible solution, we have, for $i = 1, \ldots, M$, $\|\boldsymbol{a}_i^T \widetilde{\boldsymbol{U}}\|_2^2 \leq \boldsymbol{y}(i)$. Furthermore,

$$\begin{aligned} \|\boldsymbol{a}_i^T (\boldsymbol{U}^\star \boldsymbol{O} + \boldsymbol{\Delta})\|_2^2 &= \boldsymbol{a}_i^T (\boldsymbol{U}^\star \boldsymbol{O} + \boldsymbol{\Delta})(\boldsymbol{U}^\star \boldsymbol{O} + \boldsymbol{\Delta})^T \boldsymbol{a}_i \\ &= \boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{\Delta}^T \boldsymbol{a}_i \\ &= \boldsymbol{y}(i) + \boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{\Delta}^T \boldsymbol{a}_i, \end{aligned}$$

which implies that the term $\boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{\Delta}^T \boldsymbol{a}_i$ must be nonpositive due to the fact that $\|\boldsymbol{a}_i^T \widetilde{\boldsymbol{U}}\|_2^2 \leq \boldsymbol{y}(i)$. We further simplify the terms as

$$\begin{aligned} & \boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{\Delta}^T \boldsymbol{a}_i \\ = \; & \boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i + \|\boldsymbol{a}_i^T \boldsymbol{\Delta}\|_2^2. \end{aligned}$$

In order to maintain the optimality of $\widetilde{\boldsymbol{U}}$, we must have

$$\boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i < -\|\boldsymbol{a}_i^T \boldsymbol{\Delta}\|_2^2.$$

Therefore, a sufficient condition for the optimality of $\boldsymbol{U}^\star$ is

$$\boldsymbol{a}_i^T \boldsymbol{U}^\star \boldsymbol{O} \boldsymbol{\Delta}^T \boldsymbol{a}_i + \boldsymbol{a}_i^T \boldsymbol{\Delta} \boldsymbol{O}^T \boldsymbol{U}^{\star T} \boldsymbol{a}_i < 0, \quad \Rightarrow \quad \langle \boldsymbol{U}_0, \boldsymbol{\Delta} \rangle < 0.$$

□

6.3.3 Conjecture

So far, we only showed a sufficient condition under which $\boldsymbol{U}^\star$ is the unique solution to (6.10), up to a rotational ambiguity. Further theoretical investigation remains to be done. An interesting direction is to characterize the sample complexity, i.e., the number of samples M, such that Theorem 6.12 holds and exact recovery, up to a rotational ambiguity, is guaranteed. We conjecture that, given a valid anchor matrix $\boldsymbol{U}_0$ with $\langle \boldsymbol{U}_0, \boldsymbol{U}^\star \rangle > 0$, when the sensing vectors $\{\boldsymbol{a}_i\}$ are populated with random vectors with i.i.d. Gaussian entries, and the number of quadratic measurements is on the order of $O(r^2 N)$ up to a polylogarithmic factor, $\boldsymbol{U}^\star \boldsymbol{O}$ is the unique solution to (6.10) for some orthonormal matrix $\boldsymbol{O}$.

6.4 Simulations

We conduct simulations to verify the effectiveness of the proposed SOCP. We solve all instances of program (6.10) using CVX. In the first experiment, we would like to study the reconstruction performance and phase transition of the proposed method. To do this, we set $N = 50$, $r = 2$, and construct $\boldsymbol{U}^\star$ with zero mean and unit variance random Gaussian entries. Then, the rank-r PSD matrix $\boldsymbol{X}^\star$ is constructed as $\boldsymbol{X}^\star = \boldsymbol{U}^\star \boldsymbol{U}^{\star T}$. Note that in this experiment, the number of degrees of freedom in the PSD matrix $\boldsymbol{X}^\star$ is roughly equal to $Nr = 100$. Each quadratic measurement $\boldsymbol{y}(i)$ is constructed using random sensing vectors $\{\boldsymbol{a}_i\}$ with zero mean and unit variance Gaussian entries. The number of measurements M ranges from 100 to 1200. The anchor matrix $\boldsymbol{U}_0$ is constructed according to (6.12) and the reconstruction is performed by solving (6.10). We run 20 trials for each value of M and compute the relative recovery errors $\frac{\|\widetilde{\boldsymbol{U}}\widetilde{\boldsymbol{U}}^T - \boldsymbol{U}^\star \boldsymbol{U}^{\star T}\|_F}{\|\boldsymbol{U}^\star \boldsymbol{U}^{\star T}\|_F}$. Figure 6.1 shows the median (among 20 trials) of the relative recovery errors for different oversampling ratios $\frac{M}{Nr}$. From Figure 6.1, we can see clearly that when the sampling ratio is larger than 4, the median of the relative

construction errors drops significantly. Figure 6.2 shows the percentage of the successful recovery among 20 trials as the sampling ratio $\frac{M}{N}$ changes. For each trial, we declare success if the relative reconstruction Frobenius norm error $\frac{\|\widetilde{\boldsymbol{U}}\widetilde{\boldsymbol{U}}^T - \boldsymbol{U}^\star\boldsymbol{U}^{\star T}\|_F}{\|\boldsymbol{U}^\star\boldsymbol{U}^{\star T}\|_F}$ is less than 10^{-4}. From Figure 6.2, we observe that when the oversampling ratio is larger than 7, every trial leads to a successful recovery.

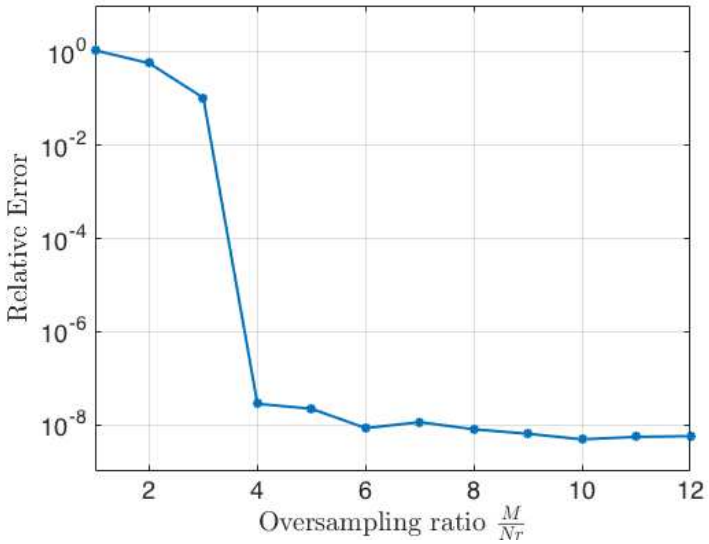

Figure 6.1: Empirical relative recovery error $\frac{\|\widetilde{\boldsymbol{U}}\widetilde{\boldsymbol{U}}^T - \boldsymbol{U}^\star\boldsymbol{U}^{\star T}\|_F}{\|\boldsymbol{U}^\star\boldsymbol{U}^{\star T}\|_F}$ by solving the second-order cone program (6.10) with different sampling ratios $\frac{M}{Nr}$ when $N = 50, r = 2$.

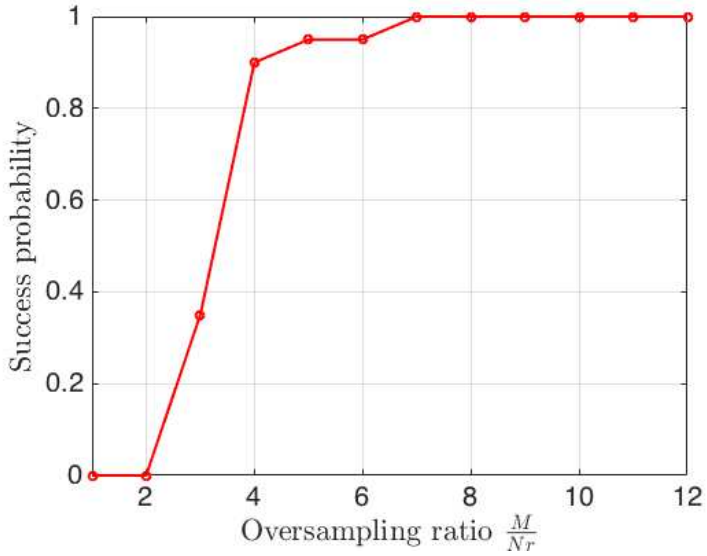

Figure 6.2: Success probability of solving the second-order cone program (6.10) with different oversampling ratios $\frac{M}{Nr}$ when $N = 50, r = 2$.

CHAPTER 7

CONCLUSION

Throughout this dissertation, we have developed new theories and algorithms for structured low-rank matrix recovery from linear measurements. Many future research questions remain and are worth further exploration.

First of all, for all of the three instances of structured low-rank matrix recovery problems presented in this dissertation, we only consider scenarios where the measurements are noiseless, an ideal setting which does not hold in practice. Therefore, additional efforts are needed to address the issues where the measurements are contaminated with various types of noise. For example, the measurements may often be contaminated with Gaussian noise. Furthermore, in many applications, some measurements can be completely faulty and data can be severely distorted, a type of corruption known as outliers. Convex programs have been shown to be robust to outliers in many related problems in the literature. We would like to develop new convex programs that provide robust guarantees in the presence of outliers as well.

