Blind Gain and Phase Calibration via Sparse Spectral Methods

Yanjun Li, Kiryung Lee, *Member, IEEE*, and Yoram Bresler, *Fellow, IEEE*

***Abstract*— Blind gain and phase calibration (BGPC) is a bilinear inverse problem involving the determination of unknown gains and phases of the sensing system, and the unknown signal, jointly. BGPC arises in numerous applications, e.g., blind albedo estimation in inverse rendering, synthetic aperture radar autofocus, and sensor array auto-calibration. In some cases, sparse structure in the unknown signal alleviates the ill-posedness of BGPC. Recently, there has been renewed interest in solutions to BGPC with careful analysis of error bounds. In this paper, we formulate BGPC as an eigenvalue/eigenvector problem and propose to solve it via power iteration, or in the sparsity or joint sparsity case, via truncated power iteration. Under certain assumptions, the unknown gains, phases, and the unknown signal can be recovered simultaneously. Numerical experiments show that power iteration algorithms work not only in the regime predicted by our main results, but also in regimes where theoretical analysis is limited. We also show that our power iteration algorithms for BGPC compare favorably with competing algorithms in adversarial conditions, e.g., with noisy measurement or with a bad initial estimate.**

***Index Terms*— Auto-calibration, greedy algorithm, inverse rendering, multichannel blind deconvolution, nonconvex optimization, power method, SAR autofocus, sensor array processing, truncated power iteration.**

I. Introduction

BLIND gain and phase calibration (BGPC), the joint recovery of the unknown gains and phases in the sensing system and the unknown signal, is a bilinear inverse problem that arises in many applications: the joint estimation of albedo and lighting conditions in inverse rendering [2]; the joint recovery of phase error and radar image in synthetic aperture radar (SAR) autofocus [3]; and auto-calibration of sensor gains and phases in array processing [4]. There exists a long line of research regarding the solutions for each application. However, fundamental sample complexities for the uniqueness of solutions to BGPC [5], [6], and error bounds for efficient algorithms [7], [8] have been established only recently. A main drawback of the guaranteed algorithms of [7] and [8] is that the recovery error is sensitive to the choice of certain linear constraints. We refer readers to Section I-D for a detailed discussion of prior art.

In this paper, we overcome the drawbacks of previous algorithms by reformulating the BGPC problem as an eigenvalue/eigenvector problem. In the subspace case, we use algorithms that find principal eigenvectors such as the power iteration algorithm (also known as the power method) [9, Sec. 8.2.1], to find the concatenation of the gain and phase vector and the vectorized signal matrix in the form of the principal component of a structured matrix. In the sparsity case, the problem resembles sparse principal component analysis (sparse PCA) [10]. We then propose to solve the sparse eigenvector problem using truncated power iteration [11].

The main contribution of this paper is the theoretical analysis of the error bounds of power iteration and truncated power iteration for BGPC in the subspace and joint sparsity cases, respectively. When the measurement matrix is random, and the signals and the noise are adversarial, our algorithms stably recover the unknown gains and phases, and the unknown signals with high probability under near optimal sample complexities. Since truncated power iteration relies on a good initial estimate, we also propose a simple initialization algorithm, and prove that the output is sufficiently good under certain technical conditions. The fundamental estimates derived in this paper can be applied to other algorithms for BGPC, and possibly to algorithms for similar problems.

We complement the theoretical results with numerical experiments, which show that the algorithms can indeed solve BGPC in the optimal regime. We also demonstrate that the algorithms are robust against noise and an inaccurate initial estimate. Experiments with different initialization schemes show that our initialization algorithm significantly outperforms the baseline. Then we apply the power iteration algorithm to inverse rendering, and showcase its effectiveness in real-world applications.

The rest of the paper is organized as follows. In the remainder of this section, we introduce the formulation of the BGPC problem, and survey related work. We then introduce the power iteration algorithms and our main theoretical results in Sections II and III, respectively. Sections IV and V give

Manuscript received December 20, 2017; revised September 8, 2018; accepted November 10, 2018. Date of publication November 28, 2018; date of current version April 19, 2019. This work was supported in part by the National Science Foundation under Grant IIS 14-47879 and Grant CCF 17-18771. This paper was presented in part at 2017 SampTA [1] and at 2017 Asilomar Conference.

Y. Li was with the University of Illinois at Urbana–Champaign, Urbana, IL 61801 USA (e-mail: yli145@illinois.edu).

K. Lee is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: lee.8763@osu.edu).

Y. Bresler is with the Coordinated Science Laboratory and the Department of Electrical and Computer Engineering, University of Illinois at Urbana–Champaign, Urbana, IL 61801 USA (e-mail: ybresler@illinois.edu).

Communicated by G. Kutyniok, Associate Editor for Signal Processing.

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

Digital Object Identifier 10.1109/TIT.2018.2883623

some fundamental estimates regarding the structured matrix in our BGPC formulation, and the proofs for our main results. We conduct some numerical experiments in Section VI, and conclude the paper with some discussion in Section VII.

A. Notations

We use $A^\top$, $\overline{A}$, and A^* to denote the transpose, the complex conjugate, and the conjugate transpose of a matrix A, respectively. The k-th entry of a vector λ is denoted by λ_k. The j-th column, the k-th row (in a column vector form), and the (k,j)-th entry of a matrix A are denoted by $a_{\cdot j}$, $a_{k\cdot}$, and a_{kj}, respectively. Upper script t in a vector $\eta^{(t)}$ denotes the iteration number in an iterative algorithm. We use I_n to denote the identity matrix of size $n \times n$, and $\mathbf{1}_{n,m}$ and $\mathbf{0}_{n,m}$ to denote the matrices of all ones and all zeros of size $n \times m$, respectively. The i-th standard basis vector is denoted by e_i, whose ambient dimension is clear in the context. The ℓ_p norm and ℓ_0 "norm" of a vector x are denoted by $\|x\|_p$ and $\|x\|_0$, respectively. The Frobenius norm and the spectral norm of a matrix A are denoted by $\|A\|_{\mathrm{F}}$ and $\|A\|$, respectively. The support of a sparse vector x is denoted by $\mathrm{supp}(x)$. The vector $\mathrm{vec}(X)$ denotes the concatenation of the columns of $X = [x_{\cdot 1}, x_{\cdot 2}, \ldots, x_{\cdot N}]$, i.e., $\mathrm{vec}(X) = [x_{\cdot 1}^\top, x_{\cdot 2}^\top, \ldots, x_{\cdot N}^\top]^\top$. A diagonal matrix with the entries of vector x on the diagonal is denoted by $\mathrm{diag}(x)$. The Kronecker product is denoted by $\otimes$. We use $\gtrsim$ to denote the relation greater than up to log factors. We use $[n]$ to denote the set $\{1, 2, \ldots, n\}$, and use $+$ to denote Minkowski addition of sets. For example, $[n]+\{m\}$ denotes $\{m+1, m+2, \ldots, m+n\}$. For an index set T, the projection operator onto T is denoted by Π_T, and the operator that restricts onto T is denoted by Ω_T. We use these operator notations for different spaces, and the ambient dimensions will be clarified in the context.

B. Problem Formulation

In this section, we introduce the BGPC problem with a subspace constraint or a sparsity constraint. Suppose $A \in \mathbb{C}^{n\times m}$ is the known measurement matrix, and $\lambda \in \mathbb{C}^n$ is the vector of unknown gains and phases, the k-th entry of which is $\lambda_k = |\lambda_k| e^{\sqrt{-1}\varphi_k}$. Here, $|\lambda_k|$ and φ_k denote the gain and phase of the k-th sensor, respectively. The BGPC problem is the simultaneous recovery of λ and the unknown signal matrix $X \in \mathbb{C}^{m\times N}$ from the following measurement:

$$Y = \mathrm{diag}(\lambda) A X + W, \tag{1}$$

where $W \in \mathbb{C}^{n\times N}$ is the measurement noise. The (k,j)-th entry in the measurement y_{kj} has the following expression:

$$y_{kj} = \lambda_k a_{k\cdot}^\top x_{\cdot j} + w_{kj}.$$

Clearly, BGPC is a bilinear inverse problem. The solution (λ, X) suffers from scaling ambiguity, i.e., $(\lambda/\sigma, \sigma X)$ generates the same measurements as (λ, X), and therefore cannot be distinguished from it. Despite the fact that the solution can have other ambiguity issues, in this paper, we consider the generic setting where the solution suffers only from scaling ambiguity [6].[1] Even in this setting, the solution is not unique, unless we exploit the structure of the signals. In this paper, we solve the BGPC problem under two scenarios – BGPC with a subspace structure, and BGPC with sparsity.

1) Subspace Case: Suppose that the known matrix A is tall ($n > m$) and has full column rank. Then the columns of AX reside in the low-dimensional subspace spanned by the columns of A. The problem is effectively unconstrained with respect to X.

2) Sparsity Case: Suppose that A is a known dictionary with $m \geq n$, while the columns of X are s_0-sparse, i.e., $\|x_{\cdot j}\|_0 \leq s_0$ for all $j \in [N]$. A variation of this setting is that the columns of X are jointly s_0-sparse, i.e., there are at most s_0 nonzero rows in X. In this case, the subspace constraint on AX no longer applies, and one must solve the problem with a sparsity (or joint sparsity) constraint.

The BGPC problem arises in applications including inverse rending, sensor array processing, multichannel blind deconvolution, and SAR autofocus. We refer the reader to our previous work [6, Sec. II.C] for a detailed account of applications of BGPC. For consistency, from now on, we use the convention in sensor array processing, and refer to n and N as the numbers of sensors and snapshots, respectively.

C. Our Contributions

We reformulate BGPC as the problem of finding the principal eigenvector of a matrix (or operator). In the subspace case, this can be solved using any eigen-solver, e.g., power iteration (Algorithm 1). In the sparsity case, we propose to solve this problem using truncated power iteration (Algorithm 2). Our main results can be summarized as follows.

Theorem 1: Under certain assumptions on A, λ, X, and W, one can solve the BGPC problem with high probability using:

1) **Subspace case:** algorithms that find the principal eigenvector of a certain matrix, e.g., power iteration, if $n \gtrsim m$ and $N \gtrsim 1$.
2) **Joint sparsity case:** truncated power iteration with a good initialization, if $n \gtrsim s_0$ and $N \gtrsim \sqrt{s_0}$.

In Table I, we compare the above results with the sample complexities for unique recovery in BGPC [6], and previous guaranteed algorithms for BGPC in the subspace case [7] and the sparsity case [8]. In the subspace case, the optimal sample complexities for unique recovery are $n > m$ and $N \geq 2$ (for $m \leq n/2$). Our power iteration method solves BGPC using optimal (up to log factors) numbers of sensors and snapshots. These sample complexities are comparable to the least squares method in [7]. Moreover, we show that power iteration is empirically more robust against noise than least squares.

