Technische Universität München

Fakultät für Mathematik

Bilinear Compressed Sensing

Dominik Stöger

Vollständiger Abdruck der von der Fakultät für Mathematik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigte Dissertation

Vorsizender: Prof. Dr. Stefan Weltge

Prüfer der Dissertation:

1. Prof. Dr. Felix Krahmer
2. Prof. Dr. Kiryung Lee
3. Prof. Dr. Yaniv Plan

Die Dissertation wurde am 25.04.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Mathematik am 07.08.2019 angenommen.

Abstract

This dissertation studies the performance of convex and non-convex algorithms for randomized bilinear inverse problems. In particular, we examine algorithmic approaches based on convex relaxation for certain models of the blind deconvolution problem. These models are especially relevant for wireless communications. Moreover, this thesis analyses whether the required number of measurements can be reduced when additional information about the signals is available such as sparsity.

Zusammenfassung

Diese Dissertation untersucht die Leistungsfähigkeit von konvexen und nicht-konvexen Algorithmen für randomisierte bilineare inverse Probleme. Insbesondere werden algorithmische Ansätze für bestimmte Modelle des Problems der blinden Entfaltung betrachtet, welche auf konvexer Relaxation basieren. Diese Modelle sind insbesondere für die Nachrichtentechnik von Bedeutung. Des Weiteren wird in dieser Arbeit untersucht, ob die Anzahl der benötigten Messungen reduziert werden kann, falls zusätzliche Informationen über die Signale vorhanden sind wie etwa deren Dünnbesetztheit.

Acknowledgments

To begin with, I would like to express my gratitude to my advisor Professor Felix Krahmer, who gave me the opportunity to do research in this exciting field. His approach to research was ideal for me. While always insisting on tackling relevant and important problems, he nevertheless gave me enough freedom so that I could follow my interests and find my own path in mathematics. I learned a lot from his advice, which often went beyond purely scientific matters.
Moreover, I have been very lucky that I got the opportunity to collaborate with Peter Jung from the very beginning of my doctoral studies. Whenever I had questions concerning research and beyond, he was willing to help. Moreover, I profited a lot from his ability to translate problems arising in communications into mathematical questions.
Furthermore, I want to thank all my scientific collaborators as well as all the numerous people I had the opportunity to meet and discuss science with. I've learned so much from you!
I would like to thank Professor Stefan Weltge for being chair of my defense. In addition, I am grateful to Professor Kiryung Lee and Professor Yaniv Plan for serving on my thesis committee.
Furthermore, I would like to thank all my colleagues in the math department at TUM for creating such a pleasant environment. In particular, I want to thank all members of the research group M15 for all the nice discussions we had during the last years.
Moreover, I was very fortunate to have great friends in Munich. I will never forget the nice weekends, where we went biking or where we enjoyed Thai (and German) food;-)
Finally, I am grateful to my family for their love and their constant support through all the good and bad times of my life. Thank you!

List of contributed articles

Core publications:

1. F. Krahmer, P. Jung, and D. Stöger:
   Blind demixing and deconvolution at near-optimal rate
   IEEE Trans. Inf. Theory 64, No. 2, 704-727 (2018)
   (Cf. bibliography entry [56])

2. J. Geppert, F. Krahmer, and D. Stöger:
   Sparse Power Factorization: Balancing peakiness and sample complexity
   Adv. Comput. Math. 45, No. 3, 1711-1728 (2019).
   (Cf. bibliography entry [43])

Further publications:

1. F. Krahmer and D. Stöger
   On the convex geometry of blind deconvolution of matrix completion
   Submitted for publication in Communications on Pure and Applied Mathematics
   (Cf. bibliography entry [68])

I, Dominik Stöger, am the principal author of these articles.

Contents

1. Introduction

There are many problems in signal processing, which can be formulated as the task of reconstructing an unknown signal $x_0 \in \mathbb{C}^n$ from measurements of the form

$$y_i = \langle a_i, x_0 \rangle \quad \text{for } i \in [m] := \{1; \dots; m\},$$

where $a_i \in \mathbb{C}^n$ represents a known measurement vector. In many of these tasks one would like to keep the number of measurements m as small as possible, as obtaining these measurements is either expensive or time-consuming. One way to reduce the number of measurements is to make use of some a-priori knownledge of the signal such as sparsity, i.e., only few entries are non-zero. For example, already in [83] the author leveraged the sparsity of the signal by ℓ_1-minimization. The foundational papers [16] and [32] could then prove that x_0 can be recovered via practical optimization algorithms using only a small amount of measurements, if the sparsity assumption on the signal x_0 is combined with the assumption that the measurements $\{a_i\}_{i=1}^m$ are random. Arguably, these two papers initiated a whole new research field, which is now often referred as *compressed sensing*. Compressed sensing has, for example, been applied to medical magnetic resonance imaging, where the scanning time could be significantly reduced [84, 114]. As pointed out in [113] most work in the field of compressed sensing is built on the assumptions that the following three requirements are fulfilled *simultaneously*.

- **Sparsity**: the signal x_0 is sparse in a known basis, i.e., only very few coefficient in this basis representation are non-zero
- **Linear measurements**: the measurements y_i depend linearly on the unknown signal x_0
- **Randomness**: the measurement vectors a_i are not completely deterministic, but possess at least some degree of randomness

However, there are many scenarios in which not all of these three assumptions are fulfilled at once. Consider, for example, the problem of phase retrieval, which is ubiquitous in many areas of science and engineering such as X-ray crystallography [50, 90], astronomical imaging [36], ptychography [102], and quantum tomography [70]. In this signal processing problem one is only able to measure the modulus and not the phase, i.e., the measurements are given by

$$y_i = |\langle a_i, x_0 \rangle|^2 \quad \text{for } i \in [m].$$

Note that the measurements are no longer linear, which makes the problem much more difficult, even if one makes no structural assumptions such as sparsity and one wants to recover the ground truth vector x_0 with approximately as many measurements as the degrees of freedom of x_0. In [17] it was shown that if one assumes that the measurement vectors a_i are random, then one can prove recovery guarantees for a convex relaxation of this problem. As in compressed sensing, this has triggered a large number of follow-up works, which study more efficient algorithms (see, e.g., [12, 116, 95, 115, 7]), more realistic measurement scenarios (see, e.g., [11, 47, 48]), or which try to incorporate structural assumptions about x_0, (see, e.g., [7, 104]).

The problem of *blind deconvolution*, which arises in many different areas of science and engineering such as astronomy [54], neuroscience [34], medical imaging [89], and wireless communications [45], is another signal processing problem, where the measurements are nonlinear. In blind deconvolution, one observes the convolution

$$y = w * x,$$

and one wants to reconstruct w and x. As blind deconvolution is important for so many different subject, several algorithms for this problem have been proposed in the last decades (see, e.g., [72, 62, 20]), often with no provable performance guarantees. However, due to the success of randomization in compressed sensing and phase retrieval, a new viewpoint has been introduced in [1]. Namely, one assumes that w and x are contained in known, but *random* subspaces. Such a scenario is especially relevant for wireless communications. As for phase retrieval this caused a lot of follow-up work on blind deconvolution (see, e.g., [80, 75, 81, 58]) and related bilinar inverse problems such as self-calibration and passive imaging (see, e.g., [79, 82, 76, 74]). Due to the bilinear structure of these problems Friedlander and Strohmer came up with the expression "bilinear compressed sensing" [41].

As we have seen there has been a lot of work on bilinear inverse problems with randomness in the last few years. However, many questions remain open. In this thesis we will provide near-optimal recovery guarantees for a generalization of the blind deconvolution problem. Moreover, we will study the noise robustness of convex relaxations for blind deconvolution problem. We will see that are many similarities with the matrix completion problem [15], a seemingly rather different problem. Moreover, we will discuss, how and to which extent structural assumptions such as sparsity can be exploited in these bilinear inverse problems.

Outline of this dissertation

This dissertation is structured as follows. In Chapter 2 we will introduce the randomized blind deconvolution model with known subspaces as in [1]. We will show how it can be modeled as a low-rank matrix recovery problem and how it can be solved via nuclear

norm minimization. We will review theoretical results and relevant literature. Moreover, we will discuss the matrix completion problem. Although this is not a bilinear inverse problem we will see that it shares some similarity with the blind deconvolution problem, most prominently the necessity of coherence terms in recovery guarantees and the appearance of seemingly suboptimal noise factors in the error bounds, when convex relaxation is used. Furthermore, we will discuss some related literature on matrix completion.
In Chapter 3 we will review basic techniques in establishing recovery guarantees for low-rank matrix recovery problems via nuclear norm minimization. In particular, we will discuss the notion of an *approximate dual certificate*. We will try to elucidate what gives rise to the suboptimal noise factors in the noise bounds for blind deconvolution and matrix completion in Chapter 2.
In Chapter 4 we will discuss the extension of blind deconvolution to the blind demixing model. That is one tries to recover signals w_i and x_i from the superposition of their convolutions $\sum_{i=1}^{r} w_i * x_i$. We will show for the first time that this signal processing problem can be solved by a tractable algorithm at a near-optimal rate in terms of the number of observations, thus improving upon results by Ling and Strohmer [80].
In Chapter 5 we will be concerned with the noise robustness of the nuclear norm minimization approach for blind deconvolution and matrix completion problem. We will show that these problems are intrinsically badly conditioned. Despite this fact we will be able to show that near-optimal error bounds can be achieved for blind deconvolution, if the noise level is not too small.
In Chapter 6 we will consider the general bilinear inverse problem $y = \mathcal{B}(u, v)$ under the assumption that additional prior information about u and v such as sparsity as available. We will give an overview about the literature. Furthermore, we will discuss an extension of the results of [75], see our results in Appendix B.

Notation

Throughout this thesis we will use the following notation. For $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1; \dots; n\}$. For a complex number $z \in \mathbb{C}$ we will denote its real part by $\mathrm{Re}(z)$. By $\log(\cdot)$ we will denote the logarithm to the base e. By $\mathbb{E}X$ we will denote the expectation of a random variable X and by $\mathbb{P}(A)$ we denote the probability of an event A. For a vector $v \in \mathbb{C}^n$ its ℓ_2-norm will be denoted by $\|v\|$. For vectors $u, v \in \mathbb{C}^n$ the Euclidean inner product will be denoted by $\langle u, v \rangle := u^* v$. Furthermore, for a matrix $Z \in \mathbb{C}^{n_1 \times n_2}$ we denote its spectral norm by $\|Z\|$ and by $\|Z\|_*$ we denote its nuclear norm, i.e., the sum of the singular values of Z. Moreover, the Frobenius norm of Z is defined by $\|Z\|_F$. The corresponding inner product is given by $\langle Z, W \rangle_F := \mathrm{Tr}\,(Z^* W)$, where $W \in \mathbb{C}^{n_1 \times n_2}$. In those parts of the thesis which deal with matrix completion, real-valued matrices $Z \in \mathbb{R}^{n_1 \times n_2}$ will be considered and the previous quantities will be defined analogously.

2. Important measurement models for low-rank matrix recovery

2.1. Blind deconvolution

As already discussed in Chapter 1 blind deconvolution refers to the problem of recovering a signal x from the convolution $w * x + e$, where w is an unknown kernel and e refers to additive noise. In this thesis we will consider for technical reasons the circular convolution, which for $w \in \mathbb{C}^L$ and $x \in \mathbb{C}^L$ is defined by

$$w * x := \left(\sum_{j=1}^{L} w_j x_{k-j} \right)_{k=1}^{L}.$$

Here the difference $k - j$ is considered modulo L. Of course, this is a highly underdetermined inverse problem. For this reason, additional assumptions on w and x are needed. In this dissertation we will follow the model proposed by Ahmed, Recht, and Romberg [1], which is especially relevant for wireless communication as we will explain below. We will assume that w as well as x are contained in some known subspaces. More precisely, we will assume that $w = Bh_0$, where $B \in \mathbb{C}^{L \times K}$ such that $B^*B = \mathrm{Id}$. Furthermore, we assume that $x = C\overline{m}_0$, where C is a random matrix with i.i.d entries with distribution $\mathcal{CN}\left(0, \frac{1}{\sqrt{L}}\right)$.
The choice of the *random* matrix C is motivated by the success of randomization in compressed sensing as well as by applications in wireless communications. Here m is a message to be transmitted and C is a coding matrix. The signal $x = C\overline{m_0}$ gets transmitted through a channel w, which can be modeled as a convolution. In many applications it is reasonable to assume that only the first few entries of w are non-zero as the path delays are often much shorten than the length of the signals x. In this case B would be the matrix which extends $h_0 \in \mathbb{C}^K$ by zeros. Hence, the receiver observes $w * x + e$, where e represents additive noise, and the goal is be to reconstruct the original message m_0.

Now let $F \in \mathbb{C}^{L \times L}$ be the unitary discrete Fourier transformation matrix. It is well-known that F diagonalizes the convolution, i.e.,

$$\widehat{w * x} := F(w * x) = \sqrt{L}\mathrm{diag}(FBh_0) FC\overline{m}_0.$$

Let b_ℓ denote the ℓth row of $\overline{FB}$ and let c_ℓ denote the ℓth row of $\sqrt{L}FC$. Note that this implies that all the entries of $\{c_\ell\}_{\ell=1}^{L}$ are jointly independent and have distribution

$\mathcal{CN}(0,1)$. Moreover, we obtain that

$$\left(\widehat{w * x}\right)_\ell = b_\ell^* h_0 m_0^* c_\ell = \langle b_\ell c_\ell^*, h_0 m_0^* \rangle_F.$$

We observe that $\widehat{w * x}$ is linear in $h_0 m_0^*$. This motivates the definition of the linear operator $\mathcal{A} : \mathbb{C}^{K \times N} \to \mathbb{C}^L$ by

$$\left(\mathcal{A}(X)\right)(\ell) := \langle b_\ell c_\ell^*, X \rangle_F \quad \text{where } \ell \in [L]. \tag{2.1}$$

Hence, we obtain the model

$$y := \widehat{w * x} + e = \mathcal{A}(X_0) + e,$$

where $X_0 = h_0 m_0^*$ and $e \in \mathbb{C}^L$ represents noise with $\|e\| \leq \tau$. Note that X_0 is a rank-one matrix. Hence, in order to estimate X_0 one might choose the solution of the following minimization problem.

$$\begin{aligned} &\text{minimize} && \text{rank}\,(X) \\ &\text{subject to} && \|\mathcal{A}(X) - y\| \leq \tau \end{aligned}$$

Unfortunately, problems of this type are NP-hard in general, as vector cardinality minimization can be considered as a special case [93]. For this reason, in [35] it was proposed to use the nuclear norm $\|\cdot\|_*$ as a proxy for the rank. This leads to the following semidefinite program (SDP), as proposed in [1].

$$\begin{aligned} &\text{minimize} && \|X\|_* \\ &\text{subject to} && \|\mathcal{A}(X) - y\| \leq \tau. \end{aligned} \tag{2.2}$$

Of course, it is crucial to understand under which conditions the solution $\tilde{X}$ to the SDP (2.2) is close to the ground truth $X_0 = h_0 m_0^*$. In particular, one is interested in knowing how large one has to choose L in comparison to K and N such that recovery is possible. Ahmed, Recht, and Romberg established the following result.

Theorem 2.1 ([1])**.** *Consider measurements of the form $y = \mathcal{A}(h_0 m_0^*) + e$ for $h_0 \in \mathbb{C}^K$, $m_0 \in \mathbb{C}^N$, $e \in \mathbb{C}^L$, and $\mathcal{A}$ as defined in (2.1). Assume that $\|e\| \leq \tau$ and*

$$L / \log^3 L \gtrsim K \mu_{\max}^2 + N \max\left\{\mu_{h_0}^2; \tilde{\mu}_{h_0}^2\right\}.$$

Then with probability at least $1 - \mathcal{O}(L^{-1})$ every minimizer $\hat{X}$ of the SDP (2.2) satisfies

$$\|\hat{X} - h_0 m_0^*\|_F \lesssim \frac{\mu_{\max}^2}{\mu_{\min}^2} \sqrt{\min\{K; N\}}\, \tau. \tag{2.3}$$

The coherence parameters $\mu_{\max}^2$, $\mu_{\min}^2$, and $\mu_{h_0}^2$ are defined by

$$\mu_{\max}^2 := \frac{L}{K} \max_{\ell \in [L]} \|b_\ell\|^2$$

$$\mu_{\min}^2 := \frac{L}{K} \min_{\ell \in [L]} \|b_\ell\|^2$$

and

$$\mu_{h_0}^2 := \frac{L}{\|h_0\|^2} \max_{\ell \in [L]} |\langle b_\ell, h_0 \rangle|^2.$$

The parameter $\tilde{\mu}_{h_0}^2$ is a rather technical term and stems from the proof technique, which is the Golfing Scheme (see [80, Remark 2.1] and [56, Remark 2.4] for a discussion of this parameter). It is an open problem, whether this this additional coherence parameter is truly necessary or not.

Note that since the number degrees of this problem is $K + N - 1$ the requirement on the number of measurements L is optimal up to logarithmic factors. Moreover, note that this results ensures exact recovery in the important special case $\tau = 0$. However, note that the appearance of the factor $\sqrt{\min\{K; N\}}$. Such a dimension factor does not appear in other low-rank matrix recovery problems like matrix sensing [91] or phase retrieval [10]. This raises the question, whether nuclear norm minimization is less stable with respect to noise for blind deconvolution than for these other problems. We will provide an answer to this question in Section 5.

It is clear that for many applications solving the SDP (2.2) will be impractical due to its high computational complexity. For this reason several works [78, 51, 81] proposed gradient-based algorithms which operate in the natural parameter domain (instead of the lifted one) and have a much lower computational complexity than the SDP for this reason. Similar performance guarantees have been established for these algorithms in terms of the number of required measurements. In terms of noise these articles considered a *random* noise model and they could derive guarantees for the estimation error without the additional dimension factor which appears in (2.3). However, since we focus here on deterministic noise instead of random noise those two results are not comparable.

2.2. Matrix completion

Suppose we observe only a few entries of a low-rank matrix $X_0 \in \mathbb{R}^{n_1 \times n_2}$. Can one recover the complete matrix from only knowing these few entries? This problem is commonly referred to as matrix completion. Since it arises in many different applications such as multiclass learning [4] and collaborative filtering [100] it has become very popular in the last decade and has been studied intensively in the statistics, machine learning, and signal processing literature.
Let us specify the model, which we will study in this dissertation. In the following we will assume with no loss of generality that $n_1 \geq n_2$. We assume that we observe $[m]$ entries of the matrix X_0, which are sampled with replacement. This means that the linear measurement operator $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \to \mathbb{R}^{n_2}$ is given by

$$\mathcal{A}(X)(i) := \sqrt{\frac{n_1 n_2}{m}} \langle X, e_{a_i} e_{b_i}^* \rangle_F, \tag{2.4}$$

where (a_i, b_i) are i.i.d. random variables with uniform distribution in $[n_1] \times [n_2]$ and e_i denotes the vector, which is equal to 1 in the ith component and equal to zero elsewhere.[1]

We are interested in the setting, where the number of observations is much smaller than the number of total entries of the matrix X_0, i.e., the setting $m \ll n_1 n_2$. However, note that if $X_0 = e_1 e_1^*$ we will with high probability obtain that $\mathcal{A}(X_0) = 0$ and it is impossible to reconstruct the ground truth matrix X_0 from these observations. For this reason, in [15] the following two coherence parameter were introduced, where $X_0 = U\Sigma V^*$ with $U \in \mathbb{R}^{n_1 \times r}$ and $V \in \mathbb{R}^{n_2 \times r}$ denotes the singular value decomposition.

$$\mu(U) := \sqrt{\frac{n_1}{r}} \max_{i \in [n_1]} \|U^* e_i\|$$
$$\mu(V) := \sqrt{\frac{n_2}{r}} \max_{i \in [n_2]} \|V^* e_i\|$$

Indeed, it was shown in [18] that $m \gtrsim n_1 \log n_1 \max\{\mu^2(U); \mu^2(V)\}$ observations are necessary for an rank-r matrix X_0 to be uniquely determined from the revealed entries.

Now assume that we are given noisy observations $y = \mathcal{A}(X_0) + e$ with $\|e\| \leq \tau$. As X_0 has low-rank one could try estimate X_0 by solving the following minimization problem

$$\begin{aligned} &\text{minimize} && \text{rank } (X_0) \\ &\text{subject to} && \|\mathcal{A}(X) - y\| \leq \tau. \end{aligned}$$

However, as already discussed for blind deconvolution, such problems are in general NP-hard. For this reason, on [15] the following approach was proposed for matrix completion.

$$\begin{aligned} &\text{minimize} && \|X\|_* \\ &\text{subject to} && \|\mathcal{A}(X) - y\| \leq \tau. \end{aligned} \tag{2.5}$$

In a series of papers the nuclear norm minimization approach has been studied [15, 18, 46, 98, 22] in the noiseless scenario, which led to the following result

Theorem 2.2 ([22])**.** *Consider measurements of the form $y = \mathcal{A}(X_0)$, where $X_0 \in \mathbb{R}^{n_1 \times n_2}$ is a rank-r matrix and $\mathcal{A}$ is given by (2.4). Assume that*

$$m \geq C \max\left\{\mu^2(U); \mu^2(V)\right\} r n_1 \log^2 n_1.$$

Then with probability at least $1 - \mathcal{O}\left(n_1^{-1}\right)$ the matrix X_0 is the unique minimizer of the SDP (2.5) with $\tau = 0$.

[1] The normalization of $\mathcal{A}$ with the factor $\sqrt{\frac{n_1 n_2}{m}}$ is different to most other works in the literature. We have decided to use this normalization factor in this thesis, because it implies $\mathbb{E}[\mathcal{A}^*\mathcal{A}] = \mathrm{Id}$, which makes the results from blind deconvolution and matrix completion more comparable.

In [30] the result was refined even a bit further. Namely the second log-factor can be removed at the cost of an r^3-dependence in the rank.

All these papers mentioned above, which establish exact recovery. Now assume that we are in the noisy scenario with deterministic noise. Then, in [13] it was shown that with high probability the minimizer $\hat{X}$ of (2.5) satisfies

$$\|\hat{X} - X_0\|_F \lesssim \tau\sqrt{n_2}, \tag{2.6}$$

whenever $m \gtrsim n_1 \text{polylog } n_1$. Similarly, as for the randomized blind deconvolution framework we observe the appearance of an additional dimension factor $\sqrt{n_2}$. In Chapter 5 we will see that these dimension factors reflect in a certain sense the true behavior of the problem.

Let us comment on some related literature. In [63, 61] similar nuclear-norm based estimators as presented in this chapter have been studied under the assumption that the noise is random. It was shown that if the noise follows a subexponential distribution and the noise level is not too small, then near-optimal recovery guarantees are possible. In [27] the authors derived near-optimal rates for an estimator based on nuclear-norm minimization for subgaussian noise, if the rank of the ground truth matrix X_0 is assumed to be constant. For adversarial noise, the strongest result we are aware of, has been shown in [60] for a non-convex algorithm based on Riemannian optimization. However, in contrast to nuclear-norm minimization, their algorithm requires precise knowledge of the rank of the ground truth matrix X_0. Let us also mention that many faster, non-convex algorithms have been proposed and studied for matrix completion (see, e.g., [38, 71, 59, 53, 110, 85, 42]).

3. Sufficient conditions for recovery via nuclear norm minimization

3.1. Outline of this chapter

In the last chapter, we have seen that both randomized blind deconvolution and matrix completion can be formulated as a low-rank matrix recovery problem. Let $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ be a linear map and assume that $y = \mathcal{A}(X_0) + e$, where $X_0 \in \mathbb{C}^{n_1 \times n_2}$ is a low-rank matrix and let $\hat{X}$ be the minimizer of

$$\begin{aligned} \text{minimize} \quad & \|X\|_* \\ \text{subject to} \quad & \|y - \mathcal{A}(X)\| \leq \tau. \end{aligned} \tag{3.1}$$

In this chapter, we will study sufficient conditions for recovery in this very general framework, i.e., when $\|\hat{X} - X_0\|_F$ is small compared to the noise level τ. We will assume throughout this chapter that the noise vector e fulfills $\|e\| \leq \tau$. Moreover, we assume that e may be adversarial noise. By this we mean that (depending on $\mathcal{A}$) e is chosen as the vector, for which the reconstruction error $\|\hat{X} - X_0\|_F$ is maximal.
This chapter has two goals. The first one is to provide mathematical foundations for the following chapters by introducing fundamental notions and reviewing previous proof techniques such as the notion of an approximate dual certificate and the tangent space of a low-rank matrix. The second goal is to understand what gives rise to dimension factors such as $\sqrt{\min\{K, N\}}$ or $\sqrt{n}_2$ in the error bounds for blind deconvolution and matrix completion (see Chapter 2).
Throughout this chapter we will analyse this general framework and, consequently, it is applicable to other low-rank matrix recovery problems such as robust spectral compressed sensing [25].
The methods, which are presented in this chapter, have been developed in several articles [15, 13, 46]. Meanwhile, they have been adapted by many researchers and they are already included in textbooks [92]. What distinguishes our presentation in this chapter to most other literature, is that we aim to understand the noisy scenario and the sharpness of these estimates.

3.2. Exact and approximate dual certificates

In order to formulate sufficient conditions, we first need a characterization of the subdifferential for which we need to introduce some additional notation. Let $X \in \mathbb{C}^{n_1 \times n_2}$ be

a rank-r matrix. Its singular value decomposition will be denoted by $X = U\Sigma V^*$, where $\Sigma \in \mathbb{R}^{r\times r}$ is a diagonal matrix with nonnegative entries and $U \in \mathbb{C}^{n_1\times r}$ and $V \in \mathbb{C}^{n_2\times r}$ are unitary matrices, i.e., $U^*U = V^*V = \mathrm{Id}_r$. We define for the matrix $X \in \mathbb{C}^{n_1\times n_2}$ the tangent space of rank-r manifolds at the point X by

$$T_X := \left\{UA^* + BV^* :\ A \in \mathbb{C}^{n_2\times r}, B \in \mathbb{C}^{n_1\times r}\right\}. \tag{3.2}$$

We will denote the orthogonal projection onto T_X by $\mathcal{P}_{T_X}$ and the orthogonal projection onto $T_X^\perp$, the orthogonal complement of T_X, by $\mathcal{P}_{T_X^\perp}$.

The subdifferential of the nuclear norm at point X will be denoted by $\partial\|\cdot\|_*(X)$. In [118] the subdifferential of the nuclear norm has been characterized by

$$\partial\|\cdot\|_*(X) = \left\{W \in \mathbb{C}^{n_1\times n_2} :\ \mathcal{P}_{T_X}W = UV^*,\ \|\mathcal{P}_{T_X^\perp}W\| \le 1\right\}. \tag{3.3}$$

Now that we have introduced the necessary background we can formulate the following lemma, which gives a sufficient condition for X_0 being the minimizer of (3.1).

Lemma 3.1 ([15])**.** *Assume that we are in the noiseless scenario, i.e., $\tau = 0$. Let $X_0 \in \mathbb{C}^{n_1\times n_2}$ such that $y = \mathcal{A}(X_0)$. Suppose that the following two conditions hold:*

1. *There exists $z \in \mathbb{C}^m$ such that $Y = \mathcal{A}^*(z)$ satisfies*

$$\mathcal{P}_{T_{X_0}}Y = UV^* \text{ and } \|\mathcal{P}_{T_{X_0}^\perp}Y\| < 1$$

2. *the linear operator $\mathcal{A}$ is injective when restricted to the subspace T_{X_0}.*

Then X_0 is the unique minimizer of (3.1)*.*

Lemma 3.1 was first formulated in [15] in the context of low-rank matrix completion, where it was also shown that such a dual certificate exists with high probability given that the amount of measurements is high enough. In [18] it was even shown that for the matrix completion problem such a dual matrix Y can be constructed if the amount of measurements is up to logarithmic factors at the order of the information-theoretic limit. However, a difficulty in constructing such an exact dual certificate is that the dual matrix Y has to fulfill the condition $\mathcal{P}_{T_{X_0}}Y = UV^*$. In [46] David Gross observed that this condition can be relaxed, if the operator $\mathcal{A}$ fulfills some additional properties, which allowed him to simplify the proofs in [15, 18] and to strengthen the results. One of these additional conditions required by Gross is that $\mathcal{A}$ is an approximate isometry when restricted to the subspace T_X.

Definition 3.2. *Let $X \in \mathbb{C}^{n_1\times n_2}$. We say that $\mathcal{A}$ fulfills the δ-restricted isometry property (δ-RIP) on T_X, if for all matrices $Z \in T_X$ it holds that*

$$(1-\delta)\|Z\|_F^2 \le \|\mathcal{A}(Z)\|^2 \le (1+\delta)\|Z\|_F^2.$$

This allows us to state the following proposition,

Proposition 3.3. *[46, 13] Let $X_0 \in \mathbb{C}^{n_1 \times n_2}$ with corresponding singular value decomposition $X_0 = U\Sigma V^*$ and suppose that $y = \mathcal{A}(X_0) + e$ with $\|e\| \leq \tau$. Suppose that the following two conditions hold.*

1. *There exists a vector $z \in \mathbb{C}^{n_1 \times n_2}$ with $\|z\| \leq 2$ such that the matrix $Y = \mathcal{A}^*(z)$ satisfies*
$$\alpha := \|UV^* - \mathcal{P}_{T_{X_0}^\perp} Y\|_F < \frac{1}{8\|\mathcal{A}\|} \text{ and } \|\mathcal{P}_{T_{X_0}^\perp} Y\| < \frac{1}{2} \tag{3.4}$$

2. *the measurement operator $\mathcal{A}$ satisfies the δ-restricted isometry property on T_{X_0} with constant $\delta = \frac{3}{4}$.*

Then every minimizer $\hat{X}$ of (3.1) satisfies

$$\|X_0 - \hat{X}\|_F \lesssim \|\mathcal{A}\|\tau.$$

As already mentioned, in the noiseless scenario, i.e., $\tau = 0$, a first version this result has been stated in [46]. In [13] dual certificates have first been used to obtain error bounds in the noisy case.[1]
Note that in the important special case $\tau = 0$ Proposition 3.3 states that the existence of a dual matrix Y guarantees exact recovery. Such a matrix Y is typically called approximate dual matrix Y. The numerical constants in Proposition 3.3 such as $\frac{1}{8}$, $\delta = \frac{3}{4}$, or $\|z\| \leq 2$ are rather arbitrary and can be replaced by different numbers.

Remark 3.4. *For blind deconvolution it can be shown that (see [56, Lemma 5.3])*

$$\|\mathcal{A}\| \lesssim \sqrt{\frac{L}{KN\mu_{\max}^2} \log(L + KN)}.$$

Hence, if a dual certificate in the form of Proposition 3.3 exists, we obtain that

$$\|X_0 - \hat{X}\|_F \lesssim \sqrt{\frac{L}{KN\mu_{\max}^2} \log(L + KN)}\tau.$$

Note that this bound is slightly sharper than the one presented in [1], see also Theorem 2.1. This is because the authors used slightly different proof techniques to establish the noise bounds than those presented here.

[1] [13] used an exact rather than an approximate dual certicate to establish error bounds. However, it is straightforward to adopt the proof to the situation where an approximate dual certificate is available.

Proof of Proposition 3.3

This section aims to present a proof of Proposition 3.3. A focus of our presentation will be to explore what gives rise to the factor $\|\mathcal{A}\|$ in the noise bound. We will start by proving the following auxiliary lemma.

Lemma 3.5. *Assume that $y = \mathcal{A}(X_0) + e$ with $\|e\| \leq \tau$. Let $\hat{X}$ be a minimizer of (3.1). Set $Z := \hat{X} - X_0$. Assume that there exists a vector $z \in \mathbb{C}^{n_1 \times n_2}$ with $\|z\| \leq 2$ such that the matrix $Y = \mathcal{A}^*(z)$ satisfies (3.4). Then it holds that*

$$\left\| \mathcal{P}_{T_{X_0}^\perp} Z \right\|_* \leq 8\tau + 2\alpha \|\mathcal{P}_{T_{X_0}} Z\|_F,$$

where $\alpha := \|UV^* - \mathcal{P}_{T_{X_0}^\perp} Y\|_F$.

Proof. We start by noticing that

$$\begin{aligned} \|\mathcal{A}(Z)\| &\leq \|\mathcal{A}\left(\hat{X}\right) - y\| + \|y - \mathcal{A}(X_0)\| \\ &\leq 2\tau, \end{aligned} \tag{3.5}$$

where in the first line we used the triangle inequality and in the second line we used that $\hat{X}$ is feasible, i.e., $\|\mathcal{A}\left(\hat{X}\right) - y\| \leq \tau$, and that $\|y - \mathcal{A}(X_0)\| = \|e\| \leq \tau$. Next, choose $\tilde{W} \in T_{X_0}^\perp$ such that $\|\mathcal{P}_{T_{X_0}^\perp} \tilde{W}\| = \frac{1}{2}$ and $\mathrm{Re}\left(\langle \tilde{W}, Z \rangle\right) = \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_*$ holds. This is possible by the duality of the spectral norm and the nuclear norm. Moreover, by the characterization of the subdifferential (3.3) it follows that $UV^* + \mathcal{P}_{T_{X_0}^\perp} Y + \tilde{W} \in \partial\|\cdot\|_*(X_0)$. Hence, we obtain that

$$\begin{aligned} \|X_0\|_* &\overset{(a)}{\geq} \|X_0 + Z\|_* \\ &\overset{(b)}{\geq} \|X_0\|_* + \mathrm{Re}\left(\langle UV^* + \mathcal{P}_{T_{X_0}^\perp} Y + \tilde{W}, Z \rangle_F\right) \\ &\overset{(c)}{=} \|X_0\|_* + \mathrm{Re}\left(\langle UV^* - \mathcal{P}_{T_{X_0}} Y, Z \rangle_F\right) + \mathrm{Re}\left(\langle Y, Z \rangle_F\right) + \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \\ &\overset{(d)}{=} \|X_0\|_* + \mathrm{Re}\left(\langle UV^* - \mathcal{P}_{T_{X_0}} Y, Z \rangle_F\right) + \mathrm{Re}\left(\langle \mathcal{A}^*(z), Z \rangle_F\right) + \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \\ &= \|X_0\|_* + \mathrm{Re}\left(\langle UV^* - \mathcal{P}_{T_{X_0}} Y, \mathcal{P}_{T_{X_0}} Z \rangle_F\right) + \mathrm{Re}\left(\langle z, \mathcal{A}(Z) \rangle\right) + \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \\ &\overset{(e)}{\geq} \|X_0\|_* - \|UV^* - \mathcal{P}_{T_{X_0}} Y\|_F \|\mathcal{P}_{T_{X_0}} Z\|_F - \|z\| \|\mathcal{A}(Z)\| + \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \\ &\overset{(f)}{\geq} \|X_0\|_* - \alpha \|\mathcal{P}_{T_{X_0}} Z\|_F - 2\|\mathcal{A}(Z)\| + \frac{1}{2}\|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \end{aligned}$$

In (a) we used that $X_0 + Z$ is a minimizer of (3.1) and in (b) we used that $UV^* + \mathcal{P}_{T_{X_0}^\perp} Y + \tilde{W} \in \partial\|\cdot\|_*(X_0)$ by construction. Equality (c) follows from our choice of $\tilde{W}$

and in (d) we used that $Y = \mathcal{A}^*(z)$ by assumption. In (e) we applied the Cauchy-Schwarz inequality several times. Inequality (f) follows from the definition of α and from our assumption on $\|z\|$. Rearranging terms combined with inequality (3.5) yields the result. □

Note that if α is sufficiently small, then Lemma 3.5 says that $\left\|\mathcal{P}_{T_{X_0}^\perp} Z\right\|_F$ is (up to a constant) bounded by the noise level. Recall that in order to prove Proposition 3.3 one needs to find an upper bound for $\|Z\|_F = \|\hat{X} - X_0\|_F$. Hence, if we can show that for Z the quantity $\left\|\mathcal{P}_{T_{X_0}^\perp} Z\right\|_F$ is always sufficiently large compared to $\left\|\mathcal{P}_{T_{X_0}} Z\right\|_F$ we would obtain a bound for $\|\hat{X} - X_0\|_F$. This is achieved by the following lemma.

Lemma 3.6. *Let $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ be a linear operator, which satisfies the δ-restricted isometry property on T_{X_0} for some $\delta > 0$ for some $X_0 \in \mathbb{C}^{n_1 \times n_2}$. Assume that $y = \mathcal{A}(X_0) + e$, where $\|e\| \leq \tau$. Let $\hat{X}$ be a feasible point for the SDP, i.e. $\|\mathcal{A}\left(\hat{X}\right) - y\| \leq \tau$. Then for $Z = \hat{X} - X_0$ it holds that*

$$\left\|\mathcal{P}_{T_{X_0}} Z\right\|_F \leq \frac{2\tau}{\sqrt{1-\delta}} + \frac{\|\mathcal{A}\|}{\sqrt{1-\delta}} \|\mathcal{P}_{T_{X_0}^\perp} Z\|_F.$$

Proof. Note that

$$\begin{aligned}
\left\|\mathcal{P}_{T_{X_0}} Z\right\|_F \leq & \frac{1}{\sqrt{1-\delta}} \left\|\mathcal{A}\left(\mathcal{P}_{T_{X_0}} Z\right)\right\| \\
\leq & \frac{1}{\sqrt{1-\delta}} \left\|\mathcal{A}(Z)\right\| + \frac{1}{\sqrt{1-\delta}} \left\|\mathcal{A}\left(\mathcal{P}_{T_{X_0}^\perp} Z\right)\right\| \\
\leq & \frac{2\tau}{\sqrt{1-\delta}} + \frac{\|\mathcal{A}\|}{\sqrt{1-\delta}} \|\mathcal{P}_{T_{X_0}^\perp} Z\|_F.
\end{aligned}$$

In the first line we used the δ-restricted isometry property and in the second line we used the triangle inequality. In the third line we used that $\|\mathcal{A}(Z)\| \leq 2\tau$, where one can argument in exactly the same way as for (3.5). This completes the proof. □

Having gathered all the necessary ingredients we can prove Proposition 3.3.

Proof of Proposition 3.3. Let $\hat{X}$ be the minimizer of (3.1) and set $Z = \hat{X} - X_0$. We note that

$$\begin{aligned}
\|\mathcal{P}_{T_{X_0}^\perp} Z\|_F &\overset{(a)}{\leq} \|\mathcal{P}_{T_{X_0}^\perp} Z\|_* \\
&\overset{(b)}{\leq} 8\tau + 2\alpha \|\mathcal{P}_{T_{X_0}} Z\|_F \\
&\overset{(c)}{\leq} 8\tau + 2\alpha \left(4\tau + 2\|\mathcal{A}\| \left\|\mathcal{P}_{T_{X_0}^\perp} Z\right\|_F\right) \\
&= 8(1+\alpha)\tau + 4\alpha \|\mathcal{A}\| \left\|\mathcal{P}_{T_{X_0}^\perp} Z\right\|_F \\
&\overset{(d)}{\leq} 20\tau + \frac{1}{2} \|\mathcal{P}_{T_{X_0}^\perp} Z\|_F
\end{aligned}$$

In (a) we used that the Frobenius norm is always upper bounded by the nuclear norm. In (b) we used Lemma 3.5 and in (c) we applied Lemma 3.6 as well as the assumption $\delta = \frac{3}{4}$. In (d) we used that $\alpha \leq \frac{1}{8\|\mathcal{A}\|} \leq \frac{1}{4}$ due to $\|\mathcal{A}\| \geq \frac{1}{2}$, which follows from the local isometry property on T_{X_0}.
Rearranging terms it follows that

$$\|\mathcal{P}_{T_{X_0}^{\perp}} Z\|_F \leq 40\tau.$$

Moreover, from Lemma 3.6 combined with the last estimate it follows that

$$\begin{aligned}\|\mathcal{P}_{T_{X_0}} Z\|_F &\leq \frac{2\tau}{\sqrt{1-\delta}} + \frac{\|\mathcal{A}\|}{\sqrt{1-\delta}} \|\mathcal{P}_{T_{X_0}^{\perp}} Z\|_F \\ &\lesssim \|\mathcal{A}\|\tau.\end{aligned}$$

The triangle inequality combined with the last two estimates yields that

$$\begin{aligned}\|Z\|_F &\leq \|\mathcal{P}_{T_{X_0}} Z\|_F + \|\mathcal{P}_{T_{X_0}^{\perp}} Z\|_F \\ &\lesssim \|\mathcal{A}\|\tau.\end{aligned}$$

This completes the proof. □

Remark 3.7. *Let us comment on what determines the estimation error. Lemma 3.5 states that the dual certificate yields for $Z = \hat{X} - X$ the upper bound $\|\mathcal{P}_{T_{X_0}^{\perp}} Z\|_* \leq \tau$, if one neglects the second term (and, indeed, α is often chosen to be rather small). Then the strength of the noise bound depends on what one can say about the relationship between $\|\mathcal{P}_{T_{X_0}^{\perp}} Z\|_*$ and $\|Z\|_F$. Note that the bound becomes stronger, if one can knows that not too much $\|\cdot\|_F$-mass of Z is concentrated on T_{X_0}. Lemma 3.6 gives a bound for how much mass can be concentrated at most on T_{X_0} compared to $T_{X_0}^{\perp}$. However, recall that by Remark 3.4 the norm $\|\mathcal{A}\|$ may be rather large. This is the reason why the additional dimension factors enter.*

3.3. The golfing scheme

In order to be able to use Proposition 3.3 one needs to show that such an approximate dual certificate exists with high probability. In this section we will describe the construction of David Gross [46], often referred to as the golfing scheme. Recall that we have to find $z \in \mathbb{C}^m$ and $Y = \mathcal{A}^*(z)$ which satisfy the following three properties.

(a)
$$\|z\| \leq 2,$$

(b)
$$\alpha = \|UV^* - \mathcal{P}_{T_{X_0}} Y\|_F \leq \frac{1}{8\|\mathcal{A}\|}, \text{ and}$$

(c)
$$\left\|\mathcal{P}_{T_{X_0}^{\perp}}\left(\mathcal{A}^*\left(z\right)\right)\right\| < \frac{1}{2}$$

In randomized blind deconvolution (and also in some other low-rank matrix recovery problems as matrix completion) it can be shown that $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = \mathrm{Id}$. This motivates the choice $\tilde{z}_1 = \mathcal{A}\left(UV^*\right)$ and $\tilde{Y}_1 = \mathcal{A}^*\left(z\right) = \mathcal{A}^*\mathcal{A}\left(UV^*\right)$ for z and Y. Indeed, this would imply that $\mathbb{E}\left[\mathcal{A}^*(\tilde{z}_1)\right] = UV^*$. Consequently, one could then hope to show properties (a), (b), and (c) by applying appropriate concentration inequalities. However, as already mentioned before, for some problems the operator norm $\|\mathcal{A}\|$ can be quite large (as mentioned before, for blind deconvolution $\|\mathcal{A}\|$ it is at the order of $\sqrt{\frac{KN}{L}}$). In particular, this means that the quantity $\left\|UV^* - \mathcal{P}_{T_{X_0}}\mathcal{A}^*\left(z\right)\right\|_F$ needs often to be smaller than what can be guaranteed by concentration inequalities. The idea behind the golfing scheme is to iteratively refine this initial guess until condition (b) is satisfied. It is set up as follows.

- **Step 1:** Choose a partition of $[m]$ into P sets $\{\Gamma_p\}_{p=1}^P$ of roughly the same size, i.e., $|\Gamma_p| \approx \frac{m}{P}$ for all $p \in [P]$, such that $P\mathbb{E}\left[(\mathcal{A}^p)^*\mathcal{A}^p\right] \approx \mathrm{Id}$, where $\mathcal{A}^p := P_{\Gamma_p}\mathcal{A}$. (Here, $P_{\Gamma_p} : \mathbb{C}^m \to \mathbb{C}^m$ denotes the coordinate projection onto Γ_p.)
- **Step 2:** Set

  $$Y_0 = 0 \quad \text{and}$$
  $$Y_p = Y_{p-1} + P(\mathcal{A}^p)^*\mathcal{A}^p\left(UV^* - \mathcal{P}_{T_{X_0}}Y_{p-1}\right) \qquad \text{where } p \in [P].$$

  The corresponding $z \in \mathbb{C}^m$ is then given by

  $$z := P\sum_{p=1}^{P}\mathcal{A}^p\left(UV^* - \mathcal{P}_{T_{X_0}}Y_{p-1}\right).$$

Note that a consequence of the sample splitting in Step 1 is that the golfing scheme is set up in a such a way that $\mathcal{A}^p$ is stochastically independent of Y_{p-1}, which simplifies the analysis or makes it even possible.

4. Blind demixing at near-optimal rate

4.1. Problem framework and main result

Assume that we are given observations of the form

$$y = \sum_{i=1}^{r} w_i * x_i + e \in \mathbb{C}^L, \tag{4.1}$$

where $e \in \mathbb{C}^L$ is additive noise and our goal is to reconstruct all individual signals w_i and x_i from only knowing y. This problem is often referred to as blind demixing and deconvolution and it can be seen as a generalization of the blind deconvolution problem. As in blind deconvolution one needs to make further assumptions on w_i and x_i to enable recovery. As for randomized blind deconvolution we will assume that the w_i and x_i are contained in some known subspaces. More precisely, we will assume that $w_i = B_i m_i$ and $x_i = C_i m_i$ for all $i \in [r]$, where $B_i \in \mathbb{C}^{L \times K_i}$ and $C_i \in \mathbb{C}^{L \times N_i}$. In analogy to the randomized blind deconvolution setting, we will assume that $B_i^* B_i = \mathrm{Id}$ and that the entries of the matrix C_i are independent and identically distributed with distribution $\mathcal{CN}(0,1)$. This framework has been proposed in [80]. It is particularly relevant for applications in wireless communications [117].

As for blind deconvolution we observe that there is unique linear operator $\mathcal{A}_i : \mathbb{C}^{K_i \times N_i} \to \mathbb{C}^L$ such that for all $u \in \mathbb{C}^{K_i}$ and $v \in \mathbb{C}^{N_i}$ it holds that

$$\mathcal{A}_i(uv^*) = Bu * C\overline{v}.$$

With this definition at hand equation 4.1 reads as

$$y = \sum_{i=1}^{r} \mathcal{A}_i(h_i m_i^*) + e.$$

Hence, we have noticed that (4.1) can be recasted as a low-rank matrix recovery problem, where the low-rank matrix to recover is given as

$$X_0 = (h_1 m_1^*, h_2 m_2^*, \ldots, h_r m_r^*).$$

This observation motivated Ling and Strohmer to propose the following SDP in [80].

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^{r} \|X_i\|_* \\ \text{subject to} \quad & \|y - \sum_{i=1}^{r} \mathcal{A}_i(X_i)\| \leq \tau, \end{aligned} \tag{4.2}$$

They could show that if

$$L \gtrsim r^2 \left(K\mu_{\max}^2 + N\mu_h^2 \right) \log^3 L \tag{4.3}$$

holds, where $K := \max_{i\in[r]} K_i$ and $N := \max_{i\in[r]} N_i$, then in the noiseless scenario, i.e., $\tau = 0$, exact recovery is possible. Here, $\mu_{\max}^2$ and μ_h^2 are coherence parameters, which are similar to the ones defined in Section 2.1. In the case of noise they could derive estimation errors with additional dimension factors as in (2.3). However, note that the number of degrees of freedom in this problem is $\sum_{i=1}^r (K_i + N_i) - r$, which raises the question, whether the second r-factor in (4.3) is necessary. Indeed, numerical experiments in [80] indicate that this additional r-factor seems to be an artifact of the proof.

Our main result in [56] states that nuclear norm minimization indeed is able to recover X_0 at near-optimal sampling rate and, hence, this additional r-factor is not necessary. In order to state it we set $N := \max_{i\in[r]} N_i$ and $\mu_{i,\max}$ and by $\mu_{h,\omega}$ we denote coherence parameters which are similar to the coherence parameters introduced in Section 2.1. (For the precise definition of these quantities we refer to [56].)

Theorem 4.1. *[56] Let $\omega \geq 1$ and let $y \in \mathbb{C}^L$ be given by (4.1) with $\|e\| \leq \tau$. Assume that*

$$L \geq C_\omega r \left(\max_{i\in[r]} \left(K_i\mu_{i,\max}^2 \log\left(K_i\mu_{i,\max}^2 \right) \right) + N\mu_{h,\omega}^2 \right) \log^3 L, \tag{4.4}$$

where C_ω is a universal constant only depending on ω. Then, with probability at least $1 - \mathcal{O}\left(L^{-\omega}\right)$ every minimizer $\hat{X} = \left(\hat{X}_1, \ldots, \hat{X}_r\right)$ of the SDP (4.2) satisfies

$$\sqrt{\sum_{i=1}^r \left\| \hat{X}_i - h_i m_i^* \right\|_F^2} \lesssim \tau \sqrt{r \max\left\{ 1; \max_{i\in[r]} \frac{rK_i\mu_{i,\max}^2 N}{L} \right\} \log L}. \tag{4.5}$$

Note that in the important, noiseless scenario with $\tau = 0$ our result states that exact recovery is possible at a near-optimal sampling rate.

To summarize our findings, we have established that there is a tractable, i.e., polynomial-time algorithm, which is able to recover the ground truth $X_0 = (h_1 m_1^*, h_2 m_2^*, \ldots, h_r m_r^*)$ at a near-optimal sampling rate.
However, as lifting increases the number of variables it is clear that the SDP (4.2) will be too slow for many practical purposes. For this reason, in [81] Ling and Strohmer proposed a (non-convex) gradient-based algorithm for this problem, which operates in the original parameter space and, hence, is considerably faster. They were able to establish that their algorithm is able to recover the ground truth if L scales as in (4.3). This means that again their recovery guarantees depend quadratically on r. As in the convex case numerical experiments seem to indicate that the true scaling seems to be

rather linear in r. We believe that our theoretical analysis in [56], which we will outline in the following, might also provide tools to overcome the r^2-dependence for the non-convex approach.

4.2. Outline of the proof

The overall proof strategy of Theorem 4.1 resembles the one provided in Chapter 3. In the first step one shows that stable recovery is guaranteed, if sufficient conditions are fulfilled, among them the existence of an appropriate dual certificate. (Note that Proposition 3.3 is not directly applicable here since we have the sum of operators $\mathcal{A}_i$ rather than a single operator $\mathcal{A}$.) In the second step we construct the dual certificate via an appropriately modified golfing scheme.

This proof strategy of constructing an approximate dual certificate was used by Ling and Strohmer as well as by us. In the following we want to describe the main differences and, in particular, how we were able to remove the second r-factor. Recall that one assumption in Proposition 3.3, which gave sufficient conditions for recovery, was that the operator $\mathcal{A}$ fulfills an approximate isometry property on the tangent space T. Similar properties are needed in [80] as well as by us in [56]. In order to state the assumption used by us, we define for $i \in [r]$ the tangent space of rank-1-matrices at $h_i m_i^*$ by

$$T_i = \left\{ h_i u_i^* + v_i m_i^* : \ u_i \in \mathbb{C}^{K_i}, v_i \in \mathbb{C}^{N_i} \right\}$$

Moreover, we set

$$\tilde{T} := \{ (X_1, \dots, X_r) : \ X_i \in T_i \text{ for all } i \in [r] \} .$$

This allows us to give the following definition of a δ-local isometry, which is analogous to Definition 3.2.

Definition 4.2 (Local isometry property)**.** *The operators $\{\mathcal{A}_i\}_{i=1}^r$ fulfill the δ-local isometry property on $\tilde{T}$ for some $\delta > 0$, if*

$$(1-\delta) \|X\|_F^2 \leq \Big\| \sum_{i=1}^{r} \mathcal{A}_i (X_i) \Big\|^2 \leq (1+\delta) \|X\|_F^2 \tag{4.6}$$

for all $X = (X_1, \dots, X_r) \in \tilde{T}$.

The main difference between [80] and our work [56] is that in [80] the restricted isometry property is not established on T directly, but individually on each T_i and it is shown that the images of the subspaces T_i under the operator $\mathcal{A}_i$ are near-orthogonal to each other. This is responsible for the appearance of the r^2 in their analysis. In contrast, our more global analysis allows us to establish that the local isometry property holds on T, if L scales linearly with r.

Setting

$$\hat{T} := \tilde{T} \bigcap \left\{ (X_1, \dots, X_r) \in \mathbb{C}^{K_1 \times N_1} \times \dots \times \mathbb{C}^{K_r \times N_r} : \sum_{i=1}^{r} \|X_i\|_F^2 = 1 \right\}$$

we observe that by its definition the δ-local isometry property is equivalent to the inequality

$$\begin{aligned} \delta &\geq \sup_{(X_1,\dots,X_r)\in\hat{T}} \left| \left\| \sum_{i=1}^{r} \mathcal{A}_i (X_i) \right\|^2 - \sum_{i=1}^{r} \|X_i\|_F^2 \right| \\ &= \sup_{(X_1,\dots,X_r)\in\hat{T}} \left| \left\| \sum_{i=1}^{r} \mathcal{A}_i (X_i) \right\|^2 - \sum_{i=1}^{r} \mathbb{E}\left[\left\| \mathcal{A}_i (X_i) \right\|^2 \right] \right|. \\ &= \sup_{(X_1,\dots,X_r)\in\hat{T}} \left| \left\| \sum_{i=1}^{r} \mathcal{A}_i (X_i) \right\|^2 - \mathbb{E}\left[\left\| \sum_{i=1}^{r} \mathcal{A}_i (X_i) \right\|^2 \right] \right|. \end{aligned}$$

The second line is due to the assumption that $B_i^* B_i = \mathrm{Id}$ for all $i \in [r]$ and the third line follows from the fact that the random operators $\mathcal{A}_i$ are independent with zero mean (see [56, Proof of Proposition 5.8] for a more detailed calculation).
Now let $\mathrm{vec}([C_1, \dots, C_r])$ denote the vectorization of the random matrices $C_1, \dots, C_r$, i.e., the vector which contains all the entries of the matrices $(C_1, \dots, C_r)$. We observe that the expression $\sum_{i=1}^{r} \mathcal{A}_i (X_i)$ is linear in $\mathrm{vec}([C_1, \dots, C_r])$, which means that we can find for each $X = (X_1, \dots, X_r) \in \hat{T}$ a matrix V_X such that the equation

$$\sum_{i=1}^{r} \mathcal{A}_i (X_i) = V_X \, \mathrm{vec}([C_1, \dots, C_r])$$

always holds. Hence, in order to show the δ-local isometry property it is sufficient to prove that

$$\delta \geq \sup_{X\in\hat{T}} \left| \left\| V_X \, \mathrm{vec}([C_1, \dots, C_r]) \right\|^2 - \mathbb{E}\left[\left\| V_X \, \mathrm{vec}([C_1, \dots, C_r]) \right\|^2 \right] \right|. \tag{4.7}$$

To obtain upper bounds suprema of this type recently developed results can be used [65]. In order to state them we will first need to introduce the notion of Talagrand's γ_2-functional (see, e.g., [111]). In a certain sense Talagrand's γ_2-functional measures the complexity of a set with respect to a given metric.

Definition 4.3. *Let $(X, |||\cdot|||)$ be a Banach space and suppose that $S \subset X$. We say that a sequence $(S_n)_{n\geq 0}$ of subsets of S is admissible, if $|S_0| = 1$ and $|S_n| \leq 2^{2^n}$ for all $n \geq 1$. Then we set*

$$\gamma_2 (S, |||\cdot|||) = \inf_{(S_n)_{n\geq 0}} \sup_{s\in S} \sum_{n=0}^{\infty} 2^{n/2} \inf_{s\in S_n} |||s - s_n|||,$$

where the infimum is taken over all admissible sequences $(S_n)_{n\geq 0}$.

Moreover, for a set $\mathcal{X}$ of matrices we define

$$d_F(\mathcal{X}) := \sup\{\|X\|_F : \ X \in \mathcal{X}\},$$
$$d_{2\to 2}(\mathcal{X}) := \sup\{\|X\| : \ X \in \mathcal{X}\}.$$

This allows us to state the following theorem by Krahmer, Mendelson, and Rauhut, which will allow us to show (4.7).

Theorem 4.4. *[66, Theorem 1.4] Let $\mathcal{X}$ be a symmetric set of matrices, i.e., $\mathcal{X} = -\mathcal{X}$ and let ξ be a random vector whose entries ξ_i are independent circular-symmetric standard normal random variables with mean 0 and variance 1. Set*

$$E = \gamma_2(\mathcal{X}, \|\cdot\|)(\gamma_2(\mathcal{X}, \|\cdot\|) + d_F(\mathcal{X}))$$
$$V = d_{2\to 2}(\mathcal{X})(\gamma_2(\mathcal{X}, \|\cdot\|) + d_F(\mathcal{X}))$$
$$U = d^2_{2\to 2}(\mathcal{X})$$

Then, for $t > 0$,

$$\mathbb{P}\left(\sup_{A\in\mathcal{X}} \big|\|A\xi\|^2 - \mathbb{E}\|A\xi\|^2\big| \geq c_1 E + t\right) \leq 2\exp\left(-c_2 \min\left(\frac{t^2}{V^2}, \frac{t}{U}\right)\right).$$

The constants c_1 and c_2 are universal.

Our goal is to apply this theorem to $\mathcal{X} := \left\{V_X : X \in \hat{T}\right\}$, which is possible as the entries of $\mathrm{vec}([C_1, \ldots, C_r])$ have independent circular-symmetric standard distribution with mean 0 and variance 1. Note that in order to make the bounds in Theorem 4.4 effective we need to get appropriate upper bounds for $\gamma_2(\mathcal{X})$. However, in many cases it is rather difficult to deal with the γ_2-functional directly. The following inequality by Dudley states the γ_2-functional can be bounded from above by an integral involving covering numbers.

Lemma 4.5. *Let $(V, |||\cdot|||)$ be a normed vector space and let $\mathcal{X} \subset V$. It holds that*

$$\gamma_2(\mathcal{X}, |||\cdot|||) \lesssim \int_0^{d_{|||\cdot|||}(S)} \sqrt{\log N(S, |||\cdot|||, \varepsilon)}\, d\varepsilon,$$

where $N(S, |||\cdot|||, \varepsilon)$ denotes the covering number of S with respect to the metric induced by $|||\cdot|||$, i.e., the minimum number of open $|||\cdot|||$-balls with radius ε, whose midpoint is in S, which are necessary in order to cover S.

Hence, we can apply Dudley's inequality to $\mathcal{X}$ (with respect to the operator norm) and are left with estimating the arising covering numbers. This is rather involved and requires different tools as Maurey's lemma [19] and a theorem by Artstein, Milman, and Szarek [5], which solved a special case of a conjecture by Pietsch [97] concerning the duality of covering numbers. For details we refer the interested reader to [56].

5. Noise robustness of blind deconvolution and matrix completion

5.1. Descent cone analysis

We have seen that in problems like blind deconvolution (see inequality (2.3)), blind demixing (see inequality (4.5)), and matrix completion (see inequality (2.6)) noise bounds of the form

$$\|\hat{X} - X_0\|_F \lesssim \sqrt{\min\{K; N\}}\tau,$$

appear, when they are solved via nuclear norm minimization. Here X_0 denotes the ground truth and $\hat{X}$ denotes the minimizer of the SDP to be solved. The appearance of additional dimension factors such as $\sqrt{\min\{K; N\}}$ is rather suprising as it does not appear in other problems as matrix sensing [99], phase retrieval [11, 10], and compressed sensing (see, e.g., [14, 40]). The goal of this chapter is to examine this phenomenon more closely and to examine whether the dimension factors are truly necessary or whether they are a mere proof artifact of the golfing scheme. Our analysis will be based on the understanding of the minimum conic singular value of the descent cone. This approach has been first studied for ℓ_1-minimization by Rudelson and Vershynin in [103] and for more general atomic norms in [21]. In recent years many new results and insights for sparse and low-rank matrix recovery have been obtained in this way (see, e.g., [2, 70, 57, 73, 31, 88]).

Let us shortly describe the main ideas of this geometric analysis. The descent cone of the nuclear norm is the set of all directions where the nuclear norm does not increase.

Definition 5.1. *For any matrix $X_0 \in \mathbb{C}^{K\times N}$ define its descent cone $\mathcal{K}_*(X_0)$ by*

$$\mathcal{K}_*(X_0) := \left\{ Z \in \mathbb{C}^{K\times N} : \ \|X_0 + \varepsilon Z\|_* \leq \|X_0\|_* \text{ for some } \varepsilon > 0 \right\}.$$

In order to understand why the descent cone is relevant assume first that $y = \mathcal{A}(X_0)$ and $\tau = 0$, i.e., we are in the noiseless scenario. Then we observe that X_0 is the unique solution of the SDP ((3.1)), if and only if the descent cone $\mathcal{K}_*(X_0)$ intersects with the null space of the linear map $\mathcal{A}$ only at 0.
To understand the noisy scenario assume that $y = \mathcal{A}(X_0) + e$, where $\|e\| \leq \tau$. Let $\hat{X}$ be a minimizer of (3.1). As X_0 is a feasible point of the SDP (3.1), i.e., $\|y - \mathcal{A}(X_0)\| \leq \tau$, it follows that $\|\hat{X}\|_* \leq \|X_0\|_*$. Moreover, note that the constraint $\|\mathcal{A}(X) - y\| \leq \tau$ defines a region around the point X_0. Hence, the minimizer $\hat{X}$ must lie in the intersection of

this region with the set $X_0 + \mathcal{K}_*(X_0)$, i.e., the shifted descent cone. The size of this region can be quantified using by the minimum conic singular value, which is defined as follows.

$$\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(X_0)) := \inf_{Z \in \mathcal{K}_*(X_0) \setminus \{0\}} \frac{\|\mathcal{A}(Z)\|}{\|Z\|_F}.$$

The size of this intersection should become smaller if the minimum conic singular value increases and, consequently, better noise estimates should be available. The following result by Chandrasekaran et al. confirms this intuition.