More broadly, this dissertation centers around capturing low-dimensional models (low-rank models, in particular) for regularizing highly ill-posed inverse problems and solving the corresponding problems using optimization methods. We would like to expand the research further under this central theme. One interesting direction would be to find new inverse problems in physical and data sciences where low-rank models occur and play a vital role. For example, in subspace clustering, the low-rank model arises as data points within the same subspace obey a low-rank model. This low-rank structure has been used for low-rank representation within subspaces, and subsequently, for subspace clustering. However, one open question remains, which is subspace clustering with missing data, where each high dimensional data point is only observed partially. Therefore, new theories and algorithmic

frameworks are needed to address the missing data challenges in subspace clustering. We plan to exploit the idea of low-rank representation for subspace clustering further and extend it to the missing data regime with proper reformulations of available convex programs. Subspace clustering is just one example, where structured low-rankness can potentially be exploited. We believe that structured low-rankness can be found in many other inverse problems in engineering beyond the ones presented in this dissertation.

Another interesting research direction is the development of new efficient algorithmic frameworks that are scalable to problems with large-scale datasets. While convex programs are easy to be optimized to the global minimum and provide robustness to perturbations, they are not efficient for large-scale problems. Recent advances in non-convex optimization show promising directions by harnessing low-rank matrix factorizations to derive non-convex schemes for global optimization in the factored domain. We would like to pursue this line of research further and propose non-convex optimization schemes for structured low-rank matrix recovery in many inverse problems across engineering and science.

REFERENCES CITED

[1] R. G. Baraniuk. Compressive sensing [lecture notes]. *IEEE signal processing magazine*, 24(4):118–121, 2007.

[2] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

[3] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.

[4] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 1990.

[5] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[6] Y. Koren, R. Bell, C. Volinsky, et al. Matrix factorization techniques for recommender systems. *IEEE Trans. on Computers*, 42(8):30–37, 2009.

[7] A. M. Bruckstein, D. L. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM review*, 51(1):34–81, 2009.

[8] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. on Information Theory*, 56(4):1982–2001, 2010.

[9] S. Sarvotham, D. Baron, M. Wakin, M. F. Duarte, and R. G. Baraniuk. Distributed compressed sensing of jointly sparse signals. In *Asilomar conference on signals, systems, and computers*, pages 1537–1541, 2005.

[10] S. Oymak, A. Jalali, M. Fazel, Y. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Trans. Information Theory*, 61(5):2886–2908, 2015.

[11] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[12] G. Tang, B. Bhaskar, P. Shah, and B. Recht. Compressed sensing off the grid. *IEEE Trans. Information Theory*, 59(11):7465–7490, 2013.

[13] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.

[14] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[15] E. Candès and B. Recht. Simple bounds for recovering low-complexity models. *Mathematical Programming*, pages 1–13, 2013.

[16] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(Dec):3413–3430, 2011.

[17] Y. Chen. Incoherence-optimal matrix completion. *IEEE Trans. on Information Theory*, 61(5):2909–2923, 2015.

[18] D. Yang, G. Tang, and M. B. Wakin. Super-resolution of complex exponentials from modulations with unknown waveforms. *IEEE Trans. on Information Theory*, 62(10): 5809–5830, 2016.

[19] D. Yang, G. Tang, and M. Wakin. Non-stationary blind super-resolution. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4727–4731, Shanghai, China, March 2016.

[20] B. Huang, W. Wang, M. Bates, and X. Zhuang. Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy. *Science*, 319(5864):810–813, 2008.

[21] A. Yildiz, J. Forkey, S. McKinney, T. Ha, Y. Goldman, and P. Selvin. Myosin V walks hand-over-hand: Single fluorophore imaging with 1.5-nm localization. *Science*, 300(5628):2061–2065, 2003.

[22] G. Margrave, M. Lamoureux, and D. Henley. Gabor deconvolution: Estimating reflectivity by nonstationary deconvolution of seismic data. *Geophysics*, 76(3):W15–W30, 2011.

[23] S. Quirin, S. Pavani, and R. Piestun. Optimal 3D single-molecule localization for superresolution microscopy with aberrations and engineered point spread functions. *Proceedings of the National Academy of Sciences*, 109(3):675–679, 2012.

[24] R. Fergus, B. Singh, A. Hertzmann, S. Roweis, and W. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graphics*, 25(3):787–794, 2006.

[25] J. Starck, E. Pantin, and F. Murtagh. Deconvolution in astronomy: A review. *Publications of the Astronomical Society of the Pacific*, 114(800):1051–1069, 2002.

[26] X. Luo and G. Giannakis. Low-complexity blind synchronization and demodulation for (ultra-) wideband multi-user ad hoc access. *IEEE Trans. Wireless Communications*, 5 (7):1930–1941, 2006.

[27] R. Heckel, V. Morgenshtern, and M. Soltanolkotabi. Super-resolution radar. *Information and Inference: A Journal of the IMA*, 2016, to appear.

[28] J. Cai, X. Qu, W. Xu, and G. Ye. Robust recovery of complex exponential signals from random Gaussian projections via low-rank Hankel matrix reconstruction. *Applied and Computational Harmonic Analysis*, 2016, to appear.

[29] E. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.

[30] B. Bhaskar, G. Tang, and B. Recht. Atomic norm denoising with applications to line spectral estimation. *IEEE Trans. Signal Processing*, 61(23):5987–5999, 2013.

[31] E. Candès and C. Fernandez-Granda. Super-resolution from noisy data. *Journal of Fourier Analysis and Applications*, 19(6):1229–1254, 2013.

[32] G. Tang, B. Bhaskar, and B. Recht. Near minimax line spectral estimation. *IEEE Trans. Information Theory*, 61(1):499–512, 2015.

[33] V. Morgenshtern and E. Candès. Super-resolution of positive sources: The discrete setup. *SIAM Journal on Imaging Sciences*, 9(1):412–444, 2016.

[34] V. Duval and G. Peyré. Exact support recovery for sparse spikes deconvolution. *Foundations of Computational Mathematics*, 15(5):1315–1355, 2015.

[35] G. Schiebinger, E. Robeva, and B. Recht. Superresolution without separation. *arXiv preprint arXiv:1506.03144*, 2015.

[36] C. Fernandez-Granda. Super-resolution of point sources via convex programming. *Information and Inference: A Journal of the IMA*, 2016, to appear.

[37] Y. Li and Y. Chi. Off-the-grid line spectrum denoising and estimation with multiple measurement vectors. *IEEE Trans. Signal Processing*, 64(5):1257–1269, 2016.

[38] Z. Yang and L. Xie. Exact joint sparse frequency recovery via optimization methods. *arXiv preprint arXiv:1405.6585*, 2014.

[39] Y. Chen and Y. Chi. Robust spectral compressed sensing via structured matrix completion. *IEEE Trans. Information Theory*, 60(10):6576–6601, 2014.

[40] A. Ahmed, B. Recht, and J. Romberg. Blind deconvolution using convex programming. *IEEE Trans. Information Theory*, 60(3):1711–1732, 2014.

[41] A. Ahmed, A. Cosse, and L. Demanet. A convex approach to blind deconvolution with diverse input. *Preprint*, 2015.

[42] S. Ling and T. Strohmer. Blind deconvolution meets blind demixing: Algorithms and performance bounds. *arXiv preprint arXiv:1512.07730*, 2015.

[43] K. Lee, Y. Wu, and Y. Bresler. Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization. *arXiv preprint arXiv:1312.0525*, 2013.

[44] G. Tang and B. Recht. Convex blind deconvolution with random masks. In *Computational Optical Sensing and Imaging (COSI)*, 2014.

[45] S. Bahmani and J. Romberg. Lifting for blind deconvolution in random mask imaging: Identifiability and convex relaxation. *SIAM Journal on Imaging Sciences*, 8(4):2203–2238, 2015.

[46] Y. Li, K. Lee, and Y. Bresler. Identifiability in blind deconvolution with subspace or sparsity constraints. *IEEE Trans. Information Theory*, 62(7):4266–4275, 2016.

[47] S. Choudhary and U. Mitra. Identifiability scaling laws in bilinear inverse problems. *arXiv preprint arXiv:1402.2637*, 2014.

[48] Y. Chi. Guaranteed blind sparse spikes deconvolution via lifting and convex optimization. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):782–794, 2016.

[49] S. Ling and T. Strohmer. Self-calibration and biconvex compressive sensing. *Inverse Problems*, 31(11):115002, 2015.

[50] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Trans. on Information Theory*, 61(4):1985–2007, 2015.

[51] M. Grant, S. Boyd, and Y. Ye. CVX: Matlab software for disciplined convex programming, 2008.

[52] E. Candès and Y. Plan. A probabilistic and RIPless theory of compressed sensing. *IEEE Trans. Information Theory*, 57(11):7235–7254, 2011.

[53] S. Li, D. Yang, G. Tang, and M. B. Wakin. Atomic norm minimization for modal analysis from random and compressed samples. *Accepted, IEEE Trans. on Signal Processing*, 2017.

[54] S. Li, D. Yang, and M. B. Wakin. Atomic norm minimization for modal analysis with random spatial compression. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3251–3255, New Orleans, Louisiana, March 2017.

[55] S. Li, D. Yang, and M. B. Wakin. Atomic norm minimization for modal analysis. In *Proceedings of IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–2, Seattle, Washington, June 2016.

[56] D. Balageas, C. Fritzen, and A. Güemes. *Structural health monitoring*, volume 90. John Wiley & Sons, 2010.

[57] H. Sohn, C. R. Farrar, F. M. Hemez, and J. J. Czarnecki. A review of structural health review of structural health monitoring literature 1996-2001. Technical report, Los Alamos National Laboratory, 2002.

[58] S. W. Doebling, C. R. Farrar, M. B. Prime, and D. W. Shevitz. Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: a literature review. 1996.

[59] J. P. Lynch. An overview of wireless structural health monitoring for civil structures. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 365(1851):345–372, 2007.

[60] J. Park, M. B. Wakin, and A. C. Gilbert. Modal analysis with compressive measurements. *IEEE Trans. Signal Processing*, 62(7):1655–1670, 2014.