In the joint sparsity case, the sample complexities for unique recovery are $n > 2s_0$ and $N \geq 2$ (for $s_0 < n/4$). Truncated power iteration solves BGPC with a joint sparsity structure, with an optimal (up to log factors) number of sensors, and a slightly suboptimal (within a factor of $\sqrt{s_0}$ and log factors) number of snapshots. In comparison, the ℓ_1 minimization

[1] An example of another ambiguity is a shift ambiguity when A is the discrete Fourier transform matrix [5], [8]. For a generic matrix A, the solution to BGPC does not suffer from shift ambiguity.

TABLE I
COMPARISON OF SAMPLE COMPLEXITIES WITH PRIOR WORK

	Subspace	Joint Sparsity	Sparsity
Unique Recovery [6]	$n > m$ $N \geq \frac{n-1}{n-m}$	$n > 2s_0$ $N \geq \frac{n-1}{n-2s_0}$	–
Least Squares [7]	$n \gtrsim m$ $N \gtrsim 1$	–	–
ℓ_1 Minimization [8]	–	–	$n \gtrsim s_0$ $N \gtrsim n$
This Paper	$n \gtrsim m$ $N \gtrsim 1$	$n \gtrsim s_0$ $N \gtrsim \sqrt{s_0}$	–

Note: n, N, m and s_0 represent the number of sensors, the number of snapshots, the subspace dimension, and the sparsity level, respectively. The above sample complexities are derived for different signal models. This paper and [7] assume that A is a random Gaussian matrix, and [8] assumes that A is the discrete Fourier transform matrix and X is random following a Bernoulli-Subgaussian model.

method for the sparsity case of BGPC uses a similar number of sensors, but a much larger number of snapshots. Numerical experiments show that truncated power iteration empirically succeed, in both the joint sparsity case and the more general sparsity case, in the optimal regime.

Theorem 1 gives a theoretical guarantee for power iteration in the subspace case, and a *local* convergence guarantee for truncated power iteration in the joint sparsity case. The success of truncated power iteration relies on a good initial estimate of X and λ, which may or may not be available depending on the application. We propose a simple initialization algorithm (Algorithm 3) with the following guarantee.

Theorem 2: Under additional assumptions on the absolute values of the nonzero entries in X, our initialization algorithm produces a sufficiently good estimate of λ and X if $n \gtrsim s_0^2$. (We do not require any additional assumption on the number N of snapshots.)

Despite the above scaling law predicted by theory, numerical experiments suggest that our initialization scheme is effective when $n \gtrsim s_0$.

D. Related Work

BGPC arises in many real-world scenarios, and previous solutions have mostly been tailored to specific applications such as sensor array processing [4], [12], [13], sensor network calibration [14], [15], synthetic aperture radar autofocus [3], and computational relighting [2]. However, the previous methods do *not* have theoretical guarantees in the forms of quantitative error bounds.

The idea of solving BGPC by reformulating it into a linear inverse problem, which is a key idea in this paper, has been proposed by many prior works [2], [3], [14]. In particular, Bilen *et al.* [16] provided a solution to BGPC with high-dimensional but sparse signals using ℓ_1 minimization. However, such methods have not been carefully analyzed until recently.

Ling and Strohmer [7] derived an error bound for the least squares solution in the subspace case of BGPC (and two similar formulations). After the conference version of our paper was submitted to SampTA 2017 [1], and while this extended version was in preparation, Ling and Strohmer independently proposed, as a variation of the least squares approach, the spectral method for the subspace case of BGPC [17]. Their spectral method and our approach to the subspace case are essentially identical, as one can be derived from the other with a few matrix manipulations.[2] In addition to bounding the error in the principal eigenvector in our formulation, we also establish convergence rate and error bounds for an efficient power iteration algorithm that finds the principal eigenvector. We show that the power iteration method has sample complexities comparable to those of the least squares method [7], but is more robust to noise than the latter, both in theory and in practice.

Wang and Chi [8] gave a theoretical guarantee for ℓ_1 minimization that solves BGPC in the sparsity case, where they assumed that A is the discrete Fourier transform (DFT) matrix and X is random following a Bernoulli-Subgaussian model. In this paper, we give a guarantee for truncated power iteration under the assumption that A is a complex Gaussian random matrix, and X is *jointly* sparse, well-conditioned, and deterministic. In this sense, we consider an adversarial scenario for the signal X. Our sample complexity results require a near optimal number n of sensors, and a much smaller number N of snapshots than [8]. Moreover, truncated power iteration is more robust against noise and inaccurate initial estimate of phases.

Very recently, Eldar *et al.* [18] proposed new methods for BGPC with signals whose sparse components may lie off the grid. Similar to earlier work on blind calibration of sensor arrays [4], these methods rely on empirical covariance matrices of the measurements and therefore need a relatively large number of snapshots.

A problem related to BGPC is multichannel blind deconvolution (MBD). Most previous works on MBD consider linear convolution with a finite impulse response (FIR) filter model (see [19]–[22], and a recent stabilized method [23], [24]). In comparison, BGPC is equivalent to MBD with *circular* convolution and a subspace model or a sparsity model, akin to some recent studies [7], [8]. BGPC is more general in the sense that: (a) linear convolution can be rewritten as circular convolution via zero-padding the signal and the filter; (b) the FIR filter model is a special case of the subspace model.

To position BGPC in a broader context, BGPC is a special bilinear inverse problem [5], which in turn is a special case of low-rank matrix recovery from incomplete measurements [25]–[28]. A resurgence of interest in bilinear inverse problems was pioneered by the recent studies in single-channel blind deconvolution of signals with subspace or sparsity structures, where both the signal and the filter are structured [29]–[33].

[2]The only difference between their spectral method and our formulation for the subspace case, is that they compute the right singular vector corresponding to the smallest singular value of a tall matrix, and we compute the principal eigenvector of a smaller square matrix.

Another related bilinear inverse problem is blind calibration via repeated measurements from multiple different sensing operators [34]–[39]. Since blind calibration with repeated measurements is in principle an easier problem than BGPC [7], we believe our methods for BGPC and our theoretical analysis can be extended to this scenario.

Also related is the phase retrieval problem [40], where there only exists uncertainty in the phases (and not the gains) of the sensing system. An active line of work solves phase retrieval with guaranteed algorithms (see [41]–[47] for a recent review).

The error bounds of power iteration and truncated power iteration have been analyzed in general settings, e.g., in [9, Sec. 8.2.1] and [11]. These previous results hinge on spectral properties of matrices such as gaps between eigenvalues, which do not translate directly to sample complexity requirements. This paper undertakes analysis specific to BGPC. We relate spectral properties in BGPC to some technical conditions on λ, A, X, and W, and derive recovery error under near optimal sample complexities. We also adapt the analysis of sparse PCA [11] to accommodate a structured sparsity constraint in BGPC.

BGPC and our proposed methods are non-convex in nature. In particular, our truncated power iteration algorithm can be interpreted as projected gradient descent for a non-convex optimization problem. There have been rapid developments in guaranteed non-convex methods [48] in a variety of domains such as matrix completion [49]–[51], dictionary learning [52], [53], blind deconvolution [32], [33], and phase retrieval [42], [43], [54]. It is a common theme that carefully crafted non-convex methods have better theoretical guarantees in terms of sample complexity than their convex counterparts, and often have faster implementations and better empirical performance. This paper is a new example of such superiority of non-convex methods.

II. Power Iteration Algorithms for BGPC

Next, we describe the algorithms we use to solve BGPC. In Section II-A, we introduce a simple trick that turns the bilinear inverse problem in BGPC to a linear inverse problem. In Sections II-B and II-C, we introduce the power iteration algorithm we use to solve BGPC with a subspace structure, and the truncated (or sparse) power iteration algorithm we use to solve BGPC with sparsity, respectively.

A. From Bilinearity to Linearity

We use a simple trick to turn BGPC into a linear inverse problem [14]. Without loss of generality, assume that $\lambda_k \neq 0$ for $k \in [n]$. Indeed, if any sensor has zero gain, then the corresponding row in Y is all zero or contains only noise, and we can simply remove the corresponding row in (1). Let γ denote the entrywise inverse of λ, i.e., $\gamma_k = 1/\lambda_k$ for $k \in [n]$. We have

$$\operatorname{diag}(\gamma) Y_{\mathrm{s}} = AX, \tag{2}$$

where $Y_{\mathrm{s}} = \operatorname{diag}(\lambda) AX$ is the noiseless measurement. Equation (2) is linear in all the entries of γ and X. The bilinear inverse problem in (λ, X) now becomes a linear inverse problem in (γ, X). In practice, since only the noisy measurement Y is available, one can solve $\operatorname{diag}(\gamma) Y \approx AX$.

This technique was widely used to solve BGPC with a subspace structure, in applications such as sensor network calibration [14], synthetic aperture radar autofocus [3], and computational relighting [2]. Recently, Ling and Strohmer [7] analyzed the least squares solution to (2). Wang and Chi [8] considered a special case where A is the DFT matrix, and analyzed the solution of a sparse X by minimizing the ℓ_1 norm of $A^{-1}\operatorname{diag}(\gamma) Y$.

We use the same trick in our algorithms. Define

$$D \coloneqq \begin{bmatrix} I_N \otimes a_{1\cdot}^{\top} \\ \vdots \\ I_N \otimes a_{n\cdot}^{\top} \end{bmatrix} \in \mathbb{C}^{Nn \times Nm}, \tag{3}$$

$$E \coloneqq \begin{bmatrix} y_{1\cdot} & & \\ & \ddots & \\ & & y_{n\cdot} \end{bmatrix} \in \mathbb{C}^{Nn \times n}. \tag{4}$$

We can decompose E into $E = E_{\mathrm{s}} + E_{\mathrm{n}}$, where

$$E_{\mathrm{s}} \coloneqq \begin{bmatrix} \lambda_1 X^{\top} a_{1\cdot} & & \\ & \ddots & \\ & & \lambda_n X^{\top} a_{n\cdot} \end{bmatrix} \in \mathbb{C}^{Nn \times n},$$

$$E_{\mathrm{n}} \coloneqq \begin{bmatrix} w_{1\cdot} & & \\ & \ddots & \\ & & w_{n\cdot} \end{bmatrix} \in \mathbb{C}^{Nn \times n}.$$

Define also

$$B \coloneqq \begin{bmatrix} D^*D & \alpha D^*E \\ \alpha E^*D & \alpha^2 E^*E \end{bmatrix} \mathbb{C}^{(Nm+n)\times(Nm+n)}, \tag{5}$$

$$B_{\mathrm{s}} \coloneqq \begin{bmatrix} D^*D & \alpha D^*E_{\mathrm{s}} \\ \alpha E_{\mathrm{s}}^*D & \alpha^2 E_{\mathrm{s}}^*E_{\mathrm{s}} \end{bmatrix} \in \mathbb{C}^{(Nm+n)\times(Nm+n)},$$

where α is a nonzero constant specified later.

Clearly, (2) can be rewritten as

$$Dx - E_{\mathrm{s}}\gamma = 0,$$

where $x = \operatorname{vec}(X)$. Equivalently, $\eta = [x^{\top}, -\gamma^{\top}/\alpha]^{\top}$ is a null vector of B_{s}. When certain sufficient conditions are satisfied, η is the unique null vector of B_{s}. For example, if λ, A, and X are in general positions in $\mathbb{C}^n$, $\mathbb{C}^{n\times m}$, and $\mathbb{C}^{m\times N}$, respectively, then $N \geq \frac{n-1}{n-m}$ snapshots are sufficient to guarantee uniqueness of the solution to BGPC in the subspace case. We refer readers to our work on the identifiability in BGPC for more details [5], [6].