Theorem 5.2. *[21, Proposition 2.2] Let $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ be a linear operator and assume that $y = \mathcal{A}(X_0) + e$ with $\|e\| \leq \tau$. Then any minimizer $\hat{X}$ of the SDP (3.1) satisfies*

$$\|\hat{X} - X_0\|_F \leq \frac{2\tau}{\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(X_0))}.$$

5.2. Bad conditioning of blind deconvolution

In problems like matrix sensing (if suitably normalized) it can be shown that with high probability it holds that $\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(X_0)) \gtrsim 1$ [21], which leads to (order-wise) optimal noise bounds via Theorem 5.2. For this reason one might expect that such a result can also be shown for blind deconvolution. This is not the case for all matrices $B \in \mathbb{C}^{K \times N}$ as shown by the following proposition.

Proposition 5.3. *[68] Let $K, N \in \mathbb{N} \setminus \{1\}$. Assume that L is an integer multiple of K and that*

$$C_1 K \leq L \leq \frac{KN}{9}. \tag{5.1}$$

Then there exists $B \in \mathbb{C}^{L \times K}$ satisfying $B^*B = Id_K$ and $\mu^2_{\max} = 1$, whose corresponding measurement operator $\mathcal{A}$ satisfies the following.

Let $h_0 \in \mathbb{C}^K \setminus \{0\}$, $m_0 \in \mathbb{C}^N \setminus \{0\}$ and set $\mu^2 := \mu_{h_0} = L \frac{\max_{\ell \in [L]} |\langle b_\ell, h_0 \rangle|^2}{\|h_0\|^2}$. Then with probability at least $1 - \mathcal{O}\left(\exp\left(-C_2 K/\mu^2\right)\right)$ it holds that

$$\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(h_0 m_0^*)) \leq C_3 \sqrt{\frac{L}{KN}}. \tag{5.2}$$

Here C_1, C_2, and C_3 are absolute constants.

As we are interested in the scenario that $L \ll KN$ this proposition says that the minimum conic singular value is indeed very small and, consequently, Theorem 5.2 will not be able to yield strong noise bounds.
One might ask whether whether the result holds true for any matrix $B \in \mathbb{C}^{K \times N}$ such that $B^*B = \mathrm{Id}$ and $\mu^2_{\max}$ is small. We think that such a result does not hold true in general. Indeed, one might choose B as a random embedding. Using a similar analysis

as in [70] we think that one should be able to show that $\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(X_0\right)\right) \gtrsim 1$ in this scenario.
Note that inequality (5.2) is equivalent to the statement that there is $Z \in \mathcal{K}_*\left(X_0\right)$ such that

$$\frac{\|\mathcal{A}\left(Z\right)\|}{\|Z\|_F} \lesssim \sqrt{\frac{L}{KN}}. \tag{5.3}$$

holds. In order to find such a Z we make use of the observation that with high probability there is a rank-one matrix $W \in \mathbb{C}^{K\times N}$, which is an element in the null-space of $\mathcal{A}$, but whose distance, measured in $\|\cdot\|_F$-norm, to the descent cone $\mathcal{K}_*\left(X_0\right)$ is relatively small. Setting $Z := W - \beta h_0 m_0^*$ and choosing β appropriately, one can show that Z possesses is an element of the descent cone of the nuclear norm at X_0. In order to show the latter, we have established the following lemma in [68], which gives a characterization of the descent cone of the nuclear norm.

Lemma 5.4. *[68] Let $X \in \mathbb{C}^{K\times N}\setminus\{0\}$ be a matrix of rank r with corresponding singular value decomposition $X = U\Sigma V^*$. Then*

$$\overline{\mathcal{K}_*\left(X\right)} = \left\{Z \in \mathbb{C}^{K\times N} : -Re\left(\langle UV^*, Z\rangle_F\right) \geq \|\mathcal{P}_{T_X^\perp}\left(Z\right)\|_*\right\},$$

where $\overline{\mathcal{K}_*\left(X\right)}$ denotes the topological closure of $\mathcal{K}_*\left(X\right)$.

A consequence of Proposition 5.3 is the following statement, which shows that blind deconvolution is unstable in the following sense.

Theorem 5.5. *[68] Let $K, N \in \mathbb{N}\setminus\{1\}$. Assume that L is an integer multiple of K and that*

$$C_1 K \leq L \leq \frac{KN}{9}.$$

Then there exists a matrix $B \in \mathbb{C}^{L\times K}$ satisfying $B^*B = Id_K$ and with FB having rows of equal norm, i.e., $\mu_{\max}^2 = 1$, such that for all $h_0 \in \mathbb{C}^K \setminus \{0\}$ and $m_0 \in \mathbb{C}^N \setminus \{0\}$ the following holds:
With probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{K}{C_2\mu^2}\right)\right)$, where $\mu^2 = \mu_{h_0}^2 = \frac{L}{\|h_0\|^2}\max\limits_{\ell\in[L]} |\langle b_\ell, h_0\rangle|^2$, there is $\tau_0 > 0$ such that for all $\tau \leq \tau_0$ there exists an adversarial noise vector $e \in \mathbb{C}^L$ with $\|e\| \leq \tau$ that admits an alternative solution $\tilde{X}$ with the following properties.

- *$\tilde{X}$ is feasible, i.e., $\|\mathcal{A}\left(\tilde{X}\right) - y\| = \tau$ for $y = \mathcal{A}\left(h_0 m_0^*\right) + e$ the noisy measurement vector*
- *$\tilde{X}$ is preferred to $X_0 = h_0 m_0^*$ by the SDP (2.2), i.e., $\|\tilde{X}\|_* \leq \|X_0\|_*$, but*
- *$\tilde{X}$ is far from the true solution in Frobenius norm, i.e.,*

$$\|\tilde{X} - X_0\|_F \geq \frac{\tau}{C_3}\sqrt{\frac{KN}{L}}. \tag{5.4}$$

The constants C_1, C_2, and C_3 are universal.

Note that the inequality dimension factor in inequality (5.4) matches with the one in (4.5) (up to a log-factor), when applied with $r = 1$. Hence, the additional dimension factors in (2.3) are no proof artifacts in the sense that there is an alternative solution, which is preferred to the ground truth, but is far from the ground truth in $\|\cdot\|_F$-distance.

However, note that the matrix $\tilde{X}$ in Theorem 5.5 might, in general, not be the minimizer of the SDP (2.2). Indeed, if the matrix $\tilde{X}$ as constructed as in [68], it will never by the minimizer of this SDP (see the discussion in [68, Remark 3,4]).

Remark 5.6. *It might be interesting to go back to the approach using dual certificates with these insights. Note that the stability proof using approximate dual certificates in Chapter 3 does not distinguish between $\hat{X}$ and $\tilde{X}$. However, from Proposition 5.8 we know that any proof, which does not make this distinction, cannot yield error estimates below the barriers set by inequality* (5.4)*.*
A closer look into the proof of Proposition 5.3 reveals that for the matrix Z, which we constructed in order to show (5.3)*, $\|\mathcal{P}_{T_{X_0}^\perp} Z\|_*$ is rather small compared to $\|Z\|_F$. This makes Lemma 3.5 ineffective for such Z. Hence, a potential strategy to improve previous noise bounds would be to show that for $Z = \hat{X} - X_0$, where $\hat{X}$ is the actual minimizer, the concentration of the mass with respect to $\|\cdot\|_F$ on T_{X_0} cannot occur.*

5.3. Noise robustness of blind deconvolution

In the last section we have seen that blind deconvolution is intrinsically badly conditioned. Nevertheless, in the situation, where the noise-level is not too small, near-optimal recovery guarantees are possible as our next result shows.

Theorem 5.7. *[68] Let $\alpha > 0$ and $B \in \mathbb{C}^{L\times K}$ such that $B^*B = Id$. Assume that*

$$L \geq C_1 \frac{\mu^2}{\alpha^2} (K + N) \log^2 L.$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{L\alpha^{4/3}}{C_2 \log^{4/3}(eL)\mu^{4/3}}\right)\right)$ the following statement holds for all $h_0 \in \mathcal{H}_\mu \setminus \{0\}$, all $m_0 \in \mathbb{C}^N \setminus \{0\}$, all $\tau > 0$, and all $e \in \mathbb{C}^L$ with $\|e\| \leq \tau$: Any minimizer $\hat{X}$ of (2.2) satisfies

$$\|\hat{X} - h_0 m_0^*\|_F \leq \frac{C_3 \mu^{2/3} \log^{2/3} L}{\alpha^{2/3}} \max\left\{\tau; \alpha \|h_0 m_0^*\|_F\right\}.$$

Here C_1, C_2, and C_3 are absolute constants.

The crucial insight for the proof of Theorem 5.7 is that the bad conditioning, i.e., $\frac{\|\mathcal{A}(Z)\|}{\|Z\|_F}$ is small, arises only for descent directions Z, which are near-orthogonal to the

ground truth. Due to this we partition the descent cone $\mathcal{K}_*(h_0 m_0)$ into two cones $\mathcal{K}_1$ and $\mathcal{K}_2$, where $\mathcal{K}_1$ contains all the directions, which are almost orthogonal to the ground truth $h_0 m_0^*$. On the one hand one can show that due to the curvature of the nuclear norm ball, only a small error can occur from these near-orthogonal directions $Z \in \mathcal{K}_1$. On the other hand one can show that $\lambda_{\min}(\mathcal{A}, \mathcal{K}_2)$ is with high probability at the order of a constant (up to log-factors and ignoring the μ-dependence), which allows on to controll the error which can arise from the directions contained in $\mathcal{K}_2$. In order to show the latter we have used Mendelson's small ball method [64, 87], a tool which allows one to bound nonnegative empirical processes from below.

5.4. Bad conditioning of matrix completion

In this section we want to discuss briefly the noise robustness of matrix completion. It turns out that similarly as randomized blind deconvolution the matrix completion problem is badly conditioned in the sense of the following proposition.

Proposition 5.8. *[68] Let $X_0 \in \mathbb{R}^{n_1 \times n_2} \setminus \{0\}$ be a rank-r matrix with corresponding singular value decomposition $X_0 = U\Sigma V^*$. Moreover, assume that*

$$C_1 r n_1 \mu^2(V) \log(2r) \leq m \leq \frac{n_1 n_2}{32}. \tag{5.5}$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{m}{C_2 r \mu^2(U)\mu^2(V)}\right)\right)$ it holds that

$$\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(X_0)) \leq C_3 \sqrt{\frac{m}{n_1 n_2 r}}. \tag{5.6}$$

The constants C_1, C_2, and C_3 are universal.

The proof of Proposition 5.8 follows the same strategy as the correspoding proposition for blind deconvolution. One tries to construct a matrix $Z \in \mathcal{K}_*(X_0)$ such that $\frac{\|\mathcal{A}(Z)\|}{\|Z\|_F}$ is small enough. For that one finds a rank-one matrix W in the null space of $\mathcal{A}$, which is close to the tangent space T_{X_0} in $\|\cdot\|_F$-distance. By an appropriate perturbation one can then construct a matrix $Z \in \mathcal{K}_*(X_0)$ with the required properties.

Proposition also leads to the following theorem, which is analogous to Theorem 5.5 in the randomized blind deconvolution framework.

Theorem 5.9. *[68] Let $n_1 \geq n_2$ and let $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \to \mathbb{R}^m$ be defined as in (2.4). Assume that $X_0 \in \mathbb{R}^{n_1 \times n_2} \setminus \{0\}$ is a rank r matrix with singular value decomposition $X_0 = U\Sigma V^*$. Moreover, assume that*

$$C_1 r n_1 \mu^2(V) \log(2r) \leq m \leq \frac{n_1 n_2}{32}.$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{m}{C_2 r \mu^2(U)\mu^2(V)}\right)\right)$ there is $\tau_0 > 0$ such that for all $\tau \leq \tau_0$ there exists an adversarial noise vector $e \in \mathbb{R}^m$ with $\|e\| \leq \tau$ that admits an alternative solution $\tilde{X} \in \mathbb{R}^{n_1 \times n_2}$ with the following properties.

- *$\tilde{X}$ is feasible, i.e., $\left\|\mathcal{A}\left(\tilde{X}\right) - y\right\| = \tau$ for $y = \mathcal{A}(X_0) + e$ the noisy measurement vector*
- *$\tilde{X}$ is preferred to X_0 by the SDP (2.5), i.e., $\|\tilde{X}\|_* \leq \|X_0\|_*$, but*
- *$\tilde{X}$ is far from the true solution in Frobenius norm, i.e.,*

$$\|\tilde{X} - X_0\|_F \geq \frac{\tau}{C_3}\sqrt{\frac{rn_1n_2}{m}}.$$

Here the constants C_1, C_2, and C_3 are universal.

6. Bilinear inverse problems with sparsity constraints

6.1. Overview

In the previous chapters we have discussed blind deconvolution, which is a bilinear inverse problem. In such problems one tries to recover vectors $u \in \mathbb{C}^{n_1}$ and $v \in \mathbb{C}^{n_2}$ simultanously from measurements of the form

$$y = \mathcal{B}(u, \overline{v}),$$

where $\mathcal{B} : \mathbb{C}^{n_1} \times \mathbb{C}^{n_2} \to \mathbb{C}^m$ is a bilinear map. We note that not only randomized blind deconvolution falls under this framework but also other signal processing problems like self-calibration [79]. In the previous chapters we have seen that blind deconvolution, can be solved via a convex relaxation in an average-case sense, if the number of measurements m is up to log-factors at the order of the number of degrees of freedom $n_1 + n_2$. However, in many applications additional information about the structure of u and v is available. For example, in blind deconvolution the channel w might be sparse. Another example would be blind sensor calibration (see, e.g., [9, 79]). Here u is an unknown tuning parameters of the sensor and v represents the signal to be measured, which might be sparse. In order to keep the amount of needed measurements as small as possible it would be desirable to develop algorithms, which take this additional information about the structure into account. Throughout this chapter we will assume that $u \in \mathbb{C}^{n_1}$ is s_1-sparse and $v \in \mathbb{C}^{n_2}$ is s_2-sparse. (We say that a vector x is s-sparse, if it contains at most s nonzero entries.) Ideally, we would like to be able to recover u and v from y, when the number of measurements m is up to log-factors at the order of $s_1 + s_2$, i.e., the number of degrees of freedom.

As in the previous chapters we denote by $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ the unique linear operator, which fulfills

$$\mathcal{A}(uv^*) = B(u, \overline{v})$$

for all $u \in \mathbb{C}^{n_1}$ and $v \in \mathbb{C}^{n_2}$. In the following we will assume that the linear operator $\mathcal{A}$ can be represented as

$$(\mathcal{A}(X))(\ell) := \frac{1}{\sqrt{m}} \langle A_i, X \rangle_F \quad \text{for all } \ell \in [m], \tag{6.1}$$

where $\{A_i\}_{i=1}^m \subset \mathbb{C}^{n_1 \times n_2}$ are random matrices, whose entries are i.i.d. with distribution $\mathcal{CN}(0, 1)$. While such a model is rather unrealistic in most practical scenarios, we will

focus on this model in the following because it often allows for an easier analysis and as we will discuss in the following even for this simplified scenario basic questions are not yet understood.

In [96, Theorem 3] it was shown that uv^* is the minimizer of a certain non-convex program with high probability whenever $m \gtrsim (s_1 + s_2) \max\left\{\frac{n_1}{s_1}; \frac{n_2}{s_2}\right\}$. However, it is not clear whether this minimizer can be found in polynomial time. While such a result says that in principle enough information is contained in the measurements to uniquely specify the signal uv^* it is also clear that for practical purposes one needs algorithms, which determine uv^* in a reasonable amount of time. For this reason our focus here will be on tractable, i.e., polynomial-time algorithms.
Note that uv^* is a rank-one matrix, whose factors u and v are s_1- and s_2-sparse. In particular, this means that uv^* has at most $s_1 s_2$ non-zero entries. As the ℓ_1-norm has been successful in recovering sparse vectors and the nuclear norm $\|\cdot\|_*$ has been successful in recovering low-rank matrices one might try to combine these two norms, which leads to the following tractable minimization problem

$$\begin{aligned} \text{minimize} \quad & \|X\|_* + \lambda \|X\|_{\ell_1} \\ \text{subject to} \quad & y = \mathcal{A}(uv^*)\,. \end{aligned}$$

However, in [96] it was shown that this approach will not perform better than just one of the two relaxations. More precisely, this means that in order to achieve recovery still $m \gtrsim \min\left\{n_1 + n_2; s_1 s_2 \log\left(\max\left(\frac{n_1}{s_1}; \frac{n_2}{s_2}\right)\right)\right\}$ measurements are necessary. In [101] a much tighter, norm-based convex relaxation was presented. While it was shown that, whenever $m \gtrsim (s_1 + s_2)\log(\max(n_1, n_2))$, the matrix uv^* will be the unique minimizer of this convex relaxation, it is not clear whether the minimizer can be computed in polynomial time.

In [52, 6] it was shown that recovery at near-optimal sampling rate can be achieved with tractable algorithms, if the measurement model (6.1) is replaced by a cleverly constructed *nested* measurement model. However, as such a measurement model is unrealistic for most practical purposes it is clear that one still would like to understand the more generic measurement model (6.1).

6.2. Non-convex approaches

We have seen that it is an open question whether there exists a convex method, which achieves recovery at near-optimal sample complexity and enjoys polynomial computational complexity at the same time. For this reason one might resort to non-convex algorithms based on variants of gradient descent. Indeed, in the last few years there have been huge advances in the study of such non-convex gradient based algorithms. In particular, it has been shown that for low-rank matrix recovery problems like matrix completion, blind deconvolution, or phase retrieval (without sparsity assumptions)

these algorithms often perform equally well in terms of sample complexity as convex programs but enjoy significantly reduced computational complexity (see the overview article [28]).
However, it is often difficult to analyse gradient descent based algorithms if one starts from an arbitrary initialization as saddle-points or local minima may exist. For this reason these algorithms often follow a two-stage approach: One uses a starting point obtained via spectral initialization, which is then subsequently refined by a gradient-descent based approach. Such a two-stage approach has been pioneered by Keshavan et al. in [59] for the matrix completion problem and since then it became very popular in the literature. It has then been adopted to study efficient non-convex algorithms for phase retrieval [94, 12, 23], matrix completion [53, 49, 110, 42, 38, 71, 85], blind deconvolution [78, 81, 51] and other problems [24].[1]
Following this line of work Lee et al. [77] proposed such a two-stage algorithm for bilinear inverse problem with sparsity constraints, which they dubbed Sparse Power Factorization. In the case of full Gaussian measurements as in (6.1) they could show that a certain, rather restricted restricted class of sparse vectors can be recovery with a sample complexity at the order of $(s_1 + s_2) \log\left(\max\left\{\frac{n_1}{s_1}; \frac{n_2}{s_2}\right\}\right)$.
We will describe the precise formulation of the algorithm and their result below. Before that let us comment on some related work. One algorithm, which is closely related to the one by Lee et. al, has been proposed by Fornasier et al. [37]. However, comparable global performance guarantees seem not to be available at this point.
Furthermore, some papers tried to replace full Gaussian measurement matrices by a more realistic measurement matrices. In [75] Lee et al. studied performance guarantees of a variation of the Sparse Power Factorization for blind deconvolution with two random subspaces. The analysis of this problem becomes significantly more challenging than the one of full Gaussian measurements as one has to deal with coherence issues. In order to establish theoretical results, the authors of [75] had to add a projection step to the algorithm, which cannot be guaranteed to be computed in polynomial time.
In [104, 7] two different tractable approaches were proposed for the related problem of sparse phase retrieval. In both cases it was shown that if a good initialization is available, then at the order of $s_1 + s_2$ measurements (up to log-factors) are enough to recover the ground truth. As we will see finding a good initialization will also be the bottleneck in the Sparse Power Factorization algorithm below.

6.3. Sparse Power Factorization

Let us describe the main idea behind the spectral initialization step, which aims to give a good first guess for the ground truth. In the case of full Gaussian measurements (as in (6.1)) it can be shown that $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = \mathrm{Id}$. In particular, this implies that $\mathbb{E}\left[\mathcal{A}^*(y)\right] = uv^*$

[1] In the last few years it has been more and more understood that one can expect convergence to the ground truth for such problems in many cases even from a random initialization as the loss landscape often enjoys favorable properties [42, 109, 26].

due to $y = \mathcal{A}(uv^*)$. Hence, if the number of measurement is high enough one could hope that $\mathcal{A}^*\mathcal{A}(uv^*) \approx uv^*$ holds and one might choose the leading singular vectors of $\mathcal{A}^*(y)$ as a first guess for u and v. However, this approach requires $m \gtrsim n_1 + n_2$ samples and recall that we assume that we are in the underdetermined case where $m \ll n_1 + n_2$. This motivates the following approach which takes the structure of the problem into account.

$$\begin{aligned} \max \quad & \mathrm{Re}\left(\tilde{u}^* \mathcal{A}^*(y)\, \tilde{v}\right) \\ \text{subject to} \quad & \|\tilde{u}\|_0 \le s_1, \|\tilde{u}\| = 1 \\ & \|\tilde{v}\|_0 \le s_2, \|\tilde{v}\| = 1, \end{aligned} \tag{6.2}$$

This problem is often referred to as sparse principal component analysis (SparsePCA). However, solving (6.2) is in general NP-hard [112]. Sparse PCA has been studied rather intensly the last decade. Several computationally tractable algorithms have been proposed [33, 55, 86], where the theoretical analysis was in particular focused on the single-spike model [69, 3, 29].
In [8, 117] it was shown that, under a certain model, a tractable algorithm, which solves the sparse PCA problem, would lead to an algorithm, which solves the k-clique problem in polynomial time. The existence of such an algorithm is an open problem in theoretical computer science and it is widely believed that such an algorithm does not exist. This indicates that an approach based on a spectral initialization might not be able yield an algorithm, which recovers sparse vectors in a bilinear model at near-optimal sampling rate.

In [77] the following algorithm was proposed as a tractable substitute for (6.2).

Algorithm 1 (Algorithm 3 in [77])**.**

Input: *Operator $\mathcal{A}$, Measurement b, Sparsity Constraints s_1, s_2,*
Output: *Initial guess v_0 for $v \in \mathbb{C}^{n_2}$.*

1. *For all $i \in [n_1]$ let ξ_i be the ℓ_2-norm of the best s_2-sparse approximation of the ith row of the matrix $\mathcal{A}^*(b) \in \mathbb{C}^{n_1 \times n_2}$.*
2. *Let $\widehat{J_1} \subset [n_1]$ be the set of the s_2 largest elements in $\{\xi_1; \xi_2; \ldots; \xi_{n_1}\}$*
3. *Choose $\widehat{J}_2$ to contain the indices of the s_2 columns of $\Pi_{\widehat{J}_1} \mathcal{A}^*(b)$ largest in ℓ_2 norm, i.e.,*

$$\widehat{J}_2 := \underset{J \subset [n_2],\ |J| = s_2}{\text{arg max}} \left\| \Pi_{\widehat{J}_1} [\mathcal{A}^*(b)] \Pi_J \right\|_{\mathrm{F}}. \tag{6.3}$$

4. ***return*** *v_0, the leading right singular vector of $\Pi_{\widehat{J}_1} [\mathcal{A}^*(b)] \Pi_{\widehat{J}_2}$.*

Having obtained such an initial guess v_0, the Sparse Power Factorization algorithm iteratively proceeds by keeping v_0 fixed. Note that this results in an underdetermined linear system for u. In order to solve this system one uses the Hard Thresholding Pursuit algorithm [39]. Having obtained a guess u_1 for the ground truth u one keeps $\tilde{u}$ fixed and solves the resulting underdetermined system for v by the Hard Thresholding Pursuit. In this way one alternates several times between $\tilde{u}$ and $\tilde{v}$ until a stopping criterion is met.

Theoretical results

Lee et al. established the following result in [77] for the Sparse Power Factorization algorithm.

Theorem 6.1 ([77, see Theorem 1 and 4])**.** *Assume that* $\mathcal{A}\colon \mathbb{C}^{n_1 \times n_2} \longrightarrow \mathbb{C}^m$ *is a Gaussian linear operator as described above. Let* $b = \mathcal{A}(uv^*) + z$*, where* u *is* s_1*-sparse and* v *is* s_2*-sparse. Suppose that* $\|u\|_\infty \geq 0.78\|u\|$*,* $\|v\|_\infty \geq 0.78\|v\|$*, and that the noise level satisfies* $\nu(z) \leq 0.04$*. Then, with probability exceeding* $1 - \exp(-c_1 m)$*, the output of the Sparse Power Factorization algorithm (with initialization given by Algorithm 1) converges linearly to* uv^* *provided that*

$$m \geq c_2 (s_1 + s_2) \log\left(\max\left\{\frac{n_1}{s_1}, \frac{n_2}{s_2}\right\}\right),$$

where $c_1, c_2 > 0$ *are absolute constants.*

Note that the requirement $\|u\|_\infty \geq 0.78\|u\|$ and $\|v\|_\infty \geq 0.78\|v\|$ means that most of the mass in ℓ_2-norm of the vectors u and v is concentrated on one singly entry each. Hence, this result covers only a rather restricted class of signals. The main reason why this restriction appears in [77] is that only with this restriction the initialization can be guaranteed to be in the basin of attraction, i.e. a region to ground truth, where convergence of the Sparse Power Factorization algorithm can be established. Our main result in [44], Theorem 2, states that this class of possible vectors can be enlargened, if one is willing to increase the number of measurements slightly. Our results allows for replacing the number 0.78 by a smaller number. Furthermore, instead of having one peaky entry it allows for having several entries, on which most of the ℓ_2-mass of u is concentrated.

Bibliography

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

[2] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: phase transitions in convex programs with random data. *Inf. Inference*, 3(3):224–294, 2014.

[3] A. A. Amini and M. J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Statist.*, 37(5B):2877–2921, 2009.

[4] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

[5] S. Artstein, V. Milman, and S. J. Szarek. Duality of metric entropy. *Ann. of Math.*, pages 1313–1328, 2004.

[6] S. Bahmani and J. Romberg. Near-optimal estimation of simultaneously sparse and low-rank matrices from nested linear measurements. *Inf. Inference*, 5(3):331–351, 2016.

[7] S. Bahmani and J. Romberg. Solving equations of random convex functions via anchored regression. *Found. Comput. Math.*, Oct 2018.

[8] Q. Berthet and P. Rigollet. Optimal detection of sparse principal components in high dimension. *Ann. Statist.*, 41(4):1780–1815, 2013.

[9] Ç. Bilen, G. Puy, R. Gribonval, and L. Daudet. Convex optimization approaches for blind sensor calibration using sparsity. *IEEE Trans. Signal Process.*, 62(18):4847–4856, 2014.

[10] E. J. Candès and X. Li. Solving quadratic equations via phaselift when there are about as many equations as unknowns. *Found. Comput. Math.*, 14(5):1017–1026, 2014.

[11] E. J. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.*, 39(2):277–299, 2015.

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

[13] E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

[14] E. J. Candes and Y. Plan. A probabilistic and ripless theory of compressed sensing. *IEEE Trans. Inform. Theory*, 57(11):7235–7254, 2011.

[15] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717, 2009.

[16] E. J. Candès, J. K. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, 52(2):489–509, 2006.

[17] E. J. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274, 2013.

[18] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2010.

[19] B. Carl. Inequalities of Bernstein-Jackson-type and the degree of compactness of operators in Banach spaces. *Ann. Inst. Fourier*, 35(3):79–118, 1985.

[20] T. F. Chan and C.-K. Wong. Total variation blind deconvolution. *IEEE Trans. Signal Process.*, 7(3):370–375, 1998.

[21] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Found. Comput. Math.*, 12(6):805–849, 2012.

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

[23] Y. Chen and E. J. Candès. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Commun. Pure Appl. Math.*, 70(5):822–883, 2017.

[24] Y. Chen and E. J. Candès. The projected power method: an efficient algorithm for joint alignment from pairwise differences. *Commun. Pure Appl. Math.*, 71(8):1648–1714, 2018.

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

[26] Y. Chen, Y. Chi, J. Fan, and C. Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Math. Program.*, pages 1–33.

[27] Y. Chen, Y. Chi, J. Fan, C. Ma, and Y. Yan. Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. *arXiv preprint arXiv:1902.07698*, 2019.

[28] Y. Chi, Y. M. Lu, and Y. Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *arXiv preprint arXiv:1809.09573*, 2018.

[29] Y. Deshpande and A. Montanari. Sparse pca via covariance thresholding. In *Advances in Neural Information Processing Systems*, pages 334–342, 2014.

[30] L. Ding and Y. Chen. The leave-one-out approach for matrix completion: Primal and dual analysis. *arXiv preprint arXiv:1803.07554*, 2018.

[31] S. Dirksen, G. Lecué, and H. Rauhut. On the gap between restricted isometry properties and sparse recovery conditions. *IEEE Trans. Inform. Theory*, 64(8):5478–5487, 2018.

[32] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.

[33] A. d'Aspremont, F. Bach, and L. E. Ghaoui. Optimal solutions for sparse principal component analysis. *J. Mach. Learn. Res.*, 9(Jul):1269–1294, 2008.

[34] C. Ekanadham, D. Tranchina, and E. P. Simoncelli. A blind sparse deconvolution method for neural spike identification. In *Advances in Neural Information Processing Systems*, pages 1440–1448, 2011.

[35] M. Fazel, H. Hindi, S. P. Boyd, et al. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American control conference*, volume 6, pages 4734–4739. Citeseer, 2001.

[36] C. Fienup and J. Dainty. Phase retrieval and image reconstruction for astronomy. *Image Recovery: Theory and Application*, 231:275, 1987.

[37] M. Fornasier, J. Maly, and V. Naumova. Sparse pca from inaccurate and incomplete measurements. *arXiv preprint arXiv:1801.06240*, 2018.

[38] M. Fornasier, H. Rauhut, and R. Ward. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM J. Optim.*, 21(4):1614–1640, 2011.

[39] S. Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM J. Numer. Anal.*, 49(6):2543–2563, 2011.

[40] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Birkhäuser Basel, 2013.

[41] B. Friedlander and T. Strohmer. Bilinear compressed sensing for array self-calibration. In *2014 48th Asilomar Conference on Signals, Systems and Computers*, pages 363–367. IEEE, 2014.

[42] R. Ge, J. D. Lee, and T. Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.

[43] J. Geppert, F. Krahmer, and D. Stöger. Sparse power factorization: Balancing peakiness and sample complexity. *Adv. Comput. Math.*, 2019.

[44] J. A. Geppert, F. Krahmer, and D. Stöger. Refined performance guarantees for sparse power factorization. In *International Conference on Sampling Theory and Applications (SampTA)*, pages 509–513. IEEE, 2017.

[45] D. Godard. Self-recovering equalization and carrier tracking in two-dimensional data communication systems. *IEEE transactions on communications*, 28(11):1867–1875, 1980.

[46] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inform. Theory*, 57(3):1548–1566, 2011.

[47] D. Gross, F. Krahmer, and R. Kueng. A partial derandomization of phaselift using spherical designs. *J. Fourier Anal. Appl.*, 21(2):229–266, 2015.

[48] D. Gross, F. Krahmer, and R. Kueng. Improved recovery guarantees for phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.*, 42(1):37–64, 2017.

[49] M. Hardt. Understanding alternating minimization for matrix completion. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 651–660. IEEE, 2014.

[50] R. W. Harrison. Phase problem in crystallography. *JOSA a*, 10(5):1046–1055, 1993.

[51] W. Huang and P. Hand. Blind deconvolution by a steepest descent algorithm on a quotient manifold. *SIAM J. Imaging Sci.*, 11(4):2757–2785, 2018.

[52] M. Iwen, A. Viswanathan, and Y. Wang. Robust sparse phase retrieval made easy. *Appl. Comput. Harmon. Anal.*, 42(1):135–142, 2017.

[53] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 665–674, New York, NY, USA, 2013. ACM.

[54] S. M. Jefferies and J. C. Christou. Restoration of astronomical images by iterative blind deconvolution. *The Astrophysical Journal*, 415:862, 1993.

[55] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.*, 11(Feb):517–553, 2010.

[56] P. Jung, F. Krahmer, and D. Stöger. Blind demixing and deconvolution at near-optimal rate. *IEEE Trans. Inform. Theory*, 64(2):704–727, 2018.

[57] M. Kabanava, R. Kueng, H. Rauhut, and U. Terstiege. Stable low-rank matrix recovery via null space properties. *Inf. Inference*, 5(4):405–441, 2016.

[58] M. Kech and F. Krahmer. Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems. *SIAM J. Appl. Algebra Geom.*, 1(1):20–37, 2017.

[59] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Trans. Inf. Theory*, 56(6):2980–2998, 2010.

[60] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *J. Mach. Learn. Res.*, 11:2057–2078, 2010.

[61] O. Klopp. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.

[62] Z. Koldovskỳ and P. Tichavskỳ. Time-domain blind audio source separation using advanced ica methods. In *Eighth Annual Conference of the International Speech Communication Association*, 2007.

[63] V. Koltchinskii, K. Lounici, A. B. Tsybakov, et al. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, 39(5):2302–2329, 2011.

[64] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN*, 2015(23):12991–13008, 2015.

[65] F. Krahmer, S. Mendelson, and H. Rauhut. Suprema of chaos processes and the restricted isometry property. *Comm. Pure Appl. Math.*, 67(11):1877–1904, 2014.

[66] F. Krahmer, S. Mendelson, and H. Rauhut. Suprema of chaos processes and the restricted isometry property. *Commun. Pure Appl. Math.*, 67(11):1877–1904, 2014.

[67] F. Krahmer and D. Stöger. Blind deconvolution: Convex geometry and noise robustness. In *52nd Annual Asilomar Conference on Signals, Systems, and Computers*, 2018.

[68] F. Krahmer and D. Stöger. On the convex geometry of blind deconvolution and matrix completion. *arXiv preprint arXiv:1902.11156*, 2019.

[69] R. Krauthgamer, B. Nadler, and D. Vilenchik. Do semidefinite relaxations solve sparse pca up to the information limit? *Ann. Statist.*, 43(3):1300–1322, 2015.

[70] R. Kueng, H. Rauhut, and U. Terstiege. Low rank matrix recovery from rank one measurements. *Appl. Comput. Harmon. Anal.*, 42(1):88–116, 2017.

[71] C. Kümmerle and J. Sigl. Harmonic mean iteratively reweighted least squares for low-rank matrix recovery. *J. Mach. Learn. Res.*, 19:49, 2018.

[72] E. Y. Lam and J. W. Goodman. Iterative statistical approach to blind image deconvolution. *JOSA A*, 17(7):1177–1184, 2000.

[73] G. Lecué and S. Mendelson. Sparse recovery under weak moment assumptions. *J. Eur. Math. Soc. (JEMS)*, 19(3):881–904, 2017.

[74] K. Lee, F. Krahmer, and J. Romberg. Spectral methods for passive imaging: Nonasymptotic performance and robustness. *SIAM J. Imaging Sci.*, 11(3):2110–2164, 2018.

[75] K. Lee, Y. Li, M. Junge, and Y. Bresler. Blind recovery of sparse signals from subsampled convolution. *IEEE Trans. Inform. Theory*, 63(2):802–821, 2017.

[76] K. Lee, N. Tian, and J. Romberg. Fast and guaranteed blind multichannel deconvolution under a bilinear system model. *IEEE Trans. Inform. Theory*, 64(7):4792–4818, 2018.

[77] K. Lee, Y. Wu, and Y. Bresler. Near-optimal compressed sensing of a class of sparse low-rank matrices via sparse power factorization. *IEEE Trans. Inform. Theory*, 64(3):1666–1698, 2018.

[78] X. Li, S. Ling, T. Strohmer, and K. Wei. Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *Appl. Comput. Harmon. Anal.*, 2018.

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

[80] S. Ling and T. Strohmer. Blind deconvolution meets blind demixing: Algorithms and performance bounds. *IEEE Trans. Inform. Theory*, 63(7):4497–4520, 2017.

[81] S. Ling and T. Strohmer. Regularized gradient descent: a non-convex recipe for fast joint blind deconvolution and demixing. *Inf. Inference*, 2017.

[82] S. Ling and T. Strohmer. Self-calibration and bilinear inverse problems via linear least squares. *SIAM J. Imaging Sci.*, 11(1):252–292, 2018.

[83] B. Logan. *Properties of high-pass signals*. PhD thesis, Columbia University.

[84] M. Lustig, D. Donoho, and J. M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

[85] C. Ma, K. Wang, Y. Chi, and Y. Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467*, 2017.

[86] Z. Ma. Sparse principal component analysis and iterative thresholding. *Ann. Statist.*, 41(2):772–801, 2013.

[87] S. Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39, 2014.

[88] S. Mendelson, H. Rauhut, R. Ward, et al. Improved bounds for sparse recovery from subsampled random convolutions. *Ann. Appl. Probab.*, 28(6):3491–3527, 2018.

[89] O. Michailovich and A. Tannenbaum. Blind deconvolution of medical ultrasound images: A parametric inverse filtering approach. *IEEE Trans. Image Process.*, 16(12):3005–3019, 2007.

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

[91] K. Mohan and M. Fazel. New restricted isometry results for noisy low-rank recovery. In *2010 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 1573–1577. IEEE, 2010.

[92] A. Moitra. *Algorithmic aspects of machine learning*. Cambridge University Press, 2018.

[93] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.

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

[95] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. *IEEE Trans. Signal Process.*, 63(18):4814–4826, 2015.

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

[97] A. Pietsch. Theorie der Operatorenideale . Jena: Eigenverlag der Friedrich-Schiller-Universität. 260 S. M 17.00 (1972)., 1972.

[98] B. Recht. A simpler approach to matrix completion. *J. Mach. Learn. Res.*, 12(Dec):3413–3430, 2011.

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

[100] J. D. M. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22Nd International Conference on Machine Learning*, ICML '05, pages 713–719, New York, NY, USA, 2005. ACM.

[101] E. Richard, G. R. Obozinski, and J.-P. Vert. Tight convex relaxations for sparse matrix factorization. In *Advances in neural information processing systems*, pages 3284–3292, 2014.

[102] J. M. Rodenburg. Ptychography and related diffractive imaging methods. *Advances in imaging and electron physics*, 150:87–184, 2008.

[103] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Commun. Pure Appl. Math.*, 61(8):1025–1045, 2008.

[104] M. Soltanolkotabi. Structured signal recovery from quadratic measurements: breaking sample complexity barriers via nonconvex optimization. *IEEE Trans. Inf. Theory*, 65(4):2374–2400, 2019.

[105] D. Stöger, J. Geppert, and F. Krahmer. Sparse power factorization with refined peakiness conditions. In *IEEE Statistical Signal Processing Workshop (SSP)*, pages 816–820. IEEE, 2018.

[106] D. Stöger, P. Jung, and F. Krahmer. Blind deconvolution and compressed sensing. In *4th International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa)*, pages 24–27. IEEE, 2016.

[107] D. Stöger, P. Jung, and F. Krahmer. Blind demixing and deconvolution with noisy data at near optimal rate. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 10394, 2017.

[108] D. Stöger, P. Jung, and F. Krahmer. Blind demixing and deconvolution with noisy data: Near-optimal rate. In *WSA 2017; 21th International ITG Workshop on Smart Antennas*, pages 1–5. VDE, 2017.

[109] J. Sun, Q. Qu, and J. Wright. A geometric analysis of phase retrieval. *Found. Comput. Math.*, 18(5):1131–1198, 2018.

[110] R. Sun and Z.-Q. Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Trans. Inform. Theory*, 62(11):6535–6579, 2016.

[111] M. Talagrand. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media, 2006.

[112] A. M. Tillmann and M. E. Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *IEEE Trans. Inf. Theory*, 60(2):1248–1259, 2014.

[113] J. A. Tropp. Book review of: S. Foucart and H. Rauhut, A mathematical introduction to compressive sensing. *Bull. Am. Math. Soc., New Ser.*, 54(1):151–165, 2017.

[114] S. Vasanawala, M. Murphy, M. T. Alley, P. Lai, K. Keutzer, J. M. Pauly, and M. Lustig. Practical parallel imaging compressed sensing mri: Summary of two years of experience in accelerating body mri of pediatric patients. In *2011 IEEE International Symposium on biomedical Imaging: From nano to macro*, pages 1039–1043. IEEE, 2011.

[115] I. Waldspurger. Phase retrieval with random Gaussian sensing vectors by alternating projections. *IEEE Trans. Inf. Theory*, 64(5):3301–3312, 2018.

[116] I. Waldspurger, A. d'Aspremont, and S. Mallat. Phase recovery, MaxCut and complex semidefinite programming. *Math. Program.*, 149(1-2 (A)):47–81, 2015.

[117] X. Wang and H. V. Poor. Blind equalization and multiuser detection in dispersive cdma channels. *IEEE Trans. Commun.*, 46(1):91–103, 1998.

[118] G. Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra Appl.*, 170:33–45, 1992.

A. Core article: Blind Demixing and Deconvolution at Near-Optimal Rate

Summary

Suppose one is given the vector

$$y = \sum_{i=1}^{r} w_i * x_i + e \in \mathbb{C}^L.$$

How can one reconstruct $\{w_i\}_i^r$ and $\{x_i\}_i^r$? This ill-posed inverse problem is commonly referred to as blind demixing and deconvolution. In order to make this problem more well-posed one typically imposes further structural constraints on w_i and x_i. In this paper we assume that $w_i = Bh_i$ and $x_i = Cm_i$ for all $i \in [r]$, where $\{B_i\}_{i=1}^r \subset \mathbb{C}^{L\times K}$ are deterministic matrices and $\{C_i\}_{i=1}^r \subset \mathbb{C}^{L\times N}$ are random matrices.

This model was first studied by Ling and Strohmer in [80]. They proposed to solve the following convex relaxation to estimate $\{h_i\}_{i=1}^r$ and $\{m_i\}_{i=1}^r$.

$$\min \sum_{i=1}^{r} \left\| X_i \right\|_* \quad \text{s.t.} \quad \left\| \sum_{i=1}^{r} \mathcal{A}_i(X_i) - y \right\|_{\ell_2} \leq \tau. \tag{A.1}$$

(See also [56, Section I.B] where this optimization problem is stated.) In [80] it was shown that this algorithm recovers the ground truth with high probability, whenever

$$L \gtrsim r^2 \left(K\mu_{\max}^2 + N\mu_h^2 \right),$$

where we have set $K = \max_{i\in[r]} K_i$ and $N = \max_{i\in[r]} N_i$. and $\mu_{\max}^2$ and μ_h^2 are coherence parameters. Note that the number of degrees of freedom is $\sum_{i=1}^r K_i + \sum_{i=1}^r N_i - r$, which means that the result above is suboptimal by a factor of r. Indeed, numerical experiments in [80] suggested that this additional factor of r might be an artifact of their proof.

Our main result in [56] confirms this observation: If

$$L \gtrsim r \left(K\mu_{\max}^2 + N\mu_h^2 \right)$$

the SDP (A.1) recovers the ground truth with high probability. (For a precise statement of our main result we refer to [56, Theorem 1]). The proof of this result relies on recently developed tools for bounding the suprema of chaos processes [65].

Statement of individual contribution

Felix Krahmer assigned me with the task of applying the results in [65] to the blind deconvolution framework as in [1] to proof a recovery result based on the restricted isometry property rather than on the Golfing Scheme. While we were not successful, I observed that this approach can be used to overcome some of the difficulties, which Ling and Strohmer encountered in their work.
I was responsible for working out the technical details, where I was stimulated from discussions with Peter Jung and Felix Krahmer. I was fully responsible for writing and preparing this manuscript except for the numerical simulations in Section III, which were performed by Peter Jung. Peter Jung, Felix Krahmer, and I proofread and polished the article together.

Conference proceedings

The results of this paper have been presented in part at the 2016 International Workshop on CoSeRa 2016 [106], in part at the 2017 21st International ITG Workshop on Smart Antenna [108], and in part at the Conference Wavelets and Sparsity XVII [107].

IEEE COPYRIGHT FORM

To ensure uniformity of treatment among all contributors, other forms may not be substituted for this form, nor may any wording of the form be changed. This form is intended for original material submitted to the IEEE and must accompany any such material in order to be published by the IEEE. Please read the form carefully and keep a copy for your files.

Blind Demixing and Deconvolution at Near-Optimal Rate
Peter Jung , Felix Krahmer , Dominik St\u00F6ger
IEEE Transactions on Information Theory

COPYRIGHT TRANSFER

The undersigned hereby assigns to The Institute of Electrical and Electronics Engineers, Incorporated (the "IEEE") all rights under copyright that may exist in and to: (a) the Work, including any revised or expanded derivative works submitted to the IEEE by the undersigned based on the Work; and (b) any associated written or multimedia components or other enhancements accompanying the Work.

GENERAL TERMS

1. The undersigned represents that he/she has the power and authority to make and execute this form.
2. The undersigned agrees to indemnify and hold harmless the IEEE from any damage or expense that may arise in the event of a breach of any of the warranties set forth above.
3. The undersigned agrees that publication with IEEE is subject to the policies and procedures of the IEEE PSPB Operations Manual.
4. In the event the above work is not accepted and published by the IEEE or is withdrawn by the author(s) before acceptance by the IEEE, the foregoing copyright transfer shall be null and void. In this case, IEEE will retain a copy of the manuscript for internal administrative/record-keeping purposes.
5. For jointly authored Works, all joint authors should sign, or one of the authors should sign as authorized agent for the others.
6. The author hereby warrants that the Work and Presentation (collectively, the "Materials") are original and that he/she is the author of the Materials. To the extent the Materials incorporate text passages, figures, data or other material from the works of others, the author has obtained any necessary permissions. Where necessary, the author has obtained all third party permissions and consents to grant the license above and has provided copies of such permissions and consents to IEEE

BY TYPING IN YOUR FULL NAME BELOW AND CLICKING THE SUBMIT BUTTON, YOU CERTIFY THAT SUCH ACTION CONSTITUTES YOUR ELECTRONIC SIGNATURE TO THIS FORM IN ACCORDANCE WITH UNITED STATES LAW, WHICH AUTHORIZES ELECTRONIC SIGNATURE BY AUTHENTICATED REQUEST FROM A USER OVER THE INTERNET AS A VALID SUBSTITUTE FOR A WRITTEN SIGNATURE.

Dominik Stoger	17-12-2017
Signature	**Date (dd-mm-yyyy)**

Information for Authors

AUTHOR RESPONSIBILITIES

The IEEE distributes its technical publications throughout the world and wants to ensure that the material submitted to its publications is properly available to the readership of those publications. Authors must ensure that their Work meets the requirements as stated in section 8.2.1 of the IEEE PSPB Operations Manual, including provisions covering originality,

authorship, author responsibilities and author misconduct. More information on IEEE's publishing policies may be found at http://www.ieee.org/publications_standards/publications/rights/authorrightsresponsibilities.html Authors are advised especially of IEEE PSPB Operations Manual section 8.2.1.B12: "It is the responsibility of the authors, not the IEEE, to determine whether disclosure of their material requires the prior consent of other parties and, if so, to obtain it." Authors are also advised of IEEE PSPB Operations Manual section 8.1.1B: "Statements and opinions given in work published by the IEEE are the expression of the authors."

RETAINED RIGHTS/TERMS AND CONDITIONS

- Authors/employers retain all proprietary rights in any process, procedure, or article of manufacture described in the Work.
- Authors/employers may reproduce or authorize others to reproduce the Work, material extracted verbatim from the Work, or derivative works for the author's personal use or for company use, provided that the source and the IEEE copyright notice are indicated, the copies are not used in any way that implies IEEE endorsement of a product or service of any employer, and the copies themselves are not offered for sale.
- Although authors are permitted to re-use all or portions of the Work in other works, this does not include granting third-party requests for reprinting, republishing, or other types of re-use.The IEEE Intellectual Property Rights office must handle all such third-party requests.
- Authors whose work was performed under a grant from a government funding agency are free to fulfill any deposit mandates from that funding agency.

AUTHOR ONLINE USE

- **Personal Servers**. Authors and/or their employers shall have the right to post the accepted version of IEEE-copyrighted articles on their own personal servers or the servers of their institutions or employers without permission from IEEE, provided that the posted version includes a prominently displayed IEEE copyright notice and, when published, a full citation to the original IEEE publication, including a link to the article abstract in IEEE Xplore. Authors shall not post the final, published versions of their papers.
- **Classroom or Internal Training Use.** An author is expressly permitted to post any portion of the accepted version of his/her own IEEE-copyrighted articles on the author's personal web site or the servers of the author's institution or company in connection with the author's teaching, training, or work responsibilities, provided that the appropriate copyright, credit, and reuse notices appear prominently with the posted material. Examples of permitted uses are lecture materials, course packs, e-reserves, conference presentations, or in-house training courses.
- **Electronic Preprints.** Before submitting an article to an IEEE publication, authors frequently post their manuscripts to their own web site, their employer's site, or to another server that invites constructive comment from colleagues. Upon submission of an article to IEEE, an author is required to transfer copyright in the article to IEEE, and the author must update any previously posted version of the article with a prominently displayed IEEE copyright notice. Upon publication of an article by the IEEE, the author must replace any previously posted electronic versions of the article with either (1) the full citation to the IEEE work with a Digital Object Identifier (DOI) or link to the article abstract in IEEE Xplore, or (2) the accepted version only (not the IEEE-published version), including the IEEE copyright notice and full citation, with a link to the final, published article in IEEE Xplore.

Questions about the submission of the form or manuscript must be sent to the publication's editor.
Please direct all questions about IEEE copyright policy to:
IEEE Intellectual Property Rights Office, copyrights@ieee.org, +1-732-562-3966

Frequently Asked Questions Regarding IEEE Permissions

- **When is permission to reuse IEEE required?**

As a general rule, IEEE requires permission be sought to reproduce any substantial part of its intellectual property, including any text, illustrations, charts, tables, photographs, or other material from previously published sources used. IEEE also requires that all references or sources used be credited, whether or not permission is required. For further guidance, please contact pubs-permissions@ieee.org.

- **From whom do I need permission?**

Permission must be sought from IEEE to reuse its intellectual property. In most cases this will mean locating the material you wish to reuse in IEEE Xplore, where you will find a "request permission" link either on the Table of Contents or on the Article Abstract Page.

- **What if I do not see the "Request Permission" link on either the Table of Contents or the Abstract Page in Xplore?**

If you do not see a permission link on the Abstract Page, we recommend you review the front cover and/or the copyright page in the document itself (often, these pages are freely available for viewing in Xplore) in order to determine copyright owner. If you are unsure, please contact pubs-permissions@ieee.org.

- **Does IEEE require individuals working on a thesis or dissertation to obtain formal permission for reuse?**

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you must follow the requirements listed below:

Textual Material

Using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.

In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.

If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Full-Text Article

If you are using the entire IEEE copyright owned article, the following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]

Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis on-line.

In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

- **If I want to republish an article in another language do I still need to obtain a license from IEEE?**

If you are republishing IEEE intellectual property, we do require you obtain a license that includes any translations. The required translation disclaimer and other translation guidelines are available in the IEEE Terms and Conditions contained in the license provided by the Copyright Clearance Center (RightsLink service).

- **How do I obtain permission to use photographs or illustrations?**

IEEE does not always own reproduction rights to photographs or illustrations; rather, such rights may have been retained by the photographer or illustrator. If the source from which the material is borrowed does not indicate who owns reproduction rights, users of these photographs or illustrations are required to locate the rightsholder, directly.

- **Do I need to obtain permission to use IEEE material posted on its website?**

Yes. As a general rule, most material found on the internet is protected by copyright law even if a notice is not displayed. IEEE does require that you inquire about such permission before using any material found on all IEEE copyright owned websites.

- **Does IEEE require certain rights when requesting permission to use material in an IEEE work?**

IEEE does allow permission to reuse small portions of text in another IEEE copyright owned document only (e.g., the equivalent of several paragraphs only) and figures. Our only requirement is that you 1) provide full credit information pertaining to the original IEEE publications (e.g., author name, paper title, publication title, month and year of original publication). Requests for permission to reuse larger portions should be sent to pubs-permissions@ieee.org .

- **What is Rightslink®?**

Rightslink® is the Copyright Clearance Center's automated permissions granting service, which is used by IEEE along with many other STM publishers such as Springer, Elsevier, and Taylor & Francis. Through this permission service, customers can request permission for IEEE Periodical and Conference content from the point of access; (normally found on the abstract page of the individual article, in IEEE Xplore).

- **Is IEEE an STM signatory publisher?**

No, IEEE is not a signatory to the STM (International Association of Scientific, Technical & Medical Publishers) Permissions Guidelines, last updated February 2012.

April 2013, nbd

Blind Demixing and Deconvolution at Near-Optimal Rate

Peter Jung, *Member, IEEE*, Felix Krahmer, and Dominik Stöger

***Abstract*— We consider simultaneous blind deconvolution of r source signals from their noisy superposition, a problem also referred to *blind demixing and deconvolution*. This signal processing problem occurs in the context of the Internet of Things where a massive number of sensors sporadically communicate only short messages over unknown channels. We show that robust recovery of message and channel vectors can be achieved via convex optimization when random linear encoding using i.i.d. complex Gaussian matrices is used at the devices and the number of required measurements at the receiver scales with the degrees of freedom of the overall estimation problem. Since the scaling is linear in r our result significantly improves over recent works.**

***Index Terms*— Blind deconvolution, channel estimation, demixing, compressed sensing, semidefinite programming.**

I. Introduction

RECENT progress regarding recovery problems for low-complexity structures in high-dimensional data have shown that a substantial reduction in sampling and storage complexity can be achieved in many relevant non–adaptive linear signal separation and estimation problems, in particular in the case of randomized strategies. This includes the recovery of sparse and compressible vectors (often referred to as *compressed sensing*) [4], [5], low–rank matrices [6], and higher–order tensors from subsampled linear measurements [7], as well as compressive demixing of multiple source signals [8]. An important step in many of such vector and matrix recovery problems is to establish computational tractability in the sense of complexity theory; a common strategy to achieve this is to show that, under appropriate assumptions on the measurement map, the reconstruction problem can be recast as a tractable convex program.

Manuscript received May 2, 2017; accepted September 21, 2017. Date of publication December 27, 2017; date of current version January 18, 2018. This work was supported by German Science Foundation (DFG) in the context of the joint project Bilinear Compressed Sensing (JU2795/3-1, KR 4512/2-1) as part of the Priority Program 1798. This paper was presented in part at the 2016 International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar, and Remote Sensing (Cosera), [1], in part at the 2017 21st International ITG Workshop on Smart Antenna [2], and in part at the Conference Wavelets and Sparsity XVII, [3].

P. Jung is with the Department of Telecommunication Systems, Technische Universität Berlin, 10587 Berlin, Germany (e-mail: peter.jung@tu-berlin.de).

F. Krahmer and D. Stöger are with the Department of Mathematics, Technische Universität München, 85748 Munich, Germany (e-mail: felix.krahmer@tum.de; dominik.stoeger@ma.tum.de).

Communicated by P. Grohs, Associate Editor for Signal Processing.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2017.2784481

In practice, however, one faces additional difficulties. Namely, the data acquisition process has to cope with uncalibrated measurement devices depending on further unknown parameters. In many such scenarios one can only sample the output of an unknown or partially known linear system. In such cases the object/signal s to recover is coupled with the unknown or partially known environment w in a multiplicative way giving rise to a *bilinear inverse problem*, i.e., solve for s and w given a bilinear combination $\mathcal{B}(w, s)$. Relevant examples are when the effective sensing matrix might be subject to uncertainties [9]–[12], or signals might have been transmitted through individual channels whose properties are not completely known [13]. Our current understanding of these *blind information retrieval* tasks is at the very beginning and usually it forces one therefore to operate at sub-optimal sensing rates, or else incur significant reconstruction errors due to model mismatch. The situation is all the more unsatisfactory, as such blind sampling problems are often much closer to practical applications than the original linear models.

A. Blind Deconvolution

The prototypical bilinear mapping, practically relevant in many applications, is the convolution

$$w * s := (\sum_{j=1}^{L} w_j s_{k-j})_{k=1}^{L}.$$

For technical reasons we will consider the circular convolution, where the index difference $k - j$ is considered modulo L. The classical convolution can be reduced to this setup by appropriate zero padding or cyclic extensions. Then the corresponding inverse problem, that is, the problem of recovering s and w from their convolution up to inherent ambiguities, is known as *blind deconvolution* [14]. The precise role of s and w depends on the underlying application. In imaging, for example, the signal vector s typically represents the image and w is an unknown blurring kernel [15]. In communication engineering, w represents the channel parameters and the task is to demodulate and decode the signal information s only having access to the *channel output* $w * s$, and the important question is how much overhead is required for coping with the unknown impulse response w of the communication channel [16].

Obviously, without further constraining s and w the convolution $(s, w) \to w * s$ has many more degrees of freedom than measurements and is hence far from being injective, exhibiting various kinds of ambiguities. The goal must then

be to eliminate these ambiguities as much as possible by imposing structural constraints on the signal and the channel paramters. It should be noted that a scaling ambiguity will always remain, as any bilinear mapping $\mathcal{B}$ satisfies $\mathcal{B}(s, w) = \mathcal{B}(\lambda s, w/\lambda)$ for any $0 \neq \lambda \in \mathbb{C}$ and can hence be injective only up to a multiplicative factor. Specific scenarios can give rise to additional ambiguities, as it has been investigated in [17]. For more detailed discussions of ambiguities in the one-dimensional case such as shifts or reflections, see [18] and [19]. In any case, additional constraints like sparsity and subspace priors, depending on the specific application, are necessary to make blind deconvolution feasible. It has been shown that sparsity in the canonical basis alone is not sufficient for these purposes [20], and for generic bases, the subspace dimensions and sparsity levels that yield injectivity have been exactly classified [21]–[23].

Even when injectivity can be established, this does not directly yield a tractable reconstruction scheme. While a number of works have studied algorithms for recovery (see, e.g., [24]–[26]), the focus has mostly been on algorithmic performance rather than on recoverability guarantees. The search for algorithms allowing for guaranteed recovery has recently shown significant progress by taking a compressed sensing viewpoint, namely aiming to choose remaining degrees of freedom to reduce the degree of ill-posedness. The first near-optimal rigourous recovery guarantees in a randomized setting have been established in [27] with high probability under the assumption that both the signal and the channel parameters lie in subspaces of small dimension, and one of them is chosen at random. The main idea was to exploit that any bilinear map $\mathcal{B}(w, s)$ can be represented as a linear map in the outer product ws^T of the two input vectors (this approach is often referred to as lifting) and hence analyzed using methods from the theory of low rank matrix recovery. More precisely, exploiting the fact that the (normalized, unitary) $L \times L$ discrete Fourier matrix F diagonalizes the circular convolution one can establish the representation

$$w * s := \sqrt{L} \cdot F^* \mathrm{diag}(Fw) Fs, \tag{I.1}$$

with $\mathrm{diag}(v)$ denoting the diagonal matrix with the entries of v on its diagonal.

Under the subspace model, where both the signal s and the vector of channel parameters are assumed to lie in a known low-dimensional subspace and hence can be represented as $w = F^*Bh$ and $s = F^*C\overline{x}/\sqrt{L}$, for given $B \in \mathbb{C}^{L\times K}$ and $C \in \mathbb{C}^{L\times N}$, this translates to

$$y := F(w * s) = \mathrm{diag}(Bh)C\overline{x} =: \mathcal{A}(hx^*), \tag{I.2}$$

where $\mathcal{A}$ is a linear map and M^* denotes the adjoint of a matrix M, that is, its conjugate transpose. This formulation yields a low rank recovery problem, as of all potential matrices giving rise to measurements y, the rank one matrix hx^* is the one of the lowest rank. Even though recovering a low rank matrix from linear measurements is known to be, in general, NP-hard [28], it has been shown that under appropriate random measurement models, one can establish recovery guarantees for tractable algorithms with high probability [29], [30]. While the results in these works require more randomness than what is available in the convolution setup due to the structure imposed by (I.2) and hence do not apply directly, Ahmed *et al.* [27] derived recovery guarantees for blind deconvolution. Their result assumes that (i) C has independent standard Gaussian entries and that (ii) $B^*B = 1$ and B is incoherent in two ways, namely that $\mu^2_{\max} := \frac{L}{K} \max_\ell \|b_\ell\|^2_{\ell_2}$ and $\mu^2_h = L \cdot \max_{1\leq \ell \leq L} |b^*_\ell h|^2$ are sufficiently small (b_ℓ are the columns of B^*). Under these assumptions, they showed that an unknown real $K \times N$–matrix hx^* can be recovered with overwhelming probability by nuclear norm minimization, that is, via the semidefinite program

$$\min \|X\|_* \quad \text{s.t. } \mathcal{A}(X) = y. \tag{I.3}$$

Here, $\|X\|_*$ denotes the nuclear norm of the matrix X, which is defined to be the sum of its singular values.

Although nuclear norm minimization is computational tractable, the lifted representation drastically increases the size of the signal to be recovered. Consequently, the resulting algorithm will be too slow for most practical applications. The theoretical analysis of nuclear norm minimization has, however, paved the way for more efficient algorithms with similar guarantees. Namely, the recent work [31] demonstrates that a gradient-based algorithm with a suitable initialization can be used without lifting and in the regime $\mu^2_h \max(K, N) \lesssim L/\log^2(L)$ which comes with considerably reduced complexity.

Finally, typical channel impulse responses h exhibit further structural properties such as sparsity, which should be used as well. However, the challenging extension of these works to sparsity models seems to be much more involved. The difficulty with such models is that the lifted representation is both sparse and of low rank, and no straightforward tractable convex relaxation is known. In particular, minimizing convex combinations of nuclear and ℓ_1-norm regularizers has been shown to yield provably suboptimal recovery performance [32]. Research regarding alternative convex surrogates as for example in [33] is only in its beginnings. For this reason, some recent approaches ignore the rank constraint, just aiming for sparsity, as investigated for the ℓ_1–approach in [34] and for the mixed ℓ_1/ℓ_2-case in [35] and [36].

On the other hand, the search for non-convex alternatives to overcome this obstacle is an active area of research. In particular, local convergence guarantees as well as global convergence guarantees for peaky signals have been derived in [37] for the sparse power factorization method, an alternating minimization approach originally introduced in [38], for the context of deconvolution. The near-optimal recovery guarantees build on some property similar to the restricted isometry property, which has been derived in [39] (for both inputs lying in random subspaces). The search for global recovery guarantees in the sparsity model without peakiness assumptions, however, remains open.