[61] D. Yang and M. B. Wakin. Modeling and recovering non-transitive pairwise comparison matrices. In *Proceedings of IEEE International Conference on Sampling Theory and Applications (SampTA)*, pages 39–43, Washington, DC, May, 2015.

[62] D. F. Gleich and L.-H. Lim. Rank aggregation via nuclear norm minimization. In *Proc. ACM SIGKDD*, pages 60–68, 2011.

[63] X. Jiang, L.H. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, pages 1–42, 2010.

[64] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proc. ACM STOC*, pages 665–674, 2013.

[65] R. A. Brualdi and IMA-ISU research group on minimum rank. The minimum rank of skew-symmetric matrices described by a graph. *Linear Algebra and Its Applications*, 432:2457–2472, 2010.

[66] A. N. Langville and C. D. Meyer. *Who's# 1?: the science of rating and ranking*. Princeton University Press, 2012.

[67] P. Jain, R. Meka, and I. S. Dhillon. Guaranteed rank minimization via singular value projection. In *Proc. NIPS*, pages 937–945, 2010.

[68] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.

[69] Y. Chen, Y. Chi, and A. J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Trans. on Information Theory*, 61 (7):4034–4059, 2015.

[70] R. Kueng, H. Rauhut, and U. Terstiege. Low rank matrix recovery from rank one measurements. *Applied and Computational Harmonic Analysis*, 42(1):88–116, 2017.

[71] R. P. Millane. Phase retrieval in crystallography and optics. *JOSA A*, 7(3):394–411, 1990.

[72] E. Candès, T. Strohmer, and V. Voroninski. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

[73] L. Demanet and P. Hand. Stable optimizationless recovery from phaseless linear measurements. *Journal of Fourier Analysis and Applications*, 20(1):199–221, 2014.

[74] E. Candès and X. Li. Solving quadratic equations via PhaseLift when there are about as many equations as unknowns. *Foundations of Computational Mathematics*, 14(5): 1017–1026, 2014.

[75] P. Hand. Phaselift is robust to a constant fraction of arbitrary errors. *Applied and Computational Harmonic Analysis*, 42(3):550–562, 2017.

[76] Y. Li, Y. Sun, and Y. Chi. Low-rank positive semidefinite matrix recovery from quadratic measurements with outliers. *arXiv preprint*, 2016.

[77] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. In *Advances in Neural Information Processing Systems*, pages 2796–2804, 2013.

[78] Q. Zheng and J. Lafferty. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In *Advances in Neural Information Processing Systems*, pages 109–117, 2015.

[79] S. Sanghavi, R. Ward, and C. White. The local convexity of solving systems of quadratic equations. *Results in Mathematics*, 71(3-4):569–608, 2017.

[80] J. Sun, Q. Qu, and J. Wright. A geometric analysis of phase retrieval. *arXiv preprint arXiv:1602.06664*, 2016.

[81] J. Lee, M. Simchowitz, M. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pages 1246–1257, 2016.

[82] S. Bahmani and J. Romberg. Phase retrieval meets statistical learning theory: A flexible convex relaxation. *arXiv preprint arXiv:1610.04210*, 2016.

[83] T. Goldstein and C. Studer. Phasemax: Convex phase retrieval via basis pursuit. *arXiv preprint arXiv:1610.07531*, 2016.

[84] P. Hand and V. Voroninski. An elementary proof of convex phase retrieval in the natural parameter space via the linear program phasemax. *arXiv preprint arXiv:1611.03935*, 2016.

[85] P. Hand and V. Voroninski. Corruption robust phase retrieval via linear programming. *arXiv preprint arXiv:1612.03547*, 2016.

[86] M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear algebra and its applications*, 284(1-3):193–228, 1998.

[87] J. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[88] A. Schaeffer. Inequalities of A. Markoff and S. Bernstein for polynomials and related functions. *Bulletin of the American Mathematical Society*, 47(8):565–579, 1941.

APPENDIX A

PROOF OF PROPOSITION 3.1

The proof strategy follows quite straightforwardly from that in Proposition II.4 of [12]. First of all, any dual vector $\boldsymbol{\lambda}$ satisfying (3.12) and (3.13) in Proposition 3.1 is dual feasible. To see this, note that

$$\begin{aligned}
\|\mathcal{B}^*(\boldsymbol{\lambda})\|_{\mathcal{A}}^* &= \sup_{\|\boldsymbol{X}\|_{\mathcal{A}}\leq 1} \langle \mathcal{B}^*(\boldsymbol{\lambda}), \boldsymbol{X}\rangle_{\mathbb{R}} \\
&= \sup_{\tau\in[0,1),\|\boldsymbol{h}\|_2=1} \langle \mathcal{B}^*(\boldsymbol{\lambda}), \boldsymbol{h}\boldsymbol{a}(\tau)^H\rangle_{\mathbb{R}} \\
&= \sup_{\tau\in[0,1),\|\boldsymbol{h}\|_2=1} \left\langle \sum_n \boldsymbol{\lambda}(n)\boldsymbol{b}_n\boldsymbol{e}_n^H, \boldsymbol{h}\boldsymbol{a}(\tau)^H \right\rangle_{\mathbb{R}} \\
&= \sup_{\tau\in[0,1),\|\boldsymbol{h}\|_2=1} \mathrm{Re}(\boldsymbol{h}^H\boldsymbol{q}(\tau)) \\
&\leq \sup_{\tau\in[0,1)} \|\boldsymbol{q}(\tau)\|_2 \\
&\leq 1.
\end{aligned}$$

Here, the second equality comes from the fact that the atoms $\left\{\boldsymbol{h}\boldsymbol{a}(\tau)^H\right\}$ comprise all of the extremal points of the atomic unit ball $\{\boldsymbol{X} : \|\boldsymbol{X}\|_{\mathcal{A}} \leq 1\}$. Furthermore, for any $\boldsymbol{\lambda}$ that satisfies (3.12) in Proposition 3.1, we have

$$\begin{aligned}
\langle \boldsymbol{\lambda}, \boldsymbol{y}\rangle_{\mathbb{R}} &= \langle \mathcal{B}^*(\boldsymbol{\lambda}), \boldsymbol{X}_o\rangle_{\mathbb{R}} \\
&= \left\langle \mathcal{B}^*(\boldsymbol{\lambda}), \sum_{j=1}^J c_j\boldsymbol{h}_j\boldsymbol{a}(\tau_j)^H \right\rangle_{\mathbb{R}} \\
&= \sum_{j=1}^J \mathrm{Re}\left(c_j^* \ \mathrm{trace}\left(\boldsymbol{a}(\tau_j)\boldsymbol{h}_j^H\mathcal{B}^*(\boldsymbol{\lambda})\right)\right) \\
&= \sum_{j=1}^J \mathrm{Re}\left(c_j^*\boldsymbol{h}_j^H\boldsymbol{q}(\tau_j)\right) \\
&= \sum_{j=1}^J \mathrm{Re}\left(c_j^*\mathrm{sign}(c_j)\right) \\
&= \sum_{j=1}^J |c_j| \geq \|\boldsymbol{X}_o\|_{\mathcal{A}}
\end{aligned}$$

where the fifth line follows from (3.12) in Proposition 3.1, and the last line follows from the definition of the atomic norm. On the other hand, by Hölder's inequality, we have

$$\begin{aligned}\langle \boldsymbol{\lambda}, \boldsymbol{y} \rangle_{\mathbb{R}} &= \langle \mathcal{B}^*(\boldsymbol{\lambda}), \boldsymbol{X}_o \rangle_{\mathbb{R}} \\ &\leq \|\mathcal{B}^*(\boldsymbol{\lambda})\|_{\mathcal{A}}^* \, \|\boldsymbol{X}_o\|_{\mathcal{A}} \\ &\leq \|\boldsymbol{X}_o\|_{\mathcal{A}} .\end{aligned}$$

This implies that any $\boldsymbol{\lambda}$ satisfying (3.12) also satisfies $\langle \boldsymbol{\lambda}, \boldsymbol{y} \rangle_{\mathbb{R}} = \|\boldsymbol{X}_o\|_{\mathcal{A}}$. Zero duality gap means that $\boldsymbol{\lambda}$ is dual optimal and $\boldsymbol{X}_o$ is primal optimal.

Finally, condition (3.13) ensures that $\boldsymbol{X}_o$ is the unique optimal solution. To see this, suppose that there exists another optimal solution $\widetilde{\boldsymbol{X}} = \sum_j \widetilde{c}_j \widetilde{\boldsymbol{h}}_j \boldsymbol{a}(\widetilde{\tau}_j)^H$. Then we can see that

$$\begin{aligned}\langle \boldsymbol{\lambda}, \boldsymbol{y} \rangle_{\mathbb{R}} &= \left\langle \mathcal{B}^*(\boldsymbol{\lambda}), \widetilde{\boldsymbol{X}} \right\rangle_{\mathbb{R}} \\ &= \left\langle \mathcal{B}^*(\boldsymbol{\lambda}), \sum_j \widetilde{c}_j \widetilde{\boldsymbol{h}}_j \boldsymbol{a}(\widetilde{\tau}_j)^H \right\rangle_{\mathbb{R}} \\ &= \sum_{\tau_j \in \mathbb{D}} \mathrm{Re}\left(\widetilde{c}_j^* \widetilde{\boldsymbol{h}}_j^H \boldsymbol{q}(\widetilde{\tau}_j) \right) + \sum_{\tau_\ell \notin \mathbb{D}} \mathrm{Re}\left(\widetilde{c}_\ell^* \widetilde{\boldsymbol{h}}_\ell^H \boldsymbol{q}(\widetilde{\tau}_\ell) \right) \\ &< \sum_{\tau_j \in \mathbb{D}} |\widetilde{c}_j| + \sum_{\tau_\ell \notin \mathbb{D}} |\widetilde{c}_\ell| \\ &= \|\widetilde{\boldsymbol{X}}\|_{\mathcal{A}}.\end{aligned}$$