Since only the noisy matrix B is accessible in practice, one can instead find the minor eigenvector, i.e., the eigenvector corresponding to the smallest eigenvalue of B. The rest of this section focuses on algorithms that find such an eigenvector of B, with no constraint (in the subspace case), or with a sparsity constraint (in the sparsity case).

B. Power Iteration for BGPC With a Subspace Structure

In the subspace case ($n > m$), we solve for the minor eigenvector of the positive definite matrix B. In Section III,

Algorithm 1 Power Iteration for BGPC

Input: $A \in \mathbb{C}^{n\times m}$, $Y \in \mathbb{C}^{n\times N}$, initial estimate $\eta^{(0)} \in \mathbb{C}^{Nm+n}$
Output: $\eta^{(t)} \in \mathbb{C}^{Nm+n}$
Parameters: α, β
Compute operator $G : \mathbb{C}^{Nm+n} \to \mathbb{C}^{Nm+n}$ by (3), (4), (5), (6)
$t \leftarrow 1$
repeat
  Compute $\eta^{(t)} = G\eta^{(t-1)}/\|G\eta^{(t-1)}\|_2$
  $t \leftarrow t+1$
until *convergence criterion is met*

we derive an upper bound on the error between this eigenvector and the true solution η.

The minor eigenvector of B can be computed by a variety of methods. Here, we propose an algorithm that remains computationally efficient for large scale problems. By eigenvalue decomposition, the null vector of B is identical to the principal eigenvector of

$$G = \beta I_{mN+n} - B, \tag{6}$$

for a large enough constant β. This eigenvector can be computed using the power iteration algorithm (see Algorithm 1).

The size of G is $(Nm+n)\times(Nm+n)$. An advantage of Algorithm 1 over an eigen-solver that decomposes G, is that one does not need to explicitly compute the entries of G to iteratively apply it to a vector. Furthermore, rather than $O((Nm+n)^2)$, by the structure of D and E, the per iteration time complexity of applying the operator G to a vector is only $O(mnN)$. This can be further reduced if A and A^* are linear operators with implementations faster than $O(mn)$.

The rule of thumb for selecting parameter α is that the ℓ_2 norms of the columns of D be close to those of αE so that G in (6) exhibits good spectral properties for power iterations. A safe choice for β is $\|B\|$, which may be conservatively large in some cases, but works well in practice. In Section III, we discuss our choice of parameters α, β under certain normalization assumptions (see Remark 8).

Algorithm 1 converges to the principal eigenvector of G, as long as the initial estimate $\eta^{(0)}$ is not orthogonal to that eigenvector. This insensitivity to initialization is a privilege not shared by the sparsity case (see Section II-C).

C. Truncated Power Iteration for BGPC With Sparsity

When $2 \le n \le m$, $[D, \alpha E] \in \mathbb{C}^{Nn\times(Nm+n)}$ is a fat matrix, and the null space of B has dimension at least 2. Therefore, there exist at least two linearly independent eigenvectors corresponding to the largest eigenvalue of G. To overcome the ill-posedness, one can leverage the sparsity structure in X to make the solution to the eigenvector problem unique.

Let $\Pi_s(x)$ denote the projection of a vector x onto the set of s-sparse vectors. It is computed by setting to zero all but the s entries of x of the largest absolute values. Let $\Pi'_s(X)$ denote the projection of a matrix X onto the set of matrices whose columns are jointly s-sparse. This projection is computed by setting to zero all but the s rows of X of the largest ℓ_2 norms. We define two projection operators on $\eta = [x^\top, -\gamma^\top/\alpha]^\top$ that will be used repeatedly in the rest of this paper:

$$\widetilde{\Pi}_s(\eta) \coloneqq [\Pi_s(x_{\cdot 1})^\top, \Pi_s(x_{\cdot 2})^\top, \ldots, \Pi_s(x_{\cdot N})^\top, -\gamma^\top/\alpha]^\top,$$
$$\widetilde{\Pi}'_s(\eta) \coloneqq [\mathrm{vec}(\Pi'_s(X))^\top, -\gamma^\top/\alpha]^\top.$$

For the sparsity case of BGPC, we adapt the eigenvector problem in Section II-B by adding a sparsity constraint:

$$\begin{aligned} \max_{\eta}\ & \eta^* G \eta \\ \text{s.t. }\ & \|\eta\|_2 = 1, \\ & \widetilde{\Pi}_{s_0}(\eta) = \eta. \end{aligned} \tag{7}$$

This nonconvex optimization is very similar to the sparse PCA problem. The only difference lies in the structure of the sparsity constraint. In sparse PCA, the principal component is s_0-sparse. In (7), the vector η consists of s_0-sparse vectors $x_{\cdot 1}, x_{\cdot 2}, \ldots, x_{\cdot N}$, and a dense vector $-\gamma/\alpha$.

To solve (7), we adopt a sparse PCA algorithm called truncated power iteration [11], and revise it to adapt to the sparsity structure of BGPC (see Algorithm 2). One can choose parameters α and β using the same rules as in Section II-B. Note that we use a sparsity level $s_1 \ge s_0$ in this algorithm, for two reasons: (a) in practice, it is easier to obtain an upper bound on the sparsity level instead of the exact number of nonzero entries in the signal; and (b) the ratio s_0/s_1 is an important constant in the main results, controlling the trade-off between the number of measurements and the rate of convergence.

For the joint sparsity case, we use essentially the same algorithm, with $\widetilde{\Pi}_{s_1}$ replaced by $\widetilde{\Pi}'_{s_1}$.

Since (7) is a nonconvex optimization problem, a good initialization $\eta^{(0)}$ is crucial to the success of Algorithm 2. Algorithm 3 outlines one such initialization. We denote by Π_{T_x} the projection onto the support set T_x, which sets to zero all rows of D^*E but the s_1 rows of the largest ℓ_2 norms in each block. (Recall that $d_{\cdot((j-1)m+\ell)}$ denotes the $((j-1)m+\ell)$-th column of D, and the j-th block of D^*E consists of m contiguous rows $\{d^*_{\cdot((j-1)m+\ell)}E\}_{\ell\in[m]}$.) Then the normalized left and right singular vectors u and v of $\Pi_{T_x}D^*E$ are computed as initial estimates for x and λ. We use $1./v$ to denote the entrywise inverse of v except for zero entries, which are kept zero. In Section III, we further comment on how to choose a proper initial estimate $\eta^{(0)}$ (see Remark 13).

D. Alternative Interpretation as Projected Gradient Descent

Algorithms 1 and 2 can be interpreted as gradient descent and projected gradient descent, respectively. Next, we explain such equivalence using the sparsity case as an example.

Recall that BGPC is linearized as $\left[D\ \alpha E\right]\eta = 0$. Relaxing the sparsity level from s_0 to s_1, the optimization in (7) is equivalent to:

$$\begin{aligned} \min_{\eta}\ & \frac{1}{2}\left\|\left[D\ \alpha E\right]\eta\right\|_2^2 \\ \text{s.t. }\ & \|\eta\|_2 = 1, \\ & \widetilde{\Pi}_{s_1}(\eta) = \eta. \end{aligned}$$

Algorithm 2 Truncated Power Iteration for BGPC With Sparsity

Input: $A \in \mathbb{C}^{n\times m}$, $Y \in \mathbb{C}^{n\times N}$, initial estimate $\eta^{(0)} \in \mathbb{C}^{Nm+n}$
Output: $\eta^{(t)} \in \mathbb{C}^{Nm+n}$
Parameters: α, β, s_1
Compute operator $G : \mathbb{C}^{Nm+n} \to \mathbb{C}^{Nm+n}$ by (3), (4), (5), (6)
$t \leftarrow 1$
repeat
  Compute $\tilde{\eta}^{(t)} = G\eta^{(t-1)}/\|G\eta^{(t-1)}\|_2$
  Compute $\eta^{(t)} = \widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})/\|\widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})\|_2$
  $t \leftarrow t+1$
until *convergence criterion is met*

The gradient of the objective function at $\eta^{(t-1)}$ is

$$\begin{bmatrix} D^* \\ \alpha E^* \end{bmatrix} \begin{bmatrix} D & \alpha E \end{bmatrix} \eta^{(t-1)} = B\eta^{(t-1)}.$$

Each iteration of projected gradient descent consists of two steps:

(i) **Gradient descent** with a step size of $1/\beta$:

$$\tilde{\eta}^{(t)} = \eta^{(t-1)} - \frac{1}{\beta} B\,\eta^{(t-1)} = \frac{1}{\beta} G\,\eta^{(t-1)}.$$

(ii) **Projection** onto the constraint set, i.e., the intersection of a cone ($\widetilde{\Pi}_{s_1}(\eta) = \eta$) and a sphere ($\|\eta\|_2 = 1$):

$$\eta^{(t)} = \widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})/\|\widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})\|_2.$$

Clearly, the two steps are identical to those in each truncated power iteration except for a different scaling in Step (i), which, due to the normalization in Step (ii), is insignificant.

III. Main Results

In this section, we give theoretical guarantees for Algorithms 1 and 2 in the subspace case and in the joint sparsity case, respectively. We also give a guarantee for the initialization by Algorithm 3.

A. Main Assumptions

We start by stating the assumptions on A, λ, X and W, which we use throughout this section.

Assumption 3: A is a complex Gaussian random matrix, whose entries are i.i.d. following $\mathcal{CN}(0, \frac{1}{n})$. Equivalently, the vectors $\{a_k.\}_{k=1}^n$ are i.i.d. following $\mathcal{CN}(\mathbf{0}_{m,1}, \frac{1}{n} I_m)$.

Assumption 4: The vector λ has "flat" gains in the sense that $1-\delta \le |\lambda_k|^2 \le 1+\delta$ for some $\delta \in (0,1)$.

