

# Provable Low Rank Phase Retrieval

Seyedehsara Nayer, Praneeth Narayananmurthy, Namrata Vaswani  
Iowa State University, Ames, IA

**Abstract**—We study the Low Rank Phase Retrieval (LRPR) problem defined as follows: recover an  $n \times q$  matrix  $\mathbf{X}^*$  of rank  $r$  from a different and independent set of  $m$  phaseless (magnitude-only) linear projections of each of its columns. To be precise, we need to recover  $\mathbf{X}^*$  from  $y_k := |\mathbf{A}_k' \mathbf{x}_k^*|, k = 1, 2, \dots, q$  when the measurement matrices  $\mathbf{A}_k$  are mutually independent. Here  $y_k$  is an  $m$  length vector,  $\mathbf{A}_k$  is an  $n \times m$  matrix, and  $'$  denotes matrix transpose. The question is when can we solve LRPR with  $m \ll n$ ? A reliable solution can enable fast and low-cost phaseless dynamic imaging, e.g., Fourier ptychographic imaging of live biological specimens. In this work, we develop the first provably correct approach for solving this LRPR problem. Our proposed algorithm, Alternating Minimization for Low-Rank Phase Retrieval (AltMinLowRaP), is an AltMin based solution and hence is also provably fast (converges geometrically). Our guarantee shows that AltMinLowRaP solves LRPR to  $\epsilon$  accuracy, with high probability, as long as  $mq \geq Cnr^4 \log(1/\epsilon)$ , the matrices  $\mathbf{A}_k$  contain i.i.d. standard Gaussian entries, and the right singular vectors of  $\mathbf{X}^*$  satisfy the incoherence assumption from matrix completion literature. Here  $C$  is a numerical constant that only depends on the condition number of  $\mathbf{X}^*$  and on its incoherence parameter. Its time complexity is only  $Cmqnr \log^2(1/\epsilon)$ .

Since even the linear (with phase) version of the above problem is not fully solved, the above result is also the first complete solution and guarantee for the linear case. Finally, we also develop a simple extension of our results for the dynamic LRPR setting.

## I. INTRODUCTION

In recent years, there has been a resurgence of interest in the classical phase retrieval (PR) problem [2], [3]. The original PR problem involved recovering an  $n$ -length signal  $\mathbf{x}^*$  from the *magnitudes* of its Discrete Fourier Transform (DFT) coefficients. Its generalized version, studied in recent literature, replaces DFT by inner products with any arbitrary design vectors,  $\mathbf{a}_i$ . Thus, the goal is to recover  $\mathbf{x}^*$  from  $\mathbf{y}_i := |\langle \mathbf{a}_i, \mathbf{x}^* \rangle|, i = 1, 2, \dots, m$ . These are commonly referred to as phaseless linear projections of  $\mathbf{x}^*$ . While practical PR methods have existed for a long time, e.g., see [2], [3], the focus of the recent work has been on obtaining correctness guarantees for these and newer algorithms. This line of work includes convex relaxation methods [4], [5] as well as non-convex methods [6]–[13]. It is easy to see that, without extra assumptions, PR requires  $m \geq n$ . The best known guarantees – see [9] and follow-up works – prove exact recovery with high probability (whp) with order-optimal number of measurements/samples:  $m = Cn$ ; and with time complexity  $Cmn \log(1/\epsilon)$  that is nearly linear in the problem size. Here and below,  $C$  is reused often to refer to a constant more than one. Most guarantees for PR assume that  $\mathbf{a}_i$ ’s are independent and identically distributed (iid) standard Gaussian vectors. When this is assumed, we refer to the PR problem as “standard PR”.

Part of this work appears in the proceedings of ICML 2019 [1].

A natural approach to reduce the sample complexity to below  $n$  is to impose structure on the unknown signal(s). In existing literature, with the exception of sparse PR which has been extensively studied, e.g., [6], [14]–[18], there is little other work on structured PR. Low rank is the other common structure. This can be used in one of two ways. One is to assume that the unknown signal/image, whose phaseless linear projections are available, can be rearranged to form a low-rank matrix. This would be valid only for very specific types of images for which different image rows or columns look similar, so that the entire image matrix can be modeled as low rank. *In general it is not a very practical model for images, and this is probably why this setting has not been explored in the literature. We do not consider this model here either.*

A more practical, and commonly used, low-rank model in biological applications [19], is for the dynamic imaging setting. It assumes that a set, e.g., a time sequence, of signals/images is generated from a lower dimensional subspace of the ambient space. For our problem, we assume that we have a set of  $m$  phaseless linear projections of each signal, with a different set of measurement vectors used for each signal. The question is can we jointly recover the signals using an  $m \ll n$  and when? This setting was first studied in our recent work [20] where we called it “Low-Rank PR” (LRPR). It is a valid model whenever the set/sequence of signals is sufficiently similar (correlated). A solution to LRPR can enable fast and low-cost phaseless dynamic imaging of live biological specimens, *in vitro*. See Sec. I-C and [21] for a detailed motivation for studying LRPR.

### A. Low Rank PR (LRPR) Problem Setting and Notation

We study the LRPR problem described above. This was first introduced and briefly studied in [20] where we developed two algorithms, evaluated them experimentally, and provided a guarantee for the initialization step of one of them. The goal is to recover an  $n \times q$  matrix  $\mathbf{X}^* := [\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_k^*, \dots, \mathbf{x}_q^*]$ , of rank  $r$ , from measurements

$$\mathbf{y}_{ik} := |\langle \mathbf{a}_{ik} \mathbf{e}_k', \mathbf{X}^* \rangle| = |\langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle|, \quad i \in [m], k \in [q], \quad (1)$$

when all the  $\mathbf{a}_{ik}$ ’s are mutually independent. For proving guarantees we assume also that they are iid standard Gaussian and real-valued. Here and below the notation  $[m] := \{1, 2, \dots, m\}$ , and  $\mathbf{e}_k$  refers to the  $k$ -th column of  $\mathbf{I}_q$  (identity matrix of size  $q \times q$ ). We are interested in the low rank setting when  $r \ll \min(n, q)$ .

By defining the  $m$ -length vector  $\mathbf{y}_k := [\mathbf{y}_{1,k}, \mathbf{y}_{2,k}, \dots, \mathbf{y}_{m,k}]'$  and the  $n \times m$  matrix  $\mathbf{A}_k := [\mathbf{a}_{1,k}, \mathbf{a}_{2,k}, \dots, \mathbf{a}_{m,k}]$ , and letting  $|\mathbf{z}|$  denote

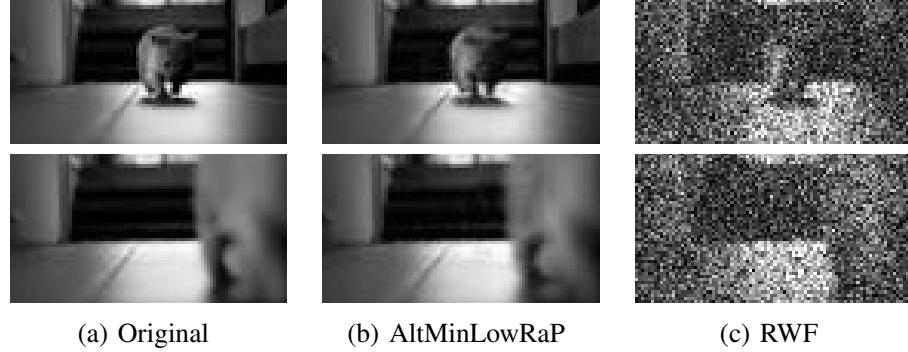


Fig. 1: Recovering a video of a moving mouse (only approximately low-rank) from simulated coded diffraction pattern (CDP) measurements. The images are shown at  $k = 60, 78$ . We describe the experiment in Sec. IV-B.

element-wise magnitude of a vector, the above measurement model can also be rewritten as

$$\mathbf{y}_k := |\mathbf{A}_k' \mathbf{x}_k^*|, \quad k = 1, 2, \dots, q \quad (2)$$

where ' denotes vector or matrix transpose.

The requirement that the measurement vectors used for imaging different  $\mathbf{x}_k^*$ 's be different and independent is what allows us to show that  $m \ll n$  suffices. To understand this point in a simple fashion, consider the  $r = 1$  setting and suppose that  $\mathbf{x}_k^* = \mathbf{x}_1^*$  (all columns are equal). We would then have  $mq$  iid Gaussian measurements of  $\mathbf{x}_1^*$  and hence  $mq \geq Cn$  would suffice [9]. If  $q = n$ , this means just a constant number of measurements  $m$  per column (signal) suffices. For  $r > 1$  but small, we will show that we can extend this idea to show that, when  $q = n$  (or is larger), the required value of  $m$  depends only on the value of  $r$  and not on  $n$ . On the other hand, if  $\mathbf{a}_{ik} = \mathbf{a}_i$  for all  $k$ , then, in the above  $r = 1$  example, only the first  $m$  measurements are useful (the others are just repeats of these). Thus we will still need  $m \geq Cn$  in this case. This  $\mathbf{a}_{ik} = \mathbf{a}_i$  case, and its linear version, is what has been studied extensively in the literature [22], [23]. In this case,  $m$  needs to be at least  $(n + q)r$ .

Let  $\mathbf{X}^* \stackrel{\text{SVD}}{=} \mathbf{U}^* \Sigma^* \mathbf{B}^*$  denote its singular value decomposition (SVD) so that  $\mathbf{U}^* \in \mathbb{R}^{n \times r}$ ,  $\mathbf{B}^* \in \mathbb{R}^{r \times q}$ , and  $\Sigma^* \in \mathbb{R}^{r \times r}$  is a diagonal matrix. Observe that this notation is a little non-standard, if the SVD was  $\mathbf{U}^* \Sigma^* \mathbf{V}^*$ , we are letting  $\mathbf{B}^* := \mathbf{V}^*$ . Thus, columns of  $\mathbf{U}^*$  and rows of  $\mathbf{B}^*$  are orthonormal. We use  $\sigma_{\max}^*$ ,  $\sigma_{\min}^*$  to denote the maximum, minimum singular values of  $\mathbf{X}^*$  and  $\kappa = \sigma_{\max}^* / \sigma_{\min}^*$  to denote its condition number. Finally, we define

$$\tilde{\mathbf{B}}^* := \Sigma^* \mathbf{B}^*.$$

We use the above non-standard notation for SVD because our solution approach will recover columns of  $\tilde{\mathbf{B}}^*$ ,  $\tilde{\mathbf{b}}_k^*$ , individually by solving an  $r$ -dimensional standard PR problem (it is more intuitive to talk about recovery of column vectors than of rows). With the above notation, the QR decomposition of an estimate of  $\tilde{\mathbf{B}}^*$ , denoted  $\hat{\mathbf{B}}$ , will be written as  $\hat{\mathbf{B}} \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}$  with  $\mathbf{B}$  being an  $r \times q$  matrix with orthonormal rows (or equivalently  $\hat{\mathbf{B}}' \stackrel{\text{QR}}{=} \mathbf{B}'(\mathbf{R}_B)'$ ).

*Right Incoherence.* Observe that we have global measurements of each column, but not of the entire matrix. Thus, in order to correctly recover  $\mathbf{X}^*$  with small  $m$ , we need an assumption that allows for correct ‘‘interpolation’’ across the rows. One way to ensure this is to borrow the right incoherence (incoherence or denseness of right singular vectors) assumption from matrix completion literature [24], [25]. In our notation, this means that we need to assume that

$$\max_k \|\mathbf{b}_k^*\|^2 \leq \mu^2 \frac{r}{q}, \quad (3)$$

with  $\mu \geq 1$  being a constant. Clearly, this implies that

$$\|\mathbf{x}_k^*\|^2 = \|\tilde{\mathbf{b}}_k^*\|^2 \leq \sigma_{\max}^{*2} \mu^2 \frac{r}{q} = \kappa^2 \sigma_{\min}^{*2} \mu^2 \frac{r}{q} \leq \kappa^2 \mu^2 \frac{\|\mathbf{X}^*\|_F^2}{q} \quad (4)$$

for each  $k$ . If we assume  $\kappa$  is a constant, up to constant factors, (4) also implies (3). Thus, up to constant factors, requiring right incoherence is the same as requiring that the maximum energy of any signal  $\mathbf{x}_k^*$  is within constant factors of the average.

1) *Notation:* We use  $\|\cdot\|$  to denote the  $l_2$ -norm of a vector or the induced 2-norm matrix and  $\|\cdot\|_F$  to denote the Frobenius norm. We use  $\mathbb{1}_{\text{statement}}$  to denote the indicator function; it takes the value one if statement is true and is zero otherwise. A tall matrix with orthonormal columns is referred to as a ‘‘basis matrix’’. For two basis matrices  $\mathbf{U}_1, \mathbf{U}_2$ , we define the subspace error (distance) as

$$\sin \Theta(\mathbf{U}_1, \mathbf{U}_2) := \|(\mathbf{I} - \mathbf{U}_1 \mathbf{U}_1') \mathbf{U}_2\|.$$

This measures the sine of the largest principal angle between the two subspaces. We often use terms like ‘‘estimate  $\mathbf{U}$ ’’ when the goal is to really estimate its column span,  $\text{Span}(\mathbf{U})$ . The phase-invariant distance is defined as

$$\text{dist}(\mathbf{x}^*, \hat{\mathbf{x}}) := \min_{\theta \in [-\pi, \pi]} \|\mathbf{x}^* - e^{-j\theta} \hat{\mathbf{x}}\|$$

For our guarantees, we work with real valued vectors and matrices, and in this case this simplifies to  $\text{dist}(\mathbf{x}^*, \hat{\mathbf{x}}) = \min(\|\mathbf{x}^* - \hat{\mathbf{x}}\|, \|\mathbf{x}^* + \hat{\mathbf{x}}\|)$ . Define the corresponding distance between two matrices as

$$\text{mat-dist}(\hat{\mathbf{X}}, \mathbf{X}^*)^2 := \sum_{k=1}^q \text{dist}^2(\mathbf{x}_k^*, \hat{\mathbf{x}}_k).$$

TABLE I: Comparing our result (first result for LRPR) with the first and the best results for (bounded time) non-convex algorithms for the three related problems to ours - LRPC, PR, Sparse PR. Here “best” refers to best sample complexity. We treat  $\kappa, \mu$  as constants.  $\mathbf{X}^*$  is an  $n \times q$  rank- $r$  matrix;  $\mathbf{x}^*$  is an  $n$ -length vector signal; and  $m$  is the number of samples needed per signal (per column of  $\mathbf{X}^*$ ). LRPR, Sparse PR, and Standard PR results assume iid Gaussian measurements, while LRPC results assume iid Bernoulli model on observed entries.

Problem	Global Measurements?	Assumptions	Sample Complexity	Time Complexity per signal (with $m$ = its lower bound)
<b>LRPR (first) (our work)</b>	No	<b>right incoherence, <math>\mathbf{X}^*</math> has rank <math>r</math></b>	$m \geq C \frac{n}{q} r^4 \log(1/\epsilon)$ , $m \geq C \max(r, \log n, \log q)$	$Cn(n/q)r^5 \log^3(1/\epsilon)$
LRMC (first) [25]	No	left & right incoherence, $\mathbf{X}^*$ has rank $r$	$m \geq C \frac{\max(n, q)}{q} r^{4.5} \log(1/\epsilon)$	$C(n/q)r^{6.5} \log n \log^2(1/\epsilon)$
LRMC (best) [26]	No	left & right incoherence $\mathbf{X}^*$ has rank $r$	$m \geq C \frac{\max(n, q)}{q} r^2 \log^2 n \log^2(1/\epsilon)$	$C(n/q)r^3 \log^n \log^2(1/\epsilon)$
Sparse PR (first) [6]	Yes	$\mathbf{x}^*$ is $s$ -sparse in canonical basis, min nonzero entry lower bounded	$m \geq Cs^2 \log n \log(1/\epsilon)$	$Cns^3 \log n \log^2(1/\epsilon)$
Sparse PR (best) [17], [18]	Yes	$\mathbf{x}^*$ is $s$ -sparse in canonical basis	$m \geq Cs^2 \log n$	$Cns^2 \log n \log(1/\epsilon)$
Standard PR (first) [6]	Yes	None	$m \geq Cn \log^3 n \log(1/\epsilon)$	$Cn^2 \log^3 n \log(1/\epsilon)$
Standard PR (best) [9], [10]	Yes	None	$m \geq Cn$	$Cn^2 \log(1/\epsilon)$

We reuse the letters  $c, C$  to denote different numerical constants in each use, with the convention  $C \geq 1$  and  $c < 1$ .

### B. Our Contributions and their Significance and Novelty

This work provides the *first* provably correct solution, AltMinLowRaP (Alternating Minimization for Low Rank PR), for Low Rank PR. AltMinLowRaP is a fast alternating minimization (AltMin) based solution approach with a carefully designed spectral initialization. We can prove that AltMinLowRaP converges geometrically to an  $\epsilon$ -accurate solution as long as (i) right incoherence stated in (3) holds, and (ii) the total number of available measurements,  $mq$ , is at least a constant (that depends on  $\kappa, \mu$ ) times  $nr^4 \log(1/\epsilon)$ . Its time complexity is order  $mqnr^2 \log^2(1/\epsilon)$ , but if we replace  $mq$  by its lowest allowed value, then the time complexity becomes  $O(n^2r^5 \log^3(1/\epsilon))$ . If  $q = n$  (or is larger), ignoring log factors, this implies that only about  $r^4$  (or lesser) samples per signal suffice when using AltMinLowRaP. Moreover, when using these many samples, the time complexity per signal is only about  $Cnr^5$ . On the other hand, standard PR approaches (to recover each signal  $\mathbf{x}_k^*$  individually) necessarily need  $m \geq Cn$  samples, and order  $Cn^2$  time, per signal [9], [10]. In the regime of small  $r$ , e.g.,  $r = \log n$ , our result provides a significant sample, and time, complexity improvement over standard PR. Moreover, in this regime, our sample complexity is also only a little worse than its order optimal value of  $(n+q)r$ . We demonstrate the practical power of AltMinLowRaP in Fig. 1, Fig. 2, and the other experiments described later.

The key insight that helps obtain the above reduction in sample complexity is the following observation: for both the initialization and the update steps for  $\mathbf{U}^*$ , conditioned on  $\mathbf{X}^*$ , we have access to  $mq$  mutually independent measurements. These are not identically distributed (because the different  $\mathbf{b}_k^*$ ’s could have different distributions), however, we can carefully use the right incoherence assumption to show that the distributions are “similar enough” so that concentration inequalities can be applied jointly for all the  $mq$  samples.

We also briefly study the dynamic setting, Phaseless Subspace Tracking, which allows the underlying signal subspace to change with time in a piecewise constant fashion.

To our best knowledge, even the linear version of our problem – recover  $\mathbf{X}^*$  from  $\mathbf{y}_k := \mathbf{A}_k' \mathbf{x}_k^*, k = 1, 2, \dots, q$  – does not have any provably correct solutions (as we explain in Sec. I-D). Thus, our work also provides the first provable solution for this linear version. Our result implies that, even in this case, we can recover  $\mathbf{X}^*$  to  $\epsilon$  accuracy using only  $m = C(n/q)r^4 \log(1/\epsilon)$  samples per column, and can do this using an alternating minimization solution. What has been studied extensively is the  $\mathbf{A}_k = \mathbf{A}$  version of both LRPR and the above linear version [22], [23]. These are completely different problems as explained earlier in Sec. I-A.

1) *Significance*: The other somewhat related problem to ours, Sparse PR, is actually quite different. This is because it involves recovery from global measurements of the sparse vector (each measurement depends on the entire unknown sparse vector), where as, in our case, the measurements are not global for the entire matrix  $\mathbf{X}^*$ . It is well known that, when studying iterative (non-convex) solutions to problems, the global measurements’ setting is easier to study, and one can obtain better sample complexity guarantees for it, as compared to its non-global counterpart [27]. For example, we can compare guarantees for iterative low rank matrix sensing (LRMS) from iid Gaussian linear projections with those for low rank matrix completion (LRMC) when assuming the iid Bernoulli model on observed entries [24], [25], [28]. LRMS can be solved using an iterative algorithm with nearly order optimal number of measurements, e.g., the approach of [28] needs  $mq \geq C(n+q)r \log \max(n, q)$ , while even the best iterative LRMC guarantee (under the iid Bernoulli measurement model) [26] needs  $mq \geq C(n+q)r^2 \log^2(nq) \log^2(1/\epsilon)$ .

In this sense, the problem closest to ours that is extensively studied is LRMC. Of course LRMC involves recovery from completely local but linear measurements of  $\mathbf{X}^*$ , while LRPR involves recovery from nonlinear but column-wise global measurements. For this reason, for LRPR, in the regime of

$q$  significantly larger than  $n$ , the required sample complexity  $m$  is very small. As an example, suppose that  $q \geq nr^4$ , then we only need  $m \geq C \max(r, \log q, \log(1/\epsilon))$ . But this does not happen for LRMC.

We provide a comparison with the first and the best guarantees for non-convex (iterative) solutions for LRMC, sparse PR, and standard PR in Table I. These and other works are discussed in detail in Sec. I-D. As can be seen from the table, the first guarantee for iterative solutions to many problems is often sub-optimal (either needs more samples or more assumptions) compared to the best one that appeared later. Moreover, in the practical regime of  $r$  being order  $\log n$  or smaller, our LRPR sample complexity is as good or better than that of the best LRMC guarantee.

2) *Novelty*: In the absence of relevant existing work for even solving the linear version of our problem, except for a convex solution for the  $a_{ik} = a_i$  case (which is a significantly different problem), developing and analyzing our approach was not a straightforward extension of existing ideas. For example, the AltMinLowRaP algorithm itself is not just alternating standard PR over  $\mathbf{U}$  and  $\mathbf{B}$ . The PR problem for recovering  $\mathbf{U}^*$  given an estimate of  $\tilde{\mathbf{B}}^*$  is significantly different from standard PR; see Sec. II-A1 below.

For the above reasons, it is also not possible to directly modify proof techniques from existing work. We borrow some ideas from LRMC [25] and standard PR results [9], [10]. But the major difference is that concentration bounds need to be applied differently than for either of these problems. (i) The LRMC guarantees use results for Bernoulli random variables (which is a much more well-developed literature that has also been studied in the context of random graphs). In our setting, the random variables are not Bernoulli and not even bounded. Hence we rely on the sub-exponential Bernstein inequality [29, Theorem 2.8.1] and the fact that the product of two sub-Gaussian random variables is a sub-exponential [29, Lemma 2.7.7]. A second difference is that LRMC results do not need to deal with the phase error term. (ii) Standard PR results do have a phase error term and do deal with unbounded random variables using results from [29]. But they do not have to prove concentration using a set of  $mq$  measurements that are not identically distributed and, on first glance, may not even be “similar enough” to get a useful result. The “similarity” that is needed is of the following form: the maximum sub-exponential norm of any of the  $mq$  random variables being summed is not much larger than its average value. For each term, we have to carefully exploit the right incoherence assumption to show that this holds.

### C. Motivation for studying Low Rank PR (LRPR)

Low rank is a commonly used model in many dynamic biomedical imaging applications since (i) such images cannot change too much from one frame to the next, and (ii) these images are taken in controlled settings and so there are no fast changing foreground occlusions to worry about<sup>1</sup>. For

<sup>1</sup>Occlusions by moving objects or persons in the foreground are a common feature in computer vision problems such as surveillance or autonomous vehicles etc; for such videos a sparse + low-rank model is more appropriate

example, it is an important part of many practically useful fast compressive dynamic MRI solutions, e.g., see [19], and follow-up works [30], [31]<sup>2</sup>. In a similar fashion, a low sample complexity solution to LRPR can enable fast or low-cost dynamic phaseless imaging in applications such as solar imaging when the sun’s surface properties gradually change over time [32], or Fourier ptychographic imaging of live biological specimens and other dynamic scenes [21], [33]. Suppose the scene resolution is  $n$  and the total number of captured frames is  $q$ . If the dynamics is approximated to be linear and slow changing, with most of the change being explained by  $r$  linearly independent factors, then the matrix formed by stacking the vectorized image frames next to each other can be modeled as a rank- $r$  matrix plus small modeling error. In typical settings,  $r \ll \min(n, q)$  is a valid assumption, making the unknown images’ matrix approximately low-rank.

In all the above applications, *measurement acquisition is either expensive or slow*. For example, Fourier ptychography is a technique for super-resolution in which each of a set of low resolution cameras measures the magnitude of a different band-pass filtered version of the target high-resolution image. To get enough measurements per image, one either needs many cameras (expensive), or one needs to move a single camera to different locations to acquire the different bands [33]. This can make the acquisition process slow. By exploiting the low-rank assumption, it is possible to get an accurate reconstruction with using fewer total measurements (fewer cameras in this example). This has been demonstrated experimentally for dynamic Fourier ptychography in our recent work [21] and its follow-up [34]. Moreover, it is indeed *practically valid to assume that a different measurement matrix  $\mathbf{A}_k$  is used for each different signal/image*. In the ptychography example, this would correspond to using a different randomly selected subset of cameras at different times  $k$ . Modified cameras can also be designed that save power by switching off a different set of pixels at different time instants. We have explored both settings in [21].

Another practical point that should be mentioned is that, often, *in practice, a very small value of rank  $r$  suffices*. For example, we used  $r = 20$  in all our experiments on image sequences with  $n = 32400$  in [21]. In follow-up work [34], we show that just  $r = 5$  suffices for the same datasets, as long as a “modeling error correction step”<sup>3</sup> is applied to the output of AltMinLowRaP.

Lastly, in comparison to sparsity or structured sparsity priors, the low rank prior is a significantly more flexible one since it does not require knowledge of the dictionary or basis in which the signal is sufficiently sparse. In Table III, we demonstrate this via a simple experiment. We compare AltMinLowRaP with the most recent provable sparse PR algorithm [18], CoPRAM, applied with assuming wavelet sparsity

<sup>2</sup>These follow-up works exploit both low-rank of the entire sequence and wavelet sparsity of each image to further reduce the number of measurements needed in practice. This is the so-called “sparse and low-rank” model which is very different from sparse+low-rank model where the sparse component models occlusions by foreground moving objects.

<sup>3</sup>This step applies a few iterations of any standard PR approach column-wise to the output of AltMinLowRaP, in order to recover some of the modeling error in the low-rank assumption.

(which is a generic choice for any piecewise smooth image, but is not necessarily the best choice for the particular image). As can be seen, AltMinLowRaP has significantly superior performance not just for the real image sequence, but also for its deliberately sparsified version. The sparsified sequence had sparsity level  $s \approx 0.1n \approx 100$  and we provided CoPRAM with this ground truth. AltMinLowRaP used just  $r = 15$  for all three results and still had much lower reconstruction error than CoPRAM. Details of this experiment are provided in Sec. IV-B. Moreover, low-rank also includes certain types of dynamic sparsity models (those with fixed or very small changes in support over time) as special cases.

#### D. Review of Related Work

##### 1) Linear version of our problem: Compressive PCA:

While one would think that the linear (with phase) version of our problem would have been extensively studied, this is not true. There have been a few algorithmic solutions for this problem in prior work [35], [36], and attempts to prove some facts theoretically. Follow-up work consists of an Asilomar 2014 paper [37] that solves the general PCA problem for any (not necessarily low rank) matrix  $\mathbf{X}^*$ , but does not discuss recovery of  $\mathbf{X}^*$ . We explain these in detail in Sec. II-C.

2) Our measurement model, but with same set of  $m$  measurement vectors used for all signals, and its linear version: The covariance sketching problem, e.g., see [23], assumes that measurements satisfying (1), but with  $\mathbf{a}_{ik} = \mathbf{a}_i$ , are available. One aggregates these over  $k$  to get  $\mathbf{y}_i := \sum_k \mathbf{y}_{ik} = \mathbf{a}_i'(\sum_k \mathbf{x}_k^* \mathbf{x}_k^{*\prime}) \mathbf{a}_i = \mathbf{a}_i' \mathbf{X}^* \mathbf{X}^{*\prime} \mathbf{a}_i$ . This aggregation is what ensures that the memory complexity of storing the measurements is order  $m$  and not  $mq$  (which is what we need). Also, the aggregated  $\mathbf{y}_i$  is a function of  $\mathbf{X}^* \mathbf{X}^{*\prime}/q$  only in the  $\mathbf{a}_{ik} = \mathbf{a}_i$  setting, otherwise it is a meaningless quantity. Assuming random zero mean iid signals  $\mathbf{x}_k^*$ ,  $\mathbf{X}^* \mathbf{X}^{*\prime}/q$  is the empirical covariance matrix of a signal. The question is can we recover  $\mathbf{X}^* \mathbf{X}^{*\prime}/q$  from the scalar sketches  $\mathbf{y}_i, i = 1, 2, \dots, m$  with using  $m$  much smaller than  $nq$ , when  $\mathbf{X}^*$  is low rank (or has other structure)? When  $\mathbf{X}^*$  is rank  $r$ , the result of [23] proves that  $m$  of order  $(n+q)r$  suffices to estimate the empirical covariance from the aggregated measurements  $\mathbf{y}_i$  if one solves an appropriately defined nuclear norm minimization problem. For solving LRPR, we need a much smaller  $m$  than this. The reason is we assume independent  $\mathbf{a}_{ik}$ 's for different  $k$ , and we assume we have access to each individual  $\mathbf{y}_{ik}$ .

The linear version of the above problem, but with random noise added, is considered in Corollary 3 of [22] and in the remark immediately below it. In our notation, its measurement model can be written as  $\mathbf{y}_{ik} = \langle \mathbf{a}_i \mathbf{e}_k', \mathbf{X}^* \rangle + \mathbf{w}_{ik}$  where  $\mathbf{w}_{ik}$  is iid zero mean Gaussian noise with variance  $\nu^2$ . This result (specialized to the exact low rank case) shows that, whp, a nuclear norm minimization based solution will recover an estimate  $\hat{\mathbf{X}}$  of  $\mathbf{X}^*$  that satisfies  $\|\mathbf{X}^* - \hat{\mathbf{X}}\|_F^2 \leq C \nu^2 r(n+q)/m$ . In this paper, the focus is on using the low rank property to achieve noise robustness. If the low rank property was not used, and one attempted to recover the columns individually, the recovery error bound would scale as  $\nu^2 nq/m$  which is much larger. This paper also studies other settings

of recovering an approximately low rank matrix from linear measurements.