Thus, $\boldsymbol{X}_o$ is the unique optimal solution of the atomic norm minimization (3.9) if conditions (3.12) and (3.13) hold. Then, we can form the following linear system of equations:

$$\begin{bmatrix} \boldsymbol{a}(\tau_1)^H \boldsymbol{e}_{-2M} \boldsymbol{b}_{-2M}^H & \cdots & \boldsymbol{a}(\tau_J)^H \boldsymbol{e}_{-2M} \boldsymbol{b}_{-2M}^H \\ \vdots & \ddots & \vdots \\ \boldsymbol{a}(\tau_1)^H \boldsymbol{e}_{2M} \boldsymbol{b}_{2M}^H & \cdots & \boldsymbol{a}(\tau_J)^H \boldsymbol{e}_{2M} \boldsymbol{b}_{2M}^H \end{bmatrix} \begin{bmatrix} c_1 \boldsymbol{h}_1 \\ \vdots \\ c_J \boldsymbol{h}_J \end{bmatrix} = \begin{bmatrix} \boldsymbol{y}(-2M) \\ \vdots \\ \boldsymbol{y}(2M) \end{bmatrix}.$$

The linearly independent condition in Proposition 3.1 ensures that $\{c_j \boldsymbol{h}_j, j = 1, \cdots, J\}$ are unique.

APPENDIX B

PROOF OF LEMMA 3.2

We need the following supporting lemmas.

Lemma B.19. *For arbitrary two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, the non-zero singular values of their Kronecker product $\boldsymbol{A}\otimes\boldsymbol{B}$ are $\sigma_i(\boldsymbol{A})\sigma_j(\boldsymbol{B})$, where $\sigma_i(\boldsymbol{A})$ and $\sigma_j(\boldsymbol{B})$ are the non-zero singular values of $\boldsymbol{A}$ and $\boldsymbol{B}$, respectively. In particular, we have*

$$\|\boldsymbol{A}\otimes\boldsymbol{B}\| = \|\boldsymbol{A}\|\,\|\boldsymbol{B}\|\,.$$

Lemma B.20. *[29] Suppose $\Delta_\tau \geq \frac{1}{M}$. Then $\boldsymbol{D}'$ is invertible and*

$$\|\boldsymbol{I}_{2J} - \boldsymbol{D}'\| \leq 0.3623$$

$$\|\boldsymbol{D}'\| \leq 1.3623$$

$$\left\|\boldsymbol{D}'^{-1}\right\| \leq 1.568$$

According to Lemma B.19, we have

$$\begin{aligned}\|\mathbb{E}\boldsymbol{D}\| &= \|\boldsymbol{D}'\otimes\boldsymbol{I}_K\| \\ &= \|\boldsymbol{D}'\| \\ &\leq 1.3623\end{aligned}$$

and

$$\begin{aligned}\sigma_{\min}(\mathbb{E}\boldsymbol{D}) &= \sigma_{\min}(\boldsymbol{D}'\otimes\boldsymbol{I}_K) \\ &= \sigma_{\min}(\boldsymbol{D}').\end{aligned}$$

As a consequence, we have

$$\begin{aligned}\|\boldsymbol{I}_{2JK} - \mathbb{E}\boldsymbol{D}\| &= \|\boldsymbol{I}_{2J} - \boldsymbol{D}'\| \\ &\leq 0.3623\end{aligned}$$

and

$$\left\|(\mathbb{E}\boldsymbol{D})^{-1}\right\| \leq 1.568.$$

APPENDIX C

PROOF OF LEMMA 3.6

We use the matrix Bernstein inequality for proving Lemma 3.6.

Lemma C.21. *[87] (Matrix Bernstein: Rectangular Case) Consider a finite sequence $\{\boldsymbol{X}_k\}$ of independent, random matrices with dimension $d_1 \times d_2$. Assume that each random matrix satisfies*

$$\mathbb{E}\boldsymbol{X}_k = 0 \quad \textit{and} \quad \|\boldsymbol{X}_k\| \leq R \quad \textit{almost surely.}$$

Define

$$\sigma^2 := \max\left\{ \|\sum_k \mathbb{E}(\boldsymbol{X}_k \boldsymbol{X}_k^H)\|, \|\sum_k \mathbb{E}(\boldsymbol{X}_k^H \boldsymbol{X}_k)\| \right\}.$$

Then, for all $t \geq 0$,

$$\begin{aligned}\mathbb{P}\left\{\|\sum_k \boldsymbol{X}_k\| \geq t\right\} &\leq (d_1+d_2)\cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right) \\ &\leq \begin{cases} (d_1+d_2)\exp\left(\frac{-3t^2}{8\sigma^2}\right), & t \leq \sigma^2/R \\ (d_1+d_2)\exp\left(\frac{-3t}{8R}\right), & t \geq \sigma^2/R. \end{cases}\end{aligned}$$

First of all, we can write

$$\begin{aligned}\boldsymbol{V}_\ell(\tau) - \mathbb{E}\boldsymbol{V}_\ell(\tau) &= \frac{1}{M}\sum_{n=-2M}^{2M} g_M(n)\left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}\right)^\ell e^{-i2\pi n\tau}\boldsymbol{E}(n)\otimes\left(\boldsymbol{b}_n\boldsymbol{b}_n^H - \boldsymbol{I}_K\right) \\ &= \sum_{n=-2M}^{2M} \boldsymbol{Y}_n^\ell,\end{aligned}$$

where we have defined

$$\boldsymbol{Y}_n^\ell = \frac{1}{M} g_M(n)\left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}\right)^\ell e^{-i2\pi n\tau}\boldsymbol{E}(n)\otimes\left(\boldsymbol{b}_n\boldsymbol{b}_n^H - \boldsymbol{I}_K\right).$$

It is easy to see that $\left\{\boldsymbol{Y}_n^\ell\right\}$ are independent random matrices with zero mean due to the isotropy properties of $\boldsymbol{b}_n$. Thus, we can apply the matrix Bernstein inequality for bounding

$\|\boldsymbol{V}_\ell(\tau) - \mathbb{E}\boldsymbol{V}_\ell(\tau)\|$. Before establishing this, we need to compute the quantities R and σ^2 in the matrix Bernstein inequality. The following elementary bound [12, 29] will be useful at this moment.

$$\|g_M(n)\|_\infty \le 1,$$

$$\left|\frac{2\pi n}{\sqrt{|K_M''(0)|}}\right| \le 4, \qquad \text{when } M \ge 2,$$

and

$$\|\boldsymbol{E}(n)\|_2^2 \le 14J, \qquad \text{when } M \ge 4.$$

Thus, we have

$$\begin{aligned}
\left\|\boldsymbol{Y}_n^\ell\right\| &= \left\|\frac{1}{M} g_M(n) \left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}\right)^\ell e^{-i2\pi n\tau} \boldsymbol{E}(n) \otimes \left(\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K\right)\right\| \\
&\le \frac{1}{M} \left\|\left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}\right)^\ell e^{-i2\pi n\tau} \boldsymbol{E}(n) \otimes \left(\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K\right)\right\| \\
&\le \frac{1}{M} 4^\ell \left\|\boldsymbol{E}(n)\right\|_2 \left\|\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K\right\| \\
&\le \frac{1}{M} 4^\ell \sqrt{14J} \max\{\mu K, 1\} \\
&\le \frac{4^{\ell+1}\sqrt{J}\mu K}{M} \\
&=: R,
\end{aligned}$$

where the second line uses the fact that $\|g_M(n)\|_\infty \le 1$, the third line follows because $|\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}| \le 4$ and $\|\boldsymbol{A} \otimes \boldsymbol{B}\| = \|\boldsymbol{A}\| \, \|\boldsymbol{B}\|$ for arbitrary two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$. The fourth line follows from the fact that $\|\boldsymbol{E}(n)\|_2^2 \le 14J$ and $\|\boldsymbol{A} - \boldsymbol{B}\| \le \max\{\|\boldsymbol{A}\|, \|\boldsymbol{B}\|\}$ for two positive semidefinite matrices $\boldsymbol{A}$ and $\boldsymbol{B}$. The fifth line uses the assumption that $\mu K \ge 1$.

Then, we compute the variance term

$$
\begin{aligned}
&\left\| \sum_n \mathbb{E} \boldsymbol{Y}_n^{\ell\, H} \boldsymbol{Y}_n^{\ell} \right\| \\
&= \frac{1}{M^2} \left\| \sum_{n=-2M}^{2M} \mathbb{E}|g_M(n)|^2 \left(\left| \frac{-i2\pi n}{\sqrt{|K_M''(0)|}} \right| \right)^{2\ell} \right. \\
&\qquad \left. \left(\boldsymbol{E}_n^H \otimes \left(\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K \right) \right) \left(\boldsymbol{E}_n \otimes \left(\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K \right) \right) \right\| \\
&\leq \frac{1}{M^2} 4^{2\ell} \left\| \sum_{n=-2M}^{2M} \|\boldsymbol{E}_n\|_2^2 \, \mathbb{E} (\boldsymbol{b}_n \boldsymbol{b}_n^H - \boldsymbol{I}_K)^2 \right\| \\
&\leq \frac{4^{2\ell} \mu K}{M^2} \left\| \sum_{n=-2M}^{2M} \|\boldsymbol{E}_n\|_2^2 \, \boldsymbol{I}_K \right\| \\
&= \frac{4^{2\ell} \mu K}{M^2} \sum_{n=-2M}^{2M} \|\boldsymbol{E}_n\|_2^2 \\
&\leq \frac{80 \cdot 4^{2\ell} \mu J K}{M} \\
&=: \sigma^2,
\end{aligned}
$$

where the second line follows from the fact that $\|g_M(n)\|_\infty \leq 1$ and $|\frac{-i2\pi n}{\sqrt{|K_M''(0)|}}| \leq 4$, the third line uses the fact that $\|\boldsymbol{b}_n\|_2^2 \, \boldsymbol{b}_n \boldsymbol{b}_n^H \preceq \mu K \boldsymbol{b}_n \boldsymbol{b}_n^H, \mu K \geq 1$ and that $\mathbb{E} \boldsymbol{b}_n \boldsymbol{b}_n^H = \boldsymbol{I}_K$ due to the incoherence property (3.15) and isotropy property (3.14), and the last inequality uses the fact that $\|\boldsymbol{E}(n)\|_2^2 \leq 14J$ when $M \geq 4$.