Assumption 5: The matrix $X \in \mathbb{C}^{m\times N}$ is normalized and has good conditioning, i.e., $\|X\|_{\mathrm{F}} = 1$, and for some $\theta \in (0,1)$ we have:

- **Subspace case:**

$$\min\{\|NX^*X - I_N\|, \|mXX^* - I_m\|\} \le \theta,$$

- **Joint sparsity case:**

$$\min\{\|NX^*X - I_N\|, \|s_0\mathbf{\Omega}_{T_0} XX^* \mathbf{\Omega}_{T_0}^* - I_{s_0}\|\} \le \theta,$$

Algorithm 3 Initialization for Truncated Power Iteration

Input: $A \in \mathbb{C}^{n\times m}$, $Y \in \mathbb{C}^{n\times N}$
Output: initial estimate $\eta^{(0)} \in \mathbb{C}^{Nm+n}$
Parameters: s_1
Compute matrix $D^*E \in \mathbb{C}^{Nm\times n}$ by (3), (4)
$T_x \leftarrow \emptyset$
for $j \in [N]$ **do**
  Compute the row norms $\|d^*_{((j-1)m+\ell)} E\|_2$ for $\ell \in [m]$
  Find subset $T_j \subset [m]$ ($|T_j| = s_1$) s.t. for $\ell \in T_j$ and $\ell' \in [m]\backslash T_j$:

$$\|d^*_{((j-1)m+\ell)} E\|_2 \ge \|d^*_{((j-1)m+\ell')} E\|_2$$

  Merge support $T_x \leftarrow T_x \bigcup \big(T_j + \{(j-1)m\}\big)$
end
Compute the principal left and right singular vectors u, v of $\Pi_{T_x} D^*E$
$\eta^{(0)} \leftarrow [u^\top, -(1./v^\top)/n]^\top$
$\eta^{(0)} \leftarrow \eta^{(0)}/\|\eta^{(0)}\|_2$

where $\mathbf{\Omega}_T$ denotes the operator that restricts a matrix to the row support T, and $T_0 \coloneqq \{i \in [m] | \|e_i^\top X\|_2 > 0\}$ ($|T_0| = s_0$) is the row support of X.

Assumptions 3 – 5 can be relaxed in practice.

- The complex Gaussian distribution in Assumption 3 can be relaxed to $\mathcal{CN}(0, \sigma_A^2)$ for any $\sigma_A > 0$. We choose the particular scaling $\sigma_A^2 = 1/n$, because then A satisfies the restricted isometry property (RIP) [55], i.e., $(1-\delta_s)\|x\|_2^2 \le \|Ax\|_2^2 \le (1+\delta_s)\|x\|_2^2$ for some $\delta_s \in (0,1)$, when n is large compared to the number s of nonzero entries in x.
- The gains can center around any $\sigma > 0$, i.e., $\sigma(1-\delta) \le |\lambda_k|^2 \le \sigma(1+\delta)$. Due to bilinearity, we may assume that λ_k's center around 1 without loss of generality by solving for $(\lambda/\sigma, \sigma X)$.
- The Frobenius norm $\|X\|_{\mathrm{F}}$ of matrix X can be any positive number. If $\|X\|_{\mathrm{F}}$ is known, one can scale X to have unit Frobenius norm before solving BGPC. In practice, the norm of X is generally unknown. However, due to Assumptions 3 (RIP) and 4 ("flat" gains), we have

$$\sqrt{(1-\delta_s)(1-\delta)} \le \frac{\|\mathrm{diag}(\lambda)AX\|_{\mathrm{F}}}{\|X\|_{\mathrm{F}}} \le \sqrt{(1+\delta_s)(1+\delta)}.$$

  Hence $\|Y\|_{\mathrm{F}}$ is a good surrogate for $\|X\|_{\mathrm{F}}$ in noiseless or low noise settings, and one can scale X by $1/\|Y\|_{\mathrm{F}}$ to achieve the desired scaling. The slight deviation of $\|X\|_{\mathrm{F}}/\|Y\|_{\mathrm{F}}$ from 1 does *not* have any significant impact on our theoretical analysis. Therefore, we assume $\|X\|_{\mathrm{F}} = 1$ to simply the constants in our derivation.
- The conditioning of X can also be relaxed. When N is large, one can choose a subset of $N' < N$ columns in Y, such that the matrix formed from the corresponding columns of X has good conditioning. When noise amplification is not of concern (noiseless or low noise settings), one can choose a preconditioning matrix $H \in \mathbb{C}^{N\times N}$

such that $X' = XH$ is well conditioned, and then solve the BGPC with $Y' = YH$.

In summary, we can manipulate the BGPC problem and make it approximately satisfy our assumptions. For example, (1) can be rewritten as:

$$\frac{1}{\|YH\|_{\mathrm{F}}} YH = \operatorname{diag}\Big(\frac{\lambda}{\sigma}\Big)\Big(\frac{1}{\sqrt{n}\sigma_A} A\Big)\Big(\frac{\sqrt{n}\sigma\sigma_A}{\|YH\|_{\mathrm{F}}} XH\Big) + \frac{1}{\|YH\|_{\mathrm{F}}} WH.$$

We can run Algorithms 1 and 2 with input $\frac{1}{\sqrt{n}\sigma_A} A$ and $\frac{1}{\|YH\|_{\mathrm{F}}} YH$, and solve for $\frac{\lambda}{\sigma}$ and $\frac{\sqrt{n}\sigma\sigma_A}{\|YH\|_{\mathrm{F}}} XH$. The above manipulations do not have any significant impact on the solution, or on our theoretical analysis. However, by making these assumptions, we eliminate some tedious and unnecessary discussions.

We also need an assumption on the noise level.

Assumption 6: The noise term W satisfies

- **Subspace case:** $\max_{k\in[n], j\in[N]} |w_{kj}| \le \frac{C_W}{\sqrt{nN}}$
- **Joint sparsity case:** $\max_{k\in[n], j\in[N]} |w_{kj}| \le \frac{C_W}{\sqrt{n}N^2}$

for an absolute constant $C_W > 0$.

In the subspace case, the assumption on the noise level is very mild. Because under Assumptions 3 – 5, $\|\operatorname{diag}(\lambda)AX\|_{\mathrm{F}} \le \sqrt{(1+\delta_s)(1+\delta)}$, the noise term W, which satisfies $\|W\|_{\mathrm{F}} \le C_W$, can be on the same order in terms of Frobenius norm as the clean signal $\operatorname{diag}(\lambda)AX$.

Finally, the following assumption is required for a theoretical guarantee of the initialization.

Assumption 7: For all $j \in [N]$, there exists $T_j' \subset \operatorname{supp}(x_{\cdot j}) \subset [m]$, such that for all $\ell \in T_j'$,

$$\frac{|x_{\ell j}|^2}{\|x_{\cdot j}\|_2^2} \ge \frac{\omega}{s_0},$$

for some absolute constant ω, and

$$\frac{\sum_{\ell' \in [m]\setminus T_j'} |x_{\ell' j}|^2}{\|x_{\cdot j}\|_2^2} \le \delta_X,$$

for some small absolute constant $\delta_X \in (0, 1)$.

Assumption 7 says that the support of $x_{\cdot j}$ can be partitioned into two subsets. The absolute values of the entries in the first subset T_j' are sufficiently large. Moreover, the total energy (sum of squares of the entries) in the second subset is small compared to the squared norm of $x_{\cdot j}$. For example, the assumption is satisfied with $\omega = 1$ and $\delta_X = 0$ in the following special case: $T_j' = \operatorname{supp}(x_{\cdot j})$ (therefore $x_{\ell' j} = 0$ for $\ell' \in [m]\setminus T_j'$), and the absolute values of the nonzero entries are all equal, i.e., $x_{\ell j} = \pm\frac{\|x_{\cdot j}\|}{\sqrt{s_0}}$.

We would like to emphasize that Assumption 7 is not very demanding, and it is satisfied by commonly used signal models. For example, if the s_0 nonzero entries of $x_{\cdot j}$ are i.i.d. standard Gaussian random variables, then Assumption 7 is satisfied with high probability, with $\omega = \frac{1}{4}$ and $\delta_X = \frac{1}{\sqrt{2\pi}}$. This can be shown as follows: First of all, by simple Chernoff bounds, $\frac{s_0}{2} \le \|x_{\cdot j}\|_2^2 \le 2s_0$ with probability at least $1 - 2\, e^{-0.09\, s_0}$. Define $T_j' \coloneqq \{\ell : |x_{\ell j}| \ge \frac{1}{2}\}$. Then, invoking again the Chernoff bound, the number of nonzero entries with absolute values less than $\frac{1}{2}$ (in support set $[m]\setminus T_j'$) is less than $\frac{2s_0}{\sqrt{2\pi}}$ with probability at least $1 - e^{-0.15 s_0}$. Therefore, with high probability,

$$\frac{|x_{\ell j}|^2}{\|x_{\cdot j}\|_2^2} \ge \frac{1/2}{2s_0} = \frac{1}{4s_0},$$

and

$$\frac{\sum_{\ell' \in [m]\setminus T_j'} |x_{\ell' j}|^2}{\|x_{\cdot j}\|_2^2} \le \frac{\frac{2s_0}{\sqrt{2\pi}} \times (\frac{1}{2})^2}{s_0/2} = \frac{1}{\sqrt{2\pi}}.$$

Before introducing our main results, we disclose the choice of parameters α and β for our theoretical analysis of Algorithms 1 and 2.

Remark 8: When Assumptions 3 – 5 are satisfied, we choose $\alpha = \sqrt{n}$ and $\beta = 3/2$.

B. A Perturbation Bound for the Eigenvector Problem

Next, we introduce a key result, a perturbation bound for the eigenvector problem, which is used to derive error bounds for power iteration algorithms.

Let $\{T_j\}_{j=1}^N$ denote subsets of $[m]$, such that $|T_j| = s$ and $\operatorname{supp}(x_{\cdot j}) \subset T_j$. We define $T_x \subset [Nm]$ and $T_\eta \subset [Nm+n]$ as follows:

$$T_x \coloneqq \bigcup_{j\in[N]} \big(T_j + \{(j-1)m\}\big), \tag{8}$$

$$T_\eta \coloneqq T_x \bigcup \big([n] + \{Nm\}\big). \tag{9}$$

Recall that Ω_T restricts a vector to the support T, and hence $\Omega_T^*\Omega_T$ is the projection operator onto the support T. Clearly, we have $x = \Omega_{T_x}^*\Omega_{T_x} x$, and $\eta = \Omega_{T_\eta}^*\Omega_{T_\eta}\eta$. In the subspace case discussed in Theorem 9, we have $s = m$, $T_j = [m]$, $T_x = [Nm]$, and $T_\eta = [Nm+n]$. In the *joint* sparsity case, we have $T_1 = T_2 = \cdots = T_N$. We set $|T_j| = s = s_0 + 2s_1$, which we justify later in the analysis of truncated power iteration.

Let

$$\dot{\eta} \coloneqq \frac{\eta}{\|\eta\|_2}$$

denote the normalized version of the ground truth η, which is the eigenvector of B_{s} and $\mathbb{E}B_{\mathrm{s}}$ corresponding to eigenvalue 0. Let $\hat{\eta}$ denote the principal eigenvector of G. In the joint sparsity case, let $\hat{\eta}_{T_\eta}$ denote the principal eigenvector of $\Omega_{T_\eta} G \Omega_{T_\eta}^*$, where $T = T_1 = \cdots = T_N$, $|T| = s$, and the support of η is a subset of T_η defined in (9).

In Algorithms 1 and 2 and in our analysis, vectors $\dot{\eta}$, $\hat{\eta}$, and $\eta^{(t)}$ are normalized to unit norm. However, multiplication by a scalar of unit modulus is a remaining ambiguity, i.e., the set $\{e^{\sqrt{-1}\varphi}\dot{\eta} : \varphi \in [0, 2\pi)\}$ is an equivalence class for $\dot{\eta}$. Our main results use $d(\eta, \eta') \coloneqq \min_\varphi \|e^{\sqrt{-1}\varphi}\eta - \eta'\|_2$ to denote the distance between η and η', which is a metric on the set of such equivalence classes.