B. Simultaneous Demixing and Blind Deconvolution

The extension of the model we shall consider here is blind deconvolution and simultaneously demixing of multiple source signals. This setting is motivated by recent challenges

in future wireless multi–terminal communication scenarios for uncoordinated sporadic communication [40], [41]. We consider the prototypical case of R transmitters each having an individual information message encoded into the vector $x_i \in \mathbb{C}^{N_i}$ for $i = 1, \dots, R$ using, for example, classical modulation alphabets and error–correcting codes. In fact, such data could be independent user data payloads or even correlated sensor readings on a common source. For reasons of simplicity, we focus on the case of independent data sources. Each transmitter generates its transmit signal $s_i = F^* C_i \overline{x}_i / \sqrt{L} \in \mathbb{C}^L$ by multiplying (linearly encoding) its complex–valued (conjugated) message vector $\overline{x}_i$ by an $L \times N_i$ matrix $F^* C_i / \sqrt{L}$ which is then transmitted into the shared channel. Note that, from the perspective of communication engineering, this procedure has been simplified to facilitate the analysis. In a more advanced setting one could consider a directly randomized mapping from bits to sequences in $\mathbb{C}^L$. Now consider a single receiver, for example a base station. Each transmitter i has its individual impulse response w_i describing the channel propagation conditions to this base station. For simplicity we consider a low–mobility scenario where, for appropriate block length L, the channel is time–invariant and can be modeled by a convolution of the transmit signal with a channel impulse response w_i. Furthermore, with cyclic extensions and/or zero-padding at the transmitter such a signal propagation can then be modeled as a circular convolution. To incorporate further structure for the channel impulse response we write it as $w_i = F^* B_i h_i$ where $B_i \in \mathbb{C}^{L \times K_i}$. A reasonable assumption for our application is that the unknown coefficients h_i are located on the first samples since the path delays in the channel are usually much shorter than the frame length L. In this case $F^* B_i$ is a truncated identity, i.e., $B_i^* B_i = \mathrm{Id}$.

In practice, since the desired deployment scenario is uncoordinated and sporadic, only a small fraction of size r of R devices are online and transmitting data. We assume for this work that the receiver is able to detect the activity pattern correctly (which can be achieved through a separate control channel, see for example [42] for a certain approach). One can even detect activity simultaneously with data. However, algorithms for blind deconvolution and demixing are usually quite complex from practical and computational aspects and it is desired to reduce the problem size as much as possible already from the beginning. This means, restricted and resorting to the active set, the receiver observes the noisy superposition

$$y = \sum_{i=1}^{r} F(w_i * s_i) + e = \sum_{i=1}^{r} \mathrm{diag}(B_i h_i) C_i \overline{x}_i + e = \sum_{i=1}^{r} \mathcal{A}_i(h_i x_i^*) + e \tag{I.4}$$

of r signal contributions where the vector $e \in \mathbb{C}^L$ denotes additive noise.

The conventional approach is (i) to design the matrices C_i in such a way that resources are used exclusively by $\mathcal{O}(R)$ devices which requires considerable processing, resource planning and allocation algorithms and (ii) estimate the channel from pilot signals during a calibration phase prior to data transmission. However, in an increasing number of new applications the typical data traffic consists only of short messages (status updates or sensor data) yielding a sporadic traffic type and then the overall communication in a network is then considerable dominated by control data.

In [43] it has therefore been proposed to consider the scenario of *simultaneous blind deconvolution and demixing* of multiple signals from its superposition y, which we will also study in this paper. Demixing by convex programming methods has been intensively investigated in the fields of "sine and spikes" (and pairs of bases) decompositions, see [44] and [45], and in the field of sparse and low–rank decomposition, see, e.g., the work [46]. More generally, as for example outlined in [47] and [48], a convex approach consists of minimizing the sum of the individual regularizers over all signal formations which are conform with the model and consistent with the observations. To this end, assuming a priori that $\|e\|_{\ell_2} \leq \tau$, we consider the convex optimization problem

$$\min \sum_{i=1}^{r} \|X_i\|_* \quad \text{s.t. } \|\sum_{i=1}^{r} \mathcal{A}_i(X_i) - y\|_{\ell_2} \leq \tau. \tag{I.5}$$

According to [47], reliable convex demixing is possible whenever (i) the signal contributions are incoherent to each other and (ii) the number of observations is sufficiently above the sum of effective dimensions of the descent cones of the individual regularizers at the unknown ground truth. Since the rank-one matrix $X_i = h_i x_i^*$ has effective dimension $K_i + N_i$ this amounts to $\mathcal{O}(r(K + N))$ observations, where $K = \max_i(K_i)$ and $N = \max_i(N_i)$. First results and guarantees, based on the incoherence between the mappings $\mathcal{A}_i$ which explicitly occur in blind deconvolution (I.4) with random C_i's are worked out in [43]. The result in this paper states that if L is in the order of $r^2 \max(K, N)$ (up to logarithmic factors) the minimizer $(\hat{X}_1, \dots, \hat{X}_r)$ of the program (I.5) satisfies with high probability that

$$\sum_{i=1}^{r} \|\hat{X}_i - X_i^0\|_F^2 \lesssim r^2 \cdot \max\{K; N\} \tau^2 \tag{I.6}$$

Hence, for $\tau = 0$ the ground truth $(\hat{X}_1^0, \dots, \hat{X}_r^0)$ is recovered exactly. However, the embedding dimension does not quite match the effective dimension, which would suggest a linear dependence on r. Ling and Strohmer suggested that this mismatch is a proof artifact, observing numerically that linear dependence on r. In this paper, we will analytically justify these observations. In the special case of partial (low-frequency) Fourier matrices B_i mentioned above, our main result, Theorem 6, reads as follows.

Theorem 1: Let $\omega \geq 1$ and set $\mu_h^2 = L \max_{i,\ell} |b_{i,\ell}^* h_i|^2$. Assume $\|e\|_{\ell_2} \leq \tau$ and that

$$L \geq C_\omega r \left(K \log K + N \mu_h^2 \right) \log^3 L, \tag{I.7}$$

where C_ω is a universal constant only depending on ω. Then with probability at least $1 - \mathcal{O}\left(L^{-\omega}\right)$ the minimizer $\hat{X}$ of the recovery program (I.5) *satisfies*

$$\sum_{i=1}^{r} \|\hat{X}_i - X_i^0\|_F^2 \lesssim r \cdot \max\left\{1; \frac{rKN}{L}\right\} \log(L)\ \tau^2. \tag{I.8}$$

The coherence parameter μ_h^2 ranges between 1 and K, [43]. For a detailed discussion we refer to [43, Sec II.D]. At this point we only want to comment that in applications in Wireless Communications this parameter is typically small.

Shortly before the completion of this manuscript Ling and Strohmer presented recovery guarantees for (considerably more efficient) nonconvex gradient (Wirtinger) based methods [49], again with quadratic scaling in r. Again they conjecture linear dependence, as observed in their numerical experiments. We also include some numerical experiments in Section III that illustrate the linear dependence. We expect that our paper at hand will pave the way to an optimized parameter dependence also for more efficient algorithms.

II. General Framework and Main Result

A. Notation

Before we describe the mathematical model we introduce some basic notation. For complex numbers $z \in \mathbb{C}$ we denote its conjugate by $\bar{z}$ and write Re z and Im z for the real and imaginary part. Similarly, for a vector $w = (w[1], \dots, w[n]) \in \mathbb{C}^n$ we use the notation Re w = (Re $w[1], \dots,$ Re $w[n]$) and Im w = (Im $w[1], \dots,$ Im $w[n]$). For a matrix $A \in \mathbb{C}^{d_1 \times d_2}$ we will denote its adjoint by A^* and (for $d_1 = d_2$) its trace by Tr (A). For matrices $A, B \in \mathbb{C}^{d_1 \times d_2}$ we will define the inner product by $\langle A, B \rangle_F = \text{Tr}\,(AB^*)$. The Frobenius norm of A is $\|A\|_F^2 = \langle A, A \rangle_F$ and $\|A\|_{2\to 2}$ denotes its operator norm. If $\mathcal{B}$ is a linear operator mapping matrices to vectors or matrices, we will denote its operator norm by $\|\cdot\|_{F\to 2}$ or $\|\cdot\|_{F\to F}$, respectively. The nuclear norm of the matrix A, which is defined as the sum of its singular values, will be denoted by $\|A\|_*$. Note that the notation for $\|\cdot\|_*$, $\|\cdot\|_F$ and $\langle \cdot, \cdot \rangle_F$ will be used later in a more generalized setting, as will be pointed out in the next section. The matrix Id_d will denote the identity matrix in $\mathbb{C}^{d\times d}$. If no confusion can arise, we will suppress d and write Id instead of Id_d. For a vector $v \in \mathbb{C}^d$ diag (v) denotes the matrix whose diagonal entries are given by v. Furthermore, $\|v\|_{\ell_2}$ denotes the ℓ_2-norm of this vector, i.e. $\|v\|_{\ell_2}^2 = \langle v, v \rangle = \text{Tr}\,(vv^*)$.

By $\mathbb{P}(E)$ we will denote the probability of an event E. For any $N \in \mathbb{N}$ we will denote the set $\{1, \dots, N\}$ by $[N]$. For a set S we will denote its cardinality by $|S|$. The notation log $(\cdot)$ will refer to the logarithm of base 2. Furthermore, during the whole manuscript C will denote positive numerical constants, which are independent of all other variables which appear in the text and whose value may change from line to line. Similarly, C_ω will denote universal numerical constants, which only depend on ω. We will write $a \lesssim b$, if $a \le Cb$ and $a \lesssim_\omega b$, if $a \le C_\omega b$. We will write $a \sim b$, if we have $a \lesssim b$ as well as $b \lesssim a$.

B. The General Model

In this paper we will work with a more general model, as also studied in [43], which includes the demixing-deconvolution scenario given above as special case. Assume that the vector $y \in \mathbb{C}^L$ of L noisy measurements corresponding to inputs $\{h_i\}_{i=1}^r$, $h_i \in \mathbb{C}^{K_i}$ and $\{x_i\}_{i=1}^r$, $x_i \in \mathbb{C}^{N_i}$, is given by

$$y = \sum_{i=1}^{r} \text{diag}\,(B_i h_i)\, C_i \overline{x}_i + e. \tag{II.1}$$

where e is additive noise, the matrices $B_i \in \mathbb{C}^{L\times K_i}$ satisfy $B_i^* B_i = \text{Id}_{K_i}$ for all $i \in [r]$, and all the entries of the random matrices $C_i \in \mathbb{C}^{L\times N_i}$ are independent and follow a standard circular-symmetric complex normal distribution $\mathcal{CN}(0, 1)$ (see Appendix B for more details). The vectors h_i are assumed to be normalized, $\|h_i\|_{\ell_2} = 1$, whereas the norms of x_i are arbitrary. (This is not restrictive as there is an inherent scaling ambiguity.) Furthermore, we set

$$K = \max_{i\in[r]} K_i \quad \text{and} \quad N = \max_{i\in[r]} N_i.$$

Let us denote by $b_{i,\ell}$ the ℓth column of B_i^* and by $c_{i,\ell}$ the ℓth column of C_i. Then, the ℓth entry of y is given by

$$y[\ell] = \sum_{i=1}^{r} b_{i,\ell}^* h_i x_i^* c_{i,\ell} + e[\ell].$$

We observe that the overall vector y only depends on the outer products $h_i x_i^*$. Thus, we may proceed by considering a lifted representation (see, e.g., [50]). Defining for each $i \in [r]$ the operator $\mathcal{A}_i : \mathbb{C}^{K_i\times N_i} \longrightarrow \mathbb{C}^L$ via

$$\mathcal{A}_i(Z) := \left(b_{i,\ell}^* Z c_{i,\ell}\right)_{\ell=1}^L$$

we obtain that

$$y = \sum_{i=1}^{r} \mathcal{A}_i\left(h_i x_i^*\right) + e.$$

In the following we will use the decomposition $x_i = \sigma_i m_i$ where $\sigma_i \ge 0$ and some $m_i \in \mathbb{C}^{N_i}$ such that $\|m_i\|_{\ell_2} = 1$. (If $x_i = 0$ we set $\sigma_i = 0$ and choose m_i arbitrarily.) Thus, the signal to be recovered may be written as

$$\begin{aligned} X^0 :&= \left(h_1 x_1^*, \dots, h_r x_r^*\right) = \left(\sigma_1 h_1 m_1^*, \dots, \sigma_r h_r m_r^*\right) \\ &= (X_1, \dots, X_r). \end{aligned}$$

Define

$$\mathcal{M} := \left\{(Z_1, \dots, Z_r) :\ Z_i \in \mathbb{C}^{K_i\times N_i} \text{ for all } i \in [r]\right\}$$

and note that $\mathcal{M}$ is naturally equipped with the algebraic structure of a vector space, as it may be regarded as the product space of the vector spaces $\mathbb{C}^{K_i\times N_i}$. The linear operator $\mathcal{A} : \mathcal{M} \to \mathbb{C}^L$ is defined by

$$\mathcal{A}(Z) := \sum_{i=1}^{r} \mathcal{A}_i(Z_i)$$

for $Z = (Z_1, \dots, Z_r) \in \mathcal{M}$. The linear space $\mathcal{M}$ will be endowed with a norm and an inner product defined by

$$\langle W, Z \rangle_F := \sum_{i=1}^{r} \langle W_i, Z_i \rangle_F$$

and

$$\|W\|_F^2 = \langle W, W \rangle_F = \sum_{i=1}^{r} \|W_i\|_F^2.$$

for all $W, Z \in \mathcal{M}$. The operator norms $\|\cdot\|_{F\to 2}$ and $\|\cdot\|_{F\to F}$ of linear maps on $\mathcal{M}$ are defined analogously to the matrix case. For the adjoint $\mathcal{A}^*$ of $\mathcal{A}$ with respect to the inner product on $\mathcal{M}$ it follows $\mathcal{A}^*(y) = \left(\mathcal{A}_1^*(y), \dots, \mathcal{A}_r^*(y)\right)$ for all $y \in \mathbb{C}^L$. Note that the adjoint operations $\mathcal{A}_i^*(y)$ itself are given by

$$\mathcal{A}_i^*(y) = \sum_{\ell=1}^{L} y[\ell] b_{i,\ell} c_{i,\ell}^* \quad \text{for all } y \in \mathbb{C}^L. \tag{II.2}$$

We will also use the norm defined by $\|W\|_* = \sum_{i=1}^r \|W_i\|_*$. For reasons which will become clear in Section VI-A we set

$$\operatorname{sgn}(X_i^0) := \begin{cases} h_i m_i^* & \sigma_i > 0 \\ 0 & \text{else} \end{cases}$$

for $i \in [r]$ (recall that $\sigma_i \geq 0$). This allows us to define

$$\operatorname{sgn}(X^0) := \left(\operatorname{sgn}(X_1^0), \dots, \operatorname{sgn}(X_r^0)\right).$$

C. Partition of Measurements and Incoherence Assumptions

As those of [27], [43], our results are based on two notions of coherence. The first is captured by the coherence parameter

$$\mu_i^2 = \max_{\ell\in[L]} \frac{L}{K_i} \|b_{i,\ell}\|_{\ell_2}^2 \quad \text{for } i \in [r]. \tag{II.3}$$

Note that $B_i^* B_i = \mathrm{Id} \in \mathbb{C}^{K_i \times K_i}$ for all $i \in [r]$ implies that $1 \leq \mu_i^2 \leq \frac{L}{K_i}$. In the (important) case that all matrices B_i are partial (low-frequency) DFT matrices, which refers to the special situation described in the introduction, we have minimal coherence $\mu_i^2 = 1$. In order to simplify notation we introduce the quantities

$$K_{i,\mu} := K_i \mu_i^2, \quad K_\mu := \max_{i\in[r]} K_{i,\mu}. \tag{II.4}$$

We observe that $K_i \leq K_{i,\mu} \leq L$. Again, in the special case that the matrices B_i are partial (low-frequency) DFT matrices we obtain that $K_{i,\mu} = K_i$.

For the proof of our results we will use the Golfing Scheme [30], see Section VI-C.1. This requires a partition $\left\{\Gamma_p\right\}_{p=1}^P$ of the set of the measurements $[L]$ with associated measurement operators $\mathcal{A}^p$. The second coherence parameter will also depend on this partition. In order to guarantee that the Golfing Scheme is successful with high probability we will need that $T_{i,p} := \frac{L}{Q} \sum_{\ell\in\Gamma_p} b_{i,\ell} b_{i,\ell}^* \approx \mathrm{Id}_{K_i}$, as it will become clear in Remark 29. Thus, we have to assure that the partition $\left\{\Gamma_p\right\}_{p=1}^P$ is chosen such that for $Q := \frac{L}{P}$ and $\nu > 0$ small enough one has

$$\max_{i\in[r],\ p\in[P]} \left\| \mathrm{Id}_{K_i} - T_{i,p} \right\|_{2\to 2} \leq \nu. \tag{II.5}$$

Furthermore, we require that $|\Gamma_p|$ is large enough for all $p \in [P]$, i.e., each operator $\mathcal{A}^p$ contains enough measurements, and also the partition consists of the right number of sets, that is, P is bounded from above and below. More precisely, we require that the partition is ω-admissible in the sense of the following definition.

Definition 2: Let $\omega \geq 1$ and let $\left\{\Gamma_p\right\}_{p=1}^P$ be a partition of $[L]$. The set $\left\{\Gamma_p\right\}_{p=1}^P$ is called ω-admissible if the following three conditions are satisfied:

1) $\frac{1}{2}Q \leq |\Gamma_p| \leq \frac{3}{2}Q$ *for all* $p \in [P]$, *where* $Q = \frac{L}{P}$.
2) *(II.5) is fulfilled with* $\nu = \frac{1}{32}$.
3) *It holds that* $\log\left(8\tilde{\gamma}\sqrt{r}\right) \geq P \geq \frac{1}{2}\log\left(8\tilde{\gamma}\sqrt{r}\right)$, *where*

$$\tilde{\gamma} = 2\sqrt{\omega \max\left\{1; \frac{r K_\mu N}{L}\right\} \log\left(L + rKN\right)}.$$

Here the parameter ω is the same that appears in Theorem 1 and in Theorem 6.

This definition gives rise to the question whether such a partition exists in general and how one can construct them. This has already been discussed in [43, Sec. 2.3] for several important special cases of matrices $B_i \in \mathbb{C}^{K_i \times N_i}$. In particular, it is proven that in the special case that the B_i's are partial (low-frequency) Fourier matrices of the same size and if $L = PQ$ one may find a partition such that $\nu = 0$. In [27], the authors discussed the construction of such a partition for $r = 1$ and for a general matrix $B \in \mathbb{C}^{K\times N}$ which satisfies $B^*B = \mathrm{Id}_K$. However, such a partition can be constructed for all matrices $B_i \in \mathbb{C}^{K_i \times N_i}$ simultanously via the following lemma.

Lemma 3: Let $P \in [L]$ and $\nu \in (0,1)$ be fixed. Set $Q = \frac{L}{P}$. There is a universal constant $C > 0$ such that if

$$Q \geq C \frac{K_\mu}{\nu^2} \log\left(\max\{r; P; K\}\right) \tag{II.6}$$

then there is a partition $\left\{\Gamma_p\right\}_{p=1}^P$ of $[L]$ such that (II.5) is satisfied and $\frac{1}{2}Q \leq |\Gamma_p| \leq \frac{3}{2}Q$ holds for all $p \in [P]$.

A proof of this result is included in Appendix A. As $P = \frac{L}{Q}$, this lemma implies the existence of an ω-admissible partitions provided that

$$L \gtrsim \sqrt{r}\log\left(8\tilde{\gamma}\sqrt{r}\right) \frac{K_\mu}{\nu^2} \log\left(\max\{r; P; K\}\right),$$

with $\tilde{\gamma}$ as in Definition 2, which is a somewhat milder assumption than what is required in our main theorem.

The second incoherence parameter will depend on the choice of such an ω-admissible partition, measuring how aligned the input h_i is with the basis vectors $b_{i,\ell}$ distorted by a family of linear maps corresponding to the different sets in the partition.

More precisely, for a fixed ω-admissible partition $\left\{\Gamma_p\right\}_{p=1}^P$ we define

$$\mu_h^2 := L \max\left\{ \max_{\ell\in[L], i\in[r]} |b_{i,\ell}^* h_i|^2, \max_{p\in[P], \ell\in[L], i\in[r]} |b_{i,\ell}^* S_{i,p} h_i|^2 \right\}, \tag{II.7}$$

where we have set $S_{i,p} = T_{i,p}^{-1}$. The proof in Section VI will yield the strongest result when μ_h^2 is small. Thus, we will choose for our proof a partition, which minimizes the quantity defined in (II.7). This motivates the introduction of the following quantity.

$$\mu_{h,\omega}^2 = L \min_{\{\Gamma_p\}_{p=1}^P \omega\text{-admissible}} \max\left\{ \max_{\ell\in[L], i\in[r]} |b_{i,\ell}^* h_i|^2, \max_{p\in[P], \ell\in[L], i\in[r]} |b_{i,\ell}^* S_{i,p} h_i|^2 \right\}. \tag{II.8}$$

Lemma 4: Let $\{\Gamma_p\}_{p=1}^P$ *be a* ω*-admissible partition of* $[L]$. *Then* $1 \leq \mu_h^2 \leq \left(\frac{32}{31}\right)^2 K_\mu$.

Proof: The lower bound follows immediately from the observation

$$\sum_{\ell=1}^{L} \|b_{i,\ell}^* h_i\|_{\ell_2}^2 = \sum_{\ell=1}^{L} h_i^* b_{i,\ell} b_{i,\ell}^* h_i = \|h_i\|_{\ell_2}^2 = 1.$$

For the upper bound it is enough to observe that $L|b_{i,\ell}^* h_i|^2 \leq L\|b_{i,\ell}\|_{\ell_2}^2 \|h_i\|_{\ell_2}^2 \leq K_\mu$ and similarly $L|b_{i,\ell}^* S_{i,p} h_i|^2 \leq L\|S_{i,p}\|_{2\to 2}^2 \|b_{i,\ell}\|_{\ell_2}^2 \|h_i\|_{\ell_2}^2$. The result follows from the observation $\|S_{i,p}\|_{2\to 2} \leq \frac{32}{31}$, which is due to $\|\mathrm{Id} - T_{i,p}\|_{2\to 2} \leq \frac{1}{32}$. □

Remark 5: As already pointed out in [43, Remark 2.1] the appearance of the second term in the definition of μ_h *is due to the modified Golfing Scheme (cf. Remark 29). Note, however, that our definition of* μ_h^2 *is slightly different to the definition of* μ_h^2 *in [43]. In our definition, the second term the maximum is over all* $\ell \in [L]$, *whereas in [43] the maximum is only over all* $\ell \in \Gamma_p$. *The reason is that of a simpler presentation and a less technical argument; it is possible to obtain our result with* μ_h^2 *as defined in [43] by a slightly more involved argument: One needs to replace the norm* $\|\cdot\|_B$, *which will be introduced in Section 22, by norms which depend on the individual partitions* Γ_p.

One may ask whether the second term in the definition of μ_h^2 can be removed. By a closer look at the proof of Lemma 3 one infers that for fixed P, which satisfies the third condition in Definition 2, a constant fraction of all partitions are μ-admissible. Thus, one might conjecture that there is at least one partitition such that the quantity $\max_{p\in[P],\ell\in[L],i\in[r]} |b_{i,\ell}^* S_{i,p} h_i|^2$ is small such that it can be neglected. We leave this problem for future work.

D. Main Result

Our main result establishes a recovery guarantee for the general measurement model (II.1). Reconstruction proceeds via nuclear norm minimization, the semidefinite program formulated in (I.5).

Theorem 6: Let $\omega \geq 1$ *and let* $y \in \mathbb{C}^L$ *be given by (II.1) with* $\|e\|_{\ell_2} \leq \tau$. *Assume that*

$$L \geq C_\omega r \left(\max_{i\in[r]} \left(K_i \mu_i^2 \log\left(K_i \mu_i^2\right)\right) + N\mu_{h,\omega}^2\right) \log^3 L, \tag{II.9}$$

where C_ω *is a universal constant only depending on* ω. *Then, with probability at least* $1 - \mathcal{O}\left(L^{-\omega}\right)$ *the minimizer* $\hat{X}$ *of the recovery program* (I.5) *satisfies*

$$\|\hat{X} - X^0\|_F \lesssim \tau \sqrt{r \max\left\{1; \max_{i\in[r]} \frac{r K_i \mu_i^2 N}{L}\right\} \log L}. \tag{II.10}$$

In the important special case of noiseless measurements, i.e., $\tau = 0$, Theorem 6 yields exact recovery with high probability, if L satisfies condition (II.9), i.e., X^0 is the unique minimizer of the semidefinite program (I.5). As already mentioned in the introduction our result significantly improves upon the result of [43] and exhibits optimal scaling in the degrees of freedom up to logarithmic factors. In the noisy case, the estimation error (II.10) is improved at least by a factor of $\sqrt{r}$ (cf. [43, Th. 3.3]). The authors believe that this is still not optimal as it might be possible to remove the dependence on K, N, and r in the estimation error. However, it seems to be likely that it is not possible to resolve this issue using our current proof technique, which relies on the construction of an approximate dual certificate. Also for the interesting problem of extending our result to matrices that are only approximately low-rank, similar to the study of compressible signals in compressed sensing, we expect that near-optimal guarantees need different techniques.

III. Outlook

Although the convex formulation in (I.5) is important for theoretical investigations it is also obvious that for many real-word applications nuclear minimization is not feasible due to its computional complexity as lifting considerably increases the number of optimization variables. For the case $r = 1$ a nonconvex approach has been proposed by [31] which has been demonstrated not only to be considerably more efficient but also to achieve a better empirical performance. Shortly before the completion of our work this line of research has been extended to $r \geq 1$ with explicit guarantees [49], but again for a number of measurements depending quadratically on r. As in [43], the dependence observed in numerical experiments is linear. We expect that the mathematical analysis conducted in this paper will also be important for establishing near-optimal performance guarantees for more efficient algorithms. For this reason we include such a nonconvex approach similar to the one analysed in [49] in our numerical experiments, comparing it to nuclear norm minimization as analyzed in this paper.

More precisely, we consider a gradient-based (Wirtinger flow) recovery algorithm minimizing the residual

$$F(h, x) := \|\mathcal{A}(h_1 x_1^*, \ldots, h_r x_r^*) - y\|_{\ell_2}^2 \tag{III.1}$$

where $h := (h_1, \ldots, h_r)$ and $x := (x_1, \ldots, x_r)$. Observe that in the noiseless case one has $F(h, x) = 0$ for the ground truth. Note that, while minimizing F has been shown empirically in [49] to have good recovery properties, where guarantees only apply to a regularized variant. As F is highly non-convex in (h, x) and may possess many local minima, it is essential to find a good initial guess to start the minimization process (cf. [31], [49]). Eq. (VI.5) motivates the initialization given in the following algorithm.

To minimize F a gradient descent approach is used. Here the gradient of a function $f : \mathbb{C}^n \to \mathbb{C}$ at $z_0 \in \mathbb{C}^n$ is given by $\nabla_z f(z_0) = \left(\frac{\partial f}{\partial z}(z_0)\right)^* \in \mathbb{C}^n$ where for $z = u + iv \in \mathbb{C}$ the Wirtinger derivatives are $\frac{\partial}{\partial z} = \frac{1}{2}\left(\frac{\partial}{\partial u} - i\frac{\partial}{\partial v}\right)$ and $\frac{\partial}{\partial \overline{z}} = \frac{1}{2}\left(\frac{\partial}{\partial u} + i\frac{\partial}{\partial v}\right)$. Since for real-valued complex functions $f : \mathbb{C}^n \to \mathbb{R}$ one has $\frac{\partial f}{\partial \overline{z}} = \overline{\frac{\partial f}{\partial z}}$, we do not need to consider $\frac{\partial f}{\partial \overline{z}}$ here. Consequently, we obtain

$$\nabla_{h_i} F(h, x) = \left(\mathrm{diag}\left(C_i \overline{x_i}\right) B\right)^* \left(\mathcal{A}(hx^*) - y\right);$$
$$\nabla_{x_i} F(h, x) = \left(\mathrm{diag}(B_i h_i) C_i\right)^T \overline{\left(\mathcal{A}(hx^*) - y\right)}$$

To estimate a suitable stepsize η

Algorithm 1 Initialization

Input: Observation y.
$(Z_1, \dots, Z_r) \leftarrow \mathcal{A}^* y$.
for $k = 1, \dots, r$ **do**
  $d_k \leftarrow$ largest singular value of Z_k.
  Let $v_k^{(0)}$ and $u_k^{(0)}$ be the corresponding left and right singular vectors, respectively.
  $v_k^{(0)} \leftarrow \sqrt{d_k} v_k^{(0)}$ and $u_k^{(0)} \leftarrow \sqrt{d_k} u_k^{(0)}$
end for
Output: Initial guesses $v^{(0)}, u^{(0)}$.

Algorithm 2 Wirtingers Gradient Descent With Backtracking

Input: Initial values $v^{(0)}, u^{(0)}$.
for $i = 1, \dots$ **do**
  $\eta \leftarrow$ LINE-SEARCH $\left(v^{(i-1)}, u^{(i-1)}\right)$
  $v^{(i)} \leftarrow v^{(i-1)} - \eta \nabla_h F\left(v^{(i-1)}, u^{(i-1)}\right)$
  $u^{(i)} \leftarrow u^{(i-1)} - \eta \nabla_x F\left(v^{(i-1)}, u^{(i-1)}\right)$
  if $\|\nabla F\left(v^{(i)}, u^{(i)}\right)\|_{\ell_2} < \varepsilon$ **then**
    return $v^{(i)}, u^{(i)}$
  end if
end for
Output: Approximate solutions $v^{(i)}, u^{(i)}$.

Numerical Results: We have investigated both nuclear norm minimization (I.5) and Algorithms 1 and 2 in the noiseless case for different values of r and L with equal channel dimensions $K = K_1 = \ldots = K_r = 8$ and signal dimensions $N = N_1 = \ldots = N_r = 8$. The success rates per device are estimated numerically and plotted as a function of $\rho = L/(r(K+N))$. The convex program (I.5) is solved using the Matlab CVX toolbox. For each experiment the matrices $C_i \in \mathbb{C}^{L \times N}$, the signal vectors $x_i^0 \in \mathbb{C}^N$, and the channel coefficients $h_i^0 \in \mathbb{C}^K$ are generated with i.i.d. complex normal distributed entries. Recovery is considered successful for a device if the corresponding signal pair (h_i, x_i) for $i \in [r]$ fullfils $\|h_i x_i^* - h_i^0 x_i^{0*}\|_F / \|h_i^0 x_i^{0*}\|_F \leq 1\%$. Furthermore, the stopping criterion for the Wirtinger approach is chosen to be $\epsilon = 10^{-4}$ and the maximal number of iterations is limited to 1000.

Our experiments confirm the findings of [43] and [49] that for both the convex and the non-convex approach the scaling is linear. The results in Figure 1 show that – almost independently of r – the phase transition for (I.5) occurs at $\rho \approx 2.75$ while the Wirtinger flow approach performs considerably better with a phase transition (for larger r) at $\rho \approx 1.17$.

IV. Preliminaries for the Proof

A. Concentration Inequalities

In our proof we will have to estimate the spectral norm of a random matrix several times. Amongst others one tool we will apply is a generalized version of the matrix Bernstein inequality, which may be seen as a corollorary from [51, Th. 4]. It is based on so-called Orlicz norms $\|\cdot\|_{\psi_\alpha}$, which may be regarded as a measure for the tail decay of random variables.

Definition 7: Let X be a complex-valued random variable. For $\alpha \geq 1$ we define the Orlicz norm $\|\cdot\|_{\psi_\alpha}$ by

$$\|X\|_{\psi_\alpha} = \inf\left\{t > 0 : \mathbb{E}\left[\exp\left(\frac{|X|^\alpha}{t^\alpha}\right)\right] \leq 2\right\}.$$

It is straightforward to check that $\|\cdot\|_{\psi_\alpha}$ is a norm (on the vector space of all complex-valued random variables X such that $\|X\|_{\psi_\alpha} < +\infty$). Furthermore, as shown in [52], any two random variables X, Y satisfy the Hoelder inequality

$$\|XY\|_{\psi_1} \leq \|X\|_{\psi_2} \|Y\|_{\psi_2}. \tag{IV.1}$$

If $\|X\|_{\psi_1} < \infty$ we will call a random variable sub-exponential. For sub-exponential random variables we state the Bernstein inequality in the version of [53, Proposition 5.16].

Theorem 8: Let $X_1, \dots, X_n$ be independent, mean zero sub-exponential random variables, i.e., $\|X_i\|_{\psi_1} < \infty$ for all $i \in [r]$. Then with probability at least $1 - 2\exp(-t)$

$$\left|\sum_{i=1}^n X_i\right| \lesssim \max\left\{\sqrt{t \sum_{i=1}^n \|X_i\|_{\psi_1}^2}; t\left(\max_{i\in[n]} \|X_i\|_{\psi_1}\right)\right\}.$$

There are powerful generalizations of the Bernstein inequality for the matrix-valued case. Those generalizations were discovered first in [54] and were refined in [55]. We will state a variant of this theorem for unbounded random matrices, which is a reformulation of a version of Koltchinskii [51, Th. 4].

Theorem 9 (Matrix Bernstein Inequality): Let $\alpha \in [1, +\infty)$ and let $X_1, X_2, \dots, X_n \in \mathbb{C}^{d_1 \times d_2}$ be independent random matrices that satisfy $\mathbb{E}[X_i] = 0$ for all $i \in [n]$. Set $R = \max_{i\in[n]} \left\| \|X_i\|_{2\to 2} \right\|_{\psi_\alpha}$ and

$$\sigma^2 = \max\left\{\left\|\sum_{i=1}^n \mathbb{E}\left[X_i X_i^*\right]\right\|_{2\to 2}; \left\|\sum_{i=1}^n \mathbb{E}\left[X_i^* X_i\right]\right\|_{2\to 2}\right\}. \tag{IV.2}$$

Set $Z = \sum_{i=1}^n X_i$. Then with probability at least $1 - \exp(-t)$

$$\|Z\|_{2\to 2} \lesssim \max\left\{\sigma\sqrt{t + \log(d_1+d_2)}; R\left(\log\left(1 + \frac{nR^2}{\sigma^2}\right)\right)^{\frac{1}{\alpha}} (t + \log(d_1+d_2))\right\}.$$

Indeed, when $d_1 = d_2$ and the matrices $X_1, X_2, \dots, X_n$ are self-adjoint, Theorem 9 can be deduced from [51, Th. 4] (by choosing $\psi_\alpha(u) = \exp(u^\alpha) - 1$ and, for example, $\delta = 1$). In order to pass from self-adjoint matrices to general matrices $X_i \in \mathbb{C}^{d_1 \times d_2}$ one may use self-adjoint dilations and argue as in [56, Sec. 4.6.5].

The matrix Bernstein inequality is a powerful tool, which works in many different situations. However, for some more specific examples of random matrices there are other tools, which yield better estimates and which are easier to apply. The following theorem is useful, when the matrix Z is the sum of matrices of the type $\gamma_i X_i$ where X_i is a fixed matrix and γ_i is a circular-symmetric complex normal distributed random variable. It is an immediate corollary of [56, Th. 4.1.1], where matrices of this type are called Matrix Gaussian Series. For completeness, we include a proof in the Appendix.

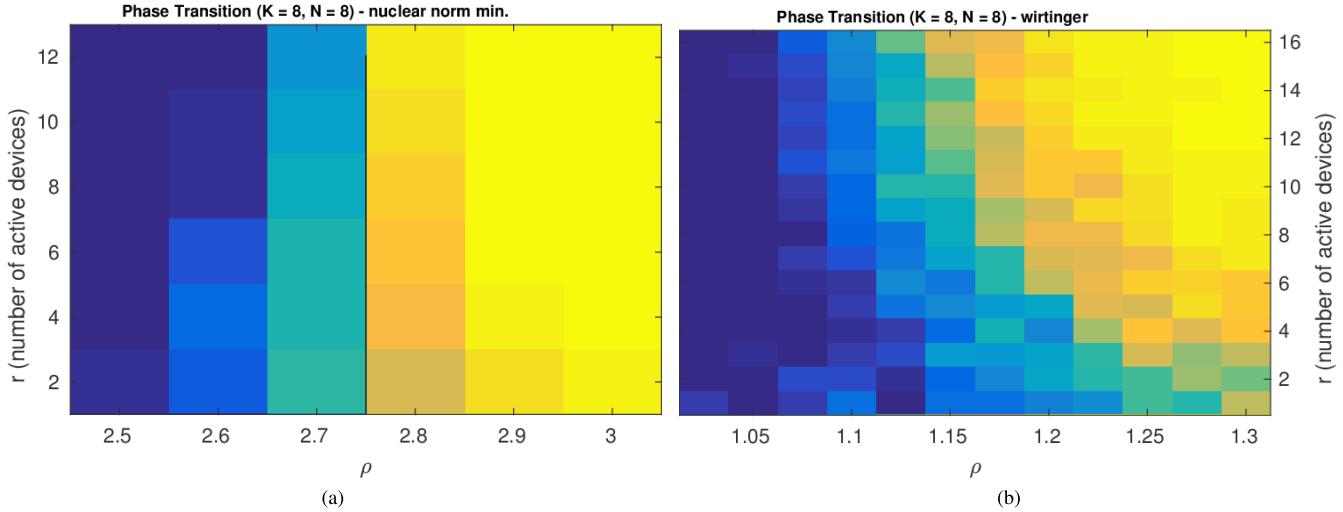

Fig. 1. Phase transition of the success rates per device for (a) the convex approach (I.5) and (b) the Wirtinger approach for $K=N=8$ where $\rho = L/(r(K+N))$.

Corollary 10 (Matrix Gaussian Series): Let $X_1, \dots, X_n \in \mathbb{C}^{d_1 \times d_2}$ *be (fixed) matrices, and let* $\gamma_1, \dots, \gamma_n$ *be independent, identically distributed random variables, where* γ_i *has circular symmetric normal distribution* $\mathcal{CN}(0,1)$. *Set* $Z = \sum_{i=1}^n \gamma_i X_i$ *and*

$$\begin{aligned}\sigma^2 &= \max\left\{\left\|\mathbb{E}\left[ZZ^*\right]\right\|_{2\to 2}, \left\|\mathbb{E}\left[Z^*Z\right]\right\|_{2\to 2}\right\} \\ &= \max\left\{\left\|\sum_{i=1}^n X_i X_i^*\right\|_{2\to 2}; \left\|\sum_{i=1}^n X_i^* X_i\right\|_{2\to 2}\right\}.\end{aligned}$$

Then, for all $t > 0$, *with probability at least* $1 - 2\exp(-t)$

$$\|Z\|_{2\to 2} \le \sigma\sqrt{2(t + \log(d_1+d_2))}.$$

B. Suprema of Chaos Processes

In addition to sums of random matrices, random variables of the form $\sup_{A\in\mathcal{X}} \|A\xi\|$, where ξ is a random vector and $\mathcal{X}$ is a class of matrices, will play an important role in this paper. To state a tail bound for such random variables, we need the γ_2-functional, a geometric quantity introduced by Talagrand (see [57]).

Definition 11: Let $(X, |||\cdot|||)$ *be a Banach space and suppose that* $S \subset X$. *We say that a sequence* $(S_n)_{n\ge 0}$ *of subsets of* S *is admissible, if* $|S_0| = 1$ *and* $|S_n| \le 2^{2^n}$ *for all* $n \ge 1$. *Then we set*

$$\gamma_2(S, |||\cdot|||) = \inf_{(S_n)_{n\ge 0}} \sup_{s\in S} \sum_{n=0}^{\infty} 2^{n/2} \inf_{s\in S_n} |||s - s_n|||,$$

where the infimum is taken over all admissible sequences $(S_n)_{n\ge 0}$.

The γ_2-functional fulfills the following inequality.

Lemma 12 [39, Lemma 2.1]: Let $(X, |||\cdot|||)$ *be an arbitrary Banach space. Suppose that* $A, B \subset X$. *Then*

$$\gamma_2(A+B, |||\cdot|||) \lesssim \gamma_2(A, |||\cdot|||) + \gamma_2(B, |||\cdot|||).$$

Let $\mathcal{X}$ be any set of matrices and define $d_F(S) = \sup_{A\in\mathcal{X}} \|A\|_F$ and $d_{2\to 2}(S) = \sup_{A\in\mathcal{X}} \|A\|_{2\to 2}$. We can now state the following theorem, which will be crucial in Section VI-B.

Theorem 13 [58, Th. 1.4]: Let $\mathcal{X}$ *be a symmetric set of matrices, i.e.,* $\mathcal{X} = -\mathcal{X}$ *and let* ξ *be a random vector whose entries* $\xi_i \sim \mathcal{CN}(0,1)$ *are independent. Set*

$$\begin{aligned}E &= \gamma_2(\mathcal{X}, \|\cdot\|_{2\to 2})(\gamma_2(\mathcal{X}, \|\cdot\|_{2\to 2}) + d_F(\mathcal{X})) \\ V &= d_{2\to 2}(\mathcal{X})(\gamma_2(\mathcal{X}, \|\cdot\|_{2\to 2}) + d_F(\mathcal{X})) \\ U &= d_{2\to 2}^2(\mathcal{X})\end{aligned}$$

Then, for $t > 0$,

$$\mathbb{P}\left(\sup_{A\in\mathcal{X}} \left|\|A\xi\|_{\ell_2}^2 - \mathbb{E}\|A\xi\|_{\ell_2}^2\right| \ge c_1 E + t\right) \le 2\exp\left(-c_2 \min\left(\frac{t^2}{V^2}, \frac{t}{U}\right)\right).$$

The constants c_1 *and* c_2 *are universal.*

Dudley's inequality yields a relation of the γ_2-functional to covering numbers. Recall that the covering number $N(S, |||\cdot|||, \varepsilon)$ is the minimum number of open $|||\cdot|||$-balls with radius ε, whose midpoint is contained in S, which are needed to cover S. More precisely, Dudley's inequality (see [57, Proposition 2.2.10], [59]) states that

$$\gamma_2(S, |||\cdot|||) \lesssim \int_0^{d_{|||\cdot|||}(S)} \sqrt{\log N(S, |||\cdot|||, \varepsilon)}\, d\varepsilon, \qquad \text{(IV.3)}$$

where $d_{|||\cdot|||}(S) = \sup_{x\in S} |||x|||$. For this reason, we will need some bounds for covering numbers, which are summarized in the following section.

C. Covering Numbers

The following lemma is a slight modification of the Maurey lemma by Carl [60]. (See also [58, Lemma 4.2] for a formulation of this lemma using notation which is closer to the notation in this paper.)

Lemma 14: Let $(X, |||\cdot|||)$ be a normed space, consider a finite set $\mathcal{U} \subset X$, and assume that for every $L \in \mathbb{N}$ and $(u_1, \ldots, u_L) \in \mathcal{U}^L$, $\mathbb{E}_\varepsilon \left|\left|\left| \sum_{j=1}^L \varepsilon_j u_j \right|\right|\right| \leq A\sqrt{L}$, where $(\varepsilon_j)_{j=1}^L$ denotes a Rademacher vector. Then, for every $u > 0$,

$$\log N\left(conv\left(\mathcal{U}\right), |||\cdot|||, u\right) \lesssim \frac{A^2}{u^2} \log |\mathcal{U}|,$$

where $|\mathcal{U}|$ denotes the cardinality of $\mathcal{U}$.

Let $V \subset \mathbb{R}^n$ be a compact, convex, and symmetric set which is absorbing, i.e. $\mathbb{R}^n = \bigcup_{t>0} tV$. We will denote by $\|\cdot\|_V$ the norm associated with V, i.e., the unique norm whose unit ball is given by V. Furthermore, denote by V° the polar body of V, i.e.,

$$V^\circ = \left\{u \in \mathbb{R}^n : \langle u, v\rangle \leq 1 \text{ for all } v \in V\right\}.$$

An elementary consequence of the definition is that the dual norm of $\|\cdot\|_V$ is given by $\|\cdot\|_{V^\circ}$. The following result about covering numbers of polar bodies solved a special instance of a conjecture by Pietsch [61].

Theorem 15 [62]: As above, assume $V \subset \mathbb{R}^n$ to be a compact, convex, symmetric, and absorbing set. Then, for all $\varepsilon > 0$

$$\log N\left(B\left(0,1\right), \|\cdot\|_V, \varepsilon\right) \lesssim \log N\left(V^\circ, \|\cdot\|_{\ell_2}, c\varepsilon\right),$$

where $c > 0$ is a universal constant and $B(0,1) := \left\{x \in \mathbb{R}^n : \|x\|_{\ell_2} \leq 1\right\}$.

V. Outline of the Proof

In this section we give a rough outline of our proof and highlight the main differences to previous work ([27] and [43]). In particular, we want to point out those parts, which enabled us to overcome the suboptimal scaling in r. The overall strategy of our proof remains similar to the one in [43] and [27]: First, we will prove sufficient conditions for recovery. These conditions will rely on the existence of a so-called inexact dual certificate. In the second step this certificate will be constructed via the Golfing Scheme, a method which has been introduced by Gross and others (see [30]).

As already mentioned, the first part of the proof consists of showing that the existence of the inexact dual certificate is a sufficient condition for recovery. This will be proven in Section VI-A. The underlying observation is that in certain cases, it suffices that standard conditions defining minimizers are only approximately satisfied. In [43], these perturbed conditions are given by [43, eq. 28]. In order for them to imply that X^0 is a minimizer, one needs that $\mathcal{A}_i$ acts approximately as an isometry on each

$$\mathcal{T}_i = \left\{h_i u_i^* + v_i m_i^* : \ u_i \in \mathbb{C}^{K_i}, v_i \in \mathbb{C}^{N_i}\right\}$$

and that the images of these operators are almost orthogonal to each other. The latter condition is responsible for the appearance of the quadratic scaling in r in [43]. Our approach will be different: We will show that the operator $\mathcal{A}$ acts as an approximate isometry on the full subspace

$$\mathcal{T} := \left\{(X_1, \ldots, X_r) : \ X_i \in \mathcal{T}_i \text{ for all } i \in [r]\right\}.$$

in the sense of the following definition.

Definition 16 (Local Isometry Property): $\mathcal{A}$ fulfills the δ-local isometry property on $\mathcal{T}$ for some $\delta > 0$, if

$$(1-\delta)\|X\|_F^2 \leq \|\mathcal{A}(X)\|_{\ell_2}^2 \leq (1+\delta)\|X\|_F^2 \tag{V.1}$$

for all $X \in \mathcal{T}$.

The main novelty in our proof is that our global viewpoint allows us to establish the local isometry property on $\mathcal{T}$ with high probability if L scales *linearly* with r. This will be achieved via Theorem 13, which involves a γ_2-functional. Thus a large part of Section VI-B is concerned with estimating those γ_2-functionals.

The local isometry property is not only needed in the first part but also in the second part of the proof, where the dual certificate is constructed via the Golfing Scheme. For that we will assume that $\left\{\Gamma_p\right\}_{p=1}^P$ is fixed ω-admissible partition (see Definition 2) which minimizes (II.8). For this partition we can introduce the operators $\mathcal{A}^p$ defined by $\mathcal{A}^p(X) = P_{\Gamma_p}(\mathcal{A}(X))$, where $P_{\Gamma_p} : \mathbb{C}^L \to \mathbb{C}^L$ denotes the (coordinate) projection of $\mathbb{C}^L$ onto the coordinates contained in the set Γ_p. Similarly, we will define $\mathcal{A}_i^p$ by $\mathcal{A}_i^p(X) = P_{\Gamma_p}(\mathcal{A}_i(X))$.

We will need that each operator $\mathcal{A}^p$ satisfies the δ-local isometry property on a subspace $\mathcal{T}^p$, which is slightly larger than $\mathcal{T}$. In order to define the space $\mathcal{T}^p$ we need to introduce the operators $\mathcal{S}^p : \mathcal{M} \to \mathcal{M}$. For that, recall $S_{i,p} = T_{i,p}^{-1}$ as defined in Section II-C.

Definition 17: Let $p \in [P]$. Then the operator $\mathcal{S}^p : \mathcal{M} \to \mathcal{M}$ is defined by

$$\mathcal{S}^p(W) = \left(S_{1,p}W_1, \ldots, S_{r,p}W_r\right) \tag{V.2}$$

for $W = (W_1, \ldots, W_r) \in \mathcal{M}$.

Then $\mathcal{T}^p$ is defined by

$$\mathcal{T}^p = \mathcal{T} + \mathcal{S}^p(\mathcal{T}). \tag{V.3}$$

Observe that we may write $\mathcal{T} = \mathcal{T}_h + \mathcal{T}_m$ and $\mathcal{T}^p = \mathcal{T}_h + \mathcal{T}_{\mathcal{S}^p h} + \mathcal{T}_m$, when the subspaces $\mathcal{T}_m$, $\mathcal{T}_h$, and $\mathcal{T}_{\mathcal{S}^p h}$ are given by

$$\begin{aligned}
\mathcal{T}_m &= \left\{\left(v_1 m_1^*, \ldots, v_r m_r^*\right) : \ v_i \in \mathbb{C}^{K_i} \text{ for all } i \in [r]\right\}, \\
\mathcal{T}_h &= \left\{\left(h_1 u_1^*, \ldots, h_r u_r^*\right) : \ u_i \in \mathbb{C}^{N_i} \text{ for all } i \in [r]\right\}, \\
\mathcal{T}_{\mathcal{S}^p h} &= \Big\{ \left(\left(S_{1,p}h_1\right)u_1^*, \ldots, \left(S_{r,p}h_r\right)u_r^*\right) : \\
&\qquad\qquad u_i \in \mathbb{C}^{N_i} \text{ for all } i \in [r]\Big\}.
\end{aligned} \tag{V.4}$$

As already mentioned, the local isometry property on $\mathcal{T}$, respectively $\mathcal{T}^p$, will be shown in Section VI-B. In Section VI-C the approximate dual certificate will be constructed via the Golfing Scheme. Finally, in Section VI-D we will prove the main result, Theorem 6.

VI. Proof of the Main Theorem

A. Sufficient Conditions for Recovery

As already mentioned in the outline of the proof, in this section we will show that the existence of an inexact dual certificate implies that the signal is approximately recovered. Therefore, we will denote in the following by $\mathcal{P}_\mathcal{T}$ the orthogonal projection onto $\mathcal{T}$. Similarly, we will denote by for all $i \in [r]$ the orthogonal projection onto $\mathcal{T}_i$.

Lemma 18: Suppose that $\mathcal{A}$ satisfies the δ-local isometry property on $\mathcal{T}$ (V.1) and set $\gamma = \|\mathcal{A}\|_{F\to 2}$, i.e., γ is the operator norm of $\mathcal{A}$. Furthermore, suppose that there is $Y = (Y_1, \dots, Y_r) = \mathcal{A}^* z$ for some $z \in \mathbb{C}^L$ such that

$$\|\mathcal{P}_{\mathcal{T}} Y - sgn\left(X^0\right)\|_F \le \alpha \tag{VI.1}$$

$$\|\mathcal{P}_{\mathcal{T}_i^\perp} Y_i\|_{2\to 2} \le \beta \text{ for all } i \in [r], \tag{VI.2}$$

where $\alpha, \beta \ge 0$ are constants such that $1 - \beta - \frac{\alpha\gamma}{\sqrt{1-\delta}} \ge \frac{1}{2}$, $\alpha \le 1$, and $\sqrt{1-\delta} \ge \frac{1}{2}$. Then if $\hat{X}$ is a minimizer of

$$\begin{aligned} &\text{minimize } \|X\|_* \\ &\text{subject to } \|\mathcal{A}(X) - \hat{y}\|_{\ell_2} \le \tau \end{aligned}$$

we have that

$$\|\hat{X} - X^0\|_F \lesssim \tau\,(1+\gamma)\left(1 + \|z\|_{\ell_2}\right). \tag{VI.3}$$

Proof: Set $V = (V_1, \dots, V_r) = \hat{X} - X^0$ and note that we seek to estimate $\|V\|_F \le \|\mathcal{P}_{\mathcal{T}}(V)\|_F + \|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F$ from above. We observe that

$$\|\mathcal{A}(V)\|_{\ell_2} \le \|\mathcal{A}(\hat{X}) - \hat{y}\|_{\ell_2} + \|\hat{y} - \mathcal{A}\left(X^0\right)\|_{\ell_2} \le 2\tau. \tag{VI.4}$$

Together with the δ-local isometry property (V.1), the definition of γ, and the triangle inequality we obtain

$$\begin{aligned} \|\mathcal{P}_{\mathcal{T}}(V)\|_F &\le \frac{1}{\sqrt{1-\delta}} \|\mathcal{A}(\mathcal{P}_{\mathcal{T}}(V))\|_{\ell_2} \\ &\le \frac{1}{\sqrt{1-\delta}} \left(\|\mathcal{A}\left(\mathcal{P}_{\mathcal{T}^\perp}(V)\right)\|_{\ell_2} + \|\mathcal{A}(V)\|_{\ell_2}\right) \\ &\le \frac{1}{\sqrt{1-\delta}} \left(\gamma \|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F + 2\tau\right). \end{aligned}$$

Thus it remains to find an upper bound for $\|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F$. For that purpose, choose $Z = (Z_1, \dots, Z_r)$ such that for all $i \in [r]$ we have that $Z_i \in \mathcal{T}_i^\perp$, $\|Z_i\|_{2\to 2} \le 1 - \beta$, and $\langle Z_i, V_i\rangle_F = (1-\beta)\|\mathcal{P}_{\mathcal{T}_i^\perp} V_i\|_*$. This is possible by duality of the norms $\|\cdot\|_{2\to 2}$ and $\|\cdot\|_*$ (see [63, Sec. 4.2]). Observe that and $\|\operatorname{sgn}\left(X_i^0\right) + \mathcal{P}_{\mathcal{T}_i^\perp} Y_i + Z_i\|_{2\to 2} \le 1$ as both the row and column spaces of $\operatorname{sgn}\left(X_i^0\right)$ and $\mathcal{P}_{\mathcal{T}_i^\perp} Y_i + Z_i$ are orthogonal. Thus, again using the duality between $\|\cdot\|_{2\to 2}$ and $\|\cdot\|_*$, we obtain

$$\begin{aligned} \|X_i^0 + V_i\|_* &= \sup_{W \in \mathbb{C}^{K_i \times N_i},\ \|W\|_{2\to 2} \le 1} |\langle W, X_i^0 + V_i\rangle_F| \\ &\ge \operatorname{Re}\left(\langle \operatorname{sgn}\left(X_i^0\right) + \mathcal{P}_{\mathcal{T}_i^\perp} Y_i + Z_i, X_i^0 + V_i\rangle_F\right) \\ &\ge \|X_i^0\|_* + \operatorname{Re}\left(\langle \operatorname{sgn}\left(X_i^0\right) + \mathcal{P}_{\mathcal{T}_i^\perp} Y_i, V_i\rangle_F\right) \\ &\quad + (1-\beta)\|\mathcal{P}_{\mathcal{T}_i} V_i\|_* \end{aligned}$$

Here, in the second inequality we used that $\mathcal{P}_{\mathcal{T}_i^\perp} Y_i + Z_i \in \mathcal{T}_i^\perp$ and $\langle \operatorname{sgn}\left(X_i^0\right), X_i^0\rangle_F = \|X_i^0\|_*$. Thus, by definition of $\|X^0 + V\|_*$ we obtain

$$\begin{aligned} &\|X^0 + V\|_* \\ &\quad \ge \sum_{i=1}^r \|X_i^0\|_* + \sum_{i=1}^r \operatorname{Re}\left(\langle \operatorname{sgn}\left(X_i^0\right) + \mathcal{P}_{\mathcal{T}_i^\perp} Y_i, V_i\rangle_F\right) \\ &\qquad + (1-\beta)\sum_{i=1}^r \|\mathcal{P}_{\mathcal{T}_i} V_i\|_* \\ &\quad = \|X^0\|_* + \operatorname{Re}\left(\langle \operatorname{sgn}\left(X^0\right) - \mathcal{P}_{\mathcal{T}} Y, V\rangle_F + \langle Y, V\rangle_F\right) \\ &\qquad + (1-\beta)\|\mathcal{P}_{\mathcal{T}^\perp} V\|_*. \end{aligned}$$

Now observe that by Cauchy-Schwarz, (VI.1) and our upper bound for $\|\mathcal{P}_{\mathcal{T}}(V)\|_{\ell_2}$

$$\begin{aligned} &\operatorname{Re}\left(\langle \operatorname{sgn}\left(X^0\right) - \mathcal{P}_{\mathcal{T}}(Y), V\rangle_F\right) \\ &\quad \ge -\|\operatorname{sgn}\left(X^0\right) - \mathcal{P}_{\mathcal{T}}(Y)\|_F \|\mathcal{P}_{\mathcal{T}}(V)\|_F \\ &\quad \ge \frac{-\alpha}{\sqrt{1-\delta}} \left(\gamma \|\mathcal{P}_{\mathcal{T}^\perp} V\|_F + 2\tau\right). \end{aligned}$$

Furthermore, we note that by Cauchy-Schwarz and (VI.4)

$$\begin{aligned} \operatorname{Re}\left(\langle Y, V\rangle_F\right) &= \operatorname{Re}\left(\langle \mathcal{A}^* z, V\rangle_F\right) = \left(\langle z, \mathcal{A}(V)\rangle_{\ell_2}\right) \\ &\ge -2\|z\|_{\ell_2}\tau. \end{aligned}$$

Combining the last three calculations and using that the nuclear norm is greater or equal than the Frobenius norm we obtain

$$\begin{aligned} \|\hat{X}\|_* \ge \|X^0\|_* &+ \left(1 - \beta - \frac{\alpha\gamma}{\sqrt{1-\delta}}\right)\|\mathcal{P}_{\mathcal{T}^\perp} V\|_* \\ &- 2\tau\left(\|z\|_{\ell_2} + \frac{\alpha}{\sqrt{1-\delta}}\right). \end{aligned}$$

As $\hat{X}$ is the nuclear norm minimizer and we have that $\|X^0\|_* \ge \|\hat{X}\|_*$ this yields

$$\left(1 - \beta - \frac{\alpha\gamma}{\sqrt{1-\delta}}\right)\|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_* \le 2\tau\left(\|z\|_{\ell_2} + \frac{\alpha}{\sqrt{1-\delta}}\right).$$

By our assumptions on α, β, and δ this implies

$$\|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F \lesssim \tau\left(\|z\|_{\ell_2} + 1\right).$$

Thus, using again the upper bound for $\|\mathcal{P}_{\mathcal{T}}(V)\|_F$, which was calculated above, and again our assumptions on α, β, and δ we obtain

$$\begin{aligned} \|V\|_F &\le \|\mathcal{P}_{\mathcal{T}}(V)\|_F + \|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F \\ &\lesssim (1+\gamma)\|\mathcal{P}_{\mathcal{T}^\perp}(V)\|_F + \tau \\ &\lesssim \tau\,(1+\gamma)\left(1 + \|z\|_{\ell_2}\right), \end{aligned}$$

which finishes the proof. □

As already mentioned in the introduction, the noiseless case is also of interest for us. Note that in this situation we may set $\tau = 0$ and Lemma 18 shows that the existence of a dual certificate implies that the convex program (I.5) recovers the signal X^0 exactly.

Remark 19: Note that we still have the freedom to choose the parameters α and β in Lemma 18. In Section VI-C we will construct a dual certificate Y for the following choice of parameters: We set $\beta = \frac{1}{4}$ and assume that $\delta \le \frac{1}{4}$. In order

to fulfill the condition $1-\beta-\frac{\alpha\gamma}{\sqrt{1-\delta}} \geq \frac{1}{2}$ *it is then enough to choose* $\alpha = \frac{1}{8\gamma}$.

Note that in the noisy case the error estimate in Lemma 18 depends linearly on the operator norm of $\mathcal{A}$ as (VI.3) states. Thus, we need an upper bound for the operator norm of $\mathcal{A}$ which holds with high probability.