Applying Lemma C.21, we can see that for a fixed ℓ,

$$
\mathbb{P} \left\{ \| \sum_n \boldsymbol{Y}_n^{\ell} \| \geq \varepsilon_2 \right\} \leq (2JK + K) \cdot \exp\left(\frac{-3\varepsilon_2^2}{8\sigma^2} \right).
$$

In order to make this failure probability less than δ_2, we require

$$
\log \left((2JK + K) \cdot \exp\left(\frac{-3\varepsilon_2^2}{8\sigma^2} \right) \right) \leq \log \delta_2,
$$

which leads to the following bound on M,

$$
M \geq \frac{640 \cdot 4^{2\ell} \mu J K}{3\varepsilon_2^2} \log\left(\frac{2JK + K}{\delta_2} \right).
$$

Applying a union bound for $\ell = 0, 1, 2, 3$, we obtain that

$$\mathbb{P}\left\{\|\sum_n \boldsymbol{Y}_n^\ell\| \geq \varepsilon_2, \ell = 0,1,2,3\right\} \leq 4\delta_2,$$

provided that $M \geq \frac{640 \cdot 4^{2\ell}\mu JK}{3\varepsilon_2^2} \log\left(\frac{2JK+K}{\delta_2}\right)$. This completes the proof.

APPENDIX D

PROOF OF LEMMA 3.7

In Lemma 3.6, we showed that for $\ell = 0, 1, 2, 3$, $\|\boldsymbol{V}_\ell(\tau) - \mathbb{E}\boldsymbol{V}_\ell(\tau)\| \leq \varepsilon_2$ with probability at least $1 - 4\delta_2$ provided $M \geq \frac{640 \cdot 4^{2\ell} \mu JK}{3\varepsilon_2^2} \log\left(\frac{2JK+K}{\delta_2}\right)$. Conditioned on the events $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$ and

$$\bigcap_{\tau_d \in \Omega_{\text{Grid}}} \{\|\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\| \leq \varepsilon_2, \ell = 0, 1, 2, 3\}$$

we have

$$\begin{aligned}
\left\|\left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L}\right\| &\leq \|\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\| \, \|\boldsymbol{L}\| \\
&\leq \|\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\| \left\|\boldsymbol{D}^{-1}\right\| \\
&\leq \varepsilon_2 2 \left\|(\mathbb{E}\boldsymbol{D})^{-1}\right\| \\
&\leq 4\varepsilon_2,
\end{aligned}$$

where the second line uses the fact that $\boldsymbol{L}$ is a submatrix of $\boldsymbol{D}^{-1}$, and the third and fourth lines follow from Lemmas 3.5 and 3.2, respectively.

Applying the union bound leads to

$$\mathbb{P}\left\{\sup_{\tau_d \in \Omega_{\text{Grid}}} \left\|\left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L}\right\| \geq 4\varepsilon_2, \ell = 0, 1, 2, 3\right\} \leq |\Omega_{\text{Grid}}| 4\delta_2 + \mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right).$$

APPENDIX E

PROOF OF LEMMA 3.9

We need the following lemma in the proof of Lemma 3.9.

Lemma E.22. *Assume that $\boldsymbol{h}_j \in \mathbb{C}^K$ are i.i.d. random samples on the complex unit sphere $\mathbb{CS}^{K-1}$. Then we have $\mathbb{E}\boldsymbol{h}_j\boldsymbol{h}_j^H = \frac{1}{K}\boldsymbol{I}_K$.*

Proof. Denote $\boldsymbol{\Sigma} = \mathbb{E}\boldsymbol{h}_j\boldsymbol{h}_j^H$. By unitary invariance, we have $\mathbb{E}\boldsymbol{U}\boldsymbol{h}_j\left(\boldsymbol{U}\boldsymbol{h}_j\right)^H = \boldsymbol{\Sigma}$, which implies that $\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^H = \boldsymbol{\Sigma}$ for any unitary matrix $\boldsymbol{U}$. This indicates that $\boldsymbol{\Sigma}$ is diagonal. Furthermore, if $\boldsymbol{U}$ is a permutation matrix, $\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^H$ permutes the diagonal entries of $\boldsymbol{\Sigma}$. This shows that the diagonal entries of $\boldsymbol{\Sigma}$ have equal values. Lastly, $\mathrm{trace}(\boldsymbol{\Sigma}) = \mathbb{E}\mathrm{trace}(\boldsymbol{h}_j\boldsymbol{h}_j^H) = 1$. Thus, we have $\boldsymbol{\Sigma} = \frac{1}{K}\boldsymbol{I}_K$. □

For any $\tau_d \in \Omega_{\mathrm{Grid}}$, define

$$\begin{aligned}\boldsymbol{Q} &:= \left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L} \\ &= \begin{bmatrix}\boldsymbol{Q}_1 & \boldsymbol{Q}_2 & \cdots & \boldsymbol{Q}_J\end{bmatrix},\end{aligned}$$

where each block $\boldsymbol{Q}_j$ is a $K \times K$ matrix. Also, define the event

$$\mathcal{E}_3 := \left\{ \sup_{\tau_d \in \Omega_{\mathrm{Grid}}} \left\| \left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L} \right\| \le 4\varepsilon_2, \ell = 0,1,2,3 \right\}.$$

We can write

$$\begin{aligned}\boldsymbol{I}_1^\ell(\tau_d) &= \left(\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right)^H \boldsymbol{L}\boldsymbol{h} \\ &= \sum_j \boldsymbol{Q}_j \mathrm{sign}(c_j)\boldsymbol{h}_j \\ &=: \sum_j \boldsymbol{Z}_j.\end{aligned}$$

Note that $\mathbb{E}\boldsymbol{Z}_j = \boldsymbol{0}_{K\times 1}$ due to the randomness assumption of $\boldsymbol{h}_j$. Before applying the matrix Bernstein inequality for bounding $\left\|\sum_j \boldsymbol{Z}_j\right\|$, we need to upper bound the operator norm $\|\boldsymbol{Z}_j\|$ and compute the variance term appearing in the expression of matrix Bernstein inequality. Conditioned on $\mathcal{E}_3$,

$$
\begin{aligned}
\|\boldsymbol{Z}_j\|_2 &= \left\|\boldsymbol{Q}_j \mathrm{sign}(c_j)\boldsymbol{h}_j\right\|_2 \\
&\le \left\|\boldsymbol{Q}_j\right\| \\
&\le \|\boldsymbol{Q}\| \\
&\le 4\varepsilon_2 \\
&=: R
\end{aligned}
$$

where the third line uses the fact that $\boldsymbol{Q}_j$ is a submatrix of $\boldsymbol{Q}$.

Next, conditioned on the event $\mathcal{E}_3$ (note that event $\mathcal{E}_3$ includes event $\mathcal{E}_{1,\varepsilon_1\in(0,\frac{1}{4}]}$ and $\mathcal{E}_{2,\varepsilon_2}$), we bound the variance term:

$$
\begin{aligned}
\left\|\sum_j \mathbb{E}\boldsymbol{Z}_j^H \boldsymbol{Z}_j\right\| &= \left\|\sum_j \mathbb{E}(\mathrm{sign}(c_j)\boldsymbol{h}_j)^H \boldsymbol{Q}_j^H \boldsymbol{Q}_j \mathrm{sign}(c_j)\boldsymbol{h}_j\right\| \\
&= \sum_j \mathbb{E}\mathrm{trace}\left(\boldsymbol{Q}_j^H \boldsymbol{Q}_j \boldsymbol{h}_j \boldsymbol{h}_j^H\right) \\
&= \sum_j \mathrm{trace}\left(\boldsymbol{Q}_j^H \boldsymbol{Q}_j \mathbb{E}\left[\boldsymbol{h}_j \boldsymbol{h}_j^H\right]\right) \\
&= \sum_j \mathrm{trace}\left(\boldsymbol{Q}_j^H \boldsymbol{Q}_j \frac{1}{K}\boldsymbol{I}_K\right) \\
&= \frac{1}{K}\|\boldsymbol{Q}\|_F^2,
\end{aligned}
$$

where the third line follows by exchanging the order of the trace operation and expectation, the fourth line uses Lemma E.22. Furthermore, we can bound $\frac{1}{K}\|\boldsymbol{Q}\|_F^2$ as follows:

$$
\begin{aligned}
\frac{1}{K}\|\boldsymbol{Q}\|_F^2 &\le \frac{1}{K}\left\|\boldsymbol{L}^H\right\|^2 \|\boldsymbol{V}_\ell(\tau_d) - \mathbb{E}\boldsymbol{V}_\ell(\tau_d)\|_F^2 \\
&\le \frac{4\cdot 1.568^2}{K} K\varepsilon_2^2 \\
&\le 12\varepsilon_2^2 \\
&=: \sigma^2,
\end{aligned}
$$

where the first line follows from the fact that $\|\boldsymbol{AB}\|_F^2 \le \|\boldsymbol{A}\|^2 \|\boldsymbol{B}\|_F^2$ for arbitrary two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, the second line follows from the fact that $\boldsymbol{L}$ is a submatrix of $\boldsymbol{D}^{-1}$ and Lemmas 3.5 and 3.2.