Theorem 9 (Subspace Case): Let $\alpha = \sqrt{n}$, and suppose Assumptions 3 – 6 are satisfied with $\delta < 1/3$ and a sufficiently

small absolute constant $C_W > 0$. Then there exist absolute constants $c, C, C' > 0$, such that if

$$\max\left\{\frac{m\log^2(Nm+n)}{n}, \frac{\log(Nm+n)}{N}, \frac{\log(Nm+n)}{m}\right\} \le C, \tag{10}$$

then with probability at least $1-2n^{-c}-e^{-cm}$,

$$d(\hat{\eta},\ \dot{\eta}) \le \Delta,$$

where

$$\Delta \coloneqq \frac{8C'}{1-3\delta}\max\{\nu,\ \nu^2\}, \tag{11}$$

and

$$\nu \coloneqq \sqrt{nN}\max_{k\in[n],j\in[N]}|w_{kj}|. \tag{12}$$

We defer the proof to Section V, and summarize the mathematical tools we use here. By the Davis-Kahan $\sin\theta$ Theorem [56], the error $d(\hat{\eta},\ \dot{\eta})$ in the eigenvector is bounded if there exists a sufficiently large spectral gap between the two largest (in terms of absolute values) eigenvalues of $G = \beta I - B$. We divide this task into two parts: 1) Show that there exists a large spectral gap in $\beta I - \mathbb{E}B$; 2) Prove that $\|B-\mathbb{E}B\|$ is small using concentration of measure inequalities, e.g., the matrix Bernstein inequality [57, Th. 1.6].

When m is large (e.g., $m \ge n$), (10) does not hold, hence the perturbation bound of the eigenvector $\hat{\eta}$ of G in Theorem 9 is no longer true. We can, however, bound the perturbation of the eigenvectors of submatrices of G *uniformly*.

Theorem 10 (Joint Sparsity Case): Let $\alpha = \sqrt{n}$ and $s = s_0 + 2s_1$, and suppose Assumptions 3 – 6 are satisfied with $\delta < 1/3$ and a sufficiently small absolute constant $C_W > 0$. Then there exist absolute constants $c, C, C' > 0$, such that if

$$\max\left\{\frac{(s+N)\log^8 n\log^2(sN+m)}{n}, \frac{\sqrt{s}\log^2 n\log(sN+m)}{N}, \frac{\log^4 n\log^2(sN+m)}{s_0}\right\} \le C, \tag{13}$$

then with probability at least $1-2n^{-c}-m^{-cs}$, for every T_η defined by (9),

$$d(\hat{\eta}_{T_\eta},\ \Omega_{T_\eta}\dot{\eta}) \le \widetilde{\Delta},$$

where

$$\widetilde{\Delta} \coloneqq \frac{8C'}{1-3\delta}\max\{N^{3/2}\nu,\ \nu^2\}, \tag{14}$$

and ν is defined in (12).

The main challenge in the joint sparsity case is that, instead of bounding the spectral norm of $B-\mathbb{E}B$, one must bound the "sparse" norm of $B-\mathbb{E}B$, i.e., the maximum spectral norm of all principal submatrices whose row (and column) support is T_η defined by (9). Since $B-\mathbb{E}B$ can be broken down into the sum of several terms, we give a uniform bound over all submatrices on each term. For any given term, we adopt one of two approaches, whichever provides a tighter bound: 1) We bound the spectral norm of an individual submatrix, and apply a union bound over all submatrices; 2) We use a variational form of the sparse norm, and apply a bound on the suprema of second-order chaos [58, Th. 2.3].

The error bounds for Algorithms 1 and 2 in the next section rely on Theorems 9 and 10, and existing analysis of power iteration [9] and truncated power iteration [11]. Additionally, the perturbation bounds in this section are of independent interest. In particular, Theorem 9 shows that if the assumptions and the prescribed sample complexities in (10) are satisfied, then with high probability the principal eigenvector $\hat{\eta}$ of G is an accurate estimate of the vector $\dot{\eta}$ that concatenates the unknown variables. It gives an error bound for any algorithm that finds the principal eigenvector of G. On the other hand, while Theorem 10 does not directly guarantee the success of any particular algorithm, it can be used to analyze other algorithms that find the sparse principal component of G, similar to the analysis of Algorithm 2 in Theorem 12.

C. Error Bounds for the Power Iteration Algorithms

In this section, we give performance guarantees for Algorithms 1 and 2 under the assumptions stated in Section III-A. Under the conditions in Theorem 11 (resp. Theorem 12), the iterates in Algorithm 1 (resp. Algorithm 2), in the noiseless case, converge linearly to the true solution. In the noisy case, the recovery error is proportional to the noise level.

Theorem 11 (Subspace Case): Suppose Assumptions 3 – 6 are satisfied with $\delta < 1/4$ and a sufficiently small absolute constant $C_W > 0$. Let $\alpha = \sqrt{n}$, and $\beta = 3/2$. Assume that $\xi \coloneqq |\hat{\eta}^*\eta^{(0)}| > 0$. Then there exist absolute constants $c, C, C' > 0$, such that if (10) is satisfied, then with probability at least $1-2n^{-c}-e^{-cm}$, the iterates in Algorithm 1 satisfy

$$d(\eta^{(t)},\ \dot{\eta}) \le \rho^t d(\eta^{(0)},\ \dot{\eta}) + 2\Delta,$$

where Δ is defined in (11), and

$$\rho \coloneqq \left\{1-\frac{1}{2}\left[1-\left(\frac{1+6\delta}{3-2\delta}\right)^2\right]\xi(1+\xi)\right\}^{1/2}. \tag{15}$$

Theorem 11 shows that the power iteration algorithm requires $n = O(m\log^2(Nm+n))$ sensors and $N = O(\log(Nm+n))$ snapshots to successfully recover X and λ. This agrees, up to log factors, with the sample complexity required for the uniqueness of (λ, X) in the subspace case, which is $n > m$ and $N \ge \frac{n-1}{n-m}$ [6].

Next, we compare Theorem 11 with a similar error bound for the least squares approach by Ling and Strohmer [7, Th. 3.5]. The sample complexity in Theorem 11 matches the numbers required by the least squares approach $n = O(m\log^2(Nm+n))$ and $N = O(\log^2(Nm+n))$ (up to one log factor). One caveat in the least squares approach is that, apart from the linear equation (2), it needs an extra linear constraint to avoid the trivial solution $\gamma = 0$, $X = 0$. Unfortunately, as revealed by [7, Th. 3.5], in the noisy setting, the recovery error by the least squares approach is sensitive to this extra linear constraint. Our numerical experiments (Section VI) show that power iteration outperforms least squares in the noisy setting.

Theorem 12 (Joint Sparsity Case): Suppose Assumptions 3 – 6 are satisfied with $\delta < 1/4$ and a sufficiently small

absolute constant $C_W > 0$. Let $\alpha = \sqrt{n}$, $\beta = 3/2$, $s_1 \geq s_0$ in Algorithm 2, and define $s = s_0 + 2s_1$. Then there exist absolute constants $c, C, C' > 0$, such that if $|\dot{\eta}^* \eta^{(0)}| \geq \xi + \widetilde{\Delta}$ for some $\xi \in (0, 1)$, and (13) is satisfied, then with probability at least $1 - 2n^{-c} - m^{-cs}$, the iterates in Algorithm 2 for the *joint* sparsity case satisfy

$$d(\eta^{(t)},\ \dot{\eta}) \leq \tilde{\rho}^t d(\eta^{(0)},\ \dot{\eta}) + \frac{2\sqrt{5}\widetilde{\Delta}}{1-\tilde{\rho}},$$

where $\widetilde{\Delta}$ is defined in (14), and $\tilde{\rho} < 1$ has the following expression:

$$\tilde{\rho} \coloneqq \rho \cdot \left(1 + 2\sqrt{\frac{s_0}{s_1}} + \frac{2s_0}{s_1}\right)^{1/2}, \tag{16}$$

and ρ is defined in (15).

Theorem 12 is only valid when $\tilde{\rho} < 1$. With the choice $s_1 = 2s_0$, when δ approaches 0, and ξ approaches 1, the convergence rate $\tilde{\rho}$ is roughly $\frac{1}{3}\sqrt{1+\sqrt{2}+2} \approx 0.62$. We discuss a more realistic scenario next.

Remark 13: A good initialization for λ alone is usually sufficient. Suppose one has a good initial estimate for the gains and phases, i.e., λ satisfies $|\lambda_k - e^{\sqrt{-1}\varphi_k}| < \sqrt{1+\delta} - 1$ for known phase estimates $\{\varphi_k\}_{k=1}^n$. One can initialize Algorithm 2 with $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, e^{-\sqrt{-1}\varphi_1}, \ldots, e^{-\sqrt{-1}\varphi_n}]^\top$, then when Δ is negligible (noiseless or low noise settings), ξ in Theorem 12 can be set to $1/\sqrt{(1+\delta)(2+\delta)}$. For example, if $\delta = 0.05$ and $s_1 \geq 10\ s_0$, then $\tilde{\rho} < 1$. Since we do not attempt to optimize the constants in this paper, the constants in this exemplary scenario are conservative.

Theorem 12 states that for Algorithm 2 to recover λ and a jointly sparse X, it is sufficient to have $n = O(s_0 \log^8 n \log^2(s_0 N + m))$ sensors and $N = O(\sqrt{s_0} \log^2 n \log(s_0 N + m))$ snapshots. In comparison, the (up to a factor of 2) optimal sample complexity for unique recovery in the joint sparsity case is $n > 2s_0$ and $N \geq \frac{n-1}{n-2s_0}$ [6]. Hence, the number of sensors required in Theorem 12 is (up to log factors) optimal, but the number of snapshots required is suboptimal. Another drawback is that these results apply only to the joint sparsity case, and not to the more general sparsity case. However, we believe these drawbacks are due to artifacts of our analysis.[3] For both the joint sparsity case and the sparsity case, we have Nn complex-valued measurements, and $Ns_0 + n - 1$ complex-valued unknowns. One may expect successful recovery when n and N are (up to log factors) on the order of s_0 and 1, respectively. In fact, numerical experiments in Section VI confirms that truncated power iteration successfully recovers λ and X in this regime for the more general sparsity case.

The assumption on the noise level in Theorem 12 (the joint sparsity case of Assumption 6) is demanding when compared to the subspace case. This is due to the limitations of our theoretical analysis, and the inherently maximally conservative nature of any worst-case guarantee. In fact, our experiments in Section VI-B show that Algorithm 2 performs well at a variety of reasonable noise levels. We think that the condition on the noise level can be relaxed by introducing more sophisticated concentration inequalities, but leave it as future work.

Wang and Chi [8] analyzed the performance of ℓ_1 minimization for BGPC in the sparsity case, where they assumed that A is the DFT matrix, and X is a Bernoulli-Subgaussian random matrix. Their sample complexity for ℓ_1 minimization is $n = O(s)$ and $N = O(n \log^4 n)$. The success of their algorithm relies on a restrictive assumption that $\lambda_k \approx 1$, which is analogous to the dependence of our algorithm on a good initialization of λ_k. In the next section, we show that such dependence can be relaxed under some additional conditions using the initialization provided by Algorithm 3.