Lemma 20: Let $\omega \geq 1$. *Then with probability at least* $1-2L^{-\omega}$ *we have that*

$$\|\mathcal{A}\|_{F\to 2} \leq 2\sqrt{\omega \max\left\{1; \frac{rK_\mu N}{L}\right\}\log\left(L+rKN\right)}.$$

Proof: The result will be proven by using Corollary 10. Indeed, we can represent each operator $\mathcal{A}_i$ as $\mathcal{A}_i = \sum_{\ell\in L}\sum_{j=1}^{K_i}\mathcal{B}_{\ell,j}$ such that each operator $\mathcal{B}_{\ell,j}$ depends linearly on the (ℓ,k)th entry of C_i, i.e., $(C_i)_{\ell,k} \sim \mathcal{CN}(0,1)$. Thus, we need to estimate the operator norms of $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right]$ and $\mathbb{E}\left[\mathcal{A}\mathcal{A}^*\right]$. Observe that

$$\mathcal{A}^*\mathcal{A} = \left(\mathcal{A}_1^*\left(\sum_{i=1}^r \mathcal{A}_i\right), \ldots, \mathcal{A}_r^*\left(\sum_{i=1}^r \mathcal{A}_i\right)\right).$$

Note that the operators $\{\mathcal{A}_i\}_{i=1}^r$ are independent with expectation $\mathbb{E}[\mathcal{A}_i] = 0$ for all $i \in [r]$. Thus $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = \left(\mathbb{E}\left[\mathcal{A}_1^*\mathcal{A}_1\right], \ldots, \mathbb{E}\left[\mathcal{A}_r^*\mathcal{A}_r\right]\right)$. Let $Z = (Z_1, \ldots, Z_r) \in \mathcal{M}$. Using (II.2) we compute

$$\begin{aligned}\mathbb{E}\left[\left(\mathcal{A}_i^*\mathcal{A}_i\right)(Z_i)\right] &= \sum_{\ell=1}^L \mathbb{E}\left[\left(\mathcal{A}_i\left(Z_i\right)(\ell)\right) b_{i,\ell}c_{i,\ell}^*\right] \\ &= \sum_{\ell=1}^L \mathbb{E}\left[b_{i,\ell}b_{i,\ell}^* Z_i c_{i,\ell}c_{i,\ell}^*\right] \\ &= \sum_{\ell=1}^L b_{i,\ell}b_{i,\ell}^* Z_i = Z_i \qquad \text{(VI.5)}\end{aligned}$$

Thus, $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}(Z)\right] = Z$ for any $Z \in \mathcal{M}$, which implies $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = \mathrm{Id}$. To compute $\mathbb{E}\left[\mathcal{A}\mathcal{A}^*\right]$ let $y \in \mathbb{C}^L$ be arbitrary. We compute with similar arguments as before

$$\begin{aligned}\mathbb{E}\left[\left(\mathcal{A}\mathcal{A}^*y\right)(\ell)\right] &= \sum_{i=1}^r \mathbb{E}\left[\left(\mathcal{A}_i\mathcal{A}_i^*y\right)(\ell)\right] \qquad \text{(VI.6)} \\ &= \sum_{i=1}^r \mathbb{E}\left[b_{i,\ell}^*\left(\mathcal{A}_i^*y\right)c_{i,\ell}\right] \\ &\overset{(II.2)}{=} \sum_{i=1}^r\sum_{\ell'=1}^L y\left(\ell'\right)\mathbb{E}\left[b_{i,\ell}^*b_{i,\ell'}c_{i,\ell'}^*c_{i,\ell}\right] \\ &= y(\ell)\sum_{i=1}^r \mathbb{E}\left[b_{i,\ell}^*b_{i,\ell}c_{i,\ell}^*c_{i,\ell}\right] \\ &= y(\ell)\sum_{i=1}^r \|b_{i,\ell}\|_{\ell_2}^2 N_i. \qquad \text{(VI.7)}\end{aligned}$$

This shows that $\mathcal{A}\mathcal{A}^*$ can be represented as a diagonal matrix with entries $\sum_{i=1}^r \|b_{i,\ell}\|_{\ell_2}^2 N_i$. Thus, by definition of $K_{i,\mu}$ (II.4), $\|\mathbb{E}\left[\mathcal{A}\mathcal{A}^*\right]\|_{2\to 2} \leq \frac{N\sum_{i=1}^r K_{i,\mu}}{L}$, which implies, together with (VI.5)

$$\begin{aligned}\sigma^2 &= \max\left\{\|\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right]\|_{F\to F};\ \|\mathbb{E}\left[\mathcal{A}\mathcal{A}^*\right]\|_{2\to 2}\right\} \\ &\leq \max\left\{1;\ \frac{N\sum_{i=1}^r K_{i,\mu}}{L}\right\}.\end{aligned}$$

Consequently, Corollary 10 with $t = \omega\log L$ yields that with probability exceeding $1-2L^{-\omega}$

$$\|\mathcal{A}\|_{F\to 2} \leq \max\left\{1;\ \sqrt{\frac{N\sum_{i=1}^r K_{i,\mu}}{L}}\right\} \cdot\sqrt{2\left(\omega\log L + \log\left(L+rKN\right)\right)},$$

which implies the result. □

Remark 21: Note that in (VI.7) *and other places below, only a weighted sum of the* $\|b_{i,\ell}\|_{\ell_2}^2$ *appears. If the summands vastly differ, this may be too crude, and one may consider attempting an averaging argument similar to the one in [64]. This would, however, require that the proof is completely reworked in some parts. To achieve condition* (VI.2)*, for example, we currently rely very much on bounding each* $K_{i,\mu}$ *individually.*

B. Local Isometry Property

In this subsection, we establish an isometry of $\mathcal{A}$, respectively of $\mathcal{A}^p$, on $\mathcal{T}$, respectively $\mathcal{T}^p$. More precisely, we establish the following theorem.

Theorem 22: Fix $\omega \geq 1$. *Suppose that*

$$Q \geq C_\omega\delta^{-2}r\left(K_\mu\log(L)\log^2\left(K_\mu\right) + N\mu_h^2\right). \qquad \text{(VI.8)}$$

Then with probability $1-\mathcal{O}\left(L^{-\omega}\right)$ *the following is true: All* $X \in \mathcal{T}$ *fulfill*

$$(1-\delta)\|X\|_F^2 \leq \left\|\mathcal{A}(X)\right\|_{\ell_2}^2 \leq (1+\delta)\|X\|_F^2 \qquad \text{(VI.9)}$$

and for all $p \in [P]$ *every* $Y \in \mathcal{T}^p = \mathcal{T}+\mathcal{S}^p\mathcal{T}$ *satisfies*

$$\begin{aligned}(1-\delta)\sum_{i=1}^r \|T_{i,p}^{1/2}Y_i\|_F^2 &\leq \frac{L}{Q}\left\|\mathcal{A}^p(Y)\right\|_{\ell_2}^2 \\ &\leq (1+\delta)\sum_{i=1}^r \|T_{i,p}^{1/2}Y_i\|_F^2, \qquad \text{(VI.10)}\end{aligned}$$

where $T_{i,p}^{1/2}$ *denotes the unique positive, self-adjoint matrix whose square is equal to* $T_{i,p}$.

The proof of this theorem is divided into several steps. For the proof we need some additional notation. Recall that the incoherence parameter μ_h^2 measures the alignment between the vectors $h_i \in \mathbb{C}^{K_i}$ and $b_{i,\ell} \in \mathbb{C}^{K_i}$. As the operators $\mathcal{A}$ and $\mathcal{A}_i$ are defined on matrices, it will to be useful to generalize the notion of incoherence from vectors to matrices. This is achieved by the following definition.

Definition 23: For all $i \in [r]$, *vectors* $z \in \mathbb{C}^{K_i}$ *and matrices* $Z_i \in \mathbb{C}^{K_i\times N_i}$ *define*

$$\|z\|_{B_i} = \sqrt{L}\max_{\ell\in[L]}|z^*b_{i,\ell}|$$

and

$$\|Z_i\|_{B_i} = \sqrt{L} \max_{\ell\in[L]} \|Z_i^* b_{i,\ell}\|_{\ell_2}.$$

For $Z = (Z_1, \dots, Z_r) \in \mathcal{M}$ *we define*

$$\|Z\|_B = \sqrt{L \max_{\ell\in[L]} \left(\sum_{i=1}^r \left\| Z_i^* b_{i,\ell} \right\|_{\ell_2}^2 \right)}.$$

All these three operations are norms as $\sum_{\ell=1}^L b_{i,\ell} b_{i,\ell}^* = \mathrm{Id}_{K_i}$ for all $i \in [r]$. The following lemma provides us with some useful estimates.

Lemma 24: Let $Z = (Z_1, \dots, Z_r) \in \mathcal{M}$, $i \in [r]$ *and* $z \in \mathbb{C}^{K_i}$. *Then*

$$\|z\|_{B_i} \le \sqrt{K_{i,\mu}}\, \|z\|_{\ell_2} \tag{VI.11}$$

$$\|Z_i\|_{B_i} \le \sqrt{K_{i,\mu}} \|Z_i\|_{2\to 2} \tag{VI.12}$$

$$\|Z\|_B \le \sqrt{\sum_{i=1}^r \|Z_i\|_{B_i}^2} \le \sqrt{K_\mu} \|Z\|_F. \tag{VI.13}$$

Proof: In order to prove (VI.12) note that for $Z_i \in \mathbb{C}^{K_i}$ and $\ell \in [L]$ due to the definition of $K_{i,\mu}$

$$\|Z_i^* b_{i,\ell}\|_{\ell_2}^2 \le \|Z_i\|_F^2 \|b_{i,\ell}\|_{\ell_2}^2 \overset{(II.4)}{\le} \frac{K_{i,\mu}}{L} \|Z_i\|_{2\to 2}^2.$$

Taking the maximum over all $\ell \in [L]$ shows (VI.12). Inequality (VI.11) can be proven analogously. (VI.13) follows from

$$\|Z\|_B^2 \le L \sum_{i=1}^r \max_{\ell\in[L]} \|Z_i^* b_{i,\ell}\|_{\ell_2}^2 = \sum_{i=1}^r \|Z_i\|_{B_i}^2$$

combined with (VI.12) and the definition of $\|Z\|_F$. □

The notion of $\|\cdot\|_B$-norms together with Theorem 13 allows us to state the following abstract isometry result, where we will use the notation $d_B(\mathcal{X}) = \sup_{X\in\mathcal{X}} \|X\|_B$.

Proposition 25: Let $\mathcal{X} = -\mathcal{X} \subset \mathcal{M}$ *be a symmetric set and consider*

$$\widehat{E} = \frac{\gamma_2(\mathcal{X}, \|\cdot\|_B)}{\sqrt{Q}} \left(\frac{\gamma_2(\mathcal{X}, \|\cdot\|_B)}{\sqrt{Q}} + d_F(\mathcal{X}) \right)$$

$$\widehat{V} = \frac{d_B(\mathcal{X})}{\sqrt{Q}} \left(\frac{\gamma_2(\mathcal{X}, \|\cdot\|_B)}{\sqrt{Q}} + d_F(\mathcal{X}) \right)$$

$$\widehat{U} = \frac{1}{Q} d_B^2(\mathcal{X}).$$

Then, for $t > 0$ *and all* $p \in [P]$,

$$\mathbb{P}\left(\sup_{X\in\mathcal{X}} \left| \frac{L}{Q} \|\mathcal{A}^p(X)\|_{\ell_2}^2 - \sum_{i=1}^r \|T_{i,p}^{1/2} X_i\|_F^2 \right| \ge \tilde{c}_1 \widehat{E} + t \right) \le 2\exp\left(-\tilde{c}_2 \min\left(\frac{t^2}{\widehat{V}^2}, \frac{t}{\widehat{U}} \right) \right) \tag{VI.14}$$

$$\mathbb{P}\left(\sup_{X\in\mathcal{X}} \left| \|\mathcal{A}(X)\|_{\ell_2}^2 - \|X\|_F^2 \right| \ge \tilde{c}_3 \widehat{E} + t \right) \le 2\exp\left(-\tilde{c}_4 \min\left(\frac{t^2}{\widehat{V}^2}, \frac{t}{\widehat{U}} \right) \right), \tag{VI.15}$$

provided $\{\Gamma_p\}_{p=1}^P$ *is a* ω*-admissible partition of* $[L]$. *The constants* $\tilde{c}_1$, $\tilde{c}_2$, $\tilde{c}_3$, *and* $\tilde{c}_4$ *are universal.*

Proof: We will start by proving (VI.14). Fix $p \in [P]$. For $X = (X_1, \dots, X_r) \in \mathcal{X}$ let $H_X \in \mathbb{C}^{L\times Q \sum_{i=1}^r N_i}$ be the block diagonal matrix, whose diagonal elements, indexed by $\ell \in \Gamma_p$ are given by row vectors of the form $\sqrt{\frac{L}{Q}} \left(b_{1,\ell}^* X_1, \dots, b_{r,\ell}^* X_r \right)$. Furthermore, set $\mathcal{H}_X = \{H_X : X \in \mathcal{X}\}$. Observe that

$$\|H_X\|_F^2 = \frac{L}{Q} \sum_{\ell\in\Gamma_p} \sum_{i=1}^r \|X_i^* b_{i,\ell}\|_{\ell_2}^2 = \sum_{i=1}^r \mathrm{Tr}\left(X_i X_i^* T_{i,p} \right) = \sum_{i=1}^r \|T_{i,p}^{1/2} X_i\|_F^2, \tag{VI.16}$$

$$\|H_X\|_{2\to 2} = \sqrt{\frac{L}{Q}} \max_{\ell\in\Gamma_p} \| \left(b_{1,\ell}^* X_1, \dots, b_{r,\ell}^* X_r \right) \|_{\ell_2} \le \frac{1}{\sqrt{Q}} \|X\|_B. \tag{VI.17}$$

Let $\xi^{(p)}$ be the concatenation of all the random base vectors $c_{i,\ell}$, where $i \in [r]$, $\ell \in \Gamma_p$. Then

$$\frac{L}{Q} \|\mathcal{A}^p(X)\|_{\ell_2}^2 = \frac{L}{Q} \sum_{\ell\in\Gamma_p} |\mathcal{A}^p(X)(\ell)|^2 = \frac{L}{Q} \sum_{\ell\in\Gamma_p} \left| \sum_{i=1}^r b_{i,\ell}^* X_i c_{i,\ell} \right|^2 = \|H_X \xi^{(p)}\|_{\ell_2}^2$$

and

$$\sum_{i=1}^r \|T_{i,p}^{1/2} X_i\|_F^2 = \|H_X\|_F^2 = \mathbb{E}[\|H_X \xi^{(p)}\|_{\ell_2}^2].$$

Consequently

$$\sup_{X\in\mathcal{X}} \left| \frac{L}{Q} \|\mathcal{A}^p(X)\|_{\ell_2}^2 - \sum_{i=1}^r \|T_{i,p}^{1/2} X_i\|_F^2 \right| = \sup_{X\in\mathcal{X}} \left| \|H_X \xi^{(p)}\|_{\ell_2}^2 - \mathbb{E}\left[\|H_X \xi^{(p)}\|_{\ell_2}^2 \right] \right|$$

and inequality (VI.14) follows from Theorem 13, equation (VI.16), (VI.17) combined with the fact that $\sum_{i=1}^r \|T_{i,p}^{1/2} X_i\|_F^2 \overset{(II.5)}{\le} 2\|X\|_F^2$. Inequality (VI.15) follows in an analogous way by letting H_X be the block diagonal matrix, whose diagonal elements, indexed by $\ell \in [L]$, are given by $\left(b_{1,\ell}^* X_1, \dots, b_{r,\ell}^* X_r \right)$. Furthermore, one uses $\sum_{\ell=1}^L b_{i,\ell} b_{i,\ell}^* = \mathrm{Id}$ instead of $\frac{L}{Q} \sum_{\ell\in\Gamma_p} b_{i,\ell} b_{i,\ell}^* = T_{i,p}$. □

Our strategy to prove Theorem 22 will now be to apply Proposition 25 with appropriately chosen sets $\mathcal{X}$. For $\mathcal{T}_m$, $\mathcal{T}_h$, and $\mathcal{T}_{\mathcal{S}^p h}$ as in (V.4), define

$$B^m = \{X \in \mathcal{T}_m : \|X\|_F \le 1\}$$

$$B^h = \{X \in \mathcal{T}_h : \|X\|_F \le 1\}$$

$$B^{\mathcal{S}^p h} = \{X \in \mathcal{T}_{\mathcal{S}^p h} : \|X\|_F \le 1\}$$

and observe that in order to prove the δ-local isometry property on $\mathcal{T}$ it is enough to apply Proposition 25 to the set $\mathcal{W}$ defined by

$$\mathcal{W} = B^h + B^m. \tag{VI.18}$$

Similarly, in order to prove the δ-local isometry property on $\mathcal{T}^p$ for $p \in [P]$ it is enough to apply Proposition 25 to the set $\mathcal{W}^p$ defined by

$$\mathcal{W}^p = B^h + B^{\mathcal{S}^p h} + B^m. \tag{VI.19}$$

That is, it remains to estimate the γ_2-functionals of $\mathcal{W}$ and $\mathcal{W}^p$ with respect to $\|\cdot\|_B$. By Dudley's inequality (IV.3) one can bound the γ_2-functional by an integral involving covering numbers. To estimate those, we need the following technical lemmas.

Lemma 26: Let B^m be the above defined set. Then

$$N\left(B^m, \|\cdot\|_B, \varepsilon\right) \leq N\left(B\left(0,1\right) \subset \mathbb{R}^r, \|\cdot\|_{\ell_2}, \frac{\varepsilon}{2\sqrt{K_\mu}}\right) \cdot \prod_{i=1}^{r} N\left(B\left(0,1\right) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right).$$

(By $B\left(0,1\right)$ we always denote the closed unit ball with respect to the $\|\cdot\|_{\ell_2}$-norm.)

This lemma is actually a slight modification of [29, Lemma 3.1]. For the convenience of the reader we have included a proof in Appendix C.

Lemma 27: For all $i \in [r]$

$$\log N\left(B\left(0,1\right) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right) \lesssim \frac{K_{i,\mu}}{\varepsilon^2} \log\left(L\right). \tag{VI.20}$$

Proof: Our goal is to apply Theorem 15 to $\log N\left(B\left(0,1\right) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)$. However, as $\|\cdot\|_{B_i}$ is a norm defined on a complex vector space we first need to transfer this setting into an appropriate real vector space framework. For that goal we will use the isometric embedding $P : \mathbb{C}^{K_i} \to \mathbb{R}^{2K_i}$ given by $x = \left(x_1, \ldots, x_{K_i}\right) \in \mathbb{C}^{K_i} \mapsto \left(\left(\mathrm{Re}\ x\right)_1, \left(\mathrm{Im}\ x\right)_1, \ldots, \left(\mathrm{Re}\ x\right)_{K_i}, \left(\mathrm{Im}\ x\right)_{K_i}\right)$. Furthermore, note that for all $x \in \mathbb{C}^{K_i}$

$$\begin{aligned}\|x\|_{B_i} &= \sqrt{L} \max_{\ell\in[L]} |\langle x, b_{i,\ell}\rangle| \\ &= \sqrt{L} \max_{\ell\in[L]} \sqrt{\left(\mathrm{Re}\ \langle x, b_{i,\ell}\rangle\right)^2 + \left(\mathrm{Im}\ \langle x, b_{i,\ell}\rangle\right)^2} \\ &\leq \sqrt{2L} \max_{\ell\in[L]} \max\left\{\left|\mathrm{Re}\ \langle x, b_{i,\ell}\rangle\right|; \left|\mathrm{Im}\ \langle x, b_{i,\ell}\rangle\ \right|\right\}.\end{aligned} \tag{VI.21}$$

Setting

$$u_\ell = \Big(\left(\mathrm{Re}\ b_{i,\ell}\right)_1, -\left(\mathrm{Im}\ b_{i,\ell}\right)_1, \left(\mathrm{Re}\ b_{i,\ell}\right)_2, \ldots, -\left(\mathrm{Im}\ b_{i,\ell}\right)_{K_i-1}, \left(\mathrm{Re}\ b_{i,\ell}\right)_{K_i}, -\left(\mathrm{Im}\ b_{i,\ell}\right)_{K_i}\Big)$$

yields $\mathrm{Re}\left(\langle x, b_{i,\ell}\rangle_{\ell_2}\right) = \langle Px, u_\ell\rangle_{\ell_2}$ for all $x \in \mathbb{C}^{K_i}$ and all $\ell \in [L]$. Similarly, setting

$$v_\ell = \Big(\left(\mathrm{Im}\ b_{i,\ell}\right)_1, \left(\mathrm{Re}\ b_{i,\ell}\right)_1, \left(\mathrm{Im}\ b_{i,\ell}\right)_2, \ldots, \left(\mathrm{Re}\ b_{i,\ell}\right)_{K_i-1}, \left(\mathrm{Im}\ b_{i,\ell}\right)_{K_i}, \left(\mathrm{Re}\ b_{i,\ell}\right)_{K_i}\Big)$$

yields $\mathrm{Im}\left(\langle x, b_{i,\ell}\rangle\right) = \langle Px, v_\ell\rangle$ for all $x \in \mathbb{C}^{K_i}$ and all $\ell \in [L]$. We define

$$\mathcal{U} = \bigcup_{\ell\in[L]} \{u_\ell; v_\ell\}$$

and observe

$$\max_{u\in\mathcal{U}} \|u\|_{\ell_2} = \max_{\ell\in[L]} \|b_{i,\ell}\|_{\ell_2} \leq \sqrt{\frac{K_{i,\mu}}{L}}. \tag{VI.22}$$

By (VI.21) and the definition of $\mathcal{U}$ we obtain

$$\begin{aligned}\|x\|_{B_i} &\leq \sqrt{2L} \max_{u\in\mathcal{U}} \langle Px, u\rangle \\ &= \sqrt{2L} \max_{u\in \mathrm{conv}\ \mathcal{U}} \langle Px, u\rangle = \sqrt{2L}\|Px\|_{(\mathrm{conv}\ \mathcal{U})^\circ}.\end{aligned} \tag{VI.23}$$

(For the definition of $\|\cdot\|_{(\mathrm{conv}\ \mathcal{U})^\circ}$ see Section IV-C.) Inequality (VI.23) together with Theorem 15 yields

$$\begin{aligned}&\log N\left(B\left(0,1\right) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right) \\ &\leq \log N\left(B\left(0,1\right) \subset \mathbb{R}^{2K_i}, \|\cdot\|_{\mathrm{conv}(\mathcal{U})^\circ}, \frac{\varepsilon}{2\sqrt{2L}}\right) \\ &\lesssim \log N\left(\mathrm{conv}\left(\mathcal{U}\right), \|\cdot\|_{\ell_2}, \frac{\tilde{c}\varepsilon}{\sqrt{L}}\right),\end{aligned}$$

for some numerical constant $\tilde{c} > 0$, due to $\mathrm{conv}\left(\mathcal{U}\right)^{\circ\circ} = \mathrm{conv}\left(\mathcal{U}\right)$. In order to estimate this covering number from above we will use Lemma 14. For that purpose let $M \in \mathbb{N}$ and assume $(u_1, \ldots, u_M) \in \mathcal{U}^M$. By Jensen's inequality

$$\begin{aligned}\mathbb{E}\Big\|\sum_{m=1}^{M} \varepsilon_m u_m\Big\|_{\ell_2} &\leq \sqrt{\mathbb{E}\Big\|\sum_{m=1}^{M} \varepsilon_m u_m\Big\|_{\ell_2}^2} \\ &= \sqrt{\sum_{m=1}^{M} \|u_m\|_{\ell_2}^2} \leq \sqrt{M} \max_{u\in\mathcal{U}} \|u\|_{\ell_2}.\end{aligned}$$

Thus, by Lemma 14 applied with $A = \max_{u\in\mathcal{U}} \|u\|_{\ell_2}$ we obtain

$$\begin{aligned}\log N\left(\mathrm{conv}\left(\mathcal{U}\right), \|\cdot\|_{\ell_2}, \frac{\tilde{c}\varepsilon}{\sqrt{L}}\right) &\lesssim \frac{L}{\varepsilon^2} \max_{u\in\mathcal{U}} \|u\|_{\ell_2}^2 \log|\mathcal{U}| \\ &\lesssim \frac{K_{i,\mu}}{\varepsilon^2} \log L,\end{aligned}$$

where in the second inequality we have used (VI.22). This completes the proof. □

The previous two lemmas allow us to find an upper bound for the γ_2-functional, which is needed to prove Theorem 22.

Lemma 28: Suppose that $\mathcal{X} = \mathcal{W}$ or $\mathcal{X} = \mathcal{W}^p$ for some $p \in [P]$. (For the definition of $\mathcal{W}$ and $\mathcal{W}^p$ see (VI.18) and (VI.19).) Then

$$\begin{aligned}d_F\left(\mathcal{X}\right) &\leq 3, \\ d_B\left(\mathcal{X}\right) &\leq 3\sqrt{K_\mu}, \\ \gamma_2\left(\mathcal{X}, \|\cdot\|_B\right) &\lesssim \sqrt{r\left(K_\mu \log\left(L\right) \log^2\left(K_\mu\right) + N\mu_h^2\right)}.\end{aligned}$$

Proof: The first inequality follows from the triangle inequality. For the second one note that for $X \in \mathcal{X}$ by (VI.13) one obtains the inequality

$$\|X\|_B \leq \sqrt{K_\mu}\|X\|_F \leq 3\sqrt{K_\mu}.$$

The last line is more involved. We will present the proof only in the case of $\mathcal{X} = \mathcal{W}^p$. If $\mathcal{X} = \mathcal{W}$ the inequality can be

proven analogously. By Lemma 12 we obtain

$$\gamma_2\left(\mathcal{W}^p, \|\cdot\|_B\right) \lesssim \gamma_2\left(B^h, \|\cdot\|_B\right) + \gamma_2\left(B^{\mathcal{S}_p h}, \|\cdot\|_B\right) + \gamma_2\left(B^m, \|\cdot\|_B\right). \tag{VI.24}$$

We will estimate the three γ_2-functionals separately.

Step 1: To bound $\gamma_2\left(B^h, \|\cdot\|_B\right)$, let $U = \left(h_1 u_1^*, \ldots, h_r u_r^*\right), V = \left(h_1 v_1^*, \ldots, h_r v_r^*\right) \in B^h$. Observe that by definition

$$\begin{aligned}\|U - V\|_B &= \max_{\ell\in[L]} \sqrt{L \sum_{i=1}^r \left\| \left(h_i u_i^* - h_i v_i^*\right)^* b_{i,\ell} \right\|_{\ell_2}^2} \\ &= \max_{\ell\in[L]} \sqrt{L \sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2 |h_i^* b_{i,\ell}|^2} \\ &\le \mu_h \sqrt{\sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2} = \mu_h \|U - V\|_F,\end{aligned}$$

where the last equality is due to $\|h_i\|_{\ell_2} = 1$ for all $i \in [r]$. This implies

$$\begin{aligned}\gamma_2\left(B^h, \|\cdot\|_B\right) &\le \mu_h \gamma_2\left(B^h, \|\cdot\|_F\right) \\ &\lesssim \mu_h \int_0^1 \sqrt{\log N\left(B^h, \|\cdot\|_F, \varepsilon\right)} d\varepsilon \\ &\lesssim \mu_h \sqrt{rN}, \end{aligned} \tag{VI.25}$$

where the second inequality follows from the Dudley inequality (IV.3). The third inequality follows from the fact that $\left(B^h, \|\cdot\|_F\right)$ is isometric to $\left(B\left(0,1\right) \subset \mathbb{R}^{2\sum_{i=1}^r N_i}, \|\cdot\|_{\ell_2}\right)$ and from a standard volumetric estimate.

Step 2: To bound $\gamma_2\left(B^{\mathcal{S}_p h}, \|\cdot\|_B\right)$ let $U = \left(S_{1,p} h_1 u_1^*, \ldots, S_{r,p} h_r u_r^*\right)$ and $V = \left(S_{1,p} h_1 v_1^*, \ldots, S_{r,p} h_r v_r^*\right) \in B_h$. Then

$$\begin{aligned}\|U - V\|_B &= \max_{\ell\in[L]} \sqrt{L \sum_{i=1}^r \left\| \left(S_{i,p} h_i u_i^* - S_{i,p} h_i v_i^*\right)^* b_{i,\ell} \right\|_{\ell_2}^2} \\ &= \max_{\ell\in[L]} \sqrt{L \sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2 |h_i^* S_{i,p} b_{i,\ell}|^2} \\ &\le \mu_h \sqrt{\sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2} \\ &= \mu_h \sqrt{\sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2 \|h_i\|_{\ell_2}^2} \\ &= \mu_h \sqrt{\sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2 \|T_{i,p} S_{i,p} h_i\|_{\ell_2}^2} \\ &\le (1+\nu)\, \mu_h \sqrt{\sum_{i=1}^r \|u_i - v_i\|_{\ell_2}^2 \|S_{i,p} h_i\|_{\ell_2}^2} \\ &\lesssim \mu_h \|U - V\|_F.\end{aligned}$$

In the third line we used that $\|h_i\|_{\ell_2} = 1$ and in the last line we used that $\|T_{i,p}\|_{2\to 2} \le 1 + \nu$ and $\nu = \frac{1}{32}$. An analogous reasoning as in (VI.25) then yields

$$\gamma_2\left(B^{S_p h}, \|\cdot\|_B\right) \lesssim \mu_h \sqrt{rN}.$$

Step 3: To bound $\gamma_2\left(B^m, \|\cdot\|_B\right)$ note that inequality (IV.3) and the fact that $d_B\left(B^m\right) \le \sqrt{K_\mu}$ imply

$$\gamma_2\left(B^m, \|\cdot\|_B\right) \lesssim \int_0^{\sqrt{K_\mu}} \sqrt{\log N\left(B_m, \|\cdot\|_B, \varepsilon\right)} d\varepsilon.$$

Thus, by Lemma 26

$$\begin{aligned}&\gamma_2\left(B^m, \|\cdot\|_B\right) \\ &\lesssim \int_0^{\sqrt{K_\mu}} \sqrt{\log N\left(B\left(0,1\right) \subset \mathbb{R}^r, \|\cdot\|_{\ell_2}, \frac{\varepsilon}{2\sqrt{K_\mu}}\right)} d\varepsilon \\ &\quad + \int_0^{\sqrt{K_\mu}} \sqrt{\sum_{i=1}^r \log\left(N\left(B(0,1) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)\right)} d\varepsilon. \\ &\le \int_0^{\sqrt{K_\mu}} \sqrt{\log N\left(B\left(0,1\right) \subset \mathbb{R}^r, \|\cdot\|_{\ell_2}, \frac{\varepsilon}{2\sqrt{K_\mu}}\right)} d\varepsilon \\ &\quad + \sqrt{r} \int_0^{\sqrt{K_\mu}} \max_{i\in[r]} \sqrt{\log\left(N\left(B(0,1) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)\right)} d\varepsilon.\end{aligned} \tag{VI.26}$$

The first integral can be bounded by

$$\begin{aligned}&\int_0^{\sqrt{K_\mu}} \sqrt{\log N\left(B\left(0,1\right) \subset \mathbb{R}^r, \|\cdot\|_{\ell_2}, \frac{\varepsilon}{2\sqrt{K_\mu}}\right)} d\varepsilon \\ &\le \sqrt{r} \int_0^{\sqrt{K_\mu}} \sqrt{\log\left(1 + \frac{4\sqrt{K_\mu}}{\varepsilon}\right)} d\varepsilon \lesssim \sqrt{rK_\mu},\end{aligned} \tag{VI.27}$$

where we have used a volumetric estimate and a change of variables. In order to deal with the second term we will split the integrals into two parts: For small ε we will use a volumetric estimate and for large ε we will apply Lemma 27. First we consider the case that $\varepsilon \in (0,1)$. Therefore, note that

$$\begin{aligned}B\left(0,1\right) &\subset \sqrt{K_{i,\mu}} B_{\|\cdot\|_{B_i}}\left(0,1\right) \\ &:= \left\{x \in \mathbb{C}^{K_i} : \|x\|_{B_i} \le \sqrt{K_{i,\mu}}\right\}\end{aligned}$$

by inequality (VI.11). This fact combined with a volumetric estimate yields

$$\begin{aligned}&\max_{i\in[r]} N\left(B\left(0,1\right) \subset \mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \varepsilon\right) \\ &\le \max_{i\in[r]} N\left(B_{\|\cdot\|_{B_i}}\left(0,1\right), \|\cdot\|_{B_i}, \frac{\varepsilon}{\sqrt{K_{i,\mu}}}\right) \\ &\le \left(1 + \frac{2\sqrt{K_\mu}}{\varepsilon}\right)^{2K}.\end{aligned}$$

By a change of variables and an elementary integral inequality (see [65, Lemma C.9]) this implies

$$\begin{aligned}\int_0^1 \max_{i\in[r]} \sqrt{\log N\left(B\left(0,1\right), \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)}d\varepsilon &\leq \sqrt{2K}\int_0^1 \sqrt{\log\left(1+\frac{2\sqrt{K_\mu}}{\varepsilon}\right)}d\varepsilon \\ &\leq \sqrt{2K\log\left(e\left(1+2\sqrt{K_\mu}\right)\right)}.\end{aligned}$$

Next, we are going to deal with the case that $\varepsilon \in \left(1, \sqrt{K_\mu}\right)$. Using Lemma 27 we get

$$\begin{aligned}\int_1^{\sqrt{K_\mu}} \max_{i\in[r]} \sqrt{\log\left(N\left(B\left(0,1\right), \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)\right)}d\varepsilon &\lesssim \int_1^{\sqrt{K_\mu}} \frac{\sqrt{K_\mu \log L}}{\varepsilon}d\varepsilon \\ &\lesssim \sqrt{K_\mu \log L}\log\left(K_\mu\right).\end{aligned}$$

Summing up the two integral inequalities yields

$$\sqrt{r}\max_{i\in[r]}\int_0^{\sqrt{K_\mu}} \sqrt{\log\left(N\left(B(0,1)\subset\mathbb{C}^{K_i}, \|\cdot\|_{B_i}, \frac{\varepsilon}{2}\right)\right)}d\varepsilon \lesssim \sqrt{rK_\mu \log\left(L\right)}\log\left(K_\mu\right).$$

This inequality together with (VI.26) and (VI.27) shows that

$$\gamma_2\left(B^m, \|\cdot\|_B\right) \lesssim \sqrt{rK_\mu \log\left(L\right)\log^2\left(K_\mu\right)}.$$

The result then follows from inequality (VI.24). □

Combining the upper bounds for the γ_2-functionals in the last lemma with the abstract isometry result Proposition 25 we are able to prove the main result in this section.

Proof of Theorem 22: Fix $p \in [P]$. Using Lemma 28 and choosing the constant C_ω in (VI.8) large enough we get for the quantities arising in Proposition 25 that $\widehat{E} \leq \frac{\delta}{2\tilde{c}_1}$, $\widehat{V} \leq \frac{\delta}{\sqrt{\tilde{c}_2\omega\log L}}$, and $\widehat{U} \leq \frac{\delta}{\tilde{c}_2\omega\log L}$, where we have set $\mathcal{X} = \mathcal{W}^p$ (see (VI.19)) and $\tilde{c}_i$ are the constants appearing in Proposition 25. Thus inequality (VI.14) of Proposition 25 for $t = \frac{\delta}{2}$ shows that (VI.10) holds with probability $1 - \mathcal{O}\left(L^{-\omega}\right)$ for fixed p.

In order to prove (VI.9) we may argue analogously (with $\mathcal{X} = \mathcal{W}$ and $t = \frac{\delta}{2}$) and apply inequality (VI.15) of Proposition 25. Thus, (VI.10) holds with probability at least $1 - \mathcal{O}\left(L^{-\omega}\right)$. Replacing ω by $\omega + 1$ in the argument above and using a union bound argument one observes that (VI.10) and (VI.9) are satisfied for all $p \in [P]$ with probability at least $1 - (P+1)\mathcal{O}\left(L^{-\omega-1}\right) = 1 - \mathcal{O}\left(L^{-\omega}\right)$, which finishes the proof. □

C. Constructing the Dual Certificate

1) The Golfing Scheme: The goal of this section is to construct $Y \in \text{Range}\left(\mathcal{A}^*\right)$ such that the conditions (VI.1) and (VI.2) in Lemma 18 are fulfilled with high probability. The construction itself will make use of the Golfing Scheme, an iterative method which has been introduced in [30] for the first time. We set

$$\begin{aligned}Y_0 &= 0 \\ Y_p &= Y_{p-1} + \frac{L}{Q}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\left(\text{sgn}\left(X^0\right) - \mathcal{P}_\mathcal{T}\left(Y_{p-1}\right)\right)\end{aligned}$$

for $p \in [P]$. We will make use of the notation

$$W_p = \text{sgn}\left(X^0\right) - \mathcal{P}_\mathcal{T}\left(Y_p\right) \quad \text{for } 0 \leq p \leq P. \tag{VI.28}$$

The individual components of W_p will be denoted by $W_{i,p}$ for $i \in [r]$, i.e., $W_p = \left(W_{1,p}, \ldots, W_{r,p}\right)$. Then the dual certificate will be given by

$$Y = Y_P = \sum_{p=1}^P \frac{L}{Q}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right).$$

Our Golfing Scheme is set up in the same way as in [43]. In particular, they also use the operator $\mathcal{S}^p$ as a corrector function as explained in the following remark.

Remark 29: The reason for the appearance of the operator $\mathcal{S}^p$ is the following: Observe that

$$\mathbb{E}\left[\left(\mathcal{A}^p\right)^*\mathcal{A}^p\left(X\right)\right] = \frac{L}{Q}\left(T_{i,p}X_1, \ldots, T_{r,p}X_r\right).$$

Recall that $T_{i,p}$ may only be approximately equal to the identity matrix (see (II.5)). Thus, $\left(\mathcal{A}^p\right)^*\mathcal{A}^p$ is not necessarily an unbiased estimator. However,

$$\begin{aligned}&\mathbb{E}\left[\frac{L}{Q}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\left(X\right)\right] \\ &\quad= \frac{L}{Q}\left(T_{1,p}S_{1,p}X_1, \ldots, T_{r,p}S_{r,p}X_r\right) \\ &\quad= \left(X_1, \ldots, X_r\right) = X.\end{aligned}$$

Thus, we get that $\mathbb{E}\left[\frac{L}{Q}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\right] = Id$. Note that $\mathcal{S}^p\left(W_{p-1}\right)$ is, in general, not an element of the subspace $\mathcal{T}$. However, due to definition of $\mathcal{T}^p$ we observe that $\mathcal{S}^p\left(W_{p-1}\right) \in \mathcal{T}^p$. This is the reason why we require the operator $\mathcal{A}^p$ to satisfy the δ-local isometry property not only on $\mathcal{T}$, but also on $\mathcal{T}^p$.

Let us check that $Y \in \text{Range}\left(\mathcal{A}^*\right)$: Recall that the $\mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right)$ is obtained by setting the vector $\mathcal{A}\mathcal{S}^p\left(W_{p-1}\right)$ zero in those components, which do not belong to Γ_p (see Section II-C). In particular, this implies that $\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right) = \mathcal{A}^*\mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right)$. Thus, setting

$$z = \sum_{p=1}^P \mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right). \tag{VI.29}$$

we get that $Y = \mathcal{A}^*z$. The vector z will also be important when we prove an upper bound for the estimation error in the presence of noise. In the remaining part of the proof we will verify that Y satisfies the conditions in Lemma 18 with the constants $\alpha = \frac{1}{8\gamma}$, $\beta = \frac{1}{4}$, and $\delta = \frac{1}{4}$ (cf. Remark 19).

2) Exponential Decay: In this section we will verify condition (VI.1) in Lemma 18. In other words, we have to show that the quantity

$$\|W_P\|_F = \|\mathrm{sgn}\left(X^0\right) - \mathcal{P}_{\mathcal{T}}(Y)\|_F$$

is small enough. An important observation, which we will need in the proof, is that $W_0 = \mathrm{sgn}\left(X^0\right)$ and one has the recurrence relation

$$W_p = W_{p-1} - \frac{L}{Q}\left(\mathcal{P}_{\mathcal{T}}\left(\mathcal{A}^p\right)^* \mathcal{A}^p \mathcal{S}^p\right)\left(W_{p-1}\right) \quad \text{for all } p \in [P], \tag{VI.30}$$

which is a direct of consequence of the definition of W_p (see equation (VI.28)). In Lemma 31, we will prove that W_p decays exponentially fast. We will need the following rather technical inequalities.

Lemma 30: Let $\nu = \frac{1}{32}$. For all $i \in [r]$ and for all $p \in [P]$ we have the inequalities

$$\left\|Id - T_{i,p}^{1/2}\right\|_{2\to 2} \le \frac{1}{32} \tag{VI.31}$$

$$\left\|\left(Id - \mathcal{S}^p\right)X\right\|_F \le \frac{1}{31}\|X\|_F \tag{VI.32}$$

$$\left\|\mathcal{S}^p X\right\|_F \le \frac{32}{31}\|X\|_F. \tag{VI.33}$$

Proof: Inequality (VI.31) follows directly from (II.5) and the observation that the square-root shifts the eigenvalues of $T_{i,p}$ closer to one. The inequalities (VI.32) and (VI.33) follow from the observation that for all $i \in [r]$, $p \in [P]$

$$\begin{aligned}&\|\mathrm{Id} - S_{i,p}\|_{2\to 2}\\ &\quad = \max\left\{1 - \sigma_{\min}\left(S_{i,p}\right); \sigma_{\max}\left(S_{i,p}\right) - 1\right\}\\ &\quad = \max\left\{1 - \sigma_{\max}^{-1}\left(T_{i,p}^{-1}\right); \sigma_{\min}^{-1}\left(T_{i,p}^{-1}\right) - 1\right\} \le \frac{1}{31}.\end{aligned}$$

□

This allows us to prove the main lemma in this section.

Lemma 31: Suppose that $\mathcal{A}^p$ satisfies the δ-local isometry property on $\mathcal{T}^p$ with $\delta = \frac{1}{32}$ for all $p \in [P]$. Then, for all $p \in [P]$,

$$\|W_p\|_F \le 4^{-p}\sqrt{r} \tag{VI.34}$$

and, in particular, if $P \ge \frac{1}{2}\log\left(8\gamma\sqrt{r}\right)$,

$$\|sgn\left(X^0\right) - Y\|_F \le \frac{1}{8\gamma}. \tag{VI.35}$$

Proof: First notice that by (VI.31) and the triangle inequality

$$(1-\nu)\|X_i\|_F \le \left\|T_{i,p}^{1/2}X_i\right\|_F \le (1+\nu)\|X_i\|_F$$

for all $X_i \in \mathbb{C}^{K_i \times N_i}$. Thus, by the local isometry property (VI.10)

$$\begin{aligned}(1-\nu)^2(1-\delta)\|X\|_F^2 &\le \frac{L}{Q}\left\|\mathcal{A}^p(X)\right\|_{\ell_2}^2\\ &\le (1+\delta)(1+\nu)^2\|X\|_F^2\end{aligned}$$

for all $X \in \mathcal{T}^p$. Together with $\delta = \nu = \frac{1}{32}$ this implies

$$\left|\frac{L}{Q}\|\mathcal{A}^p(X)\|_{\ell_2}^2 - \|X\|_F^2\right| \le \frac{1}{8}\|X\|_F^2$$

for all $X \in \mathcal{T}^p$, which in turn is equivalent to

$$\left\|\mathcal{P}_{\mathcal{T}^p} - \frac{L}{Q}\mathcal{P}_{\mathcal{T}^p}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{P}_{\mathcal{T}^p}\right\|_{F\to F} \le \frac{1}{8}, \tag{VI.36}$$

where $\mathcal{P}_{\mathcal{T}^p}$ denotes the orthogonal projection onto $\mathcal{T}^p$. Now note that $\|W_{p-1} - \mathcal{P}_{\mathcal{T}}(X)\|_F \le \|W_{p-1} - \mathcal{P}_{\mathcal{T}^p}(X)\|_F$ for all $X \in \mathcal{M}$ due to $W_{p-1} \in \mathcal{T}$ and $\mathcal{T} \subset \mathcal{T}^p$. This fact together with (VI.30) implies that

$$\begin{aligned}\|W_p\|_F &\le \left\|W_{p-1} - \left(\frac{L}{Q}\mathcal{P}_{\mathcal{T}^p}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)\left(W_{p-1}\right)\right\|_F\\ &= \left\|W_{p-1} - \left(\frac{L}{Q}\mathcal{P}_{\mathcal{T}^p}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{P}_{\mathcal{T}^p}\mathcal{S}^p\right)\left(W_{p-1}\right)\right\|_F,\end{aligned}$$

where in the second line we use that $\mathcal{S}^p W_{p-1} \in \mathcal{T}^p$ by the definition of $\mathcal{T}^p$ (see (V.3)) and because of $W_{p-1} \in \mathcal{T}$. Using this computation and (VI.32), (VI.33), (VI.36) we obtain

$$\begin{aligned}\|W_p\|_F &\le \left\|\left(\mathrm{Id} - \frac{L}{Q}\mathcal{P}_{\mathcal{T}^p}\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{P}_{\mathcal{T}^p}\right)\left(\mathcal{S}^p W_{p-1}\right)\right\|_F\\ &\quad + \left\|\left(\mathrm{Id} - \mathcal{S}^p\right)W_{p-1}\right\|_F\\ &\le \frac{1}{8}\|\mathcal{S}^p W_{p-1}\|_F + \frac{1}{16}\|W_{p-1}\|_F \le \frac{1}{4}\|W_{p-1}\|_F.\end{aligned}$$

Thus, the previous estimate yields

$$\|W_p\|_F \le \left(\frac{1}{4}\right)^p \|W_0\|_F = \left(\frac{1}{4}\right)^p \sqrt{r}.$$

This shows ($VI.34$) and, in particular, we obtain $\|W_P\|_F \le 4^{-P}\sqrt{r}$. The assumption $P \ge \frac{1}{2}\log\left(8\gamma\sqrt{r}\right)$ and the definition of W_P imply (VI.35), which finishes the proof. □

3) Bounding the Operator Norm on $\mathcal{T}^\perp$: To apply Lemma 18 we need in addition to controlling the share of Y in $\mathcal{T}$ also a bound on $\mathcal{T}_i^\perp$ for all $i \in [r]$. For that, recall from [43] that

$$\begin{aligned}&\left\|\mathcal{P}_{\mathcal{T}_i^\perp}\left(Y_i^P\right)\right\|_{2\to 2}\\ &\le \sum_{p=1}^P \left\|\mathcal{P}_{\mathcal{T}_i^\perp}\left(\frac{L}{Q}\left(\left(\mathcal{A}^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)\left(W_{p-1}\right) - W_{i,p-1}\right)\right\|_{2\to 2}\\ &\le \sum_{p=1}^P \left\|\frac{L}{Q}\left(\left(\mathcal{A}_i^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)\left(W_{p-1}\right) - W_{i,p-1}\right\|_{2\to 2}\\ &= \sum_{p=1}^P \|W_{i,p}\|_{2\to 2},\end{aligned}$$

where one uses the fact that $W_{i,p-1} \in \mathcal{T}_i$. Thus to establish the bound $\left\|\mathcal{P}_{\mathcal{T}_i^\perp}\left(Y_i^P\right)\right\|_{2\to 2} < \frac{1}{4}$ it remains to show that

$$\left\|\frac{L}{Q}\left(\left(\mathcal{A}_i^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)\left(W_{p-1}\right) - W_{i,p-1}\right\|_{2\to 2} \le \frac{1}{4^{p+1}}.$$

To proceed, set for $p \in \{0; 1; \ldots; P-1\}$

$$\mu_p = \sqrt{L}\max_{\ell \in \Gamma_{p+1}, k \in [r]}\left\|W_{k,p}^* S_{k,p+1} b_{k,\ell}\right\|_{2\to 2}. \tag{VI.37}$$

This allows us to state the following lemma.

Lemma 32: Fix $i \in [r]$ and let $\omega \ge 1$. Assume that

$$\mu_p \le 4^{-p}\mu_h \text{ and } \|W_p\|_F \le 4^{-p}\sqrt{r}. \tag{VI.38}$$

If

$$Q \gtrsim_\omega r\left(K_\mu + N\mu_h^2\right)(\log L)^2, \tag{VI.39}$$

then with probability $1-\mathcal{O}\left(L^{-\omega}\right)$ *the inequality*

$$\left\| \frac{L}{Q}\left(\mathcal{A}_i^p\right)^* \mathcal{A}^p\mathcal{S}^p W_{p-1} - W_{i,p-1}\right\|_{2\to 2} \le \frac{1}{4^{p+1}} \tag{VI.40}$$

is true for all $p\in[P]$ *and for all* $i\in[r]$.

Remark 33: The validity of assumption (VI.38) is assured by Lemma 31 and Lemma 34 below.

Proof: The proof follows the same strategy as [43, Lemma 5.12]. Fix $p\in[P]$ and $i\in[r]$. First, we will decompose $W_{i,p}$ as a sum of independent random matrices such that the matrix Bernstein inequality can be applied. For that purpose, observe that for all $y\in\mathbb{C}^L$ and for all $\ell\in\Gamma_p$ by definition of $\mathcal{S}^p$ (Definition 17) and $\mathcal{A}^p$

$$\left(\mathcal{A}^p\mathcal{S}^p W_{p-1}\right)(\ell) = \sum_{k=1}^r b_{k,\ell}^* S_{k,p} W_{k,p-1} c_{k,\ell}.$$

(For $\ell\in[L]\backslash\Gamma_p$ the left-hand side is equal to zero as $\mathcal{A}^p(X) = P_{\Gamma_p}(\mathcal{A}(X))$.) Using (II.2) one obtains

$$\left(\left(\mathcal{A}_i^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)W_{p-1} = \sum_{\ell\in\Gamma_p}\sum_{k=1}^r b_{i,\ell}b_{k,l}^* S_{k,p}W_{k,p-1}c_{k,\ell}c_{i,\ell}^*.$$

With $S_{i,p} = T_{i,p}^{-1}$ and the definition of $T_{i,p}$ (see equation (II.5)) this implies

$$W_{i,p-1} = T_{i,p}S_{i,p}W_{i,p-1} = \frac{L}{Q}\sum_{\ell\in\Gamma_p} b_{i,\ell}b_{i,\ell}^* S_{i,p}W_{i,p-1}.$$

In order to simplify notation we introduce the vectors $w_{k,\ell}$ defined by

$$w_{k,\ell} = W_{k,p-1}^* S_{k,p} b_{k,\ell}. \tag{VI.41}$$

Using this definition we may write (as $S_{k,p}$ is self-adjoint)

$$\begin{aligned} W_{i,p} &= \frac{L}{Q}\left(\left(\mathcal{A}_i^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)W_{p-1} - W_{i,p-1} && \text{(VI.42)}\\ &= \frac{L}{Q}\sum_{\ell\in\Gamma_p}\sum_{k=1}^r b_{i,\ell}w_{k,\ell}^* c_{k,\ell}c_{i,\ell}^* - \frac{L}{Q}\sum_{\ell\in\Gamma_p} b_{i,\ell}w_{i,\ell}^* && \text{(VI.43)}\\ &= \frac{L}{Q}\sum_{\ell\in\Gamma_p} b_{i,\ell}w_{i,\ell}^*\left(c_{i,\ell}c_{i,\ell}^* - \mathrm{Id}\right)\\ &\quad + \frac{L}{Q}\sum_{\ell\in\Gamma_p}\sum_{k\neq i} b_{i,\ell}w_{k,\ell}^* c_{k,l}c_{i,\ell}^*\\ &= \sum_{\ell\in\Gamma_p} Z_\ell, && \text{(VI.44)}\end{aligned}$$

where we have set

$$Z_\ell = \frac{L}{Q}\left(\sum_{k=1}^L b_{i,\ell}w_{k,\ell}^*\left(c_{k,\ell}c_{i,\ell}^* - \mathbb{E}\left[c_{k,\ell}c_{i,\ell}^*\right]\right)\right).$$

Note that until the last step of the proof i is assumed to be fixed which is why we refrain from indicating the i-dependence in every step for reasons of notational simplicity. Observe that each summand of Z_ℓ and hence the the cross terms in $Z_\ell Z_\ell^*$ and $Z_\ell^* Z_\ell$ have expectation zero. Thus using basic properties of circular symmetric normal random variables, Lemma 37 and Lemma 38 we compute

$$\mathbb{E}\left[Z_\ell Z_{\ell'}^*\right] = \frac{L^2}{Q^2}\sum_{k=1}^r N_k\left\|w_{k,\ell}\right\|_{\ell_2}^2 b_{i,\ell}b_{i,\ell'}^*. \tag{VI.45}$$

$$\mathbb{E}\left[Z_\ell^* Z_\ell\right] = \frac{L^2}{Q^2}\|b_{i,\ell}\|_{\ell_2}^2\sum_{k=1}^r\left\|w_{k,\ell}\right\|_{\ell_2}^2\mathrm{Id}. \tag{VI.46}$$

We have to find an upper bound for the spectral norms of these quantities. First, observe that

$$\begin{aligned}&\left\|\sum_{\ell\in\Gamma_p}\mathbb{E}\left[Z_\ell Z_\ell^*\right]\right\|_{2\to 2}\\ &\quad\le \frac{L^2N}{Q^2}\left(\max_{k\in[r],\ell\in\Gamma_p}\|w_{k,\ell}\|_2^2\right)\left\|\sum_{k=1}^r\sum_{\ell\in\Gamma_p} b_{i,\ell}b_{i,\ell}^*\right\|_{2\to 2}\\ &\quad\le \frac{rN}{Q}\mu_{p-1}^2\|T_{i,p}\|_{2\to 2} \overset{\text{(VI.38)}}{\lesssim} \frac{16^{-p+1}rN\mu_h^2}{Q}.\end{aligned}$$

By a similar computation we obtain

$$\begin{aligned}&\left\|\sum_{\ell\in\Gamma_p}\mathbb{E}\left[Z_\ell^* Z_\ell\right]\right\|_{2\to 2}\\ &\quad\le \frac{L^2}{Q^2}\left(\max_{\ell\in\Gamma_p}\|b_{i,\ell}\|_{\ell_2}^2\right)\sum_{k=1}^r\sum_{\ell\in\Gamma_p}\|w_{k,\ell}\|_{\ell^2}^2\\ &\quad\lesssim \frac{LK_{i,\mu}}{Q^2}\sum_{k=1}^r\sum_{\ell\in\Gamma_p}\mathrm{Tr}\left(W_{k,p-1}^* S_{k,p}b_{k,\ell}b_{k,\ell}^* S_{k,p}W_{k,p-1}\right)\\ &\quad= \frac{K_{i,\mu}}{Q}\sum_{k=1}^r\|S_{k,p}^{1/2}W_{k,p-1}\|_F^2 \lesssim \frac{K_{i,\mu}}{Q}\|W_{p-1}\|_F^2\\ &\quad\le 16^{-p+1}\frac{rK_{i,\mu}}{Q}.\end{aligned}$$

Thus, we have obtained

$$\begin{aligned}\sigma^2 &:= \max\left\{\left\|\sum_{\ell\in\Gamma_p}\mathbb{E}\left[Z_\ell^* Z_\ell\right]\right\|_{2\to 2}, \left\|\sum_{\ell\in\Gamma_p}\mathbb{E}\left[Z_\ell Z_\ell^*\right]\right\|_{2\to 2}\right\}\\ &\lesssim 16^{-p}\frac{r}{Q}\max\left\{K_{i,\mu}, N\mu_h^2\right\}. && \text{(VI.47)}\end{aligned}$$

Observe that a lower bound for σ^2 is given by

$$\sigma^2 \ge \left\|\sum_{\ell\in\Gamma_p}\mathbb{E}\left[Z_\ell^* Z_\ell\right]\right\|_{2\to 2} = \frac{L^2}{Q^2}\sum_{k=1}^r\sum_{\ell\in\Gamma_p}\|b_{i,\ell}\|_{\ell_2}^2\|w_{k,\ell}\|_{\ell_2}^2. \tag{VI.48}$$

Next we have to estimate $R = \max_{\ell\in\Gamma_p}\left\|\|Z_\ell\|_{2\to 2}\right\|_{\psi_1}$. By Lemma 39 and inequality (IV.1) we have that

$$\begin{aligned}\left\|\|Z_\ell\|_{2\to 2}\right\|_{\psi_1} &\le \frac{L}{Q}\Big(\sum_{k\neq i}\|b_{i,\ell}\|_{\ell_2}\left\||w_{k,\ell}^* c_{k,\ell}|\|c_{i,\ell}\|_{\ell_2}\right\|_{\psi_1}\\ &\quad + \|b_{i,\ell}\|_{\ell_2}\left\|\left\|\left(c_{i,\ell}c_{i,\ell}^* - \mathrm{Id}\right)w_{i,\ell}\right\|_{\ell_2}\right\|_{\psi_1}\Big)\\ &\lesssim \frac{L\sqrt{N_i}}{Q}\|b_{i,\ell}\|_{\ell_2}\sum_{k=1}^r\|w_{k,\ell}\|_{\ell_2} && \text{(VI.49)}\end{aligned}$$

$$\lesssim \frac{r\sqrt{K_{i,\mu}N_i}\mu_{p-1}}{Q} \lesssim 4^{-p}\frac{r\sqrt{K_{i,\mu}N_i}\mu_h}{Q} \lesssim 4^{-p}\frac{r\left(K_{i,\mu}+N_i\mu_h^2\right)}{Q} \tag{VI.50}$$

and, consequently, $R \lesssim 4^{-p}\frac{r\left(K_{i,\mu}+N_i\mu_h^2\right)}{Q}$. Moreover, combining (VI.48) and (VI.49) we obtain

$$\frac{|\Gamma_p|R^2}{\sigma^2} \lesssim QN\frac{\max_{\ell\in\Gamma_p}\left(\sum_{k=1}^r \|b_{i,\ell}\|_{\ell_2}\|w_{k,\ell}\|_{\ell_2}\right)^2}{\max_{\ell\in\Gamma_p}\left(\sum_{k=1}^r \|b_{i,\ell}\|_{\ell_2}^2\|w_{k,\ell}\|_{\ell_2}^2\right)} \leq QNr. \tag{VI.51}$$

As $Q \leq L$ by definition (VI.39) implies that $\log\left(1+\frac{|\Gamma_p|R^2}{\sigma^2}\right) \lesssim \log L$. Thus, setting $t=(\omega+2)\log L$ we obtain from Theorem 9 applied with $\alpha=1$ and combined with (VI.47) that with probability $1-\mathcal{O}\left(L^{-\omega-2}\right)$

$$\Big\|\sum_{\ell\in\Gamma_p} Z_\ell\Big\|_{2\to 2} \lesssim_\omega 4^{-p}\max\left\{\sqrt{\frac{r\left(K_{i,\mu}+N\mu_h^2\right)}{Q}\log L}, \frac{r\left(K_{i,\mu}+N\mu_h^2\right)}{Q}(\log L)^2\right\}.$$

Thus, by choosing the constant in (VI.39) large enough it holds that $\left\|\sum_{\ell\in\Gamma_p} Z_\ell\right\|_{2\to 2} \leq 4^{-p-1}$ with probability $1-\mathcal{O}\left(L^{-\omega-2}\right)$ for fixed $p\in[P]$ and for fixed $i\in[r]$. By taking the union bound over all $i\in[r]$ and over all $p\in[P]$ we obtain that with probability $1-rP\mathcal{O}\left(L^{-\omega-2}\right)=1-\mathcal{O}\left(L^{-\omega}\right)$ equation (VI.40) is true for all $p\in[P]$ and for all $i\in[r]$. This finishes the proof. □

4) Proof That $\mu_p \leq \frac{1}{4}\mu_{p-1}$: Lemma 32 additionaly required that $\mu_p \leq \frac{1}{4}\mu_{p-1}$ for all $p\in[P-1]$. In this section we will verify that this property holds with high probability.

Lemma 34: Let $\omega \geq 1$. If

$$Q \gtrsim_\omega r\left(K_\mu + N\mu_h^2\right)\log^2 L\ , \tag{VI.52}$$

then with probability at least $1-\mathcal{O}\left(L^{-\omega}\right)$ it holds that $\mu_p \leq \frac{1}{4}\mu_{p-1}$ for all $p\in[P-1]$.

A similar lemma was established in [43]. However, it was required that L scales quadratically with r. Thus, we need to refine the argument in order to achieve a linear scaling in r.