3) *Tangentially related work:* Some other tangentially related work includes: (i) computing the approximate rank  $r$  approximation of any matrix (need not be low rank) from its random sketches [38], [39] (sketched SVD); (ii) compressed covariance estimation using different sketching matrices for each data vector, but without the low-rank assumption [40]; and (iii) a generalization of low-rank covariance sketching [41]: this attempts to recover an  $n \times r$  matrix  $\mathbf{U}^*$  from measurements  $\mathbf{y}_i = \|\mathbf{a}_i' \mathbf{U}^*\|^2$  with  $r \ll n$ . When  $r = 1$ , this is the standard PR problem. In the general case, this is related to covariance sketching described above, but not to our problem.

##### 4) Linear low-rank matrix recovery – LRMS and LRMC:

Low-rank matrix recovery problems with linear measurements that have been extensively studied can be split into two kinds - those with “global measurements” and those without. “Global measurements” means that each measurement contains information about the entire structured quantity-of-interest, here the low-rank matrix. Such problems are called “affine rank minimization problems” or “low-rank matrix sensing” (LRMS) and involve recovery of  $\mathbf{X}^*$  from  $\mathbf{y}_i = \langle \mathbf{A}_i, \mathbf{X}^* \rangle$  with  $\mathbf{A}_i$  being dense matrices (typically iid Gaussian), see for example, [25], [28], [42]–[44]. More recent work studies the case of  $\mathbf{A}_i = \mathbf{a}_i \mathbf{a}_i'$  [45], [46]. Low-rank Matrix Completion (LRMC) is the completely local measurements’ setting that involves recovering  $\mathbf{X}^*$  from measurements of a randomly (iid Bernoulli) selected subset of its entries [12], [24]–[26], [47], [48]. Thus  $\mathbf{A}_i$ ’s are one-sparse matrices.

A precursor to LRMS is compressive sensing (CS) of sparse signals. This has the same property as LRMS, it involves recovering a sparse  $\mathbf{x}^*$  from  $\mathbf{y}_i = \langle \mathbf{a}_i, \mathbf{x}^* \rangle$  with  $\mathbf{a}_i$  being dense (sub-)Gaussian random vectors. Similar to CS, even for LRMS, it is typically possible to prove a simple (sparse or low-rank) restricted isometry property which simplifies the rest of the analysis. Our problem setting is different from, and more difficult than, LRMS. There are no “global measurements” of the entire  $\mathbf{X}^*$  and, moreover, the measurements’ phase/sign is unknown. In this sense it is closer to LRMC than to LRMS. But, unlike LRMC, we do have column-wise global measurements. This is why, for our problem only incoherence of right singular vectors suffices, while LRMC needs incoherence of both left and right singular vectors.

In summary, the problem closest to ours that is well studied is LRMC. The first iterative solution to LRMC was [47]. However, its guarantee does not bound the required number of algorithm iterations, and thus its time complexity cannot be bounded. The first iterative LRMC solution with bounded time complexity, AltMinComplete [25], needs a sample complexity of about  $C\kappa^4 \mu^2 nr^{4.5} \log(1/\epsilon)$  and assumes sample-splitting (a different independent set of measurements is used at each iteration). The most recent work on LRMC [12] removes the sample-splitting requirement and has bounded time complexity, but its sample complexity was  $Cnr^3 \log^6(n)$ . The best iterative LRMC solution in terms of sample complexity [26] needs  $Cnr^2 \log^2 n \log^2 \kappa/\epsilon$  samples but needs sample-splitting. In the practical regime  $r \in O(\log n)$ , clearly our

sample complexity for LRPR is comparable to the best LRMC guarantee [26]. For all values of  $r$ , it is slightly better than the first bounded time iterative LRMC solution [25]. Time-wise the AltMin algorithm for LRMC is faster than our AltMinLowRaP algorithm for LRPR (see Table I). This is because the LS problem to be solved in each AltMin step of LRMC involves a matrix with a large number of zeros but this is not the case for LRPR.

5) *Sparse PR*: Sparse PR is a somewhat related problem to ours since it involves PR with a different type of structural assumption on the signals. But, as noted earlier, there is a major difference. Sparse PR involves recovery from global measurements of the entire sparse vector. It can be understood as the phaseless version of Compressive Sensing with random Gaussian measurements. The global measurements' setting is typical easier than its non-global counterparts. Provably correct sparse PR approaches include convex relaxation approaches such as  $\ell_1$ -PhaseLift [14]; older combinatorial methods [49]; and a series of fast iterative approaches: (i) AltMinSparse [6], (ii) Sparse Truncated Amplitude Flow (SPARTA) [16], (iii) Thresholded WF [17] and CoPRAM [18]. The first two fast nonconvex approaches – AltMinSparse and SPARTA – needed to assume a lower bound on the minimum nonzero entry of  $\mathbf{x}$ . In follow-up work on Thresholded WF and then on CoPRAM, this extra assumption was removed. All four results need at least order  $s^2 \log n$  measurements. A summary of comparison of our work with LRMC, sparse PR, and standard PR is provided in Table I.

### E. Organization

We present our algorithm and guarantee along with a detailed discussion of the novelty of our proof techniques in Sec. II given next. The overall proof is given in Sec. III. The lemmas introduced in Sec. III are proved in Appendix A. Numerical experiments are provided in Sec. IV. We develop extensions to phaseless subspace tracking in Sec. V. We conclude in Sec. VI with a detailed discussion of ongoing and future work.

## II. LOW RANK PR: ALGORITHM AND GUARANTEE

### A. AltMinLowRaP algorithm

The complete algorithm is summarized in Algorithm 1. We explain its main idea next and then explain the details.

1) *Main idea*: AltMinLowRaP minimizes the following

$$\sum_{k=1}^q \|\mathbf{y}_k - |\mathbf{A}_k' \mathbf{U} \mathbf{b}_k| \|^2 \quad (5)$$

alternatively over  $\mathbf{U}, \mathbf{B}$  with the constraint that  $\mathbf{U}$  is a basis matrix. To initialize, we develop a spectral initialization for  $\text{Span}(\mathbf{U}^*)$  explained below. At a top level, the alternating minimization (AltMin) can be understood as alternating PR: minimize (5) over  $\mathbf{B}$  keeping  $\mathbf{U}$  fixed at its current value and then vice versa. But there are important differences between the two PR problems and how they can be solved.

- Given an estimate of  $\text{Span}(\mathbf{U}^*)$ , denoted  $\mathbf{U}$ , and assuming that  $\mathbf{U}$  contains orthonormal columns and is independent of the measurement vectors, the recovery of each  $\tilde{\mathbf{b}}_k^*$

is an  $r$ -dimensional “standard PR” problem. We can use either of [9], [10] to solve it. The estimate that we get, denoted  $\hat{\mathbf{b}}_k$ , is actually an estimate of  $\mathbf{g}_k := \mathbf{U}' \mathbf{U}^* \tilde{\mathbf{b}}_k^*$  which is a rotated version of  $\tilde{\mathbf{b}}_k^*$ . If  $\sin \Theta(\mathbf{U}, \mathbf{U}^*) \leq \delta$ , then, by the noisy PR result of [10], we can see that  $\|\hat{\mathbf{b}}_k - \mathbf{g}_k\| \leq C\delta \|\tilde{\mathbf{b}}_k^*\|$ .

- Given a previous estimate of  $\tilde{\mathbf{B}}^*$ , the update of  $\mathbf{U}^*$ , or equivalently of its vectorized version,  $\mathbf{U}_{vec}^*$ , is a significantly non-standard PR problem for two reasons. First, the “measurement vectors” for this PR problem are no longer independent or identically distributed. Second, and more importantly, by using the previous estimates of  $\mathbf{U}^*$  and of  $\tilde{\mathbf{b}}_k^*$ , with accuracy level  $\delta$ , we can get an estimate,  $\hat{\mathbf{x}}_k = \mathbf{U} \hat{\mathbf{b}}_k$ , of  $\mathbf{x}_k^*$  with the same accuracy level. With this, we can also get an estimate of the phase/sign of the measurements,  $\mathbf{c}_{ik} := \text{phase}(\mathbf{a}_{ik}' \mathbf{x}_k^*)$  with the same accuracy level. As a result, obtaining a new estimate of  $\mathbf{U}_{vec}^*$  becomes a much simpler Least Squares (LS) problem rather than a PR problem.
- We stress here that an argument similar to the above *does not apply* when recovering  $\tilde{\mathbf{b}}_k^*$ 's. The reason is, with a new estimate of  $\mathbf{U}^*$ , denoted  $\mathbf{U}^+$ , the previous estimate of  $\tilde{\mathbf{b}}_k^*$  becomes useless: (i) it is close to  $\mathbf{U}' \mathbf{U}^* \tilde{\mathbf{b}}_k^*$ , and not to  $\mathbf{U}^+ \mathbf{U}^* \tilde{\mathbf{b}}_k^*$ ; and (ii) we only estimate the span of the columns of  $\mathbf{U}^*$  accurately, so  $\mathbf{U}$  or  $\mathbf{U}^+$  are close to  $\mathbf{U}^*$  (and hence to each other) only in the subspace error  $\sin \Theta$  (and not in spectral or Frobenius norm). As a result an estimate of the form  $\mathbf{U}^+ \hat{\mathbf{b}}_k$  cannot be shown to be close to  $\mathbf{x}_k^*$ .<sup>4</sup>

2) *Details*: A different way to understand AltMinLowRaP is to split it into a three-way AltMin problem over  $\mathbf{U}^*$ ,  $\tilde{\mathbf{b}}_k^*$ 's, and  $\mathbf{c}_{ik}$ 's. This discussion assumes “sample-splitting”: a new set of  $mq$  measurements is used for each update of  $\mathbf{U}^*$  and another new set for each update of  $\tilde{\mathbf{B}}^*$ . Thus the total number of measurements used is  $2mq$  times the number of iterations.

- 1) At each new iteration, we first obtain a new estimate of  $\mathbf{U}^*$  using previous estimates of  $\tilde{\mathbf{b}}_k^*$ 's and of the measurements' phases  $\mathbf{c}_{ik}$ 's. This is a LS problem, see line 10 of Algorithm 1. The output of the LS step may not have orthonormal columns; this is easily resolved by a QR decomposition step after it (line 11).
- 2) Given a new estimate of  $\mathbf{U}^*$ , we recover each  $\tilde{\mathbf{b}}_k^*$ , by solving easy individual  $r$ -dimensional standard PR problems. These are easy because  $m$  of order  $r$  suffices.
- 3) Given a good estimate of  $\tilde{\mathbf{b}}_k^*$  and of  $\mathbf{U}^*$ , we can get an equally good estimate of  $\mathbf{x}_k^*$  and hence of the signs/phases  $\mathbf{c}_{ik}$ 's.

Consider the PR step to update  $\tilde{\mathbf{b}}_k^*$ 's. Observe that we can rewrite  $\mathbf{y}_{ik}$  as  $\mathbf{y}_{ik} = |\langle \mathbf{a}_{ik}, \mathbf{U}^* \tilde{\mathbf{b}}_k^* \rangle| = |\langle (\mathbf{U}^* \mathbf{a}_{ik}), \tilde{\mathbf{b}}_k^* \rangle|$ . If

<sup>4</sup>To understand this point easily, suppose both  $\mathbf{U}$  and  $\mathbf{U}^+$  are perfect estimates of  $\mathbf{U}^*$  in terms of the subspaces they span. Suppose  $\mathbf{U}^+ = \mathbf{U}^* \mathbf{R}_2$  and  $\mathbf{U} = \mathbf{U}^* \mathbf{R}_1$  where  $\mathbf{R}_1, \mathbf{R}_2$  could be any rotation matrices. Then it is easy to see that  $\|\mathbf{x}_k^* - \mathbf{U}^+ \hat{\mathbf{b}}_k\| \geq \|(\mathbf{U} - \mathbf{U}^+) \hat{\mathbf{b}}_k\| - \|\mathbf{x}_k^* - \mathbf{U} \hat{\mathbf{b}}_k\| \geq \|\mathbf{U}^* (\mathbf{R}_1 - \mathbf{R}_2) \hat{\mathbf{b}}_k\| - \delta \|\mathbf{x}_k^*\|$ . Since  $\mathbf{R}_1, \mathbf{R}_2$  can be any rotation matrices, e.g. one could have  $\mathbf{R}_2 = -\mathbf{R}_1$ ; in this case, the above error is lower bounded by  $2\|\hat{\mathbf{b}}_k\| - \delta \|\mathbf{x}_k^*\| \geq (2 - C\delta) \|\mathbf{x}_k^*\|$ . The last inequality follows since  $\|\hat{\mathbf{b}}_k\| = \|\hat{\mathbf{x}}_k\|$  and  $\text{dist}(\mathbf{x}_k^*, \hat{\mathbf{x}}_k) \leq C\delta \|\mathbf{x}_k^*\|$  by Lemma 3.2. The error can thus be even higher than using a zero vector to estimate  $\mathbf{x}_k^*$ .

$U^*$  were known, we would have a noise-free standard PR problem. If, at the  $t$ -th iteration, instead, we have a good estimate of  $\text{Span}(U^*)$ , denoted  $U^t$ , we can still recover the  $\tilde{b}_k^*$ 's by solving a noisy version of the same problem. Due to sample splitting,  $U^t$  is independent of  $a_{ik}$ 's and so the design vectors  $(U^t a_{ik})$  are still iid standard Gaussian. Any standard PR solution can be used, here we use Reshaped Wirtinger flow (RWF) [10]. The noise seen by RWF is proportional to  $\sin \Theta(U^t, U^*)$ . The error in the output of RWF cannot be lower than this value [10]. Thus, one needs to use just enough iterations of RWF so that the error at the end of the final RWF iteration is proportional to  $\sin \Theta(U^t, U^*)$ . Since we prove geometric convergence of  $\sin \Theta(U^t, U^*)$ , we can let  $T_{RWF,t}$  grow linearly with  $t$ .

3) *Initialization:* To obtain the initialization, we develop a careful modification of the truncated spectral initialization idea from [9], [20]. First assume that  $r$  is known. We initialize  $\hat{U}$  as the top  $r$  left singular vectors of the following matrix:

$$Y_U = \frac{1}{mq} \sum_{k=1}^q \sum_{i=1}^m y_{ik}^2 a_{ik} a_{ik}' \mathbb{1}_{\{y_{ik}^2 \leq C_Y \frac{1}{mq} \sum_{i,k} y_{ik}^2\}}. \quad (6)$$

where  $C_Y$  is a constant that decides the truncation threshold (which measurements are too large in magnitude compared to the mean energy of the measurements and should be discarded). For our guarantee, we set it equal to  $9\kappa^2\mu^2$ . In practice a good value can be chosen by experimentation and cross-validation or by using the ideas in [50]. An alternative approach is to set it differently for each  $k$  as done in [20]. This approach does not require knowledge of  $\kappa$  or  $\mu$ , but it results in a worse lower bound on just  $m$ . We discuss the effect of this choice in Remark 2.2 and also in Sec. II-D3.

To understand why the above approach works, first consider the above matrix *with the indicator function removed*. Then it is not hard to see that its expected value equals  $(1/q)[U^*(\Sigma^{*2})U^{*'} + 2\text{trace}(\Sigma^{*2})I]$ , and so its span of top  $r$  singular vectors equals  $\text{Span}(U^*)$ . Hence, with large enough  $mq$ , the same should approximately hold for the original matrix. However, when using  $Y_U$  with the indicator function removed, a few “bad” measurements (those with very large magnitude  $y_{ik}^2$  compared to their empirical mean over  $i, k$ ) can heavily bias its value. To mitigate this effect, and get a good initialization in spite of it, we will need a larger value of  $mq$ . Using the indicator function helps truncate the summation to only sum over the “good” measurements, and as a result a smaller value of  $mq$  suffices. Mathematically, this helps ensure that  $Y_U$  is close to a matrix that can be written as  $\sum_{ik} w_{ik} w_{ik}'$  with  $w_{ik}$ 's being iid sub-Gaussian vectors (instead of sub-exponential in the case without truncation) [9].

We can also use  $Y_U$  to correctly estimate  $r$  whp by using the fact that, when  $m$  and  $q$  are large, the gap between its  $r$ -th and  $(r+1)$ -th singular value is close to  $\sigma_{\min}^* / q$ . With this idea, we estimate  $r$  as given in the first step of Algorithm 1. As explained in [20], another way of estimating the rank is to set  $\hat{r} = \arg \max_j (\sigma_j(Y_U) - \sigma_{j+1}(Y_U))$ . This approach does not require knowledge of any model parameters. Hence it is easily applicable for real data (even without training samples being available). However, it works under the assumption that

---

**Algorithm 1** AltMin-LowRaP: Alt-Min for Phaseless Low Rank Recovery

---

- 1: Parameters:  $T, T_{RWF,t}, \omega$ .
- 2: Partition the  $m_{\text{tot}}$  measurements and design vectors for each  $x_k^*$  into one set for initialization and  $2T$  disjoint sets for the main loop.
- 3: Set  $\hat{r}$  as the largest index  $j$  for which  $\lambda_j(Y_U) - \lambda_n(Y_U) \geq \omega$  where  $Y_U$  is in (6).
- 4:  $U^0 \leftarrow \hat{U}^0 \leftarrow$  top  $\hat{r}$  singular vectors of  $Y_U$  defined in (6).
- 5: **Initialize  $U$**
- 6:  $\hat{b}_k^t \leftarrow \text{RWF}(\{y_k^{(t)}, U^{t'} A_k^{(t)}\}, T_{RWF,t})$  for each  $k = 1, 2, \dots, q$  (RWF: Reshaped WF [10]).
- 7: **Update  $\hat{B}$**
- 8:  $\hat{C}_k \leftarrow \text{Phase}(\hat{A}_k^{(T+t)'} \hat{x}_k^t)$  for each  $k = 1, 2, \dots, q$ .
- 9: **Update  $\hat{C}_k$ 's**
- 10: Get  $B^t$  by QR decomp:  $\hat{B}^t \stackrel{\text{QR}}{=} R_B^t B^t$ .
- 11:  $\hat{U}^{t+1} \leftarrow \arg \min_{\hat{U}} \sum_{k=1}^q \|\hat{C}_k y_k^{(T+t)} - \hat{A}_k^{(T+t)'} \hat{U} \hat{b}_k^t\|^2$ .
- 12: **Line 10, 11: Update  $U$**
- 12: **end for**

---

consecutive nonzero singular values of  $X^*$  are close (do not have significant gap), see [20, Corollary 3.7] for one precise statement of this claim.

The main idea of the initialization step explained above was first developed in our previous work [20], we explain the difference later in Sec. II-D3. Also, as pointed out by an anonymous reviewer, a matrix that is related to  $Y_U$  was used in earlier work [35], [36] to try to solve what can be called the linear version of our problem; see Sec. II-C. We were not aware of this work when developing our approach. Our approach was developed independently in [20] as a modification of the truncated spectral initialization idea from PR literature [9] and then modified here.

We summarize the complete algorithm in Algorithm 1. As is commonly done in existing literature, e.g., see [6], [25], in order to obtain a provable guarantee in a simple fashion, we assume sample-splitting. Since we prove geometric convergence of the iterates, this increases the required sample complexity by a factor of only  $\log(1/\epsilon)$ . In our empirical evaluations, we reuse the same set of measurements.

### B. Guarantee

We have the following guarantee.

**Theorem 2.1.** *Consider Algorithm 1. Assume that the  $y_{ik}$ 's satisfy (1) with  $a_{ik}$  being iid standard Gaussian; and  $X^*$  is an  $n \times q$  rank- $r$  matrix that satisfies right-incoherence with parameter  $\mu$ . Set  $T := C \log(1/\epsilon)$ ,  $T_{RWF,t} = C(\log r + \log \kappa + t(\log(0.7) / \log(1 - c)))$ ,  $\omega = 1.3\sigma_{\min}^* / q$ , and  $C_Y = 9\kappa^2\mu^2$  in (6). Assume that, for the initialization step and for each new update, we use a new set of  $m$  measurements with  $m$  satisfying*

$mq \geq C\kappa^{12}\mu^4 \cdot nr^4$  and  $m \geq C \max(r, \log q, \log n)$ . Then, with probability (w.p.) at least  $1 - Cn^{-10}$ ,

$$\sin \Theta(\mathbf{U}^*, \mathbf{U}^T) \leq \epsilon, \text{ mat-dist}(\hat{\mathbf{X}}^T, \mathbf{X}^*) \leq \epsilon \|\mathbf{X}^*\|_F$$

and  $\text{dist}(\hat{\mathbf{x}}_k^T, \mathbf{x}_k^*) \leq \epsilon \|\mathbf{x}_k^*\|$  for each  $k$ . Moreover, after the  $t$ -th iteration,

$$\sin \Theta(\mathbf{U}^*, \mathbf{U}^t) \leq 0.7^t \delta_{\text{init}}, t = 0, 1, 2, \dots, T$$

where  $\delta_{\text{init}} = \frac{c}{\kappa^2 r}$ . Similar bounds also hold on the error in estimating  $\mathbf{x}_k^*$ 's. The time complexity is  $mqnr \log^2(1/\epsilon)$ .

Proof: We prove this in Sec. III.

**Remark 2.2.** If we are willing to tolerate another lower bound on just  $m$  of  $m \geq Cr^4$ , we can improve the dependence on  $\kappa, \mu$  to  $\kappa^8\mu^2$ . This will require the following change: define  $\mathbf{Y}_U$  as done in [20]: use  $9 \sum_{i=1}^m y_{ik}^2/m$  as the threshold inside the indicator function. A second advantage of using this is that it allows one to use  $C_Y = 9$  instead of  $C_Y$  that depends on  $\kappa, \mu$ . The disadvantage is of course a more stringent lower bound on  $m$ .

**Remark 2.3.** To understand the time complexity, observe that the most expensive step at each iteration is the update of  $\mathbf{U}^*$ . This requires solving an LS problem of recovering an  $nr$  length vector from  $mq$  measurements. One can solve this by conjugate gradient descent with a cost of  $mq \cdot nr \cdot \log(1/\epsilon)$  [6]. This times the total number of iterations,  $T = C \log(1/\epsilon)$  gives the complexity of the algorithm after initialization. Consider the initialization step. Observe that the top  $r$  singular vectors of  $\mathbf{Y}_U$  are also the top  $r$  singular vectors of an  $mq \times n$  matrix  $\mathbf{G}$  whose columns are given by  $a_{ik} y_{ik}$  times the indicator function used in (6). Thus  $\mathbf{Y}_U = \mathbf{G}\mathbf{G}'$ . Since computing the  $r$ -SVD of an  $a \times b$  matrix to  $\delta$  accuracy needs time of order  $abr \log(1/\delta)$  [51], the SVD needed for initialization can be computed in time  $Cmq \cdot n \cdot r \cdot \log(1/\delta_{\text{init}}) = Cmqnr \log r$  where  $\delta_{\text{init}}$  is the error level to which the initialization needs to be accurate. As explained later,  $\delta_{\text{init}} = c/r$  suffices.

Theorem 2.1 implies that one can achieve geometric convergence as long as the sample complexity  $m_{\text{tot}} := (2T + 1)m$  satisfies  $m_{\text{tot}}q \geq C\kappa^{12}\mu^4nr^4 \log(1/\epsilon)$  along with  $m_{\text{tot}} \geq C \max(r, \log q, \log n) \log(1/\epsilon)$ . The second lower bound is very small and essentially redundant<sup>5</sup> except when  $q \geq Cnr^4$ .

Notice that the LRPR sample complexity is significantly better than that of standard (unstructured) PR methods which necessarily need  $m = Cn$  samples per signal (matrix column). For fixed  $m$  and  $q$ , LRPR time complexity is about  $r$  times worse than that of standard PR. But, if we use the smallest value of  $mq$  needed by each method to get an  $\epsilon$ -accurate estimate, AltMinLowRaP is actually faster when  $r$  is small: its needs time of order  $n^2r^5 \log(1/\epsilon)$  while standard PR methods need time of order  $n^2q$ . We demonstrate this fact experimentally in Fig 2.

The minimum number of samples needed to recover an  $n \times q$  matrix of rank  $r$  is  $(q + n)r$ . Thus, in general, our sample

<sup>5</sup>We need this lower bound because we recover the  $q \tilde{\mathbf{b}}_k^*$ 's individually by solving a standard PR problem for each. This step works correctly w.p. at least  $1 - 2q \exp(r - cm)$ .

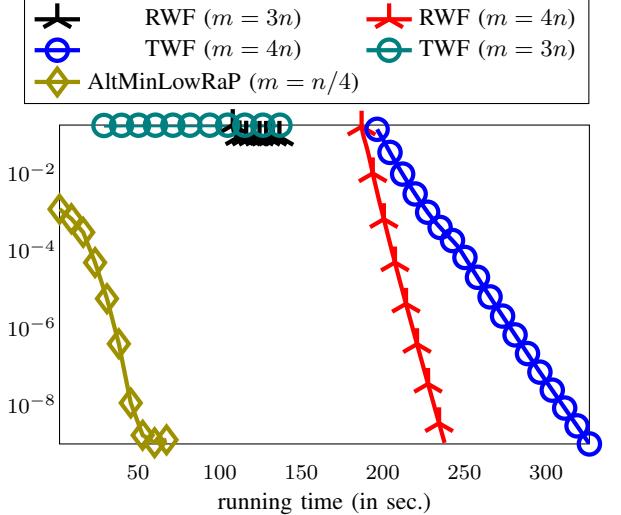


Fig. 2: Recovery error versus time-taken plot with time in seconds. The “time” here is the time taken in seconds to reach a certain error level. We generate this plot as explained in Sec. IV-A: plot the time taken until end of iteration  $t$  on the x-axis and plot the error at the end of iteration  $t$  on the y-axis. This figure illustrates the fact that if we use the lowest allowed value of  $m$  for each approach, AltMinLowRaP is in fact faster than RWF or TWF, for any given value of desired recovery error. Notice that RWF and TWF fail for  $m = 3n$ , but work for  $m = 4n$ . AltMinLowRaP works with just  $m = n/4$  and with this value of  $m$  it is roughly 5-times faster than both RWF (with  $m = 4n$ ) and TWF (with  $m = 4n$ ) for any desired level of error  $\epsilon$ .

complexity is  $r^3 \log(1/\epsilon)$  times worse than its order-optimal value. As noted earlier, in problem settings like ours, where the measurements are not global, non-convex algorithms typically do need more than the order-optimal number of samples. Since neither our problem nor its linear version have any complete provable guarantees for correct recovery in existing work, LRMC is the closest problem to ours that has been extensively studied and that also uses non-global measurements. Sparse PR is the other somewhat related problem to ours but it is easier because it involves recovery from global measurements of a sparse vector. Table I provides a comparison of our guarantee with the first and best results for both problems. More details are in Sec. I-D. As can be seen from the table, our sample complexity compares favorably with that for the first non-convex solution for LRMC that has bounded time complexity (bounds the required number of iterations) [25]. In the practical regime of  $r$  being order  $\log n$ , it even compares with the best iterative LRMC result [26]. Also, the the first guarantee for iterative solutions to both problems is sub-optimal (needs more samples or more assumptions) compared to the best one that appeared later.

We should reiterate that, (i) in many practical applications, a small value of  $r$  suffices, e.g., we used  $r = 5$  for images with  $n = 32400$  in [34] followed by a few iterations of model error correction via standard PR; and (ii) low rank is a more flexible model for dynamic imaging than sparsity. Also see Tables II and III.

### C. Linear version of our problem: Compressive PCA

Consider the linear (with phase) version of our problem – recover a low rank matrix  $\mathbf{X}^*$ , or its column span, from

$$\mathbf{y}_{ik} := \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle, \quad i = 1, 2, \dots, m, \quad k = 1, 2, \dots, q.$$

This problem can also be understood as “compressive PCA” although compressive PCA typically allows  $\mathbf{X}^*$  to also be only approximately low rank [35]–[37]. Clearly, both our algorithm and our guarantee, Theorem 2.1, directly apply to this simpler special case as long as the  $\mathbf{a}_{ik}$ ’s are iid Gaussian. Even when phases are available one could use the magnitude measurements and solve our harder problem instead. We can state the following corollary.

**Corollary 2.4.** *Consider the problem of recovering an  $n \times q$  rank- $r$  matrix  $\mathbf{X}^*$  from  $\mathbf{y}_{ik} := \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle, i = 1, 2, \dots, m, k = 1, 2, \dots, q$ . As long as right incoherence given in (3) holds, one needs  $m \geq C \frac{n}{q} r^4 \log(1/\epsilon)$  to recover  $\mathbf{X}^*$  and its column span to  $\epsilon$  accuracy (precisely defined in Theorem 2.1).*

Thus our work also provides a simple and fast AltMin solution that provably converges geometrically to the compressive PCA solution. Of course, when the phase/sign is known, a simpler version of our algorithm should work and a better guarantee should be obtainable. We discuss this further in the conclusions’ section.

In terms of existing related work, [35] provided an approach for solving the above problem when  $\mathbf{X}^*$  is only approximately low-rank. A follow-up paper [36] studied its modification where  $\mathbf{a}_{ik}$ ’s are sparse random vectors, e.g., “sparse Bernoulli” as they call it (each entry takes values  $-1, 0, 1$  with probabilities  $1/2s, 1-1/s, 1/2s$ ). Their proposed approach is related to our initialization step: one computes an estimate of the principal subspace of the column span of  $\mathbf{X}^*$  by computing the top  $r$  singular vectors of a matrix that is similar to  $\mathbf{Y}_U$  defined in (6). The difference is that the indicator function is absent, and they also sum over terms of the form  $\mathbf{a}_{ik} \mathbf{y}_{ik} \mathbf{y}_{jk} \mathbf{a}_{jk}'$  for all  $j \neq i$  which, in this linear setting, have nonzero expected value. To be precise, they compute the top  $r$  singular vectors of  $\sum_k \mathbf{A}_k \mathbf{y}_k \mathbf{A}_k' \mathbf{y}_k'$ . Denote such a matrix by  $\mathbf{Y}_{U,modified}$ . These works treat the columns  $\mathbf{x}_k^*$  as random vectors and assume that they have *nonzero mean* and provide a simple intuitive approach to estimate the mean. This is of course only possible because the measurements are linear, and cannot be done in our setting. Both works [35], [36] show that the expected value of  $\mathbf{Y}_{U,modified}$  is equal to  $c_1 \mathbf{I} + c_2 \mathbb{E}[\mathbf{X}^* \mathbf{X}^{* \prime} / q]$ , thus, the span of top  $r$  eigenvectors of the expected value equals the desired subspace: principal subspace of the population covariance,  $\mathbb{E}[\mathbf{X}^* \mathbf{X}^{* \prime} / q]$ . They also claim that as  $q \rightarrow \infty$ ,  $\mathbf{Y}_{U,modified}$  converges to its expected value, but do not provide a rate of convergence. Finally, they also upper bound the expected value of the Frobenius norm of the error between  $\mathbf{Y}_{U,modified}$  and its expected value for any value of  $q$ . Since this result also only bounds the expected value of the error, it cannot be used to get any useful information about the sample complexity  $m$  that is required even for the initialization step. Moreover, of course these works do not provide any guarantee on how to recover the entire matrix  $\mathbf{X}^*$ .