D. A Theoretical Guarantee of the Initialization

The next theorem shows that, under certain conditions, Algorithm 3 recovers the locations of the large entries in X correctly, and yields an initial estimate $\eta^{(0)}$ that satisfies $|\dot{\eta}^* \eta^{(0)}| > 1 - 2\delta$ (close to 1).

Theorem 14 (Initialization): Suppose Assumptions 3 – 7 are satisfied. Then there exist absolute constants $C'', c'' > 0$, such that if

$$n > C'' s_0^2 \log^6(nmN),$$

then with probability at least $1 - n^{-c''}$, for all $j \in [N]$ the set T_j' in Assumption 7 is a subset of T_j in Algorithm 3. Additionally, in the joint sparsity case, if sample complexity (13) is satisfied with a sufficiently large C, Assumption 6 is satisfied with a sufficiently small C_W, and Assumption 7 is satisfied with a sufficiently small δ_X, then η_0 produced by Algorithm 3 will satisfy that $|\dot{\eta}^* \eta^{(0)}|$ is arbitrarily close to

$$\frac{n^{3/2} + \|\lambda\|_2 \|\gamma\|_2^2}{\sqrt{n^2 + \|\lambda\|_2^2 \|\gamma\|_2^2}\sqrt{n + \|\gamma\|_2^2}} > 1 - 2\delta.$$

By Theorem 14, the constant ξ in Theorem 12 can be set to $1 - 2\delta$ in a low noise setting. For $\delta < 0.19$, this constant ξ is larger than the one in Remark 13, and allows $\tilde{\rho} < 1$ for more choices of s_1.

Our guarantee for the initialization requires that the number n of sensors scales quadratically (up to log factors) in the sparsity s_0, which seems suboptimal. Similar suboptimal sampling complexities show up in sparse PCA [59] and sparse phase retrieval [42], [44], [60].

In the joint sparsity case, instead of estimating the supports of $x_{\cdot 1}, x_{\cdot 2}, \ldots, x_{\cdot N}$ separately, one can estimate the row support of X directly by sorting $\sum_{j \in [N]} \|d^*_{\cdot((j-1)m+\ell)} E\|_2^2$ for $\ell \in [m]$ and finding the s_1 largest. In this case, Assumption 7 can be changed to: There exists a subset T' of large rows (in terms of ℓ_2 norm), such that for all $\ell \in T'$,

$$\frac{\sum_{j \in [N]} |x_{\ell j}|^2}{\|X\|_{\mathrm{F}}^2} \geq \frac{\omega}{s_0},$$

and

$$\frac{\sum_{j \in [N], \ell' \in [m] \setminus T'} |x_{\ell' j}|^2}{\|X\|_{\mathrm{F}}^2} \leq \delta_X.$$

[3] Our analysis of the joint sparsity case does not trivially generalize to the sparsity case. A key result – the supremum of the second-order chaos (defined in the proof of Lemma 20) – is bounded by a small constant for large n and N in the joint sparsity case. However, in the general sparsity case, the supremum must be taken over a much larger set, and grows unbounded.

In this case, the subset T' can be identified and an initialization $\eta^{(0)}$ can be computed under the same conditions as in Theorem 14, which can be proved using the same arguments.

IV. Fundamental Estimates

To prove the main results, we must first establish some fundamental estimates specific to BGPC. Proofs of some lemmas in this section can be found in the appendix.

A. A Gap in Eigenvalues

A key component in establishing a perturbation bound for an eigenvector problem (e.g., Theorem 9) is bounding the gap between eigenvalues. Lemma 15 gives us such a bound.

Lemma 15: Suppose Assumptions 3 – 5 are satisfied and $\alpha = \sqrt{n}$. Then the smallest eigenvalue of $\mathbb{E}\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^*$ is 0, and the rest of the eigenvalues reside in the interval $[\frac{(1-\delta)^2}{1+\delta}, 2(1+\delta)]$.

B. Perturbation Due to Randomness in A

Next, we show that $\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^*$, whose randomness comes from A, is close to its mean $\mathbb{E}\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^*$ under certain conditions.

Lemma 16: Suppose Assumptions 3 – 5 are satisfied, and $\alpha = \sqrt{n}$. For any constant $\delta_B > 0$, there exist absolute constants $C, c > 0$, such that:

- **Subspace case:** If (10) is satisfied with C, then
$$\|B_{\mathrm{s}} - \mathbb{E}B_{\mathrm{s}}\| \le \delta_B$$
with probability at least $1 - n^{-c} - e^{-cm}$.
- **Joint sparsity case:** If (13) is satisfied with C, then
$$\|\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^* - \mathbb{E}\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^*\| \le \delta_B$$
for all $T_1 = \cdots = T_N$ and T_η defined in (9), with probability at least $1 - n^{-c} - m^{-cs}$.

Proof of Lemma 16: Recall that
$$\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^* = \begin{bmatrix} \Omega_{T_x} D^* D \Omega_{T_x}^* & \sqrt{n}\Omega_{T_x} D^* E_{\mathrm{s}} \\ \sqrt{n} E_{\mathrm{s}}^* D \Omega_{T_x}^* & n E_{\mathrm{s}}^* E_{\mathrm{s}} \end{bmatrix}.$$
It follows that
$$\begin{aligned} &\|\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^* - \mathbb{E}\Omega_{T_\eta} B_{\mathrm{s}} \Omega_{T_\eta}^*\| \\ &\quad \le \|\Omega_{T_x} D^* D \Omega_{T_x}^* - \mathbb{E}\Omega_{T_x} D^* D \Omega_{T_x}^*\| && (17) \\ &\qquad + n\|E_{\mathrm{s}}^* E_{\mathrm{s}} - \mathbb{E}E_{\mathrm{s}}^* E_{\mathrm{s}}\| && (18) \\ &\qquad + 2\sqrt{n}\|\Omega_{T_x} D^* E_{\mathrm{s}} - \mathbb{E}\Omega_{T_x} D^* E_{\mathrm{s}}\|. && (19) \end{aligned}$$
Lemma 16 follows from the bounds on the spectral norms in (17)–(19) in Lemmas 17–20, respectively. ∎

Lemma 17: Suppose Assumption 3 is satisfied, then there exist absolute constants $C_1, c_1 > 0$, such that:

- **Subspace case:**
$$\|D^* D - \mathbb{E}D^* D\| \le C_1\sqrt{\frac{m}{n}},$$
with probability at least $1 - e^{-c_1 m}$.
- **Joint sparsity case:** For any $\{T_j\}_{j=1}^N$ and T_x defined in (8),
$$\|\Omega_{T_x} D^* D \Omega_{T_x}^* - \mathbb{E}\Omega_{T_x} D^* D \Omega_{T_x}^*\| \le C_1\sqrt{\frac{s}{n}\log m},$$
with probability at least $1 - m^{-c_1 s}$.

Lemma 18: Suppose Assumptions 3 – 5 are satisfied, then there exist absolute constants $C_2, c_2 > 0$, such that

- **Subspace case:**
$$\|E_{\mathrm{s}}^* E_{\mathrm{s}} - \mathbb{E}E_{\mathrm{s}}^* E_{\mathrm{s}}\| \le \frac{C_2}{n}\max\Big\{\sqrt{\frac{\log n}{N}}, \sqrt{\frac{\log n}{m}}, \frac{\log n}{N}, \frac{\log n}{m}\Big\}$$
- **Joint sparsity case:**
$$\|E_{\mathrm{s}}^* E_{\mathrm{s}} - \mathbb{E}E_{\mathrm{s}}^* E_{\mathrm{s}}\| \le \frac{C_2}{n}\max\Big\{\sqrt{\frac{\log n}{N}}, \sqrt{\frac{\log n}{s_0}}, \frac{\log n}{N}, \frac{\log n}{s_0}\Big\}$$

with probability at least $1 - n^{-c_2}$.

Lemma 19 (Subspace Case): Suppose Assumptions 3 – 5 are satisfied, and $\min\{N, m\} > \log n$, then there exist absolute constants $C_3, c_3 > 0$, such that
$$\|D^* E_{\mathrm{s}} - \mathbb{E}D^* E_{\mathrm{s}}\| \le C_3\max\Big\{\sqrt{\frac{\log(Nm+n)}{nN}}, \sqrt{\frac{\log(Nm+n)}{nm}}, \frac{\sqrt{m}\log(Nm+n)}{n}\Big\}$$
with probability at least $1 - n^{-c_3}$.

Lemma 20 (Joint Sparsity Case): Suppose Assumptions 3 – 5 are satisfied, then there exist absolute constants $C_3, c_3 > 0$, such that for all $T_1 = \cdots = T_N$,
$$\begin{aligned} &\|\Omega_{T_x} D^* E_{\mathrm{s}} - \mathbb{E}\Omega_{T_x} D^* E_{\mathrm{s}}\| \\ &\quad \le \frac{C_3\, s_0^{1/4}(s+N)^{1/4}(\sqrt{n}+\sqrt{s+N})^{1/2}}{n\min\{\sqrt{s_0}, \sqrt{N}\}} \\ &\qquad \times \log^3 n\log(sN+m), \end{aligned}$$
with probability at least $1 - n^{-c_3}$.

C. Perturbation Due to Noise

We established some fundamental estimates regarding B_{s} in Sections IV-A and IV-B. In this section, we turn to perturbation caused by noise. By the definitions of B, B_{s}, E, E_{s}, and E_{n}, we have
$$B = B_{\mathrm{s}} + B_{\mathrm{n}},$$
where
$$B_{\mathrm{n}} \coloneqq \begin{bmatrix} 0 & \alpha D^* E_{\mathrm{n}} \\ \alpha E_{\mathrm{n}}^* D & \alpha^2(E_{\mathrm{s}}^* E_{\mathrm{n}} + E_{\mathrm{n}}^* E_{\mathrm{s}} + E_{\mathrm{n}}^* E_{\mathrm{n}}) \end{bmatrix}.$$
Therefore,
$$\Omega_{T_\eta} B_{\mathrm{n}} \Omega_{T_\eta}^* = \begin{bmatrix} 0 & \alpha\Omega_{T_x} D^* E_{\mathrm{n}} \\ \alpha E_{\mathrm{n}}^* D\Omega_{T_x}^* & \alpha^2(E_{\mathrm{s}}^* E_{\mathrm{n}} + E_{\mathrm{n}}^* E_{\mathrm{s}} + E_{\mathrm{n}}^* E_{\mathrm{n}}) \end{bmatrix}.$$

Lemma 21 gives an upper bound on the spectral norm of the perturbation from noise.