Proof of Lemma 34: First, we will show the claim for fixed $p\in\{0;1;\ldots;P-1\}$. Observe that it is enough to show that for all $\ell\in\Gamma_{p+1}$ and all $i\in[r]$

$$\sqrt{L}\|w_{i,\ell}\|_{\ell_2} \leq \frac{1}{4}\mu_{p-1} \tag{VI.53}$$

with $w_{i,\ell} := W_{i,p}S_{i,p+1}b_{i,\ell}$ as in (VI.41). Furthermore, observe that from the recurrence relation (VI.30) we obtain

$$W_{i,p} = W_{i,p-1} - \frac{L}{Q}\left(\mathcal{P}_{\mathcal{T}_i}\left(\mathcal{A}_i^p\right)^*\mathcal{A}^p\mathcal{S}^p\right)\left(W_{p-1}\right).$$

Due to the definition of $\mathcal{T}_i$ and $\|h_i\|_{\ell_2}=\|m_i\|_{\ell_2}=1$ we may write for all $Z\in\mathbb{C}^{K_i\times N_i}$

$$\mathcal{P}_{\mathcal{T}_i}Z = h_ih_i^*Z + \left(\mathrm{Id}-h_ih_i^*\right)Zm_im_i^*.$$

Together with (VI.42, VI.44) this implies

$$\begin{aligned}W_{i,p} &= \frac{L}{Q}\sum_{j\in\Gamma_p}\Big[h_ih_i^*b_{i,j}w_{i,j}^*\left(\mathrm{Id}-c_{i,j}c_{i,j}^*\right)\\ &\quad+\left(\mathrm{Id}-h_ih_i^*\right)b_{i,j}w_{i,j}^*\left(\mathrm{Id}-c_{i,j}c_{i,j}^*\right)m_im_i^*\Big]\\ &\quad-\frac{L}{Q}\sum_{k\neq i}\sum_{j\in\Gamma_p}\Big[h_i^*h_ib_{i,j}w_{k,j}^*c_{k,j}c_{i,j}^*\\ &\quad+(\mathrm{Id}-h_ih_i^*)b_{i,j}w_{k,j}^*c_{k,j}c_{i,j}^*m_im_i^*\Big].\end{aligned}$$

We define for all $j\in\Gamma_p$

$$\begin{aligned}\mathbf{z}_{i,j} &= \frac{L}{Q}\left(\mathrm{Id}-c_{i,j}c_{i,j}^*\right)w_{i,j}b_{i,j}^*h_ih_i^*S_{i,p+1}b_{i,\ell},\\ z_{i,j} &= \frac{L}{Q}m_i^*\left(\mathrm{Id}-c_{i,j}c_{i,j}^*\right)w_{i,j}b_{i,j}^*\left(\mathrm{Id}-h_ih_i^*\right)S_{i,p+1}b_{i,\ell}\end{aligned}$$

and for all $k\neq i$ and for all $j\in\Gamma_p$

$$\begin{aligned}\mathbf{z}_{k,j} &= \frac{L}{Q}c_{i,j}c_{k,j}^*w_{k,j}b_{i,j}^*h_ih_i^*S_{i,p+1}b_{i,\ell},\\ z_{k,j} &= \frac{L}{Q}m_i^*c_{i,j}c_{k,j}^*w_{k,j}b_{i,j}^*(\mathrm{Id}-h_ih_i^*)S_{i,p+1}b_{i,\ell}.\end{aligned}$$

Hence, to establish (VI.53) by the triangle inequality it is sufficient to prove that with high probability

$$\Big\|\sum_{j\in\Gamma_p}\mathbf{z}_{i,j}\Big\|_{\ell_2} \leq \frac{1}{16\sqrt{L}}\mu_{p-1}, \tag{VI.54}$$

$$\Big|\sum_{j\in\Gamma_p} z_{i,j}\Big| \leq \frac{1}{16\sqrt{L}}\mu_{p-1}, \tag{VI.55}$$

$$\Big\|\sum_{k\neq i}\sum_{j\in\Gamma_p}\mathbf{z}_{k,j}\Big\|_{\ell_2} \leq \frac{1}{16\sqrt{L}}\mu_{p-1}, \tag{VI.56}$$

$$\Big|\sum_{k\neq i}\sum_{j\in\Gamma_p} z_{k,j}\Big| \leq \frac{1}{16\sqrt{L}}\mu_{p-1}. \tag{VI.57}$$

Step 1 (Proof of (VI.54)): In order to apply Theorem 9 we compute using Lemma 38

$$\begin{aligned}&\Big\|\mathbb{E}\Big[\sum_{j\in\Gamma_p}\mathbf{z}_{i,j}\mathbf{z}_{i,j}^*\Big]\Big\|_{2\to 2}\\ &= \frac{L^2}{Q^2}|h_i^*S_{i,p+1}b_{i,\ell}|^2\sum_{j\in\Gamma_p}|b_{i,j}^*h_i|^2\|w_{i,j}\|_{\ell_2}^2\\ &\leq \frac{1}{QL}\mu_h^2\mu_{p-1}^2\|T_{i,p}^{1/2}h_i\|_{\ell_2}^2 \lesssim \frac{1}{QL}\mu_h^2\mu_{p-1}^2.\end{aligned}$$

Analogously, using Lemma 37

$$\begin{aligned}&\mathbb{E}\Big[\sum_{j\in\Gamma_p}\mathbf{z}_{i,j}^*\mathbf{z}_{i,j}\Big]\\ &= \frac{L^2N_i}{Q^2}\sum_{j\in\Gamma_p}\|w_{i,j}\|_{\ell_2}^2|b_{i,j}^*h_i|^2|b_{i,\ell}^*S_{i,p+1}^*h_i|^2\\ &\lesssim \frac{N_i}{QL}\mu_{p-1}^2\mu_h^2.\end{aligned}$$

Next, we estimate $R = \max_{j\in\Gamma_p} \left\| \|\mathbf{z}_{i,j}\|_{\ell_2} \right\|_{\psi_1}$. For that purpose we apply Lemma 39 to observe that

$$\begin{aligned} R &= \max_{j\in\Gamma_p} \left\| \|\mathbf{z}_{i,j}\|_{\ell_2} \right\|_{\psi_1} && \text{(VI.58)}\\ &= \frac{L}{Q} \max_{j\in\Gamma_p} \left(|b^*_{i,j}h_i| |h_i^* S_{i,p+1} b_{i,\ell}| \| \left(\mathrm{Id} - c_{i,j} c^*_{i,j} \right) w_{i,j} \|_{\psi_1} \right) \\ &\lesssim \frac{L\sqrt{N_i}}{Q} \max_{j\in\Gamma_p} \left(|h_i^* S_{i,p+1} b_{i,\ell}| |b^*_{i,j} h_i| \|w_{i,j}\|_{\ell_2} \right) && \text{(VI.59)}\\ &\lesssim \frac{\sqrt{N_i}\mu_h^2}{Q\sqrt{L}} \mu_{p-1}. \end{aligned}$$

Furthermore, (VI.59) yields, analogously to the derivation of (VI.51), that

$$\begin{aligned} \frac{|\Gamma_p| R^2}{\sigma^2} &\le |\Gamma_p| \frac{\max_{j\in\Gamma_p} |h_i^* S_{i,p+1} b_{i,\ell}|^2 |b^*_{i,j} h_i|^2 \|w_{i,j}\|^2_{\ell_2}}{\sum_{j\in\Gamma_p} \|w_{i,j}\|^2_{\ell_2} |b^*_{i,j} h_i|^2 |b^*_{i,\ell} S^*_{i,p+1} h_i|^2} \\ &\lesssim Q \le L. && \text{(VI.60)} \end{aligned}$$

Applying Theorem 9 with $t = (\omega+2)\log L$ and $\alpha = 1$ we obtain that with probability $1 - \mathcal{O}\left(L^{-\omega-2}\right)$

$$\left\| \sum_{j\in\Gamma_p} \mathbf{z}_{i,j} \right\|_{\ell_2} \lesssim_\omega \frac{\mu_{p-1}}{\sqrt{L}} \max\left\{ \sqrt{\frac{N_i\mu_h^2}{Q}\log L}; \frac{\sqrt{N_i}\mu_h^2}{Q} (\log L)^2 \right\},$$

which implies (VI.54), if the numerical constant in (VI.52) is chosen large enough.

Step 2 (Proof of (VI.55)): By Lemma 39 we obtain that

$$\begin{aligned} &\left\| |z_{i,j}| \right\|_{\psi_1} \\ &\lesssim \frac{L}{Q} |b^*_{i,j} \left(\mathrm{Id} - h_i h_i^* \right) S_{i,p+1} b_{i,\ell}| \|w_{i,j}\|_{\ell_2} \\ &\le \frac{L}{Q} \|b_{i,j}\|_{\ell_2} \|\mathrm{Id} - h_i h_i^*\|_{2\to 2} \|S_{i,p+1}\|_{2\to 2} \|b_{i,\ell}\|_{\ell_2} \|w_{i,j}\|_{\ell_2} \\ &\lesssim \frac{L}{Q} \|b_{i,j}\|_{\ell_2} \|b_{i,\ell}\|_{\ell_2} \|w_{i,j}\|_{\ell_2} \lesssim \frac{K_{i,\mu}}{Q\sqrt{L}} \mu_{p-1} \end{aligned}$$

and

$$\begin{aligned} &\sum_{j\in\Gamma_p} \left\| |z_{i,j}| \right\|^2_{\psi_1} \\ &\lesssim \frac{L^2}{Q^2} \left(\max_{j\in\Gamma_p} \|w_{i,j}\|^2_{\ell_2} \right) \sum_{j\in\Gamma_p} |b^*_{i,j} \left(\mathrm{Id} - h_i h_i^* \right) S_{i,p+1} b_{i,\ell}|^2 \\ &= \frac{L}{Q} \left(\max_{j\in\Gamma_p} \|w_{i,j}\|^2_{\ell_2} \right) \|T^{\frac{1}{2}}_{i,p} \left(\mathrm{Id} - h_i h_i^* \right) S_{i,p+1} b_{i,\ell}\|^2_{\ell_2} \\ &\lesssim \frac{L}{Q} \|b_{i,\ell}\|^2_{\ell_2} \|w_{i,j}\|^2_{\ell_2} \lesssim \frac{K_{i,\mu}}{QL} \mu^2_{p-1}. \end{aligned}$$

Consequently, Theorem 8 applied with $t = (\omega+2)\log L$ yields that

$$\left| \sum_{j\in\Gamma_p} z_{i,j} \right| \lesssim_\omega \frac{\mu_{p-1}}{\sqrt{L}} \max\left\{ \sqrt{\frac{K_{i,\mu}\log L}{Q}}; \frac{K_{i,\mu}}{Q}\log L \right\}$$

with probability $1 - \mathcal{O}\left(L^{-\omega-2}\right)$, which shows (VI.55).

Step 3 (Proof of (VI.56)): As for $k_1 \neq i, k_2 \neq i$ the vectors $\mathbf{z}_{k_1,j}$ and $\mathbf{z}_{k_2,j}$ are not independent, we will condition on the random variables $\left\{c_{i,j}\right\}_{j\in\Gamma_p}$, use that the random variables $\left\{\mathbf{z}_{k,j}\right\}_{k,j}$ are conditionally independent, and then apply Corollary 10. For that, we bound

$$\begin{aligned} &\left| \sum_{k\neq i} \sum_{j\in\Gamma_p} \mathbb{E}\left[\mathbf{z}^*_{k,j} \mathbf{z}_{k,j} \middle| \left\{c_{i,j}\right\}_{j\in\Gamma_p} \right] \right| && \text{(VI.61)}\\ &= \frac{L^2}{Q^2} \sum_{k\neq i} \sum_{j\in\Gamma_p} \|w_{k,j}\|^2_{\ell_2} \|c_{i,j}\|^2_{\ell_2} |h_i^* b_{i,j}|^2 |h_i^* S_{i,p+1} b_{i,\ell}|^2 \\ &\le \mu^2_{p-1} \frac{\mu_h^2}{Q^2} \left(\max_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \right) \sum_{k\neq i} \sum_{j\in\Gamma_p} |h_i^* b_{i,j}|^2 \\ &\le \mu^2_{p-1} \frac{\mu_h^2}{LQ} \left(\max_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \right) \sum_{k\neq i} \|T^{1/2}_{i,p} h_i\|^2_{\ell_2} \\ &\lesssim \mu^2_{p-1} \frac{r\mu_h^2}{QL} \left(\max_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \right). && \text{(VI.62)} \end{aligned}$$

Analogously, using the triangle inequality,

$$\begin{aligned} &\left\| \sum_{k\neq i} \sum_{j\in\Gamma_p} \mathbb{E}\left[\mathbf{z}_{k,j} \mathbf{z}^*_{k,j} \middle| \left\{c_{i,j}\right\}_{j\in\Gamma_p} \right] \right\|_{2\to 2} \\ &= \frac{L^2}{Q^2} \Big\| \sum_{k\neq i} \sum_{j\in\Gamma_p} c_{i,j} c^*_{i,j} \mathbb{E}\left[|c^*_{k,j} w_{k,j}|^2 \right] \\ &\quad \cdot |h_i^* b_{i,j}|^2 |h_i^* S_{i,p+1} b_{i,\ell}|^2 \Big\|_{2\to 2} \\ &\le \frac{L^2}{Q^2} \sum_{k\neq i} \sum_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \|w_{k,j}\|^2_{\ell_2} |h_i^* b_{i,j}|^2 |h_i^* S_{i,p+1} b_{i,\ell}|^2 \\ &\overset{(VI.62)}{\lesssim} \mu^2_{p-1} \frac{r\mu_h^2}{QL} \left(\max_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \right). \end{aligned}$$

Conditionally on $\left\{c_{i,j}\right\}_{j\in\Gamma_p}$, we can now apply Corollary 10 with $t = (\omega+2)\log L$. Together with the last two estimates this yields that with probability $1 - \mathcal{O}\left(L^{-\omega-2}\right)$

$$\left\| \sum_{k\neq i} \sum_{j\in\Gamma_p} \mathbf{z}_{k,j} \right\|_{\ell_2} \lesssim_\omega \mu_{p-1} \sqrt{\frac{r\mu_h^2 \left(\max_{j\in\Gamma_p} \|c_{i,j}\|^2_{\ell_2} \right) \log L}{QL}}.$$

Then, we can truncate the random variables $\left\{c_{i,j}\right\}_{j\in\Gamma_p}$ by Lemma 40 and obtain that inequality (VI.56) holds with probability $1 - \mathcal{O}\left(L^{-\omega-2}\right)$, if the constant in (VI.52) is chosen large enough.

Step 4 (Proof of (VI.57)): Note that conditionally on $\left\{c_{i,j}\right\}_{j\in\Gamma_p}$ $\sum_{k\neq i} \sum_{j\in\Gamma_p} z_{k,j}$ is a circular symmetric random variable with variance

$$\begin{aligned} &\mathbb{E}\left[\sum_{k\neq i} \sum_{j\in\Gamma_p} |z_{k,j}|^2 \middle| \left\{c_{i,j}\right\}_{j\in\Gamma_p} \right] \\ &= \frac{L^2}{Q^2} \sum_{k\neq i} \sum_{j\in\Gamma_p} |b^*_{i,\ell} S_{i,p+1} \left(\mathrm{Id} - h_i h_i^* \right) b_{i,j}|^2 \|w_{k,j}\|^2_{\ell_2} |c^*_{i,j} m_i|^2 \\ &\le \mu^2_{p-1} \frac{1}{Q} \left(\max_{j\in\Gamma_p} |c^*_{i,j} m_i|^2 \right) \end{aligned}$$

$$\cdot \sum_{k\neq i} \|T_{i,p}^{1/2}\left(\mathrm{Id}-h_i h_i^*\right) S_{i,p+1} b_{i,\ell}\|_{\ell_2}^2$$
$$\lesssim \mu_{p-1}^2 \frac{r K_{i,\mu}}{QL}.$$

Consequently, one obtains that with probability at least $1-\mathcal{O}\left(L^{-\omega-2}\right)$

$$\left|\sum_{k\neq i}\sum_{j\in\Gamma_p} z_{k,j}\right| \lesssim_\omega \mu_{p-1}\sqrt{\frac{\left(\max_{j\in\Gamma_p}|c_{i,j}^* m_i|^2\right) r K_{i,\mu}\log L}{QL}}.$$

Thus, by Lemma 40 inequality (VI.57) holds with probability at least $1-\mathcal{O}\left(L^{-\omega-2}\right)$, if the constant in (VI.52) is chosen large enough.

Union bound: By the previous four steps we see that for fixed $p\in[P]$, $\ell\in\Gamma_{p+1}$, and $i\in[r]$ the inequalities (VI.54), (VI.55), (VI.56), (VI.57) hold with probability $1-\mathcal{O}\left(L^{-\omega-2}\right)$. Thus, by (VI.53) and a union bound we have $\mu_{p-1}\leq\frac{1}{4}\mu_p$ with probability $1-rQ\ \mathcal{O}\left(L^{-\omega-2}\right)$ for fixed $p\in[P-1]$. Thus, with probability at most $1-rPQ\ \mathcal{O}\left(L^{-\omega-2}\right)$ we obtain $\mu_{p-1}\leq\frac{1}{4}\mu_p$ for all $p\in[P-1]$. We obtain the desired result as we find $r\lesssim Q\leq L$ and $PQ=L$. □

5) An Upper Bound for $\|z\|_{\ell_2}$: In the case of noise, the error bound given by Lemma 18 is proportional to $\|z\|_{\ell_2}$, where z is the dual certificate as constructed in (VI.29). Thus, one needs an upper bound for $\|z\|_{\ell_2}$. This will be accomplished by the following lemma.

Lemma 35: Let $z\in\mathbb{C}^L$ be given by (VI.29) and assume that $\|W_p\|_F\leq 4^{-p}\sqrt{r}$. Furthermore, suppose that $\mathcal{A}^p$ satisfies the δ-local isometry property (VI.10) with $\delta\leq\frac{1}{4}$ on $\mathcal{T}^p$ for all $p\in[P]$. Then

$$\|z\|_{\ell_2}\lesssim\sqrt{r}.$$

Proof: Observe that

$$\|z\|_{\ell_2}\leq\sum_{p=1}^{P}\|\mathcal{A}^p\mathcal{S}^p\left(W_{p-1}\right)\|_{\ell_2}\lesssim\sum_{p=1}^{P}\|W_{p-1}\|_F$$
$$\lesssim\sum_{p=0}^{P-1}4^{-p}\sqrt{r}\lesssim\sqrt{r},$$

where the first equality follows from the definition of z (VI.29) and the triangle inequality. The second inequality is due to the local isometry property (VI.10) and (VI.33). We derive by (VI.34) the desired bound. □

D. Proof of Theorem 6

First of all, recall that by Lemma 20 with probability at least $1-2\exp(-t)$ it holds that

$$\gamma=\|\mathcal{A}\|_{F\to 2}\leq 2\sqrt{\omega\max\left\{1;\frac{rK_\mu N}{L}\right\}\log\left(L+rKN\right)}. \tag{VI.63}$$

In the following, let $\left\{\Gamma_p\right\}_{p=1}^P$ be an ω-admissible partition of $[L]$ (see Definition 2), which is a minimizer of (II.8). From Definition 2 combined with the assumptions on L (see (II.9)) we infer that

$$Q=\frac{L}{P}\gtrsim r\left(K_\mu\log\left(K_\mu\right)+N\mu_h^2\right)(\log L)^2 \tag{VI.64}$$
$$P\geq\frac{1}{2}\log\left(8\gamma\sqrt{r}\right). \tag{VI.65}$$

Note that due to Theorem 22 and our assumptions on L and Q (and also $\log K_\mu\leq\log L$) we may assume that the inequalities (VI.9) and (VI.10) hold with probability $1-\mathcal{O}\left(L^{-\omega}\right)$ and constant $\delta=\frac{1}{32}$. Thus, by Lemma 18 applied with $\alpha=\frac{1}{8\gamma}$, $\beta=\frac{1}{4}$, and $\delta=\frac{1}{4}$ it is enough to construct $Y\in\mathrm{Range}\left(\mathcal{A}^*\right)$ which satisfies (VI.1) and (VI.2). This is achieved by the Golfing Scheme as explained in Section VI-C.1: Note that the assumption of Lemma 31 is given by (VI.65) and (VI.10). Thus, it holds that $\|W_p\|_F\leq 4^{-p}\sqrt{r}$ for all $p\leq P$ and, by (VI.28), $Y=Y_P$ satisfies Condition (VI.1). Furthermore, observe that Lemma 34 implies that with probability $1-\mathcal{O}\left(L^{-\omega}\right)$ one has $\mu_p\leq\frac{1}{4}\mu_{p-1}$ for all $p\in[P-1]$. Using this fact and $\|W_p\|_F\leq 4^{-p}\sqrt{r}$ it follows from Lemma 32 that Condition (VI.2) is fulfilled. Using a union bound we conclude that with probability $1-\mathcal{O}\left(L^{-\omega}\right)$ the approximate dual certificate $Y=Y_P$ satisfies the assumptions in Lemma 18. Thus, if $\hat{X}$ is a minimizer of (I.5) it satisfies the estimation error (VI.3).

It remains to prove the upper bound for the estimation error in order to obtain inequality (II.10). Note that by Lemma 35 we have that $\|z\|_{\ell_2}\lesssim\sqrt{r}$. Thus, in combination with (VI.63) we derive

$$\|\hat{X}-X^0\|_F\ \lesssim\ (1+\gamma)\left(1+\|z\|_{\ell_2}\right)\tau$$
$$\lesssim_\omega\ \tau\sqrt{r\max\left\{1;\frac{rK_\mu N}{L}\right\}\log L}.$$

This finishes the proof. □

APPENDIX A
CONSTRUCTION OF THE PARTITION $\left\{\Gamma_p\right\}_{p\in[P]}$

A. Proof of Lemma 3

The goal of this section is to prove Lemma 3. Our proof will rely on the following lemma.

Lemma 36: Fix $i\in[r]$ and let $Q\in(0,L)$, $\delta>0$ and $\nu\in(0,1)$. Assume that

$$Q\geq C\frac{K_{i,\mu}}{\nu^2}\log\frac{K_i}{\delta}, \tag{A.1}$$

where $C>0$ is an absolute constant and let $\hat{\delta}_1,\ldots,\hat{\delta}_L$ be independent, identically distributed random variables such that

$$\mathbb{P}\left(\hat{\delta}_1=1\right)=\frac{Q}{L}\quad\text{and}\quad\mathbb{P}\left(\hat{\delta}_1=0\right)=1-\frac{Q}{L}.$$

Then with probability exceeding $1-\delta$ we have that

$$\left\|\frac{L}{Q}\sum_{\ell=1}^{L}\hat{\delta}_\ell b_{i,\ell}b_{i,\ell}^*-Id\right\|_{2\to 2}\leq\nu.$$

A proof of this lemma can be obtained using arguments contained in the proof of [66, Th. 1.2]. For the sake of

completeness we will give a proof below (relying on different techniques). Our proof of Lemma 3 will use essentially the same ideas as in [27], but has been slightly refined.

Proof of Lemma 3: Let $\hat{\delta}_1, \dots, \hat{\delta}_k$ be independent, uniformly distributed random variables which take values in $[P]$. For $p \in [P]$ we define

$$\Gamma_p = \left\{ \ell \in [L] : \ \hat{\delta}_\ell = p \right\}.$$

Thus, $\left\{\Gamma_p\right\}_{p\in[P]}$ is a partition of $[L]$. To finish the proof it is enough to show that with positive probability the partition $\left\{\Gamma_p\right\}_{p\in[P]}$ has the required properties, i.e., for all $p \in [P]$, (II.5) holds and $\frac{1}{2}Q \leq |\Gamma_p| \leq \frac{3}{2}Q$. For $i \in [r]$ and $p \in [P]$ we define the event

$$A_{i,p} = \{(II.5) \text{ fails}\} = \left\{ \left\| \frac{L}{Q} \sum_{\ell \in \Gamma_p} b_{i,\ell} b_{i,\ell}^* - \mathrm{Id} \right\|_{2\to 2} > \nu \right\}.$$

Set $\delta = \frac{1}{3rP}$ and note that $\log(\frac{K}{\delta}) = \log(3rPK) \lesssim \log(\max\{r; P; K\})$. Thus, by Lemma 36 we get that $\mathbb{P}\left(A_{i,p}\right) \leq \frac{1}{3rP}$, if the constant in inequality (II.6) is chosen large enough. By a union bound over all choices of i and p, (II.5) follows with probability at least $\frac{1}{3}$. It remains to control the size of the sets $\{\Gamma_p\}_{p\in[P]}$. By the Bernstein inequality for bounded random variables (e.g., [65, Corollary 7.31]) we obtain that for fixed $p \in [P]$ one has $\frac{Q}{2} \leq |\Gamma_p| \leq \frac{3Q}{2}$ with probability at least $1 - 2\exp\left(\frac{-Q}{10}\right) \geq 1 - \frac{1}{2P}$, where the last inequality follows from (II.6), if the constant C is chosen large enough. Thus, by another union bound we observe

$$\mathbb{P}\left(\frac{Q}{2} \leq |\Gamma_p| \leq \frac{3Q}{2} \text{ for all } p \in [P] \right) > \frac{1}{2}.$$

Thus with positive probability the partition $\left\{\Gamma_p\right\}_{p\in[P]}$ has the required properties. In particular, this implies the existence of a partition $\left\{\Gamma_p\right\}_{p\in[P]}$ with the properties stated in Lemma 3.□

B. Proof of Lemma 36

As already mentioned before this lemma can be proven using arguments from the proof [66, Th. 1.2]. The arguments in this article are based on Talagrand's inequality [67] and Rudelson's Lemma [68]. Recent technical advances (see [56]) allow us to give a simplified proof.

Proof: The goal is to use the matrix Bernstein inequality to estimate the spectral norm of

$$Y = \frac{L}{Q} \sum_{\ell=1}^{L} \hat{\delta}_\ell b_{i,\ell} b_{i,\ell}^* - \mathrm{Id}.$$

We will decompose Y into a sum of independent random matrices with mean zero. Thus, by setting

$$Y_\ell = \left(\hat{\delta}_\ell - \frac{Q}{L} \right) \frac{L}{Q} b_{i,\ell} b_{i,\ell}^*$$

we obtain $Y = \sum_{\ell=1}^{L} Y_\ell$ and $\mathbb{E}Y_\ell = 0$ for all $\ell \in [L]$ due to $\mathrm{Id} = \sum_{\ell=1}^{L} b_{i,\ell} b_{i,\ell}^*$. To apply the matrix Bernstein inequality we need first to obtain an upper bound for $\|\mathbb{E}Y^2\|_{2\to 2}$. For that purpose note that

$$\mathbb{E}Y^2 = \sum_{\ell=1}^{L} \mathbb{E}Y_\ell^2 = \sum_{\ell=1}^{L} \mathbb{E}\left[\left(\hat{\delta}_\ell - \frac{Q}{L} \right)^2 \right] \frac{L^2}{Q^2} \|b_{i,\ell}\|_{\ell_2}^2 b_{i,\ell} b_{i,\ell}^*.$$

Observe that $\mathbb{E}\left[\left(\hat{\delta}_\ell - \frac{Q}{L} \right)^2 \right] = \frac{Q(L-Q)}{L^2}$, which implies

$$\mathbb{E}Y^2 = \frac{L-Q}{L} \sum_{\ell=1}^{L} \frac{L\|b_{i,\ell}\|_{\ell_2}^2}{Q} b_{i,\ell} b_{i,\ell}^*.$$

Thus, by $\sum_{\ell=1}^{L} b_{i,\ell} b_{i,\ell}^* = \mathrm{Id}$ and the definition of $K_{i,\mu}$ we get

$$\begin{aligned} \|\mathbb{E}Y^2\|_{2\to 2} &\leq \frac{L-Q}{L} \left(\max_{\ell\in[L]} L\|b_{i,\ell}\|_{\ell_2}^2 \right) \left\| \sum_{\ell=1}^{L} b_{i,\ell} b_{i,\ell}^* \right\|_{2\to 2} \\ &\leq \frac{K_{i,\mu}}{Q}. \end{aligned}$$

Furthermore, for all $\ell \in [L]$ we have

$$\begin{aligned} \|Y_\ell\|_{2\to 2} &\leq \max\left\{ \frac{Q}{L}; \frac{L-Q}{L} \right\} \frac{L}{Q} \|b_{i,\ell}\|_{\ell_2}^2 \\ &\leq \frac{L}{Q} \|b_{i,\ell}\|_{\ell_2}^2 \leq \frac{K_{i,\mu}}{Q} \quad \text{almost surely.} \end{aligned}$$

Thus, we can apply the matrix Bernstein inequality in the version of [56, Th. 6.6.1] to obtain

$$\begin{aligned} \mathbb{P}\left(\|Y\|_{2\to 2} \geq \nu \right) &\leq K \exp\left(\frac{-\nu^2/2}{\left(1 + \frac{\nu}{3}\right) K_{i,\mu}/Q} \right) \\ &\overset{(A.1)}{\leq} K \exp\left(\frac{-C \log(K/\delta)}{2\left(1 + \frac{\nu}{3}\right)} \right). \end{aligned}$$

As we have $0 < \nu < 1$ this yields the claim if the constant $C > 0$ in (A.1) is chosen large enough.

□

APPENDIX B
CIRCULAR-SYMMETRIC COMPLEX NORMAL RANDOM VARIABLES

In this section we will recall some useful facts concerning random variables which have a circular–symmetric complex normal distribution $\mathcal{CN}(0, \sigma^2)$ with zero mean and variance σ^2. This means that their real and imaginary parts are uncorrelated jointly Gaussian with zero mean and variance $\sigma^2/2$ (and are therefore independent). For more details concerning this probability distribution we refer to [69, Sec. A.1.3]. The following two well-known lemmas are concerned with two useful identities. A proof of them can be found for example in [27, Lemma 11 and 12].

Lemma 37: Assume that $c \in \mathbb{C}^n$ is a random vector with independent entries $c_i \sim \mathcal{CN}(0, 1)$. Then we have

$$\mathbb{E}\left[\left(Id - cc^* \right)^2 \right] = nId.$$

Lemma 38: Let $q \in \mathbb{C}^n$ be any deterministic vector. Furthermore, assume that $c \in \mathbb{C}^n$ is a random vector with independent entries $c_i \sim \mathcal{CN}(0, 1)$. Then we have

$$\mathbb{E}\left[\left(cc^* - Id \right) qq^* \left(cc^* - Id \right) \right] = \|q\|_{\ell_2}^2 Id.$$

The following lemma summarizes well-known facts regarding the tail decay of certain quantities which involve circular-symmetric normal random variables. For the sake of completeness we include a proof.

Lemma 39: Suppose that $c \in \mathbb{C}^N$ is a random vector with independent entries $c_i \sim \mathcal{CN}(0,1)$. Let $p, q \in \mathbb{C}^N$ be arbitrary. Then we have the following inequalities:

$$\left\| \|c\|_{\ell_2} \right\|_{\psi_2} \lesssim \sqrt{N} \tag{B.1}$$

$$\left\| |c^* q| \right\|_{\psi_2} \lesssim \|q\|_{\ell_2} \tag{B.2}$$

$$\left\| \| \left(cc^* - Id\right) q \|_{\ell_2} \right\|_{\psi_1} \lesssim \sqrt{N} \|q\|_{\ell_2} \tag{B.3}$$

$$\left\| p^* \left(cc^* - Id\right) q \right\|_{\psi_1} \lesssim \|p\|_{\ell_2} \|q\|_{\ell_2} \tag{B.4}$$

Proof: In order to prove (B.1) note that

$$\left\| \|c\|_{\ell_2} \right\|_{\psi_2}^2 \lesssim \left\| \|c\|_{\ell_2}^2 \right\|_{\psi_1} \leq \sum_{i=1}^N \left\| |c_i|^2 \right\|_{\psi_1} \lesssim N.$$

The first inequality follows from [53, Lemma 5.14] and for the second one we used the triangle inequality. In order to prove (B.2) it is enough to note that $c^* q \sim \mathcal{CN}\left(0, \|q\|_{\ell_2}^2\right)$. (B.3) follows from the inequality chain

$$\begin{aligned} \left\| \| \left(cc^* - \mathrm{Id}\right) q \|_{\ell_2} \right\|_{\psi_1} &\leq \left\| \|c\|_{\ell_2} |c^* q| + \|q\|_{\ell_2} \right\|_{\psi_1} \\ &\leq \left\| \|c\|_{\ell_2} \right\|_{\psi_2} \left\| |c^* q| \right\|_{\psi_2} + \left\| \|q\|_{\ell_2} \right\|_{\psi_1} \\ &\lesssim \sqrt{N} \|q\|_{\ell_2} + \|q\|_{\ell_2} \lesssim \sqrt{N} \|q\|_{\ell_2}. \end{aligned}$$

In the second inequality we have used the Hoelder inequality (IV.1) and the second line follows directly from (B.1) and (B.2). In a similar way one proves (B.4). □

We will also need the following standard fact, which follows from a union bound.

Lemma 40: Let $\omega, L \geq 1$ and Γ a finite set. For all $i \in [r]$ let $m_i \in \mathbb{C}^{N_i}$ such that $\|m_i\|_{\ell_2} = 1$. Furthermore, assume that $c_{i,j} \in \mathbb{C}^{N_i}$, $i \in [r]$, $j \in \Gamma$, are independent random vectors with i.i.d. entries distributed according to $\mathcal{CN}(0,1)$. Then with probability at least $1 - \mathcal{O}\left(L^{-\omega}\right)$ one has

$$\max_{i \in [r], j \in \Gamma} \|c_{i,j}\|_{\ell_2} \lesssim_\omega \max\left\{ \sqrt{N \log\left(r|\Gamma|\right)};\ \sqrt{N \log L} \right\}$$

$$\max_{i \in [r], j \in \Gamma} |c_{i,j}^* m_i| \lesssim_\omega \max\left\{ \sqrt{\log\left(r|\Gamma|\right)}; \sqrt{\log L} \right\}.$$

We conclude this section with a proof of Corollary 10.

Proof of Corollary 10: Observe that

$$\left\| Z \right\|_{2 \to 2} \leq \left\| \sum_{i=1}^n \mathrm{Re}\left(\gamma_i\right) X_i \right\|_{2 \to 2} + \left\| \sum_{i=1}^n \mathrm{Im}\left(\gamma_i\right) X_i \right\|_{2 \to 2}.$$

By Theorem [70, Th. 4.1.1] we obtain that with probability at least $1 - \exp\left(-t\right)$

$$\left\| \sum_{i=1}^n \mathrm{Re}\left(\gamma_i\right) X_i \right\|_{2 \to 2} \leq \frac{1}{\sqrt{2}} \sigma \sqrt{t + \log\left(d_1 + d_2\right)}$$

and with probability at least $1 - \exp\left(-t\right)$

$$\left\| \sum_{i=1}^n \mathrm{Im}\left(\gamma_i\right) X_i \right\|_{2 \to 2} \leq \frac{1}{\sqrt{2}} \sigma \sqrt{t + \log\left(d_1 + d_2\right)}.$$

Combining these facts yields the result. □

APPENDIX C
PROOF OF LEMMA 26

For $i \in [r]$ let $\mathcal{N}_i$ be an $\frac{\varepsilon}{2}$-cover of $B\left(0,1\right) \subset \mathbb{C}^{K_i}$ with respect to the $\|\cdot\|_{B_i}$-norm. Furthermore, let $\mathcal{O}$ be an $\frac{\varepsilon}{2\sqrt{K_\mu}}$-cover of $B\left(0,1\right) \subset \mathbb{R}^r$ with respect to the $\|\cdot\|_{\ell_2}$-norm. We will show that any $Z = \left(u_1 m_1^*, \ldots, u_r m_r^*\right) \in B^m$ can be approximated by $Y = \left(\sigma_1 y_1 m_1^*, \ldots, \sigma_r y_r m_r^*\right)$, where $\sigma = \left(\sigma_1, \ldots, \sigma_r\right) \in \mathcal{O}$ and $y_i \in \mathcal{N}_i$. This proves the claim, as the number of such Y's is bounded by the right-hand side. For that choose $\sigma = \left(\sigma_1, \ldots, \sigma_r\right) \in \mathcal{O}$ such that

$$\sqrt{\sum_{i=1}^r \left(\|u_i\|_{\ell_2} - \sigma_i\right)^2} \leq \frac{\varepsilon}{2\sqrt{K_\mu}} \tag{C.1}$$

and $y_i \in \mathcal{N}_i$ such that

$$\left\| \frac{1}{\|u_i\|_{\ell_2}} u_i - y_i \right\|_{B_i} \leq \frac{\varepsilon}{2}. \tag{C.2}$$

Then one has for $\hat{Y} = \left(\|u_1\|_{\ell_2} y_1 m_1^*, \ldots, \|u_r\|_{\ell_2} y_r m_r^*\right)$

$$\begin{aligned} \left\| Z - \hat{Y} \right\|_B^2 &\leq \sum_{i=1}^r \left\| u_i m_i^* - \|u_i\|_{\ell_2} y_i m_i^* \right\|_{B_i}^2 \\ &= \sum_{i=1}^r \left\| u_i - \|u_i\|_{\ell_2} y_i \right\|_{B_i}^2 \\ &\leq \frac{\varepsilon^2}{4} \sum_{i=1}^r \|u_i\|_{\ell_2}^2 = \frac{\varepsilon^2}{4} \|Z\|_F^2 \leq \frac{\varepsilon^2}{4}. \end{aligned}$$

The first inequality follows from (VI.13) and the next equality follows from

$$\|m_i \left(u_i - \|u_i\|_{\ell_2} y_i\right)^* b_{i,\ell}\|_{\ell_2} = |\left(u_i - \|u_i\|_{\ell_2} y_i\right)^* b_{i,\ell}|$$

which is due to $\|m_i\|_{\ell_2} = 1$. The subsequent inequality is a consequence of (C.2). The second equality again follows from $\|m_i\|_{\ell_2} = 1$ for all $i \in [r]$. Similarly,

$$\begin{aligned} \|\hat{Y} - Y\|_B &\leq \sqrt{\sum_{i=1}^r \left\| \left(\|u_i\|_{\ell_2} - \sigma_i\right) y_i m_i^* \right\|_{B_i}^2} \\ &= \sqrt{\sum_{i=1}^r \left(\|u_i\|_{\ell_2} - \sigma_i\right)^2 \|y_i\|_{B_i}^2} \\ &\leq \sqrt{K_\mu \sum_{i=1}^r \left(\|u_i\|_{\ell_2} - \sigma_i\right)^2} \leq \frac{\varepsilon}{2}. \end{aligned}$$

Here the second inequality follows from

$$\|y_i\|_{B_i} = \sqrt{L} \max_{\ell \in [L]} |y_i^* b_{i,\ell}| \leq \sqrt{L} \|y_i\|_{\ell_2} \max_{\ell \in [L]} \|b_{i,\ell}\|_{\ell_2} \leq \sqrt{K_\mu}$$

and the last inequality is a consequence of (C.1). Combining the two inequalities gives $\|Z - Y\|_B \leq \varepsilon$ which finishes the proof. □

ACKNOWLEDGMENT

The three authors acknowledge support by the Hausdorff Institute for Mathematics (HIM), where part of this work was completed in the context of the HIM Trimester Program *Mathematics of Signal Processing*. Furthermore, the authors want to thank David Gross, Richard Kueng, Kiryung Lee, Shuyang Ling, and Tom Szollmann for fruitful discussions.

REFERENCES

[1] D. Stöger, P. Jung, and F. Krahmer, "Blind deconvolution and compressed sensing," in *Proc. Cosera*, 2016, pp. 24–27.

[2] D. Stöger, P. Jung, and F. Krahmer, "Blind demixing and deconvolution with noisy data: Near-optimal rate," in *Proc. 21st Int. ITG Workshop Smart Antenna*, 2017, pp. 1–5.

[3] D. Stöger, P. Jung, and F. Krahmer, "Blind demixing and deconvolution with noisy data at near optimal rate," *Proc. SPIE*, vol. 10394, p. 103941E, Aug. 2017. [Online]. Available: http://dx.doi.org/10.1117/12.2271571

[4] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[5] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[6] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.

[7] H. Rauhut, R. Schneider, and Ž. Stojanac, "Low rank tensor recovery via iterative hard thresholding," *Linear Algebra Appl.*, vol. 523, pp. 220–262, Jun. 2017.

[8] M. B. McCoy and J. A. Tropp, "Sharp recovery bounds for convex demixing, with applications," *Found. Comput. Math.*, vol. 14, no. 3, pp. 503–567, 2014.

[9] L. Balzano and R. Nowak, "Blind calibration of sensor networks," in *Proc. 6th Int. Conf. Inf. Process. Sens. Netw.*, 2007, pp. 79–88.

[10] M. A. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 342–349, Apr. 2010.

[11] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2182–2195, May 2011.

[12] S. Gleichman and Y. C. Eldar, "Blind compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6958–6975, Oct. 2011.

[13] X. Wang and H. V. Poor, "Blind equalization and multiuser detection in dispersive CDMA channels," *IEEE Trans. Commun.*, vol. 46, no. 1, pp. 91–103, Jan. 1998.

[14] S. Haykin, *Blind Deconvolution*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1994.

[15] T. G. Stockham, T. M. Cannon, and R. B. Ingebretsen, "Blind deconvolution through digital signal processing," *Proc. IEEE*, vol. 63, no. 4, pp. 678–692, Apr. 1975.

[16] G. H. Godard, "Self-recovering equalization and carrier tracking in two dimensional data communication systems," *IEEE Trans. Commun.*, vol. COM-28, no. 11, pp. 1867–1875, Nov. 1980.

[17] S. Choudhary and U. Mitra. (2014). "Identifiability scaling laws in bilinear inverse problems." [Online]. Available: https://arxiv.org/abs/1402.2637

[18] S. Choudhary and U. Mitra. (2014). "Fundamental limits of blind deconvolution part I: Ambiguity kernel." [Online]. Available: https://arxiv.org/abs/1411.3810

[19] P. Walk, P. Jung, G. Pfander, and B. Hassibi, "Ambiguities on convolutions with applications to phase retrieval," in *Proc. Asilomar*, 2016, pp. 1228–1234.

[20] S. Choudhary and U. Mitra. (2015). "Fundamental limits of blind deconvolution part II: Sparsity-ambiguity trade-offs." [Online]. Available: https://arxiv.org/abs/1503.03184

[21] Y. Li, K. Lee, and Y. Bresler. (2015). "A unified framework for identifiability analysis in bilinear inverse problems with applications to subspace and sparsity models." [Online]. Available: https://arxiv.org/abs/1501.06120

[22] Y. Li, K. Lee, and Y. Bresler, "Identifiability and stability in blind deconvolution under minimal assumptions," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4619–4633, Jul. 2017.

[23] M. Kech and F. Krahmer, "Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems," *SIAM J. Appl. Algebra Geometry*, vol. 1, no. 1, pp. 20–37, 2017.

[24] T. F. Chan and C. K. Wong, "Convergence of the alternating minimization algorithm for blind deconvolution," *Linear Algebra Appl.*, vol. 316, nos. 1–3, pp. 259–285, Sep. 2000.

[25] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 1964–1971.

[26] M. S. C. Almeida and M. A. T. Figueiredo, "Blind image deblurring with unknown boundaries using the alternating direction method of multipliers," in *Proc. 20th IEEE Int. Conf. Image Process.*, Sep. 2013, pp. 586–590.

[27] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *IEEE Trans. Inf. Theory*, vol. 60, no. 3, pp. 1711–1732, Mar. 2014.

[28] A. L. Chistov and D. Y. Grigor, "ev, "Complexity of quantifier elimination in the theory of algebraically closed fields," in *Proc. Int. Symp. Math. Found. Comput. Sci.*, 1984, pp. 17–31.

[29] E. J. Candés and Y. Plan, "Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2342–2359, Apr. 2011.

[30] D. Gross, "Recovering low-rank matrices from few coefficients in any basis," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 1548–1566, Mar. 2011.

[31] X. Li, S. Ling, T. Strohmer, and K. Wei. (2016). "Rapid, robust, and reliable blind deconvolution via nonconvex optimization." [Online]. Available: http://arxiv.org/abs/1606.04933

[32] S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi, "Simultaneously structured models with application to sparse and low-rank matrices," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2886–2908, May 2015.

[33] E. Richard, G. R. Obozinski, and J.-P. Vert, "Tight convex relaxations for sparse matrix factorization," in *Proc. NIPS*, 2014, pp. 3284–3292.

[34] S. Ling and T. Strohmer, "Self-calibration and biconvex compressive sensing," *Inverse Problems*, vol. 31, no. 11, p. 115002, 2015.

[35] A. Flinth, "Sparse blind deconvolution and demixing through $\ell_{1,2}$-minimization," *Adv. Comput. Math.*, pp. 1–21, Apr. 2017. [Online]. Available: https://doi.org/10.1007/s10444-017-9533-0

[36] M. Genzel and P. Jung. (2017). "Recovering structured data from superimposed non-linear measurements." [Online]. Available: https://arxiv.org/abs/1708.07451

[37] K. Lee, Y. Li, M. Junge, and Y. Bresler, "Blind recovery of sparse signals from subsampled convolution," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 802–821, Feb. 2017.

[38] K. Lee, Y. Wu, and Y. Bresler. (2013). "Near optimal compressed sensing of a class of sparse low-rank matrices via sparse power factorization." [Online]. Available: https://arxiv.org/abs/1312.0525

[39] K. Lee and M. Junge. (2015). "RIP-like properties in subsampled blind deconvolution." [Online]. Available: https://arxiv.org/abs/1511.06146

[40] G. Wunder, H. Boche, T. Strohmer, and P. Jung, "Sparse signal processing concepts for efficient 5G system design," *IEEE Access*, vol. 3, pp. 195–208, 2015.

[41] P. Jung and P. Walk, "Sparse model uncertainties in compressed sensing with application to convolutions and sporadic communication," in *Compressed Sensing and its Applications*, H. Boche, R. Calderbank, G. Kutyniok, and J. Vybiral, Eds. Berlin, Germany: Springer, 2015, pp. 1–29. [Online]. Available: http://arxiv.org/abs/1404.0218

[42] R. Kueng and P. Jung, "Robust nonnegative sparse recovery and the nullspace property of 0/1 measurements," *IEEE Trans. Inf. Theory*, to be published, doi: 10.1109/TIT.2017.2746620.

[43] S. Ling and T. Strohmer, "Blind deconvolution meets blind demixing: Algorithms and performance bounds," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4497–4520, Jul. 2017.

[44] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, Nov. 2001.

[45] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, "Living on the edge: Phase transitions in convex programs with random data," *Inf. Inference*, vol. 3, no. 3, pp. 224–294, Sep. 2014.

[46] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, "Sparse and low-rank matrix decompositions," *IFAC Proc. Vols.*, vol. 42, no. 10, pp. 1493–1498, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474667016388632

[47] M. B. McCoy and J. A. Tropp, "The achievable performance of convex demixing," Appl. Comput. Math.,California Inst. Technol., Pasadena, CA, USA, Tech. Rep. 2017-02, 2017.
[48] J. Wright, A. Ganesh, K. Min, and Y. Ma, "Compressive principal component pursuit," *Inf. Inference*, vol. 2, no. 1, pp. 32–68, 2013. [Online]. Available: http://imaiai.oxfordjournals.org/content/2/1/32.short
[49] S. Ling and T. Strohmer. (2017). "Regularized gradient descent: A nonconvex recipe for fast joint blind deconvolution and demixing." [Online]. Available: https://arxiv.org/abs/1703.08642
[50] B. G. Bodmann, P. G. Casazza, D. Edidin, and R. Balan, "Frames for linear reconstruction without phase," in *Proc. 42nd Annu. Conf. Inf. Sci. Syst.*, Mar. 2008, pp. 721–726.
[51] V. Koltchinskii, *A Remark Low Rank Matrix Recovery Noncommutative Bernstein Type Inequalities* (Collections), vol. 9. Beachwood, OH, USA: Inst. Math. Statist., 2013, pp. 213–226. [Online]. Available: http://dx.doi.org/10.1214/12-IMSCOLL915
[52] M. A. Krasnosel'skii and Y. Rutickii, *Convex Functions and Orlicz Spaces*. Gröningen, The Netherlands: Noordhoff, 1961.
[53] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," in *Compressed Sensing*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2012, pp. 210–268. [Online]. Available: http://dx.doi.org/10.1017/CBO9780511794308.006
[54] R. Ahlswede and A. Winter, "Strong converse for identification via quantum channels," *IEEE Trans. Inf. Theory*, vol. 48, no. 3, pp. 569–579, Mar. 2002.
[55] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Found. Comput. Math.*, vol. 12, no. 4, pp. 389–434, Aug. 2012.
[56] J. A. Tropp, "An introduction to matrix concentration inequalities," *Found. Trends Mach. Learn.*, vol. 8, nos. 1–2, pp. 1–230, 2015.
[57] M. Talagrand, *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems*. Berlin, Germany: Springer-Verlag, 2014.
[58] F. Krahmer, S. Mendelson, and H. Rauhut, "Suprema of chaos processes and the restricted isometry property," *Commun. Pure Appl. Math.*, vol. 67, no. 11, pp. 1877–1904, Nov. 2014.
[59] R. M. Dudley, "The sizes of compact subsets of Hilbert space and continuity of Gaussian processes," *J. Funct. Anal.*, vol. 1, no. 3, pp. 290–330, 1967. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0022123667900171
[60] B. Carl, "Inequalities of Bernstein-Jackson-type and the degree of compactness of operators in Banach spaces," *Ann. Inst. Fourier*, vol. 35, no. 3, pp. 79–118, 1985. [Online]. Available: http://eudml.org/doc/74689
[61] A. Pietsch, *Theorie der Operatorenideale* (Wissenschaftliche Beiträge der). Jena, Germany: Univ. Jena, 1972.
[62] S. Artstein, V. Milman, and S. J. Szarek, "Duality of metric entropy," *Ann. Math.*, vol. 159, no. 3, pp. 1313–1328, 2004.
[63] R. Bhatia, *Matrix Analysis*. New York, NY, USA: Springer, 1996.
[64] F. Krahmer and R. Ward, "Stable and robust sampling strategies for compressive imaging," *IEEE Trans. Image Process.*, vol. 23, no. 2, pp. 612–622, Feb. 2014.
[65] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. New York, NY, USA: Springer, 2013.
[66] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, p. 969, 2007.
[67] M. Talagrand, "New concentration inequalities in product spaces," *Invent. Math.*, vol. 126, no. 3, pp. 505–563, 1996.
[68] M. Rudelson, "Random vectors in the isotropic position," *J. Funct. Anal.*, vol. 164, no. 1, pp. 60–72, 1999.
[69] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. New York, NY, USA: Cambridge Univ. Press, 2005.
[70] J. A. Tropp, "Convex recovery of a structured signal from independent random linear measurements," in *Sampling Theory, a Renaissance. Applied and Numerical Harmonic Analysis*, G. Pfander, Ed. Cham, Swizerland: Birkhäuser, 2015.

Peter Jung (M'05) received the Dipl.-Phys. degree in high energy physics from Humboldt University, Berlin, Germany, in 2000, in cooperation with DESY, Hamburg, and the Dr.rer.nat. (Ph.D.) degree in Weyl–Heisenberg representations in communication theory from the Technical University of Berlin (TUB), Germany, in 2007. Since 2001, he has been with the Department of Broadband Mobile Communication Networks, Fraunhofer Institute for Telecommunications, Heinrich-Hertz-Institut, and the Fraunhofer German-Sino Laboratory for Mobile Communications since 2004. He is currently working under DFG grants JU 2795/3 with TUB in the field of signal processing and information and communication theory. His current research interests are in the area of compressed sensing, time–frequency analysis, dimension reduction, and randomized algorithms. He is giving lectures in compressed sensing and estimation theory. He is a member of VDE/ITG.

Felix Krahmer received his BSc in Mathematics from International University Bremen and his MSc as well as his PhD in Mathematics from New York University under the supervision of Percy Deift and Sinan Güntürk. He was a HCM postdoc in the group of Holger Rauhut at the University of Bonn, Germany from 2009-2012. In 2012 he joined the University of Göttingen as an assistant professor for mathematical data analysis, where he has been awarded an Emmy Noether Junior Research Group. Since 2015 he has been a tenure track assistant professor for optimization and data analysis in the department of mathematics at the Technical University of Munich. His research interests include mathematical signal and data processing, applied harmonic analysis, and random matrix theory.

Dominik Stöger received his BSc and his MSc both in mathematics, at the Techical University of Munich (TUM). He is currently a PhD student in the department of mathematics at TUM under the supervision of Felix Krahmer. His research interests include random matrix theory, mathematical signal and data processing, and calculus of variations.

B. Core article: Sparse Power Factorization: Balancing peakiness and sample complexity

Summary

Suppose one is given a signal

$$y = \mathcal{B}(u, v),$$

where $\mathcal{B} : \mathbb{C}^{n_1} \times \mathbb{C}^{n_2} \to \mathbb{C}^m$ is a sesquilinear map and $u \in \mathbb{C}^{n_1}$ is s_1-sparse and $v \in \mathbb{C}^{n_2}$ is s_2-sparse. Our manuscript [43] deals with the question under which circumstances one can recover u and v using a tractable (i.e., polynomial-time) algorithm.
Note that there is a unique linear map $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ such that $\mathcal{A}(uv^*) = \mathcal{B}(u, v)$ for all $u \in \mathbb{C}^{n_1}$ and $v \in \mathbb{C}^{n_2}$. In this paper we assume as in [77] that $\mathcal{A}$ can be represented as

$$(\mathcal{A}(X))(\ell) = \mathrm{Tr}\,(A_\ell^*, X),$$

where $\{A_\ell\}_{\ell=1}^m$ are i.i.d. random Gaussian matrices.

Lee, Wu, and Bresler proposed the follow two-step approach to recover u and v. First one uses a modified spectral initialization to obtain an initial guess v' for v. Having obtained this initial guess one uses an alternating minimization algorithm to iterativey refine the initial guess. They could show the following result for a fixed constant $0 < \mu < 1$. If $\|u\|_\infty \geq \mu\|u\|$, $\|v\|_\infty \geq \mu\|v\|$, and

$$m \gtrsim (s_1 + s_2) \log\left(\max\left\{\frac{n_1}{s_1}; \frac{n_2}{s_2}\right\}\right),$$

then their proposed algorithm is able to recover u and v with high probability. The condition on u and v means that a large fraction of the ℓ_2-mass of then vector is concentrated on one entry, which is arguably a rather restrictive condition. The question we examined in the article [43] was, whether we one can relaxe these conditions at the price of a slightly increased amount of measurements.

In our main result we have shown that this is indeed the case. Namely we could derive recovery guantees under the assumption that the parameter μ varies and under the assumption is concentrated on a few, but possibly more than one entry. A precise statement of our result can be found in [43, Theorem 2]

Statement of individual contribution

This paper builds upon a Master's thesis by Jakob Geppert that made initial steps in the direction of an appropriate initialization procedure but encountered some technical difficulties. Under my lead, we could overcome these difficulties subsequently, which resulted in the publication [43]. Felix Krahmer and Jakob Geppert provided technical input in joint discussions. The ingredients and understanding that had to be added account for the core of the paper. I was fully responsible for writing and preparing this manuscript. Felix Krahmer and I proofread and polished the article together.

Conference proceedings

The results of this paper have been presented in part at the 12th International Conference on Sampling Theory and Applications, July 3-7, 2017, Tallinn, Estonia [44] and in the IEEE Statistical Signal Processing Workshop 2018, June 10-13, Freiburg [105].

Remark B.1. *The DOI listed in the attached article will change (see the correspondence with Mr. Mercullo).*

Betreff: ACOM: Your manuscript entitled Sparse Power Factorization: Balancing peakiness and sample complexity
Von: "Karsten Urban (ACOM Managing EiC)" <em@editorialmanager.com>
Datum: 07.01.19, 1:28 nachm.
An: Dominik Stöger <dominik.stoeger@ma.tum.de>

CC: acom@uni-ulm.de

Ref.: Ms. No. ACOM-D-18-00123R1
Sparse Power Factorization: Balancing peakiness and sample complexity
Advances in Computational Mathematics

Dear Mr. Stöger,

I am pleased to inform you that your work has now been accepted for publication in Advances in Computational Mathematics.

It was accepted on 07 Jan 2019.

Thank you for submitting your work to this journal.

With kind regards

Karsten Urban (ACOM Managing EiC)
Managing Editor-in-Chief
Advances in Computational Mathematics

Please, take into account the minor comments by reviewer #1 when submitting the final files of your manuscript.

Reviewer #1: Thank you for improving the reading of Lemma 6. Please see some little suggestions attached.

Reviewer #2: The author has addressed my comments properly. I have no other comments to add now. I would like to suggest the acceptance of this paper.

--
There is additional documentation related to this decision letter. To access the file(s), please click the link below. You may also login to the system and click the 'View Attachments' link in the Action column.
https://acom.editorialmanager.com/l.asp?i=33513&l=VGSD03RA

Recipients of this email are registered users within the Editorial Manager database for this journal. We will keep your information on file to use in the process of submitting, evaluating and publishing a manuscript. For more information on how we use your personal details please see our privacy policy at https://www.springernature.com/production-privacy-policy. If you no longer wish to receive messages from this journal or you have questions regarding database management, please email our publication office, stating the journal name(s) and your email address(es):
PublicationOfficeSPI@springernature.com

Betreff: Copyright Transfer Statement (CTS) for your article in Advances in Computational Mathematics (9698)
Von: Springer <springerauthorquery@springeronline.com>
Datum: 11.04.19, 1:02 nachm.
An: dominik.stoeger@ma.tum.de

Copyright Transfer

11.04.2019

Besuchen Sie uns: springer.com

Copyright Transfer Statement

Lieber Autor, liebe Autorin,

Hinweis: Sie haben dem Copyright Transfer Statement (CTS) zugestimmt und erhalten diese E-Mail für Ihre persönliche Ablage.

1. Publication
The copyright to this article, (including any supplementary information and graphic elements therein (e.g. illustrations, charts, moving images) (the 'Article'), is hereby assigned for good and valuable consideration to Springer Science+Business Media, LLC, part of Springer Nature (the 'Assignee'). Headings are for convenience only.

2. Grant of Rights
In consideration of the Assignee evaluating the Article for publication, the Author(s) grant the Assignee without limitation the exclusive (except as set out in clauses 3, 4 and 5 a) iv), assignable and sub-licensable right, unlimited in time and territory, to copy-edit, reproduce, publish, distribute, transmit, make available and store the Article, including abstracts thereof, in all forms of media of expression now known or developed in the future, including pre- and reprints, translations, photographic reproductions and extensions. Furthermore, to enable additional publishing services, such as promotion of the Article, the Author(s) grant the Assignee the right to use the Article (including the use of any graphic elements on a stand-alone basis) in whole or in part in electronic form, such as for display in databases or data networks (e.g. the Internet), or for print or download to stationary or portable devices. This includes interactive and multimedia use as well as posting the Article in full or in part or its abstract on social media, and the right to alter the Article to the extent necessary for such use. The Assignee may also let third parties share the Article in full or in part or its abstract on social media and may in this context sub-license the Article and its abstract to social media users. Author(s) grant to Assignee the right to re-license Article metadata without restriction (including but not limited to author name, title, abstract, citation, references, keywords and any additional information as determined by Assignee).

3. Self-Archiving
Author(s) are permitted to self-archive a pre-print and an author's accepted manuscript version of their Article.

a. A pre-print is the author's version of the Article before peer-review has taken place ("Pre-Print"). Prior to acceptance for publication, Author(s) retain the right to make a Pre-Print of their Article

available on any of the following: their own personal, self-maintained website; a legally compliant, non-commercial pre-print server such as but not limited to arXiv and bioRxiv. Once the Article has been published, the Author(s) should update the acknowledgement and provide a link to the definitive version on the publisher's website: "This is a pre-print of an article published in [insert journal title]. The final authenticated version is available online at: https://doi.org/[insert DOI]".

b. An Author's Accepted Manuscript (AAM) is the version accepted for publication in a journal following peer review but prior to copyediting and typesetting that can be made available under the following conditions:
    a. Author(s) retain the right to make an AAM of their Article available on their own personal, self-maintained website immediately on acceptance,
    b. Author(s) retain the right to make an AAM of their Article available for public release on any of the following 12 months after first publication ("Embargo Period"): their employer's internal website; their institutional and/or funder repositories. AAMs may also be deposited in such repositories immediately on acceptance, provided that they are not made publicly available until after the Embargo Period.

An acknowledgement in the following form should be included, together with a link to the published version on the publisher's website: "This is a post-peer-review, pre-copyedit version of an article published in [insert journal title]. The final authenticated version is available online at: http://dx.doi.org/[insert DOI]".

4. Authors' Retained Rights
Author(s) retain the following non-exclusive rights for the published version provided that, when reproducing the Article or extracts from it, the Author(s) acknowledge and reference first publication in the Journal:

a. to reuse graphic elements created by the Author(s) and contained in the Article, in presentations and other works created by them;
b. they and any academic institution where they work at the time may reproduce the Article for the purpose of course teaching (but not for inclusion in course pack material for onward sale by libraries and institutions); and
c. to reproduce, or to allow a third party Assignee to reproduce the Article in whole or in part in any printed volume (book or thesis) written by the Author(s).

5. Warranties
The Author(s) warrant and represent that:

a. (i) the Author(s) are the sole copyright owners or have been authorised by any additional copyright owner(s) to assign the rights defined in clause 2, (ii) the Article does not infringe any intellectual property rights (including without limitation copyright, database rights or trade mark rights) or other third party rights and no licence from or payments to a third party are required to publish the Article, (iii) the Article has not been previously published or licensed, (iv) if the Article contains material from other sources (e.g. illustrations, tables, text quotations), Author(s) have obtained written permissions to the extent necessary from the copyright holder(s), to license to the Assignee the same rights as set out in Clause 2 but on a non-exclusive basis and without the right to use any graphic elements on a stand-alone basis and have cited any such material correctly;
b. all of the facts contained in the Article are according to the current body of science true and accurate;
c. nothing in the Article is obscene, defamatory, violates any right of privacy or publicity, infringes any other human, personal or other rights of any person or entity or is otherwise unlawful and that informed consent to publish has been obtained for all research participants;
d. nothing in the Article infringes any duty of confidentiality which any of the Author(s) might owe to anyone else or violates any contract, express or implied, of any of the Author(s). All of the institutions in which work recorded in the Article was created or carried out have authorised and approved such research and publication; and
e. the signatory (the Author or the employer) who has signed this agreement has full right, power and authority to enter into this agreement on behalf of all of the Author(s).

6. Cooperation
The Author(s) shall cooperate fully with the Assignee in relation to any legal action that might arise from the publication of the Article, and the Author(s) shall give the Assignee access at reasonable times to any relevant accounts, documents and records within the power or control of the Author(s). The Author(s) agree

that the distributing entity is intended to have the benefit of and shall have the right to enforce the terms of this agreement.

7. Author List
After signing, changes of authorship or the order of the authors listed will not be accepted unless formally approved in writing by the Assignee.

8. Edits & Corrections
The Author(s) agree(s) that the Assignee may retract the Article or publish a correction or other notice in relation to the Article if the Assignee considers in its reasonable opinion that such actions are appropriate from a legal, editorial or research integrity perspective.

Diese E-Mail wurde automatisch gesendet; Bitte antworten Sie nicht auf diese Adresse. Falls Sie Fragen haben, besuchen Sie bitte unsere Hilfe-Seiten.

Vielen Dank.

Mit freundlichen Grüßen,

Springer Autorenservice

Details des Artikels

Zeitschrift
Advances in Computational Mathematics

Titel des Artikels
Sparse power factorization: balancing peakiness and sample complexity

DOI
10.1007/s10444-019-09698-6

Korrespondenzautor
Dominik Stöger

Übertragung des Urheberrechts an
Springer Science+Business Media, LLC, part of Springer Nature

Übertragen am
Thu Apr 11 13:02:26 CEST 2019

Service Contacts

Springer Nature Customer Service Center
Tiergartenstr. 15-17
69121 Heidelberg
Germany
phone: +49 6221 345 0
fax: +49 6221 345 4229
customerservice@springernature.com

Springer New York, LCC
233 Spring Street
New York, NY 10013
USA
phone: +1 212 460 1500 or 800-SPRINGER
(Weekdays 8:30am - 5:30pm ET)
fax: +1 212-460-1700
customerservice@springernature.com

Betreff: RE: Fwd: Correction for Advances in Computational Mathematics, Article 9664
Von: Kenneth Mercullo <Kenneth.Mercullo@springer.com>
Datum: 10.04.19, 9:14 vorm.
An: Dominik Stöger <dominik.stoeger@ma.tum.de>, "CorrAdmin3@spi-global.com" <CorrAdmin3@spi-global.com>
Kopie (CC): Jakob Alexander Geppert <j.geppert@math.uni-goettingen.de>, Felix Krahmer <felix.krahmer@tum.de>

Dear Dr. Stöger,

I contacted our IT to check your article since it was done in typesetting last February. I regret to tell you that there was an error during transmission that caused some of the files to get corrupted.

As of writing, I have only one option to get this published at the soonest. I have to assign it a new DOI number and reprocess, this would mean that we have to send you online the copyright transfer forms to fill up again. You wont receive galleys as this process is already done.

I regret the inconvenience and the long delay this error has caused you and your colleagues.

I have initiated the process and the corresponding author will receive the e-mail for the online forms tomorrow.