A later work in Asilomar 2014 [37] attempted to solve a generalization of the above problem: it provided an approach (which is again somewhat related to only our initialization step) and a guarantee for recovering the top  $r$  singular vectors of a general matrix  $\mathbf{X}^*$ . Specialized to the exactly rank  $r$  case, their result proves that, if each column of  $\mathbf{X}^*$  is bounded, and if  $m \geq \frac{1}{\epsilon} \max \left( \frac{n}{\sqrt{q}} \sqrt{r\kappa}, \frac{n}{q} r^2 \kappa^2 \right)$ , one can obtain an  $\epsilon$ -accurate recovery of the subspace  $\text{Span}(\mathbf{U}^*)$ . This is a much weaker result than ours: (i) its sample complexity  $m$  depends on  $1/\epsilon$  instead of  $\log(1/\epsilon)$ ; and (ii)  $m$  needs to grow as  $n/\sqrt{q}$  instead of as  $n/q$ . The reason it is weaker of course, is because they used a single step approach instead of an iterative algorithm.

### D. Discussion of Proof Techniques and Reason for Worse Dependence on $\kappa, r$

1) *Proof techniques:* Since even the linear version of our problem has not been studied theoretically (except for a convex solution for the  $\mathbf{a}_{ik} = \mathbf{a}_i$  case which is a significantly different problem), it is not possible to directly modify proof techniques from existing work. We do borrow some ideas from LRMC [25] or from standard PR results [9], [10]. But, as explained earlier in Sec. I-B2, concentration bounds need to be applied in a significantly different way than for either of these problems. The algebra for obtaining an expression that bounds the subspace recovery error between  $\hat{\mathbf{U}}^t$  and  $\mathbf{U}^*$  uses the overall approach of [25]. After this, the details are different because, for LRMC, the random variables are Bernoulli, while in our case they are unbounded sub-Gaussian or sub-exponential. Also, in LRPR, the sign/phase are unknown and this introduces an extra term that needs to be bounded, see Term2 defined in Lemma 3.9. To bound this we first use Cauchy-Schwarz to bound it by a product of two terms. One term is easy to deal with. For the second term, we borrow a lemma from the RWF paper [10] on standard PR, but the rest of our approach is different because of the need to prove concentration of a set of  $mq$  measurements that are not identically distributed (we need to carefully exploit right concurrence to ensure that they are “similar enough” to apply the concentration bounds for sums of products of sub-Gaussians).

Our initialization step uses the truncated spectral initialization approach, first introduced for PR in [9]. The proof for it also uses the overall approach of [9] but there are many important differences in proving the concentration bounds (see above). Our initialization is also significantly different from that of LRMC or LRMS. In these cases, the phase is known and thus one can come up with a matrix whose expected value equals  $\mathbf{X}^*$ . This is not possible for LRPR. The matrix whose top  $r$  singular vectors we compute has expected value  $2\beta_2 \mathbf{X}^* \mathbf{X}^{* \prime} + \beta_1 \mathbf{I}$ . This point has important implications for the sample complexity dependence on  $\kappa, r$ , we discuss this next.

2) *Worse dependence on  $\kappa, r$ :* Our sample complexity is comparable to that of [25] in terms of its dependence on  $n$  and  $r$ . However, our result has a worse dependence on  $\kappa$  because we have access only to phaseless measurements. Because of this, (1) our initialization step needs to use the matrix  $\mathbf{Y}_U$

and find its top  $r$  eigenvectors in order to get an initial estimate of the column span  $\mathbf{U}^*$ . For simplicity, consider  $\mathbf{Y}_U$  without the truncation (without the indicator function). Then, its expected value is  $\mathbf{U}^* \Sigma^{*2} \mathbf{U}^{*'} + 2\text{trace}(\Sigma^{*2}) \mathbf{I}$ . Notice that the condition number of the first term of this matrix (the term of interest) is  $\kappa^2$ . Because of this, when analyzing this step, we end up with a dependence of  $mq$  on  $\kappa^8 \mu^4 nr^2 / \delta_{\text{init}}^2$  (see Claim 3.1 and Sec. III-B where this claim is proved). Here  $\delta_{\text{init}}$  is the subspace error after the initialization step. Instead, the expected value of the matrix used for initialization of AltMinComplete [25] has expected value equal to  $\mathbf{X}^*$  and thus its condition number is just  $\kappa$ . (2) A second issue is as follows: because of magnitude-only measurements, we are having to deal with phase error (sign error) in each LS step that updates the estimate of  $\mathbf{U}^*$ . In bounding the phase error term – Term2 defined in Lemma 3.9 in Sec. III-C – we need to use the Cauchy-Schwarz inequality; see proof of Lemma 3.12. Because of this, when using the bound on Term2 from Lemma 3.12 to prove the main descent claim, Claim 3.4, we end up with a bound of the form  $\sin \Theta(\mathbf{U}^{t+1}, \mathbf{U}^*) \leq C\delta_t \sqrt{\delta_t} \sqrt{r\kappa}$  where  $\delta_t$  is the bound on the subspace error from the previous step. Thus, to ensure that this  $(t+1)$ -th step error is below  $0.7\delta_t$  (decays geometrically), we need to set  $\delta_t \leq c/(\kappa^2 r)$  for each iteration  $t$ , including the initialization. This is why we need  $\delta_{\text{init}} = c/(\kappa^2 r)$ . This, along with the fact that the init step needs  $mq \geq \kappa^8 \mu^4 nr^2 / \delta_{\text{init}}^2$ , implies that our sample complexity per iteration becomes  $C\kappa^{12} \mu^4 nr^4$ .

If somehow the Cauchy-Schwarz were not used, we would only need  $mq \geq \kappa^{10} \mu^4 r^3$ . Similarly, if we somehow did not use the loose bound  $\|\mathbf{X}\| \leq \|\mathbf{X}\|_F$  for rank  $r$  matrices at a few different places in the proof, we could remove another factor of  $r$ .

3) *Our previous work:* The only other work that also studies our problem is our previous work [20]. This introduced a series of heuristics and evaluated them experimentally. It also provided a guarantee for the initialization step of one of them. If we compare their main result (their Theorem 3.2) with ours, it required the following lower bound on just  $m$ :  $m \geq C \max(\sqrt{n}, r^4) / \epsilon^2$  in addition to a lower bound on  $mq$  that also depends on  $1/\epsilon^2$ . We remove the  $1/\epsilon^2$  dependence by analyzing the complete algorithm.

The first two requirements on just  $m$  are also significantly relaxed in our work because we study a significantly modified version of our previous algorithm. The most important algorithmic difference is that, both for initialization and for later iterations, we recover  $\tilde{\mathbf{b}}_k^*$ 's by solving the standard PR problem either fully (or, at least for enough iterations so that the error in recovering  $\tilde{\mathbf{b}}_k^*$ 's is of the same order as the subspace error in the estimate of  $\mathbf{U}^*$ ). This is what allows us to replace the strong requirement  $m \geq C\sqrt{n}$  that [20] needed by just  $m \geq Cr$ . This is also what enables us to get a complete guarantee for the entire algorithm. The algorithm in [20] used only one iteration of AltMinPhase [6] for obtaining a new estimate of  $\tilde{\mathbf{b}}_k^*$ 's using a new estimate of  $\mathbf{U}^*$ . With this, it was not possible to show that the recovery error of  $\tilde{\mathbf{b}}_k^*$ 's is of the same order as that of  $\mathbf{U}^*$ .

Our approach for initializing  $\mathbf{U}^*$  is taken from [20], but with a simple, but important, difference: the threshold in the

indicator function used for defining  $\mathbf{Y}_U$  in (6) now takes an average over all  $mq$  measurements (instead of over only the  $m$  measurements of the  $k$ -th column in [20]). This simple change allows us to use concentration over all the  $mq$  measurements (and design vectors) in every step of deriving the initialization guarantee for  $\mathbf{U}^*$ . This is what helps us eliminate the requirement of  $m \geq Cr^4$  on just  $m$  that was needed in [20].

### III. PROOF OF THEOREM 2.1

The proof borrows ideas from past works – [9], [20] (for initialization of  $\mathbf{U}^*$ ), [25] (the overall approach for getting a subspace error bound given in the Appendix), [52] for careful  $\epsilon$ -net arguments for unit Frobenius norm matrices, and [10] (for recovering  $\tilde{\mathbf{b}}_k^*$ 's, and in one step of trying to show that the phase error is small). We cite the relevant reference again where it is used. In Sec. III-A next, we provide the two main claims (one for initialization and one for the descent), the two other auxiliary lemmas needed for proving Theorem 2.1 and the theorem's proof using these. In Sec. III-B, we give the key lemmas needed for proving the initialization claim and also prove it. The same is done for the descent claim in Sec. III-C. Each of these subsections also provides the main ideas (intuition) used for proving the lemmas. We then prove all but one of the lemmas from this entire section in Appendix A. Lemma 3.9 which uses the overall approach of [25] is proved in Appendix B.

#### A. Overall lemmas and proof of Theorem 2.1

**Claim 3.1** (Rank estimation and Initialization of  $\mathbf{U}^*$ ). *Let  $\mathbf{U}_{\text{init}} = \hat{\mathbf{U}}^0$ . Pick a  $\delta_{\text{init}} < 0.25$ . Assume  $mq \geq \kappa^8 \mu^4 nr^2 / \delta_{\text{init}}^2$ . Set the rank estimation threshold  $\omega = 1.3\sigma_{\min}^{*2} / q$  (we can actually set the multiplier to any number between 0.025 and 1.5). Then, w.p. at least  $1 - 6n^{-10}$ , the rank is correctly estimated and*

$$\sin \Theta(\mathbf{U}_{\text{init}}, \mathbf{U}^*) \leq \delta_{\text{init}}.$$

Define

$$\mathbf{g}_k^t := (\mathbf{U}^t)' \mathbf{x}_k^* \text{ and } \mathbf{e}_k^t := (\mathbf{I} - \mathbf{U}^t \mathbf{U}^{t'}) \mathbf{x}_k^*. \quad (7)$$

It is easy to see that  $\mathbf{x}_k^* = \mathbf{U}^t \mathbf{g}_k^t + \mathbf{e}_k^t$  and so  $\mathbf{y}_{ik} = |((\mathbf{U}^t)' \mathbf{a}_{ik})' \mathbf{g}_k^t + \mathbf{a}_{ik}' \mathbf{e}_k^t|$ . Thus, we have a noisy PR problem to solve with the noise magnitude proportional to  $\|\mathbf{e}_k^t\|$  with  $\|\mathbf{e}_k^t\| \leq \sin \Theta(\mathbf{U}^t, \mathbf{U}^*) \|\mathbf{x}_k^*\|$ . We use RWF to solve it. RWF provides an estimate,  $\hat{\mathbf{b}}_k^t$ , of  $\mathbf{g}_k^t$ . Observe that

$$\mathbf{g}_k^t = (\mathbf{U}^t)' \mathbf{x}_k^* = ((\mathbf{U}^t)' \mathbf{U}^*) \tilde{\mathbf{b}}_k^t$$

is just a rotated version of  $\tilde{\mathbf{b}}_k^t$ . We show in the next lemma that, whp, the error in the RWF estimate,  $\text{dist}(\mathbf{g}_k^t, \hat{\mathbf{b}}_k^t)$ , is proportional to  $\sin \Theta(\mathbf{U}^t, \mathbf{U}^*)$ ; and the same is true for the error in  $\hat{\mathbf{x}}_k^t := \mathbf{U}^t \hat{\mathbf{b}}_k^t$ .

**Lemma 3.2** (Recovery of  $\tilde{\mathbf{b}}_k^*$ 's). *At iteration  $t$ , assume that  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^t) \leq \delta_t$ . Pick a  $\delta_b < 1$ . If  $m \geq Cr$ , and if we*

set  $T_{RWF,t} = C \log \delta_t / \log(1 - c)$ , then, w.p. at least  $1 - 2q \exp(-c\delta_b^2 m)$ , the following is true for each  $k = 1, 2, \dots, q$

$$\begin{aligned} \text{dist}(\mathbf{g}_k^t, \hat{\mathbf{b}}_k^t) &\leq C\delta_t \|\tilde{\mathbf{b}}_k^*\| = C\delta_t \|\mathbf{x}_k^*\|, \\ \text{mat-dist}(\mathbf{G}^t, \hat{\mathbf{B}}^t) &\leq C\delta_t \|\tilde{\mathbf{B}}^*\|_F = C\delta_t \|\mathbf{X}^*\|_F, \\ \text{dist}(\hat{\mathbf{x}}_k^t, \mathbf{x}_k^*) &\leq (C + 1)\delta_t \|\mathbf{x}_k^*\|. \end{aligned} \quad (8)$$

with  $C = \sqrt{1 + \delta_b} + 1$ .

Thus, if  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ , then the above bounds hold w.p. at least  $1 - n^{-10}$ .

From above,  $\hat{\mathbf{b}}_k^t$  is close to  $\mathbf{g}_k^t$  (which is a rotated version of  $\tilde{\mathbf{b}}_k^*$ ) for each  $k$ . We thus expect  $\hat{\mathbf{B}}^t$ , or equivalently  $\mathbf{B}^t$ , to also satisfy the incoherence assumption. We show next that this is indeed true if  $\delta_t$  is small enough. Recall that  $\hat{\mathbf{B}}^t \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}^t$ .

**Lemma 3.3** (Incoherence of  $\mathbf{B}^t$  implies incoherence of  $\mathbf{B}^t$ ). *Pick  $a \delta_b < 1/10$  and assume that  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . At iteration  $t$ , assume that  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^t) \leq \delta_t$  with  $\delta_t \leq \frac{0.25}{C\sqrt{r}\kappa}$ . If  $\mathbf{B}^t$  is  $\mu$ -incoherent, then, w.p. at least  $1 - n^{-10}$ ,  $\mathbf{B}^t$  is  $\hat{\mu}$ -incoherent with  $\hat{\mu} = C\kappa\mu$ .*

Finally, the next claim shows that the LS step to update  $\mathbf{U}$  reduces its error by a factor of 0.7 at each iteration. Its proof relies on the previous two lemmas and the fact that  $\hat{\mathbf{x}}_k^t$  close to  $\mathbf{x}_k^*$  implies that, with large probability, the phases (signs) of  $(\mathbf{a}_{ik}' \hat{\mathbf{x}}_k^t)$  and  $(\mathbf{a}_{ik}' \mathbf{x}_k^*)$  are equal too.

**Claim 3.4** (Descent Lemma). *At iteration  $t$ , assume that  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^t) \leq \delta_t$ . If  $\delta_t \leq \frac{c}{r\kappa^2}$ ,  $mq \geq C\kappa^3\mu^2nr^2/\delta_t^2$  and  $m \geq C \max(r, \log n, \log q)$  then w.p. at least  $1 - Cn^{-10}$ ,*

$$\sin \Theta(\mathbf{U}^{t+1}, \mathbf{U}^*) \leq 0.7\delta_t := \delta_{t+1}.$$

*Proof of Theorem 2.1.* The  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^t)$  bounds are an immediate consequence of Claims 3.1 and 3.4, along with setting  $\delta_{\text{init}} = c/\kappa^2 r$ ,  $\delta_t = 0.7^t \delta_{\text{init}}$  and  $\delta_b = 1/11$ . With these, we require  $mq \geq C\kappa^{12}\mu^4nr^4$  (for initialization) and  $mq \geq C\kappa^7\mu^2nr^4$  (for the descent steps), along with  $m \geq C \max(r, \log n, \log q)$ . The first lower bound on  $mq$  dominates. The other bounds of Theorem 2.1 follow by Lemma 3.2.  $\square$

We prove Claims 3.1 and 3.4 next in Sec. III-B and III-C. The proof of Lemmas 3.2 and 3.3 and of the lemmas needed for proving these two claims is postponed to Appendix A.

### B. Proof of Claim 3.1

In this section, we let  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(0)}$  and  $\mathbf{y}_{ik} := \mathbf{y}_{ik}^{(0)}$ .

The overall idea for proving this is inspired by the approach in [20] which itself borrows ideas from [9]. But there are many important differences because we define  $\mathbf{Y}_U$  differently in this work, see (6): the threshold in the indicator function now takes an average over all  $mq$  measurements (instead of over only the  $m$  measurements of the  $k$ -th column as in [20]). This simple change enables us to get a significantly improved result. It lets us use concentration over all the  $mq$  measurements (and design vectors) in each of the three steps of the proof. This is what helps eliminate the lower bound  $m \geq Cr^4$  on just  $m$  that

was needed in [20]. However, this also means that the proofs are much more involved (more quantities now vary with  $k$ ).

Recall the expression for  $\mathbf{Y}_U$  from earlier, and define matrices  $\mathbf{Y}_-(\epsilon_1)$  and  $\mathbf{Y}_+(\epsilon_1)$  as

$$\begin{aligned} \mathbf{Y}_U &= \frac{1}{mq} \sum_{ik} |\mathbf{a}_{ik}' \mathbf{x}_k^*|^2 \mathbf{a}_{ik} \mathbf{a}_{ik}' \mathbb{1}_{\{|\mathbf{a}_{ik}' \mathbf{x}_k^*|^2 \leq \frac{9\mu^2\kappa^2}{mq} \sum_{ik} |\mathbf{a}_{ik}' \mathbf{x}_k^*|^2\}} \\ \mathbf{Y}_-(\epsilon_1) &= \\ &\frac{1}{mq} \sum_{ik} |\mathbf{a}_{ik}' \mathbf{x}_k^*|^2 \mathbf{a}_{ik} \mathbf{a}_{ik}' \mathbb{1}_{\{|\mathbf{a}_{ik}' \mathbf{x}_k^*|^2 \leq \frac{9\mu^2\kappa^2(1-\epsilon_1)}{q} \|\mathbf{X}^*\|_F^2\}}. \end{aligned}$$

Define  $\mathbf{Y}_+(\epsilon_1)$  similarly but with  $(1 - \epsilon_1)$  replaced by  $(1 + \epsilon_1)$  in the indicator function. We will show that  $\mathbf{Y}_U$  is sandwiched between  $\mathbf{Y}_-$  and  $\mathbf{Y}_+$ . This, along with showing that  $\mathbf{Y}_-$  and  $\mathbf{Y}_+$  are close, will help us show that  $\mathbf{Y}_U$  is close to  $\mathbf{Y}_-$  and, hence, also to its expected value. After this, use of the  $\sin \theta$  theorem will give us the desired bound.

Adapting the approach of [9], [20],

$$\mathbb{E}[\mathbf{Y}_-(\epsilon_1)] = \frac{1}{q} \left\{ \sum_k \beta_{1,k}^- \mathbf{x}_k^* \mathbf{x}_k^{*'} + \left( \sum_k \beta_{2,k}^- \|\mathbf{x}_k^*\|^2 \right) \mathbf{I} \right\} \quad (9)$$

where

$$\begin{aligned} \beta_{1,k}^-(\epsilon_1) &:= \mathbb{E}[(\xi^4 - \xi^2) \mathbb{1}_{\xi^2 \leq (1-\epsilon_1)\gamma_k}] \\ \beta_{2,k}^-(\epsilon_1) &:= \mathbb{E}[\xi^2 \mathbb{1}_{\xi^2 \leq (1-\epsilon_1)\gamma_k}], \\ \gamma_k &:= \frac{9\mu^2\kappa^2 \|\mathbf{X}^*\|_F^2}{q \|\mathbf{x}_k^*\|^2}, \end{aligned}$$

and  $\xi$  is a scalar standard Gaussian random variable. The expression for  $\mathbb{E}[\mathbf{Y}_+(\epsilon_1)]$  is similar but with  $(1 - \epsilon_1)$  replaced by  $(1 + \epsilon_1)$  in the expression for  $\beta_{1,k}^+$ ,  $\beta_{2,k}^+$ .

Observe that  $\mathbb{E}[\mathbf{Y}_-(\epsilon_1)]$  can be simplified as

$$\mathbb{E}[\mathbf{Y}_-(\epsilon_1)] = \frac{1}{q} [\mathbf{U}^* \left( \sum_k \beta_{1,k}^- \tilde{\mathbf{b}}_k^* \tilde{\mathbf{b}}_k^{*'} \right) \mathbf{U}^{*'} + \left( \sum_k \beta_{2,k}^- \|\tilde{\mathbf{b}}_k^*\|^2 \right) \mathbf{I}]$$

Thus, the span of its top  $r$  eigenvectors (same as singular vectors) equals  $\text{Span}(\mathbf{U}^*)$ . Hence, we can use the  $\sin \Theta$  theorem [53] stated below in a fashion similar to [20] (Sec 6).

**Lemma 3.5** (Davis-Kahan  $\sin \Theta$  theorem). *Given two symmetric matrices  $\mathbf{D}$  and  $\hat{\mathbf{D}}$ . Let  $\mathbf{U}^*$  ( $\mathbf{U}$ ) be the matrix of top eigenvectors of  $\mathbf{D}$  ( $\hat{\mathbf{D}}$ ). If  $\lambda_r(\mathbf{D}) - \lambda_{r+1}(\mathbf{D}) - \|\mathbf{D} - \hat{\mathbf{D}}\| > 0$ , then*

$$\sin \Theta(\mathbf{U}, \mathbf{U}^*) \leq \frac{\|\mathbf{D} - \hat{\mathbf{D}}\|}{\lambda_r(\mathbf{D}) - \lambda_{r+1}(\mathbf{D}) - \|\mathbf{D} - \hat{\mathbf{D}}\|}.$$

Using Lemma 3.5 with  $\hat{\mathbf{D}} = \mathbf{Y}_U$  and  $\mathbf{D} = \mathbb{E}[\mathbf{Y}_-]$ ,

$$\begin{aligned} \sin \Theta(\mathbf{U}_{\text{init}}, \mathbf{U}^*) &\leq \\ &\leq \frac{\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-(\epsilon_1)]\|}{\lambda_r(\mathbb{E}[\mathbf{Y}_-(\epsilon_1)]) - \lambda_{r+1}(\mathbb{E}[\mathbf{Y}_-(\epsilon_1)]) - \|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-(\epsilon_1)]\|} \end{aligned}$$

Moreover,

$$\begin{aligned}\lambda_r(\mathbb{E}[\mathbf{Y}_-]) - \lambda_{r+1}(\mathbb{E}[\mathbf{Y}_-]) &= \frac{1}{q} \lambda_{\min} \left( \sum_k \beta_{1,k}^- \tilde{\mathbf{b}}_k^* \tilde{\mathbf{b}}_k^{*\prime} \right) \\ &\geq \left( \min_k \beta_{1,k}^- \right) \frac{\sigma_{\min}^{*2}}{q}\end{aligned}$$

Now we just need to upper bound  $\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-(\epsilon_1)]\|$  and lower bound  $\min_k \beta_{1,k}^-$ . Both these follow by combining the three lemmas given next and triangle inequality.

**Lemma 3.6.** *We have that, w.p. at least  $1 - \exp(-\epsilon_1^2 \frac{mq}{\mu^2 \kappa^2})$ ,*

$$\mathbf{Y}_-(\epsilon_1) \preceq \mathbf{Y}_U \preceq \mathbf{Y}_+(\epsilon_1)$$

and so  $\|\mathbf{Y}_U - \mathbf{Y}_-(\epsilon_1)\| \leq \|\mathbf{Y}_+(\epsilon_1) - \mathbf{Y}_-(\epsilon_1)\|$ .

**Lemma 3.7.** *Let  $\mathbf{Y}_+ = \mathbf{Y}_+(\epsilon_1)$  and  $\mathbf{Y}_- = \mathbf{Y}_-(\epsilon_1)$ . We have*

$$\|\mathbb{E}[\mathbf{Y}_+] - \mathbb{E}[\mathbf{Y}_-]\| \leq \frac{9\epsilon_1 \mu^2 \kappa^2 \|\mathbf{X}^*\|_F^2}{q} \leq \frac{9\epsilon_1 \mu^2 \kappa^2 r \sigma_{\max}^{*2}}{q}$$

and, assuming  $\epsilon_1 < 0.01$ ,

$$\min_k \beta_{1,k}^-(\epsilon_1) \geq 1.5.$$

**Lemma 3.8.** *We have that,*

w.p. at least  $1 - 2 \exp(n \log 9 - c\epsilon_2^2 mq)$ ,

$$\|\mathbf{Y}_-(\epsilon_1) - \mathbb{E}[\mathbf{Y}_-(\epsilon_1)]\| \leq \frac{1.5\epsilon_2 \mu^2 \kappa^2 r \sigma_{\max}^{*2}}{q}$$

We get the exact same claim also for  $\|\mathbf{Y}_+ - \mathbb{E}[\mathbf{Y}_+]\|$ .

We prove the above lemmas in Sec. A-B.

Using triangle inequality,  $\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \leq \|\mathbf{Y}_U - \mathbf{Y}_-\| + \|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\|$ . Moreover, using Lemma 3.6,  $\|\mathbf{Y}_U - \mathbf{Y}_-\| \leq \|\mathbf{Y}_+ - \mathbf{Y}_-\|$ . Using these and again using triangle inequality,  $\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \leq 2\|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\| + \|\mathbf{Y}_+ - \mathbb{E}[\mathbf{Y}_+]\| + \|\mathbb{E}[\mathbf{Y}_+] - \mathbb{E}[\mathbf{Y}_-]\|$ . Thus, combining bounds from the above lemmas and setting  $\epsilon_1 = \epsilon_2 = \frac{\delta_{\text{init}}}{C(\kappa^2 \mu^2) \kappa^2 r}$  for a  $\delta_{\text{init}} < 1$ , we conclude that w.p.  $1 - 2 \exp(n - \frac{c\delta_{\text{init}}^2 mq}{\kappa^8 \mu^4 r^2}) - 2 \exp(-\frac{c\delta_{\text{init}}^2 mq}{\kappa^8 \mu^4 r^2})$ ,

$$\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \leq \frac{0.25\delta_{\text{init}}\sigma_{\min}^{*2}}{q}. \quad (10)$$

Using  $\min_k \beta_{1,k}^- \geq 1.5$ , from Lemma 3.7, and (10), since  $\delta_{\text{init}} < 1$ ,

$$\begin{aligned}\sin \Theta(\mathbf{U}_{\text{init}}, \mathbf{U}^*) &\leq \frac{\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\|}{\lambda_r(\mathbb{E}[\mathbf{Y}_-]) - \lambda_{r+1}(\mathbb{E}[\mathbf{Y}_-]) - \|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\|} \\ &\leq \frac{0.25\delta_{\text{init}}\sigma_{\min}^{*2}/q}{1.5\sigma_{\min}^{*2}/q - 0.25\delta_{\text{init}}\sigma_{\min}^{*2}/q} < \delta_{\text{init}}.\end{aligned}$$

1) *Proof that rank is correctly estimated:* Consider the rank estimation step. This requires lower bounding  $\lambda_r(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U)$  and upper bounding  $\lambda_{r+1}(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U)$ . Both bounds follow using (i) (10), along with Weyl's inequality, and (ii) the lower bound on  $\beta^- := \min_k \beta_{1,k}^-$  from Lemma 3.7:  $\beta^- > 1.5$ .

We have  $\lambda_r(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U) \geq \lambda_r(\mathbb{E}[\mathbf{Y}_-]) - \lambda_n(\mathbb{E}[\mathbf{Y}_-]) - 2\|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\| \geq \beta^- \lambda_r(\mathbf{U}^* \Sigma^{*2} \mathbf{U}^*) - 2\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \geq$

$1.5(\sigma_{\min}^{*2}/q) - 0.25\delta_{\text{init}}\sigma_{\min}^{*2}/q = (1.5 - 0.25\delta_{\text{init}})\sigma_{\min}^{*2}/q > 1.4(\sigma_{\min}^{*2}/q)$  as long as  $\delta_{\text{init}} < 0.1$ .

Also,  $\lambda_{r+1}(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U) \leq \lambda_{r+1}(\mathbb{E}[\mathbf{Y}_-]) - \lambda_n(\mathbb{E}[\mathbf{Y}_-]) + 2\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| = 2\|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\| \leq 0.25\delta_{\text{init}}(\sigma_{\min}^{*2}/q) < 0.025(\sigma_{\min}^{*2}/q)$  as long as  $\delta_{\text{init}} < 0.1$ .

In summary, as long as  $\delta_{\text{init}} < 0.1$ ,  $\lambda_r(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U) \geq 1.4\sigma_{\min}^{*2}/q$  and  $\lambda_{r+1}(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U) \leq 0.025\sigma_{\min}^{*2}/q$ . Thus by setting the threshold  $\omega = C\sigma_{\min}^{*2}/q$  with  $C$  being any constant between 0.025 and 1.4, we can ensure that the rank is correctly estimated whp.

### C. Proof of Claim 3.4

In this section, we remove the superscript  $t$  except where essential. Also, we let  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(T+t)}$  and  $\mathbf{y}_{ik} := \mathbf{y}_{ik}^{(T+t)}$ .

We first use the overall approach of [25] to get the following deterministic bound on the subspace error of the  $(t+1)$ -th estimate of  $\mathbf{U}^*$ ,  $\mathbf{U}^{t+1}$ . The proof requires some messy algebra and hence we give it in Appendix B.