Applying the matrix Bernstein inequality and the union bound, we get

$$\mathbb{P}\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\boldsymbol{I}_1^\ell(\tau_d)\right\|_2\geq\varepsilon_4\middle|\mathcal{E}_3\right\}\leq|\Omega_{\text{Grid}}|\mathbb{P}\left\{\|\sum_j\boldsymbol{Z}_j\|\geq\varepsilon_4\middle|\mathcal{E}_3\right\}$$
$$\leq|\Omega_{\text{Grid}}|(K+1)\cdot\exp\left(\frac{-\varepsilon_4^2/2}{\sigma^2+R\varepsilon_4/3}\right)$$
$$\leq\begin{cases}|\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_4^2}{8\sigma^2}\right), & \varepsilon_4\leq\sigma^2/R\\ |\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_4}{8R}\right), & \varepsilon_4\geq\sigma^2/R.\end{cases}$$

Taking $\varepsilon_2^2=\frac{640\cdot4^{2\ell}\mu JK}{3M}\log\left(\frac{2JK+K}{\delta_2}\right)$ and applying Lemma 3.7 yield

$$\mathbb{P}\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\boldsymbol{I}_1^\ell(\tau_d)\right\|_2\geq\varepsilon_4\right\}$$
$$\leq\begin{cases}|\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_4^2}{8\sigma^2}\right)+|\Omega_{\text{Grid}}|4\delta_2+\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right), & \varepsilon_4\leq\sigma^2/R\\ |\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_4}{8R}\right)+|\Omega_{\text{Grid}}|4\delta_2+\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right), & \varepsilon_4\geq\sigma^2/R.\end{cases}$$

According to Lemmas 3.6 and 3.7, for the second term $|\Omega_{\text{Grid}}|4\delta_2\leq\delta$, it is sufficient to have

$$M\geq\frac{640\cdot4^{2\ell}\mu JK}{3\varepsilon_2^2}\log\left(\frac{4|\Omega_{\text{Grid}}|(2JK+K)}{\delta}\right).$$

To make the failure probability smaller than δ for the first term, we choose

$$\begin{cases}96\varepsilon_2^2=\frac{3\varepsilon_4^2}{\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right)}, & \varepsilon_4\leq\sigma^2/R,\\ 32\varepsilon_2=\frac{3\varepsilon_4}{\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right)}, & \varepsilon_4\geq\sigma^2/R.\end{cases}$$

Equivalently, when $\varepsilon_4\leq\sigma^2/R$, one has

$$M\geq\frac{640\cdot4^{2\ell}\mu JK}{3\varepsilon_2^2}\log\left(\frac{4|\Omega_{\text{Grid}}|(2JK+K)}{\delta}\right)$$
$$=\frac{640\cdot96\cdot4^{2\ell}\mu JK}{9\varepsilon_4^2}\log\left(\frac{4|\Omega_{\text{Grid}}|(2JK+K)}{\delta}\right)\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right).$$

When $\varepsilon_4\geq\sigma^2/R$, one has

$$M\geq\frac{640\cdot4^{2\ell}\mu JK}{3\varepsilon_2^2}\log\left(\frac{4|\Omega_{\text{Grid}}|(2JK+K)}{\delta}\right)$$
$$=\frac{32^2\cdot640\cdot4^{2\ell}\mu JK}{27\varepsilon_4^2}\log\left(\frac{4|\Omega_{\text{Grid}}|(2JK+K)}{\delta}\right)\log^2\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right).$$

Finally, according to Lemma 3.3, for the third term $\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right)\leq\delta$, we have

$$M\geq\frac{80\mu JK}{\varepsilon_1^2}\log\left(\frac{4JK}{\delta}\right).$$

Setting $\varepsilon_1 = \frac{1}{4}$, absorbing all of the constants into one and applying the union bound for $\ell = 0, 1, 2, 3$, we can see that

$$\mathbb{P}\left\{ \sup_{\tau_d \in \Omega_{\text{Grid}}} \left\| \boldsymbol{I}_1^{\ell}(\tau_d) \right\|_2 \geq \varepsilon_4, \ell = 0, 1, 2, 3 \right\} \leq 12\delta$$

provided

$$M \geq C\mu JK \max\left\{ \frac{1}{\varepsilon_4^2} \log\left(\frac{|\Omega_{\text{Grid}}|JK}{\delta} \right) \log^2\left(\frac{|\Omega_{\text{Grid}}|K}{\delta} \right), \log\left(\frac{JK}{\delta} \right) \right\}$$

for some constant C.

APPENDIX F

PROOF OF LEMMA 3.10

Observe that

$$\begin{aligned}
\|\mathbb{E}\boldsymbol{V}_\ell(\tau)\|_F^2 &= \left\| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \begin{bmatrix} K_M^\ell(\tau-\tau_1)^* \\ \vdots \\ K_M^\ell(\tau-\tau_J)^* \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_1)^* \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_J)^* \end{bmatrix} \otimes \boldsymbol{I}_K \right\|_F^2 \\
&= K \left\| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \begin{bmatrix} K_M^\ell(\tau-\tau_1)^* \\ \vdots \\ K_M^\ell(\tau-\tau_J)^* \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_1)^* \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_J)^* \end{bmatrix} \right\|_2^2 \\
&\le CK
\end{aligned}$$

for some numerical constant C, where the inequality above follows from Lemma IV.9 of [12]. The key to being able to obtain such a bound of order $O(K)$ is because $\{\tau_j\}$ are well separated, implying that the sequence $\left\{K_M^\ell(\tau-\tau_j)\right\}$ decreases rapidly if properly ordered. Then, conditioned on the event $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$, we have

$$\begin{aligned}
\left\| \left(\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\right)^H \mathbb{E}\boldsymbol{V}_\ell(\tau) \right\|_F^2 &\le \left\| \left(\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\right)^H \right\|^2 \|\mathbb{E}\boldsymbol{V}_\ell(\tau)\|_F^2 \\
&\le (2 \cdot 1.568^2 \varepsilon_1)^2 CK \\
&=: CK\varepsilon_1^2
\end{aligned}$$

for some redefined numerical constant C. The first line above uses the inequality $\|\boldsymbol{AB}\|_F^2 \le \|\boldsymbol{A}\|^2 \|\boldsymbol{B}\|_F^2$, and the second line above follows from the fact that $\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K$ is a submatrix of $\boldsymbol{D}^{-1} - \mathbb{E}\boldsymbol{D}^{-1}$ and from Lemmas 3.5 and 3.2.

APPENDIX G

PROOF OF LEMMA 3.11

To begin with, for any $\tau_d \in \Omega_{\text{Grid}}$, define

$$\begin{aligned}\widetilde{\boldsymbol{Q}} &:= \left[\mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right]^H \left(\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\right) \\ &= \left[\widetilde{\boldsymbol{Q}}_1 \quad \widetilde{\boldsymbol{Q}}_2 \quad \ldots \quad \widetilde{\boldsymbol{Q}}_J\right],\end{aligned}$$

where each block $\widetilde{\boldsymbol{Q}}_j$ is a $K \times K$ matrix. Then, we have

$$\begin{aligned}\boldsymbol{I}_2^\ell(\tau_d) &= \left[\mathbb{E}\boldsymbol{V}_\ell(\tau_d)\right]^H \left(\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\right) \boldsymbol{h} \\ &= \sum_j \widetilde{\boldsymbol{Q}}_j \text{sign}(c_j) \boldsymbol{h}_j \\ &=: \sum_j \widetilde{\boldsymbol{Z}}_j.\end{aligned}$$

Again, we bound $\left\|\boldsymbol{I}_2^\ell(\tau_d)\right\|_2$ using the matrix Bernstein inequality. First of all, we have $\mathbb{E}\widetilde{\boldsymbol{Z}}_j = \boldsymbol{0}_{K\times 1}$ due to the randomness assumption of $\boldsymbol{h}_j$. Conditioned on $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$,

$$
\begin{aligned}
\left\|\widetilde{\boldsymbol{Z}}_j\right\| &= \left\|\widetilde{\boldsymbol{Q}}_j \operatorname{sign}(c_j)\boldsymbol{h}_j\right\| \\
&\le \|\widetilde{\boldsymbol{Q}}\| \\
&\le \|\boldsymbol{L} - \boldsymbol{L}' \otimes \boldsymbol{I}_K\| \, \|\mathbb{E}\boldsymbol{V}_\ell(\tau_d)\| \\
&\le 2 \cdot 1.568^2 \varepsilon_1 \left\| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \begin{bmatrix} K_M^\ell(\tau-\tau_1)^* \\ \vdots \\ K_M^\ell(\tau-\tau_J)^* \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_1)^* \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_J)^* \end{bmatrix} \otimes \boldsymbol{I}_K \right\| \\
&= 2 \cdot 1.568^2 \varepsilon_1 \left\| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \begin{bmatrix} K_M^\ell(\tau-\tau_1)^* \\ \vdots \\ K_M^\ell(\tau-\tau_J)^* \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_1)^* \\ \vdots \\ \frac{1}{\sqrt{|K_M''(0)|}} K_M^{\ell+1}(\tau-\tau_J)^* \end{bmatrix} \right\|_2 \\
&\le C\varepsilon_1 \\
&=: R.
\end{aligned}
$$

for some small universal constant C. For the variance term, we have

$$
\begin{aligned}
\left\|\sum_j \mathbb{E}\widetilde{\boldsymbol{Z}}_j^H \widetilde{\boldsymbol{Z}}_j\right\| &= \left\|\sum_j \mathbb{E}(\operatorname{sign}(c_j)\boldsymbol{h}_j)^H \widetilde{\boldsymbol{Q}}_j^H \widetilde{\boldsymbol{Q}}_j \operatorname{sign}(c_j)\boldsymbol{h}_j\right\| \\
&= \sum_j \mathbb{E}\operatorname{trace}\left(\widetilde{\boldsymbol{Q}}_j^H \widetilde{\boldsymbol{Q}}_j \boldsymbol{h}_j \boldsymbol{h}_j^H\right) \\
&= \sum_j \operatorname{trace}\left(\widetilde{\boldsymbol{Q}}_j^H \widetilde{\boldsymbol{Q}}_j \mathbb{E}\left[\boldsymbol{h}_j \boldsymbol{h}_j^H\right]\right) \\
&= \sum_j \operatorname{trace}\left(\widetilde{\boldsymbol{Q}}_j^H \widetilde{\boldsymbol{Q}}_j \frac{1}{K}\boldsymbol{I}_K\right) \\
&\le \frac{1}{K} CK\varepsilon_1^2 \\
&= C\varepsilon_1^2 \\
&=: \sigma^2
\end{aligned}
$$

where the third line follows by exchanging the trace operation and expectation, the fourth line uses Lemma E.22, and the fifth line follows from Lemma 3.10.