Thank you.

Sincerely,

(Mr.) Kenneth Xavier U. Mercullo
Production Coordinator
Journals Production

Springer Nature
Tel: +63 (035) 532-9912 loc. 23618
Kenneth.Mercullo@springer.com
www.springernature.com

-----Original Message-----
From: Dominik Stöger [mailto:dominik.stoeger@ma.tum.de]
Sent: Tuesday, April 2, 2019 5:58 PM
To: Kenneth Mercullo; CorrAdmin3@spi-global.com
Cc: Jakob Alexander Geppert; Felix Krahmer
Subject: Re: Fwd: Correction for Advances in Computational Mathematics, Article 9664

Dear all,

A few months ago we were in contact because of our paper "Sparse Power Factorization: Balancing peakiness and sample complexity", which has been accepted for publication in Advances in Computational Mathematics.
I just wanted to ask what the current status is. Can you roughly estimate when the online-first version will be available?

Advances in Computational Mathematics
https://doi.org/10.1007/s10444-019-09664-2

Sparse power factorization: balancing peakiness and sample complexity

Jakob Geppert[1] · **Felix Krahmer**[2] · **Dominik Stöger**[2]

Received: 26 April 2018 / Accepted: 11 January 2019 /

Abstract

In many applications, one is faced with an inverse problem, where the known signal depends in a bilinear way on two unknown input vectors. Often at least one of the input vectors is assumed to be sparse, i.e., to have only few non-zero entries. Sparse power factorization (SPF), proposed by Lee, Wu, and Bresler, aims to tackle this problem. They have established recovery guarantees for a somewhat restrictive class of signals under the assumption that the measurements are random. We generalize these recovery guarantees to a significantly enlarged and more realistic signal class at the expense of a moderately increased number of measurements.

Keywords Bilinear inverse problems · Sparse power factorization · Compressed sensing

Mathematics Subject Classification (2010) 94A12

Communicated by: Holger Rauhut

Jakob Geppert is supported by the German Science Foundation (DFG) in the Collaborative Research Centre "SFB 755: Nanoscale Photonic Imaging" and partially in the framework of the Research Training Group "GRK 2088: Discovering Structure in Complex Data: Statistics meets Optimization and Inverse Problems." Felix Krahmer and Dominik Stöger have been supported by the German Science Foundation (DFG) in the context of the joint project "SPP 1798: Bilinear Compressed Sensing" (KR 4512/2-1). The results of this paper have been presented in part at the 12th International Conference on Sampling Theory and Applications, July 3–7, 2017, Tallinn, Estonia [12], and in the IEEE Statistical Signal Processing Workshop 2018, June 10-13, Freiburg, Germany [34].

✉ Dominik Stöger

Jakob Geppert

Felix Krahmer

[1] Göttingen, Germany

[2] Garching, Germany

1 Introduction

Many measurement operations in signal and image processing as well as in communication follow a bilinear model. Namely, in addition to the measurements depending linearly on the unknown signal, also certain parameters of the measurement procedure enter in a linear fashion. Hence, one cannot employ a linear model (for example, in connection compressed sensing techniques [7]) unless one has an accurate estimate of these parameters.

When such estimates are not available or too expensive to obtain, there are certain asymmetric scenarios when one of the inputs can be recovered even though the other one is out of reach (e.g., [21, 37], this scenario is sometimes referred to as passive imaging). In most cases, however, the natural aim will be to recover both the signal and the parameters, that is, to solve the associated bilinear inverse problem. Even when some estimates of the parameters are available, such a unified approach will be preferred in many situations, especially when information is limited. Consequently, the study of bilinear inverse problems, including but not limited to the important problem of blind deconvolution, has been an active area of research for many years [13].

Observing that bilinear maps admit a representation as a linear map in the rank-one outer product of the unknown signal and the parameter vector, one can approach such problems using tools from the theory of low-rank recovery (see, e.g., [1, 17, 26]). Under sparsity assumptions, that is, when the signals and/or parameter vectors admit an approximate representation using just a small (but unknown) subset of an appropriate basis (for more details regarding when such assumptions appear in bilinear inverse problems, see [25]), however, the direct applicability of these approaches is limited, as two competing objectives arise: one aims to simultaneously minimize rank and sparsity. As a consequence, the problem becomes considerably more difficult; Oymak et al., for example, have demonstrated that minimizing linear combinations of the nuclear norm (a standard convex proxy for the rank) and the ℓ_1 norm (the corresponding quantity for sparsity) exhibits suboptimal scaling [31]. In fact, it is not even clear if without additional assumptions efficient recovery is at all possible for a near-linear number of measurements (as it would be predicted identifiability considerations [18]).

Recently, a number of non-convex algorithms for bilinear inverse problems have been proposed. For example, for such problems without sparsity constraints, several such algorithms have been analyzed for blind deconvolution and related problems [24, 27] with near-optimal recovery guarantees. In contrast, our understanding of bilinear inverse problems with sparsity constraints is only in its beginning. Recently, several algorithms have been analyzed for sparse phase retrieval [4, 33] or blind deconvolution with sparsity constraints [32]. The recovery guarantees for these algorithms, however, are either suboptimal in the number of necessary measurements or only local convergence guarantees are available, i.e., one relies on the existence of a good initialization. (A noteworthy exception is the two related papers [3, 14], where a two-stage approach for (sparsity) constrained bilinear inverse problems is proposed,

which achieves recovery at near-optimal rate. However, the algorithm relies on a special nested structure of the measurements, which is not feasible for many practical applications.)

In [23], Lee, Wu, and Bresler introduced the *sparse power factorization* (SPF) method together with a tractable initialization procedure based on alternating minimization. They also provide a first performance analysis of their method for random bilinear measurements in the sense that their lifted representation is a matrix with independent Gaussian entries. That is, they work with linear operators $\mathcal{A}\colon \mathbb{C}^{n_1\times n_2} \longrightarrow \mathbb{C}^m$ that admit a representation as

$$(\mathcal{A}(X))(\ell) = \operatorname{trace}\left(A_\ell^* X\right)$$

for i. i. d. Gaussian matrices $A_\ell \in \mathbb{C}^{n_1\times n_2}$.

For such measurements, they show that with high probability, SPF converges locally to the right solution, i.e., one has convergence for initializations not too far from the signal to be recovered.

For signals that have a very large entry, they also devise a tractable initialization procedure—they call it thresholding initialization—such that one has global convergence to the right solution. Local convergence has also been shown for the multi-penalty approach *A-T-LAS*$_{1,2}$ [10], but to our knowledge, comparable global recovery guarantees are not available to date. This is why we focus on SPF in this paper, using the results of [23] as our starting point.

The precise condition for their guarantee to hold is that both (normalized) input signals need to be larger than some $c > \frac{1}{2}$ in supremum norm—more than one quarter of its mass needs to be located in just one entry, that is, the signals must have a very high peak-to-average power ratio.

In this paper, we considerably weaken this rather strong restriction in two ways. Firstly, we show that similar results hold for smaller lower bounds c at the expense of a moderately increased number of measurements. Secondly, we show that similar results can be obtained when the mass of one of the signals is concentrated in more than one, but still a small number of entries.

The SPF algorithm, the thresholding initialization, and the resulting recovery guarantees are reviewed in Section 2 before we discuss and prove our results in Sections 4 and 5.

1.1 Notation

Throughout the paper, we will use the following notation. By $[n]$, we will denote the set $\{1; \ldots; n\}$. For any set J, we will denote its cardinality by $|J|$. For a vector $v \in \mathbb{C}^m$, we will denote by $\|v\|$ its ℓ_2-norm and by $\|v\|_\infty$ the modulus of its largest entry. If $J \subset [n]$, we will by v_J denote the restriction of v to elements indexed by J. For matrices $A \in \mathbb{C}^{n_1\times n_2}$, we will denote by $\|A\|_F$ its Frobenius norm and by $\|A\|$ its spectral norm, i.e., the largest singular value of A.

2 Sparse power factorization: algorithm and initialization

2.1 Problem formulation

Let $b \in \mathbb{C}^m$ be given by

$$b := B(u, v) + z,$$

where $B\colon \mathbb{C}^{n_1} \times \mathbb{C}^{n_2} \to \mathbb{C}^m$ is a bilinear map and $z \in \mathbb{C}^m$ is noise. Recall that one can represent the bilinear map $B\colon \mathbb{C}^{n_1} \times \mathbb{C}^{n_2} \to \mathbb{C}^m$ by a linear map $\mathcal{A}\colon \mathbb{C}^{n_1 \times n_2} \longrightarrow \mathbb{C}^m$, which satisfies

$$B(u, v) = \mathcal{A}(uv^*).$$

for all vectors $u \in \mathbb{C}^{n_1}$ and all $v \in \mathbb{C}^{n_2}$. Note that such a linear map $\mathcal{A}$ is characterized by a (unique) set of matrices $\{A_\ell\}_{\ell=1}^m \subset \mathbb{C}^{n_1 \times n_2}$ such that the ℓth entry of $\mathcal{A}(X)$ is given by

$$(\mathcal{A}(X))(\ell) = \operatorname{trace}\left(A_\ell^* X\right). \tag{1}$$

In this notation, our goal will be to reconstruct u and v from linear measurements given by

$$b_\ell = \operatorname{trace}\left(A_\ell^* uv^*\right)$$

At the core of the Sparse Power Factorization Algorithm, as introduced in [23], are the linear operators $F\colon \mathbb{C}^{n_2} \longrightarrow \mathbb{C}^{m \times n_1}$ and $G\colon \mathbb{C}^{n_1} \longrightarrow \mathbb{C}^{m \times n_2}$ defined by

$$F(y) := \begin{pmatrix} y^* A_1^* \\ \vdots \\ y^* A_m^* \end{pmatrix}, \quad G(x) := \begin{pmatrix} x^* A_1 \\ \vdots \\ x^* A_m \end{pmatrix}.$$

A direct consequence of this definition is that

$$\mathcal{A}(xy^*) = [F(y)]x = \overline{[G(x)]y}$$

for all $x \in \mathbb{C}^{n_1}$ and all $y \in \mathbb{C}^{n_2}$.

2.2 Sparse power factorization

The idea of sparse power factorization is to iteratively update estimates u_t and v_t for u and v in an alternating fashion. That is, in each iteration, one keeps one of v_t and u_t fixed and updates the respective other one by solving a (underdetermined) linear system. Solving each of these linear systems then amounts to solving a linear inverse problem with sparsity constraints. Hence, many pursuit algorithms proposed in the context of compressed sensing can be applied such as CoSaMP [30], Hard Thresholding Pursuit [11], or Basis Pursuit. In [23], the authors used Hard Thresholding Pursuit (HTP) for their analysis and in this paper, we will also restrict ourselves to HTP. With this, the Sparse Power Factorization Algorithm reads as follows.

Algorithm 1 Algorithm 1 in [23].

Input: Operator $\mathcal{A}$, Measurement b, Sparsity Constraints s_1, s_2, Initialisation v_0.
Output: Estimate $\widehat{X}$.

$t \leftarrow 0$
while stop condition not satisfied **do**
  $t \leftarrow t + 1$
  $v_{t-1} \leftarrow \frac{v_{t-1}}{\left\| v_{t-1} \right\|}$
  if $s_1 < n_1$ **then**
    $u_t \leftarrow \mathrm{HTP}(\mathrm{F}(v_{t-1}), b, s_1)$
  else
    $u_t \leftarrow \arg\min_x \left\| b - [\mathrm{F}(v_{t-1})]x \right\|^2$
  end if
  $u_t \leftarrow \frac{u_t}{\left\| u_t \right\|}$
  if $s_2 < n_2$ **then**
    $v_t \leftarrow \mathrm{HTP}(\mathrm{G}(u_t), \bar{b}, s_2)$
  else
    $v_t \leftarrow \arg\min_b \left\| \bar{b} - [\mathrm{G}(u_t)]b \right\|^2$
  end if
end while
return $\widehat{X} \leftarrow u_t v_t^*$

The Hard Thresholding Pursuit Algorithm is defined as follows:

Algorithm 2 HTP(A, b, s).

Input: Measurement matrix $A \in \mathbb{C}^{m \times n}$, measurement $b \in \mathbb{C}^m$, sparsity constraint $s \in \mathbb{N}$.
Output: $\hat{x} \in \mathbb{C}^n$.

$t \leftarrow 0$
while stop condition not satisfied **do**
  $t \leftarrow t + 1$
  $w = x_{t-1} + A^* (b - A x_{t-1})$
  $J \leftarrow \underset{J \subset [n],\ |J| = s}{\arg\max} \ \| w_J \|$
  $x_t \leftarrow \underset{x : \mathrm{supp}(x) \subset J}{\arg\min} \ \| Ax - b \|$
end while
return $\hat{x} \leftarrow x$

2.3 Initialization

As for many other non-convex algorithms (e.g., [6, 15]), the convergence properties of sparse power factorization depend crucially on the choice of the starting

point. In [6, 15], the starting point is chosen via a spectral initialization. That is, one chooses the leading left- and right-singular vectors of $\mathcal{A}^*(b)$ as the starting point, where $\mathcal{A}^*$ denotes the adjoint of $\mathcal{A}$. The general idea behind this approach is that, if the map $\mathcal{A}$ is random, under appropriate assumptions, the leading eigenvectors of $\mathbb{E}\left[\mathcal{A}^*(y)\right]$ are given by u and v. Indeed, under the assumptions in this manuscript (see Section 3), we have that $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = Id$. If $z = 0$, this implies that $\mathbb{E}\left[\mathcal{A}^*(y)\right] = \mathbb{E}\left[\mathcal{A}^*\mathcal{A}(uv^*)\right] = uv^*$. However, spectral initialization requires that the number of measurements is at the order of $\max\{n_1, n_2\}$, as otherwise the matrix $\mathcal{A}^*(y)$ will in general not concentrate around its expectation. Hence, spectral initialization will in general not be optimal as it does not take into account the sparsity of the vectors u and v. One way to incorporate the sparsity assumption would be to solve the sparse principal component analysis (SparsePCA) problem.

$$\begin{array}{ll} \max & \mathrm{Re}\,(\tilde{u}^*\mathcal{A}^*(b)\,\tilde{v}) \\ \text{subject to} & \|\tilde{u}\|_0 \le s_1,\ \|\tilde{u}\| = 1 \\ & \|\tilde{v}\|_0 \le s_2,\ \|\tilde{v}\| = 1, \end{array} \tag{2}$$

where $\|\cdot\|_0$ denotes the number of non-zero entries. As it was shown in [23, Proposition 2], Algorithm 1, if initialized by a solution of (2) is able to recover the solution u and v from a number of measurements at the order of $(s_1 + s_2)\max\left\{\frac{s_1}{n_1}, \frac{s_2}{n_2}\right\}$. However, the SparsePCA problem has been shown to be NP-hard [35]. Nevertheless, in the last 15 years, there has been a lot of research on the SparsePCA problem and, in particular, on tractable (i.e., polynomial time) algorithms, which yield good approximations to the true solution. Several computationally tractable algorithms have been proposed for solving (2), e.g., thresholdings algorithms [28], a general version of the power method [16] and semidefinite programs [8]. From the statistical perspective, a particular emphasis has been put for computationally efficient or at least tractable algorithms on the analysis of the single spike model [2, 9, 19]. These approaches, however, require that the number of samples scales with the square of the number of non-zero entries of the signal to estimate (up to log-factors). This raised the question whether there are fundamental barriers preventing the SparsePCA problem to be solved in polynomial time at a sampling rate close to the information theoretic limit. Indeed, it has been shown that an algorithm that achieves this would also allow for an algorithm which solves the k-clique problem in polynomial time [5, 36]. However, a widely believed conjecture in theoretical computer science states that this is not the case, which indicates that this approach will not be suited for initializing bilinear recovery problems either.

In this manuscript, we will analyze the following initialization algorithm, which is the one proposed in [23]. For a set $J_1 \subset [n]$, respectively $J_2 \subset [n_2]$ in the following, we will denote by Π_{J_1}, respectively Π_{J_2} the matrix, which projects a vector onto the components which belong to J_1, respectively J_2.

Algorithm 3 Algorithm 3 in [23].

Input: Operator $\mathcal{A}$, Measurement b, Sparsity Constraints s_1, s_2,
Output: Initial guess v_0 for $v \in \mathbb{C}^{n_2}$.

1: For all $i \in [n_1]$ let ξ_i be the ℓ_2-norm of the best s_2-sparse approximation of the ith row of the matrix $\mathcal{A}^*(b) \in \mathbb{C}^{n_1 \times n_2}$.
2: Let $\widehat{J_1} \subset [n_1]$ be the set of the s_2 largest elements in $\{\xi_1; \xi_2; \ldots; \xi_{n_1}\}$
3: Choose $\widehat{J_2}$ to contain the indices of the s_2 columns of $\Pi_{\widehat{J_1}} \mathcal{A}^*(b)$ largest in ℓ_2 norm, i.e.,

$$\widehat{J_2} := \underset{J \subset [n_2],\ |J| = s_2}{\arg\max} \left\| \Pi_{\widehat{J_1}} [\mathcal{A}^*(b)] \Pi_J \right\|_{\mathrm{F}}. \tag{3}$$

4: **return** v_0, the leading right singular vector of $\Pi_{\widehat{J_1}} [\mathcal{A}^*(b)] \Pi_{\widehat{J_2}}$.

3 Previous results

In the following, we will work with model (1), i.e., we observe

$$\operatorname{trace}\left(A_\ell^* u v^*\right) + z_\ell$$

where $u \in \mathbb{C}^{n_1}$ is s_1-sparse, $v \in \mathbb{C}^{n_2}$ is s_2-sparse, and $z \in \mathbb{C}^m$ is noise. As in [23], $\nu(z)$ will quantify the noise-to-signal ratio by

$$\nu(z) := \frac{\|z\|}{\|\mathcal{A}(uv^*)\|}. \tag{4}$$

In our results, we will make no assumptions on the noise z except that $\nu(z)$ needs to be smaller than a certain threshold. (Our results will hold for adversarial noise, i.e., the recovery guarantees will hold uniformly for all $z \in \mathbb{C}^m$, such that $\nu(z)$ satisfies this threshold.) For our analysis, $\mathcal{A}$ will be a Gaussian linear operator, that is, all the entries of the matrices $A_1, \ldots, A_m$ are independent with distribution $\mathcal{CN}\left(0, \frac{1}{m}\right)$. (Here, a complex-valued random variable X has distribution $\mathcal{CN}\left(0, \frac{1}{m}\right)$ if its real and complex part are independent Gaussians with expectation 0 and variance $\sqrt{\frac{\sigma}{2}}$.)

In [23], the authors derived that Algorithm 1, if initialized by Algorithm 3, is able to recover both u and v (up to scale ambiguity), if both u and v belong to a certain restricted class of signals. More precisely, they proved the following result.

Theorem 1 ([23, see Theorem 1 and 4]) *Assume that $\mathcal{A}: \mathbb{C}^{n_1 \times n_2} \longrightarrow \mathbb{C}^m$ is a Gaussian linear operator as described above. Let $b = \mathcal{A}(uv^*) + z$, where u is s_1-sparse and v is s_2-sparse. Suppose that $\|u\|_\infty \geq 0.78\|u\|$, $\|v\|_\infty \geq 0.78\|v\|$, and that the noise level satisfies $\nu(z) \leq 0.04$. Then, with probability exceeding*

$1 - \exp(-c_1 m)$, *the output of the Algorithm 1, initialized by Algorithm 3, converges linearly to* uv^* *provided that*

$$m \geq c_2 (s_1 + s_2) \log \left(\max \left\{ \frac{n_1}{s_1}, \frac{n_2}{s_2} \right\} \right),$$

where $c_1, c_2 > 0$ *are absolute constants.*

Note that in order to apply Theorem 1 to signals u and v, one needs to require that more than half of the mass of u and v are located in one single entry, which is a severe restriction, which can be prohibitive for many applications. Our goal in the following will be to considerably relax this assumption by slightly increasing the amount of required measurements. We will relax this assumption in two different ways: On the one hand, we will show that one can replace 0.78 by an arbitrary small constant that will then show up in the number of measurements. On the other hand, we generalize the result to the case that a significant portion of mass of u is concentrated on a small number of entries k, rather than just one of them.

4 Main result

In this section, we will state the main result of this article, Theorem 2. For that, we need to define the quantity

$$\|x\|_{[k]} := \max_{I \subset [n_1],\ |I|=k} \left(\sum_{i \in I} |x_i|^2 \right)^{1/2} = \left(\sum_{i=1}^{k} \left(x_i^*\right)^2 \right)^{1/2},$$

for any $x \in \mathbb{C}^{n_1}$, where $\left(x_i^*\right)_{i=1}^{n_1}$ denotes the non-increasing rearrangement of $(|x_i|)_{i=1}^{n_1}$. (This is in fact a norm, cf. [29].) Our main requirement on the vector u will be that a significant amount of its mass is located in the largest k entries, i.e., that $\frac{\|u\|_{[k]}}{\|u\|}$ is large enough.

Theorem 2 *Let* $k \in [n_1]$ *and* $0 < \xi < 1, 0 < \mu < 1$. *Then, there are absolute constants* $C_1, C_2, C_3 > 0$ *such that if*

$$m \geq C_1 \max \left\{ \frac{1}{\xi^4 \mu^4}, \frac{k}{\xi^2} \right\} (s_1 + s_2) \log \left(\max \left\{ \frac{n_1}{s_1}, \frac{n_2}{s_2} \right\} \right), \tag{5}$$

then with probability at least $1 - \exp(-C_2 m)$, *the following holds.*

For all s_1*-sparse* $u \in \mathbb{C}^{n_1}$ *with* $\|u\|_{[k]} \geq \xi \|u\|$, *all* s_2*-sparse* $v \in \mathbb{C}^{n_2}$ *with* $\|v\|_\infty \geq \mu \|v\|$, *and all* $z \in \mathbb{C}^m$ *with* $\nu(z) \leq C_3 \min \left\{ \xi^2 \mu^2; \frac{\xi}{\sqrt{k}} \right\}$, *the iterates* $\{X_t\}_{t \in \mathbb{N}}$ *generated by applying Algorithm 1, initialized by Algorithm 3, satisfy*

$$\limsup_{t \to \infty} \frac{\|X_t - uv^*\|_{\mathrm{F}}}{\|uv^*\|_{\mathrm{F}}} \leq 8.3\nu.$$

Furthermore, the convergence is linear, i.e., for all $t \gtrsim \log\left(\frac{1}{\varepsilon}\right)$, *we have that*

$$\frac{\|X_t - uv^*\|_{\mathrm{F}}}{\|uv^*\|_{\mathrm{F}}} \leq 8.3\nu + \varepsilon. \tag{6}$$

In the following, we will discuss some important special cases of Theorem 2.

- **Peaky signals:** In [23], the authors discuss recovery guarantees for signals u and v with $\frac{\|u\|_\infty}{\|u\|}$ and $\frac{\|v\|_\infty}{\|v\|}$, both bounded below by an absolute constant $\mu \approx 0.78$. The case $k = 1$ of our theorem yields a direct improvement of this result in the sense that μ can be chosen arbitrarily small with the number of required measurements only increasing by a factor of order μ^{-8}. Hence, even when this constant decays logarithmically in the dimension, the required number of measurements will only increase by logarithmic factors.
- **Signals with multiple large entries:** When one of the input signals has multiple large entries, using the $\|\cdot\|_{[k]}$ norm improves upon the resulting guarantee as compared to the scenario just discussed. As an example, assume that $s_1 = s_2 = s$, that u and v are normalized with $\|v\|_\infty \geq c_1 s^{-1/8}$, and that $k = c_2 s^{1/2}$ of the entries of u are of absolute value at least $c_3 s^{-1/4}$. Then, $\|u\|_{[k]} \geq \sqrt{c_2} c_3$. Using Theorem 2, we obtain that the vectors u and v can be recovered if the number of measurements is on the order of $s^{3/2}$, thus below the order of s^2 that has been established for arbitrary sparse signals in [25] (cf. next item). In contrast, applying Theorem 2 with $k = 1$ would yield that the number of measurements would have to be on the order of $s^{5/2}$, which is worse than the state-of-the-art.
- **Arbitrary sparse signals:** Applying Theorem 2 to non-peaky signals yields suboptimal results. Indeed, let $u \in \mathbb{C}^{n_1}$ s_1-sparse and $v \in \mathbb{C}^{n_2}$ s_2-sparse be generic vectors. Observe that $\|v\|_\infty \asymp \frac{1}{\sqrt{s_2}}\|v\|$. Consequently, Theorem 2 applied with $\xi = 1$, $k = s_1$, and $\mu = \frac{1}{\sqrt{s_2}}$ yields that with high probability a generic s_1-sparse u and a generic s_2-sparse v can be recovered from $y = \mathcal{A}(uv^*) + z$, if the number of measurements satisfies

$$m \geq C \max\left\{s_1; s_2^2\right\}(s_1 + s_2) \log\left(\max\left\{\frac{n_1}{s_1}, \frac{n_2}{s_2}\right\}\right),$$

and if the noise level ν is on the order of $\mathcal{O}\left(\max\left\{\frac{1}{s_2}; \frac{1}{\sqrt{s_1}}\right\}\right)$. Previous results (see, e.g., [25]), in contrast, require $m \geq C \max\left\{s_1^2; s_2^2\right\} \log\left(\max\left\{\frac{n_1}{s_1}, \frac{n_2}{s_2}\right\}\right)$ samples.

Remark 1 The peakiness assumptions in Theorem 2 may seem arbitrary at first sight but in certain applications they are reasonable. Namely, when u is the signal transmitted via a wireless channel and v is the unknown vector of channel parameters, it is natural to assume that v has a large entry, as the direct path will always carry most of the energy. The signal u can be modified by the sender, so some large entries can be artificially introduced. In this regard, being able to consider multiple entries of comparable size is of advantage as adding a single very large entry will result in a dramatic increase of the peak-to-average power ratio.

5 Proofs

5.1 Technical tools

The goal of this section is to prove Theorem 2. We will start by recalling the following variant of the well-known restricted isometry property.

Definition 1 (see [23]) A linear operator $\mathcal{A}$ has the (s_1, s_2, r)-restricted isometry property with constant $\delta > 0$ if

$$(1-\delta)\,\|X\|_F^2 \leq \|\mathcal{A}(X)\|^2 \leq (1+\delta)\,\|X\|_F^2 \tag{7}$$

for all matrices $X \in \mathbb{C}^{n_1 \times n_2}$ of rank at most r with at most s_1 non-zero rows and at most s_2 non-zero columns.

The following lemma tells us that this property holds with high probability for a number of measurements close to the information-theoretic limit.

Lemma 1 (See, e.g., Theorem 2 in [23]) *There are absolute constants $c_1, c_2 > 0$, such that if*

$$m \geq \frac{c_1}{\delta^2} r\,(s_1+s_2) \log\left(\max\left\{\frac{n_1}{s_1}, \frac{n_2}{s_2}\right\}\right), \tag{8}$$

for some $\delta > 0$, then with probability at least $1 - \exp(-c_2 m)$ $\mathcal{A}$ has the (s_1, s_2, r)-restricted isometry property with restricted isometry constant δ.

As in [23, Lemma 6], we will need the following quantity, which depends on δ and ν.

$$\omega_{\sup} := \sup\left\{\omega \in [0, \tfrac{\pi}{2}) : \sin(\omega) \geq C_\delta[\delta\tan(\omega) + (1+\delta)\nu\sec(\omega)]\right\} \tag{9}$$

Here, the constant C_δ is given by the expression

$$C_\delta = 1.01\frac{\sqrt{\frac{2}{1-\delta^2}} + \frac{1}{1-\delta}}{1 - \sqrt{\frac{2}{1-\delta^2}}\delta},$$

as it can be seen by an inspection of the proof of Lemma 3 in [23]. The precise value of C_δ will not be important in the following, we will only use that $2 \leq C_\delta \leq 5$ for $\delta \leq 0.04$.

A simple estimate for $\omega_{\sup}$ is given by the following lemma.

Lemma 2 *Assume that $0 < \delta \leq 0.04$ and $\nu \leq 0.04$. Then, it holds that*

$$\frac{1}{2} \leq \sin(\omega_{\sup}) \leq 1.$$

Proof We observe that in order to show the claim, it is enough to verify that $\omega = \arcsin\frac{1}{2}$ fulfills the inequality in (9). Indeed, using $\cos\omega = \sqrt{\frac{3}{4}}$ and $C_\delta \leq 5$, we obtain that

$$C_\delta\left[\delta\tan\left(\arcsin\frac{1}{2}\right)+(1+\delta)\,\nu\sec\left(\arcsin\frac{1}{2}\right)\right] \leq C_\delta\left[0.04\frac{1/2}{\sqrt{3/4}}+\frac{1.04\cdot 0.04}{\sqrt{3/4}}\right] \leq \frac{1}{2}.$$

□

The quantity ω_{sup} controls the maximal angle between the initialization v_0 and the ground truth v such that the sparse power factorization is guaranteed to converge as captured by the following theorem.

Theorem 3 (Theorem 3 in [23]) *Assume that:*

1) $\mathcal{A}$ *has the* $(3s_1, 3s_2, 2)$*-RIP with isometry constant* $\delta \leq 0.08$.
2) $\nu \leq 0.08$.
3) *The initialization* v_0 *satisfies* $\sin(\angle(v_0, v)) < \sin\left(\omega_{\text{sup}}\right)$.

Then, the iterates $\{X_t\}_{t\in\mathbb{N}}$ *generated by Algorithm 1, initialized via Algorithm 3, satisfy*

$$\limsup_{t\to\infty}\frac{\|X_t - uv^*\|_{\text{F}}}{\|uv^*\|_{\text{F}}} \leq 8.3\nu.$$

Furthermore, the convergence is linear in the sense of (6).

Thus, it remains to verify that the initialization satisfies $\sin(\angle(v_0, v)) < \sin\left(\omega_{\text{sup}}\right)$. The following lemma gives an upper bound on $\sin(\angle(v_0, v))$.

Lemma 3 (Lemma 8 in [23]) *Assume that the* $(3s_1, 3s_2, 2)$*-restricted isometry property holds for some constant* $\delta > 0$. *Furthermore, assume that* $\|u\| = \|v\| = 1$. *Let* $\widehat{J}_1 \subseteq [n_1]$ *and* $\widehat{J}_2 \subseteq [n_2]$ *denote the output resulting from Algorithm 3. Denote by* v_0 *the leading right-singular vector of* $\Pi_{\widehat{J}_1}[\mathcal{A}^*(b)]\Pi_{\widehat{J}_2}$. *Then, it holds that*

$$\sin(\angle(v_0, v)) \leq \frac{\left\|\Pi_{\widehat{J}_1}u\right\|\left\|\Pi^{\perp}_{\widehat{J}_2}v\right\| + (\delta+\nu+\delta\nu)}{\left\|\Pi_{\widehat{J}_1}u\right\| - (\delta+\nu+\delta\nu)}. \tag{10}$$

Furthermore, we will need the following two lemmas for our proof.

Lemma 4 (Lemma 10 in [23]) *Let* u *and* v *be as in Lemma 5 and assume that the measurement operator* $\mathcal{A}$ *satisfies the* $(3s_1, 3s_2, 2)$*-restricted isometry property with constant* δ. *Recall that* $\widehat{J}_1 \subset [n_1]$ *is the support estimate for* v_0 *given by the initialization algorithm 3. Define*

$$\widetilde{J}_1 := \left\{j \in [n_1] : \ |u_j| \geq 2\,(\delta+\nu+\delta\nu)\right\}. \tag{11}$$

Then, we have that $\widetilde{J}_1 \subset \widehat{J}_1$.

Lemma 5 *Assume that $\mathcal{A}$ has the $(3s_1, 3s_2, 2)$-restricted isometry property with isometry constant $\delta > 0$ and assume that u, respectively v, are s_1-sparse, respectively s_2-sparse, and satisfy $\|u\| = \|v\| = 1$. Let $\widetilde{J}_1$ be defined as in (11). Then, it holds that*

$$\left\|\Pi_{\widehat{J}_1} u\right\| \left\|\Pi_{\widehat{J}_2} v\right\| \geq \left\|\Pi_{\widetilde{J}_1} u\right\| \|v\|_\infty - 2\left(\delta + \nu + \delta\nu\right).$$

Lemma 5 is actually a slight generalization of what has been shown in [23, p. 1685]. For completeness, we have included a proof in the Appendix, which closely follows the proof in [23].

5.2 Proof of our main result

We will now piece together these ingredients to obtain a sufficient condition; in the remainder of the section, we will then show that the condition holds in our measurement setup. First note that in order to apply Theorem 3, we need to check that $\sin(\angle(v_0, v)) < \sin\left(\omega_{\text{sup}}\right)$ is satisfied. By Lemma 3, it is sufficient to show that the right-hand side of inequality (10) is strictly smaller than $\sin\left(\omega_{\text{sup}}\right)$. Combining this with the equality $\left\|\Pi_{\widehat{J}_2}^{\perp} v\right\| = \sqrt{1 - \left\|\Pi_{\widehat{J}_2} v\right\|^2}$, we obtain the sufficient condition

$$\left\|\Pi_{\widehat{J}_1} u\right\| \sqrt{1 - \left\|\Pi_{\widehat{J}_2} v\right\|^2} < \sin\left(\omega_{\text{sup}}\right) \left(\left\|\Pi_{\widehat{J}_1} u\right\| - (\delta + \nu + \delta\nu)\right) - (\delta + \nu + \delta\nu)$$

Further manipulations yield that this is equivalent to

$$\begin{aligned} \left\|\Pi_{\widehat{J}_1} u\right\|^2 < & \left(\sin\left(\omega_{\text{sup}}\right) \left\|\Pi_{\widehat{J}_1} u\right\| - \left(1 + \sin\left(\omega_{\text{sup}}\right)\right)(\delta + \nu + \delta\nu)\right)^2 \\ + & \left\|\Pi_{\widehat{J}_1} u\right\|^2 \left\|\Pi_{\widehat{J}_2} v\right\|^2. \end{aligned} \tag{12}$$

Hence, in the following, our goal will be to verify (12). We already noticed that the angle ω_{sup} measures how much the vector v_0 given by the initializiation has to be aligned with the ground truth v in order for the sparse power factorization to converge. Consequently, it is natural to expect that the smaller the constant δ and the noise-to-signal ratio ν, the less the initializiation vector has to be aligned with the ground truth, i.e., the larger ω_{sup} can be. This fact is captured by the following lemma.

Lemma 6 *Let $0 < \delta \leq 0.04$ and $\nu \leq 0.04$. Then, it holds that*

$$\sin(\omega_{\text{sup}}) \geq 1 - C_\delta^2 \left(\delta + 2\delta\nu + 2\nu\right)^2.$$

Proof We will first show that it holds that

$$\sin\left(\omega_{\text{sup}}\right) = C_\delta \left[\delta \tan\left(\omega_{\text{sup}}\right) + (1 + \delta)\, \nu \sec\left(\omega_{\text{sup}}\right)\right]. \tag{13}$$

For that, we define the set

$$\Omega := \{\omega \in [0, \pi/2) : \ \sin(\omega) \geq C_\delta[\delta \tan(\omega) + (1 + \delta)\nu \sec(\omega)]\}$$

and note that $\omega_{\sup} = \sup \Omega$. (Note that from Lemma 2 and its proof, it follows that Ω is nonempty.) We observe that

$$\lim_{\omega \uparrow \pi/2} C_\delta \left[\delta \tan(\omega) + (1+\delta)\, \nu \sec(\omega)\right] = +\infty. \tag{14}$$

Hence, Ω is a compact set. In particular, $\omega_{\sup}$ is contained in this set. Now assume by contradiction that

$$\sin\left(\omega_{\sup}\right) > C_\delta[\delta \tan(\omega_{\sup}) + (1+\delta)\nu \sec(\omega_{\sup})].$$

Because of (14) and as the sine is bounded in $[0, \pi/2)$, this implies by continuity that there must be an $\omega \in [0, \pi/2)$ such that $\omega > \omega_{\sup}$ and $\sin(\omega) = C_\delta[\delta \tan(\omega) + (1+\delta)\nu \sec(\omega)$. Hence, $\omega \in \Omega$, which is a contradiction to the definition of $\omega_{\sup}$. Hence, we have proven equation (13). Using trigonometric identities, we obtain from (13) that

$$\sin\left(\omega_{\sup}\right) = C_\delta \left[\delta \frac{\sin\left(\omega_{\sup}\right)}{\sqrt{1-\sin\left(\omega_{\sup}\right)^2}} + (1+\delta)\,\nu \frac{1}{\sqrt{1-\sin\left(\omega_{\sup}\right)^2}}\right].$$

Lemma 2 implies that

$$\sin\left(\omega_{\sup}\right) \le \frac{\sin\left(\omega_{\sup}\right)}{\sqrt{1-\sin\left(\omega_{\sup}\right)^2}} C_\delta \left(\delta + 2\,(1+\delta)\,\nu\right).$$

Rearranging terms yields that

$$\sin\left(\omega_{\sup}\right) \ge \sqrt{1 - C_\delta^2\,(\delta + 2\delta\nu + 2\nu)^2}.$$

The claim follows then using the fact that $\sqrt{x} \ge x$ for all $x \in [0, 1]$. □

With these preliminary lemmas, we can now prove the following proposition, which is a slightly more general form of Theorem 2.

Proposition 1 *There are absolute constants $c_1, c_2, c_3 > 0$ such that if*

$$m \ge c_1 \delta^{-2}\,(s_1 + s_2) \log\left(\max\left\{\frac{n_1}{s_1}, \frac{n_2}{s_2}\right\}\right), \tag{15}$$

for some $0 < \delta < 0.01$, then with probability at least $1 - \exp(-c_2 m)$, the following statement holds uniformly for all s_1-sparse $u \in \mathbb{C}^{n_1}$, s_2-sparse $v \in \mathbb{C}^{n_2}$, and $z \in \mathbb{C}^m$ such that $\|u\| = \|v\| = 1$ and $\nu(z) \le 0.01$:
Let the measurements be given by $b = \mathcal{A}(uv^*) + z$ for $\mathcal{A}$ Gaussian as above and let $\widetilde{J}_1$ be defined by

$$\widetilde{J}_1 := \left\{j \in [n_1] : \ |u_j| \ge M_{\delta,\nu}\right\}, \tag{16}$$

where

$$M_{\delta,\nu} := 2\,(\delta + \nu + \delta\nu).$$

Then, whenever

$$\left\|\Pi_{\widetilde{J}_1} u\right\| \|v\|_\infty > c_3 \sqrt{M_{\delta,\nu}}, \tag{17}$$

the iterates $\{X_t\}_{t\in\mathbb{N}}$ *generated by Algorithm 1 initialized via Algorithm 3 satisfy*

$$\limsup_{t\to\infty} \|X_t - uv^*\|_{\mathrm{F}} \leq 8.3\nu.$$

Furthermore, the convergence is linear in the sense of (6).

Proof of Proposition 1 Assumption (15) and Lemma 1 yield that with probability at least $1 - \exp(-cm)$, the $(3s_1, 3s_2, 2)$-restricted isometry property holds with constant δ. For the remainder of the proof, we will consider the event that the restricted isometry property holds for such δ. We obtain

$$\left\|\Pi_{\widetilde{J}_1} u\right\| \|v\|_\infty \geq \left(\sqrt{C_\delta^2 + 1} + 1\right) \sqrt{M_{\delta,\nu}}$$

from $2 \leq C_\delta \leq 5$ and by choosing the constant c_3 in assumption (17) large enough. Combining this with Lemma 5, we obtain that

$$\begin{aligned} \left\|\Pi_{\widehat{J}_1} u\right\| \left\|\Pi_{\widehat{J}_2} v\right\| &\geq \left\|\Pi_{\widetilde{J}_1} u\right\| \|v\|_\infty - M_{\delta,\nu}. \\ &> \sqrt{\left(C_\delta^2 + 1\right) M_{\delta,\nu}}, \end{aligned} \tag{18}$$

where we used that $\sqrt{x} \geq x$ for all $x \in [0, 1]$. This yields a lower bound for the second summand of the right-hand side of (12). To bound the first summand, we estimate

$$\begin{aligned} &\sin(\omega_{\sup})\|\Pi_{\widehat{J}_1} u\| - \left(\sin(\omega_{\sup}) + 1\right)(\delta + \nu + \delta\nu) \\ &\geq \left(1 - C_\delta^2 (\delta + 2\nu + 2\delta\nu)^2\right) \|\Pi_{\widehat{J}_1} u\| - 2(\delta + \nu + \delta\nu) \\ &\geq \|\Pi_{\widehat{J}_1} u\| - C_\delta^2 (\delta + 2\nu + 2\delta\nu)^2 - 2(\delta + \nu + \delta\nu) \\ &\geq \|\Pi_{\widehat{J}_1} u\| - \left(C_\delta^2 + 1\right) M_{\delta,\nu} \\ &\geq 0. \end{aligned} \tag{19}$$

In the first line, we used Lemma 6 and the fact that $\sin(\omega_{\sup}) \leq 1$. The second line is due to $\|\Pi_{\widehat{J}_1} u\| \leq 1$ and the third inequality is due to $\delta \geq 0$, $\nu \geq 0$. In order to verify the last inequality, it is enough to observe that due to Lemma 4 and due to assumption (17) with c_3 large enough

$$\|\Pi_{\widehat{J}_1} u\| \geq \|\Pi_{\widetilde{J}_1} u\| \geq \|\Pi_{\widetilde{J}_1} u\| \|v\|_\infty \geq \left(C_\delta^2 + 1\right) M_{\delta,\nu},$$

where the last inequality uses that $C_\delta \leq 5$ and $0 \leq \delta, \nu \leq 0.01$. Hence, by squaring (19), we obtain that

$$\begin{aligned} &\left(\sin(\omega_{\sup})\|\Pi_{\widehat{J}_1} u\| - \left(\sin(\omega_{\sup}) + 1\right)(\delta + \nu + \delta\nu)\right)^2 \\ &\geq \left(\|\Pi_{\widehat{J}_1} u\| - \tfrac{1}{2}\left(C_\delta^2 + 1\right) M_{\delta,\nu}\right)^2 \\ &\geq \|\Pi_{\widehat{J}_1} u\|^2 - \left(C_\delta^2 + 1\right) M_{\delta,\nu} \|\Pi_{\widehat{J}_1} u\| \\ &\geq \|\Pi_{\widehat{J}_1} u\|^2 - \left(C_\delta^2 + 1\right) M_{\delta,\nu}, \end{aligned} \tag{20}$$

where in the last line, we again used that $\|\Pi_{\widehat{J}_1} u\| \leq 1$. Together with (18), this yields (12), as desired. □

Finally, we will deduce Theorem 2 from Proposition 1.

Proof of Theorem 2 We will prove this result by applying Proposition 1 with

$$\delta = \min\left\{\frac{\xi}{6\sqrt{2k}}; \frac{\xi^2\mu^2}{8c_3^2}\right\}. \tag{21}$$

Let $u \in \mathbb{C}^{n_1}$ s_1-sparse, $v \in \mathbb{C}^{n_2}$ s_2-sparse, and $z \in \mathbb{C}^m$ such that the assumptions of Theorem 2 are satisfied. Without loss of generality, we may assume in the following that $\|u\| = \|v\| = 1$. First, we note that invoking $\delta, \nu < 0.01$ and potentially decreasing the size of C_3, we have that

$$2\left(\delta + \nu\left(z\right) + \delta\nu\left(z\right)\right) < 2\left(\delta + 2\nu\left(z\right)\right) \leq \frac{\xi}{\sqrt{2k}}.$$

Hence, we obtain that

$$\check{J}_1 := \left\{j \in [n_1] : \ |u_j| \geq \frac{\xi}{\sqrt{2k}}\right\} \subset \widetilde{J}_1, \tag{22}$$

where $\widetilde{J}_1$ is the set defined in (16). Note that

$$\sum_{i\in[k]\setminus\check{J}_1} \left(u_i^*\right)^2 < \sum_{i\in[k]\setminus\check{J}_1} \frac{\xi^2}{2k} \leq \frac{\xi^2}{2},$$

where in the first inequality, we have used that $u_i^* < \frac{\xi}{\sqrt{2k}}$ for all $i \in [k]\setminus\check{J}_1$. By the assumption $\|u\|_{[k]} \geq \xi$, this yields that $\sum_{i\in[k]\cap\check{J}_1}\left(u_i^*\right)^2 \geq \frac{\xi^2}{2}$, which in turn implies that $\|\Pi_{\check{J}_1}u\| \geq \frac{\xi}{\sqrt{2}}$. By the inclusion (22), we obtain that $\|\Pi_{\widetilde{J}_1}u\| \geq \frac{\xi}{\sqrt{2}}$. Hence, using the assumption $\|v\|_\infty \geq \mu$, our choice of δ, the assumption on the noise level $\nu\left(z\right)$ and potentially again decreasing the value of the constant C_3, we obtain that

$$\|\Pi_{\widetilde{J}_1}u\|\|v\|_\infty \geq \frac{\xi\mu}{\sqrt{2}} \geq c_3\sqrt{M_{\delta,\nu}}.$$

This shows that (17) is satisfied. Hence, we can apply Proposition 1 and by inserting our choice of δ into (15), so choosing the constant C_1 large enough, we obtain the main result. □

6 Outlook

We see many interesting directions for follow-up work. Most importantly, it remains to explore whether additional constraints on the signals to be recovered are truly necessary (cf. our discussion on to SparsePCA in Section 2.3). Even if this is the case, there is substantial room for improvement with respect to the noise-dependence of the recovery results. A direction to proceed could be to consider stochastic noise models instead of adversarial noise. Also in this work, we exclusively considered operators $\mathcal{A}$ constructed using Gaussian matrices. However, in many applications of interest, the

measurement matrices possess a significantly reduced amount of randomness. For example, in blind deconvolution, one typically encounters rank-one measurements. That is, the restricted isometry property as used in this paper does not hold. Thus, one needs additional insight to study whether there exists a computationally tractable initialization procedure at a near-optimal sampling rate. First steps in this direction were taken in [20, 22], but a lot of questions remain open.

Acknowledgements The authors want to thank Yoram Bresler and Kiryung Lee for helpful discussions. Furthermore, we would like to thank the referees for the ir careful reading and their helpful suggestions, which improved the manuscript.

Appendix: Proof of Lemma 5

For the proof of Lemma 5, we will use the following result.

Lemma 7 (Lemma 17 and Lemma 18 in [23]) *Assume that the* $(3s_1, 3s_2, 2)$*-restricted isometry property is fulfilled for some restricted isometry constant* $\delta > 0$*. Assume that the cardinality of* $\widetilde{J}_1 \subseteq [n_1]$*, respectively* $\widetilde{J}_2 \subseteq [n_2]$ *is at most* $2s_1$*, respectively* $2s_2$*. Then, whenever* $u \in \mathbb{C}^{n_1}$ *is at most* $2s_1$*-sparse and* $v \in \mathbb{C}^{n_2}$ *is at most* $2s_2$*-sparse, we have that*

$$\|\Pi_{\widetilde{J}_1}[(\mathcal{A}^*\mathcal{A} - I)(uv^*)]\Pi_{\widetilde{J}_2}\| \leq \delta\|uv^*\|_{\mathrm{F}}.$$

Furthermore for all $z \in \mathbb{C}^n$ *and for all* $\widetilde{J}_1 \subseteq [n_1]$*, respectively* $\widetilde{J}_2 \subseteq [n_2]$*, with cardinality at most* s_1*, respectively* s_2*, we have that*

$$\|\Pi_{\widetilde{J}_1}[\mathcal{A}^*(z)]\Pi_{\widetilde{J}_2}\| \leq \sqrt{1+\delta}\|z\|_{\ell_2}.$$

Proof of Lemma 5 Recall that $b = \mathcal{A}(X) + z$ and define k_1 and k_2 by

$$\begin{aligned} k_1 &:= \operatorname*{arg\,max}_{k\in[n_2]} |v_k| \\ k_2 &:= \operatorname*{arg\,max}_{k\in[n_2]} \left\|\Pi_{\widehat{J}_1}[\mathcal{A}^*(b)]\Pi_{\{k\}}\right\|_{\mathrm{F}}. \end{aligned} \tag{23}$$

The starting point of our proof is the observation that

$$\left\|\Pi_{\widehat{J}_1}[\mathcal{A}^*(b)]\Pi_{\{k_2\}}\right\|_{\mathrm{F}} \geq \left\|\Pi_{\widehat{J}_1}[\mathcal{A}^*(b)]\Pi_{\{k_1\}}\right\|_{\mathrm{F}} \geq \left\|\Pi_{\widetilde{J}_1}[\mathcal{A}^*(b)]\Pi_{\{k_1\}}\right\|_{\mathrm{F}}, \tag{24}$$

where the first inequality is due to the definition of k_2 and the second one follows from $\widetilde{J}_1 \subset \widehat{J}_1$, which is due to Lemma 4. The right-hand side of the inequality chain can be estimated from below by

$$\begin{aligned} &\left\|\Pi_{\widetilde{J}_1}[\mathcal{A}^*(b)]\Pi_{\{k_1\}}\right\|_{\mathrm{F}} \\ \geq\ &\left\|\Pi_{\widetilde{J}_1}uv^*\Pi_{\{k_1\}}\right\|_{\mathrm{F}} - \left\|\Pi_{\widetilde{J}_1}\left[(\mathcal{A}^*\mathcal{A} - I)(uv^*)\right]\Pi_{\{k_1\}}\right\|_{\mathrm{F}} - \left\|\Pi_{\widetilde{J}_1}\mathcal{A}^*(z)\,\Pi_{\{k_1\}}\right\|_{\mathrm{F}} \\ \geq\ &\left\|\Pi_{\widetilde{J}_1}uv^*\Pi_{\{k_1\}}\right\|_{\mathrm{F}} - \left(\delta\|uv^*\|_F + \sqrt{1+\delta}\|z\|\right) \\ \geq\ &\left\|\Pi_{\widetilde{J}_1}u\right\|\|v\|_\infty - (\delta + \nu + \delta\nu)\,. \end{aligned} \tag{25}$$

In the first inequality, we used $b = \mathcal{A}(uv^*) + z$ and the triangle inequality. The second inequality follows from Lemma 7. The last line follows from $\|uv^*\|_F = 1$ and $\|z\| = \nu$. Next, we will estimate the left-hand side of (24) by

$$\begin{aligned} & \left\| \Pi_{\widehat{J}_1} [\mathcal{A}^*(b)] \Pi_{\{k_2\}} \right\|_{\mathrm{F}} \\ \leq & \left\| \Pi_{\widehat{J}_1} uv^* \Pi_{\{k_2\}} \right\|_{\mathrm{F}} + \left(\delta \|uv^*\|_F + \sqrt{1+\delta} \|z\| \right) \\ \leq & \left\| \Pi_{\widehat{J}_1} u \right\| \left\| \Pi_{\{k_2\}} v \right\| + (\delta + \nu + \delta\nu) \\ \leq & \left\| \Pi_{\widehat{J}_1} u \right\| \left\| \Pi_{\widehat{J}_2} v \right\| + (\delta + \nu + \delta\nu) \,. \end{aligned} \tag{26}$$

The first two lines are obtained by an analogous reasoning as for (25). The last line is due to $\{k_2\} \subset \widehat{J}_2$, which is a consequence of the definition of $\widehat{J}_2$ (3) and the definition of $\{k_2\}$ (23). We finish the proof by combining the inequality chains (24), (25), and (26). □

References

1. Ahmed, A., Recht, B., Romberg, J.: Blind deconvolution using convex programming. IEEE Trans. Inform. Theory **60**(3), 1711–1732 (2014)
2. Amini, A.A., Wainwright, M.J.: High-dimensional analysis of semidefinite relaxations for sparse principal components. Ann. Stat. **37**(5B), 2877–2921 (2009)
3. Bahmani, S., Romberg, J.: Near-optimal estimation of simultaneously sparse and low-rank matrices from nested linear measurements. Inf. Inference **5**(3), 331–351 (2016)
4. Bahmani, S., Romberg, J.: Solving equations of random convex functions via anchored regression. arXiv:1702.05327 (2017)
5. Berthet, Q., Rigollet, P.: Optimal detection of sparse principal components in high dimension. Ann. Stat. **41**(4), 1780–1815 (2013)
6. Candes, E.J., Li, X., Soltanolkotabi, M.: Phase retrieval via Wirtinger flow: theory and algorithms. IEEE Trans. Inform. Theory **61**(4), 1985–2007 (2015)
7. Candès, E.J., Romberg, J., Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inform. Theory **52**(2), 489–509 (2006)
8. d'Aspremont, A., Bach, F., Ghaoui, L.E.: Optimal solutions for sparse principal component analysis. J. Mach. Learn. Res. **9**, 1269–1294 (2008)
9. Deshpande, Y., Montanari, A.: Sparse PCA via covariance thresholding. In: Advances in Neural Information Processing Systems, pp. 334–342 (2014)
10. Fornasier, M., Maly, J., Naumova, V.: At-las _ {2, 1}: A multi-penalty approach to compressed sensing of low-rank matrices with sparse decompositions. arXiv:1801.06240 (2018)
11. Foucart, S.: Hard thresholding pursuit: an algorithm for compressive sensing. SIAM J. Numer. Anal. **49**(6), 2543–2563 (2011)
12. Geppert, J.A., Krahmer, F., Stöger, D.: Refined performance guarantees for sparse power factorization. In: 2017 International Conference on Sampling Theory and Applications (SampTA), pp. 509–513. IEEE (2017)
13. Haykin, S.: Blind Deconvolution. Prentice Hall, New Jersey (1994)
14. Iwen, M., Viswanathan, A., Wang, Y.: Robust sparse phase retrieval made easy. Appl. Comput. Harmon. Anal. **42**(1), 135–142 (2017)
15. Jain, P., Netrapalli, P., Sanghavi, S.: Low-rank matrix completion using alternating minimization. In: Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13, pp. 665–674. ACM, New York (2013)
16. Journée, M., Nesterov, Y., Richtárik, P., Sepulchre, R.: Generalized power method for sparse principal component analysis. J. Mach. Learn. Res. **11**, 517–553 (2010)

17. Jung, P., Krahmer, F., Stöger, D.: Blind demixing and deconvolution at near-optimal rate. IEEE Trans. Inform. Theory **64**(2), 704–727 (2018)
18. Kech, M., Krahmer, F.: Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems. SIAM J. Appl. Alg. Geom. **1**(1), 20–37 (2017). https://doi.org/10.1137/16M1067469
19. Krauthgamer, R., Nadler, B., Vilenchik, D.: Do semidefinite relaxations solve sparse PCA up to the information limit. Ann. Statist. **43**(3), 1300–1322 (2015)
20. Lee, K., Junge, M.: Rip-like properties in subsampled blind deconvolution. arXiv:1511.06146 (2015)
21. Lee, K., Krahmer, F., Romberg, J.: Spectral methods for passive imaging: non-asymptotic performance and robustness. arXiv:1708.04343 (2017)
22. Lee, K., Li, Y., Junge, M., Bresler, Y.: Blind recovery of sparse signals from subsampled convolution. IEEE Trans. Inform. Theory **63**(2), 802–821 (2017)
23. Lee, K., Wu, Y., Bresler, Y.: Near optimal compressed sensing of a class of sparse low-rank matrices via sparse power factorization. IEEE Trans. Inform Theory (2017)
24. Li, X., Ling, S., Strohmer, T., Wei, K.: Rapid, robust, and reliable blind deconvolution via nonconvex optimization. arXiv:1606.04933 (2016)
25. Ling, S., Strohmer, T.: Self-calibration and biconvex compressive sensing. Inverse Probl **31**(11), 115,002 (2015)
26. Ling, S., Strohmer, T.: Blind deconvolution meets blind demixing: algorithms and performance bounds. IEEE Trans. Inform. Theory **63**(7), 4497–4520 (2017)
27. Ling, S., Strohmer, T.: Regularized gradient descent: a nonconvex recipe for fast joint blind deconvolution and demixing. arXiv:1703.08642 (2017)
28. Ma, Z.: Sparse principal component analysis and iterative thresholding. Ann. Statist. **41**(2), 772–801 (2013)
29. Mendelson, S., Rauhut, H., Ward, R., et al.: Improved bounds for sparse recovery from subsampled random convolutions. Ann. Appl. Probab. **28**(6), 3491–3527 (2018)
30. Needell, D., Tropp, J.A.: Cosamp: iterative signal recovery from incomplete and inaccurate samples. Appl. Comput. Harmon. Anal. **26**(3), 301–321 (2009)
31. Oymak, S., Jalali, A., Fazel, M., Eldar, Y.C., Hassibi, B.: Simultaneously structured models with application to sparse and low-rank matrices. IEEE Trans. Inform. Theory **61**(5), 2886–2908 (2015)
32. Qu, Q., Zhang, Y., Eldar, Y.C., Wright, J.: Convolutional phase retrieval via gradient descent. arXiv:1712.00716 (2017)
33. Soltanolkotabi, M.: Structured signal recovery from quadratic measurements: breaking sample complexity barriers via nonconvex optimization. arXiv:1702.06175 (2017)
34. Stöger, D., Geppert, J.A., Krahmer, F.: Sparse power factorization with refined peakiness conditions. In: IEEE Statistical Signal Processing Workshop 2018. IEEE (2018)
35. Tillmann, A.M., Pfetsch, M.E.: The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. IEEE Trans. Inform. Theory **60**(2), 1248–1259 (2014)
36. Wang, T., Berthet, Q., Samworth, R.J.: Statistical and computational trade-offs in estimation of sparse principal components. Ann. Statist. **44**(5), 1896–1930 (2016)
37. Xu, G., Liu, H., Tong, L., Kailath, T.: A least-squares approach to blind channel identification. IEEE Trans. Signal Process. **43**(12), 2982–2993 (1995)

C. Article: On the convex geometry of blind deconvolution and matrix completion

Summary

We have seen in previous sections that both blind deconvolution and matrix completion can be formulated as problems of the form

$$y = \mathcal{A}(X_0) + e \in \mathbb{C}^m,$$

where X_0 is a low-rank matrix, $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ is a linear operator, and $\|e\| \leq \tau$ is additive, deterministic noise. A widely-studied approach is to solve this problem is the semidefinite program

$$\begin{aligned} \text{minimize} \quad & \|X\|_* \\ \text{subject to} \quad & \|y - \mathcal{A}(X)\| \leq \tau. \end{aligned} \tag{C.1}$$

While previous results show that exact recovery is possible for this approach for a near-optimal amount of measurements, the corresponding noise bounds are rather of the type

$$\|\hat{X} - X_0\|_F \lesssim \sqrt{n_2}\tau,$$

where $\hat{X}$ denotes a minimizer of (C.1). This noise bound seems to be suboptimal as such an additional dimension factor as $\sqrt{n_2}$ does not appear in other problems like matrix sensing [99] or phase retrieval [17, 10]. The goal of the following article is to obtain a better understanding of this phenomenon. Its two main results are as follows.

(a) Matrix completion and blind deconvolution are unstable in the following sense. Let the ground truth matrix X_0 be given. Then with high probability there is a noise level $\tau_0 > 0$ such that for all $\tau \leq \tau_0$ there exists a noise vector e with $\|e\| \leq \tau$ such that there is an alternative solution $\hat{X}$ to (C.1) with the following properties.

- $\hat{X}$ is feasible, i.e., $\|\mathcal{A}(X) - y\| \leq e$
- $\hat{X}$ is preferred to X_0 by the SDP (C.1), i.e., $\|\hat{X}\|_* \leq \|X_0\|_*$,
- but $\hat{X}$ is far from X_0 in Frobenius norm, i.e., $\|\hat{X} - X_0\|_F$ is large.

For a precise statement of these results we refer to [68, Theorem 3.1] and [68, Theorem 3.5].

(b) If the noise level τ is sufficiently large near-optimal noise bounds are possible in the blind deconvolution framework. For a precise statement of this result we refer to [68, Theorem 3.7].

Statement of individual contribution

The project started with my observation that in blind deconvolution Mendelson's small method can be applied successfully to a certain subset of matrices in the descent cone of the nuclear norm. This observation allowed me to find a proof of [68, Theorem 3.7], where Felix Krahmer and I had several research discussions about this topic.
While we were still trying to prove that $\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(h_0 m_0^*\right)\right) \gtrsim 1$, which was our first conjecture, my advisor Felix Krahmer suggested me to look at certain special matrices B and to also take into account the possibility that such a result might just not be true in general. From this discussions the proof idea for a first version of [68, Proposition 3.3] arose, which showed a special case. After I worked out the proof details, Felix Krahmer and I found in joint discussions a way to simplify the proof. I could use this insight to show [68, Proposition 3.3] in full generality.
I was mainly responsible for tranferring our results from the randomized blind deconvolution setting to the matrix completion framework, where Felix Krahmer provided technical input. Furthermore, I was fully responsible for writing and preparing this manuscript. After that Felix Krahmer and I both proofread and polished the article together.

Conference proceedings

The results on blind deconvolution have been presented in part at the 52nd Annual Asilomar Conference on Signals, Systems, and Computers [67].

On the convex geometry of blind deconvolution and matrix completion*

Felix Krahmer†, Dominik Stöger†

March 13, 2019

Low-rank matrix recovery from structured measurements has been a topic of intense study in the last decade and many important problems like matrix completion and blind deconvolution have been formulated in this framework. An important benchmark method to solve these problems is to minimize the nuclear norm, a convex proxy for the rank. A common approach to establish recovery guarantees for this convex program relies on the construction of a so-called approximate dual certificate. However, this approach provides only limited insight in various respects. Most prominently, the noise bounds exhibit seemingly suboptimal dimension factors. In this paper we take a novel, more geometric viewpoint to analyze both the matrix completion and the blind deconvolution scenario. We find that for both these applications the dimension factors in the noise bounds are not an artifact of the proof, but the problems are intrinsically badly conditioned. We show, however, that bad conditioning only arises for very small noise levels: Under mild assumptions that include many realistic noise levels we derive near-optimal error estimates for blind deconvolution under adversarial noise.