**Lemma 3.9.** *We have*

$$\sin \Theta(\mathbf{U}^{t+1}, \mathbf{U}^*) \leq \frac{\text{MainTerm}}{\sigma_{\min}(\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}') - \text{MainTerm}} \quad (11)$$

where  $\text{MainTerm} :=$

$$\frac{\max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term1}(\mathbf{W})| + \max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term2}(\mathbf{W})|}{\min_{\mathbf{W} \in \mathcal{S}_W} \text{Term3}(\mathbf{W})},$$

$$\text{Term1}(\mathbf{W}) := \sum_{ik} \mathbf{b}_k' \mathbf{W}' \mathbf{a}_{ik} \mathbf{a}_{ik}' \mathbf{U}^* (\tilde{\mathbf{B}}^* \mathbf{B}' \mathbf{b}_k - \tilde{\mathbf{b}}_k^*),$$

$$\text{Term2}(\mathbf{W}) := \sum_{ik} (\mathbf{c}_{ik} \hat{\mathbf{c}}_{ik} - 1) (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k) (\mathbf{a}_{ik}' \mathbf{x}_k^*),$$

$$\text{Term3}(\mathbf{W}) := \sum_{ik} (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k)^2,$$

$$\mathcal{S}_W := \{\mathbf{W} \in \mathbb{R}^{n \times r} : \|\mathbf{W}\|_F = 1\}$$

is the space of all  $n \times r$  matrices with unit Frobenius norm, and  $\mathbf{c}_{ik}, \hat{\mathbf{c}}_{ik}$  are the phases (signs) of  $\mathbf{a}_{ik}' \mathbf{x}_k^*$  and  $\mathbf{a}_{ik}' \hat{\mathbf{x}}_k$ .

We obtain high probability bounds on the three terms above in the three lemmas that follow, Lemmas 3.10, 3.11, 3.12. All three lemmas first bound the terms for a fixed  $\mathbf{W}$ , followed by using a carefully developed epsilon-net argument to extend the bounds for all unit Frobenius norm  $\mathbf{W}$ 's. This is inspired by similar arguments in [52].

Consider a fixed  $\mathbf{W}$ . To bound Term1, we first show  $\mathbb{E}[\text{Term1}] = 0$ . Next, we use Lemma 3.2 to show that  $\|\tilde{\mathbf{B}}^* (\mathbf{B}' \mathbf{B} - \mathbf{I})\|_F \leq C\delta_t \|\mathbf{X}^*\|_F$ . Finally, we use these two facts and a simple modification of Lemma 5.16 of [54] for sums of products of sub-Gaussian random variables (Lemma A.1), along with careful linear algebra to show that, if  $mq$  is large enough, whp,  $|\text{Term1}| \leq Cm\delta_t^2 \|\mathbf{X}^*\|_F$  for any  $\delta_t < 0.1$ . This is followed by a careful epsilon-net argument to extend the bound for all unit Frobenius norm  $\mathbf{W}$ 's.

To bound Term2 for a fixed  $\mathbf{W}$ , we first use Cauchy-Schwarz. This implies that

$$|\text{Term2}(\mathbf{W})| \leq \sqrt{\text{Term3}(\mathbf{W})} \sqrt{\text{Term22}}, \text{ where}$$

$$\text{Term22} := \sum_{ik} (\mathbf{c}_{ik} \hat{\mathbf{c}}_{ik} - 1)^2 (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2$$

We explain how to upper bound Term3( $\mathbf{W}$ ) in the next paragraph. Consider Term22. Notice that  $(c_{ik}\hat{c}_{ik} - 1)^2$  takes only two values - zero or four. It is zero when the signs are equal, else it is four. To start bounding  $\mathbb{E}[\text{Term22}]$ , we can use Lemma 1 of [10]. This shows that the probability that the signs are unequal is upper bounded by a term that is directly proportional to the ratio  $\text{dist}^2(\mathbf{x}_k^*, \hat{\mathbf{x}}_k) / (\mathbf{a}_{ik}'\mathbf{x}_k^*)^2$ . The probability bound is thus large when this ratio is large and small otherwise. Moreover, it is easy to see that  $\mathbb{E}[(\mathbf{a}_{ik}'\mathbf{x}_k^*)^2] = \|\mathbf{x}_k^*\|^2$  and by Lemma 3.2, whp,  $\text{dist}^2(\mathbf{x}_k^*, \hat{\mathbf{x}}_k) \leq \delta_t^2 \|\mathbf{x}_k^*\|^2$ . Relying on these ideas, we can argue that, on average, Term22 is very small:  $\mathbb{E}[\text{Term22}] \leq Cm\delta_t^3 \|\mathbf{X}^*\|_F^2$ . Careful use of concentration bounds then implies that, if  $mq$  is large enough, the same order bound holds whp.

To upper and lower bound Term3, notice first that  $\mathbb{E}[\text{Term3}] = m\|\mathbf{W}\mathbf{B}\|_F^2 = m$ . Also, each summand in this term is sub-exponential with sub-exponential norm bounded by  $\|\mathbf{W}\mathbf{b}_k\|^2$ ; and  $\|\mathbf{W}\mathbf{b}_k\|^2 \leq \|\mathbf{b}_k\|^2 \leq \hat{\mu}^2 r/q$  (by Lemma 3.3), and  $\sum_{ik} \|\mathbf{W}\mathbf{b}_k\|^2 = m$ . Using these facts and Lemma 3.2 (Bernstein-like inequality similar to Lemma 5.16 of [54]), we can show that Term3 concentrates around  $m$  whp.

**Lemma 3.10.** *Pick a  $\delta_b < 1/10$  and assume that  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . Under the conditions of Theorem 2.1, for a  $\delta_t < 1/10$ , w.p. at least  $1 - 2 \exp\left(nr(\log 17) - c\frac{\delta_t^2 mq}{\hat{\mu}^2 r}\right) - n^{-10}$ ,*

$$\min_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term3}(\mathbf{W}) \geq 0.5(1 - \delta_t)m$$

and

$$\max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term3}(\mathbf{W}) \leq 1.5(1 + \delta_t)m.$$

**Lemma 3.11.** *Pick a  $\delta_b < 1/10$  and assume that  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . Under the conditions of Theorem 2.1 and assuming that  $\sin \Theta(\mathbf{U}^*, \mathbf{U}) \leq \delta_t$ , with  $\delta_t < 1/10$ , w.p. at least  $1 - 2 \exp\left(nr(\log 17) - c\frac{\delta_t^2 mq}{\kappa^3 \mu^2 r}\right) - n^{-10}$ .*

$$\max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term1}(\mathbf{W}) \leq m\delta_t^2 \|\mathbf{X}^*\|_F$$

In proving the above, we also show that

$$\|\tilde{\mathbf{B}}^* (\mathbf{I} - \mathbf{B}'\mathbf{B})\|_F \leq C\delta_t \|\mathbf{X}^*\|_F$$

(we will use this in proving Lemma 3.13).

**Lemma 3.12.** *Pick a  $\delta_b < 1/10$  and assume that  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . Under the conditions of Theorem 2.1 and assuming  $\sin \Theta(\mathbf{U}^*, \mathbf{U}) \leq \delta_t$  with  $\delta_t < 1/10$ , w.p. at least  $1 - 2 \exp\left(nr(\log 17) - c\frac{\delta_t^2 mq}{\hat{\mu}^2 r}\right) - 2 \exp(-c\delta_t^2 mq) - n^{-10}$ ,*

$$\max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term2}(\mathbf{W}) \leq m \sqrt{1 + \delta_t} \sqrt{\delta_t} \delta_t \|\mathbf{X}^*\|_F.$$

Finally, we lower bound the first denominator term of (11). This can be done by using the bound on  $\|\tilde{\mathbf{B}}^* (\mathbf{I} - \mathbf{B}'\mathbf{B})\|_F$  from Lemma 3.11. This, in turn, implies a lower bound on the minimum singular value of  $\mathbf{B}'\mathbf{B}$ , and hence the following.

**Lemma 3.13.** *Pick a  $\delta_b < 1/10$  and assume that  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . Under the conditions of Theorem*

2.1, if  $\sin \Theta(\mathbf{U}^*, \mathbf{U}) \leq \delta_t$  with  $\delta_t \leq \frac{1}{4C\sqrt{r}\kappa}$ , then, w.p. at least  $1 - n^{-10}$ ,  $\sigma_{\min}(\mathbf{U}^* \mathbf{\Sigma}^* \mathbf{B}' \mathbf{B}') \geq 0.9\sigma_{\min}^*$ .

We prove the above lemmas in Sec. A-E.

*Proof of Claim 3.4.* Combining Lemmas 3.9 and 3.13, if  $\delta_t < c/\sqrt{r}\kappa$ ,

$$\sin \Theta(\mathbf{U}^{t+1}, \mathbf{U}^*) \leq \frac{\text{MainTerm}}{0.9\sigma_{\min}^* - \text{MainTerm}}. \quad (12)$$

Set  $\delta_b = 1/11$ . Combining Lemmas 3.11, 3.12 and 3.10, and using  $\|\mathbf{X}^*\|_F \leq \sqrt{r}\sigma_{\max}^*$ , we conclude that,

$$\text{MainTerm} \leq C(\delta_t + \sqrt{\delta_t})\delta_t \sqrt{r}\sigma_{\max}^*.$$

In the above bound, the  $\sqrt{\delta_t}$  term dominates. In order to ensure that  $\text{MainTerm} \leq 0.7\delta_t\sigma_{\min}^*$ , we need to set  $\sqrt{\delta_t} = c/\sqrt{r}\kappa$ . Doing this and using (12),

$$\sin \Theta(\mathbf{U}^{t+1}, \mathbf{U}^*) \leq \frac{C(\delta_t + \sqrt{\delta_t})\delta_t \sqrt{r}\sigma_{\max}^*}{0.9\sigma_{\min}^* - C(\delta_t + \sqrt{\delta_t})\delta_t \sqrt{r}\sigma_{\max}^*} \leq 0.7\delta_t$$

□

#### IV. NUMERICAL EVALUATION

In this section we provide detailed description of the numerical evaluation of our algorithms on synthetic and real data. All time comparisons are performed on a single Desktop Computer with Intel® Xeon E3-1240 8-core CPU @ 3.50GHz and 32GB RAM. We must mention that for all the algorithms, (a) we compare on the same system, and we do not run other memory and compute intensive programs while performing the time comparison, (b) we use the most efficient subroutines that are provided by the respective authors and these implementations rely on the “anonymous functions” feature of MATLAB. This provides a significant speed up since the major chunk of computation time involves matrix-vector/matrix-matrix products. This ensures uniformity to a large extent for all algorithms.

##### A. Synthetic Data Experiments

We demonstrate the effectiveness of AltMinLowRaP over existing work using two synthetic experiments. For both, we generate an error versus time-taken plot as follows: for each  $t = 0, 1, \dots, T$ , we plot the matrix recovery error (at the end of that iteration) and time-taken (until the end of that iteration) on the y- and x-axes respectively. All synthetic data experiments are performed for 100 independent trials, and for each algorithm, we plot the average error over the best 90 trials (drop the 10 trials with the largest error). This is done because all algorithms are guaranteed to work well only with high probability.

In the *first* experiment, we compare the performance of several algorithms for different values of  $m$ . We generated data as follows:  $\mathbf{X}^* = \mathbf{U}^* \tilde{\mathbf{B}}^*$  where  $\mathbf{U}^* \in \mathbb{R}^{n \times r}$  is generated by orthonormalizing a iid standard Gaussian matrix. The entries of  $\tilde{\mathbf{B}}^* \in \mathbb{R}^{r \times q}$  are chosen from another iid standard Gaussian distribution. Thus, in this setting,  $\sigma_{\min}^* / q \approx \sigma_{\max}^* / q \approx 1$ . Measurements were generated using (1) with  $\mathbf{a}_{ik}$ 's being iid

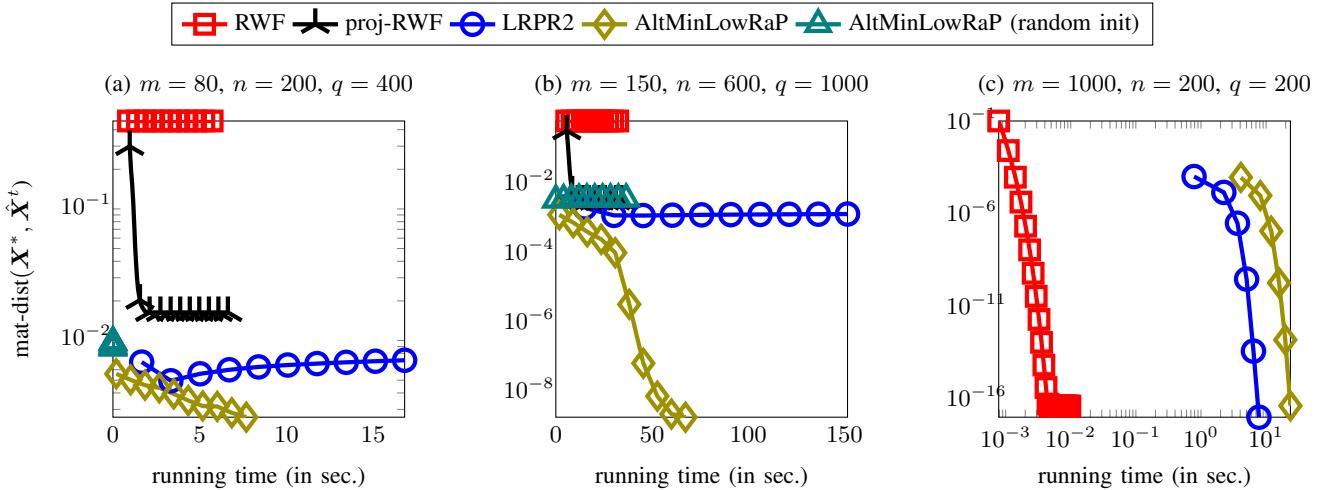


Fig. 3: Error versus time plot with time in seconds. We compare with LRPR2 which is the only other existing Low-Rank Phase Retrieval algorithm [20], RWF [10] and projected RWF. The first step (initialization) of LRPR2 is slower than our proposed method. This is likely due to the fact that the estimates of the rank are different for the two algorithms and thus the errors are also slightly different in the two cases. For the purpose of better illustration, we only plot the error and time at the end of every 10 iterations for RWF and proj-RWF.

Gaussian. We compare AltMinLowRaP with LRPR2 (best heuristic from [20]), RWF [10], and with what we call projected-RWF or proj-RWF (for all  $t$ , after the  $t$ -th RWF iteration, we project the matrix  $\hat{X}_{RWF}^t$  onto the space of rank- $r$  matrices, with  $r$  known). We provide results for three cases: (i)  $n = 200$ ,  $q = 400$ ,  $r = 4$ , and  $m = 0.4n$ ; (ii)  $n = 600$ ,  $q = 1000$ ,  $r = 4$ , and  $m = 0.25n$ ; and (iii)  $n = 200$ ,  $m = 5n = 1000$ ,  $q = 200$  and  $r = 4$ . The results are summarized in Fig. 3. In the first two settings (Figs. 3(a),(b)), we show that by exploiting the low-rank structure, AltMinLowRaP is able to outperform the unstructured PR methods. Notice that, proj-RWF performs better than RWF (which does not work at all since  $m \ll n$ ). Secondly, AltMinLowRaP is significantly better than proj-RWF, and for reasons explained earlier, it is also better than LRPR2. AltMinLowRaP and LRPR2 estimate the rank, whereas proj-RWF is provided the true rank. We observed that, in practice, the estimated rank is higher than the true rank in all the cases. In both these settings,  $m \ll n$ , and thus, unstructured algorithms fail. The third setting is a setting with a very large value of  $m$ : we are using  $m = 5n$  measurements. This is a case where all of AltMinLowRaP, TWF, and RWF work, but AltMinLowRaP is much slower as shown in Fig. 3(c). Additionally, for the first two cases, based on a reviewer's comment, we also empirically evaluate a random initialization scheme for the AltMinLowRaP algorithm. We observed that although the initial estimates are approximately in the order of  $10^{-2}$ , the algorithm itself fails to improve this with subsequent iterations. A possible workaround to this could be through using a gradient descent based algorithm (inspired by [12]) but this requires a detailed analysis, and we will study this as part of future work..

Our second experiment illustrates the time complexity discussion given in Sec. II-B. For a given  $m$ , AltMinLowRaP is about  $r$  times slower than the best provably correct regular (unstructured) PR methods - TWF and RWF. But, if for each algorithm, we use the minimum  $m$  needed for the algorithm

to achieve  $\epsilon$  accuracy, then, theoretically, AltMinLowRaP should be faster if the rank  $r$  is small enough. We tested this empirically as follows. We generated data as before with  $n = 600$ ,  $q = 1000$  and  $r = 4$ . We implemented TWF and RWF using two values of  $m$ ,  $m = 3n$  and  $m = 4n$ . We evaluated AltMinLowRaP with using  $m = n/4$ . The error-at-iteration- $t$  versus time-taken-until-iteration- $t$  plot is shown in Fig. 2 for all these cases. As can be seen, using  $3n$  measurements, neither of TWF or RWF works. Using  $m = 4n$ , both work. But if we compare the time taken (x-axis value) for any value of error level  $\epsilon$ , both are at least 5 times slower than AltMinLowRaP ( $m = n/4$ ). For all algorithms, we repeat the experiments for 100 independent trials, and plot the mean taken over the best 90 trials to illustrate the high probability results.

1) *Algorithm parameters:* For both experiments, AltMinLowRaP was implemented as Algorithm 1 but using the same set of measurements (does not require sample-splitting), and with the following parameters:  $T_{RWF,t}$  scales linearly from 5 to 30,  $\omega = 1.3\sigma_{\min}^{*2}/q \approx 1.3$  and  $C_Y = 9$ . All parameters are as suggested in the theorem. For LRPR2 we used the default parameters mentioned in the documentation. We set the maximum number of outer-loop iterations,  $T_{\max} = 10$  for both. For RWF and TWF, we used the default parameters suggested by the authors with the exception that we let the maximum number of iterations  $T_{\max} = 300$  (to try to see if its error reduces with more iterations). However, as we plot the time-taken at the end of each iteration, this is not an unfair implementation of RWF; it only means that we have 300 data points to plot on our graph. We should point out that the recovery error for AltMinLowRaP is sensitive to the choice of  $\omega$ . If  $\omega$  is too small (for given values of  $m$  and of  $C_Y$  used in computing  $\mathbf{Y}_U$ ), the algorithm will significantly overestimate the rank. This is especially problematic when  $m$  is small (or  $C_Y$  is large for a given  $m$ ). Thus, as a thumb rule, for a lower value of  $m/n$ , the threshold  $\omega$  should be larger.

Of course if it is too large, it will underestimate the rank<sup>6</sup>

### B. Real Videos with Simulated CDP measurements: Small $r$ suffices

We demonstrate the effectiveness of AltMinLowRaP for recovering two real video sequences (these are only approximately low-rank) from simulated Coded Diffraction Pattern (CDP) measurements. These measurements can be represented as  $\mathbf{Y} = |\mathcal{F}(\mathbf{DX}^*)|$  where  $\mathcal{F}(\cdot)$  is the DFT operation and the matrix  $\mathbf{D}$  represents a diagonal mask matrix whose diagonal entries are chosen uniformly at random from  $\{\pm 1, \pm \sqrt{-1}\}$  and modulate the intensity of the input. We generate CDP measurements of each frame of the video (the  $k$ -th frame vectorized is  $\mathbf{x}_k^*$ ). We compared our algorithm with LRPR2 and RWF. We present the quantitative results in Table II and the visual comparisons in Fig. 1 (given in the beginning), and in Fig. 4. Notice that, in this case, even with  $m = 5n$  measurements, RWF is unable to accurately recover the video and AltMinLowRaP has a slightly better performance w.r.t. LRPR2. The algorithm parameters are set as in the synthetic data experiments, with the exception that we now set  $T_{max} = 30$  for AltMinLowRaP and LRPR2. AltMinLowRaP implementation used all the speed-up ideas for Fourier measurements explained in [20] for LRPR2 and so did LRPR2.

We tested AltMinLowRaP with three possible values of rank,  $r = 15, 20, 25$ . As can be seen, even  $r = 15$  suffices to get a significantly better reconstruction error than RWF for  $m = 5n$  CDP measurements. For the plane video,  $n = 6912$  and  $q = 105$  and for the mouse video,  $n = 5182$  and  $q = 90$ , and thus in both cases,  $r \approx 0.003n$ . If  $q$  is larger, a natural idea would be to use similar parameter settings, but instead implement the tracking variant of AltMinLowRaP (Algorithm 2).

### C. Real Videos with Simulated CDP measurements: Low-Rank versus (Wavelet) Sparse Models

To justify the low-rank assumption on videos, we compare with CoPRAM [18], a state-of-the-art, provable algorithm for compressive phase retrieval. Since the videos are not sparse in the spatial domain, as suggested in [18], we use the Haar wavelet as the sparsifying basis<sup>7</sup>.

As can be seen from Table III and Fig. 5, the low-rank prior gives a much better reconstruction error in all three cases in this table including the exact sparse case. Since the video is not exactly wavelet sparse, we also performed a comparison on the sparsified video, wherein, for each image frame, we truncate the wavelet coefficients such that approximately 90% of the energy in each frame is preserved. We refer to this as the *sparsified video* in this experiment. Sparsifying the

<sup>6</sup>This is a bigger problem when  $\kappa = 1$  as in the simulated data above. It is a lesser problem for real approximately low-rank data (e.g., slow changing videos) with a larger  $\kappa$ , since in those cases, the missed directions will be the ones with smaller singular values.

<sup>7</sup>We also experimented with the Daubhechies-3 wavelet as the sparsifying basis in our experiments. However, we noticed that, for the plane video, irrespective of the wavelet basis, the number of coefficients necessary to preserve  $\approx 90\%$  of the energy of the video required  $s \approx 0.1n$  and thus the choice of wavelet basis is not detrimental in this experiment.

TABLE II:  $\text{mat-dist}(\hat{\mathbf{X}}, \mathbf{X}^*)$  and time comparison for the *mouse* and *plane* videos. We generate the measurements using the CDP model and consider two different number of settings. Notice that AltMinLowRaP is slightly better than LRPR2 but is slower than in the simulated data experiments.

Algorithm	$\text{mat-dist}(\hat{\mathbf{X}}, \mathbf{X}^*)$ (Running Time in sec)		
	$m = 5n$ (mouse)	$m = 5n$ (plane)	$m = 2n$ (mouse)
RWF [10]	0.65 (0.35s)	0.65 (0.35s)	1.36 (0.25s)
LRPR2 [20] ( $r = 25$ )	0.48 (81.8s)	0.10 (122.3s)	0.61 (31.0s)
AltMinLowRaP ( $r = 25$ )	0.39 (297.6s)	0.09 (467.6s)	0.52 (122.1s)
AltMinLowRaP ( $r = 15$ )	0.57 (277.2s)	0.15 (418.9s)	0.60 (113.6s)
AltMinLowRaP ( $r = 10$ )	0.70 (262.7s)	0.23 (409.0s)	0.72 (97.4s)

TABLE III: Low rank versus Sparse PR: We compare with a recent state of the art algorithm for provable sparse phase retrieval.  $\text{mat-dist}(\hat{\mathbf{X}}, \mathbf{X}^*)$  and time comparison for the *mouse* and *plane* videos using *real valued* CDP measurements. We conclude that low-rank is a better model than wavelet-sparse for slowly changing videos.

Algorithm	Video		
	$m = 5n$ (plane)	$m = 10n$ (plane)	$m = 5n$ (sparsified-plane)
CoPRAM [18]	2.113	1.019	0.3104
RWF	0.6531	0.4134	0.6514
AltMinLowRaP ( $r = 15$ )	0.111	0.109	0.1427

video significantly improves the performance of CoPRAM, but AltMinLowRaP ( $r = 15$ ) is still better. Even with  $m = 10n$  measurements, the sparse model is unable to capture the finer details in the video. We also observed that a standard complex mask does not work very well for CoPRAM and hence for this experiment, we report the results when the entries of the CDP mask,  $\mathbf{D}$  are chosen uniformly at random from  $\{\pm 1\}$ . We reshaped each video frame into size  $32 \times 32$  since the online implementation of their code only works for small sized data. We provide the quantitative results in Table III and qualitative results for the plane video in Fig. 5.

## V. PHASELESS SUBSPACE TRACKING

When the matrix  $\mathbf{X}^*$  consists of a time sequence of signals  $\mathbf{x}_k^*$ , then the column-wise measurements appear one column at a time (sequentially). Hence, there is benefit in trying to develop a mini-batch algorithm that works with measurements of short batches of  $\alpha$  consecutive columns. Moreover, for long data sequences, the subspace from which the data are generated could itself change with time. Detecting and being able to track such subspace changes is important for long sequences. Interestingly the algorithm that works for this purpose is a simple modification of the static case idea along with a carefully designed subspace change detection step.

### A. Problem setting

The low-rank assumption is equivalent to assuming that  $\mathbf{x}_k^* = \mathbf{U}^* \tilde{\mathbf{b}}_k^*$  where  $\mathbf{U}^*$  specifies a fixed  $r$ -dimensional subspace. For long signal/image sequences, a better model (one that allows the required subspace dimension  $r$  to be smaller) is to let the subspace change with time. As is common in time-series analysis, the simplest model for time-varying quantities is to assume that they are piecewise constant with time. We

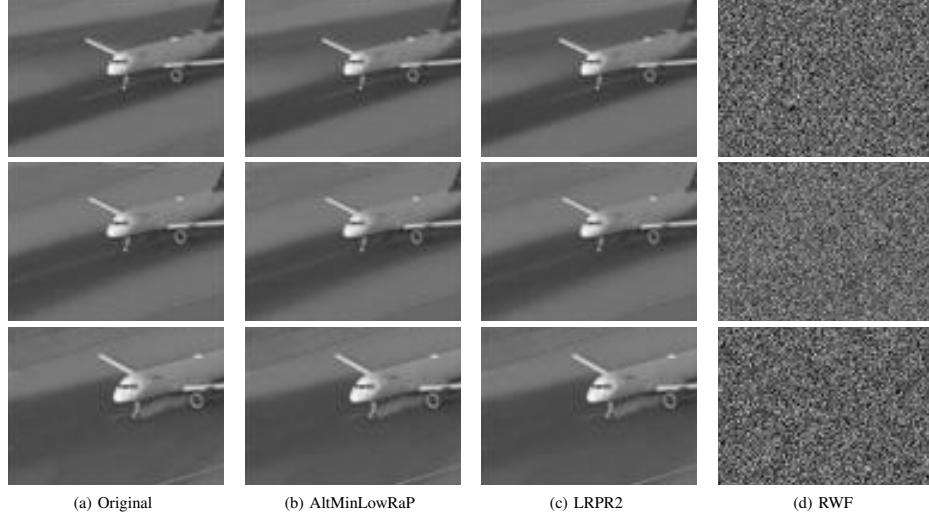


Fig. 4: Comparison of visual performance for the plane video with  $m = 2n$ . The images are shown at  $k = 20, 60, 100$ .

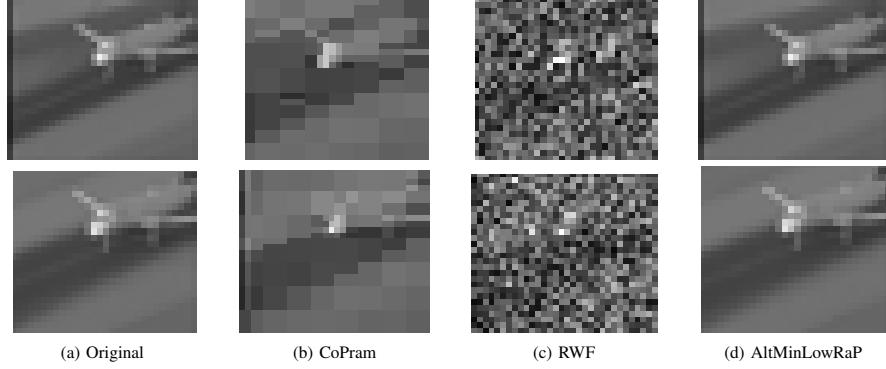


Fig. 5: Comparison of visual performance for the plane video for  $m = 5n$  at  $k = 10, 60$ . We observe that RWF fails visually although numerical error is better, and CoPram is able to recover the global features but cannot recover the fine details due to the high level of sparsity required. AltMinLowRaP outperforms both methods, confirming the validity of the low-rank model.

adopt this approach here. Moreover, in order to easily borrow ideas from the static setting, we will assume that we now have a total of  $q_{\text{full}}$  signals (matrix columns) and we will denote the  $n \times q_{\text{full}}$  matrix formed by all these columns by  $\mathbf{X}^*_{\text{full}}$ . Our algorithm will operate on measurements of  $\alpha$ -consecutive column sub-matrices of  $\mathbf{X}^*_{\text{full}}$ .

Let  $k_0 = 1$ , and let  $k_j$  denote the  $j$ -th subspace change time, for  $j = 1, 2, \dots, J$  and let  $k_{J+1} = q_{\text{full}}$ . We have the following model

$$\mathbf{x}_k^* = \mathbf{U}_{\text{sub},(j)}^* \tilde{\mathbf{d}}_k^*, \text{ for all } k_j \leq k \leq k_{j+1} \quad (13)$$

where  $\mathbf{U}_{\text{sub},(j)}^*$  is an  $n \times r$  “basis matrix” for the  $j$ -th subspace and  $\tilde{\mathbf{d}}_k^*$  is the coefficients’ vector at time  $k$ .

The goal is to track the subspaces  $\text{Span}(\mathbf{U}_{\text{sub},(j)}^*)$  on-the-fly; of course, “on-the-fly” for subspace tracking means with a delay of at least  $r$ . Once this can be done accurately enough, it is easy to also recover the matrix columns  $\mathbf{x}_k^*$  (by solving a simple  $r$ -dimensional PR problem to recover the  $\tilde{\mathbf{d}}_k^*$ ’s).

The reason we use a different notation here (the subscript sub and use of  $\tilde{\mathbf{d}}_k^*$  instead of  $\tilde{\mathbf{b}}_k^*$ ) is as follows. Consider an

$\alpha$ -column sub-matrix formed by  $\alpha$  consecutive signals. Let us call it  $\mathbf{X}^*$  and let  $\mathbf{X}^* \stackrel{\text{SVD}}{=} \mathbf{U}^* \Sigma^* \mathbf{B}^*$ . If all the  $\mathbf{x}_k^*$ ’s forming this matrix are generated from the same subspace, say  $\mathbf{U}_{\text{sub},(j)}^*$ , then  $\text{Span}(\mathbf{U}^*) = \text{Span}(\mathbf{U}_{\text{sub},(j)}^*)$  and there is no need for a different notation. However, if a subspace change occurred inside this interval, then we cannot say anything simple like this. All we can say is that  $\mathbf{X}^* = [\mathbf{U}_{\text{sub},(j-1)}^* \mathbf{D}_{(j-1)}, \mathbf{U}_{\text{sub},(j)}^* \mathbf{D}_{(j)}]$  and so  $\text{Span}(\mathbf{U}^*) \subseteq \text{Span}(\mathbf{U}_{\text{sub},(j-1)}^*) \cup \text{Span}(\mathbf{U}_{\text{sub},(j)}^*)$ .