Lemma 21: Suppose Assumptions 3 – 5 are satisfied. Let $\alpha = \sqrt{n}$ and let ν be defined by (12). Then there exist absolute constants $c, C, C' > 0$ such that:

- **Subspace case:** If (10) is satisfied, then with probability at least $1 - n^{-c}$

$$\|B_{\mathrm{n}}\| \leq C' \max\{\nu,\ \nu^2\}.$$

  Additionally, for any constant $\delta_W > 0$, there exists an absolute constant $C_W > 0$, if Assumption 6 is satisfied with C_W, then the above bound becomes

$$\|B_{\mathrm{n}}\| \leq \delta_W.$$

- **Joint sparsity case:** If (13) is satisfied, then with probability at least $1 - n^{-c}$

$$\|\Omega_{T_\eta} B_{\mathrm{n}} \Omega_{T_\eta}^*\| \leq C' \max\{N^{3/2}\nu,\ \nu^2\}$$

  for all $T_1 = \cdots = T_N$ and T_η defined in (9). Additionally, for any constant $\delta_W > 0$, there exists an absolute constant $C_W > 0$, if Assumption 6 is satisfied with C_W, then the above bound becomes

$$\|\Omega_{T_\eta} B_{\mathrm{n}} \Omega_{T_\eta}^*\| \leq \delta_W.$$

Proof: To complete the proof, we bound the spectral norms of $\Omega_{T_x} D^* E_{\mathrm{n}}$, $E_{\mathrm{s}}^* E_{\mathrm{n}}$, and $E_{\mathrm{n}}^* E_{\mathrm{n}}$ in Lemmas 22, 24, and 25, respectively. ∎

Lemma 22 (Subspace Case): Suppose Assumption 3 is satisfied, and $m > \log n$, then there exist absolute constants $C_4, c_4 > 0$, such that

$$\|D^* E_{\mathrm{n}}\| \leq C_4 \max\left\{\sqrt{\log(Nm+n)}, \sqrt{\frac{Nm}{n}}\log(Nm+n)\right\} \max_{k\in[n], j\in[N]} |w_{kj}|,$$

with probability at least $1 - n^{-c_4}$.

Lemma 23 (Joint Sparsity Case): Suppose Assumption 3 is satisfied, then there exist absolute constants $C_4, c_4 > 0$, such that for all $T_1 = \cdots = T_N$,

$$\|\Omega_{T_x} D^* E_{\mathrm{n}}\| \leq C_4(\sqrt{s}N + \sqrt{sN\log m} + \sqrt{N\log^3 n}) \times \sqrt{\log n} \max_{k\in[n], j\in[N]} |w_{kj}|,$$

with probability at least $1 - n^{-c_4}$.

Lemma 24: Suppose Assumptions 3 – 5 are satisfied, then there exist absolute constants $C_5, c_5 > 0$, such that

- **Subspace case:**

$$\|E_{\mathrm{s}}^* E_{\mathrm{n}}\| \leq C_5\sqrt{\frac{N}{n}} \max\left\{1, \sqrt{\frac{\log n}{N}}, \sqrt{\frac{\log n}{m}}\right\} \times \max_{k\in[n], j\in[N]} |w_{kj}|,$$

- **Joint sparsity case:**

$$\|E_{\mathrm{s}}^* E_{\mathrm{n}}\| \leq C_5\sqrt{\frac{N}{n}} \max\left\{1, \sqrt{\frac{\log n}{N}}, \sqrt{\frac{\log n}{s_0}}\right\} \times \max_{k\in[n], j\in[N]} |w_{kj}|,$$

  with probability at least $1 - n^{-c_5}$.

Lemma 25:

$$\|E_{\mathrm{n}}^* E_{\mathrm{n}}\| \leq N \max_{k\in[n], j\in[N]} |w_{kj}|^2,$$

D. Scalar Concentration

We now introduce a few scalar concentration bounds that are useful in the proof of Theorem 14.

Lemma 26: Suppose Assumptions 3 – 6 is satisfied, then there exist absolute constants $C_6, c_6 > 0$, such that for all $j \in [N]$ and $\ell \in [m]$, we have

$$\left|\sum_{k\in[n]} \left(|\lambda_k \overline{a_{k\ell}} a_{k\cdot}^\top x_{\cdot j}|^2 - \mathbb{E}|\lambda_k \overline{a_{k\ell}} a_{k\cdot}^\top x_{\cdot j}|^2\right)\right| \leq \frac{C_6 \|x_{\cdot j}\|_2^2 \log^3(nmN)}{n^{3/2}}, \tag{20}$$

$$\left|\sum_{k\in[n]} \lambda_k a_{k\ell} \overline{a_{k\ell}} a_{k\cdot}^\top x_{\cdot j} \overline{w_{kj}}\right| \leq \frac{C_6 \|x_{\cdot j}\|_2 \log^2(nmN)}{n} \max_{k\in[n], j\in[N]} |w_{kj}| \leq \frac{C_6 C_W \|x_{\cdot j}\|_2^2 \log^2(nmN)}{\sqrt{1-\theta} n^{3/2}}, \tag{21}$$

and

$$\left|\sum_{k\in[n]} \left(|\overline{a_{k\ell}} w_{kj}|^2 - \mathbb{E}|\overline{a_{k\ell}} w_{kj}|^2\right)\right| \leq \frac{C_6 \log^2(nmN)}{n^{1/2}} \max_{k\in[n], j\in[N]} |w_{kj}|^2 \leq \frac{C_6 C_W^2 \|x_{\cdot j}\|_2^2 \log^2(nmN)}{(1-\theta) n^{3/2}}, \tag{22}$$

with probability at least $1 - n^{-c_6}$.

V. Proofs of the Main Results

A. Proof of the Perturbation Bound for the Eigenvector Problem

In this section, we prove Theorem 9. The proof centers around the Davis-Kahan $\sin\theta$ Theorem [56], which bounds the error in the principal eigenvector $\dot{\eta}$ of G using the perturbation of G. The spectral norm of the perturbed matrix is in turn bounded by the lemmas in Section IV. Theorem 10 can be proved similarly.

Proof of Theorem 9: First,

$$G = \beta I_{Nm+n} - B = (\beta I_{Nm+n} - \mathbb{E}B_{\mathrm{s}}) - (B_{\mathrm{s}} - \mathbb{E}B_{\mathrm{s}}) - B_{\mathrm{n}}. \tag{23}$$

Lemma 15 establishes a gap in the eigenvalues of the matrix $\mathbb{E}B_{\mathrm{s}}$ – the smallest and the second-smallest eigenvalues of $\mathbb{E}B_{\mathrm{s}}$ are separated by a gap of at least

$$\frac{(1-\delta)^2}{1+\delta} \geq 1 - 3\delta > 0.$$

Therefore, the gap between the largest and the second-largest eigenvalues of $\beta I_{Nm+n} - \mathbb{E}B_{\mathrm{s}}$ is at least $1 - 3\delta$. By Lemmas 16 and 21, there exist absolute constants $c, C, C', C_W > 0$ such

that if all the assumptions are satisfied, then with probability at least $1-2n^{-c}-e^{-cm}$,

$$\|(B_{\rm s}-\mathbb{E}B_{\rm s})+B_{\rm n}\| \le \|B_{\rm s}-\mathbb{E}B_{\rm s}\|+\|B_{\rm n}\| \le \frac{1-3\delta}{4}, \quad (24)$$

$$\|B_{\rm n}\| \le C'\max\{\nu,\ \nu^2\}. \quad (25)$$

Recall that $\dot\eta$ is the principal eigenvector of $\beta I_{Nm+n}-\mathbb{E}B_{\rm s}$. By the Davis-Kahan $\sin\theta$ Theorem ([56]; see also [9, Th. 8.1.12]), (24) and (25) imply

$$\begin{aligned}\sin\angle(\dot\eta,\hat\eta) &\le \frac{4}{1-3\delta}\|(B_{\rm s}-\mathbb{E}B_{\rm s}+B_{\rm n})\dot\eta\|_2\\ &\le \frac{4}{1-3\delta}\|B_{\rm n}\|\\ &\le \frac{4C'}{1-3\delta}\max\{\nu,\ \nu^2\},\end{aligned}$$

where the second inequality is due to $B_{\rm s}\dot\eta=\mathbb{E}B_{\rm s}\dot\eta=0$.

Theorem 9 follows from the above bound, and the fact that

$$\begin{aligned}d(\dot\eta,\hat\eta)=\sqrt{2-2\cos\angle(\dot\eta,\hat\eta)}&=2\sin\frac{\angle(\dot\eta,\hat\eta)}{2}\\ &\le 2\sin\angle(\dot\eta,\hat\eta).\end{aligned}$$

■

One can prove Theorem 10 using the same steps as in the proof of Theorem 9, by restricting rows and columns of matrices to the support T_η and applying the corresponding uniform bounds on submatrices.

B. Proof of the Error Bound for Algorithm 1

We prove Theorem 11 by following a standard analysis of power iteration. The specific convergence rates and error bounds for the BGPC problem follow from the lemmas in Section IV and Theorem 9.

Proof of Theorem 11: Recall that the largest eigenvalue of $\beta I_{Nm+n}-\mathbb{E}B_{\rm s}$ is $\beta-0=\frac{3}{2}$, and all other eigenvalues reside in the interval $[\frac{3}{2}-2(1+\delta),\ \frac{3}{2}-\frac{(1-\delta)^2}{1+\delta}]$. By Lemmas 16 and 21, there exist constants $c, C, C_W>0$ such that

$$\begin{aligned}\|(B_{\rm s}-\mathbb{E}B_{\rm s})+B_{\rm n}\| &\le \|B_{\rm s}-\mathbb{E}B_{\rm s}\|+\|B_{\rm n}\|\\ &\le \min\Big\{\delta,\ \frac{(1-\delta)^2}{1+\delta}+3\delta-1\Big\},\end{aligned}$$

with probability at least $1-2n^{-c}-e^{-cm}$. By (23), the largest eigenvalue of G is $\|G\|\ge\frac{3}{2}-\delta$, the corresponding eigenvector is $\hat\eta$, and all the other eigenvalues of G reside in the interval $[-\frac{1}{2}-3\delta,\ \frac{1}{2}+3\delta]$.