The matrix Bernstein inequality and the union bound yield

$$\mathbb{P}\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\boldsymbol{I}_2^{\ell}(\tau_d)\right\|_2\geq\varepsilon_5\middle|\mathcal{E}_{1,\varepsilon_1}\right\}\leq|\Omega_{\text{Grid}}|\mathbb{P}\left\{\|\sum_j\widetilde{\boldsymbol{Z}}_j\|\geq\varepsilon_5\middle|\mathcal{E}_{1,\varepsilon_1}\right\}$$
$$\leq|\Omega_{\text{Grid}}|(K+1)\cdot\exp\left(\frac{-\varepsilon_5^2/2}{\sigma^2+R\varepsilon_5/3}\right)$$
$$\leq\begin{cases}|\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_5^2}{8\sigma^2}\right), & \varepsilon_5\leq\sigma^2/R\\ |\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_5}{8R}\right), & \varepsilon_5\geq\sigma^2/R.\end{cases}$$

Therefore, we can write

$$\mathbb{P}\left\{\sup_{\tau_d\in\Omega_{\text{Grid}}}\left\|\boldsymbol{I}_2^{\ell}(\tau_d)\right\|_2\geq\varepsilon_5\right\}$$
$$\leq\begin{cases}|\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_5^2}{8\sigma^2}\right)+\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right), & \varepsilon_5\leq\sigma^2/R\\ |\Omega_{\text{Grid}}|(K+1)\exp\left(\frac{-3\varepsilon_5}{8R}\right)+\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right), & \varepsilon_5\geq\sigma^2/R.\end{cases}$$

When $\varepsilon_5\leq\sigma^2/R$, to ensure the first term less than δ, it is sufficient to choose

$$\varepsilon_1^2=\frac{3\varepsilon_5^2}{8C\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right)}$$

for the same constant C that appears in the variance bound above.

To make the second term $\mathbb{P}\left(\mathcal{E}_{1,\varepsilon_1}^c\right)$ less than δ, according to Lemma 3.3, we have

$$M\geq\frac{80\mu JK}{\varepsilon_1^2}\log\left(\frac{4JK}{\delta}\right)$$
$$=\frac{8C\cdot80\mu JK}{3\varepsilon_5^2}\log\left(\frac{4JK}{\delta}\right)\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right)$$
$$=:C\frac{\mu JK}{\varepsilon_5^2}\log\left(\frac{JK}{\delta}\right)\log\left(\frac{|\Omega_{\text{Grid}}|K}{\delta}\right)$$

with a redefined numerical constant C.

Similarly, when $\varepsilon_5\geq\sigma^2/R$, to ensure the first term less than δ, we can take

$$\varepsilon_1=\frac{3\varepsilon_5}{8C\log\left(\frac{|\Omega_{\text{Grid}}|(K+1)}{\delta}\right)},$$

for the same constant C shown in the bound of $\left\|\widetilde{\boldsymbol{Z}}_j\right\|$.

To make the second term less than δ, we have

$$
\begin{aligned}
M &\geq \frac{80\mu JK}{\varepsilon_1^2} \log\left(\frac{4JK}{\delta}\right) \\
&= \frac{(8C)^2 \cdot 80\mu JK}{9\varepsilon_5^2} \log\left(\frac{4JK}{\delta}\right) \log^2\left(\frac{|\Omega_{\mathrm{Grid}}|(K+1)}{\delta}\right) \\
&=: C\frac{\mu JK}{\varepsilon_5^2} \log\left(\frac{JK}{\delta}\right) \log^2\left(\frac{|\Omega_{\mathrm{Grid}}|K}{\delta}\right)
\end{aligned}
$$

for a redefined numerical constant C.

Combining the two different cases above and applying the union bound with respect to $\ell = 0, 1, 2, 3$ complete the proof.

APPENDIX H

PROOF OF LEMMA 3.12

Denote the pth column of $\boldsymbol{V}_\ell(\tau)$ as $\boldsymbol{V}_\ell(\tau;p)$, whose ℓ_2 norm can be bounded as follows:

$$\begin{aligned}\|\boldsymbol{V}_\ell(\tau;p)\|_2 &= \left\| \frac{1}{M} \sum_{n=-2M}^{2M} g_M(n) \left(\frac{-i2\pi n}{\sqrt{|K_M''(0)|}} \right)^\ell e^{-i2\pi n\tau} \boldsymbol{E}(n) \otimes \boldsymbol{b}_n \boldsymbol{b}_n^*(p) \right\|_2 \\ &\le \frac{(4M+1)4^\ell}{M} \mu\sqrt{K} \, \|\boldsymbol{E}(n)\|_2 \\ &\le \frac{\sqrt{14}\cdot(4M+1)4^\ell}{M} \mu\sqrt{JK} \\ &=: C\mu\sqrt{JK}\end{aligned}$$

for some constant C, where we have used the fact that $|\boldsymbol{b}_n(p)| \le \sqrt{\mu}$ in the first inequality and the fact that $\|\boldsymbol{E}(n)\|_2^2 \le 14J$ when $M \ge 4$ in the second inequality.

We define the pth entry of $\boldsymbol{q}(\tau)$ as $\boldsymbol{q}(\tau;p)$. Conditioned on the event $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$, we have

$$\begin{aligned}\left| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \boldsymbol{q}^\ell(\tau;p) \right| &= \left| \boldsymbol{V}_\ell(\tau;p)^H \boldsymbol{L}\boldsymbol{h} \right| \\ &\le \|\boldsymbol{V}_\ell(\tau;p)\|_2 \, \|\boldsymbol{L}\| \, \|\boldsymbol{h}\|_2 \\ &\le C\mu J\sqrt{K}\end{aligned}$$

for some constant C.

Applying Bernstein's polynomial inequality [12, 88], we have

$$\begin{aligned}&\left| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \boldsymbol{q}^\ell(\tau_a;p) - \frac{1}{\sqrt{|K_M''(0)|}^\ell} \boldsymbol{q}^\ell(\tau_b;p) \right| \\ &\le |e^{i2\pi\tau_a} - e^{i2\pi\tau_b}| \sup_{z=e^{i2\pi\tau}} \left| \frac{d \frac{1}{\sqrt{|K_M''(0)|}^\ell} \boldsymbol{q}^\ell(\tau;p)}{dz} \right| \\ &\le 4\pi|\tau_a - \tau_b| 2M \sup_\tau \left| \frac{1}{\sqrt{|K_M''(0)|}^\ell} \boldsymbol{q}^\ell(\tau;p) \right| \\ &\le C\mu J\sqrt{K} M|\tau_a - \tau_b|\end{aligned}$$

for some numerical constant C. Therefore, we have

$$\left\| \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau_a) - \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau_b) \right\|_2 \leq C\mu JKM|\tau_a - \tau_b| \leq CM^2|\tau_a - \tau_b|,$$

where the second line follows when $M \geq \mu JK$. We choose Ω_{Grid} such that for any $\tau \in [0,1)$, there exists a point $\tau_d \in \Omega_{\text{Grid}}$ with $|\tau - \tau_d| \leq \frac{\varepsilon}{3CM^2}$. Note that $|\Omega_{\text{Grid}}| \leq \frac{3CM^2}{\varepsilon}$.

Using such a choice of Ω_{Grid} and conditioned on the event $\mathcal{E}_{1,\varepsilon_1}$ with $\varepsilon_1 \in (0, \frac{1}{4}]$ and event $\mathcal{E}$, we have

$$\begin{aligned}
&\left\| \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau) - \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \overline{\boldsymbol{q}}^{\ell}(\tau) \right\|_2 \\
&\leq \left\| \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau) - \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau_d) \right\|_2 \\
&\quad + \left\| \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \boldsymbol{q}^{\ell}(\tau_d) - \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \overline{\boldsymbol{q}}^{\ell}(\tau_d) \right\|_2 \\
&\quad + \left\| \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \overline{\boldsymbol{q}}^{\ell}(\tau_d) - \frac{1}{\sqrt{|K_M''(0)|}^{\ell}} \overline{\boldsymbol{q}}^{\ell}(\tau) \right\|_2 \\
&\leq CM^2|\tau - \tau_d| + \frac{\varepsilon}{3} + CM^2|\tau - \tau_d| \\
&\leq \varepsilon, \quad \forall\ \tau \in [0,1)
\end{aligned}$$

where the first inequality follows from the triangle inequality, the second inequality follows from Proposition 3.2. With such a choice of grid size, we can immediately get the following bound

$$M \geq C\mu JK \max\left\{ \frac{1}{\varepsilon^2} \log\left(\frac{MJK}{\varepsilon\delta}\right) \log^2\left(\frac{MK}{\varepsilon\delta}\right), \log\left(\frac{JK}{\delta}\right) \right\}$$

with a redefined numerical constant C. This finishes the proof of Lemma 3.12.