#### B. Basic PST algorithm and extensions

As noted earlier the PST algorithm is a simple modification of the static case algorithm (AltMinLowRaP) along with a carefully designed change detection strategy. In the static case, in each iteration, we used a set of  $mq$  measurements of a single  $n \times q$  matrix  $\mathbf{X}^*$ . For obtaining the guarantees, we assumed a new (independent) set of  $2mq$  measurements of the same matrix  $\mathbf{X}^*$  were used in each iteration ( $mq$  for updating the estimate of  $\mathbf{B}^*$  and another  $mq$  for  $\mathbf{U}^*$ ). For the tracking setting, using a mini-batch size of  $\alpha$ , we proceed as follows:

each new update iteration uses  $2m\alpha$  measurements of a *new*  $\alpha$ -consecutive-column sub-matrix of  $\mathbf{X}^*_{\text{full}}$ . The input to the update iteration is the subspace estimate from the previous iteration. Under the assumption that the subspace remains constant for at least  $T\alpha$  time instants after a subspace change has been detected, this approach works: with  $T = C \log(1/\epsilon)$ , we can show that, after  $T\alpha$  time instants, we get an  $\epsilon$ -accurate estimate of the  $j$ -th subspace.

We summarize the algorithm in Algorithm 2. This toggles between a “detect” and an “update” mode. It starts in the “update” mode (described above) and remains in it for the first  $T\alpha$  time instants. At this time it enters the “detect” mode. We are able to guarantee that, when the algorithm enters this detect mode, the previous subspace has been estimated to  $\epsilon$  error whp. In the detect mode, the algorithm *does not* perform any subspace updates. This is done to simplify our analysis; it ensures that, in the interval during which the subspace change occurs, the subspace is not updated. This is what allows us to use our previous two main claims (Claims 3.1 and 3.4) without change to analyze the update mode. Practically, this is of course wasteful. We develop an improvement below.

To understand the change detection strategy, let  $\hat{k}_j$  denote the estimated change times. Consider an  $\alpha$ -length interval,  $\mathcal{J}_\alpha$ , contained in  $[k_j, k_{j+1})$ . Assume that an  $\epsilon$ -accurate estimate of the previous subspace  $\mathbf{U}_{\text{sub},(j-1)}^*$  has been obtained by  $\hat{k}_{j-1} + T\alpha$  and that this time is before  $k_j$ . Let  $\mathbf{U}_{\text{sub},(j-1)}$  denote this estimate. Define the matrix

$$\begin{aligned} \mathbf{Y}_{U,\text{det},\text{big}} &:= \\ & (I - \mathbf{U}_{\text{sub},(j-1)} \mathbf{U}_{\text{sub},(j-1)}') \mathbf{Y}_U (I - \mathbf{U}_{\text{sub},(j-1)} \mathbf{U}_{\text{sub},(j-1)}') \end{aligned}$$

with  $\mathbf{Y}_U = \mathbf{Y}_U(\mathcal{J}_\alpha)$ . This means that  $\mathbf{Y}_U$  is as defined earlier in (6) with the  $k$  summation being over all  $k \in \mathcal{J}_\alpha$  (it is using measurements for all the columns within this  $\alpha$ -length interval). With a little bit of work (see Lemma C.1 and its proof), one can show that, in this interval, the matrix  $\mathbf{Y}_{U,\text{det}} := \mathbf{U}_{\text{sub},(j-1),\perp}' \mathbf{Y}_{U,\text{det},\text{big}} \mathbf{U}_{\text{sub},(j-1),\perp}$  is close to a matrix  $\mathbf{E}_{\text{det}}$  whose eigenvalues satisfy

$$\begin{aligned} \lambda_{\max}(\mathbf{E}_{\text{det}}) - \lambda_{\min}(\mathbf{E}_{\text{det}}) \\ \geq 1.5(\sin \Theta(\mathbf{U}_{\text{sub},(j-1)}, \mathbf{U}_{\text{sub},(j)}^*) - 2\epsilon)^2 \frac{\sigma_{\min}^*}{\alpha}^2. \end{aligned}$$

On the other hand, in an  $\alpha$ -length interval contained in  $[\hat{k}_{j-1} + T\alpha, k_j)$ ,

$$\begin{aligned} \lambda_{\max}(\mathbf{E}_{\text{det}}) - \lambda_{\min}(\mathbf{E}_{\text{det}}) \\ \leq \sin^2 \Theta(\mathbf{U}_{\text{sub},(j-1)}, \mathbf{U}_{\text{sub},(j-1)}^*) \sigma_{\max}^* \frac{\sigma_{\max}^*}{\alpha} \leq \epsilon^2 \frac{\sigma_{\max}^*}{\alpha}^2. \end{aligned}$$

Thus, this quantity is small when the  $j$ -th change has not occurred (before  $k_j$ ), and is large when the subspace has changed (after  $k_j$ ). By using a large enough lower bound on the product  $m\alpha$ , the same can be shown for the difference between the maximum and minimum eigenvalues of  $\mathbf{Y}_{U,\text{det}}$  (these are equal to the maximum and  $(n-r)$ -th eigenvalues of  $\mathbf{Y}_{U,\text{det},\text{big}}$ ).

Once we have an  $\epsilon$ -accurate estimate of the current subspace, it is straightforward to also recover the corresponding signals  $\mathbf{x}_k^*$ . This can simply be done by solving a standard

PR problem to recover the coefficients vector. See last line of Algorithm 2. This borrows a similar idea from [55].

1) *Improved algorithm: PST-all*: Notice from Theorem 5.1 that Algorithm 2 can only provably detect and track subspace changes that are larger than a small threshold. While this makes sense for detection, it should be possible to track all types of changes. By including a simple modification in Algorithm 2 (include the “update” step during the detection mode as well), we can empirically demonstrate that this is indeed true. We demonstrate this in Fig 6(a). Moreover, PST-all also removes the other limitation of basic PST (not using the detect phase samples for improving the subspace estimate). Thus, even for large changes that basic PST can detect, PST-all has better tracking performance; see Fig 6(b).

### C. Guarantee for basic PST

We can prove the following about Algorithm 2 (basic PST).

**Corollary 5.1** (PST algorithm). *Consider Algorithm 2. Pick any value of  $m \geq C \max(r, \log n, \log q_{\text{full}})$ . For this  $m$ , set  $\alpha = \frac{C\kappa^{12}\mu^4 \cdot nr^4}{m}$ . Set  $T := C \log(1/\epsilon)$ , and the detection threshold  $\omega_{\text{det}} = c/(\kappa^2 r)$ . Assume that  $k_{j+1} - k_j \geq (T+3)\alpha$  and that  $\sin \Theta(\mathbf{U}_{\text{sub},(j-1)}^*, \mathbf{U}_{\text{sub},(j)})^2 > \frac{2c}{\kappa^2 r}$ . Then, w.p. at least  $1 - Cn^{-10}$ , no itemsep*

- 1) *we can detect the change with a delay of at most  $2\alpha$ , while ensuring no false detections, i.e.,  $k_j \leq \hat{k}_j \leq k_j + 2\alpha$ ;*
- 2) *for any  $\epsilon > 0$ , we can get an  $\epsilon$ -accurate estimate of the  $j$ -th subspace with a delay of at most  $(T+3)\alpha$  from  $k_j$  (when the subspace changed);*
- 3) *we have the following subspace error bounds: let  $\mathbf{U}_{\text{sub},(j)}^{(-1)} = \mathbf{U}_{\text{sub},(j-1)} := \mathbf{U}_{\text{sub},(j-1)}^{(T)}$ , and let  $\mathbf{U}_{\text{sub},(j)}^{(\ell)}$ ,  $\ell = 0, 1, \dots, T$ , be the  $\ell$ -th estimate;*

$$\begin{aligned} \sin \Theta(\mathbf{U}_{\text{sub},(j)}^{(\ell)}, \mathbf{U}_{\text{sub},(j)}^*) &\leq \\ & \begin{cases} \sin \Theta(\mathbf{U}_{\text{sub},(j-1)}^*, \mathbf{U}_{\text{sub},(j)}^*) + \epsilon & \text{if } \ell = -1 \\ (0.7)^{\ell-1} \frac{c}{\kappa^2 r} & \text{if } \ell = 0, 1, 2, \dots, T, \\ \epsilon & \text{if } \ell = T \end{cases} \end{aligned}$$

*Offline PST returns  $\hat{\mathbf{X}}$  that satisfies  $\text{mat-dist}(\hat{\mathbf{X}}, \mathbf{X}^*) \leq \epsilon$ .*

We provide a proof sketch in Appendix C.

The above result shows that, if the subspace remains constant for at least  $\alpha \log(1/\epsilon)$  time instants, and if the amount of subspace change (largest principal angle of subspace change) is of order  $1/\sqrt{r}$  or larger, then we can both detect the change and track the changed subspace to  $\epsilon$  error within a delay of order  $\alpha \log 1/\epsilon$ . Moreover, for only at most  $3\alpha$  time instants after a change, the subspace error does not reduce and is essentially bounded by the amount of change. After this, it decays exponentially every  $\alpha$  time instants.

Notice from the expression for  $\alpha$  that, if we pick the smallest allowed value of  $m$ , then the required  $\alpha$  (and hence the required delays) will be large. However, we are allowed to tradeoff  $m$  and  $\alpha$ . If we let  $m$  grow linearly with  $n$ , then we will only need  $\alpha \approx r^4$ , which is, in fact, close to the minimum required delay of  $r$ . This also matches what is seen

in existing works on provable subspace tracking (ST) in other settings (e.g., robust ST, ST with missing data, or streaming PCA with missing data) [55]–[57]. These are able to allow close to optimal detection and tracking delays but all these assume that  $m$  increases linearly with  $n$ . We can also pick any value of  $m$  in between the two extremes of  $m = Cr$  or  $m = Cn$ . For example, if  $m = Cn/r$ , then  $\alpha = r^5$  and so on.

**Algorithm 2** PST: detect and track large subspace changes

```

1: Set  $r$  equal to the largest index  $j$  for which  $\lambda_j(\mathbf{Y}_U) - \lambda_n(\mathbf{Y}_U) \geq \omega$ .
2:  $\hat{k}_0 \leftarrow 0, j \leftarrow 0, \ell \leftarrow 0$ 
3: Mode  $\leftarrow$  update
4: for  $k \geq 0$  do
5:   if Mode = update then
6:     if  $k = \hat{k}_j + (\ell + 1)\alpha$  then
7:       if  $\ell = 0$  then
8:          $\mathbf{U}_{\text{sub},(j)}^\ell \leftarrow$  top  $r$  singular vectors of  $\mathbf{Y}_U$ .
9:       end if
10:       $\hat{\mathbf{b}}_\tau \leftarrow \text{RWF}((\mathbf{y}_\tau, \mathbf{U}_{\text{sub},(j)}^\ell)' \mathbf{A}_\tau), T_{\text{RWF},\ell}$ , for  $\tau \in [k - \alpha + 1, k]$ 
11:      QR decomposition  $\hat{\mathbf{B}} \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}$ 
12:       $\hat{\mathbf{C}}_\tau \leftarrow \text{Phase}(\mathbf{A}_\tau' \mathbf{U}_{\text{sub},(j)}^\ell \hat{\mathbf{b}}_\tau), \text{for } \tau \in [k - \alpha + 1, k]$ 
13:       $\hat{\mathbf{U}}_{\text{sub},(j)}^{\ell+1} \leftarrow \arg \min_{\mathbf{U}} \sum_{\tau \in [k - \alpha + 1, k]} \|\hat{\mathbf{C}}_\tau \mathbf{y}_\tau - \mathbf{A}_\tau' \mathbf{U} \hat{\mathbf{b}}_\tau\|^2$ 
14:      QR decomposition  $\hat{\mathbf{U}}_{\text{sub},(j)}^{\ell+1} \stackrel{\text{QR}}{=} \mathbf{U}_{\text{sub},(j)}^{\ell+1} \mathbf{R}_U$ 
15:       $\ell \leftarrow \ell + 1$ 
16:    end if
17:    if  $\ell = T$  then
18:       $\mathbf{U}_{\text{sub},(j)} \leftarrow \mathbf{U}_{\text{sub},(j)}^T$ , Mode  $\leftarrow$  detect
19:    end if
20:  end if
21:  if Mode = detect then
22:    if  $\lambda_{\max}(\mathbf{Y}_{U,\text{det},\text{big}}) - \lambda_{n-r}(\mathbf{Y}_{U,\text{det},\text{big}}) \geq \omega_{\text{det}}$  then
23:       $j \leftarrow j + 1, \hat{k}_j \leftarrow k, \ell \leftarrow 0$ , Mode  $\leftarrow$  update
24:    end if
25:  end if
26:  Output  $\mathbf{U}_{\text{sub},(j)}^\ell$ 
27: end for

Offline PST: For each  $k \in [\hat{k}_j, \hat{k}_{j+1})$ , output  $\hat{\mathbf{x}}_k = \mathbf{U} \hat{\mathbf{d}}_k^*$  where  $\hat{\mathbf{d}}_k^*$  is a (at most)  $2r$ -length vector obtained by RWF applied on  $\{\mathbf{y}_{ik}, (\mathbf{U}' \mathbf{a}_{ik}), i = 1, 2, \dots, m\}$  with  $\mathbf{U} = \text{basis}([\mathbf{U}_{\text{sub},(j)}, \mathbf{U}_{\text{sub},(j+1)}])$ . Here  $\text{basis}(\mathbf{U}_1, \mathbf{U}_2)$  means a matrix with orthonormal columns that span the subspace spanned by the columns of  $\mathbf{U}_1$  and  $\mathbf{U}_2$ . We need to use the union of both subspace estimates because the actual subspace change time,  $k_{j+1}$ , is not known. Corollary 5.1 implies that, whp, it is contained in  $[\hat{k}_j, \hat{k}_{j+1})$ .
```

1) *Related Work:* For Phaseless Subspace Tracking (PST) the only works before this work was our first heuristic versions [58], and [59]. Other subspace tracking (ST) problems that have been extensively studied include dynamic compressive sensing [60] (a special case of ST where the subspace is

defined by the span of a subset of  $r$  vectors from a known dictionary matrix), dynamic robust PCA (or robust ST), see [55], [56] and references therein, streaming PCA with missing data [57], [61], and ST with missing data [62]–[66]. In terms of works with complete provable guarantees, there is the nearly optimal robust subspace tracking via recursive projected compressive sensing approach [55], [56], [66] and its precursors; recent papers on streaming PCA with missing data [57], [61], and older work on dynamic compressive sensing (CS) [60]. For robust ST, the problem setting itself implies  $m = n/2$ . In the streaming PCA case, the availability of  $m = \rho n$  measurements, with  $\rho < 1$ , is assumed. This is why both achieve close to optimal tracking delays (at least when the added unstructured noise is nearly zero). As noted earlier, our method can also achieve a delay of order  $r^4$  if we let  $m$  grow linearly with  $n$ .

Dynamic CS (like basic CS) is able to detect support changes (with sufficiently nonzero magnitude) immediately even with a small value of  $m = Cr \log n$  measurements; here  $r$  is the sparsity level (support size). This is because it is a much simpler special case of ST: in this case, one just needs to be finding the correct subset of basis vectors from a large provided set (dictionary matrix).

#### D. Numerical experiments

This experiment evaluates the PST algorithm (Algorithm 2) and PST-all algorithms from Sec. V. We generate the true data for the first subspace  $\mathbf{X}^* = \mathbf{U}_{\text{sub},(0)}^* \mathbf{D}_0^*$  where  $\mathbf{U}_{\text{sub},(0)}^* \in \mathbb{R}^{n \times r}$  with  $n = 300$ ,  $r = 2$  is generated by orthonormalizing the columns of a  $n \times r$  iid standard normal matrix. The entries of  $\mathbf{D}_0^* \in \mathbb{R}^{r \times t_1}$  with  $t_1 = 2992$  are also generated from an i.i.d. standard normal distribution. We generate the true data from the second subspace similarly and set  $\mathbf{X}^* = \mathbf{U}_{\text{sub},(1)}^* \mathbf{D}_1^*$  and we set  $q = 6000$ . Notice that  $\kappa \approx 1$ . The subspace  $\mathbf{U}_{\text{sub},(1)}^*$  is generated using the idea of [55] as  $\mathbf{U}_{\text{sub},(1)}^* = e^{-\gamma \mathbf{M}} \mathbf{U}_{\text{sub},(0)}^*$  in order to control the subspace error. Here  $\mathbf{M}$  is a skew-symmetric matrix and  $\gamma$  controls the amount of subspace change. We study two cases in which we set  $\gamma = 0.08, 0.001$  which roughly translates to  $\sin \Theta(\mathbf{U}_{\text{sub},(0)}^*, \mathbf{U}_{\text{sub},(1)}^*) = 0.8, 0.01$ . We generate the measurement matrices  $\mathbf{A}_k (\in \mathbb{R}^{n \times m}) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathbf{I})$  with  $m = 100$  for  $i = 1, \dots, q$ . We then implemented PST (Algorithm 2) and PST-all. PST requires large-enough change in order to ensure good results, and PST-all which works even with small changes. We chose the algorithm parameters as follows. We set  $\alpha = 250$  and  $L = 8$ . For the detection, and initialization steps of both algorithms we set  $m_{\text{init}} = m$ . We set the threshold for detection,  $\omega = 0.6$  through cross-validation. The results for the two algorithms are shown in Fig. 6. Notice that for the small change case, since PST is always in the detect mode, it does not improve the estimation error whereas PST-all does. However, when the change is large enough, both algorithms converge to a small error. The results are averaged over 100 independent trials.

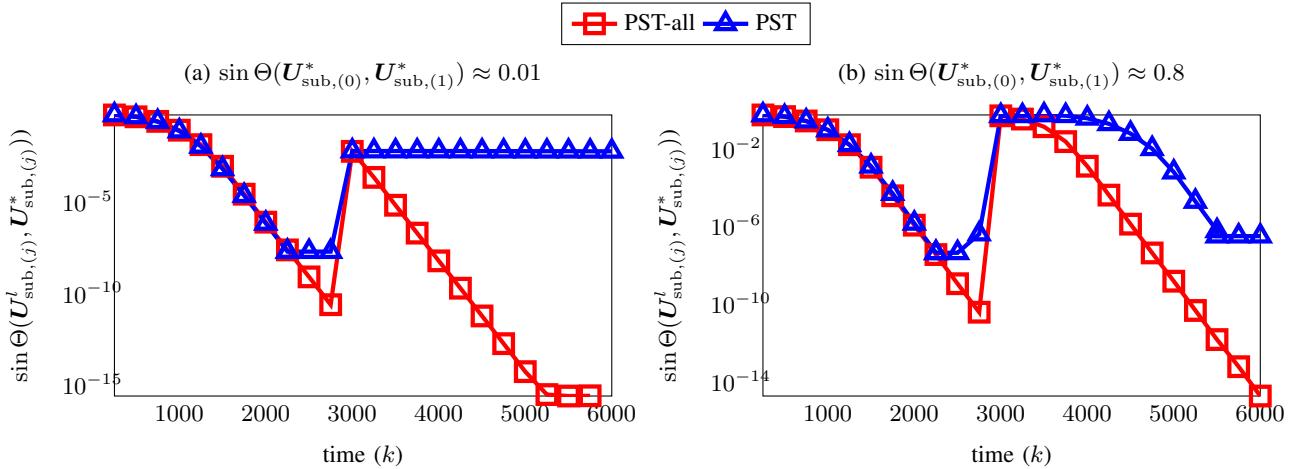


Fig. 6: Plot of subspace error versus time at each  $\alpha$  frames. Notice that for the cases where  $\sin \Theta(\mathbf{U}_{\text{sub},(0)}^*, \mathbf{U}_{\text{sub},(1)}^*) = 0.8$  both algorithms are able to detect and track changes whereas when  $\sin \Theta(\mathbf{U}_{\text{sub},(0)}^*, \mathbf{U}_{\text{sub},(1)}^*) = 0.01$  only the PST-all algorithm works. We perform the experiment for 100 independent trials, and plot the average taken over the best 90 trials.

## VI. CONCLUSIONS AND FUTURE WORK

This work introduced the first simple, fast, and provably correct, algorithm for Low Rank Phase Retrieval – low-rank matrix recovery from different (mutually independent) column-wise phaseless linear projections – AltMinLowRaP. Moreover, since, even the linear version of our problem has not been studied, this work also provides the first fast and provable solution to the linear version, “Compressive” or “Sketched PCA”. AltMinLowRaP relies on a careful spectral initialization followed by alternating minimization. We showed that its required sample complexity is about  $r^3$  times the order-optimal value of  $nr$ . We also developed its dynamic extension that is relevant for datasets where we would like to develop a mini-batch solution that recovers the current sub-matrix of  $\mathbf{X}^*$  without waiting for all the measurements of all the signals (columns) to arrive.

In ongoing work we are (i) exploring how to remove the sample-splitting requirement by studying an alternating gradient descent solution, and attempting to borrow the leave-one-out ideas from [12]; (ii) how to reduce the dependence of our sample complexity on  $r$  and on  $\kappa$ . Some thoughts are provided in Sec. II-D; and (iii) how to analyze an easy modification of AltMinLowRaP to get a better guarantee for the linear version of our problem. The algorithm modification is easy, it just involves replacing the standard PR step for recovering  $\mathbf{b}_k^*$ ’s by a simpler LS step, and of course remove the phase/sign estimation step before updating  $\mathbf{U}$ . In terms of analysis, (a) we can possibly simplify the analysis of the initialization step because in the linear case,  $\mathbb{E}[\mathbf{a}_i k \mathbf{y}_i k] = \mathbf{x}_k^*$ ; and (b) in the iterations, there will be no phase error term, Term2, and hence, no need for Cauchy-Schwarz. This latter change itself will reduce the sample complexity to  $nr^3$  instead of  $nr^4$ .

Open questions for future work include (i) study if we can exploit the right incoherence assumption in the algorithm itself, for example, by using a projected GD approach inspired by [67]; (ii) can phaseless LRMS be solved (this would be the other possible LRPR problem alluded to in the introduction),

and (iii) develop a fast algorithm and a guarantee for exploiting both low-rank (as we do) and column-wise sparsity. This type of modeling been used very successfully in the MRI literature to come up with practical algorithms to reduce the sample complexity empirically, see for example, [30], [31]. It has also been studied theoretically in the linear setting [68]. (iv) Another open question is how to analyze the improved tracking algorithm PST-all that we currently only empirically evaluate. In experiments, it is clearly much better than the simpler version we analyze.

## APPENDIX A PROOFS OF THE LEMMAS FROM SEC. III

### A. Simple facts for various proofs

Our proofs will use the following facts: for two arbitrary matrices  $\mathbf{A}, \mathbf{H}$ ,

- 1)  $\sigma_{\max}(\mathbf{A} + \mathbf{H}) \leq \sigma_{\max}(\mathbf{A}) + \|\mathbf{H}\|$
- 2)  $\sigma_{\min}(\mathbf{A} + \mathbf{H}) \geq \sigma_{\min}(\mathbf{A}) - \|\mathbf{H}\|$ .
- 3)  $\sigma_{\min}(\mathbf{A}\mathbf{H}) \geq \sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{H})$ .
- 4) For two basis matrices,  $\mathbf{U}_1, \mathbf{U}_2$ ,  $\sigma_{\min}^2(\mathbf{U}_1' \mathbf{U}_2) = 1 - \sin \Theta(\mathbf{U}_1, \mathbf{U}_2)$ .
- 5) For any matrix  $\mathbf{P}$ ,  $\|\mathbf{P}\mathbf{M}\|_F \leq \|\mathbf{P}\| \|\mathbf{M}\|_F$  and  $\|\mathbf{M}\mathbf{P}\|_F \leq \|\mathbf{P}\| \|\mathbf{M}\|_F$ .
- 6) For an invertible matrix  $\mathbf{P}$ ,  $\|\mathbf{M}\|_F = \|\mathbf{P}^{-1}\mathbf{P}\mathbf{M}\|_F \leq \|\mathbf{P}^{-1}\| \|\mathbf{P}\mathbf{M}\|_F = \frac{1}{\sigma_{\min}(\mathbf{P})} \|\mathbf{P}\mathbf{M}\|_F$ .

The following lemma is a simple but useful modification of [Lemma 5.16] [54]. The proof follows by combining Lemma 2.7.7. and Theorem 2.8.1. of [29]. The sub-Gaussian norm for a random vector  $X$ , in this lemma, can be defined as follows:

$$\|X\|_{\psi_2} = \sup_{x \in \mathcal{S}^{n-1}} \|\langle X, x \rangle\|_{\psi_2},$$

where  $\|z\|_{\psi_2}$  is the sub-Gaussian norm of a scalar  $z$  [54, Definition 5.7].

**Lemma A.1.** Let  $X_i, Y_i$  be independent sub-Gaussian random variables with sub-Gaussian norm  $K_{X_i}$  and  $K_{Y_i}$  respectively and with  $\mathbb{E}[X_i Y_i] = 0$ . Then

$$\Pr \left\{ \left| \sum_i X_i Y_i \right| \geq t \right\} \leq 2 \exp \left( -c \min \left( \frac{t^2}{\sum_i K_{X_i}^2 K_{Y_i}^2}, \frac{t}{\max_i |K_{X_i} K_{Y_i}|} \right) \right)$$

When  $X_i = Y_i$ , this simplifies to Lemma 5.16 of [54].

### B. Proof of the lemmas for Claim 3.1

In this section, we let  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(0)}$  and  $\mathbf{y}_{ik} := \mathbf{y}_{ik}^{(0)}$ .

*Proof of Lemma 3.6.* Observe that we will be done if we can show that, whp,  $\frac{1}{mq} \sum_{ik} (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2$  lies in the interval  $[(1 - \epsilon_1) \|\mathbf{X}^*\|_F^2 / q, (1 + \epsilon_1) \|\mathbf{X}^*\|_F^2 / q]$ . Using Lemma A.1, with  $K_{X_{ik}} = K_{Y_{ik}} = \|\mathbf{x}_k\|$ , with probability more than  $1 - 2 \exp \left( \frac{-C\epsilon_1^2 mq}{\mu^2 \kappa^2} \right)$ , we have

$$\left| \sum_{ik} (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2 - m \|\mathbf{X}^*\|_F^2 \right| \leq \epsilon_1 m \|\mathbf{X}^*\|_F^2.$$

Details for obtaining this bound: using  $\sum_k \|\mathbf{x}_k^*\|^4 \leq \max_k \|\mathbf{x}_k^*\|^2 \sum_k \|\mathbf{x}_k^*\|^2$  and right incoherence,

$$\begin{aligned} \frac{t^2}{\sum_{ik} K_{X_{ik}}^4} &\geq \frac{\epsilon_1^2 m^2 \|\mathbf{X}^*\|_F^4}{m \max_k \|\mathbf{x}_k^*\|^2 \|\mathbf{X}^*\|_F^2} \geq \frac{\epsilon_1^2 mq}{\mu^2 \kappa^2}, \\ \frac{t}{\max K_{X_{ik}}^2} &= \frac{\epsilon_1 m \|\mathbf{X}^*\|_F^2}{\|\mathbf{x}_k^*\|^2} \geq \frac{\epsilon_1 mq}{\mu^2 \kappa^2}. \end{aligned}$$

□

*Proof of Lemma 3.7.* It is easy to see that

$$\begin{aligned} \|\mathbb{E}[\mathbf{Y}_+] - \mathbb{E}[\mathbf{Y}_-]\| &\leq \frac{1}{q} \sum_k \left( \beta_{1,k}^+ - \beta_{1,k}^- \right) \|\mathbf{x}_k^*\|^2 \\ &\quad + \frac{1}{q} \sum_k \left( \beta_{2,k}^+ - \beta_{2,k}^- \right) \|\mathbf{x}_k^*\|^2. \end{aligned}$$

Recall  $\gamma_k = 9 \|\mathbf{X}^*\|_F^2 \mu^2 \kappa^2 / (q \|\mathbf{x}_k^*\|^2)$ . Using  $x^3 e^{-x^2/2} \leq 3\sqrt{3}e^{-3/2}$ , we have

$$\begin{aligned} \beta_{1,k}^+ - \beta_{1,k}^- &= \mathbb{E} [(\xi^4 - \xi^2) \mathbb{1}_{\{(1-\epsilon_1)\gamma_k \leq \xi^2 \leq (1+\epsilon_1)\gamma_k\}}] \\ &= \frac{2}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\sqrt{(1+\epsilon_1)\gamma_k}} x^2 (x^2 - 1) e^{-x^2/2} dx \\ &\leq \frac{2}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\sqrt{(1+\epsilon_1)\gamma_k}} x^4 e^{-x^2/2} dx \\ &\leq \frac{6\sqrt{3}e^{-3/2}}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\sqrt{(1+\epsilon_1)\gamma_k}} x dx \\ &= \frac{6\sqrt{3}e^{-3/2}}{\sqrt{2\pi}} \gamma_k \epsilon_1 \leq \gamma_k \epsilon_1. \end{aligned}$$

Similarly, using  $xe^{-x^2/2} \leq e^{-1/2}$ ,

$$\begin{aligned} \beta_{2,k}^+ - \beta_{2,k}^- &= \mathbb{E} [\xi^2 \mathbb{1}_{\{(1-\epsilon_1)\gamma_k \leq \xi^2 \leq (1+\epsilon_1)\gamma_k\}}] \\ &= \frac{2}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\sqrt{(1+\epsilon_1)\gamma_k}} x^2 e^{-x^2/2} dx \\ &\leq \frac{2e^{-1/2}}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\sqrt{(1+\epsilon_1)\gamma_k}} x dx \\ &= \frac{2e^{-1/2}}{\sqrt{2\pi}} \gamma_k \epsilon_1 \leq \gamma_k \epsilon_1. \end{aligned}$$

Therefore,

$$\begin{aligned} \|\mathbb{E}[\mathbf{Y}_+] - \mathbb{E}[\mathbf{Y}_-]\| &\leq \frac{\epsilon_1}{q} \sum_k \gamma_k \|\mathbf{x}_k^*\|^2 \\ &= 9 \frac{\epsilon_1 \mu^2 \kappa^2}{q} \sum_k \frac{\|\mathbf{X}^*\|_F^2}{q} = 9\epsilon_1 \frac{\mu^2 \kappa^2 \|\mathbf{X}^*\|_F^2}{q}. \end{aligned}$$

To lower bound  $\beta_{1,k}^-$ , we will use right incoherence which implies that  $\gamma_k \geq 9$ .

$$\begin{aligned} \beta_{1,k}^- &= \mathbb{E} [\xi^2 (\xi^2 - 1) \mathbb{1}_{\{\xi^2 \leq (1-\epsilon_1)\gamma_k\}}] \\ &= \mathbb{E} [(\xi^4 - \xi^2)] - \mathbb{E} [\xi^2 (\xi^2 - 1) \mathbb{1}_{\{\xi^2 \geq (1-\epsilon_1)\gamma_k\}}] \\ &= 2 - 2 \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\infty} x^2 (x^2 - 1) \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\ &\geq 2 - \frac{2}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\infty} x^4 e^{-x^2/2} dx \\ &\geq 2 - \frac{7}{\sqrt{2\pi}} \int_{\sqrt{(1-\epsilon_1)\gamma_k}}^{\infty} x e^{-x^2/4} dx \\ &= 2 - \frac{14}{\sqrt{2\pi}} \exp(-(1-\epsilon_1)\gamma_k/4) > 1.5, \end{aligned}$$

where we used the fact that  $x^3 e^{-x^2/4} \leq 3.5$  for any  $x$ ;  $\gamma_k \geq 9$  (follows by right incoherence); and  $\epsilon_1 < 0.01$ . □

*Proof of Lemma 3.8.* Let us define

$$\mathbf{w}_{ik} = |\mathbf{a}_{ik}' \mathbf{x}_k^*| \mathbf{a}_{ik} \mathbb{1}_{\left\{ (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2 \leq \frac{9(1-\epsilon_1) \|\mathbf{X}^*\|_F^2 \mu^2 \kappa^2}{q} \right\}}.$$

As argued in [20], which itself borrows the key idea from [9], we can show that the  $\mathbf{w}_{ik}$ s are sub-Gaussian random variables with sub-Gaussian norm  $K = C\mu\kappa\|\mathbf{X}^*\|_F/\sqrt{q}$ . Notice that we have defined  $\mathbf{Y}_U$  differently in this paper (in order to be able to exploit concentration over  $mq$ ) as compared to that in [20] and hence only the above argument is similar.