Keywords: Convex relaxation, matrix completion, blind deconvolution, noise robustness, nuclear norm minimization

1. Introduction

A number of recent works have explored the observation that various ill-posed inverse problems in signal processing, imaging, and machine learning can be naturally formulated as the task of recovering a low-rank matrix $X_0 \in \mathbb{C}^{n_1 \times n_2}$ from an underdetermined system of structured linear measurements

$$y = \mathcal{A}(X_0) + e \in \mathbb{C}^m,$$

*The results of this paper have been presented in part at the 52nd Annual Asilomar Conference on Signals, Systems, and Computers, October 28-31, 2018, Pacific Grove, USA [37]

†Department of Mathematics, Technische Universität München, 85748 Garching/Munich, Germany

where $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ is a linear map and $e \in \mathbb{C}^m$, $\|e\| \leq \tau$, represents additive noise. Such problems include, for example, matrix completion [8], phase retrieval [9], blind deconvolution [1], robust PCA [6], and demixing [47]. In this paper, we aim to analyze the worst case scenario, that is, we do not make any assumptions on the noise except for the bound on its Euclidean norm (this scenario is sometimes referred to as *adversarial noise*, as it allows for noise specifically designed to be most harmful in a given situation). A natural first approach to recover X_0 that remains an important benchmark is to solve the semidefinite program

$$\begin{aligned} \text{minimize} \quad & \|X\|_* \\ \text{subject to} \quad & \|y - \mathcal{A}(X)\| \leq \tau, \end{aligned}$$

where $\|\cdot\|_*$ denotes the nuclear norm, i.e., the sum of the singular values. Recovery guarantees have been shown under the assumption that the measurement operator $\mathcal{A}$ possesses a certain degree of randomness. To establish such guarantees various proof strategies have been proposed, including approaches via the restricted isometry property [52, 45], descent cone analysis [11], and so-called approximate dual certificates [24, 23]. While the latter approach remains state of the art for many structured problems including the highly relevant problems of randomized blind deconvolution and matrix completion, it seemingly has some disadvantages. Most prominently, the resulting recovery guarantees take the form

$$\|\hat{X} - X_0\|_F \lesssim \sqrt{n_1}\tau, \tag{1}$$

where $\hat{X}$ denotes a minimizer of the semidefinite program above and $\|\cdot\|_F$ denotes the Frobenius norm, whereas under comparable normalization, the first two approaches, when applicable, give rise to superior recovery guarantees of the form

$$\|\hat{X} - X_0\|_F \lesssim \tau.$$

Before this paper it was open whether the additional dimension scaling factor in (1) is a proof artifact. Similarly, for randomized blind deconvolution one of the coherence terms appearing in the result was believed to arise only from the proof technique (cf. [44, Remark 2]).

Another drawback of proceeding via an approximate dual certificate is that it gives only limited insight into geometric properties of the problems such as the null-space property [15], which is also an important ingredient for the study of some more efficient non-convex algorithms [19, 40].

Approaches via descent cone analysis [11], in contrast, provide much more geometrical insight. The underlying idea of such approaches is to study the minimum conic singular value defined by

$$\lambda_{\min}(\mathcal{A}, \mathcal{K}) := \inf_{Z \in \mathcal{K} \setminus \{0\}} \frac{\|\mathcal{A}(Z)\|}{\|Z\|_F}$$

for $\mathcal{K}$ the descent cone of the underlying atomic norm – the nuclear norm in case of low-rank matrix recovery. For a more detailed review of this approach including a precise

definition of the descent cone we refer to Section 2.3 below. Through the study of the minimum conic singular value many superior results were obtained for low-rank recovery problems, most importantly in the context of phase retrieval [39, 38]. Furthermore, minimum conic singular values can also help to understand certain nonlinear measurement models [50].

For all these reasons, it would be desirable to apply this approach also for matrix completion and blind deconvolution. A challenge that one faces, however, is that for both problems one cannot hope to recover all low-rank matrices; rather, only matrices that satisfy certain coherence constraints are admissible (cf. the discussion in [57, Section 5.4]). In this article we address this challenge, providing the first geometric analysis of these problems. We find that the dimensional factors appearing in the error bounds are the true scaling of the minimum conic singular value and hence intrinsically relate to the underlying geometry. Nevertheless for blind deconvolution, near-optimal recovery is possible, if the noise level is not too small.

1.1. Organization of the paper and our contribution

In Section 2 we will review blind deconvolution, matrix completion, as well as some techniques related to descent cone analysis. In Section 3 we will present the main results of this paper. Theorems 3.1 and 3.5 establish that for both blind deconvolution and matrix completion, nuclear norm minimization is intrinsically ill-conditioned. In contrast, Theorem 3.7 provides a near-optimal error bound for blind deconvolution when the noise level is not too small, implying that the conditioning problems only take effect for very small noise levels. The upper bounds for the minimum conic singular value which are the main ingredients of Theorems 3.1 and 3.5 are derived in Section 4. In Section 5 we prove the stability results for blind deconvolution.
We believe that not only our results, but also the proof techniques and geometric insights in this manuscript will be of general interest and help to obtain further understanding of low-rank matrix recovery models, in particular under coherence constraints. We discuss interesting directions for future research in Section 6.

2. Background and related work

2.1. Blind deconvolution

Blind deconvolution problems arise in a number of different areas in science and engineering such as astronomy, imaging, and communications. The goal is to recover both an unknown signal and an unknown kernel from their convolution. In this paper we work with the circular convolution, which is defined by

$$w * x := \left(\sum_{j=1}^{L} w_j x_{k-j} \right)_{k=1}^{L},$$

where the index difference $k-j$ is considered modulo L. Without further assumptions on w and x this bilinear map is far from injective. Consequently, it is crucial to impose structural constraints on both w and x. Arguably, the simplest such model is given by linear constraints, that is, both w and x are constrained to known subspaces. Such a model is reasonable in many applications. In wireless communication, for example, it makes sense to assume that the channel behaviour is dominated by the most direct paths and for the signal x a subspace model can be enforced by embedding the message via a suitable coding map into a higher-dimensional space before transmission.

The first rigorous recovery guarantees for such a model were derived by Ahmed, Recht, and Romberg [1]. More precisely, they assume that $w = Bh$, where $B \in \mathbb{C}^{L\times K}$ is a fixed, deterministic matrix such that $B^*B = Id_K$ (i.e., B is an isometry) and they model $x = C\overline{m_0}$, where $\overline{m}_0$ denotes the complex-conjugate of m_0. Here, the matrix $C \in \mathbb{C}^{L\times K}$ is a random matrix, whose entries are independent and identically distributed with circular symmetric normal distribution $\mathcal{CN}\left(0, \frac{1}{\sqrt{L}}\right)$. In this paper we also adopt this model.

Using the well-known fact that the Fourier transform diagonalizes the circular convolution one can rewrite

$$w * x = \sqrt{L}F^*\mathrm{diag}\left(Fw\right)Fx,$$

where $F \in \mathbb{C}^{L\times L}$ denotes the normalized, unitary discrete Fourier matrix, and

$$\widehat{w * x} := F\left(w * x\right) = \sqrt{L}\mathrm{diag}\left(FBh_0\right)FC\overline{m}_0.$$

Denoting by b_ℓ the ℓth row of the matrix $\overline{FB}$, and by c_ℓ the ℓth row of the matrix $\sqrt{L}FC$, one observes that

$$\left(\widehat{w * x}\right)_\ell = b_\ell^* h_0 m_0^* c_\ell = \mathrm{Tr}\left(h_0 m_0^* c_\ell b_\ell^*\right) = \langle b_\ell c_\ell^*, h_0 m_0^*\rangle_F.$$

Furthermore, because of the rotation invariance of the circular symmetric normal distribution all the entries of the vectors $\{c_\ell\}_{\ell=1}^L$ are (jointly) independent and identically distributed with distribution $\mathcal{CN}\left(0,1\right)$. Noting that the expression $\langle h_0 m_0^*, b_\ell c_\ell^*\rangle_F$ is linear in $h_0 m_0^*$, Ahmed, Recht, and Romberg [1] defined the operator $\mathcal{A} : \mathbb{C}^{K\times N} \to \mathbb{C}^L$ by

$$\left(\mathcal{A}\left(X\right)\right)\left(\ell\right) := \langle b_\ell c_\ell^*, X\rangle_F \tag{2}$$

obtaining the measurement model

$$y = \widehat{w * x} + e = \mathcal{A}\left(X_0\right) + e,$$

where $e \in \mathbb{C}^L$ is additive noise and $X_0 = h_0 m_0^*$. The goal is then to determine h_0 and m_0 from $y \in \mathbb{C}^L$ up to the inherent scaling ambiguity, or, equivalently, to find the rank-one matrix $X_0 = h_0 m_0^*$.

For $e = 0$, among all solutions of the equation $y = \mathcal{A}\left(X_0\right)$, the matrix X_0 is the one with the smallest rank. For this reason, Ahmed, Recht, and Romberg [1] suggested to

minimize a natural proxy for the rank, the nuclear norm $\|\cdot\|_*$, defined as the sum of the singular values of a matrix.

$$\begin{array}{ll} \text{minimize} & \|X\|_* \\ \text{subject to} & \|\mathcal{A}(X) - y\| \leq \tau. \end{array} \tag{3}$$

Here $\tau > 0$ is an a priori bound for the noise level, that is, we assume that $\|e\| \leq \tau$. For this semidefinite program, they establish the following recovery guarantee.

Theorem 2.1 ([1]). *Consider measurements of the form $y = \mathcal{A}(h_0 m_0^*) + e$ for $h_0 \in \mathbb{C}^K$, $m_0 \in \mathbb{C}^N$, $e \in \mathbb{C}^L$, and $\mathcal{A}$ as defined in (2). Assume that $\|e\| \leq \tau$ and*

$$L/\log^3 L \gtrsim K\mu_{\max}^2 + N \max\left\{\mu_{h_0}^2; \tilde{\mu}_{h_0}^2\right\}.$$

Then with probability exceeding $1 - \mathcal{O}\left(L^{-1}\right)$ every minimizer $\hat{X}$ of the SDP (3) satisfies

$$\|\hat{X} - h_0 m_0^*\|_F \lesssim \sqrt{K+N}\tau. \tag{4}$$

Here $\mu_{\max}^2$ and $\mu_{h_0}^2$ are coherence parameters, which are defined via

$$\mu_{\max}^2 := \frac{L}{K} \max_{\ell \in [L]} \|b_\ell\|^2.$$

and

$$\mu_{h_0}^2 := \frac{L}{\|h_0\|^2} \max_{\ell \in [L]} |\langle b_\ell, h_0 \rangle|^2.$$

The third coherence factor $\tilde{\mu}_{h_0}$ is a technical term corresponding to a partition that is constructed as a part of the proof of Theorem 2.1, which is based on the Golfing Scheme [23].

To put the impact of the coherence factors into perspective, observe that if all vectors b_ℓ have the same ℓ_2-norm, one obtains that $\mu_{\max} = 1$; this will be the case, for example, when B is a low-frequency Fourier matrix, as it appears for applications in wireless communication. The second coherence factor always satisfies $1 \leq \mu_{h_0}^2 \leq K\mu_{\max}^2$. If μ_{h_0} is smaller, this indicates that the mass $\|h_0\|^2 = \sum_{\ell=1}^L |\langle b_\ell, h_0 \rangle|^2$ is distributed fairly evenly among $|\langle b_\ell, h_0 \rangle|$. For example, if $\mu_{h_0} = 1$, then $|\langle b_\ell, h_0 \rangle| = \frac{1}{\sqrt{L}}\|h_0\|$ for all $\ell \in [L]$. Numerical simulations in [1] confirm that many h_0 corresponding to large μ_{h_0} show worse performance, indicating that this factor may be necessary.
The last coherence factor $\tilde{\mu}_{h_0}$, in contrast, will no longer appear in our result below, which is why we refrain from detailed discussion. We refer the interested reader to [44, Remark 2.1] and [28, Section 2.3] for details.
For generic h_0 the parameters μ_{h_0} and $\tilde{\mu}_{h_0}$ are reasonably small. For example, if h_0 is chosen from the uniform distribution on the sphere, one can show that with high probability $\mu_{h_0} = \mathcal{O}\left(\sqrt{\log L}\right)$.
For the noiseless case, i.e., $\tau = 0$, Theorem 2.1 yields exact recovery, and the required

sample complexity $L/\log^3 L \gtrsim K+N$ is optimal up to logarithmic factors, as the number of degrees of freedom is $K+N-1$ (see [30] for an exact identifiability analysis based on algebraic geometry.) However, if there is noise, the bound for the reconstruction error scales with $\sqrt{K+N}$, in contrast to other measurement scenarios such as low-rank matrix recovery from Gaussian measurements (see, e.g., [11]).

Let us comment on some related work. The foundational paper [1] has triggered a number of follow-up works on the problem of randomized blind deconvolution. A first line of works extended the result to recovering signals from their superposition $\sum_{i=1}^{r} w_i * x_i$, a problem often referred to as blind demixing [44, 28]. Another line of works investigated non-convex (gradient-descent based) algorithms [43, 46, 26], which have the advantage that they are computationally less expensive, as they operate in the natural parameter space. It has been shown that they require a near-optimal number of measurements for recovery. For such an algorithm, [43] derived near-optimal noise-bounds for a Gaussian noise model. However, as in this paper, we focus on the scenario of adversarial noise (instead of random noise) the resulting guarantees are not comparable to ours below.

2.2. Matrix completion

The matrix completion problem of reconstructing a low-rank matrix $X_0 \in \mathbb{R}^{n_1 \times n_2}$ (we assume that w.l.o.g. $n_1 \geq n_2$) from only a part of its entries arises in many different applications such as in collaborative filtering [53] and multiclass learning [3]. For this reason one could observe a flurry of work on this problem in the last decade, and we will only be able to give a very selective overview of this topic. The precise sampling model that we consider is that m entries of X_0 are sampled uniformly at random with replacement. Denoting by e_i the standard coordinate vectors in $\mathbb{R}^{n_1}$ and $\mathbb{R}^{n_2}$, respectively, the corresponding measurement operator $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \to \mathbb{R}^m$ can be written as

$$\mathcal{A}(X)(i) := \sqrt{\frac{n_1 n_2}{m}} \langle X, e_{a_i} e_{b_i}^* \rangle_F, \tag{5}$$

where $(a_i, b_i) \in [n_1] \times [n_2]$ is chosen uniformly at random for each $i \in [m]$ (and independently from all other measurements). The scaling factor $\sqrt{\frac{n_1 n_2}{m}}$ in the definition of the measurement operator $\mathcal{A}$ is chosen to ensure that $\mathbb{E}\left[\|\mathcal{A}(X)\|^2\right] = \|X\|_F^2$. (Some other papers on matrix completion choose a different scaling. We have chosen this normalization because in this way the results for the matrix completion problem can be better compared to those for the blind deconvolution scenario.) Alternative sampling models analyzed in other works include sampling a subset Ω uniformly from $[n_1] \times [n_2]$ (i.e., without replacement, see, e.g., [10]), or sampling using random selectors.

Again we aim to recover X_0 from noisy observations $y = \mathcal{A}(X_0) + e$, with a noise vector $e \in \mathbb{R}^m$ that satisfies $\|e\| \leq \tau$ via the SDP

$$\begin{array}{ll} \text{minimize} & \|X\|_* \\ \text{subject to} & \|\mathcal{A}(X) - y\| \leq \tau. \end{array} \tag{6}$$

For matrix completion, this approach has first been studied in [8].

It is well known that similarly to the blind deconvolution problem, some incoherence assumptions are necessary to allow for successful recovery. Indeed, suppose that $X_0 = e_1 e_1^*$. Then, if $m \ll n_1 n_2$ with high probability it holds that $\mathcal{A}(X_0) = 0$ and one cannot hope to recover X_0. To avoid such special cases, one needs to ensure that the mass of the Frobenius norm of X_0 is spread out over all entries rather evenly. This property is captured by the following coherence parameters [23]

$$\mu(U) := \sqrt{\frac{n_1}{r}} \max_{i \in [n_1]} \|U^* e_i\|$$
$$\mu(V) := \sqrt{\frac{n_2}{r}} \max_{i \in [n_2]} \|V^* e_i\|.$$

For these coherence parameters, a series of works [8, 10, 23, 51, 12] lead to the following recovery guarantee for the noiseless scenario.

Theorem 2.2 ([12] [1]). *Consider measurements of the form $y = \mathcal{A}(X_0)$, where $X_0 \in \mathbb{R}^{n_1 \times n_2}$ is a rank-r matrix and $\mathcal{A}$ is given by (5). Assume that*

$$m \geq C \max\left\{\mu^2(U); \mu^2(V)\right\} r n_1 \log^2 n_1$$

Then with probability at least $1 - \mathcal{O}\left(n_1^{-1}\right)$ the matrix X_0 is the unique minimizer of the SDP (6) with $\tau = 0$.

As for blind deconvolution, this result has been shown using an approximate dual certificate. In [7] this result has been generalized to the case of adversarial noise, showing that with high probability the minimizer $\hat{X}$ of (6) satisfies

$$\|\hat{X} - X_0\|_F \lesssim \tau \sqrt{n_2}, \tag{7}$$

whenever $m \gtrsim n_1 \text{polylog } n_1$. As in the blind deconvolution framework, this error bound differs from the case of full Gaussian measurements as discussed, for example, in [11], and also from oracle estimates [6, Section III.B] by a dimensional scaling factor, which will be addressed in this paper.

Also random noise models for matrix completion have been studied in a number of works. In particular, we would like to mention [34, 49], which derive near-optimal rates (both in sample size and estimation error) for matrix completion under subexponential noise with a slightly different nuclear-norm penalized estimator than the one we consider as long as the noise-level is not too small. Similar bounds have also been obtained in [33] using an estimator, which is closer to the one in this work.

Apart from convex methods also many nonconvex algorithms have been proposed and analysed, for example a number of variants of gradient descent (see, e.g., [31, 27, 25, 56,

[1] For the case of very small ranks this result can be refined further [17]. Namely one can remove one of the two log-factors at the cost of an r^3-dependence on the rank.

21, 19, 40, 46]). Arguably the strongest result for matrix completion under adversarial noise has been shown in [31, 32]. These works propose a non-convex algorithm based on Riemannian optimization and show that if the number of measurements is larger than $r^2 n_1 \mathrm{polylog}(n_1)$ the true matrix can be reconstructed up to an estimation error superior to the one in [7]. Namely for κ denoting the condition number of the matrix X_0 they show that the output $\hat{X}$ of their algorithm satisfies (in our notation)

$$\|\hat{X} - X_0\| \lesssim \kappa^2 \sqrt{rm} \|e\|_\infty, \tag{8}$$

provided the noise level is below a certain, small threshold that scales with the smallest singular value of X_0. For error vectors e that are spread out evenly and matrices that are well conditioned, one has that $\sqrt{m}\|e\|_\infty \approx \|e\|_2$, so this bound is superior to (7) in the sense that the scaling factors that appear only scale with the rank r and not the dimension. It should be noted though that in contrast to nuclear norm minimization the underlying algorithm requires precise knowledge of the true rank of the matrix to be recovered.

Just before completion of this manuscript, Chen et al. [14] bridged convex and nonconvex approaches, using nonconvex methods to analyze a convex recovery scheme. Their results provide near optimal recovery guarantees for the matrix completion problem via nuclear norm minimization under a subgaussian random noise model for a much larger range of admissible noise levels than the aforementioned works. More precisely, the proof is based on the observation that in their scenario the minimizer of the convex problem is very close to an approximate critical point of a non-convex gradient based method. This allows them to transfer existing stability results [46] for non-convex optimization to the convex problem. However, the required sample complexity scales suboptimally in the rank r of the matrix and similarly to (8), the error bound depends on the condition number κ.

2.3. Descent cone analysis

In recent years a number of works have studied low-rank matrix recovery and compressed sensing via a descent cone analysis. This approach has been pioneered for ℓ_1-norm minimization in [55] and for more general (atomic) norms in [11]. Here the descent cone of a norm at a point $X_0 \in \mathbb{C}^{K\times N}$ is the set of all possible directions $Z \in \mathbb{C}^{K\times N}$ such that the norm does not increase. For the nuclear norm, this leads to the following definition.

Definition 2.3. *For any matrix $X_0 \in \mathbb{C}^{K\times N}$ define its descent cone $\mathcal{K}_*(X_0)$ by*

$$\mathcal{K}_*(X_0) := \left\{ Z \in \mathbb{C}^{K\times N} : \ \|X_0 + \varepsilon Z\|_* \leq \|X_0\|_* \text{ for some } \varepsilon > 0 \right\}.$$

To understand its relevance for recovery guarantees assume for a moment that we are in the noiseless scenario, i.e., $\tau = 0$ and $e = 0$. Then the matrix $X_0 \in \mathbb{C}^{K\times N}$ is the unique minimizer the semidefinite program (3), if and only if the null space of $\mathcal{A}$ does not intersect the descent cone $\mathcal{K}_*(X_0)$. In the case of noise, the constraint $\|y - \mathcal{A}(X_0)\| \leq \tau$

in the SDPs (3) and (6) defines a region around $X_0 + \ker\mathcal{A}$, i.e., the affine subspace consistent with the observed measurements in the noiseless scenario. The intersection of this region with the set of all signals that have a smaller nuclear norm than the ground truth X_0 is the set of feasible solutions that are preferred to X_0. The following quantity for a matrix X_0, which is often referred to as minimum conic singular value, quantifies the size of this intersection

$$\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(X_0\right)\right) := \inf_{Z \in \mathcal{K}_*(X_0)\setminus\{0\}} \frac{\|\mathcal{A}\left(Z\right)\|}{\|Z\|_F}.$$

If $\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(X_0\right)\right)$ becomes larger, this intersection becomes smaller, which translates into stronger recovery guarantees. The following theorem confirms this intuition.

Theorem 2.4. *[11, Proposition 2.2] Let $\mathcal{A} : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^m$ be a linear operator and assume that $y = \mathcal{A}\left(X_0\right) + e$ with $\|e\| \leq \tau$. Then any minimizer $\hat{X}$ of the SDP (3) satisfies*

$$\|\hat{X} - X_0\|_F \leq \frac{2\tau}{\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(X_0\right)\right)}.$$

When measurement matrices of the operator $\mathcal{A}$ are full Gaussian matrices (in contrast to rank-1 measurements as in this paper) and $\mathcal{A}$ is normalized such that $\mathbb{E}\left[\mathcal{A}^*\mathcal{A}\right] = \mathrm{Id}$, for an arbitrary low-rank matrix X_0 one has with high probability that $\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(X_0\right)\right) \asymp 1$. Consequently, Theorem 2.4 yields an optimal estimation error even for adversarial noise. As we will show this is no longer the case for blind deconvolution and matrix completion.

The geometric analysis of linear inverse problems via the descent cone and the minimum conic singular value has lead to many new results and insights in compressed sensing and low-rank matrix recovery. For convex programs the phase transition of the success rate could be precisely predicted [2]. As the proofs are specific to full Gaussian measuements, they do not apply for a number of important structured and heavy-tailed measurement scenarios. Stronger results [41, 18, 38, 29, 39] were subsequently obtained using Mendelson's small ball method [35, 48], a powerful tool for bounding a nonnegative empirical process from below, now often refereed to as Mendelson's small ball method.

2.4. Notation

For $n \in \mathbb{N}$ we will write $[n]$ to denote the set $\{1; \ldots; n\}$. For any set A we will denote its cardinality by $|A|$. For a complex number z we will denote its real part by $\mathrm{Re}\left(z\right)$ and its imaginary part by $\mathrm{Im}\left(z\right)$. By $\log\left(\cdot\right)$ we will denote the logarithm to the base e. By $\mathbb{E}X$ we will denote the expectation of a random variable X and by $\mathbb{P}\left(A\right)$ we denote the probability of an event A. If $v \in \mathbb{C}^n$ we will denote its ℓ_2-norm by $\|v\|$ and its Hermitian transpose by v^*. For $u, v \in \mathbb{C}^n$ the (Euclidean) inner product is defined by $\langle u, v\rangle := u^*v$. Furthermore, for $Z \in \mathbb{C}^{n_1 \times n_2}$ its spectral norm is given by $\|Z\|$, i.e., the dual norm of the nuclear norm $\|Z\|_*$. Moreover, the Frobenius norm of Z is defined by $\|Z\|_F$ with corresponding inner product $\langle Z, W\rangle_F := \mathrm{Tr}\left(Z^*W\right)$, where $W \in \mathbb{C}^{n_1 \times n_2}$.

When we study matrix completion, we will work with matrices $Z \in \mathbb{R}^{n_1 \times n_2}$ and the previous quantities will be defined analogously. Moreover, in that scenario we will use the notation $\|Z\|_{\ell_\infty} := \max_{(i,j)\in[n_1]\times[n_2]} |Z_{i,j}|$, where $\{Z_{i,j}\}_{i,j=1}^{n_1,n_2}$.

3. Our results

3.1. Instability of low-rank matrix recovery

3.1.1. Blind deconvolution

Our first main result states that randomized blind deconvolution can be unstable under adversarial noise.

Theorem 3.1. *Let $K, N \in \mathbb{N} \setminus \{1\}$. Assume that L is an integer multiple of K and that*

$$C_1 K \leq L \leq \frac{KN}{9}.$$

Then there exists a matrix $B \in \mathbb{C}^{L\times K}$ satisfying $B^*B = Id_K$ and with FB having rows of equal norm, i.e., $\mu^2_{\max} = 1$, such that for all $h_0 \in \mathbb{C}^K \setminus \{0\}$ and $m_0 \in \mathbb{C}^N \setminus \{0\}$ the following holds:
With probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{K}{C_2\mu^2}\right)\right)$, where $\mu^2 = \mu^2_{h_0} = \frac{L}{\|h_0\|^2}\max_{\ell\in[L]} |\langle b_\ell, h_0\rangle|^2$, there is $\tau_0 > 0$ such that for all $\tau \leq \tau_0$ there exists an adversarial noise vector $e \in \mathbb{C}^L$ with $\|e\| \leq \tau$ that admits an alternative solution $\tilde{X}$ with the following properties.

- *$\tilde{X}$ is feasible, i.e., $\|\mathcal{A}\left(\tilde{X}\right) - y\| = \tau$ for $y = \mathcal{A}(h_0 m_0^*) + e$ the noisy measurement vector*
- *$\tilde{X}$ is preferred to $X_0 = h_0 m_0^*$ by the SDP (3), i.e., $\|\tilde{X}\|_* \leq \|X_0\|_*$, but*
- *$\tilde{X}$ is far from the true solution in Frobenius norm, i.e.,*

$$\|\tilde{X} - X_0\|_F \geq \frac{\tau}{C_3}\sqrt{\frac{KN}{L}}.$$

The constants C_1, C_2, and C_3 are universal.

Remark 3.2. *The matrix B in the above result exactly fits into the framework of Theorem 2.1, and also one can check that for our choice of B (see the proof of Proposition 3.3 for its definition) one has that $\mu_{2,h_0} = \mu_{1,h_0}$. That is, the assumptions of Theorem 2.1 cannot be enough to deduce stability.*

We do not expect, however, that this kind of instability is observed for arbitrary isometric embeddings $B \in \mathbb{C}^{L\times K}$. In particular, if B is a random embedding we expect that a similar result as in [38] applies.

To put our results in perspective note that for $L \asymp (K+N)\,\mathrm{polylog}\,(K+N)$, which is the minimal number of measurements required for noiseless recovery, it holds that $\sqrt{\frac{KN}{L}} \asymp \sqrt{\frac{\min\{K,N\}}{\mathrm{polylog}(K+N)}}$. Up to logarithmic factors, this coincides with the rate predicted by (4), whenever $K \asymp N$.

Theorem 3.1 is a direct consequence of the following proposition, which we think is interesting in its own right.

Proposition 3.3. *Let $K, N \in \mathbb{N} \setminus \{1\}$. Assume that L is an integer multiple of K and that*

$$C_1 K \leq L \leq \frac{KN}{9}. \tag{9}$$

Then there exists $B \in \mathbb{C}^{L\times K}$ satisfying $B^*B = Id_K$ and $\mu_{\max}^2 = 1$, whose corresponding measurement operator $\mathcal{A}$ satisfies the following.

Let $h_0 \in \mathbb{C}^K \setminus \{0\}$, $m_0 \in \mathbb{C}^N \setminus \{0\}$ and set $\mu^2 := \mu_{h_0} = L\frac{\max_{\ell\in[L]}|\langle b_\ell, h_0\rangle|^2}{\|h_0\|^2}$. Then with probability at least $1 - \mathcal{O}\left(\exp\left(-C_2 K/\mu^2\right)\right)$ it holds that

$$\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(h_0 m_0^*\right)\right) \leq C_3\sqrt{\frac{L}{KN}}. \tag{10}$$

Here C_1, C_2, and C_3 are absolute constants.

The proof of Proposition 3.3 will be provided in Section 4. Note that by definition of the minimum conic singular value $\lambda_{\min}\left(\mathcal{A}, \mathcal{K}_*\left(h_0 m_0^*\right)\right)$ Proposition 3.3 is equivalent to the statement that with high probability there is $Z \in \mathcal{K}_*\left(h_0 m_0^*\right) \setminus \{0\}$ such that

$$\|\mathcal{A}(Z)\| \lesssim \sqrt{\frac{L}{KN}}\|Z\|_F.$$

Our construction of such $Z \in \mathcal{K}_*\left(h_0 m_0^*\right)$ relies on the observation that with high probability there is a rank-one matrix $W \in \mathbb{C}^{K\times N}$ in the null-space of $\mathcal{A}$ which is relatively close to the descent cone (with respect to the $\|\cdot\|_F$-distance). Perturbing W by $-\beta h_0 m_0^*$ for a suitable β one can then obtain a matrix $Z \in \mathcal{K}_*\left(h_0 m_0^*\right)$, which fulfills (3.1.1).
The existence of such a matrix $W \in \ker\mathcal{A}$ also reveals a fact about the geometry of the problem, which we find somewhat surprising: while the null space of $\mathcal{A}$ does not intersect the descent cone (otherwise exact recovery would not be possible), the angle between those objects is very small. This is very different from the behavior for measurement matrices $\mathcal{A}$ with i.i.d. Gaussian entries (instead of $b_\ell c_\ell^*$).

Remark 3.4. *While $\tilde{X}$ is preferred to the true solution by the SDP (3) $\tilde{X}$ is typically not a minimizer of (3). To see this, assume that without noise exact recovery is possible, which is the case with high probability by Theorem 2.1. Then consider $\tilde{X} = X_0 + tZ$ for $Z \in \mathcal{K}_*\left(h_0 m_0^*\right)$ of the form $Z = W - \beta h_0 m_0^*$ with $W \in ker\mathcal{A}$ and $\beta > 0$ such that*

$\frac{\|\mathcal{A}(Z)\|}{\|Z\|_F} \lesssim \sqrt{\frac{L}{KN}}$, *as in the proof of Proposition 3.3. As* $W \notin \mathcal{K}_*(h_0 m_0^*)$ *(otherwise exact recovery would not be possible) it follows that for* $t > 0$

$$\begin{aligned}\|\tilde{X}\|_* &= \|X_0 + tZ\|_* \\ &= \|(1-t\beta)X_0 + tW\|_* \\ &> \|(1-t\beta)X_0\|_*\end{aligned}$$

where the last line is due to $\mathcal{K}_*(X_0) = \mathcal{K}_*((1-t\beta)X_0)$.

On the other hand, we also have that $\mathcal{A}\left(\hat{X}\right) = \mathcal{A}((1-t\beta)X_0)$ *due to* $\mathcal{A}(W) = 0$ *and, hence,* $(1-t\beta)X_0$ *is admissible whenever* $\tilde{X}$ *is admissible. Consequently, the SDP (3) will always prefer* $(1-\beta t)h_0 m_0^*$ *to* $\tilde{X}$ *and* $\tilde{X}$ *will never be a minimizer. It remains an open problem what one can say about the minimizer* $\hat{X}$ *of (3), see also Section 6. Even if the minimizer of (3)* $\hat{X}$ *is closer to the ground truth (in* $\|\cdot\|_F$*-distance) than* $\tilde{X}$*, however, the nuclear norms of* X *and* $\tilde{X}$ *will be very close, which can easily lead to numerical instabilities.*

3.1.2. Matrix completion

Our second main result states that for arbitrary incoherent low-rank matrices, matrix completion is unstable with high probability. Note that in contrast to Theorem 3.1 which is based on a specific choice of parameters the following result holds for an arbitrary incoherent matrix X_0.

Theorem 3.5. *Let* $n_1 \geq n_2$ *and let* $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \to \mathbb{R}^m$ *be defined as in (5). Assume that* $X_0 \in \mathbb{R}^{n_1 \times n_2} \setminus \{0\}$ *is a rank* r *matrix with singular value decomposition* $X_0 = U\Sigma V^*$*. Moreover, assume that*

$$C_1 r n_1 \mu^2(V) \log(2r) \leq m \leq \frac{n_1 n_2}{32}.$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{m}{C_2 r \mu^2(U)\mu^2(V)}\right)\right)$ *there is* $\tau_0 > 0$ *such that for all* $\tau \leq \tau_0$ *there exists an adversarial noise vector* $e \in \mathbb{R}^m$ *with* $\|e\| \leq \tau$ *that admits an alternative solution* $\tilde{X} \in \mathbb{R}^{n_1 \times n_2}$ *with the following properties.*

- $\tilde{X}$ *is feasible, i.e.,* $\left\|\mathcal{A}\left(\tilde{X}\right) - y\right\| = \tau$ *for* $y = \mathcal{A}(X_0) + e$ *the noisy measurement vector*
- $\tilde{X}$ *is preferred to* X_0 *by the SDP (6), i.e.,* $\|\tilde{X}\|_* \leq \|X_0\|_*$ *, but*
- $\tilde{X}$ *is far from the true solution in Frobenius norm, i.e.,*

$$\|\tilde{X} - X_0\|_F \geq \frac{\tau}{C_3}\sqrt{\frac{r n_1 n_2}{m}}.$$

Here the constants C_1*,* C_2*, and* C_3 *are universal.*

Again, to put our results in perspective note that for $m \asymp n_1 \text{polylog}\,(n_1)$, which is the minimal number of measurements required for noiseless recovery, it holds that $\sqrt{\frac{rn_1n_2}{m}} \asymp \sqrt{\frac{n_2}{\text{polylog}(n_1)}}$. Up to logarithmic factors, this coincides with the rate predicted by (7).

Theorem 3.5 is a direct consequence of the following proposition, which in our opinion is of independent interest, as it provides a negative answer to a question by Tropp [57, Section 5.4].

Proposition 3.6. *Let $X_0 \in \mathbb{R}^{n_1 \times n_2} \setminus \{0\}$ be a rank-r matrix with corresponding singular value decomposition $X_0 = U\Sigma V^*$. Moreover, assume that*

$$C_1 r n_1 \mu^2(V) \log(2r) \leq m \leq \frac{n_1 n_2}{32}. \tag{11}$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{m}{C_2 r \mu^2(U)\mu^2(V)}\right)\right)$ it holds that

$$\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(X_0)) \leq C_3 \sqrt{\frac{m}{n_1 n_2 r}}. \tag{12}$$

The constants C_1, C_2, and C_3 are universal.

Proposition 3.6 corresponds to Proposition 3.3 for blind deconvolution and will be proved analogously. We will again show that with high probability there is $W \in \mathbb{R}^{n_1 \times n_2}$ such that $\mathcal{A}(W) = 0$ and W is relatively close to the descent cone of X_0 in $\|\cdot\|_F$-distance. Setting $Z := W - \beta UV^*$ for a suitable $\beta > 0$ yields an element of $\mathcal{K}_*(X_0)$ with

$$\frac{\|\mathcal{A}(Z)\|}{\|Z\|_F} \leq C_3 \sqrt{\frac{m}{n_1 n_2 r}}.$$

3.2. Stable recovery

A geometric interpretation of Theorems 3.1 and 3.5 is that the nuclear norm ball is near-tangential to both the kernels of matrix completion and randomized blind deconvolution. Given that tangent spaces only provide local approximation, these results leave open, what happens in some distance, i.e., for larger noise levels – this will depend on the curvature of the nuclear norm ball.

Our third main result concerns exactly this problem for the randomized blind deconvolution setup. As it turns out, the descent directions $Z \in \mathcal{K}_*(h_0 m_0^*)$ with $\|\mathcal{A}(Z)\|/\|Z\|_F$ very small correspond to directions of significant curvature. That is, only a very short segment in this direction will have smaller nuclear norm than $h_0 m_0^*$, and the corresponding alternative solutions all correspond to very small e. For noise levels τ large enough, in contrast, these directions can be excluded and one can obtain near-optimal error bounds. In order to precisely formulate this observation, we denote the set of μ-incoherent vectors $h \in \mathbb{C}^K$ with respect to $B \in \mathbb{C}^{L \times K}$ for $\mu \geq 1$ by

$$\mathcal{H}_\mu := \left\{ h_0 \in \mathbb{C}^K : \ \sqrt{L}|\langle b_\ell, h_0\rangle| \leq \mu \|h_0\| \text{ for all } \ell \in [L] \right\}.$$

With this notation, our result reads as follows.

Theorem 3.7. *Let $\alpha > 0$ and $B \in \mathbb{C}^{L\times K}$ such that $B^*B = Id$. Assume that*

$$L \geq C_1 \frac{\mu^2}{\alpha^2} (K+N) \log^2 L.$$

Then with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{L\alpha^{4/3}}{C_2 \log^{4/3}(eL)\mu^{4/3}}\right)\right)$ the following statement holds for all $h_0 \in \mathcal{H}_\mu \setminus \{0\}$, all $m_0 \in \mathbb{C}^N \setminus \{0\}$, all $\tau > 0$, and all $e \in \mathbb{C}^L$ with $\|e\| \leq \tau$: Any minimizer $\hat{X}$ of (3) satisfies

$$\|\hat{X} - h_0 m_0^*\|_F \leq \frac{C_3 \mu^{2/3} \log^{2/3} L}{\alpha^{2/3}} \max\left\{\tau; \alpha \|h_0 m_0^*\|_F\right\}.$$

Here C_1, C_2, and C_3 are absolute constants.

In words, this theorem establishes linear scaling in the noise level τ with only a logarithmic dimensional factor for $\tau \geq \alpha \|h_0 m_0^*\|_F$, in contrast to the polynomial factor required for small noise levels as a consequence of Theorem 3.1. Here the value of α can be chosen arbitrarily small, at the expense of an increased number of measurements. For example when one is interested in noise levels $\tau = \epsilon \mu^{-2} \log^{-2} L$ for some $\epsilon > \epsilon_0$ (this is the largest order to expect meaningful error bounds despite the additional logarithmic factors) one should choose $\alpha \asymp \epsilon_0 \mu^{-2} \log^{-2} L$, and near-linear error bounds will be guaranteed for a sample complexity of

$$L \geq C_1 \frac{\mu^6}{\epsilon_0^2} (K+N) \log^6 L.$$

Remark 3.8. *A similar approach to the proof of Theorem 3.7 also yields a corresponding result for rank-one matrix completion. Arguably, however, matrix completion is mainly of interest for ground truth matrices of rank higher than one, so we decided to omit the proof details.*

4. Upper bounds for the minimum conic singular values

4.1. Characterization of the descent cone of the nuclear norm

The goal of this section is to prove Proposition 3.3 and Proposition 3.6, from which we will then be able to deduce Theorem 3.1 and Theorem 3.5. For that we first discuss a characterization of the descent cone $\mathcal{K}_*(X)$. In order to state this characterization, Lemma 4.1, we need to introduce some additional notation. Let $X \in \mathbb{C}^{n_1\times n_2}$ be a matrix of rank r. We will denote its corresponding singular value decomposition by $X = U\Sigma V^*$, where $\Sigma \in \mathbb{R}^{r\times r}$ is a diagonal matrix with nonnegative entries and $U \in \mathbb{C}^{n_1\times r}$ and $V \in \mathbb{C}^{n_2\times r}$ are unitary matrices, i.e., $U^*U = V^*V = \mathrm{Id}_r$. This allows us to define the tangent space of the manifold of rank-r matrices at the point X by

$$T_X := \left\{UA^* + BV^* : \ A \in \mathbb{C}^{n_2\times r}, B \in \mathbb{C}^{n_1\times r}\right\}. \tag{13}$$

By $\mathcal{P}_{T_X}$ we will denote the orthogonal projection onto T_X, by $\mathcal{P}_{T_X^\perp} = \mathrm{Id} - \mathcal{P}_{T_X}$ the projection onto its orthogonal complement.

Lemma 4.1. *Let $X \in \mathbb{C}^{n_1 \times n_2} \setminus \{0\}$ be a matrix of rank r with corresponding singular value decomposition $X = U\Sigma V^*$. Then*

$$\overline{\mathcal{K}_* (X)} = \left\{ Z \in \mathbb{C}^{n_1 \times n_2} : -Re\left(\langle UV^*, Z \rangle_F\right) \geq \|\mathcal{P}_{T_X^\perp}(Z)\|_* \right\},$$

where $\overline{\mathcal{K}_* (X)}$ denotes the topological closure of $\mathcal{K}_* (X)$.

Remark 4.2. *Lemma 4.1 is similar to well-known results in convex optimization and may be known to the community. As we could not find it in the literature in this form, we decided to include a proof for completeness.*

The proof of Lemma 4.1 relies on the duality between the descent cone and the subdifferential of a convex function. In the following we will denote by $\partial \|\cdot\|_* (X)$ the subdifferential of the nuclear norm at the point $X \in \mathbb{C}^{n_1 \times n_2}$. We will use that a characterization of $\partial \|\cdot\|_*$ is well-known [61]. Namely, for all $X \in \mathbb{C}^{n_1 \times n_2}$ with corresponding singular value decomposition $X = U\Sigma V^*$ it holds that

$$\partial \|\cdot\|_* (X) = \left\{ W \in \mathbb{C}^{n_1 \times n_2} : \ \mathcal{P}_{T_X} W = UV^*, \ \|\mathcal{P}_{T_X^\perp} W\| \leq 1 \right\}. \tag{14}$$

Proof. Recall that for a set of matrices $\mathcal{V} \subset \mathbb{C}^{n_1 \times n_2}$ its polar cone $\mathcal{V}^\circ$ is defined by

$$\mathcal{V}^\circ := \left\{ Z \in \mathbb{C}^{n_1 \times n_2} : \mathrm{Re}\left(\langle W, Z \rangle_F\right) \leq 0 \text{ for all } W \in \mathcal{V} \right\}.$$

For all $X \in \mathbb{C}^{n_1 \times n_2} \setminus \{0\}$ we have the following polarity relation between the descent cone and the subdifferential

$$\mathcal{K}_* (X)^\circ = \overline{\{\lambda W : \ \lambda \geq 0, \ W \in \partial \|\cdot\|_* (X)\}}.$$

For sets and functions defined in $\mathbb{R}^n$ with the usual Euclidean inner product, this is [54, Theorem 23.7]. The complex case directly follows, as $\mathbb{C}^{n_1 \times n_2}$ with the inner product $\mathrm{Re}\left(\langle \cdot, \cdot \rangle_F\right)$ can be identified with an $2n_1 n_2$-dimensional real-valued vector space with standard Euclidean inner product.

It follows from the bipolar theorem (see, e.g., [5, p. 53]) that

$$\overline{\mathcal{K}_* (X)} = \left(\partial \|\cdot\|_* (X)\right)^\circ.$$

Hence, in order to complete the proof it is sufficient to show that

$$\left\{ Z \in \mathbb{C}^{n_1 \times n_2} : -\mathrm{Re}\left(\langle UV^*, Z \rangle_F\right) \geq \|\mathcal{P}_{T_X^\perp}(Z)\|_* \right\} = \left(\partial \|\cdot\|_* (X)\right)^\circ = \overline{\mathrm{cone}\left(\partial \|\cdot\|_* (X)\right)}. \tag{15}$$

First, suppose that $Z \in \mathbb{C}^{n_1 \times n_2}$ satisfies $-\mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) \geq \|\mathcal{P}_{T_X^\perp}(Z)\|_*$. We have to show that $\mathrm{Re}\left(\langle W, Z\rangle_F\right) \leq 0$ for all $W \in \partial\|\cdot\|_*(X)$. Indeed,

$$\begin{aligned}
\mathrm{Re}\left(\langle W, Z\rangle_F\right) &= \mathrm{Re}\left(\langle \mathcal{P}_{T_X} W, Z\rangle_F\right) + \mathrm{Re}\left(\langle \mathcal{P}_{T_X^\perp} W, Z\rangle_F\right) \\
&= \mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) + \mathrm{Re}\left(\langle \mathcal{P}_{T_X^\perp} W, \mathcal{P}_{T_X^\perp} Z\rangle_F\right) \\
&\leq \mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) + \|\mathcal{P}_{T_X^\perp} W\| \|\mathcal{P}_{T_X^\perp} Z\|_* \\
&\leq \mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) + \|\mathcal{P}_{T_X^\perp} Z\|_* \\
&\leq 0.
\end{aligned}$$

In the first inequality we have used that the spectral norm is the dual norm of the nuclear norm. The second inequality follows from $\|\mathcal{P}_{T_X^\perp} W\| \leq 1$. Hence, we have shown that $Z \in (\partial\|\cdot\|_*(X))^\circ$. Next, let $Z \in (\partial\|\cdot\|_*(X))^\circ$ be arbitrary. Choose $\tilde{W} \in T_X^\perp$ such that $\mathrm{Re}\left(\langle \tilde{W}, Z\rangle_F\right) = \|\mathcal{P}_{T_X^\perp}(Z)\|_*$ and $\|\tilde{W}\| \leq 1$. Then by (14) it follows that $UV^* + \tilde{W} \in \partial\|\cdot\|_*(X)$ and as $Z \in (\partial\|\cdot\|_*(X))^\circ$ we obtain that

$$\begin{aligned}
0 &\geq \mathrm{Re}\left(\langle UV^* + \tilde{W}, Z\rangle_F\right) \\
&= \mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) + \|\mathcal{P}_{T_X^\perp}(Z)\|_*.
\end{aligned}$$

This shows that $-\mathrm{Re}\left(\langle UV^*, Z\rangle_F\right) \geq \|\mathcal{P}_{T_X^\perp}(Z)\|_*$. Hence, we have verified (15), which completes the proof. □

4.2. Upper bound for blind deconvolution

The goal of this section is to prove Proposition 3.3. For that we need the following lemma, which is a consequence of the concentration of measure theorem for Lipschitz functions. (For a proof of the real-valued case see, e.g., [60, Lemma 5.3.2]. The complex-case can be shown analogously.)

Lemma 4.3. *Let $P : \mathbb{C}^n \to \mathbb{C}^n$ be a random projection onto a k-dimensional subspace, which is uniformly distributed in the Grassmannian $Gr(k, \mathbb{C}^n)$. Fix $z \in \mathbb{C}^n$. Then for all $\varepsilon > 0$ with probability at least $1 - 2e^{-\tilde{c}k\varepsilon^2}$ we have that*

$$(1-\varepsilon)\frac{k}{n}\|z\|^2 \leq \|Pz\|^2 \leq (1+\varepsilon)\frac{k}{n}\|z\|^2,$$

where $\tilde{c} > 0$ is absolute some constant.

Proof of Proposition 3.3. By assumption we may write $L = \eta K$ for some $\eta \in \mathbb{N}$. Let $\{e_i\}_{i=1}^K$ an arbitrary orthonormal basis of $\mathbb{C}^K$. Choose $B \in \mathbb{C}^{L \times K}$ such that the matrix $\overline{FB}$ consists of rows $\{b_{k,i}\}_{k,i}^{K,\eta}$ defined by $b_{k,i} = \frac{1}{\sqrt{\eta}} e_k$ for all $k \in [K]$ and all $i \in [\eta]$. We observe that $(FB)^*(FB) = \mathrm{Id}_K$, which implies $B^*B = \mathrm{Id}_K$. Furthermore, we have $\|b_{k,i}\| = \sqrt{\frac{1}{\eta}}$ for all $k \in [K]$ and $i \in [\eta]$, which implies that also all rows of FB have

equal norm.

Without loss of generality we assume that $\|h_0\| = \|m_0\| = 1$ as rescaling does not change the descent cone $\mathcal{K}_*(h_0 m_0^*)$. For the proof we will condition on two events. The first event states that

$$\|\mathcal{A}(h_0 m_0^*)\| \leq 2\|h_0 m_0^*\|_F, \tag{16}$$

which by the Bernstein inequality (see, e.g., [60]) is fulfilled with probability at least $1 - \exp\left(-cL/\mu^2\right)$, where $c > 0$ is some numerical constant. To formulate the second event, we define for all $i \in [K]$ the subspaces

$$D_i := \operatorname{span}\{c_{i,1}; \ldots; c_{i,\eta}\} \subset \mathbb{C}^N$$

and denote by $m_i^{\|}$ the orthogonal projection of m_0 onto D_i and by $m_i^{\perp}$ the projection onto $D_i^{\perp}$, the orthogonal complement of D_i. Note that $D_i \subset \mathbb{C}^N$ is a random subspace of dimension $\frac{L}{K}$, distributed uniformly over the Grassmannian $Gr\left(\frac{L}{K}, \mathbb{C}^N\right)$ due to the rotation invariance of $\mathcal{CN}(0,1)$. Hence, as $\|m_0\| = 1$ Lemma 4.3 yields that for fixed $i \in [K]$ with probability at least $1 - 2\exp\left(-\frac{\tilde{c}L}{4K}\right)$ one has

$$\frac{L}{2KN} \leq \left\|m_i^{\|}\right\|^2 \leq \frac{3L}{2KN}. \tag{17}$$

As the matrix C is Gaussian, the different subspaces D_i's and hence also the random vectors $\left\{m_i^{\|}\right\}_{i=1}^K$ are independent, so with probability at least

$$1 - \left(2\exp\left(-\frac{\tilde{c}L}{4K}\right)\right)^K = 1 - \exp\left(K\log 2 - \frac{\tilde{c}L}{4}\right)$$

there exists at least one $k \in [K]$ such that (17) holds (with $k = i$). Also note that

$$1 - \exp\left(K\log 2 - \frac{\tilde{c}L}{4}\right) \geq 1 - \exp\left(-\frac{C_1}{2}L\right),$$

which for $C_1 = \frac{8\log 2}{\tilde{c}}$ follows from assumption (9).

To summarize, we have shown that the two events $\mathcal{E}_1 := \{\|\mathcal{A}(h_0 m_0^*)\| \leq 2\|h_0 m_0^*\|_F\}$ and

$$\mathcal{E}_2 := \left\{\exists k \in [K] : \ \frac{L}{2KN} \leq \left\|m_k^{\|}\right\|^2 \leq \frac{3L}{2KN}\right\}$$

happen with probability at least $1 - \mathcal{O}\left(\exp\left(-C_2 L/\mu^2\right)\right)$, where $C_2 > 0$ is an appropriately chosen constant.

Conditional on $\mathcal{E}_1$ and $\mathcal{E}_2$, we will construct $Z \in \mathbb{C}^{K\times N}$ (depending on the vectors $(c_{i,k})_{i,k}$) such that $Z \in \overline{\mathcal{K}_*(h_0 m_0^*)} \setminus \{0\}$ and such that the inequality

$$\|\mathcal{A}(Z)\| < 8\sqrt{\frac{L}{KN}}\|Z\|_F \tag{18}$$

is satisfied. Note that this will complete the proof. Indeed, by definition of the closure and the continuity of $\mathcal{A}$ this implies that there exists $\tilde{Z} \in \mathcal{K}_*(X_0)$ such that

$$\frac{\|\mathcal{A}\left(\tilde{Z}\right)\|}{\|\tilde{Z}\|_F} \leq 8\sqrt{\frac{L}{KN}},$$

which by the definition of $\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(h_0 m_0^*))$ implies that (10) holds with constant $C_3 = 8$.

To construct Z satisfying (18), define

$$W := -\frac{\langle h_0, e_i\rangle}{\|m_i^\perp\||\langle e_i, h_0\rangle|} e_i \left(m_i^\perp\right)^*,$$

where $i \in [K]$ is chosen to satisfy (17). It follows directly from the definition of W that $\|W\|_F = 1$. We observe that $\mathcal{A}(W) = 0$ as for each $k \in [K]$ and $i \in [\eta]$ we either have $\langle e_k, b_{k,i}\rangle = 0$ or $\langle m_{k,i}^\perp, c_{k,i}\rangle = 0$. Denote by $T = T_{X_0}$ the tangent space of the manifold of rank-one matrices at $X_0 = h_0 m_0^*$ as defined in (13) and by $\mathcal{P}_T$ and $\mathcal{P}_{T^\perp}$ the corresponding orthogonal projections. It follows that

$$\begin{aligned}
\|\mathcal{P}_{T^\perp} W\|_* &= \left\|\mathcal{P}_{T^\perp}\left(e_i\left(\frac{m_i^\perp}{\|m_i^\perp\|}\right)^*\right)\right\|_F \\
&= \left\|P_{h_0^\perp} e_i\right\| \left\|P_{m_0^\perp}\left(\frac{m_i^\perp}{\|m_i^\perp\|}\right)\right\| \\
&= \sqrt{1 - |\langle h_0, e_i\rangle|^2}\sqrt{1 - \left|\langle m_0, \frac{m_i^\perp}{\|m_i^\perp\|}\rangle\right|^2} \\
&\leq \sqrt{1 - \left|\langle m_0, \frac{m_i^\perp}{\|m_i^\perp\|}\rangle\right|^2} \\
&= \sqrt{1 - \|m_i^\perp\|^2} \\
&= \|m_i^\parallel\| \\
&\overset{(17)}{\leq} \sqrt{\frac{3L}{2KN}}.
\end{aligned} \tag{19}$$

Thus we have shown that W, an element of the null space of $\mathcal{A}$, is close to the tangent space T. We will now show that for $\beta = 2\sqrt{\frac{L}{KN}}$

$$Z := -\beta h_0 m_0^* + W,$$

lies in the closure of the descent cone $\overline{\mathcal{K}_*(h_0 m_0^*)}$. For that, we observe that

$$\begin{aligned}
-\mathrm{Re}\left(\langle Z, h_0 m_0^* \rangle\right) &= \beta - \mathrm{Re}\left(\langle W, h_0 m_0^* \rangle_F\right) \\
&= \beta + \mathrm{Re}\left(\frac{\langle h_0, e_i \rangle}{\|m_i^\perp\| |\langle h_0, e_i \rangle|} \langle e_i \left(m_i^\perp\right)^*, h_0 m_0^* \rangle_F\right) \\
&= \beta + \mathrm{Re}\left(\frac{\langle h_0, e_i \rangle}{\|m_i^\perp\| |\langle h_0, e_i \rangle|} \langle e_i, h_0 \rangle \langle m_0, m_i^\perp \rangle\right) \\
&= \beta + |\langle h_0, e_i \rangle| \|m_i^\perp\| \\
&\geq \beta \\
&\overset{(19)}{\geq} \|\mathcal{P}_{T^\perp} W\|_* \\
&= \|\mathcal{P}_{T^\perp} Z\|_*
\end{aligned}$$

and hence Lemma 4.1 entails that $Z \in \overline{\mathcal{K}_*(h_0 m_0^*)}$. Moreover, note that by the triangle inequality and by the assumption $L \leq \frac{1}{16} KN$ it holds that

$$\begin{aligned}
\|Z\|_F \geq & \|W\|_F - \beta \\
= & 1 - 2\sqrt{\frac{L}{KN}} \\
> & \frac{1}{2}.
\end{aligned} \tag{20}$$

These observations together with $\mathcal{A}(W) = 0$ yield that

$$\|\mathcal{A}(Z)\| = \|\mathcal{A}(\beta h_0 m_0^*)\| \overset{(16)}{\leq} 2\beta = 4\sqrt{\frac{L}{KN}} \overset{(20)}{<} 8\sqrt{\frac{L}{KN}} \|Z\|_F.$$

This shows (18), as desired.

□

4.3. Upper bound for matrix completion

In this section, we prove Proposition 3.6. For that we introduce sets $\mathcal{N}_a$, $a \in [n_1]$, via

$$\mathcal{N}_a := \{b \in [n_2] : a = a_i \text{ and } b = b_i \text{ for some } i \in [m]\}.$$

That is, $\mathcal{N}_a$ contains all the indices of the ath row of the matrix X_0, which are observed by the measurements. Furthermore, define by $P_{\mathcal{N}_a} \in \mathbb{R}^{n_2 \times n_2}$ the projection onto the coordinates, which are contained in $\mathcal{N}_a$, i.e. $P_{\mathcal{N}_a} = \sum_{b \in \mathcal{N}_a} e_b e_b^*$. By $P_{\mathcal{N}_a^\perp} = \sum_{b \in [m] \setminus \mathcal{N}_a} e_b e_b^*$ we denote the coordinate projection onto $[m] \setminus \mathcal{N}_a$.

We need the following technical lemma.

Lemma 4.4. *Let $V \in \mathbb{R}^{n_2 \times r}$ be an isometry, i.e. $V^* V = Id$. Assume that*

$$m \geq C_1 r n_1 \mu^2(V) \log(2r). \tag{21}$$

Then with probability at least $1-\mathcal{O}\left(\exp\left(-\frac{m}{C_2 r\mu^2(V)}\right)\right)$ *there exists* $a\in[n_1]$ *such that*

$$\|P_{\mathcal{N}_a}V\|\leq\sqrt{\frac{2m}{n_1n_2}}.$$

C_1 *and* C_2 *are universal constants.*

Proof. For each $a\in[n_1]$ we set $\mathcal{I}_a:=\{i\in[m]:a_i=a\}$ and define the event

$$\mathcal{E}_a:=\left\{\|P_{\mathcal{N}_a}V\|^2\leq\frac{2m}{n_1n_2}\right\}.$$

We will first derive a lower bound for $\mathbb{P}\left(\mathcal{E}_a\middle|\mathcal{I}_a\right)$. For that we note that $\|P_{\mathcal{N}_a}V\|^2=\|V^*P_{\mathcal{N}_a}V\|$. Let $v_1,v_2,\ldots,v_{n_2}$ denote the rows of the matrix V. By definition of $\mathcal{N}_a$ and $\mathcal{I}_a$ it follows that for $X_i:=v_{b_i}v_{b_i}^*$

$$V^*P_{\mathcal{N}_a}V=\sum_{b\in\mathcal{N}_a}v_bv_b^*\preceq\sum_{i\in\mathcal{I}_a}v_{b_i}v_{b_i}^*=\sum_{i\in\mathcal{I}_a}X_i. \tag{22}$$

Here we write $A\preceq B$ for two symmetric matrices A and B, if and only if $B-A$ is positive semidefinite. By (22) it is sufficient to bound the probability of the event

$$\left\{\left\|\sum_{i\in\mathcal{I}_a}X_i\right\|\leq\frac{2m}{n_1n_2}\right\}\subset\mathcal{E}_a \tag{23}$$

conditionally on $\mathcal{I}_a$. To bound $\|\sum_{i\in\mathcal{I}_a}X_i\|$ we will use the matrix Bernstein inequality (see, e.g., [58, Theorem 6.1.1]) conditionally on $\mathcal{I}_a$, which requires as ingredients

$$\mathbb{E}\left[\sum_{i\in\mathcal{I}_a}X_i\,\middle|\mathcal{I}_a\right]=\frac{|\mathcal{I}_a|}{n_2}\sum_{b=1}^{n_2}v_bv_b^*=\frac{|\mathcal{I}_a|}{n_2}V^*V=\frac{|\mathcal{I}_a|}{n_2}\mathrm{Id},$$

an upper bound for $\sigma^2(\mathcal{I}_a):=\left\|\mathbb{E}\left[\sum_{i\in\mathcal{I}_a}\left(X_i-\mathbb{E}[X_i]\right)^2\middle|\mathcal{I}_a\right]\right\|$ and a constant $K>0$ such that $\|X_i-\mathbb{E}[X_i]\|\leq K$ almost surely. To bound $\sigma^2(\mathcal{I}_a)$ we note that

$$\begin{aligned}\mathbb{E}\left[\sum_{i\in\mathcal{I}_a}\left(X_i-\mathbb{E}[X_i]\right)^2\middle|\mathcal{I}_a\right]&\preceq\mathbb{E}\left[\sum_{i\in\mathcal{I}_a}X_i^2\middle|\mathcal{I}_a\right]\\&\preceq\left(\max_{b\in[n_2]}\|v_b\|^2\right)\mathbb{E}\left[\sum_{i\in\mathcal{I}_a}X_i\middle|\mathcal{I}_a\right]\\&=\frac{|\mathcal{I}_a|\left(\max_{b\in[n_2]}\|v_b\|^2\right)}{n_2}\mathrm{Id}\\&=\frac{|\mathcal{I}_a|\mu^2(V)\,r}{n_2^2}\mathrm{Id},\end{aligned}$$

where the fourth line is due to the definition of $\mu^2(V)$. This implies that

$$\sigma^2(\mathcal{I}_a) \leq \frac{|\mathcal{I}_a|\mu^2(V)\, r}{n_2^2}.$$

To find an appropriate $K > 0$ note that almost surely

$$\begin{aligned}
\left\|X_i - \mathbb{E}X_i\right\| &= \left\|X_i - \frac{1}{n_1 n_2}\mathrm{Id}\right\| \\
&\leq \max_{i\in[n_2]}\left\|v_i v_i^* - \frac{1}{n_1 n_2}\mathrm{Id}\right\| \\
&\leq \frac{1}{n_1 n_2} + \max_{i\in[n_2]}\left\|v_i\right\|^2 \\
&\leq \left(\frac{1}{n_1 r} + 1\right)\max_{i\in[n_2]}\left\|v_i\right\|^2 \\
&\leq \frac{2r}{n_2}\mu^2(V) =: K,
\end{aligned}$$

where in the fourth line we used that $\max_{i\in[n_2]}\|v_i\|^2 \geq \frac{r}{n_2} = \frac{1}{n_2}\sum_{i\in[n_2]}\|v_i\|^2$. Finally to apply Bernstein inequality we need that the X_i's are independent conditionally on $\mathcal{I}_a$, which follows from the fact that the a_i's and b_i's are drawn independently. With these ingredients the matrix Bernstein inequality yields that

$$\begin{aligned}
\mathbb{P}\left(\left\|\sum_{i\in\mathcal{I}_a} X_i - \frac{|\mathcal{I}_a|}{n_2}\mathrm{Id}\right\| \leq t \,\middle|\, \mathcal{I}_a\right) &\geq 1 - 2r\exp\left(-c\min\left(\frac{t^2}{\sigma^2(\mathcal{I}_a)}; \frac{t}{K}\right)\right) \\
&\geq 1 - 2r\exp\left(-c\min\left\{\frac{n_2^2 t^2}{|\mathcal{I}_a| r\mu^2(V)}; \frac{n_2 t}{2r\mu^2(V)}\right\}\right).
\end{aligned}$$

Setting $t = \frac{m}{2n_1 n_2}$ this implies that for fixed $a \in [n_1]$ it holds that

$$\begin{aligned}
\mathbb{P}\left(\left\|\sum_{i\in\mathcal{I}_a} X_i\right\| \leq \frac{m}{2n_1 n_2} + \frac{|\mathcal{I}_a|}{n_2}\,\middle|\, \mathcal{I}_a\right) &\geq \mathbb{P}\left(\left\|\sum_{i\in\mathcal{I}_a} X_i - \frac{|\mathcal{I}_a|}{n_2}\mathrm{Id}\right\| \leq \frac{m}{2n_1 n_2}\,\middle|\, \mathcal{I}_a\right) \\
&\geq 1 - 2r\exp\left(-\frac{cm}{4rn_1\mu^2(V)}\min\left\{\frac{m}{|\mathcal{I}_a|n_1}; 1\right\}\right).
\end{aligned} \tag{24}$$

To complete the proof we restrict our attention to $A := \left\{a \in [n_1] :\ |\mathcal{I}_a| \leq \frac{4m}{3n_1}\right\}$ as it follows from (23) that

$$\left\{|\mathcal{I}_a| \leq \frac{4m}{3n_1}\right\} \cap \left\{\left\|\sum_{i\in\mathcal{I}_a} X_i\right\| \leq \frac{m}{2n_1 n_2} + \frac{|\mathcal{I}_a|}{n_2}\right\} \subset \mathcal{E}_a, \tag{25}$$

and, consequently, for $a \in A$ we obtain that

$$\begin{aligned}\mathbb{P}\left(\mathcal{E}_a \middle| \mathcal{I}_a\right) &\overset{(25),(24),}{\geq} 1 - 2r\exp\left(-\frac{cm}{4rn_1\mu^2\left(V\right)}\min\left\{\frac{m}{|\mathcal{I}_a|n_1};1\right\}\right) \\ &\overset{a\in A}{\geq} 1 - 2r\exp\left(-\frac{3cm}{16rn_1\mu^2\left(V\right)}\right).\end{aligned} \tag{26}$$

As the $\mathcal{E}_a$'s only depend on $\{b_i\}_{i\in\mathcal{I}_a}$ and are hence independent conditionally on $\mathcal{I}_\alpha$, this implies that

$$\begin{aligned}\mathbb{P}\left(\bigcap_{a\in A}\mathcal{E}_a^c \;\middle|\; \{\mathcal{I}_a\}_{a=1}^{n_1}\right) &= \prod_{a\in A}\mathbb{P}\left(\mathcal{E}_a^c \;\middle|\; \{\mathcal{I}_a\}_{a=1}^{n_1}\right) \\ &= \prod_{a\in A}\mathbb{P}\left(\mathcal{E}_a^c \;\middle|\mathcal{I}_a\right) \\ &\overset{(26)}{\leq} \prod_{a\in A}\left(2r\exp\left(-\frac{3cm}{16rn_1\mu^2\left(V\right)}\right)\right) \\ &= \left(2r\exp\left(-\frac{3cm}{16rn_1\mu^2\left(V\right)}\right)\right)^{|A|} \\ &= \left(\exp\left(\log\left(2r\right) - \frac{3cm}{16rn_1\mu^2\left(V\right)}\right)\right)^{|A|} \\ &\leq \left(\exp\left(-\frac{m}{C_1'rn_1\mu^2\left(V\right)}\right)\right)^{|A|},\end{aligned}$$

where in the last line we have used assumption (21) with C_1 large enough. Furthermore, note that

$$m = \sum_{a=1}^{n_1}|\mathcal{I}_a| \geq \sum_{i\in[n_1]\setminus A}|\mathcal{I}_a| \geq \frac{4m}{3n_1}\left|\left[n_1\right]\setminus A\right| = \frac{4m}{3n_1}\left(n_1 - |A|\right)$$

implies that $|A| \geq \frac{n_1}{4}$ almost surely. Hence, it follows that

$$\mathbb{P}\left(\bigcap_{a\in A}\mathcal{E}_a^c \;\middle|\; \{\mathcal{I}_a\}_{a=1}^{n_1}\right) \leq \exp\left(-\frac{m}{C_2r\mu^2\left(V\right)}\right).$$

This shows that conditional on $\{\mathcal{I}_a\}_{a=1}^{n_1}$ we have that almost surely

$$\mathbb{P}\left(\bigcup_{a\in[n_1]}\mathcal{E}_a\middle|\;\{\mathcal{I}_a\}_{a=1}^{n_1}\right) \geq \mathbb{P}\left(\bigcup_{a\in A}\mathcal{E}_a\middle|\;\{\mathcal{I}_a\}_{a=1}^{n_1}\right) \geq 1 - \exp\left(-\frac{m}{C_2r\mu^2\left(V\right)}\right).$$

Taking expectations yields the claim.