APPENDIX I

SUPPLEMENTARY MATERIALS FOR LEMMA 3.14

First of all, we record some useful results from the proof of Lemma 2.3 and Lemma 2.4 in [29].

Assume that $\Delta_\tau \geq \frac{1}{M}, M \geq 64$. Then, we have

$$0.9539 \leq K(\tau) \leq 1, \quad \tau \in [-\tau_{b,1}, \tau_{b,1}],$$

$$-13.572M^2 \leq K''(\tau) \leq -11.692M^2, \quad \tau \in [-\tau_{b,1}, \tau_{b,1}],$$

$$\sum_{\tau_j \in \mathbb{D}} |K'(\tau - \tau_j)| \leq 1.2722M, \quad \tau \in [-\tau_{b,1}, \tau_{b,1}],$$

$$\sum_{\tau_j \in \mathbb{D}} |K'''(\tau - \tau_j)| \leq 194.0560M^3, \quad \tau \in [-\tau_{b,1}, \tau_{b,1}],$$

$$\sum_{\tau_j \in \mathbb{D}\setminus 0} |K(\tau - \tau_j)| \leq 6.279 \times 10^{-3}, \quad \tau \in \Omega_{\text{near}} \setminus [-\tau_{b,1}, \tau_{b,1}],$$

$$\sum_{\tau_j \in \mathbb{D}\setminus 0} |K''(\tau - \tau_j)| \leq 4.2200M^2, \quad \tau \in \Omega_{\text{near}} \setminus [-\tau_{b,1}, \tau_{b,1}].$$

We know that

$$\overline{\boldsymbol{q}}(\tau) = \sum_{j=1}^{J} K_M(\tau - \tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{j=1}^{J} K'_M(\tau - \tau_j)\overline{\boldsymbol{\beta}}_j.$$

It was shown in [38] that

$$1 - 8.824 \times 10^{-3} = \alpha_{\min} \leq \left\|\overline{\boldsymbol{\alpha}}_j\right\|_2 \leq \alpha_{\max} = 1 + 8.824 \times 10^{-3}$$

and

$$\left\|\overline{\boldsymbol{\beta}}_j\right\|_2 \leq \beta_{\max} = \frac{1.647}{M} \times 10^{-2}.$$

Without loss of generality, in the following, we assume that $0 \in \mathbb{D}$. Thus, we have

$$
\begin{aligned}
\|\overline{\boldsymbol{q}}(\tau)\|_2 &= \left\| \sum_{\tau_j \in \mathbb{D}} K(\tau - \tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{\tau_j \in \mathbb{D}} K'(\tau - \tau_j)\overline{\boldsymbol{\beta}}_j \right\|_2 \\
&\leq \alpha_{\max} \sum_{\tau_j \in \mathbb{D}} |K(\tau - \tau_j)| + \beta_{\max} \sum_{\tau_j \in \mathbb{D}} |K'(\tau - \tau_j)| \\
&\leq 1.008824 \times (1 + 6.279 \times 10^{-3}) + \frac{1.647 \times 10^{-2}}{M} \times (1.2722M) \\
&= 1.0361,
\end{aligned}
$$

$$
\begin{aligned}
\|\overline{\boldsymbol{q}}'(\tau)\|_2 &= \left\| \sum_{\tau_j \in \mathbb{D}} K'(\tau - \tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{\tau_j \in \mathbb{D}} K''(\tau - \tau_j)\overline{\boldsymbol{\beta}}_j \right\|_2 \\
&\leq \alpha_{\max} \sum_{\tau_j \in \mathbb{D}} |K'(\tau - \tau_j)| + \beta_{\max}|K''(\tau)| + \beta_{\max} \sum_{\tau_j \in \mathbb{D}\backslash 0} |K''(\tau - \tau_j)| \\
&\leq 1.008824 \times (1.2722M) + \frac{1.647 \times 10^{-2}}{M} \times (13.572M^2 + 4.2200M^2) \\
&= 1.5765M,
\end{aligned}
$$

$$
\begin{aligned}
\|\overline{\boldsymbol{q}}''(\tau)\|_2 &= \left\| \sum_{\tau_j \in \mathbb{D}} K''(\tau - \tau_j)\overline{\boldsymbol{\alpha}}_j + \sum_{\tau_j \in \mathbb{D}} K'''(\tau - \tau_j)\overline{\boldsymbol{\beta}}_j \right\|_2 \\
&\leq \alpha_{\max}|K''(\tau)| + \alpha_{\max} \sum_{\tau_j \in \mathbb{D}\backslash 0} |K''(\tau - \tau_j)| + \beta_{\max} \sum_{\tau_j \in \mathbb{D}} |K'''(\tau - \tau_j)| \\
&\leq 1.008824 \times (13.572M^2 + 4.2200M^2) + \frac{1.647 \times 10^{-2}}{M} \times (194.0560M^3) \\
&= 21.1451M^2,
\end{aligned}
$$

$$
\begin{aligned}
(\overline{\boldsymbol{q}}''(\tau))^H \overline{\boldsymbol{q}}(\tau) &= \left(\sum_{\tau_j \in \mathbb{D}} K''(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j + \sum_{\tau_j \in \mathbb{D}} K'''(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j \right)^H \\
&\quad \cdot \left(\sum_{\tau_j \in \mathbb{D}} K(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j + \sum_{\tau_j \in \mathbb{D}} K'(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j \right) \\
&= \|\overline{\boldsymbol{\alpha}}_k\|_2^2 K''(\tau) K(\tau) + \overline{\boldsymbol{\alpha}}_k^H K''(\tau) \sum_{\tau_j \in \mathbb{D} \setminus 0} K(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j \\
&\quad + \overline{\boldsymbol{\alpha}}_k^H K''(\tau) \sum_{\tau_j \in \mathbb{D}} K'(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j + \left(\sum_{\tau_j \in \mathbb{D} \setminus 0} K''(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j \right)^H \overline{\boldsymbol{q}}(\tau) \\
&\quad + \left(\sum_{\tau_j \in \mathbb{D}} K'''(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j \right)^H \overline{\boldsymbol{q}}(\tau).
\end{aligned}
$$

Now we upper bound each term:

$$
\begin{aligned}
\|\overline{\boldsymbol{\alpha}}_k\|_2^2 K''(\tau) K(\tau) &\leq \alpha_{\min}^2 \times (-11.692 M^2) \times 0.9539 \\
&= -10.9570 M^2,
\end{aligned}
$$

$$
\begin{aligned}
\overline{\boldsymbol{\alpha}}_k^H K''(\tau) \sum_{\tau_j \in \mathbb{D} \setminus 0} K(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j &\leq \alpha_{\max}^2 |K''(\tau)| \sum_{\tau_j \in \mathbb{D} \setminus 0} |K(\tau - \tau_j)| \\
&\leq 1.008824^2 \times (13.572 M^2) \times (6.279 \times 10^{-3}) \\
&= 0.0867 M^2,
\end{aligned}
$$

$$
\begin{aligned}
\overline{\boldsymbol{\alpha}}_k^H K''(\tau) \sum_{\tau_j \in \mathbb{D}} K'(\tau - \tau_j) \overline{\boldsymbol{\beta}}_j &\leq \alpha_{\max} \beta_{\max} |K''(\tau)| \sum_{\tau_j \in \mathbb{D}} |K'(\tau - \tau_j)| \\
&\leq 1.008824 \times \frac{1.647 \times 10^{-2}}{M} \times (13.572 M^2) \times (1.2722 M) \\
&= 0.2869 M^2,
\end{aligned}
$$

$$
\begin{aligned}
\left(\sum_{\tau_j \in \mathbb{D} \setminus 0} K''(\tau - \tau_j) \overline{\boldsymbol{\alpha}}_j \right)^H \overline{\boldsymbol{q}}(\tau) &\leq \alpha_{\max} \sum_{\tau_j \in \mathbb{D} \setminus 0} |K''(\tau - \tau_j)| \, \|\overline{\boldsymbol{q}}(\tau)\|_2 \\
&\leq 1.008824 \times (4.2200 M^2) \times 1.0361 \\
&= 4.4109 M^2,
\end{aligned}
$$

$$\left(\sum_{\tau_j\in\mathbb{D}} K'''(\tau-\tau_j)\overline{\boldsymbol{\beta}}_j\right)^H \overline{\boldsymbol{q}}(\tau) \le \beta_{\max}\sum_{\tau_j\in\mathbb{D}} |K'''(\tau-\tau_j)|\,\|\overline{\boldsymbol{q}}(\tau)\|_2$$
$$\le (194.0560M^3)\times\frac{0.01647}{M}\times 1.0361$$
$$= 3.3115M^2.$$

Combining the above upper bounds, we have

$$(\overline{\boldsymbol{q}}''(\tau))^H\overline{\boldsymbol{q}}(\tau) \le (-10.9570+0.0867+0.2869+4.4109+3.3115)M^2$$
$$= -2.8610M^2.$$

Consequently, we have

$$\|\overline{\boldsymbol{q}}'(\tau)\|_2^2 + \mathrm{Re}(\overline{\boldsymbol{q}}''(\tau))^H\overline{\boldsymbol{q}}(\tau) \le (1.5765M)^2 - 2.8610M^2$$
$$= -0.3756M^2.$$