Observe that

$$mq \|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\| = \max_{\mathbf{z}: \|\mathbf{z}\|=1} |\mathbf{z}' \sum_{ik} (\mathbf{w}_{ik} \mathbf{w}_{ik}' - \mathbb{E}[\mathbf{w}_{ik} \mathbf{w}_{ik}']) \mathbf{z}|.$$

First consider a fixed unit vector  $\mathbf{z}$ . Observe that  $\mathbf{z}' \mathbf{w}_{ik}$  is sub-Gaussian with sub-Gaussian norm  $K = C\mu\kappa\|\mathbf{X}^*\|_F/\sqrt{q}$ . Thus, using Lemma A.1 with  $t = \epsilon_2 m \mu^2 \kappa^2 \|\mathbf{X}^*\|_F^2$ , and  $K_{X_{ik}} = K_{Y_{ik}} = \frac{\mu\kappa\|\mathbf{X}^*\|_F}{\sqrt{q}}$ , we can conclude that w.p. at least  $\geq 1 - 2 \exp(-c\epsilon_2^2 mq)$ ,

$$|\mathbf{z}' (\sum_{ik} \mathbf{w}_{ik} \mathbf{w}_{ik}' - mq \mathbb{E}[\mathbf{Y}_-]) \mathbf{z}| \leq \epsilon_2 m \mu^2 \kappa^2 \|\mathbf{X}^*\|_F^2,$$

After this, we can use a standard epsilon-net argument to extend the bound to all unit vectors  $\mathbf{z}$ . With it, we can conclude that, w.p. at least  $1 - 2 \exp(n \log 9 - c\epsilon_2^2 mq)$ ,

$$\|\mathbf{Y}_- - \mathbb{E}[\mathbf{Y}_-]\| \leq \frac{1.5\epsilon_2\mu^2\kappa^2\|\mathbf{X}^*\|_F^2}{q}$$

□

### C. Clarifying the sign inconsistency issue

Recall that we had defined  $\mathbf{g}_k^t := (\mathbf{U}^t)' \mathbf{x}_k^*$  in (7). Since the solution of phase retrieval always comes with a phase (sign) ambiguity, at each iteration  $t$ , for each  $k$ , the output of RWF,  $\hat{\mathbf{b}}_k^t$ , may be closer to either  $\mathbf{g}_k^t$  or  $-\mathbf{g}_k^t$ . This is what decides whether  $\text{dist}(\mathbf{g}_k^t, \hat{\mathbf{b}}_k^t)$  equals  $\|\mathbf{g}_k^t - \hat{\mathbf{b}}_k^t\|$  or  $\|\mathbf{g}_k^t + \hat{\mathbf{b}}_k^t\|$ . However, bound both in each proof is cumbersome. Instead we can proceed as follows. *Re-define*  $\mathbf{g}_k^t$  as

$$\mathbf{g}_k^t = \begin{cases} +(\mathbf{U}^t)' \mathbf{x}_k^* & \text{if } \|(\mathbf{U}^t)' \mathbf{x}_k^* - \hat{\mathbf{b}}_k^t\| \leq \|(\mathbf{U}^t)' \mathbf{x}_k^* + \hat{\mathbf{b}}_k^t\| \\ -(\mathbf{U}^t)' \mathbf{x}_k^* & \text{otherwise} \end{cases}$$

and define the matrix

$$\mathbf{G}^t := [\mathbf{g}_1^t, \mathbf{g}_2^t, \dots, \mathbf{g}_q^t].$$

With these new definitions,  $\text{dist}(\mathbf{g}_k^t, \hat{\mathbf{b}}_k^t) = \|\mathbf{g}_k^t - \hat{\mathbf{b}}_k^t\|$  and  $\text{mat-dist}(\mathbf{G}^t, \hat{\mathbf{B}}^t) = \|\mathbf{G}^t - \hat{\mathbf{B}}^t\|_F$ .

As an aside, we should point out that, even if some columns of a matrix change sign (are multiplied by  $(-1)$ ), its singular values do not change. Thus, the minimum singular value of  $\mathbf{G}^t$  remains the same with or without the above re-definition.

We need to do something similar to the above for  $\mathbf{x}_k^*$ 's as well. Define  $\tilde{\mathbf{x}}_k^t = \mathbf{x}_k^*$  if  $\|\tilde{\mathbf{x}}_k^t - \mathbf{x}_k^*\| \leq \|\tilde{\mathbf{x}}_k^t + \mathbf{x}_k^*\|$  and  $\tilde{\mathbf{x}}_k^t = -\mathbf{x}_k^*$  otherwise. Define the corresponding matrix  $\tilde{\mathbf{X}}^t$ .

Clearly  $\text{mat-dist}(\tilde{\mathbf{X}}^t, \mathbf{X}^*) = 0$ . So, in the rest of the writing in this section, to reduce notation, we will *re-define*

$$\mathbf{X}^* := \tilde{\mathbf{X}}^t.$$

With this, we can define the error/perturbation in  $\hat{\mathbf{x}}_k$  as just

$$\mathbf{h}_k := \hat{\mathbf{x}}_k^t - \mathbf{x}_k^*$$

and we have  $\text{dist}(\hat{\mathbf{x}}_k^t, \mathbf{x}_k^*) = \|\hat{\mathbf{x}}_k^t - \mathbf{x}_k^*\| = \|\mathbf{h}_k\|$ .

### D. Proof of Lemmas 3.2 and 3.3

In this section, we let  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(t)}$  and  $\mathbf{y}_{ik} := \mathbf{y}_{ik}^{(t)}$ . Also, everywhere below, we remove the superscripts  $t$  for ease of notation. Recall that  $\mathbf{x}_k^* = \mathbf{U}\mathbf{g}_k + \mathbf{e}_k$  with  $\mathbf{e}_k := (\mathbf{I} - \mathbf{U}\mathbf{U}')\mathbf{x}_k^*$ .

*Proof of Lemma 3.2.* To estimate  $\mathbf{b}_k$ , we first need to estimate  $\mathbf{g}_k$  which requires measurements of the form  $\mathbf{a}_{ik}'\mathbf{U}\mathbf{g}_k$ . Our measurements satisfy

$$\mathbf{y}_{ik} = |\mathbf{a}_{ik}'\mathbf{U}\mathbf{g}_k| + \nu_{ik},$$

where  $\nu_{ik} = |\mathbf{a}_{ik}'\mathbf{x}_k^*| - |\mathbf{a}_{ik}'\mathbf{U}\mathbf{g}_k|$  is the noise. We use these to obtain the estimate  $\hat{\mathbf{b}}_k$  using RWF. By Theorem 2 of [10], if  $m \geq Cr$ , w.p. at least  $1 - \exp(-cm)$ ,

$$\text{dist}(\mathbf{g}_k, \hat{\mathbf{b}}_k) \leq \frac{\|\nu_k\|}{\sqrt{m}} + (1 - c_1)^{T_{\text{RWF},t}} \|\mathbf{g}_k\|,$$

where  $c_1$  is a constant less than one. For our problem,

$$|\nu_{ik}| \leq |\mathbf{a}_{ik}'\mathbf{U}\mathbf{g}_k + \mathbf{a}_{ik}'\mathbf{e}_k| - |\mathbf{a}_{ik}'\mathbf{U}\mathbf{g}_k| \leq |\mathbf{a}_{ik}'\mathbf{e}_k|$$

$$\|\nu_k\|^2 = \sum_i \nu_{ik}^2 = \sum_i |\mathbf{a}_{ik}'\mathbf{e}_k|^2.$$

Clearly,  $\mathbb{E}[\|\nu_k\|^2] = m\|\mathbf{e}_k\|^2$  and  $\|\mathbf{e}_k\|^2 \leq \sin \Theta(\mathbf{U}, \mathbf{U}^*) \|\tilde{\mathbf{b}}_k^*\|$ . Using Lemma A.1 with  $t = m\delta_b \|\mathbf{e}_k\|^2$ ,  $K_{X_i} = K_{Y_i} = \|\mathbf{e}_k\|$ , and summing over  $i = 1, 2, \dots, m$ , we conclude that, w.p. at least  $1 - \exp(-c\delta_b^2 m)$ ,

$$\|\nu_k\|^2 \leq m(1 + \delta_b) \|\mathbf{e}_k\|^2 \leq m(1 + \delta_b) \delta_t^2 \|\tilde{\mathbf{b}}_k^*\|^2.$$

where the last inequality used  $\|\mathbf{e}_k\| \leq \delta_t \|\tilde{\mathbf{b}}_k^*\|$ . Thus, using the above and  $\|\mathbf{g}_k\| \leq \|\tilde{\mathbf{b}}_k^*\|$ ,

$$\text{dist}(\mathbf{g}_k, \hat{\mathbf{b}}_k) \leq \sqrt{1 + \delta_b} \delta_t \|\tilde{\mathbf{b}}_k^*\| + (1 - c_1)^{T_{\text{RWF},t}} \|\tilde{\mathbf{b}}_k^*\|.$$

By setting  $T_{\text{RWF},t}$  so that  $(1 - c)^{T_{\text{RWF},t}} \leq \delta_t$ , we get that  $\text{dist}(\mathbf{g}_k, \hat{\mathbf{b}}_k) \leq C\delta_t \|\tilde{\mathbf{b}}_k^*\| = C\delta_t \|\mathbf{x}^*\|$  with  $C = (\sqrt{1 + \delta_b} + 1)$ . The above bound holds w.p. at least  $1 - \exp(-c\delta_b^2 m)$  for a given  $k$ . By union bound, it holds for all  $k = 1, 2, \dots, q$ , w.p. at least  $1 - q \exp(-c\delta_b^2 m)$ . Hence, with this probability,

$$\text{mat-dist}(\mathbf{G}, \hat{\mathbf{B}}) \leq C\delta_t \|\mathbf{X}^*\|_F.$$

For proving the third claim, recall that  $\hat{\mathbf{B}} \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}$  and so  $\mathbf{b}_k = \mathbf{R}_B^{-1} \hat{\mathbf{b}}_k$ . Let  $\mathbf{h}_k := \mathbf{x}_k^* - \hat{\mathbf{x}}_k$ . We can rewrite  $\mathbf{h}_k$  as  $\mathbf{h}_k = \mathbf{x}_k^* - \mathbf{U}\mathbf{g}_k + \mathbf{U}\mathbf{g}_k - \hat{\mathbf{b}}_k$ . Thus, by triangle inequality, and using  $\|\mathbf{e}_k\| \leq \sin \Theta(\mathbf{U}, \mathbf{U}^*) \|\mathbf{g}_k\| \leq \delta_t \|\tilde{\mathbf{b}}_k^*\|$ ,

$$\|\mathbf{h}_k\| \leq \|\mathbf{e}_k\| + \|\mathbf{U}\|\|\mathbf{g}_k - \hat{\mathbf{b}}_k\| \leq (1 + C)\delta_t \|\tilde{\mathbf{b}}_k^*\|.$$

□

*Proof of Lemma 3.3.* Recall that  $\hat{\mathbf{B}} \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}$  and so  $\mathbf{b}_k = \mathbf{R}_B^{-1} \hat{\mathbf{b}}_k$ . Using Lemma 3.2,  $\|\mathbf{g}_k\| \leq \|\tilde{\mathbf{b}}_k^*\|$ , and right incoherence (which implies that  $\|\tilde{\mathbf{b}}_k^*\|^2 \leq \sigma_{\max}^* \mu^2 r/q$ ),

$$\begin{aligned} \|\mathbf{b}_k\| &= \|\mathbf{R}_B^{-1}(\mathbf{g}_k - \hat{\mathbf{b}}_k + \mathbf{g}_k)\| \\ &\leq \|\mathbf{R}_B^{-1}\| (\text{dist}(\hat{\mathbf{b}}_k, \mathbf{g}_k) + \|\mathbf{g}_k\|) \\ &\leq \|\mathbf{R}_B^{-1}\| (1 + C\delta_t) \|\tilde{\mathbf{b}}_k^*\| \\ &\leq \frac{(1 + C\delta_t)\sigma_{\max}^* \mu \sqrt{r/q}}{\sigma_{\min}(\mathbf{R}_B)} \leq \frac{1.5\sigma_{\max}^* \mu \sqrt{r/q}}{\sigma_{\min}(\mathbf{R}_B)} \end{aligned}$$

To lower bound  $\sigma_{\min}(\mathbf{R}_B)$ , observe that  $\sigma_{\min}(\mathbf{R}_B) = \sigma_{\min}(\hat{\mathbf{B}})$ . Using Lemma 3.2, the discussion of Sec. A-C, facts from Sec. A-A, and  $\sin \Theta(\mathbf{U}, \mathbf{U}^*) \leq \delta_t$ ,

$$\begin{aligned} \sigma_{\min}(\hat{\mathbf{B}}) &\geq \sigma_{\min}(\mathbf{G}) - \|\mathbf{G} - \hat{\mathbf{B}}\| \\ &\geq \sigma_{\min}(\mathbf{U}'\mathbf{U}^*) \sigma_{\min}(\tilde{\mathbf{B}}^*) - \|\mathbf{G} - \hat{\mathbf{B}}\|_F \\ &\geq \sqrt{1 - \sin \Theta^2(\mathbf{U}, \mathbf{U}^*)} \sigma_{\min}^* - C\delta_t \|\tilde{\mathbf{B}}^*\|_F \\ &\geq \sqrt{1 - \delta_t^2} \sigma_{\min}^* - C\delta_t \sqrt{r} \sigma_{\max}^*. \end{aligned}$$

Using  $\delta_t \leq c/\kappa\sqrt{r}$ ,  $\sigma_{\min}(\mathbf{R}_B) = \sigma_{\min}(\hat{\mathbf{B}}) \geq 0.9\sigma_{\min}^*$ . Thus,

$$\|\mathbf{b}_k\| \leq \frac{1.5\sigma_{\max}^* \mu \sqrt{r/q}}{0.9\sigma_{\min}^*} \leq 2\kappa\mu\sqrt{r/q} := \hat{\mu}\sqrt{r/q}.$$

All of the above bounds used the bound from Lemma 3.2. Thus the above bounds hold w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log q, \log n)$ . □

### E. Proof of the lemmas for Claim 3.4

In this section, we let  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(T+t)}$  and  $\mathbf{y}_{ik} := \mathbf{y}_{ik}^{(T+t)}$ . Also, at almost all places, we remove the superscript  $t$ .

All the proofs in this section use incoherence of  $\mathbf{B}$  with parameter  $\hat{\mu} = C\kappa\mu$  (by Lemma 3.3). This holds w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ .

*Proof of Lemma 3.10.* Recall that  $\text{Term3}(\mathbf{W}) := \sum_{ik} (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k)^2$ . We have

$$\mathbb{E} \left[ \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 \right] = m \|\mathbf{W} \mathbf{B}\|_F^2 = m.$$

Let  $X_{ik} = |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|$ .  $X_{ik}$  is sub-Gaussian with sub-Gaussian norm  $\|\mathbf{W} \mathbf{b}_k\|$ . We use Lemma A.1 with  $Y_{ik} = X_{ik}$  and  $t = \delta_t m$ , along with the following facts:

- 1)  $\sum_k \|\mathbf{W} \mathbf{b}_k\|^4 \leq \max_k \|\mathbf{W} \mathbf{b}_k\|^2 \sum_k \|\mathbf{W} \mathbf{b}_k\|^2$ ,
- 2)  $\sum_k \|\mathbf{W} \mathbf{b}_k\|^2 = \|\mathbf{W} \mathbf{B}\|_F^2 = \text{trace}(\mathbf{W} \mathbf{B} \mathbf{B}' \mathbf{W}') = \|\mathbf{W}\|_F^2 = 1$ , and
- 3)  $\max_k \|\mathbf{W} \mathbf{b}_k\|^2 \leq \|\mathbf{W}\|^2 \max_k \|\mathbf{b}_k\|^2 \leq \|\mathbf{W}\|_F^2 \max_k \|\mathbf{b}_k\|^2 \leq \max_k \|\mathbf{b}_k\|^2 \leq \hat{\mu}^2 r / q$ , w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log n, \log q) / \delta_b^2$  by Lemma 3.3.

Using Lemma A.1 and the above facts, for a fixed  $\mathbf{W}$ ,

$$\Pr \left\{ \left| \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 - m \right| \geq \delta_t m \right\} \leq 2 \exp \left( -c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right).$$

Now we develop an epsilon-net argument to complete the proof. This is inspired by similar arguments in [52]. Recall that  $\mathcal{S}_{\mathbf{W}} = \{\mathbf{W} \in \mathbb{R}^{n \times r}, \|\mathbf{W}\|_F = 1\}$ . By [54](Lemma 5.2), there is a set (called  $\epsilon$ -net),  $\bar{\mathcal{S}}_{\mathbf{W}} \subseteq \mathcal{S}_{\mathbf{W}}$  so that for any  $\mathbf{W}$  in  $\mathcal{S}_{\mathbf{W}}$ , there is a  $\bar{\mathbf{W}} \in \bar{\mathcal{S}}_{\mathbf{W}}$ , such that

$$\|\bar{\mathbf{W}} - \mathbf{W}\|_F \leq \epsilon$$

and

$$|\bar{\mathcal{S}}_{\mathbf{W}}| \leq \left( 1 + \frac{2}{\epsilon} \right)^{nr}.$$

Pick  $\epsilon = 1/8$  so that  $|\mathcal{S}_{\mathbf{W}}| \leq 17^{nr}$ . Also, define

$$\Delta \mathbf{W} := \bar{\mathbf{W}} - \mathbf{W}$$

so that  $\|\Delta \mathbf{W}\|_F \leq \epsilon = 1/8$ .

Using a union bound over all entries in the finite set  $\bar{\mathcal{S}}_{\mathbf{W}}$ ,

$$\begin{aligned} \Pr \left( \left| \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 - m \right| \leq \delta_t m, \text{ for all } \bar{\mathbf{W}} \in \bar{\mathcal{S}}_{\mathbf{W}} \right) \\ \geq 1 - 2|\bar{\mathcal{S}}_{\mathbf{W}}| \exp \left( -c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right) \\ \geq 1 - 2 \exp \left( nr(\log 17) - c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right). \end{aligned} \quad (14)$$

Next we extend the above to obtain lower and upper bounds over the entire hyper-sphere,  $\mathcal{S}_{\mathbf{W}}$ . Define

$$\theta_W = \max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2,$$

as the maximum of  $\text{Term3}(\mathbf{W})$  over  $\mathcal{S}_{\mathbf{W}}$ . Since  $\frac{\Delta \mathbf{W}}{\|\Delta \mathbf{W}\|_F} \in \mathcal{S}_{\mathbf{W}}$ ,

$$\sum_{ik} |\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k|^2 \leq \theta_W \|\Delta \mathbf{W}\|_F^2 \leq \theta_W \epsilon^2.$$

Using this, (14), and Cauchy-Schwarz,

$$\begin{aligned} & \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 \\ &= \sum_{ik} |\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k|^2 + \sum_{ik} |\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k|^2 \\ &+ 2 \sum_{ik} (\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k) (\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k) \\ &\leq (1 + \delta_t)m + \epsilon^2 \theta_W + 2\sqrt{m(1 + \delta_t)} \sqrt{\theta_W \epsilon} \\ &= (1 + \delta_t)m + (1/64)\theta_W + (1/4)m\sqrt{1 + \delta_t} \sqrt{\theta_W/m} \end{aligned} \quad (15)$$

w.p. at least  $1 - 2 \exp \left( nr(\log 17) - c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right)$ . The last equality just used  $\epsilon = 1/8$  and re-arranged the third term.

If  $\theta_W/m < 1$ , we are done because then  $\theta_W \leq m$ . Otherwise,  $\theta_W/m \geq 1$  and so  $\sqrt{\theta_W/m} \leq \theta_W/m$ . Using this and taking  $\max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}}$  of (15),

$$\theta_W \leq (1 + \delta_t)m + \theta_W((1/64) + (1/4)\sqrt{1 + \delta_t}).$$

By assumption,  $\delta_t < 1/10$ , and so the above implies that  $\theta_W \leq 1.25(1 + \delta_t)m$ .

Thus, w.p.  $1 - 2 \exp \left( nr(\log 17) - c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right)$ ,

$$\theta_W := \max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term3} \leq 1.25(1 + \delta_t)m \leq 1.5m.$$

We now obtain the lower bound on the minimum of Term3 over the entire hyper-sphere. This uses (14), Cauchy-Schwarz, and the upper bound on  $\theta_W$  from above.

$$\begin{aligned} & \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 \\ &\geq \sum_{ik} |\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k|^2 + 2 \sum_{ik} (\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k) (\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k) \\ &\geq \sum_{ik} |\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k|^2 - 2 \left| \sum_{ik} \mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k \mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k \right| \\ &\geq m(1 - \delta_t) - 2 \sqrt{\sum_{ik} |\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k|^2} \sqrt{\sum_{ik} |\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k|^2} \\ &\geq m(1 - \delta_t) - 2\sqrt{m(1 + \delta_t)} \sqrt{\theta_W \|\Delta \mathbf{W}\|_F^2} \\ &\geq m(1 - \delta_t) - 2m(1 + \delta_t)\sqrt{1.5\epsilon} \geq m(0.9 - 0.26) = 0.64m \end{aligned}$$

w.p.  $1 - 2 \exp \left( nr(\log 17) - c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right)$ . In the last line we substituted  $\epsilon = 1/8$  and used  $\delta_t < 1/10$ .

All of the above bounds hold on the event in which  $\mathbf{B}$  is  $\hat{\mu}$  incoherent. This holds w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log n, \log q) / \delta_b^2$  (this follows by Lemma 3.3).

Thus, if  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ , w.p.  $1 - 2 \exp \left( nr(\log 17) - c \frac{\delta_t^2 mq}{\hat{\mu}^2 r} \right) - n^{-10}$ ,

$$\min_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term3}(\mathbf{W}) \geq 0.64m.$$

□

*Proof of Lemma 3.11.* Define

$$\mathbf{p}_k := \tilde{\mathbf{B}}^* \mathbf{B}' \mathbf{b}_k - \tilde{\mathbf{b}}_k^* = \Sigma^* (\mathbf{B}^* \mathbf{B}' \mathbf{b}_k - \mathbf{b}_k^*)$$

Recall that

$$\text{Term1}(\mathbf{W}) = \sum_{ik} \mathbf{b}_k' \mathbf{W}' \mathbf{a}_{ik} \mathbf{a}_{ik}' \mathbf{U}' \mathbf{p}_k.$$

Define the matrix  $\mathbf{P} := [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_q]$ . Observe that

$$\mathbf{P} = \tilde{\mathbf{B}}^* (\mathbf{B}' \mathbf{B} - \mathbf{I}).$$

We first upper bound  $\|\mathbf{P}\|_F$  by using Lemma 3.2. In Lemma 3.2, we bounded  $\text{mat-dist}(\mathbf{G}, \hat{\mathbf{B}}) = \|\mathbf{G} - \hat{\mathbf{B}}\|_F$ . Recall also that  $\hat{\mathbf{B}} \stackrel{\text{QR}}{=} \mathbf{R}_B \mathbf{B}$  where  $\mathbf{B}$  is a matrix with orthonormal rows. Thus,  $\hat{\mathbf{B}}(\mathbf{I} - \mathbf{B}' \mathbf{B}) = 0$ . Using this and facts from Sec. A-A,

$$\begin{aligned} \|\mathbf{P}\|_F &= \|\tilde{\mathbf{B}}^* (\mathbf{I} - \mathbf{B}' \mathbf{B})\|_F \\ &\leq \frac{1}{\sigma_{\min}(\mathbf{U}' \mathbf{U}^*)} \|\mathbf{U}' \mathbf{U}^* \tilde{\mathbf{B}}^* (\mathbf{I} - \mathbf{B}' \mathbf{B})\|_F \\ &= \frac{1}{\sigma_{\min}(\mathbf{U}' \mathbf{U}^*)} \|(\mathbf{G} - \hat{\mathbf{B}} + \hat{\mathbf{B}})(\mathbf{I} - \mathbf{B}' \mathbf{B})\|_F \\ &\leq \frac{1}{\sigma_{\min}(\mathbf{U}' \mathbf{U}^*)} \|\mathbf{G} - \hat{\mathbf{B}}\|_F + 0 \\ &\leq \frac{1}{\sigma_{\min}(\mathbf{U}' \mathbf{U}^*)} C \delta_t \|\mathbf{X}^*\|_F. \end{aligned}$$

Furthermore  $\sigma_{\min}^2(\mathbf{U}' \mathbf{U}^*) = 1 - \sin \Theta(\mathbf{U}, \mathbf{U}^*)^2 \geq 1 - \delta_t^2$ . By assumption,  $\delta_t \leq 1/10$  and so

$$\|\mathbf{P}\|_F \leq \sqrt{10/9} C \delta_t \|\mathbf{X}^*\|_F. \quad (16)$$

Next we show that  $\mathbb{E}[\text{Term1}(\mathbf{W})] = 0$ .

$$\begin{aligned} \mathbb{E}[\text{Term1}(\mathbf{W})] &= m \sum_k \mathbf{b}_k' \mathbf{W}' \mathbf{U}' \Sigma^* (\mathbf{B}^* \mathbf{B}' \mathbf{b}_k - \mathbf{b}_k^*) \\ &= m \sum_k \text{tr}(\mathbf{W}' \mathbf{U}' \Sigma^* (\mathbf{B}^* \mathbf{B}' \mathbf{b}_k - \mathbf{b}_k^*) \mathbf{b}_k') \\ &= m \text{tr}(\mathbf{W}' \mathbf{U}' \Sigma^* (\mathbf{B}^* \mathbf{B}' \mathbf{B} \mathbf{B}' - \mathbf{B}^* \mathbf{B}')) \\ &= 0 \end{aligned}$$

where we used  $\mathbf{B} \mathbf{B}' = \mathbf{I}_r$ . Now we use the above and (16) to show that it also concentrates around zero. Let  $X_{ik} = \mathbf{a}'_{ik} \mathbf{W} \mathbf{b}_k$  and  $Y_{ik} = \mathbf{a}'_{ik} \mathbf{U}' \mathbf{p}_k$ . Both are sub-Gaussian, and so we can apply the Bernstein-type lemma, Lemma A.1, for sums of products of sub-Gaussian r.v.'s. Observe that  $K_{X_{ik}} = \|\mathbf{W} \mathbf{b}_k\| \leq \|\mathbf{W}\|_F \|\mathbf{b}_k\| \leq \|\mathbf{b}_k\|$ , and  $K_{Y_{ik}} \leq \|\mathbf{p}_k\|$ .