□

Now we are prepared to give a proof of Proposition 3.6.

Proof of Proposition 3.6. For the proof we will condition on two events $\mathcal{E}_1$ and $\mathcal{E}_2$, which we will define in the following. The event $\mathcal{E}_1$ is defined by

$$\mathcal{E}_1 := \left\{ \|\mathcal{A}\left(UV^*\right)\|^2 \leq 2\|UV^*\|_F^2 = 2r \right\}. \tag{27}$$

Observe that

$$\|\mathcal{A}\left(UV^*\right)\|^2 = \frac{n_1 n_2}{m} \sum_{i=1}^{m} |\left(UV^*\right)_{a_i,b_i}|^2 = \sum_{i=1}^{m} X_i,$$

where we have set $X_i := \frac{n_1 n_2}{m} |\left(UV^*\right)_{a_i,b_i}|^2$. Note that one has almost surely that

$$\begin{aligned} X_i &= \frac{n_1 n_2}{m} \left(\sum_{k=1}^{r} U_{a_i,k} V_{b_i,k} \right)^2 \\ &\leq \frac{n_1 n_2}{m} \left(\sum_{k=1}^{r} |U_{a_i,k}|^2 \right) \left(\sum_{k=1}^{r} |V_{b_i,k}|^2 \right) \\ &\leq \frac{\mu^2\left(U\right) \mu^2\left(V\right) r^2}{m} \end{aligned}$$

where we have applied Cauchy-Schwarz and the definition of $\mu\left(U\right)$ and $\mu\left(V\right)$. Hence, one has

$$\begin{aligned} \mathbb{E}X_i &= \frac{\|UV^*\|_F^2}{m} = \frac{r}{m}, \\ \mathbb{E}X_i^2 &\leq \frac{\mu^2\left(U\right) \mu^2\left(V\right) r^2}{m} \mathbb{E}X_i = \frac{r^3 \mu^2\left(U\right) \mu^2\left(V\right)}{m^2} \end{aligned}$$

where we used $\|UV^*\|_F^2 = r$ and the previous estimate. Hence, by the Bernstein inequality (see, e.g., [60, Theorem 2.8.4]) we obtain that

$$\mathbb{P}\left(\left| \sum_{i=1}^{m} X_i - \|UV^*\|_F^2 \right| \leq t \right) \leq 2\exp\left(-c \min\left\{ \frac{t^2 m^2}{r^3 \mu^2\left(U\right) \mu^2\left(V\right)}; \frac{tm}{r^2 \mu^2\left(U\right) \mu^2\left(V\right)} \right\} \right).$$

By setting $t = \|UV^*\|_F^2 = r$ we observe that $\mathcal{E}_1$ holds with probability at least $1 - 2\exp\left(-\frac{cm}{r\mu^2(U)\mu^2(V)}\right)$.

The second event $\mathcal{E}_2$ is defined by

$$\mathcal{E}_2 := \left\{ \exists a \in [n_1] \text{ such that } \|P_{\mathcal{N}_a} V\| \leq \sqrt{\frac{2m}{n_1 n_2}} \right\}.$$

For C_1 in assumption (11) chosen large enough Lemma 4.4 then entails that $\mathbb{P}\left(\mathcal{E}_2\right) \geq 1 - \mathcal{O}\left(\exp\left(-\frac{cm}{r\mu^2(V)}\right)\right)$. Consequently, we can find $a \in [n_1]$ (depending on the random

sampling pattern) such that the condition defining $\mathcal{E}_2$ is satisfied.

Note that in order to prove Proposition 3.6 it is enough to find $Z \in \overline{\mathcal{K}_*(X_0)} \setminus \{0\}$ such that

$$\|\mathcal{A}(Z)\| < 8\sqrt{\frac{m}{rn_1n_2}}\|Z\|_F, \tag{28}$$

because by definition of the closure and the continuity of $\mathcal{A}$ this implies that there is a matrix $\tilde{Z} \in \mathcal{K}_*(X_0) \setminus \{0\}$ such that

$$\frac{\left\|\mathcal{A}\left(\tilde{Z}\right)\right\|}{\|\tilde{Z}\|_F} \leq 8\sqrt{\frac{m}{rn_1n_2}},$$

which implies (12) with constant $C_3 = 8$. In the following we will construct such a matrix Z. Let $x \in \mathbb{R}^r$ be a vector such that $\|x\| = 1$. Then for $a \in [n_1]$ as above we define the vector $w_a \in \mathbb{R}^{n_2}$ by

$$w_a := P_{\mathcal{N}_a^\perp} V x$$

and set

$$W := -\frac{\langle e_a w_a^*, UV^* \rangle_F}{|\langle e_a w_a^*, UV^* \rangle_F|} e_a w_a^* = \frac{\langle e_a w_a^*, UV^* \rangle_F}{|\langle e_a w_a^*, UV^* \rangle_F|} e_a x^* V^* P_{\mathcal{N}_a^\perp}.$$

It follows directly from the definition of $\mathcal{N}_a$ that $\mathcal{A}(W) = 0$. In the following let T be the tangent space of the manifold of rank-r matrices at X_0 as defined in (13). Furthermore, denote by $P_U = UU^*$ the orthogonal projection onto the column space of U and, analogously, by $P_V = VV^*$ the orthogonal projection onto the column space of V. Then we obtain that

$$\begin{aligned}
\|\mathcal{P}_{T^\perp} W\|_* &= \|\mathcal{P}_{T^\perp}(e_i w_a^*)\|_* \\
&= \|P_{U^\perp} e_i w_a^* P_{V^\perp}\|_* \\
&= \|P_{U^\perp} e_a w_a^* P_{V^\perp}\|_F \\
&= \|P_{U^\perp} e_a\| \|P_{V^\perp} w_a\| \\
&\leq \|P_{V^\perp} w_a\|,
\end{aligned}$$

where in the second equality we have used that $\mathcal{P}_{T^\perp} M = P_{U^\perp} M P_{V^\perp}$ for all $M \in \mathbb{R}^{n_1 \times n_2}$ and in the last line we used that $\|P_{U^\perp} e_i\| \leq 1$. Plugging in $w_a = P_{\mathcal{N}_a^\perp} V m$ it follows that

$$\begin{aligned}
\|\mathcal{P}_{T^\perp} W\|_* &\leq \|P_{V^\perp} P_{\mathcal{N}_a^\perp} V x\| \\
&= \|P_{V^\perp} P_{\mathcal{N}_a} V x\|,
\end{aligned}$$

where the last line is due to $P_{V^\perp} V x = 0$. The fact that $\|P_{V^\perp}\| \leq 1$ then yields that

$$\begin{aligned}
\|\mathcal{P}_{T^\perp} W\|_* &\leq \|P_{\mathcal{N}_a} V x\| \\
&\leq \|P_{\mathcal{N}_a} V\| \|x\| \\
&\leq \sqrt{\frac{2m}{n_1 n_2}}
\end{aligned} \tag{29}$$

where for the last line we used that $a \in [n_1]$ was chosen such that the condition in $\mathcal{E}_2$ holds. This shows that W is relatively close to T. Based on W we now aim to find $Z \in \overline{\mathcal{K}_* (X)}$ of the form

$$Z := W - \beta UV^*,$$

where $\beta > 0$ will be chosen in the following such that $Z \in \overline{\mathcal{K}_* (X_0)}$, which by Lemma 4.1 is equivalent to $-\langle UV^*, Z\rangle_F \geq \|\mathcal{P}_{T^\perp} Z\|_*$. First, we note that

$$\|\mathcal{P}_{T^\perp} Z\|_* = \|\mathcal{P}_{T^\perp} W\|_* \leq \sqrt{\frac{2m}{n_1 n_2}} \tag{30}$$

due to $\mathcal{P}_{T^\perp}(UV^*) = 0$ and the inequality chain (29). Furthermore, we have that

$$\begin{aligned} -\langle UV^*, Z\rangle_F &= r\beta - \langle UV^*, W\rangle_F \\ &= r\beta + \langle UV^*, \frac{\langle e_a w_a^*, UV^*\rangle_F}{|\langle e_a w_a^*, UV^*\rangle_F|} e_a w_a^*\rangle_F \\ &= r\beta + |\langle e_a w_a^*, UV^*\rangle_F| \\ &\geq r\beta. \end{aligned} \tag{31}$$

Hence, setting $\beta = 2\sqrt{\frac{m}{r^2 n_1 n_2}}$ and combining (30) and (31) it follows that $Z \in \overline{\mathcal{K}_* (X_0)}$. This Z also satisfies (28). To see that we observe that

$$\begin{aligned} \|Z\|_F &\geq \|e_a w_a^*\|_F - \beta \|UV^*\|_F \\ &= \|w_a\| - \beta\sqrt{r} \\ &= \|P_{\mathcal{N}_a^\perp} Vx\| - \beta\sqrt{r} \\ &= \sqrt{\|Vx\|^2 - \|P_{\mathcal{N}_a} Vx\|^2} - \beta\sqrt{r} \\ &\geq \sqrt{1 - \frac{2m}{n_1 n_2}} - 2\sqrt{\frac{m}{r n_1 n_2}} \\ &> \frac{1}{2}, \end{aligned}$$

where in the last line we used the assumption that $m \leq \frac{n_1 n_2}{32}$. Furthermore, from $\mathcal{A}(W) = 0$ it follows that

$$\|\mathcal{A}(Z)\| = \beta \|\mathcal{A}(UV^*)\| \overset{(27)}{\leq} \beta\sqrt{2r} = 2\sqrt{\frac{2m}{r n_1 n_2}} < 8\sqrt{\frac{m}{r n_1 n_2}} \|Z\|_F.$$

Combining the last two inequality chains implies (28), which completes the proof. □

4.4. Proof of Theorem 3.1 and Theorem 3.5

As already mentioned Theorem 3.1 can be deduced from Proposition 3.3 and Theorem 3.5 can be deduced from Proposition 3.6. We only show how to prove Theorem 3.1, as the proof of Theorem 3.5 is analogous.

Proof of Theorem 3.1. By Proposition 3.6 and the definition of the minimum conic singular value $\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(h_0 m_0^*))$ with probability at least $1 - \mathcal{O}\left(\exp\left(-\frac{K}{C_2\mu^2}\right)\right)$ there is a matrix $Z \in \mathcal{K}_*(h_0 m_0^*) \setminus \{0\}$ such that

$$\|\mathcal{A}(Z)\| \leq C_3 \sqrt{\frac{L}{KN}} \|Z\|_F. \tag{32}$$

and such that $\tilde{X}_t := h_0 m_0^* + tZ$ obeys $\|\tilde{X}_t\|_* \leq \|h_0 m_0^*\|_*$ for all $0 < t \leq 1$. Next, set $e_t = \frac{t}{2}\mathcal{A}(Z)$. Then for $y_t = \mathcal{A}(h_0 m_0^*) + e_t$ we have that

$$\|\mathcal{A}\left(\tilde{X}_t\right) - \tilde{y}_t\| = \|\mathcal{A}(tZ) - e\| = \frac{t}{2}\|\mathcal{A}(Z)\|.$$

Hence, by setting $\tau_0 := \frac{\|\mathcal{A}(Z)\|}{2}$ we observe that $\|\mathcal{A}\left(\hat{X}_t\right) - y_t\| = t\tau_0$. Furthermore, note that

$$\|\tilde{X}_t - X_0\|_F = \|tZ\|_F \overset{(32)}{\geq} \frac{t}{C_3}\sqrt{\frac{KN}{L}}\|\mathcal{A}(Z)\| = \frac{2t\tau_0}{C_3}\sqrt{\frac{KN}{L}}$$

Now let $0 < \tau \leq \tau_0$. Then by setting $t = \frac{\tau}{\tau_0}$, $\tilde{X} := \tilde{X}_t$, $y := y_t$, and $e := e_t$ the desired claim follows. □

5. Stability of blind deconvolution

5.1. Outline of the proof and main ideas

The goal of this section is to prove Theorem 3.7. We first give a proof sketch and present the key ideas. We have seen in Proposition 3.3 that for certain isometries $B \in \mathbb{C}^{L\times K}$ with high probability one has that $\lambda_{\min}(\mathcal{A}, \mathcal{K}_*(h_0 m_0^*)) \lesssim \sqrt{\frac{L}{KN}}$. Hence, applying Theorem 2.4 cannot lead to very strong error estimates. However, if we closely inspect the proof of Proposition 3.3 we observe the following. Again, denote by T the tangent space of the manifold of rank-1 matrices at point $h_0 m_0^*$ as defined in (13) and assume that $\|h_0\| = \|m_0\| = 1$. By construction we have that $Z = W - \beta h_0 m_0^*$, where $W = \frac{\langle h_0, e_i\rangle}{\|m_i^\perp\||\langle h_0, e_i\rangle|} e_i \left(m_i^\perp\right)^*$. This implies that

$$\begin{aligned}
|\langle Z, h_0 m_0^*\rangle_F| &\leq \left|\langle \frac{\langle h_0, e_i\rangle}{\|m_i^\perp\||\langle h_0, e_i\rangle|} e_i \left(m_i^\perp\right)^*, h_0 m_0^*\rangle_F\right| + \beta\|h_0 m_0^*\|_F^2 \\
&= |\langle e_i, h_0\rangle|\|m_i^\perp\| + \beta\|h_0 m_0^*\|_F^2 \\
&\leq \frac{\mu}{\sqrt{L}} + \beta \lesssim \frac{\mu}{\sqrt{L}} + \sqrt{\frac{L}{KN}},
\end{aligned}$$

where we have used the triangle inequality in the first line and the definition of $m_i^\perp$ in the second line. In the third line we used that $\|m_i^\perp\| \leq 1$, $|\langle e_i, h_0\rangle| \leq \frac{\mu}{\sqrt{L}}$, and $\beta = \sqrt{\frac{L}{KN}}$.

As $\|Z\|_F \gtrsim 1$ this implies that $\frac{\langle Z, h_0 m_0^* \rangle_F}{\|Z\|_F \|h_0 m_0^*\|_F}$ is quite small, meaning that Z and $h_0 m_0^*$ are almost orthogonal to each other. All the descent directions Z with this property, however, have in common that the admissible descent step size

$$t_0 := \max \{t > 0 : \|h_0 m_0^* + tZ\|_* \leq \|h_0 m_0^*\|_*\}$$

is necessarily very small. Geometrically this corresponds to the fact that the nuclear norm ball is curved near X_0, which is why its near-tangential behaviour only holds locally. This will be made precise in Lemma 5.7 below. For this reason the idea of the proof of Theorem 3.7 is to split the descent cone into two parts. One part will consist of all the matrices aligned with $h_0 m_0^*$. The second part will consist of all remaining matrices, which are almost orthogonal to $h_0 m_0^*$. As mentioned above, these matrices must necessarily be close to T. The first part is captured by the set $\mathcal{E}_{\mu,\delta}$ with

$$\mathcal{E}_{\mu,\delta} := \bigcup_{h_0 \in \mathcal{H}_\mu, m_0 \in \mathbb{C}^N} \left\{ Z \in \mathcal{K}_* \left(h_0 m_0^*\right) : \delta \leq \frac{-\mathrm{Re}\left(\langle Z, h_0 m_0^* \rangle_F\right)}{\|h_0 m_0^*\|_F} \text{ and } \|Z\|_F = 1 \right\},$$

where $\delta > 0$. In Section 5.2 we will show that with high probability it holds that

$$\inf_{Z \in \mathcal{E}_{\mu,\delta}} \|\mathcal{A}(Z)\| \gtrsim \frac{\delta^2}{\log^2(L)\mu^2}. \tag{33}$$

Hence, if we have for the minimizer $\hat{X}$ of (3) that $\hat{X} - h_0 m_0^*$ is an element of the conic hull of $\mathcal{E}_{\mu,\delta}$ we can proceed similarly as in [11] to obtain near-optimal error bounds. Let us briefly explain which property of $\mathcal{E}_{\mu,\delta}$ allows us to show (33). We define for any matrix $W \in \mathbb{C}^{K \times N}$ its $\|\cdot\|_{B_1}$-norm by

$$\|W\|_{B_1} := \sum_{\ell=1}^{L} \|W^* b_\ell\|.$$

In other words $\|W\|_{B_1}$ is the ℓ_1-norm of the vector $(\|W^* b_\ell\|)_{\ell=1}^L$. We show in Lemma 5.5 below that all $Z \in \mathcal{E}_{\mu,\delta}$ have rather large $\|\cdot\|_{B_1}$-norm, which entails that the mass of the vector $(\|W^* b_\ell\|)_{\ell=1}^L$ cannot be concentrated on only very few entries. This in turn will allow us to employ a non-i.i.d. version of Mendelson's smalll bal method [35], allowing us to show a lower bound for (33), see Lemma 5.6 below. To understand the behaviour on the second part recall from Proposition 3.3 that for matrices $Z/\|Z\|_F \in \mathcal{K}_*\left(h_0 m_0^*\right) \setminus \mathcal{E}_{\mu,\delta}$ the quantity $\|\mathcal{A}(Z)\|$ may be quite small, so a uniform bound is not feasible. However, as Z is almost orthogonal to $h_0 m_0^*$, also $\|\mathcal{P}_{T^\perp} Z\|_*$ must be rather small because of the characterization of the descent cone, Lemma 4.1. Hence, Z is close to the tangent space and is almost orthogonal to $h_0 m_0^*$. For that reason, whenever $\|h_0 m_0^* + tZ\|_* \leq \|h_0 m_0^*\|_*$ holds, the cylindrical shape of the nuclear norm ball implies that $t > 0$ is small. This fact is captured by Lemma 5.7 below. Theorem 3.3 can then be proven by combining inequality (33) and Lemma 5.7, see Section 5.3.

5.2. A lower bound for the minimum conic singular value

First we recall the notion of Gaussian width (see, e.g., [60]).

Definition 5.1. *For a set $\mathcal{E} \subset \mathbb{C}^{K\times N}$ its Gaussian width is defined by*

$$\omega(\mathcal{E}) := \mathbb{E}\left[\sup_{X\in\mathcal{E}} Re\left(\langle X, G\rangle_F\right)\right],$$

where $G \in \mathbb{C}^{K\times N}$ is a matrix, whose entries are independent and identically distributed random variables with distribution $\mathcal{CN}(0,1)$.

This definition allows us to state the following lemma, which is important for our analysis of the conic singular value. It relies on a uniform lower bound on the number of measurements whose magnitude is larger than a certain constant and is a variant of Theorem 2.1 in [35].

Lemma 5.2. *Let $\mathcal{E} \subset \mathbb{C}^{K\times N}$ be a symmetric set, i.e., $\mathcal{E} = -\mathcal{E}$. For all $\xi > 0$ and $t > 0$ it holds with probability at least $1 - \exp\left(-2t^2\right)$ that*

$$\inf_{X\in\mathcal{E}} \Big| \{\ell \in [L] : \ |\langle b_\ell c_\ell^*, X\rangle_F| \geq \xi\} \Big| \geq \inf_{X\in\mathcal{E}} \left(\sum_{\ell=1}^{L} \mathbb{P}\left(|\langle b_\ell c_\ell^*, X\rangle_F| \geq 2\xi\right)\right) - \frac{4\omega(\mathcal{E})}{\xi} - t\sqrt{L}. \tag{34}$$

Here $\varepsilon_1, \ldots, \varepsilon_L$ are independent Rademacher variables, i.e., random variables which take the two values ± 1 each with probability $\frac{1}{2}$.

The proof of the Lemma 5.2 is based on a variant of Mendelson's small-ball method and proceeds in analogy to [35]. We have deferred a detailed proof to Appendix A. In order to apply Lemma 5.2 we need to estimate the first term of the right-hand side of Lemma 5.2. Such an estimate can be derived using the Paley-Zygmund inequality as in [35]. For the sake of completeness we have included a proof in Appendix B.

Lemma 5.3. *Let $X \in \mathbb{C}^{K\times N}$ be an arbitrary matrix. Then for all $\xi > 0$*

$$\sum_{\ell=1}^{L} \mathbb{P}\left(|\langle b_\ell c_\ell^*, X\rangle_F| \geq 2\xi\right) \geq \frac{9}{32} \Big| \{\ell \in [L] : \ \|X^* b_\ell\| \geq 4\xi\} \Big|.$$

In order to use Lemma 5.3 we need a lower bound for $|\{\ell \in [L] : \ \|X^* b_\ell\| \geq \xi\}|$. This will be achieved by the next lemma. For the statement of this lemma and its proof we will need to introduce the following notion. We define for any matrix $W \in \mathbb{C}^{K\times N}$ its $\|\cdot\|_{B_{1,w}}$-quasinorm by

$$\|W\|_{B_{1,w}} := \sup_{\xi\geq 0} \xi \Big| \{\ell \in [L] : \|W^* b_\ell\| \geq \xi\} \Big|.$$

That is, $\|W\|_{B_{1,w}}$ is the weak ℓ_1-norm of the vector $(\|W^* b_\ell\|)_{\ell=1}^L$. (For a more detailed discussion of the weak ℓ_1-norm see, e.g., [22].) A direct consequence of this interpretation is the inequality (see, e.g., [20, Proposition 2.10 and Exercise 2.4])

$$\|W\|_{B_{1,w}} \leq \|W\|_{B_1} \leq \log(eL)\, \|W\|_{B_{1,w}}. \tag{35}$$

Lemma 5.4. *Let $Z \in \mathbb{C}^{K\times N}$ such that $\|Z\|_F = 1$. Then it holds that*

$$\Big|\left\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \frac{\|Z\|_{B_1}}{L\log(eL)}\right\}\Big| \geq \frac{\|Z\|_{B_1}^2}{\log^2(eL)}.$$

Proof of Lemma 5.4. Choose ξ^* such that

$$\|Z\|_{B_{1,w}} = \xi^* |\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}|. \tag{36}$$

As $|\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}| \leq L$ it follows that

$$\xi^* \geq \frac{\|Z\|_{B_{1,w}}}{L} \geq \frac{\|Z\|_{B_1}}{L\log(eL)}, \tag{37}$$

where we also used inequality (35). We observe that

$$\begin{aligned}
1 &= \|Z\|_F^2 \\
&= \sum_{\ell=1}^{L} \|Z^* b_\ell\|^2 \\
&\geq \xi^{*2} |\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}| \\
&\overset{(36)}{=} \frac{\|Z\|_{B_{1,w}}^2}{|\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}|} \\
&\overset{(35)}{\geq} \frac{\|Z\|_{B_1}^2}{\log^2(eL)\,|\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}|},
\end{aligned}$$

where for the second equality we used the identity $\sum_{\ell=1}^{L} b_\ell b_\ell^* = \mathrm{Id}$. Using inequality (37) and rearranging terms it follows that

$$\Big|\left\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \frac{\|Z\|_{B_1}}{L\log(eL)}\right\}\Big| \geq \big|\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \xi^*\}\big| \geq \frac{\|Z\|_{B_1}^2}{\log^2(eL)},$$

which completes the proof. □

The next lemma gives a bound on $\inf_{Z\in\mathcal{E}_{\mu,\delta}} \|Z\|_{B_1}$.

Lemma 5.5. *It holds that*

$$\inf_{Z\in\mathcal{E}_{\mu,\delta}} \|Z\|_{B_1} \geq \frac{\delta\sqrt{L}}{\mu}.$$

Proof. Let $Z \in \mathcal{E}_{\mu,\delta}$ be arbitrary. By definition of $\mathcal{E}_{\mu,\delta}$ there is $h_0 \in \mathcal{H}_\mu$ and $m_0 \in \mathbb{C}^N$ such that $Z \in \mathcal{K}_*(h_0 m_0^*)$ and such that the inequality

$$\delta \leq \frac{-\mathrm{Re}\left(\langle Z, h_0 m_0^*\rangle_F\right)}{\|h_0 m_0^*\|_F}$$

holds. It follows that

$$\begin{aligned}\delta &\leq \frac{-\mathrm{Re}\left(\langle Zm_0, h_0\rangle\right)}{\|h_0 m_0^*\|_F} \\ &= \frac{-\sum_{\ell=1}^{L} \mathrm{Re}\left(\langle Zm_0, b_\ell\rangle\langle b_\ell, h_0\rangle\right)}{\|h_0 m_0^*\|_F} \\ &\leq \frac{\left(\max\limits_{\ell\in[L]} |\langle h_0, b_\ell\rangle|\right) \sum_{\ell=1}^{L} |\langle Zm_0, b_\ell\rangle|}{\|h_0 m_0^*\|_F},\end{aligned}$$

where for the second equality we have used that $\sum_{\ell=1}^{L} b_\ell b_\ell^* = \mathrm{Id}$. Note that for all $\ell \in [L]$ it holds that

$$\frac{|\langle Zm_0, b_\ell\rangle|}{\|m_0\|} \leq \|Z^* b_\ell\|.$$

Hence, by the previous inequality chain it follows that

$$\delta \leq \frac{\max\limits_{\ell\in[L]} |\langle h_0, b_\ell\rangle|}{\|h_0\|} \sum_{\ell=1}^{L} \|Z^* b_\ell\| \leq \frac{\mu}{\sqrt{L}} \|Z\|_{B_1},$$

where in the last inequality we used the definition of μ and $\|Z\|_{B_1}$. Rearranging terms and taking the infimum over all $Z \in \mathcal{E}_{\mu,\delta}$ yields the desired inequality. □

Having gathered all the necessary ingredients we can state and prove the main lemma of this section.

Lemma 5.6. *Let $\delta > 0$. Assume that*

$$L \geq C_1 \left(\frac{\mu}{\delta}\right)^6 (K+N) \ \log^6(eL), \tag{38}$$

Then with probability at least $1 - \exp\left(-\frac{L\delta^4}{C_2 \log^4(eL)\mu^4}\right)$ it holds that

$$\inf_{Z\in\mathcal{E}_{\mu,\delta}} \|\mathcal{A}(Z)\| \gtrsim \frac{\delta^2}{\log^2(L)\,\mu^2}. \tag{39}$$

C_1 and C_2 are absolute constants.

Proof. Our goal is to apply Lemma 5.2. In order to apply it we first derive a lower bound for the first term on the right-hand side of inequality (34). For that recall that by Lemma 5.5 it holds that

$$\inf_{Z\in\mathcal{E}_{\mu,\delta}} \|Z\|_{B_1} \geq \frac{\delta\sqrt{L}}{\mu}. \tag{40}$$

Thus, for any $Z \in \mathcal{E}_{\mu,\delta}$ we obtain that

$$\begin{aligned}\left|\left\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \frac{\delta}{\mu\sqrt{L}\log(eL)}\right\}\right| &\geq \left|\left\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \frac{\|Z\|_{B_1}}{L\log(eL)}\right\}\right| \\ &\geq \frac{\|Z\|_{B_1}^2}{\log^2(eL)} \\ &\geq \frac{\delta^2 L}{\mu^2 \log^2(eL)},\end{aligned}$$

where the first inequality follows from (40), the second one is due to Lemma 5.4, and the third one follows again from (40). Hence, by Lemma 5.3 applied with $\xi = \frac{\delta}{4\sqrt{L}\ln(eL)\mu}$ we finally obtain that

$$\begin{aligned}&\inf_{Z\in\mathcal{E}_{\mu,\delta}} \left(\sum_{\ell=1}^{L} \mathbb{P}\left(|\langle b_\ell c_\ell^*, Z\rangle_F| \geq \frac{\delta}{2\sqrt{L}\ln(eL)\,\mu}\right)\right) \\ \geq &\frac{9}{32}\inf_{Z\in\mathcal{E}_{\mu,\delta}} \left|\left\{\ell \in [L]:\ \|Z^* b_\ell\| \geq \frac{\delta}{\sqrt{L}\ln(eL)\,\mu}\right\}\right| \\ \geq &\frac{9\delta^2 L}{32\mu^2\log^2(eL)}.\end{aligned} \tag{41}$$

Next, we need an upper bound for the Gaussian width. For that, we first observe that

$$\mathcal{E}_{\mu,\delta} \subset \left(\bigcup_{h_0\in\mathbb{C}^K, m_0\in\mathbb{C}^N} \mathcal{K}_*(h_0 m_0^*)\right) \cap \left\{Z \in \mathbb{C}^{K\times N}:\ \|Z\|_F = 1\right\} =: \mathcal{E}.$$

The Gaussian width of $\mathcal{E}$ has been bounded in [29, Lemma 4.1], combined with the monotonicity of the Gaussian width their results yields that

$$\omega(\mathcal{E}_{\mu,\delta}) \leq \omega(\mathcal{E}) \leq 2\sqrt{(K+N)}. \tag{42}$$

Thus for $\xi = \frac{\delta}{4\sqrt{L}\log(eL)\mu}$ we obtain from Lemma 5.2 together with (41), (42) that with probability at least $1 - \exp\left(-2t^2\right)$ it holds that

$$\begin{aligned}\inf_{Z\in\mathcal{E}_{\mu,\delta}} \left|\left\{\ell \in [L]:\ |\langle b_\ell c_\ell^*, Z\rangle_F| \geq \xi\right\}\right| &\geq \frac{9L\delta^2}{32\log^2(eL)\,\mu^2} - \frac{2\log(eL)\,\mu\sqrt{L(K+N)}}{\delta} - t\sqrt{L} \\ &\geq \frac{9L\delta^2}{64\log^2(eL)\,\mu^2} - t\sqrt{L},\end{aligned}$$

where the second inequality follows from assumption (38), if the constant $C_1 > 0$ is chosen large enough. Consequently, setting $t = \frac{9\delta^2\sqrt{L}}{128\log^2(eL)\mu^2}$ and recalling that $(\mathcal{A}(Z))(\ell) = \langle b_\ell c_\ell^*, Z\rangle_F$ we have that with probability at least $1 - \exp\left(-\frac{L\delta^4}{C_2\log^4(eL)\mu^4}\right)$ with C_2 chosen appropriately

$$\inf_{Z\in\mathcal{E}_{\mu,\delta}} \left|\left\{\ell \in [L]:\ |\mathcal{A}(Z)(\ell)| \geq \frac{\delta}{4\sqrt{L}\log(eL)\,\mu}\right\}\right| \geq \frac{9L\delta^2}{128\log^2(L)\,\mu^2}.$$

Summing up we obtain that with probability at least $1-\exp\left(-\frac{L\delta^4}{C_2\log^4(eL)\mu^4}\right)$

$$\begin{aligned}\inf_{Z\in\mathcal{E}_{\mu,\delta}}\|\mathcal{A}(Z)\| &\geq \inf_{Z\in\mathcal{E}_{\mu,\delta}}\frac{\delta}{4\sqrt{L}\log(eL)\mu}\sqrt{\left|\left\{\ell\in[L]:\ |\mathcal{A}(Z)(\ell)|\geq\frac{\delta}{4\sqrt{L}\log(eL)\mu}\right\}\right|}\\ &\gtrsim \frac{\delta^2}{\log^2(L)\mu^2}.\end{aligned}$$

This shows the claim. □

5.3. Proof of Theorem 3.7

As already mentioned in Section 4, in order to control all matrices $Z\in\mathcal{K}_*(h_0m_0^*)$, which are almost orthogonal to $h_0m_0^*$, we need the following key lemma.

Lemma 5.7. *Let $h_0m_0^*\in\mathbb{C}^{n_1\times n_2}$ be a rank-1 matrix. Assume that $Z\in\mathcal{K}_*(h_0m_0^*)\setminus\{0\}$. Then, whenever $\|h_0m_0^*+Z\|_*\leq\|h_0m_0^*\|_*$, it holds that*

$$\|Z\|_F\leq -2Re\left(\langle h_0m_0^*,\frac{1}{\|Z\|_F}Z\rangle_F\right).$$

Proof. We observe that

$$\begin{aligned}\|h_0m_0^*\|_F^2&=\|h_0m_0^*\|_*^2\\ &\geq\|h_0m_0^*+Z\|_*^2\\ &\geq\|h_0m_0^*+Z\|_F^2\\ &=\|h_0m_0^*\|_F^2+\|Z\|_F^2+2\mathrm{Re}\left(\langle h_0m_0^*,Z\rangle_F\right).\end{aligned}$$

Rearranging terms yields the result. □

Now we have gathered all tools which are needed to prove Theorem 3.7.

Proof of Theorem 3.7. Having introduced all necessary tools in the last two sections we can now give a proof of Theorem 3.7. We set $\delta:=(\log eL)^{2/3}\mu^{2/3}\alpha^{1/3}$. Throughout the proof we will assume that inequality (39) holds, which by Lemma 5.6 holds with probability at least

$$1-\exp\left(-\frac{L\delta^4}{C_2\log^4(eL)\mu^4}\right)=1-\exp\left(-\frac{L\alpha^{4/3}}{C_2\log^{4/3}(eL)\mu^{4/3}}\right).$$

Let $h_0\in\mathcal{H}_\mu$ and $m_0\in\mathbb{C}^N$. Furthermore, let $\hat{X}$ be a minimizer of (3) and set $Z:=\hat{X}-h_0m_0^*$. Note that from the minimality of $\hat{X}$ it follows that $\|\hat{X}\|_*\leq\|h_0m_0^*\|_*$. This implies that $Z\in\mathcal{K}_*(h_0m_0^*)$. To prove the lemma it remains to derive an appropriate

upper bound on $\|Z\|_F$. For that we will distinguish two cases, namely $\frac{Z}{\|Z\|_F} \in \mathcal{E}_{\mu,\delta}$ and $\frac{Z}{\|Z\|_F} \notin \mathcal{E}_{\mu,\delta}$. If $\frac{Z}{\|Z\|_F} \in \mathcal{E}_{\mu,\delta}$, it follows from inequality (39) that

$$\begin{aligned}\|Z\|_F &\lesssim \frac{\log^2(L)\,\mu^2}{\delta^2}\left\|\mathcal{A}(Z)\right\| \\ &\leq \frac{\log^2(L)\,\mu^2}{\delta^2}\left(\|\mathcal{A}\left(\hat{X}\right)-y\|+\|e\|\right) \\ &\leq 2\frac{\log^2(L)\,\mu^2}{\delta^2}\tau \\ &= 2\frac{\log^{2/3}(L)\,\mu^{2/3}}{\alpha^{2/3}}\tau,\end{aligned} \tag{43}$$

where in the second inequality we used the triangle inequality as well as $Z = \hat{X} - h_0 m_0^*$ and $y = \mathcal{A}(h_0 m_0^*) + e$. In the third inequality we used that $\hat{X}$ is feasible and $\|e\| \leq \tau$. If $\frac{Z}{\|Z\|_F} \notin \mathcal{E}_{\mu,\delta}$, it follows directly from the definition of $\mathcal{E}_{\mu,\delta}$ that $-\mathrm{Re}\left(\langle \frac{h_0 m_0^*}{\|h_0 m_0^*\|_F}, \frac{Z}{\|Z\|_F}\rangle_F\right) < \delta$. By Lemma 5.7 we obtain that

$$\begin{aligned}\|Z\|_F &\leq -2\mathrm{Re}\left(\langle h_0 m_0^*, \frac{1}{\|Z\|_F} Z\rangle_F\right) \\ &< 2\delta\|h_0 m_0^*\|_F \\ &< 2\left(\log L\right)^{2/3}\mu^{2/3}\alpha^{1/3}\|h_0 m_0^*\|_F.\end{aligned} \tag{44}$$

Combining the estimates (43) and (44) we obtain that

$$\|\hat{X} - h_0 m_0^*\|_F = \|Z\|_F \lesssim \frac{\mu^{2/3}\log^{2/3} L}{\alpha^{2/3}} \max\left\{\tau; \alpha\|h_0 m_0^*\|_F\right\}.$$

which completes the proof. □

6. Outlook

In this paper we have analyzed two important cases of structured low-rank matrix recovery problems, blind deconvolution and matrix completion, through an inspection of the descent cone of the nuclear norm and its interaction with the measurement operator $\mathcal{A}$. We have shown that the conic singular value is typically quite small and, consequently, previous analysis approaches cannot give strong recovery guarantees. For the example of blind deconvolution we have presented a new approach based on a refined analysis of the descent cone, showing that the nuclear norm minimization approach is stable against adversarial noise in certain important parameter regimes and allows for uniform recovery guarantees in the presence of noise. In our opinion our results give rise to a number of interesting follow-up questions.

- **Stability for small noise-levels:** Until now, our stability result only covers the situation that the noise level τ is of constant order (up to logarithmic factors). For small τ, Theorem 3.1, respectively Theorem 3.5, put some barriers on

what performance can be expected. Nevertheless, it will be interesting to examine the transitional case, that τ is rather small, even further. For example, while the bad conditioning for small noise levels has been established, it remains open whether one can construct a noise vector e such that the true minimizer behaves like the alternative (but non-minimal) solutions constructed in Theorems 3.1 and 3.5. Also the transitional behavior of the minimum conic singular values between very small noise levels (where we established bad conditioning) and larger noise levels (where at least for randomized blind deconvolution, we proved stability) will be an interesting question to study.

- **Extension to the rank r case:** Understanding nuclear norm recovery for matrix completion under adversarial noise remains an important open problem in the field. While our result established that recovery guarantees for arbitrary noise levels are not feasible, our considerations for the rank one scenario give hope that for sufficiently large noise levels, near optimal guarantees are within reach also for matrices of arbitrary rank.

  Similarly, a natural generalization of blind deconvolution is the problem of blind demixing [44, 28], where one observes a noisy superposition of several convolutions, that is, $y = \sum_{i=1}^{r} w_i * x_i + e$. The corresponding low-rank matrix formulation can be interpreted as a rank r version of the randomized blind deconvolution problem.

  We expect that a rank r version of Theorem 3.7 will apply to both these scenarios, which is why we consider this a very promising direction for future research.

- **Extension to other low-rank matrix recovery models**: Various other low-rank matrix models also involve incoherence in some way, for example, robust PCA ([6]) and spectral compressed sensing via matrix completion [13]. Also for these problems, recovery results are typically proven via the Golfing Scheme and lead to a seemingly suboptimal noise bound (see, e.g., [62, Section VI]). Can these problems be analyzed with the methods developed in this paper?
  Moreover, [36] provided an incoherence based analysis of the phase retrieval problem under random Bernoulli measurements. It will be interesting to analyze this setup with similar methods as in this manuscript.

Acknowledgements

The authors would like to thank Richard Kueng for discussions related to this topic, which sparked their interest in the problem. Furthermore, the authors want to thank Peter Jung, Marius Junge, and Kiryung Lee for fruitful discussions. This work has been supported by the German Science Foundation (DFG) in the context of the project *Bilinear Compressed Sensing* (KR 4512/2-1) as a part of the Priority Program 1798, and the Emmy-Noether Junior Research Group *Randomized Sensing and Quantization of Signals and Images* (KR 4512/1-1).

References

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

[2] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: phase transitions in convex programs with random data. *Inf. Inference*, 3(3):224–294, 2014.

[3] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

[4] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

[5] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

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

[7] E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

[8] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717, 2009.

[9] E. J. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274, 2013.

[10] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2010.

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

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

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

[14] Y. Chen, Y. Chi, J. Fan, C. Ma, and Y. Yan. Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. *arXiv preprint arXiv:1902.07698*, 2019.

[15] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k-term approximation. *J. Am. Math. Soc.*, 22(1):211–231, 2009.

[16] V. de la Peña and E. Giné. *Decoupling. From dependence to independence. Randomly stopped processes, U-statistics and processes, martingales and beyond.* New York, NY: Springer, 1998.

[17] L. Ding and Y. Chen. The leave-one-out approach for matrix completion: Primal and dual analysis. *arXiv preprint arXiv:1803.07554*, 2018.

[18] S. Dirksen, G. Lecué, and H. Rauhut. On the gap between restricted isometry properties and sparse recovery conditions. *IEEE Trans. Inform. Theory*, 64(8):5478–5487, 2018.

[19] M. Fornasier, H. Rauhut, and R. Ward. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM J. Optim.*, 21(4):1614–1640, 2011.

[20] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Birkhäuser Basel, 2013.

[21] R. Ge, J. D. Lee, and T. Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.

[22] L. Grafakos. *Classical Fourier analysis. 2nd ed.* New York, NY: Springer, 2nd ed. edition, 2008.

[23] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inform. Theory*, 57(3):1548–1566, 2011.

[24] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Physical review letters*, 105(15):150401, 2010.

[25] M. Hardt. Understanding alternating minimization for matrix completion. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 651–660. IEEE, 2014.

[26] W. Huang and P. Hand. Blind deconvolution by a steepest descent algorithm on a quotient manifold. *SIAM J. Imaging Sci.*, 11(4):2757–2785, 2018.

[27] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 665–674, New York, NY, USA, 2013. ACM.

[28] P. Jung, F. Krahmer, and D. Stöger. Blind demixing and deconvolution at near-optimal rate. *IEEE Trans. Inform. Theory*, 64(2):704–727, 2018.

[29] M. Kabanava, R. Kueng, H. Rauhut, and U. Terstiege. Stable low-rank matrix recovery via null space properties. *Inf. Inference*, 5(4):405–441, 2016.

[30] M. Kech and F. Krahmer. Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems. *SIAM J. Appl. Algebra Geom.*, 1(1):20–37, 2017.

[31] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Trans. Inf. Theory*, 56(6):2980–2998, 2010.

[32] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *J. Mach. Learn. Res.*, 11:2057–2078, 2010.

[33] O. Klopp. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.

[34] V. Koltchinskii, K. Lounici, A. B. Tsybakov, et al. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, 39(5):2302–2329, 2011.

[35] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN*, 2015(23):12991–13008, 2015.

[36] F. Krahmer and Y.-K. Liu. Phase retrieval without small-ball probability assumptions. *IEEE Trans. Inform. Theory*, 64(1):485–500, 2018.

[37] F. Krahmer and D. Stöger. Blind deconvolution: Convex geometry and noise robustness. In *52nd Annual Asilomar Conference on Signals, Systems, and Computers*, 2018.

[38] R. Kueng, H. Rauhut, and U. Terstiege. Low rank matrix recovery from rank one measurements. *Appl. Comput. Harmon. Anal.*, 42(1):88–116, 2017.

[39] R. Kueng, H. Zhu, and D. Gross. Low rank matrix recovery from clifford orbits. *arXiv preprint arXiv:1610.08070*, 2016.

[40] C. Kümmerle and J. Sigl. Harmonic mean iteratively reweighted least squares for low-rank matrix recovery. *J. Mach. Learn. Res.*, 19:49, 2018.

[41] G. Lecué and S. Mendelson. Sparse recovery under weak moment assumptions. *J. Eur. Math. Soc. (JEMS)*, 19(3):881–904, 2017.

[42] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.

[43] X. Li, S. Ling, T. Strohmer, and K. Wei. Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *Appl. Comput. Harmon. Anal.*, 2018.

[44] S. Ling and T. Strohmer. Blind deconvolution meets blind demixing: Algorithms and performance bounds. *IEEE Trans. Inform. Theory*, 63(7):4497–4520, 2017.

[45] Y.-K. Liu. Universal low-rank matrix recovery from pauli measurements. In *Advances in Neural Information Processing Systems*, pages 1638–1646, 2011.

[46] C. Ma, K. Wang, Y. Chi, and Y. Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467*, 2017.

[47] M. B. McCoy and J. A. Tropp. Sharp recovery bounds for convex demixing, with applications. *Found. Comput. Math.*, 14(3):503–567, 2014.

[48] S. Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39, 2014.

[49] S. Negahban and M. J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *J. Mach. Learn. Res.*, 13(May):1665–1697, 2012.

[50] Y. Plan and R. Vershynin. The generalized lasso with non-linear observations. *IEEE Trans. Inf. Theory*, 62(3):1528–1537, 2016.

[51] B. Recht. A simpler approach to matrix completion. *J. Mach. Learn. Res.*, 12(Dec):3413–3430, 2011.

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

[53] J. D. M. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22Nd International Conference on Machine Learning*, ICML '05, pages 713–719, New York, NY, USA, 2005. ACM.

[54] R. Rockafellar. *Convex analysis.* 1970.

[55] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Commun. Pure Appl. Math.*, 61(8):1025–1045, 2008.

[56] R. Sun and Z.-Q. Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Trans. Inform. Theory*, 62(11):6535–6579, 2016.

[57] J. A. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling theory, a renaissance. Compressive sensing and other developments*, pages 67–101. Cham: Birkhäuser/Springer, 2015.

[58] J. A. Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.

[59] A. W. Van Der Vaart and J. A. Wellner. Weak convergence. In *Weak Convergence and Empirical Processes*, pages 16–28. Springer, 1996.

[60] R. Vershynin. High-dimensional probability.

[61] G. Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra Appl.*, 170:33–45, 1992.

[62] Z. Zhou, X. Li, J. Wright, E. J. Candès, and Y. Ma. Stable principal component pursuit. In *2010 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 1518–1522. IEEE, 2010.

A. Proof of Lemma 5.2

The proof of Lemma 5.2 will rely on the following two lemmas. The first lemma is a version of Mendelson's small-ball method for non-i.i.d. measurements. In order to state it let $X_1, \dots, X_L$ be independent, matrix-valued random variables defined on a probability space (Ω, μ). For every measurable, real-valued function f and for every $\xi > 0$ we define the quantity

$$Q_\xi(f) = \sum_{\ell=1}^{L} \mathbb{P}\left(f(X_\ell) \geq \xi\right).$$

Lemma A.1. *Let $X_1, \dots, X_L \in \mathbb{C}^{K \times N}$ be independent random variables and $\mathcal{F}$ be a set of real-valued functions, which are measurable with respect to (Ω, μ). Let $t > 0$ and $\xi > 0$. Then with probability at least $1 - \exp\left(-2t^2\right)$ it holds that*

$$\inf_{f \in \mathcal{F}} \left| \{\ell \in [L] : \ f(X_\ell) \geq \xi\} \right| \geq \inf_{f \in \mathcal{F}} Q_{2\xi}(f) - \frac{2}{\xi} \mathbb{E}\left[\sup_{f \in \mathcal{F}} \sum_{\ell=1}^{L} \varepsilon_\ell f(X_\ell)\right] - t\sqrt{L},$$

where $\varepsilon_1, \dots, \varepsilon_L$ are independent Rademacher variables, i.e., random variables which take the two values ± 1 each with probability $\frac{1}{2}$.

The proof of Lemma A.1 is exactly analogous as the proof of the original small-ball method [35], see Section A.1. The second auxiliary lemma, proved in Section A.2, relates the quantity $\mathbb{E}\left[\sup_{f \in \mathcal{F}} \sum_{\ell=1}^{L} \varepsilon_\ell f(X_\ell)\right]$ in the blind deconvolution framework to the Gaussian width (cf. Definition 5.1).

Lemma A.2. *Let $\mathcal{E} \subset \mathbb{C}^{K \times N}$. Then it holds that*

$$\mathbb{E}\left[\sup_{X \in \mathcal{E}} Re\left(\sum_{\ell=1}^{L} b_\ell^* X c_\ell\right)\right] = \omega(\mathcal{E}).$$

With these lemmas we can now prove Lemma 5.2.

Proof of Lemma 5.2. Set $X_\ell := b_\ell c_\ell^*$ for all $\ell \in [L]$ and define

$$\mathcal{F} := \{|\langle M, \cdot \rangle_F| : M \in \mathcal{E}\}.$$

Then, by a direct application of Lemma A.1 we obtain that with probability at least $1-\exp\left(-2t^2\right)$ it holds that

$$\begin{aligned}&\inf_{M\in\mathcal{E}}\left|\left\{\ell\in[L]:\ |\langle M,b_\ell c_\ell^*\rangle_F|\geq\xi\right\}\right|\\ \geq&\inf_{M\in\mathcal{E}}\left(\sum_{\ell=1}^{L}\mathbb{P}\left(|\langle b_\ell c_\ell^*,M\rangle_F|\geq 2\xi\right)\right)-\frac{2}{\xi}\mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell|\langle b_\ell c_\ell^*,M\rangle_F|\right]-t\sqrt{L}.\end{aligned}\tag{45}$$

To bound the second summand, we observe that

$$\begin{aligned}&\mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell|\langle b_\ell c_\ell^*,M\rangle_F|\right]\\ \leq&\mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell|\mathrm{Re}\left(\langle b_\ell c_\ell^*,M\rangle_F\right)|\right]+\mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell|\mathrm{Im}\left(\langle b_\ell c_\ell^*,M\rangle_F\right)|\right]\\ =&2\ \mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell|\mathrm{Re}\left(\langle b_\ell c_\ell^*,M\rangle_F\right)|\right]\\ =&2\ \mathbb{E}\left[\sup_{M\in\mathcal{E}}\sum_{\ell=1}^{L}\varepsilon_\ell\mathrm{Re}\left(\langle b_\ell c_\ell^*,M\rangle_F\right)\right]\\ =&2\omega\left(\mathcal{E}\right)\end{aligned}\tag{46}$$

where in the third line we used that $\mathrm{Re}\left(\langle b_\ell c_\ell^*,X\rangle_F\right)$ and $\mathrm{Im}\left(\langle b_\ell c_\ell^*,X\rangle_F\right)$ have the same distribution. The fourth line follows from the symmetry of the set $\mathcal{E}$ and the last line is due to Lemma A.2. Combining (45) and (46) finishes the proof. □

A.1. Proof of Lemma A.1

We directly trace the steps of the proof of Theorem 1.5 in [35]. In the following $\mathbb{1}_A$ denotes the indicator function, which takes the value 1, if the event A occurs and the value 0 otherwise. Note that

$$\xi\left|\left\{\ell\in[L]:\ f\left(X_\ell\right)\geq\xi\right\}\right|=\xi\sum_{\ell=1}^{L}\mathbb{1}_{\{f(X_\ell)\geq\xi\}}.$$

Taking the infimum we observe that by the definition of $Q_{2\xi}$

$$\begin{aligned}&\xi\inf_{f\in\mathcal{F}}\left|\left\{\ell\in[L]:\ f\left(X_\ell\right)\geq\xi\right\}\right|\\ \geq&\xi\inf_{f\in\mathcal{F}}Q_{2\xi}\left(f\right)-\xi\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{P}\left(f\left(X_\ell\right)\geq 2\xi\right)-\mathbb{1}_{\{f(X_\ell)\geq\xi\}}\right).\end{aligned}\tag{47}$$

The bounded difference inequality (see, for example, [4]) implies that with probability at least $1-\exp\left(-2t^2\right)$ it holds that

$$\begin{aligned}
&\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{P}\left(f\left(X_\ell\right)\geq 2\xi\right)-\mathbb{1}_{\{f(X_\ell)\geq\xi\}}\right)\\
\leq&\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{P}\left(f\left(X_\ell\right)\geq 2\xi\right)-\mathbb{1}_{\{f(X_\ell)\geq\xi\}}\right)\right]+t\sqrt{L} \qquad (48)\\
=&\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{E}\left[\mathbb{1}_{\{f(X_\ell)\geq 2\xi\}}\right]-\mathbb{1}_{\{f(X_\ell)\geq\xi\}}\right)\right]+t\sqrt{L}
\end{aligned}$$

To deal with the expectation we will use the function $\Psi_\xi:[0,+\infty)\longrightarrow\mathbb{R}$ defined by

$$\Psi_\xi\left(u\right)=\begin{cases}0 & 0\leq u\leq\xi\\ \frac{1}{\xi}\left(u-\xi\right) & \xi\leq u\leq 2\xi\\ 1 & u\geq 2\xi\end{cases}.$$

We observe that Ψ_ξ is Lipschitz continuous with Lipschitz constant $1/\xi$. Furthermore, for all $u\in[0,+\infty)$ it holds that $\mathbb{1}_{\{u\geq 2\xi\}}\leq\Psi_\xi\left(u\right)\leq\mathbb{1}_{\{u\geq\xi\}}$. Combining this monotonicity relation with Gine-Zinn symmetrization (see, e.g., [59, Lemma 2.3.1]) and the Rademacher comparison principle for Lipschitz continuous functions (see, e.g., [42, Equation (4.20)]), we obtain that

$$\begin{aligned}
&\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{E}\left[\mathbb{1}_{\{f(X_\ell)\geq 2\xi\}}\right]-\mathbb{1}_{\{f(X_\ell)\geq\xi\}}\right)\right]\\
\leq&\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\left(\mathbb{E}\left[\Psi_\xi\left(f\left(X_\ell\right)\right)\right]-\Psi_\xi\left(f\left(X_\ell\right)\right)\right)\right]\\
\leq&2\,\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\varepsilon_\ell\Psi_\xi\left(f\left(X_\ell\right)\right)\right]\\
\leq&\frac{2}{\xi}\,\mathbb{E}\left[\sup_{f\in\mathcal{F}}\sum_{\ell=1}^{L}\varepsilon_\ell f\left(X_\ell\right)\right].
\end{aligned}$$

Together with the inequality chains (47) and (48), this completes the proof. □

A.2. Proof of Lemma A.2

First, we observe that

$$\mathbb{E}\left[\sup_{X\in\mathcal{E}}\operatorname{Re}\left(\sum_{\ell=1}^{L}b_\ell^*Xc_\ell\right)\right]=\mathbb{E}\left[\sup_{X\in\mathcal{E}}\operatorname{Re}\left(\langle X,\sum_{\ell=1}^{L}b_\ell c_\ell^*\rangle_F\right)\right]$$

Note that due to the definition of $\omega(\mathcal{E})$ in order to finish the proof it is enough to show that the entries of the matrix $X = \sum_{\ell=1}^{L} b_\ell c_\ell^*$ are independent and identically distributed with distribution $\mathcal{CN}(0,1)$. For that, let $(i,j) \in [K] \times [N]$ and compute that

$$\mathbb{E}\left[|e_i^*\left(\sum_{\ell=1}^{L} b_\ell c_\ell^*\right)e_j|^2\right] = \sum_{\ell=1}^{L} e_i^* b_\ell b_\ell^* e_i \mathbb{E}\left[|c_\ell^* e_j|^2\right]^2 = \sum_{\ell=1}^{L} e_i^* b_\ell b_\ell^* e_i = 1.$$

This implies that $e_i^*\left(\sum_{\ell=1}^{L} b_\ell c_\ell^*\right)e_j \in \mathcal{CN}(0,1)$. It remains to show that the individual entries of the matrix $\sum_{\ell=1}^{m} b_\ell c_\ell^*$ are independent. For that, we set

$$X_{i,j} := \left(\sum_{\ell=1}^{m} b_\ell c_\ell^*\right)_{i,j} = e_i^*\left(\sum_{\ell=1}^{L} b_\ell c_\ell^*\right)e_j.$$

Now let $(i,j),(i',j') \in [K] \times [N]$ such that $(i,j) \neq (i',j')$. Our goal is to show that $\mathbb{E}\left[X_{i,j}\overline{X}_{i',j'}\right] = 0$. If $j \neq j'$ this follows immediately from the observation that $c_\ell^* e_j$ and $c_\ell^* e_{j'}$ are independent for all $\ell \in [L]$. Now assume that $j = j'$. Then we can compute that

$$\mathbb{E}\left[X_{i,j}\overline{X}_{i',j'}\right] = \sum_{\ell=1}^{L} e_i^* b_\ell b_\ell^* e_{i'} |c_\ell^* e_j|^2 = \sum_{\ell=1}^{L} e_i^* b_\ell b_\ell^* e_{i'} = 0.$$

Hence, we have shown that all entries of the matrix X are uncorrelated. As the entries of X are jointly Gaussian this implies that they are independent, which completes the proof.

□

B. Proof of Lemma 5.3

Proof of Lemma 5.3. Let $\ell \in [L]$ such that $\|X^* b_\ell\| \geq 4\xi$. Using the Paley-Zygmund inequality (see, e.g., [16]) we obtain that

$$\begin{aligned}
\mathbb{P}\left(|\langle b_\ell c_\ell^*, X\rangle_F| \geq 2\xi\right) &\geq \mathbb{P}\left(|\langle b_\ell c_\ell^*, X\rangle_F| \geq \frac{1}{2}\|X^* b_\ell\|\right) \\
&\geq \frac{\left(\mathbb{E}\left[|\langle b_\ell c_\ell^*, X\rangle_F|^2\right] - \frac{1}{4}\|X^* b_\ell\|^2\right)^2}{\mathbb{E}\left[|\langle b_\ell c_\ell^*, X\rangle_F|^4\right]} \\
&= \frac{\left(\|X^* b_\ell\|^2 - \frac{1}{4}\|X^* b_\ell\|^2\right)^2}{2\|X^* b_\ell\|^4} = \frac{9}{32}.
\end{aligned}$$

(We used that $\mathbb{E}|\langle b_\ell c_\ell^*, X\rangle_F|^2 = \|X^* b_\ell\|^2$ and $\mathbb{E}|\langle b_\ell c_\ell^*, X\rangle_F|^4 = 2\|X^* b_\ell\|^4$, which is due to $\langle b_\ell c_\ell^*, X\rangle_F \sim \mathcal{CN}(0, \|X^* b_\ell\|)$.) Summing over all $\ell \in [L]$ such that $\|X^* b_\ell\| \geq 4\xi$ yields the claim. □