Using Lemma 3.3,  $\mathbf{b}_k$ 's are incoherent, i.e.,

$$\|\mathbf{b}_k\|^2 \leq \hat{\mu}^2 r / q = C \kappa^2 \mu^2 r / q.$$

w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ . Using the above, we can also show that  $\mathbf{p}_k$ 's are incoherent as follows. Using  $(a+b)^2 \leq 2(a^2 + b^2)$ , and  $\|\mathbf{B}\| = \|\mathbf{B}^*\| = 1$ ,

$$\begin{aligned} \|\mathbf{p}_k\|^2 &\leq 2\sigma_{\max}^* (\|\mathbf{B}^*\|^2 \|\mathbf{B}\|^2 \|\mathbf{b}_k\|^2 + \|\mathbf{b}_k^*\|^2) \\ &= 2\sigma_{\max}^* (\|\mathbf{b}_k\|^2 + \|\mathbf{b}_k^*\|^2) \\ &\leq 2\sigma_{\max}^* (\hat{\mu}^2 + \mu^2) r / q = C \kappa^2 \mu^2 \sigma_{\max}^* r / q \end{aligned}$$

We will now apply Lemma A.1 with  $t = m \delta_t^2 \|\mathbf{X}^*\|_F$ .

$$\begin{aligned} \frac{t^2}{\sum_{ik} K_{X_{ik}}^2 K_{Y_{ik}}^2} &= \frac{m^2 \delta_t^4 \|\mathbf{X}^*\|_F^2}{m \sum_k \|\mathbf{b}_k\|^2 \|\mathbf{p}_k\|^2} \\ &\geq \frac{m \delta_t^4 \|\mathbf{X}^*\|_F^2}{\max_k \|\mathbf{b}_k\|^2 \sum_k \|\mathbf{p}_k\|^2} \\ &= \frac{m \delta_t^4 \|\mathbf{X}^*\|_F^2}{\max_k \|\mathbf{b}_k\|^2 \|\mathbf{P}\|_F^2} \\ &\geq \frac{m \delta_t^4 \|\mathbf{X}^*\|_F^2}{\max_k \|\mathbf{b}_k\|^2 C^2 \delta_t^2 \|\mathbf{X}^*\|_F^2} \\ &\geq \frac{mq \delta_t^2}{C \kappa^2 \mu^2 r}, \text{ and} \\ \frac{t}{\max_{ik} K_{X_{ik}} K_{Y_{ik}}} &= \frac{m \delta_t^2 \|\mathbf{X}^*\|_F}{\max_k \|\mathbf{b}_k\| \|\mathbf{p}_k\|} \\ &\geq \frac{mq \delta_t^2 \sqrt{r} \sigma_{\min}^*}{\sigma_{\max}^* C^2 \kappa^2 \mu^2 r} \\ &= \frac{mq \delta_t^2 \sqrt{r}}{C \kappa^3 \mu^2 r} \end{aligned}$$

The second inequality used (16), while the third used incoherence of  $\mathbf{b}_k$ 's. The last inequality used incoherence of  $\mathbf{b}_k$ 's and of  $\mathbf{p}_k$ 's (proved above) and  $\|\mathbf{X}^*\|_F \geq \sqrt{r} \sigma_{\min}^*$ . Thus

$$\min \left( \frac{t^2}{\sum_{ik} K_{X_{ik}}^2 K_{Y_{ik}}^2}, \frac{t}{\max_{ik} K_{X_{ik}} K_{Y_{ik}}} \right) \geq \frac{mq \delta_t^2}{C \kappa^3 \mu^2 r}$$

and so

$$\Pr\{|\text{Term1}(\mathbf{W})| \leq m \delta_t^2 \|\mathbf{X}^*\|_F\} \geq 1 - \exp \left( -c \frac{mq \delta_t^2}{\kappa^3 \mu^2 r} \right)$$

Now we just need to extend our bound for all  $\mathbf{W} \in \mathcal{S}_W$ . We first extend it to all  $\mathbf{W}$  in an epsilon-net of  $\mathcal{S}_W$ . By [54](Lemma 5.2), there is an  $\epsilon$ -net,  $\bar{\mathcal{S}}_W$  so that for any  $\mathbf{W}$  in  $\mathcal{S}_W$ , there is a  $\bar{\mathbf{W}}$  in  $\bar{\mathcal{S}}_W$ , such that  $\|\bar{\mathbf{W}} - \mathbf{W}\|_F \leq \epsilon$  and  $|\bar{\mathcal{S}}_W| \leq (1 + \frac{2}{\epsilon})^{nr}$ . Pick  $\epsilon = 1/8$ . With this,  $|\bar{\mathcal{S}}_W| \leq 17^{nr}$ . Define  $\Delta \mathbf{W} := \bar{\mathbf{W}} - \mathbf{W}$ . We have  $\|\Delta \mathbf{W}\|_F \leq \epsilon = 1/8$ . Using union bound on the set  $\bar{\mathcal{S}}_W$ ,

$$\begin{aligned} &\Pr\{|\text{Term1}(\mathbf{W})| \leq m \delta_t^2 \|\mathbf{X}^*\|_F \text{ for all } \bar{\mathbf{W}} \in \bar{\mathcal{S}}_W\} \\ &\geq 1 - 2|\bar{\mathcal{S}}_W| \exp \left( -c \frac{mq \delta_t^2}{\kappa^3 \mu^2 r} \right) \\ &\geq 1 - 2 \exp \left( nr(\log 17) - c \frac{mq \delta_t^2}{\kappa^3 \mu^2 r} \right) \end{aligned} \quad (17)$$

To extend the claim to all  $\mathbf{W} \in \mathcal{S}_W$ , define

$$\theta_W := \max_{\mathbf{W} \in \mathcal{S}_W} \sum_{ik} (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k) (\mathbf{a}'_{ik} \mathbf{U}' \mathbf{p}_k).$$

$$\begin{aligned} \text{Since } \frac{\Delta \mathbf{W}}{\|\Delta \mathbf{W}\|_F} \in \mathcal{S}_W, \sum_{ik} (\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k) (\mathbf{a}'_{ik} \mathbf{U}' \mathbf{p}_k) &\leq \theta_W \|\Delta \mathbf{W}\|_F \leq \theta_W \epsilon. \text{ Thus, using (17), for any } \mathbf{W} \in \mathcal{S}_W, \\ &\sum_{ik} (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k) (\mathbf{a}_{ik}' \mathbf{U}' \mathbf{p}_k) \\ &= \sum_{ik} (\mathbf{a}_{ik}' \bar{\mathbf{W}} \mathbf{b}_k) (\mathbf{a}_{ik}' \mathbf{U}' \mathbf{p}_k) + \sum_{ik} (\mathbf{a}_{ik}' \Delta \mathbf{W} \mathbf{b}_k) (\mathbf{a}'_{ik} \mathbf{U}' \mathbf{p}_k) \\ &\leq m \delta_t^2 \|\mathbf{X}^*\|_F + \theta_W \epsilon = m \delta_t^2 \|\mathbf{X}^*\|_F + (1/8) \theta_W \end{aligned}$$

w.p. at least  $1 - 2 \exp\left(nr(\log 17) - c \frac{mq\delta_t^2}{\kappa^3 \mu^2 r}\right)$ . Thus, taking the  $\max_{\mathbf{W} \in \mathcal{S}_W}$  of the above equation and solving for  $\theta_W$ , with the above probability,

$$\theta_W \leq m\delta_t^2 \|\mathbf{X}^*\|_F / (1 - \epsilon) = (8/7)m\delta_t^2 \|\mathbf{X}^*\|_F.$$

All of the above bounds hold on the event in which  $\mathbf{B}$  is  $\hat{\mu}$  incoherent. This holds w.p. at least  $1 - n^{-10}$  as long as  $m \geq C \max(r, \log n, \log q) / \delta_b^2$  (this follows by Lemma 3.3).

Thus, if  $m \geq C \max(r, \log n, \log q) / \delta_b^2$ , w.p. at least  $1 - 2 \exp\left(nr(\log 17) - c \frac{mq\delta_t^2}{\kappa^3 \mu^2 r}\right) - n^{-10}$ ,  $\max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term1}(\mathbf{W})| \leq (8/7)m\delta_t^2 \|\mathbf{X}^*\|_F$ .  $\square$

*Proof of Lemma 3.12.* Recall that  $\mathbf{a}_{ik} := \mathbf{a}_{ik}^{(T+t)}$  and same for  $\mathbf{y}_{ik}$ . Thus, these are independent of the current  $\hat{\mathbf{x}}_k$ 's.

By Cauchy-Schwarz,

$$\begin{aligned} \text{Term2}(\mathbf{W}) &:= \sum_{ik} (\mathbf{c}_{ik} \hat{\mathbf{c}}_{ik} - 1) (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k) (\mathbf{a}_{ik}' \mathbf{x}_k^*) \\ &\leq \sqrt{\sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2} \sqrt{\sum_{ik} |\mathbf{c}_{ik} \hat{\mathbf{c}}_{ik} - 1|^2 |\mathbf{a}_{ik}' \mathbf{x}_k^*|^2} \end{aligned} \quad (18)$$

We can bound the first term using Lemma 3.10. Consider the second term. Since  $\mathbf{c}_{ik} = \text{sign}(\mathbf{a}_{ik}' \mathbf{x}_k^*)$  and  $\hat{\mathbf{c}}_{ik} = \text{sign}(\mathbf{a}_{ik}' \hat{\mathbf{x}}_k)$ , clearly  $(\mathbf{c}_{ik} \hat{\mathbf{c}}_{ik} - 1)^2 = (4) \mathbb{1}_{\{\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}\}}$ . To bound this term we use the following result.

**Lemma A.2** (Lemma 1 of [10]). *Let  $\mathbf{a}_i$  be standard Gaussian random vectors. For any given  $\mathbf{x}^*$ , and  $\hat{\mathbf{x}}$  independent from  $\mathbf{a}_i, i = 1, \dots, m$ ,*

$$\begin{aligned} \Pr\left(\text{sign}(\mathbf{a}_i' \mathbf{x}^*) \neq \text{sign}(\mathbf{a}_i' \hat{\mathbf{x}}) \mid (\mathbf{a}_i' \mathbf{x}^*)^2 = z^2, \hat{\mathbf{x}}\right) \\ \leq \text{erfc}\left(\frac{z}{2\|\mathbf{x}^* - \hat{\mathbf{x}}\|}\right), \end{aligned}$$

for all  $\hat{\mathbf{x}}$  that satisfy  $\text{dist}(\hat{\mathbf{x}}, \mathbf{x}^*) \leq 0.4$ . Here  $\text{erfc}(u) := \frac{2}{\sqrt{\pi}} \int_u^\infty \exp(-\tau^2) d\tau$  is the complementary error function.

Let  $Q_{ik} := \mathbb{1}_{\{\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}\}} \cdot (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2$  and let  $Z_{ik} := \mathbf{a}_{ik}' \mathbf{x}_k^*$ . Recall from Sec. A-C that  $\text{dist}(\mathbf{x}_k, \mathbf{x}_k^*) = \|\mathbf{h}_k\|$  with  $\mathbf{h}_k = \mathbf{x}_k^* - \hat{\mathbf{x}}_k$ .

We first upper bound  $\mathbb{E}[Q_{ik} | \hat{\mathbf{x}}_k]$ . For simplicity, we remove the subscripts  $i$  and  $k$  wherever these are not needed. Consider  $\mathbb{E}[Q] = \mathbb{E}[\mathbb{1}_{\mathbf{c} \neq \hat{\mathbf{c}}} Z^2]$ . Observe that  $\mathbf{c} = \text{sign}(Z)$  is a function of  $Z$  and  $Z$  depends on  $\mathbf{a}$ . Also  $\hat{\mathbf{c}}$  is a function of  $\mathbf{a}$ . Thus both of  $\mathbf{c}, \hat{\mathbf{c}}$  are dependent on  $Z$ . Moreover  $\hat{\mathbf{c}}$  also depends on  $\hat{\mathbf{x}}$ . We first bound  $\mathbb{E}[Q | Z^2, \hat{\mathbf{x}}]$  using the lemma stated above. For any  $\hat{\mathbf{x}}$  that satisfies  $\text{dist}(\hat{\mathbf{x}}, \mathbf{x}^*) \leq 0.4$ ,

$$\begin{aligned} \mathbb{E}[Q | Z^2 = z^2, \hat{\mathbf{x}}] &= \mathbb{E}[\mathbb{1}_{\mathbf{c} \neq \hat{\mathbf{c}}} z^2 | Z^2 = z^2, \hat{\mathbf{x}}] \\ &= z^2 \Pr(\mathbf{c} \neq \hat{\mathbf{c}} | Z^2 = z^2, \hat{\mathbf{x}}) \\ &\leq z^2 \text{erfc}\left(\frac{z}{2\|\mathbf{x}^* - \hat{\mathbf{x}}\|}\right) \\ &\leq z^2 \exp\left(-\frac{z^2}{4\|\mathbf{x}^* - \hat{\mathbf{x}}\|^2}\right). \end{aligned}$$

The first inequality follows using Lemma A.2, the second is a standard upper bound on the erfc function [69]. Thus, for any  $\hat{\mathbf{x}}$  that satisfies  $\text{dist}(\hat{\mathbf{x}}, \mathbf{x}^*) \leq 0.4$ ,

$$\mathbb{E}[Q | \hat{\mathbf{x}}] = \mathbb{E}[\mathbb{E}[Q | Z^2, \hat{\mathbf{x}}]] \leq \mathbb{E}\left[Z^2 \exp\left(-\frac{Z^2}{4\|\mathbf{x}^* - \hat{\mathbf{x}}\|^2}\right) | \hat{\mathbf{x}}\right].$$

Since  $Z$  is zero mean Gaussian with variance  $\|\mathbf{x}^*\|^2$ ,  $Y := Z^2 / \|\mathbf{x}^*\|^2$  is standard chi-squared with one degree of freedom. Thus,  $\mathbb{E}[Y] = 1$ . Using this and  $\exp(-y/2) < 1$ , we get

$$\begin{aligned} \mathbb{E}[Q | \hat{\mathbf{x}}] &\leq \int_0^\infty y \|\mathbf{x}^*\|^2 \exp\left(-\frac{y\|\mathbf{x}^*\|^2}{4\|\mathbf{x}^* - \hat{\mathbf{x}}\|^2}\right) \frac{\exp(-y/2)}{\sqrt{2y\Gamma(1/2)}} dy \\ &\leq \int_0^\infty y \|\mathbf{x}^*\|^2 \exp\left(-\frac{y\|\mathbf{x}^*\|^2}{4\|\mathbf{x}^* - \hat{\mathbf{x}}\|^2}\right) \frac{1}{\sqrt{2y\Gamma(1/2)}} dy \\ &= 2\sqrt{2} \frac{\|\mathbf{x}^* - \hat{\mathbf{x}}\|^3}{\|\mathbf{x}^*\|} = 2\sqrt{2} \frac{\text{dist}(\mathbf{x}^*, \hat{\mathbf{x}})^3}{\|\mathbf{x}^*\|} \end{aligned}$$

where  $\Gamma(1/2)$  is the Gamma function evaluated at 1/2 (can treat it as a constant). We will now use Lemma 3.2 to average over  $\hat{\mathbf{x}}$ . Let  $E$  be the event that  $\text{dist}(\mathbf{x}_k^*, \hat{\mathbf{x}}_k) \leq \delta_t \|\mathbf{x}_k^*\|$  for all  $k = 1, 2, \dots, q$ . By Lemma 3.2, under the lower bound on  $m$ , this event occurs w.p. at least  $1 - n^{-10}$ . On the complement event, we do not have any tight bounds on  $Q$ . However, trivially,  $Q \leq (\mathbf{a}' \mathbf{x}^*)^2$  always and so, for any  $\hat{\mathbf{x}}$ ,  $\mathbb{E}[Q | \hat{\mathbf{x}}] \leq \|\mathbf{x}^*\|^2$ . Thus

$$\mathbb{E}[Q] \leq 4\delta_t^3 \|\mathbf{x}^*\|^2 \cdot (1 - n^{-10}) + \|\mathbf{x}^*\|^2 \cdot n^{-10}$$

We will eventually set  $\delta_t = c/(\kappa^2 r)$  which is larger than  $n^{-10}$ . Thus,

$$\mathbb{E}[Q] \leq C \max(\delta_t^3, n^{-10}) \|\mathbf{x}^*\|^2 = C\delta_t^3 \|\mathbf{x}^*\|^2$$

and so, if  $m \geq C \max(r, \log n, \log q)$ ,

$$\sum_{ik} \mathbb{E}[Q_{ik}] \leq m C \delta_t^3 \|\mathbf{X}^*\|_F^2. \quad (19)$$

As a side-note, we should clarify here that the dependence on  $\hat{\mathbf{x}}$  matters in only the above expected value computation because this is the only term where we have upper bounded the expectation using Lemma 3.2. The expected value of the other two terms is the same for all values of  $\hat{\mathbf{x}}$  and hence we ignore the dependence there. Everywhere else, we use this lemma only while obtaining the high probability error bounds, and of course we assume the bounds hold on the intersection of the desired event with  $E$ .

Next we show that, on the event  $E$ , whp,  $\sum_{ik} Q_{ik}$  is of the same order. As shown in the proof of Theorem 1 of [10],  $\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}$  implies that<sup>8</sup>  $(\mathbf{a}_{ik}' \mathbf{x}_k^*)^2 \leq (\mathbf{a}_{ik}' \mathbf{h}_k)^2$ . Here  $\mathbf{h}_k = \mathbf{x}_k^* - \hat{\mathbf{x}}_k$ . Thus,  $Q_{ik} = \mathbb{1}_{\{\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}\}} (\mathbf{a}_{ik}' \mathbf{x}_k^*)^2 \leq \mathbb{1}_{\{\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}\}} (\mathbf{a}_{ik}' \mathbf{h}_k)^2 \leq (\mathbf{a}_{ik}' \mathbf{h}_k)^2$ . Thus, it is a sub-exponential r.v., or equivalently it is a product of sub-Gaussian r.v.'s  $\sqrt{Q_{ik}}$ . Thus, we can

<sup>8</sup>This follows from  $(\mathbf{a}' \mathbf{x}^*)^2 = (\mathbf{a}' \mathbf{h})^2 + (\mathbf{a}' \hat{\mathbf{x}})^2 + 2(\mathbf{a}' \mathbf{h})(\mathbf{a}' \hat{\mathbf{x}}) = (\mathbf{a}' \mathbf{h})^2 - (\mathbf{a}' \hat{\mathbf{x}})^2 + 2(\mathbf{a}' \mathbf{x}^*)(\mathbf{a}' \hat{\mathbf{x}}) \leq (\mathbf{a}' \mathbf{h})^2$ . The inequality holds because  $\mathbf{c}_{ik} \neq \hat{\mathbf{c}}_{ik}$  means that the last term is negative.

apply Lemma A.1 with  $K_{X_{ik}} = K_{Y_{ik}} = \|\mathbf{h}_k\|$ . Pick  $t = m\delta_t\delta_t^2\|\mathbf{X}^*\|_F^2$ . Observe that

$$\begin{aligned} \frac{t^2}{\sum_{ik} \|\mathbf{h}_k\|^4} &= \frac{m^2\delta_t^2(\delta_t^2\|\mathbf{X}^*\|_F^2)^2}{m\sum_k \|\mathbf{h}_k\|^4} \\ &\geq \frac{m\delta_t^6\|\mathbf{X}^*\|_F^4}{\max_k \|\mathbf{h}_k\|^2\delta_t^2\|\mathbf{X}^*\|_F^2} \\ &\geq \frac{m\delta_t^2\|\mathbf{X}^*\|_F^2}{\max_k \|\mathbf{x}_k^*\|^2} \\ &\geq \frac{mq\delta_t^2}{\kappa^2\mu^2}, \text{ and} \\ \frac{t}{\max_k \|\mathbf{h}_k\|^2} &\geq \frac{m\delta_t\|\mathbf{X}^*\|_F^2}{\max_k \|\mathbf{x}_k^*\|^2} \geq \frac{\delta_t mq}{\kappa^2\mu^2}. \end{aligned}$$

The second inequality used  $\|\mathbf{h}_k\|^2 \leq \delta_t^2\|\mathbf{x}_k^*\|^2$  which holds on the event  $E$ . The third inequality used the fact that right incoherence implies  $\|\mathbf{x}_k^*\|^2 \leq \mu^2\kappa^2\|\mathbf{X}^*\|_F^2/q$ . Thus, on the event  $E$ ,

$$\begin{aligned} \Pr \{ |\sum_{ik} Q_{ik} - \sum_{ik} \mathbb{E}[Q_{ik}]| \geq m\delta_t\delta_t^2\|\mathbf{X}^*\|_F^2 \} \\ \leq 2 \exp(-c\delta_t^2 mq/\kappa^2\mu^2) \end{aligned}$$

Thus, using (19), and  $\Pr(E) > 1 - n^{-10}$  (by Lemma 3.2), if  $m \geq C \max(r, \log n, \log q)$ , w.p. at least  $1 - n^{-10} - 2 \exp(-c\delta_t^2 mq/\kappa^2\mu^2)$ ,

$$\sum_{ik} Q_{ik} \leq (C+1)m\delta_t\delta_t^2\|\mathbf{X}^*\|_F^2. \quad (20)$$

Finally, combining (18), (20), and Lemma 3.10, if  $m \geq C \max(r, \log n, \log q)/\delta_b^2$ , w.p. at least  $1 - n^{-10} - 2 \exp(nr - c\delta_t^2 \frac{mq}{\mu^2 r}) - 2 \exp(-c\delta_t^2 \frac{mq}{\kappa^2\mu^2})$ ,

$$\max_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \text{Term2}(\mathbf{W}) \leq Cm\sqrt{1 + \delta_t}\sqrt{\delta_t}\delta_t\|\mathbf{X}^*\|_F.$$

□

*Proof of Lemma 3.13.* Using facts from Sec. A-A, and using  $\sigma_i(\mathbf{U}^*) = 1$ ,

$$\begin{aligned} \sigma_{\min}(\mathbf{U}^*\Sigma^*\mathbf{B}^*\mathbf{B}') &\geq \sigma_{\min}(\mathbf{U}^*)\sigma_{\min}(\Sigma^*)\sigma_{\min}(\mathbf{B}^*\mathbf{B}') \\ &\geq \sigma_{\min}^* \sqrt{1 - \sin \Theta(\mathbf{B}^{*\prime}, \mathbf{B}')} \end{aligned}$$

To upper bound  $\sin \Theta(\mathbf{B}^{*\prime}, \mathbf{B}')$ , first notice that  $\mathbf{B}^{*\prime}$  and  $\mathbf{B}'$  are basis matrices. Thus,  $\sin \Theta(\mathbf{B}^{*\prime}, \mathbf{B}') = \|\mathbf{B}^*(\mathbf{I} - \mathbf{B}'\mathbf{B})\|$ . We have upper bounded  $\|\tilde{\mathbf{B}}^*(\mathbf{I} - \mathbf{B}'\mathbf{B})\|$  in Lemma 3.11. Also recall that  $\tilde{\mathbf{B}}^* = \Sigma^*\mathbf{B}^*$  and  $\|\tilde{\mathbf{B}}^*\|_F = \|\mathbf{X}^*\|_F \leq \sqrt{r}\sigma_{\max}^*$ . Thus,

$$\begin{aligned} \sin \Theta(\mathbf{B}^{*\prime}, \mathbf{B}') &= \|(\Sigma^*)^{-1}\tilde{\mathbf{B}}^*(\mathbf{I} - \mathbf{B}'\mathbf{B})\| \leq \frac{C\delta_t\|\mathbf{X}^*\|_F}{\sigma_{\min}^*} \\ &\leq C\sqrt{r}\kappa\delta_t \end{aligned}$$

Thus, using  $\delta_t \leq 0.7/C\sqrt{r}\kappa$ , we will have  $\sigma_{\min}(\mathbf{U}^*\Sigma^*\mathbf{B}^*\mathbf{B}') \geq 0.95\sigma_{\min}^*$ . □

## APPENDIX B PROOF OF LEMMA 3.9

We begin by defining a few quantities needed for writing an expression of  $\hat{\mathbf{U}}^{t+1}$  in closed form.

- 1) We use the subscript  $M_{vec}$  to refer to the vectorized version of matrix  $M$ . Thus, for example,  $\mathbf{U}_{vec}^*$  is an  $nr \times 1$  vectorized version of the basis matrix  $\mathbf{U}^*$ . When updating its estimate by LS, we first obtain an expression for the vectorized version  $\hat{\mathbf{U}}_{vec}^{t+1}$  and then rearrange it as an  $n \times r$  matrix.
- 2) Scalars such as  $b(q)$  refer to the  $q$ -th entry of vector  $\mathbf{b}$ .
- 3) Define the diagonal matrix  $\mathbf{C}_k := \text{diag}(\mathbf{c}_{ik})$ , and recall from the algorithm that  $\hat{\mathbf{C}}_k = \text{diag}(\hat{\mathbf{c}}_{ik})$ . Here  $\mathbf{c}_{ik} := \text{sign}(\mathbf{a}_{ik}'\mathbf{x}_k^*)$  and  $\hat{\mathbf{c}}_{ik} := \text{sign}(\mathbf{a}_{ik}'\hat{\mathbf{x}}_k)$  as defined in Lemma 3.9.
- 4) For  $k = 1, \dots, q$ , define

$$\begin{aligned} \mathbf{B}_{k,mat} &:= [b_k(1)\mathbf{I}_n, \dots, b_k(r)\mathbf{I}_n]' \\ \mathbf{d} &:= \sum_k \mathbf{B}_{k,mat} \mathbf{A}_k \hat{\mathbf{C}}_k \mathbf{y}_k \\ \mathbf{M} &:= \sum_k \mathbf{B}_{k,mat} \mathbf{A}_k \mathbf{A}_k' \mathbf{B}_{k,mat}' \end{aligned}$$

Observe that  $\mathbf{B}_{k,mat}$  and  $\mathbf{M}$  are  $nr \times n$  and  $nr \times nr$  matrices respectively while  $\mathbf{d}$  is an  $nr \times 1$  vector. Moreover, for any  $\mathbf{w} \in \mathbb{R}^{nr \times 1}$  it is easy to see that  $\mathbf{w}'\mathbf{B}_{k,mat} = \mathbf{b}_k'\mathbf{W}'$  where  $\mathbf{W}$  is matrix version of  $\mathbf{w}$  with  $\mathbf{W} \in \mathbb{R}^{n \times r}$ .

Recall from the algorithm that

$$\hat{\mathbf{U}}_{vec}^{t+1} = \underset{\hat{\mathbf{U}}_{vec}}{\text{argmin}} \sum_k \|\hat{\mathbf{C}}_k \mathbf{y}_k - \mathbf{A}_k' \mathbf{B}_{k,mat}' \hat{\mathbf{U}}_{vec}\|^2.$$

This is an LS problem, it can be solved in closed form to give

$$\hat{\mathbf{U}}_{vec}^{t+1} = \mathbf{M}^{-1} \mathbf{d}$$

which is  $nr$ -length vector. We get the matrix  $\hat{\mathbf{U}}^{t+1}$  by reshaping this vector into an  $n \times r$  matrix.

To simplify the above expression, first recall that  $\mathbf{y}_k = |\mathbf{A}_k' \mathbf{x}_k^*|$  and  $\mathbf{C}_k$  is the diagonal matrix containing the signs of  $(\mathbf{A}_k' \mathbf{x}_k^*)_i$ : its  $i, i$ -th entry is the sign of  $(\mathbf{A}_k' \mathbf{x}_k^*)_i$ . Thus,

$$\mathbf{y}_k = \mathbf{C}_k \mathbf{A}_k' \mathbf{x}_k^*.$$

Since  $\mathbf{x}_k^* = \mathbf{U}^* \Sigma^* \mathbf{b}_k^*$ , we can rewrite  $\mathbf{d}$  as

$$\mathbf{d} = \sum_k \mathbf{B}_{k,mat} \mathbf{A}_k \hat{\mathbf{C}}_k \mathbf{C}_k \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^*.$$

Before proceeding further, we define a few more quantities.

- 1) Define

$$\mathbf{S} = \sum_k \mathbf{B}_{k,mat} \mathbf{B}_{k,mat}'.$$

- 2) Define the “expanded” singular value matrix which is of size  $nr \times nr$ ,

$$\Sigma_{big}^* := \text{diag}(\sigma_1^* \mathbf{I}_n, \dots, \sigma_r^* \mathbf{I}_n)$$

where  $\sigma_i^*$  are the singular values of  $\mathbf{X}^*$ .

3) In order to separate the contribution of phase error from the rest, split  $\mathbf{d}$  as  $\mathbf{d} = \mathbf{d}^{(1)} + \mathbf{d}^{(2)}$  where

$$\begin{aligned}\mathbf{d}^{(1)} &= \sum_k \mathbf{B}_{k,mat} \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ \mathbf{d}^{(2)} &= \sum_k \mathbf{B}_{k,mat} (\hat{\mathbf{C}}_k \mathbf{C}_k - \mathbf{I}) \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^*,\end{aligned}$$

Thus,

$$\hat{\mathbf{U}}_{vec}^{t+1} = \mathbf{M}^{-1} \mathbf{d} = \mathbf{M}^{-1} (\mathbf{d}^{(1)} + \mathbf{d}^{(2)}).$$

4) Define the  $nr$ -length vector  $\mathbf{F}_{vec}$  as follows

$$\mathbf{F}_{vec} = \mathbf{M}^{-1} (\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{d}^{(1)}) - \mathbf{M}^{-1} \mathbf{d}^{(2)}.$$

and let  $\mathbf{F} \in \mathbb{R}^{n \times r}$  be the reshaped matrix formed from  $\mathbf{F}_{vec}$ .

We will now show that

$$\hat{\mathbf{U}}^{t+1} = \mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}' - \mathbf{F}. \quad (21)$$

This will be useful because when we try to bound  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^{t+1})$ , the first term will disappear. To do this, we add and subtract the vector  $\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^*$  from  $\mathbf{d}^{(1)}$ . This gives

$$\hat{\mathbf{U}}_{vec}^{t+1} = \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{F}_{vec}$$

Next we explain why the  $n \times r$  reshaped matrix version of the vector  $\mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^*$  equals  $\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}'$ . We have

$$\begin{aligned}\mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* &= \sum_k \mathbf{B}_{k,mat} \mathbf{B}_{k,mat}' \Sigma_{big}^* \mathbf{U}_{vec}^* \\ &= \sum_k \mathbf{B}_{k,mat} \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ &= \sum_k \begin{bmatrix} b_k(1) \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ b_k(2) \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ \vdots \\ b_k(r) \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \end{bmatrix}\end{aligned}$$

Matrix version of the above vector has  $p$ -th column as  $\mathbf{U}^* \Sigma^* (\sum_k \mathbf{b}_k^* b_k(p))$ . This implies that matrix version of this vector is  $\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}'$ . Moreover, it is easy to see that  $\mathbf{B}_{k,mat}' \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* = \mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}' \mathbf{b}_k$ .

In the rest of this section, we use (21) to obtain the desired bound on  $\sin \Theta(\mathbf{U}^*, \mathbf{U}^{t+1})$ . Recall that  $\hat{\mathbf{U}}_{t+1} \stackrel{\text{QR}}{=} \mathbf{U}^{t+1} \mathbf{R}_U$ . Thus,  $\mathbf{U}^{t+1} = \hat{\mathbf{U}}^{t+1} (\mathbf{R}_U)^{-1} = (\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}' - \mathbf{F}) (\mathbf{R}_U)^{-1}$  and so

$$\begin{aligned}\sin \Theta(\mathbf{U}^*, \mathbf{U}^{t+1}) &= \|\mathbf{U}_\perp^* \mathbf{F} (\mathbf{R}_U)^{-1}\| \leq \|\mathbf{F} (\mathbf{R}_U)^{-1}\| \\ &\leq \|\mathbf{F}\|_F \|\mathbf{R}_U^{-1}\| = \frac{\|\mathbf{F}_{vec}\|}{\sigma_{\min}(\mathbf{R}_U)}\end{aligned} \quad (22)$$

Since  $\sigma_{\min}(\mathbf{R}_U) = \sigma_{\min}(\mathbf{U}^{t+1})$ , we have  $\sigma_{\min}(\mathbf{R}_U) = \sigma_{\min}(\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}' - \mathbf{F}) \geq \sigma_{\min}(\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}') - \|\mathbf{F}\| \geq \sigma_{\min}(\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}') - \|\mathbf{F}_{vec}\|$ . Thus,

$$\sin \Theta(\mathbf{U}^*, \mathbf{U}^{t+1}) \leq \frac{\|\mathbf{F}_{vec}\|}{\sigma_{\min}(\mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}') - \|\mathbf{F}_{vec}\|} \quad (23)$$

In the rest of this proof, we show that  $\|\mathbf{F}_{vec}\|$  is upper bounded by MainTerm. We have

$$\|\mathbf{F}_{vec}\| \leq \|\mathbf{M}^{-1}\| \left( \|\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{d}^{(1)}\| + \|\mathbf{d}^{(2)}\| \right) \quad (24)$$

Consider the first term,  $\mathbf{M}^{-1}$ . Since  $\mathbf{M}$  is a symmetric positive semidefinite matrix

$$\sigma_{\min}(\mathbf{M}) = \min_{\mathbf{w} \in \mathbb{R}^{nr \times 1}, \|\mathbf{w}\|=1} \mathbf{w}' \mathbf{M} \mathbf{w}$$

For all  $\mathbf{w} \in \mathbb{R}^{nr \times 1}, \|\mathbf{w}\|^2 = 1$ , we can write  $\mathbf{w}' \mathbf{M} \mathbf{w}$

$$\begin{aligned}\sum_k \mathbf{w}' \mathbf{B}_{k,mat} \mathbf{A}_k \mathbf{A}_k' \mathbf{B}_{k,mat}' \mathbf{w} &= \sum_k \mathbf{b}_k' \mathbf{W}' \mathbf{A}_k \mathbf{A}_k' \mathbf{W} \mathbf{b}_k \\ &= \sum_{ik} |\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k|^2 = \text{Term3}(\mathbf{W})\end{aligned}$$

where  $\mathbf{W} \in \mathbb{R}^{n \times r}$  is the matrix version of  $\mathbf{w}$  and  $\mathbf{w} = \mathbf{W}_{vec}$ . Recall that  $\mathcal{S}_W = \{\mathbf{W} \in \mathbb{R}^{n \times r}, \|\mathbf{W}\|_F = 1\} = \{\mathbf{w} \in \mathbb{R}^{nr \times 1}, \|\mathbf{w}\| = 1\}$ . Thus,

$$\|\mathbf{M}^{-1}\| = \frac{1}{\sigma_{\min}(\mathbf{M})} = \frac{1}{\min_{\mathbf{W} \in \mathcal{S}_W} |\text{Term3}(\mathbf{W})|} \quad (25)$$

Now consider the first term inside the parenthesis. Using the variational definition,

$$\begin{aligned}\|\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{d}^{(1)}\| &= \max_{\mathbf{w} \in \mathbb{R}^{nr \times 1}, \|\mathbf{w}\|=1} |\mathbf{w}' (\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{d}^{(1)})|.\end{aligned}$$

It follows from definitions that

$$\begin{aligned}\mathbf{w}' \mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* &= \sum_k (\mathbf{w}' \mathbf{B}_{k,mat}) \mathbf{A}_k \mathbf{A}_k' (\mathbf{B}_{k,mat}' \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^*) = \\ &\sum_k \mathbf{b}_k' \mathbf{W}' \mathbf{A}_k \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{B}^* \mathbf{B}' \mathbf{b}_k.\end{aligned}$$

Similarly

$$\mathbf{w}' \mathbf{d}^{(1)} = \sum_k \mathbf{b}_k' \mathbf{W}' \mathbf{A}_k \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^*,$$

and thus

$$\|\mathbf{M} \mathbf{S} \Sigma_{big}^* \mathbf{U}_{vec}^* - \mathbf{d}^{(1)}\| = \max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term1}(\mathbf{W})|. \quad (26)$$

For the final term,  $\|\mathbf{d}^{(2)}\|$ , by variational definition we have,

$$\|\mathbf{d}^{(2)}\| = \max_{\mathbf{w} \in \mathbb{R}^{nr \times 1}, \|\mathbf{w}\|=1} \mathbf{w}' \mathbf{d}^{(2)}.$$

From definitions we know that

$$\begin{aligned}\mathbf{w}' \mathbf{d}^{(2)} &= \sum_k \mathbf{w}' \mathbf{B}_{k,mat} \mathbf{A}_k (\hat{\mathbf{C}}_k \mathbf{C}_k - \mathbf{I}) \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ &= \sum_k \mathbf{b}_k' \mathbf{W}' \mathbf{A}_k (\hat{\mathbf{C}}_k \mathbf{C}_k - \mathbf{I}) \mathbf{A}_k' \mathbf{U}^* \Sigma^* \mathbf{b}_k^* \\ &= \sum_{ik} (\hat{\mathbf{c}}_{ik} \mathbf{c}_{ik} - 1) (\mathbf{a}_{ik}' \mathbf{W} \mathbf{b}_k) (\mathbf{a}_{ik}' \mathbf{x}_k^*) = \text{Term2}(\mathbf{W}),\end{aligned}$$

and thus

$$\|\mathbf{d}^{(2)}\| = \max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term2}(\mathbf{W})|. \quad (27)$$

Combining (24) - (27),

$$\begin{aligned}\|\mathbf{F}_{vec}\| &\leq \frac{\max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term1}(\mathbf{W})| + \max_{\mathbf{W} \in \mathcal{S}_W} |\text{Term2}(\mathbf{W})|}{\min_{\mathbf{W} \in \mathcal{S}_W} |\text{Term3}(\mathbf{W})|}\end{aligned}$$

Combining the above bound with (23) proves the lemma.

## APPENDIX C PROOF SKETCH OF COROLLARY 5.1

Suppose first that the subspace change times  $k_j$  were known. By our assumption,  $k_{j+1} - k_j > T\alpha$ . Then the proof is almost exactly the same as that for the static case. The only difference is that, in the current case, every  $\alpha$  time instants, we are using measurements corresponding to a new set of  $\alpha$  signals (columns of  $\mathbf{X}^*_{\text{full}}$ ) but we use the estimate of the subspace obtained from the measurements for the previous  $\alpha$  time instants. As long as the subspace has not changed between the two intervals, Claims 3.1 and 3.4 apply without change. Combining them, we can again conclude that  $\sin \Theta(\mathbf{U}_{\text{sub},(j)}^0, \mathbf{U}_{\text{sub},(j)}^*) \leq \delta_{\text{init}}$  at  $k = k_j + \alpha$ , and that the bound decreases 0.7 times after each  $\alpha$ -length epoch so that  $\sin \Theta(\mathbf{U}_{\text{sub},(j)}^T, \mathbf{U}_{\text{sub},(j)}^*) \leq \epsilon$  at  $k = k_j + \alpha T$ . By our assumption,  $k_{j+1} > k_j + \alpha T$  so this happens before the next change.

The proof in the unknown  $k_j$  case follows if we can show that, whp,  $k_j \leq \hat{k}_j \leq \hat{k}_j + 2\alpha$ . This can be done using Lemma C.1 given below along with the following argument borrowed from [55], [70]. Consider the  $\alpha$ -length interval in which  $k_j$  lies. Assume that, before this interval, we have an  $\epsilon$ -accurate estimate of the previous subspace. In this interval, the first some data vectors satisfy  $\mathbf{x}_k^* = \mathbf{U}_{\text{sub},(j-1)}^* \tilde{\mathbf{d}}_k^*$ , while the rest satisfy  $\mathbf{x}_k^* = \mathbf{U}_{\text{sub},(j)}^* \tilde{\mathbf{d}}_k^*$ . By our assumption, this interval lies in the “detect phase”. We cannot guarantee whether the change will get detected in this interval, but it may. However, in the interval after this interval, all data vectors satisfy  $\mathbf{x}_k^* = \mathbf{U}_{\text{sub},(j)}^* \tilde{\mathbf{d}}_k^*$ . In this interval, Lemma C.1 given below can be used to show that the change gets detected whp. Thus, either the change is detected in the first interval itself (the one that contains  $k_j$ ), or it is not. If it is not, then, by Lemma C.1, whp, it will get detected in the second interval (in which all signals are generated from the  $j$ -th subspace). Thus,  $\hat{k}_j \leq k_j + 2\alpha$ . See Appendix A of [55] for a precise proof of this idea. The key point to note here is that we are never updating the subspace in the interval that contains  $k_j$  and hence we do not have to prove a new descent lemma that deals with the interval in which the subspace changes.

We will replace  $\alpha$  by  $q$  in the following lemma and its proof, in order to able to use bounds from earlier proofs. Thus in this lemma, we are considering a  $q$ -frame epoch.

**Lemma C.1.** Consider the  $(n-r) \times (n-r)$  matrix

$$\mathbf{Y}_{U,\text{det}} := \mathbf{U}_{\text{sub},(j-1),\perp}' \mathbf{Y}_U(\mathcal{J}_q) \mathbf{U}_{\text{sub},(j-1),\perp},$$

Assume that  $\|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \leq \frac{\delta_{\text{init}} \sigma_{\min}^*}{q} \frac{\sigma_{\min}^*}{q}$ . This is true by (10). Assume that  $\sin \Theta(\mathbf{U}_{\text{sub},(j-1)}, \mathbf{U}_{\text{sub},(j-1)}^*) \leq \epsilon$ . Then,

1) If  $\mathcal{J}_q \subseteq [k_j, k_{j+1})$  (change has occurred), then

$$\begin{aligned} & \lambda_{\max}(\mathbf{Y}_{U,\text{det}}) - \lambda_{\min}(\mathbf{Y}_{U,\text{det}}) \\ & \geq \frac{\sigma_{\min}^*}{q} \left( 1.5 \sin \Theta^2(\mathbf{U}_{\text{sub},(j)}^*, \mathbf{U}_{\text{sub},(j-1)}) - 2\delta_{\text{init}} \right) \\ & \geq \frac{\sigma_{\min}^*}{q} (1.5(\sin \Theta(\mathbf{U}_{\text{sub},(j)}^*, \mathbf{U}_{\text{sub},(j-1)}) - 2\epsilon)^2 - 2\delta_{\text{init}}) \end{aligned}$$

2) If  $\mathcal{J}_q \subseteq [\hat{k}_{j-1} + Tq, k_j)$  (change has not occurred), then

$$\begin{aligned} & \lambda_{\max}(\mathbf{Y}_{U,\text{det}}) - \lambda_{\min}(\mathbf{Y}_{U,\text{det}}) \\ & \leq \frac{1}{q} \sigma_{\max}^* \frac{\sigma_{\min}^*}{q} \sin \Theta^2(\mathbf{U}_{\text{sub},(j-1)}^*, \mathbf{U}_{\text{sub},(j-1)}) + \frac{2\delta_{\text{init}} \sigma_{\min}^*}{q} \\ & \leq \frac{\sigma_{\min}^*}{q} (\kappa^2 \epsilon^2) + 2\delta_{\text{init}} \end{aligned}$$

*Proof.* This proof uses the following fact: For basis matrices  $\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3$  of the same size,  $\sin \Theta(\mathbf{P}_1, \mathbf{P}_2) - 2\sin \Theta(\mathbf{P}_2, \mathbf{P}_3) \leq \sin \Theta(\mathbf{P}_1, \mathbf{P}_3) \leq \sin \Theta(\mathbf{P}_1, \mathbf{P}_2) + \sin \Theta(\mathbf{P}_2, \mathbf{P}_3)$ .

Define the  $(n-r) \times (n-r)$  matrix

$$\mathbf{E}_{\text{det}} = \mathbf{U}_{\text{sub},(j-1),\perp}' \mathbb{E}[\mathbf{Y}_-] \mathbf{U}_{\text{sub},(j-1),\perp}.$$

### Proof of item 1

$$\begin{aligned} \lambda_{\max}(\mathbf{Y}_{U,\text{det}}) & \geq \lambda_{\max}(\mathbf{E}_{\text{det}}) - \|\mathbf{Y}_{U,\text{det}} - \mathbf{E}_{\text{det}}\| \\ & \|\mathbf{Y}_{U,\text{det}} - \mathbf{E}_{\text{det}}\| \\ & \leq \|\mathbf{Y}_U - \mathbb{E}[\mathbf{Y}_-]\| \leq \frac{\delta_{\text{init}} \sigma_{\min}^*}{q}. \end{aligned}$$

Also we have

$$\lambda_{\min}(\mathbf{Y}_{U,\text{det}}) \leq \lambda_{\min}(\mathbf{E}_{\text{det}}) + \|\mathbf{Y}_{U,\text{det}} - \mathbf{E}_{\text{det}}\|$$

Thus using the facts from Sec. A-A and  $\min_k \beta_{1,k}^- \geq 1.5$  (proved while proving Claim 3.1 for initializing  $\mathbf{U}^*$ ),

$$\begin{aligned} & \lambda_{\max}(\mathbf{Y}_{U,\text{det}}) - \lambda_{\min}(\mathbf{Y}_{U,\text{det}}) \\ & \geq \lambda_{\max} \left( \frac{1}{q} \mathbf{U}_{\text{sub},(j-1),\perp}' \mathbf{U}_{\text{sub},(j)}^* \left( \sum_k \beta_{1,k}^- \tilde{\mathbf{b}}_k^* \tilde{\mathbf{b}}_k^{*\prime} \right) \times \right. \\ & \quad \left. \mathbf{U}_{\text{sub},(j)}^* \mathbf{U}_{\text{sub},(j-1),\perp} \right) - \frac{2\delta_{\text{init}} \sigma_{\min}^*}{q} \\ & \geq \frac{(\sigma_{\min}^*)^2}{q} \left( 1.5 \sin \Theta(\mathbf{U}_{\text{sub},(j)}^*, \mathbf{U}_{\text{sub},(j-1)})^2 - 2\delta_{\text{init}} \right). \end{aligned}$$

### Proof of item 2

$$\begin{aligned} \lambda_{\max}(\mathbf{Y}_{U,\text{det}}) - \lambda_{\min}(\mathbf{Y}_{U,\text{det}}) & \leq \lambda_{\max}(\mathbf{E}_{\text{det}}) - \lambda_{\min}(\mathbf{E}_{\text{det}}) \\ & + 2\|\mathbf{Y}_{U,\text{det}} - \mathbf{E}_{\text{det}}\|. \end{aligned}$$

It is easy to see that

$$\begin{aligned} & \lambda_{\max}(\mathbf{E}_{\text{det}}) - \lambda_{\min}(\mathbf{E}_{\text{det}}) \\ & \leq \frac{\max_k \beta_{1,k}^-}{q} \sigma_{\max}^* \frac{\sigma_{\min}^*}{q} \sin \Theta^2(\mathbf{U}_{\text{sub},(j-1)}^*, \mathbf{U}_{\text{sub},(j-1)}) \\ & \leq \frac{\sigma_{\max}^*}{q} \sin \Theta^2(\mathbf{U}_{\text{sub},(j-1)}^*, \mathbf{U}_{\text{sub},(j-1)}) \end{aligned}$$

□

## REFERENCES

- [1] S. Nayer, P. Narayananurthy, and N. Vaswani, “Phaseless pca: Low-rank matrix recovery from column-wise phaseless measurements,” in *Intnl. Conf. Machine Learning (ICML)*, 2019.
- [2] James R Fienup, “Phase retrieval algorithms: a comparison,” *Applied optics*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [3] R. W. Gerchberg and W. O. Saxton, “A practical algorithm for the determination of phase from image and diffraction plane pictures,” *Optik*, 1972.

[4] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski, “Phase retrieval via matrix completion,” *SIAM J. Imaging Sci.*, vol. 6, no. 1, pp. 199–225, 2013.

[5] E. J. Candes, T. Strohmer, and V. Voroninski, “Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming,” *Comm. Pure Appl. Math.*, 2013.

[6] P. Netrapalli, P. Jain, and S. Sanghavi, “Phase retrieval using alternating minimization,” in *NIPS*, 2013, pp. 2796–2804.

[7] Irène Waldspurger, “Phase retrieval with random gaussian sensing vectors by alternating projections,” *IEEE Transactions on Information Theory*, vol. 64, no. 5, pp. 3301–3312, 2018.

[8] E. J. Candes, X. Li, and M. Soltanolkotabi, “Phase retrieval via Wirtinger flow: Theory and algorithms,” *IEEE Trans. Info. Th.*, vol. 61, no. 4, pp. 1985–2007, 2015.

[9] Y. Chen and E. Candes, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” in *NIPS*, 2015, pp. 739–747.

[10] H. Zhang, Y. Zhou, Y. Liang, and Y. Chi, “Reshaped Wirtinger flow and incremental algorithm for solving quadratic system of equations,” in *NIPS*, 2016.

[11] G. Wang, G. B. Giannakis, and Y. C. Eldar, “Solving systems of random quadratic equations via truncated amplitude flow,” *arXiv preprint arXiv:1605.08285*, 2016.

[12] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen, “Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution,” in *Intl. Conf. Machine Learning (ICML)*, 2018.

[13] Yuxin Chen, Yuejie Chi, Jianqiang Fan, and Cong Ma, “Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval,” *arXiv preprint arXiv:1803.07726*, 2018.

[14] Xiaodong Li and Vladislav Voroninski, “Sparse signal recovery from quadratic measurements via convex programming,” *SIAM Journal on Mathematical Analysis*, vol. 45, no. 5, pp. 3019–3033, 2013.

[15] Kishore Jagannathan, Samet Oymak, and Babak Hassibi, “Sparse phase retrieval: Uniqueness guarantees and recovery algorithms,” *arXiv preprint arXiv:1311.2745*, 2013.

[16] G. Wang, L. Zhang, G. B. Giannakis, M. Akcakaya, and J. Chen, “Sparse phase retrieval via truncated amplitude flow,” *arXiv preprint arXiv:1611.07641*, 2016.

[17] T.T. Cai, X. Li, and Z. Ma, “Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow,” *The Annals of Statistics*, vol. 44, no. 5, pp. 2221–2251, 2016.

[18] Gauri Jagatap and Chinmay Hegde, “Sample-efficient algorithms for recovering structured signals from magnitude-only measurements,” *IEEE Transactions on Information Theory*, 2019.

[19] Zhi-Pei Liang, “Spatiotemporal imaging with partially separable functions,” in *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2007, pp. 988–991.

[20] N. Vaswani, S. Nayer, and Y. C. Eldar, “Low rank phase retrieval,” *IEEE Trans. Sig. Proc.*, August 2017.

[21] G. Jagatap, Z. Chen, S. Nayer, C. Hegde, and N. Vaswani, “Sample efficient Fourier ptychography for structured data,” *IEEE Trans. Comput. Imaging*, 2019.

[22] Sahand Negahban, Martin J Wainwright, et al., “Estimation of (near) low-rank matrices with noise and high-dimensional scaling,” *The Annals of Statistics*, vol. 39, no. 2, pp. 1069–1097, 2011.

[23] Yuxin Chen, Yuejie Chi, and Andrea J Goldsmith, “Exact and stable covariance estimation from quadratic sampling via convex programming,” *IEEE Transactions on Information Theory*, vol. 61, no. 7, pp. 4034–4059, 2015.

[24] E. J. Candes and B. Recht, “Exact matrix completion via convex optimization,” *Found. of Comput. Math.*, no. 9, pp. 717–772, 2008.

[25] P. Netrapalli, P. Jain, and S. Sanghavi, “Low-rank matrix completion using alternating minimization,” in *STOC*, 2013.

[26] Y. Cherapanamjeri, K. Gupta, and P. Jain, “Nearly-optimal robust matrix completion,” *ICML*, 2016.

[27] Yuejie Chi, Yue M Lu, and Yuxin Chen, “Nonconvex optimization meets low-rank matrix factorization: An overview,” *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5239–5269, 2019.

[28] K. Lee and Y. Bresler, “Admira: Atomic decomposition for minimum rank approximation,” *IEEE Transactions on Information Theory*, vol. 56, no. 9, September 2010.

[29] Roman Vershynin, *High-dimensional probability: An introduction with applications in data science*, vol. 47, Cambridge University Press, 2018.

[30] Sajan Goud Lingala, Yue Hu, Edward DiBella, and Mathews Jacob, “Accelerated dynamic MRI exploiting sparsity and low-rank structure: kt slr,” *IEEE transactions on medical imaging*, vol. 30, no. 5, pp. 1042–1054, 2011.

[31] Jiawen Yao, Zheng Xu, Xiaolei Huang, and Junzhou Huang, “An efficient algorithm for dynamic MRI using low-rank and total variation regularizations,” *Medical image analysis*, vol. 44, pp. 14–27, 2018.

[32] M.D. Butala, R.A. Frazin, Y. Chen, and F. Kamalabadi, “A monte carlo technique for large-scale dynamic tomography,” in *IEEE Intl. Conf. Acoustics, Speech, Sig. Proc. (ICASSP)*, 2007.

[33] J. Holloway, M. S. Asif, M. K. Sharma, N. Matsuda, R. Horstmeyer, O. Cossairt, and A. Veeraraghavan, “Toward long-distance subdiffraction imaging using coherent camera arrays,” *IEEE Trans. Comput. Imaging*, vol. 2, no. 3, pp. 251–265, 2016.

[34] Z. Chen, S. Nayer, and N. Vaswani, “Fast and sample-efficient low rank Fourier ptychography,” in *IEEE Intl. Conf. Image Proc. (ICIP)*, submitted, 2020.

[35] Hanchao Qi and Shannon M Hughes, “Invariance of principal components under low-dimensional random projection of the data,” in *2012 19th IEEE International Conference on Image Processing*. IEEE, 2012, pp. 937–940.

[36] Farhad Pourkamali Anaraki and Shannon Hughes, “Memory and computation efficient PCA via very sparse random projections,” in *International Conference on Machine Learning*, 2014, pp. 1341–1349.

[37] Akshay Krishnamurthy, Martin Azizyan, and Aarti Singh, “Subspace learning from extremely compressed measurements,” in *Asilomar Conference*, 2014.

[38] David P. Woodruff et al., “Sketching as a tool for numerical linear algebra,” *Foundations and Trends® in Theoretical Computer Science*, vol. 10, no. 1–2, pp. 1–157, 2014.

[39] Anna C. Gilbert, Jae Young Park, and Michael B. Wakin, “Sketched SVD: Recovering spectral features from compressive measurements,” *arXiv preprint arXiv:1211.0361*, 2012.

[40] Martin Azizyan, Akshay Krishnamurthy, and Aarti Singh, “Extreme compressive sampling for covariance estimation,” *IEEE Trans. Info. Th.*, 2018.

[41] Sujay Sanghavi, Rachel Ward, and Chris D. White, “The local convexity of solving systems of quadratic equations,” *Results in Mathematics*, vol. 71, no. 3–4, pp. 569–608, 2017.

[42] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM review*, vol. 52, no. 3, pp. 471–501, 2010.

[43] Jared Tanner and Ke Wei, “Normalized iterative hard thresholding for matrix completion,” *SIAM Journal on Scientific Computing*, vol. 35, no. 5, pp. S104–S125, 2013.

[44] Prateek Jain, Raghu Meka, and Inderjit S. Dhillon, “Guaranteed rank minimization via singular value projection,” in *Advances in Neural Information Processing Systems*, 2010, pp. 937–945.

[45] S. Tu, R. Boczar, M. Soltanolkotabi, and B. Recht, “Low-rank solutions of linear matrix equations via Procrustes flow,” *arXiv preprint arXiv:1507.03566*, 2015.

[46] Q. Zheng and J. Lafferty, “A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements,” in *NIPS*, 2015.

[47] R.H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” *IEEE Trans. Info. Th.*, vol. 56, no. 6, pp. 2980–2998, 2010.

[48] R. Sun and Z-Q. Luo, “Guaranteed matrix completion via non-convex factorization,” *IEEE Trans. Info. Th.*, vol. 62, no. 11, pp. 6535–6579, 2016.

[49] K. Jagannathan, S. Oymak, and B. Hassibi, “Recovery of sparse 1-d signals from the magnitudes of their Fourier transform,” in *IEEE Intl. Symp. Info. Th. (ISIT)*. IEEE, 2012, pp. 1473–1477.

[50] Mehryar Mohri and Ameet Talwalkar, “Can matrix coherence be efficiently and accurately estimated?,” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011, pp. 534–542.

[51] P. Netrapalli, U. N. Niranjani, S. Sanghavi, A. Anandkumar, and P. Jain, “Non-convex robust PCA,” in *NIPS*, 2014.

[52] E. Candes and Y. Plan, “Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. to appear,” *IEEE Trans. Info. Theor.*, 2009.

[53] C. Davis and W. M. Kahan, “The rotation of eigenvectors by a perturbation. iii,” *SIAM J. Numer. Anal.*, vol. 7, pp. 1–46, Mar. 1970.

[54] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” *Compressed sensing*, pp. 210–268, 2012.

[55] P. Narayananamurthy and N. Vaswani, “Nearly optimal robust subspace tracking,” in *Intl. Conf. Machine Learning (ICML)*, 2018, pp. 3701–3709.

[56] P. Narayananamurthy and N. Vaswani, “Provable dynamic robust PCA or robust subspace tracking,” *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1547–1577, 2019.

- [57] I. Mitliagkas, C. Caramanis, and P. Jain, “Streaming pca with many missing entries,” *Preprint*, 2014.
- [58] S. Nayer and N. Vaswani, “Phaseless subspace tracking,” *IEEE conference GlobalSIP (arXiv preprint arXiv:1809.04176)*, 2018.
- [59] S. Nayer and N. Vaswani, “Phast: Model-free phaseless subspace tracking,” *IEEE Intl. Conf. Acoustics, Speech, Sig. Proc. (ICASSP)*, 2019.
- [60] J. Zhan and N. Vaswani, “Time invariant error bounds for modified-CS based sparse signal sequence recovery,” *IEEE Trans. Info. Th.*, vol. 61, no. 3, pp. 1389–1409, 2015.
- [61] A. Gonen, D. Rosenbaum, Y. C. Eldar, and S. Shalev-Shwartz, “Subspace learning with partial information,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1821–1841, 2016.
- [62] L. Balzano, B. Recht, and R. Nowak, “Online Identification and Tracking of Subspaces from Highly Incomplete Information,” in *Allerton Conf. Comm., Control, Comput.*, 2010.
- [63] Y. Chi, Y. C. Eldar, and R. Calderbank, “Petrels: Parallel subspace estimation and tracking by recursive least squares from partial observations,” *IEEE Trans. Sig. Proc.*, December 2013.
- [64] L. Balzano and S. Wright, “Local convergence of an algorithm for subspace identification from partial data,” *Found. Comput. Math.*, vol. 15, no. 5, 2015.
- [65] L. Balzano, Y. Chi, and Y. M. Lu, “Streaming pca and subspace tracking: The missing data case,” *Proceedings of IEEE*, 2018.
- [66] P. Narayanamurthy, V. Daneshpajoh, and N. Vaswani, “Provable subspace tracking from missing data and matrix completion,” *IEEE Trans. Sig. Proc.*, pp. 4245–4260, 2019.
- [67] X. Yi, D. Park, Y. Chen, and C. Caramanis, “Fast algorithms for robust pca via gradient descent,” in *NIPS*, 2016.
- [68] Kiryung Lee, Yihong Wu, and Yoram Bresler, “Near-optimal compressed sensing of a class of sparse low-rank matrices via sparse power factorization,” *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1666–1698, 2017.
- [69] Natalia Ermolova and Sven-Gustav Haggman, “Simplified bounds for the complementary error function: application to the performance evaluation of signal-processing systems,” in *Signal Processing Conference, 2004 12th European*. IEEE, 2004, pp. 1087–1090.
- [70] C. Qiu and N. Vaswani, “Support predicted modified-cs for recursive robust principal components’ pursuit,” in *IEEE Intl. Symp. Info. Th. (ISIT)*, 2011.



**Namrata Vaswani** (Email: namrata@iastate.edu) received a B.Tech from the Indian Institute of Technology (IIT), Delhi, in 1999 and a Ph.D. from UMD in 2004, both in electrical engineering. Since Fall 2005, she has been with the Iowa State University where she is currently the Anderlik Professor of Electrical and Computer Engineering. Her research interests lie in a data science, with a particular focus on Statistical Machine Learning, Signal Processing, and Computer Vision. She has served two terms as an Associate Editor for the IEEE Transactions on Signal Processing; as a lead guest-editor for a Proceedings of the IEEE Special Issue (Rethinking PCA for modern datasets); and is currently serving as an Area Editor for the IEEE Signal Processing Magazine.

Vaswani is a recipient of the Iowa State Early Career Engineering Faculty Research Award (2014), the Iowa State University Mid-Career Achievement in Research Award (2019), University of Maryland’s ECE Distinguished Alumni Award (2019), as well as the 2014 IEEE Signal Processing Society Best Paper Award. This major award recognized the contributions of her 2010 IEEE Transactions on Signal Processing paper co-authored with her student Wei Lu on *Modified-CS: Modifying compressive sensing for problems with partially known support*. She is a Fellow of the IEEE Fellow (class of 2019).

## AUTHOR BIOGRAPHIES



**Seyedehsara Nayer** (Email: sarana@iastate.edu) is currently a PhD student in the Department of Electrical and Computer Engineering at Iowa State University. Previously she was working at Signal Processing research laboratory, Sharif University of Technology. Her research interests are around various aspects of information science and focuses on Computer Vision, Signal Processing, and Statistical Machine Learning..



**Praneeth Narayananurthy** (Email: pkurpadn@iastate.edu) (S’ 18) is a Ph.D. student in the Department of Electrical and Computer Engineering at Iowa State University (Ames, IA). He previously obtained his B.Tech degree in Electrical and Electronics Engineering from National Institute of Technology Karnataka (Surathkal, India) in 2014. His research interests include the algorithmic and theoretical aspects of High-Dimensional Statistical Signal Processing, and Machine Learning.