Next, we establish the convergence rate of power iterations for BGPC. By the eigenvalue decomposition of G and the Pythagorean theorem,

$$\begin{aligned}&G\hat\eta=\|G\|\hat\eta,\\ &\|G\eta^{(t-1)}\|\le\sqrt{\|G\|^2|\hat\eta^*\eta^{(t-1)}|^2+\Big(\frac{1}{2}+3\delta\Big)^2(1-|\hat\eta^*\eta^{(t-1)}|^2)}.\end{aligned}$$

Therefore,

$$\begin{aligned}|\hat\eta^*\eta^{(t)}| &= \frac{|\hat\eta^*G\eta^{(t-1)}|}{\|G\eta^{(t-1)}\|_2}\\ &\ge \frac{\|G\||\hat\eta^*\eta^{(t-1)}|}{\sqrt{\|G\|^2|\hat\eta^*\eta^{(t-1)}|^2+(\frac{1}{2}+3\delta)^2(1-|\hat\eta^*\eta^{(t-1)}|^2)}}\\ &\ge |\hat\eta^*\eta^{(t-1)}|\frac{1}{\sqrt{|\hat\eta^*\eta^{(t-1)}|^2+(\frac{1+6\delta}{3-2\delta})^2(1-|\hat\eta^*\eta^{(t-1)}|^2)}}\\ &= |\hat\eta^*\eta^{(t-1)}|\frac{1}{\sqrt{1-\big(1-(\frac{1+6\delta}{3-2\delta})^2\big)(1-|\hat\eta^*\eta^{(t-1)}|^2)}}\\ &\ge |\hat\eta^*\eta^{(t-1)}|\Big[1+\frac{1}{2}\big(1-(\frac{1+6\delta}{3-2\delta})^2\big)(1-|\hat\eta^*\eta^{(t-1)}|^2)\Big],\end{aligned}$$

where the last inequality is due to $\frac{1}{\sqrt{1-z}}\ge 1+\frac{1}{2}z$ for $z\in(0,1)$. It follows that

$$\begin{aligned}&[1-|\hat\eta^*\eta^{(t)}|]\\ &\le [1-|\hat\eta^*\eta^{(t-1)}|]\\ &\quad\times\Big[1-\frac{1}{2}\big(1-(\frac{1+6\delta}{3-2\delta})^2\big)|\hat\eta^*\eta^{(t-1)}|(1+|\hat\eta^*\eta^{(t-1)}|)\Big]. \quad (26)\end{aligned}$$

Clearly, $\{|\hat\eta^*\eta^{(\tau)}|\}_{\tau=0}^t$ is monotonically increasing unless $|\hat\eta^*\eta^{(0)}|=0$. By the definition $\xi:=|\hat\eta^*\eta^{(0)}|$, the convergence rate in (26) is bounded by $\rho^2<1$. It follows that

$$\begin{aligned}[1-|\hat\eta^*\eta^{(t)}|] &\le \rho^2[1-|\hat\eta^*\eta^{(t-1)}|]\\ &\le \rho^{2t}[1-|\hat\eta^*\eta^{(0)}|].\end{aligned}$$

Hence

$$d(\hat\eta,\eta^{(t)})\le\rho^t d(\hat\eta,\eta^{(0)}).$$

By Theorem 9, for $\tau=0,\dots,t$

$$d(\dot\eta,\hat\eta)\le\Delta.$$

It follows from the triangle inequality that

$$d(\dot\eta,\eta^{(t)})\le\rho^t d(\dot\eta,\eta^{(0)})+2\Delta.$$

■

C. Proof of the Error Bound for Algorithm 2

We prove Theorem 12 using the perturbation bound in Theorem 10, and by following steps similar to the theoretical analysis of truncated power iteration for a generic sparse eigenvector problem [11]. The proof consists of two steps: (1) the estimate after the power iteration step $\tilde\eta^{(t)}$ is closer to the ground truth $\dot\eta$ than the last sparse estimate $\eta^{(t-1)}$, by a factor of ρ; (2) the truncation step amplifies the estimation error by a factor no larger than $(1+2\sqrt{\frac{s_0}{s_1}+\frac{2s_0}{s_1}})^{1/2}$. Therefore, the estimation error of the iterates in Algorithm 2 decays at the rate specified in Theorem 12.

Proof of Theorem 12: In the joint sparsity case, any iterate $\eta^{(\tau)}=[x^{(\tau)\top},-\gamma^{(\tau)\top}/\alpha]^\top$ satisfies that $x^{(\tau)}$ is the concatenation of jointly sparse $\{x_{\cdot j}^{(\tau)}\}_{j=1}^N$. In the t-th iteration, we define a support set $T^{(t)}$ that has cardinality $s=s_0+2s_1$, and satisfies

$$\mathrm{supp}(x_{\cdot j})\bigcup\mathrm{supp}(x_{\cdot j}^{(t-1)})\bigcup\mathrm{supp}(x_{\cdot j}^{(t)})\subset T^{(t)},$$

for all $j \in [N]$. Define $T_\eta^{(t)}$ using (8) and (9) with $T_1 = \cdots = T_N = T^{(t)}$. Next, we focus on the submatrix $\Omega_{T_\eta^{(t)}} G \Omega^*_{T_\eta^{(t)}}$ and subvectors $\Omega_{T_\eta^{(t)}} \dot{\eta}$ and $\Omega_{T_\eta^{(t)}} \eta^{(t)}$, etc. Since the supports of $\eta^{(t)}$ and $\dot{\eta}$ are subsets of $T_\eta^{(t)}$, we have $|\dot{\eta}^* \Omega^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t)}| = |\dot{\eta}^* \eta^{(t)}|$.

We prove by induction that $\{|\dot{\eta}^* \eta^{(\tau)}|\}_{\tau=0}^t$ is monotonically increasing (until it crosses a threshold specified later in the proof). Suppose $\{|\dot{\eta}^* \eta^{(\tau)}|\}_{\tau=0}^{t-1}$ is monotonically increasing. Next, we prove

$$|\dot{\eta}^* \eta^{(t)}| > |\dot{\eta}^* \eta^{(t-1)}|.$$

By the assumption that $|\dot{\eta}^* \eta^{(0)}| \geq \xi + \widetilde{\Delta}$ and Theorem 10 (which applies to all T_η defined by (9), including $T_\eta^{(t)}$ for all t), we have

$$\begin{aligned} |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}| &\geq |\dot{\eta}^* \eta^{(t-1)}| - d(\Omega_{T_\eta^{(t)}} \dot{\eta},\ \hat{\eta}_{T_\eta^{(t)}}) \\ &\geq \xi + \widetilde{\Delta} - \widetilde{\Delta} \\ &= \xi. \end{aligned}$$

Following the same steps in the proof of Theorem 11, we obtain a bound for $\tilde{\eta}^{(t)}$ similar to (26):

$$\begin{aligned} &[1 - |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \tilde{\eta}^{(t)}|] \\ &\leq [1 - |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}|]\Big[1 - \frac{1}{2}\big(1 - (\frac{1+6\delta}{3-2\delta})^2\big) \\ &\quad \times |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}|(1 + |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}|)\Big] \\ &\leq [1 - |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}|]\Big[1 - \frac{1}{2}\big(1 - (\frac{1+6\delta}{3-2\delta})^2\big)\xi(1+\xi)\Big] \\ &= \rho^2 [1 - |\hat{\eta}^*_{T_\eta^{(t)}} \Omega_{T_\eta^{(t)}} \eta^{(t-1)}|], \end{aligned}$$

where ρ is defined in (15). It follows that

$$d(\hat{\eta}_{T_\eta^{(t)}},\ \Omega_{T_\eta^{(t)}} \tilde{\eta}^{(t)}) \leq \rho \cdot d(\hat{\eta}_{T_\eta^{(t)}},\ \Omega_{T_\eta^{(t)}} \eta^{(t-1)}).$$

We use the perturbation bound in Theorem 10 one more time:

$$d(\Omega_{T_\eta^{(t)}} \dot{\eta},\ \Omega_{T_\eta^{(t)}} \tilde{\eta}^{(t)}) \leq \rho \cdot d(\Omega_{T_\eta^{(t)}} \dot{\eta},\ \Omega_{T_\eta^{(t)}} \eta^{(t-1)}) + 2\widetilde{\Delta}.$$

Equivalently,

$$\sqrt{1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|} \leq \rho \sqrt{1 - |\dot{\eta}^* \eta^{(t-1)}|} + \sqrt{2}\widetilde{\Delta}. \tag{27}$$

Next, we show that the truncation step amplifies the error only by a small factor. The vector $\widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})$ is the projection of $\tilde{\eta}^{(t)}$ onto the set of structured sparse vectors, and $\eta^{(t)}$ is the normalized version. We define three index sets

$$\begin{aligned} T_a &= \mathrm{supp}(\dot{\eta}) \backslash \mathrm{supp}(\eta^{(t)}), \\ T_b &= \mathrm{supp}(\dot{\eta}) \bigcap \mathrm{supp}(\eta^{(t)}), \\ T_c &= \mathrm{supp}(\eta^{(t)}) \backslash \mathrm{supp}(\dot{\eta}). \end{aligned}$$

By the Cauchy-Schwarz inequality,

$$\begin{aligned} |\dot{\eta}^* \tilde{\eta}^{(t)}|^2 &\leq \|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2^2 + \|\Omega_{T_b} \tilde{\eta}^{(t)}\|_2^2 \\ &\leq 1 - \|\Omega_{T_c} \tilde{\eta}^{(t)}\|_2^2 \\ &\leq 1 - \frac{|T_c|}{|T_a|} \|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2^2, \end{aligned}$$

where the last inequality is due to projection rule, i.e., $\widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})$ keeps the largest entries of $\tilde{\eta}^{(t)}$ (in the part corresponding to x). Since $|T_c|/|T_a| \geq s_1/s_0$, we have

$$\|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 \leq \sqrt{\frac{s_0}{s_1}(1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|^2)}. \tag{28}$$

Also by the Cauchy-Schwarz inequality,

$$\begin{aligned} |\dot{\eta}^* \tilde{\eta}^{(t)}|^2 &\leq (\|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 \|\Omega_{T_a} \dot{\eta}\|_2 + \|\Omega_{T_b} \tilde{\eta}^{(t)}\|_2 \|\Omega_{T_b} \dot{\eta}\|_2)^2 \\ &\leq \Big(\|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 \|\Omega_{T_a} \dot{\eta}\|_2 \\ &\qquad + \sqrt{1 - \|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2^2} \sqrt{1 - \|\Omega_{T_a} \dot{\eta}\|_2^2}\ \Big)^2 \\ &\leq 1 - (\|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 - \|\Omega_{T_a} \dot{\eta}\|_2)^2. \end{aligned}$$

It follows that

$$\|\Omega_{T_a} \dot{\eta}\|_2 \leq \|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 + \sqrt{1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|^2}. \tag{29}$$

By (28) and (29),

$$\begin{aligned} |\dot{\eta}^* \tilde{\eta}^{(t)}| - |\dot{\eta}^* \widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})| &\leq |\dot{\eta}^* \big(\tilde{\eta}^{(t)} - \widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})\big)| \\ &= \|\Omega_{T_a} \tilde{\eta}^{(t)}\|_2 \|\Omega_{T_a} \dot{\eta}\|_2 \\ &\leq \Big(\sqrt{\frac{s_0}{s_1}} + \frac{s_0}{s_1}\Big)(1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|^2). \end{aligned} \tag{30}$$

By (27) and (30),

$$\begin{aligned} &\sqrt{1 - |\dot{\eta}^* \eta^{(t)}|} \\ &\leq \sqrt{1 - |\dot{\eta}^* \widetilde{\Pi}_{s_1}(\tilde{\eta}^{(t)})|} \\ &\leq \sqrt{1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|} \sqrt{1 + \Big(\sqrt{\frac{s_0}{s_1}} + \frac{s_0}{s_1}\Big)(1 + |\dot{\eta}^* \tilde{\eta}^{(t)}|)} \\ &\leq \sqrt{1 - |\dot{\eta}^* \tilde{\eta}^{(t)}|} \sqrt{1 + 2\Big(\sqrt{\frac{s_0}{s_1}} + \frac{s_0}{s_1}\Big)} \\ &\leq \rho \sqrt{1 + 2\sqrt{\frac{s_0}{s_1}} + \frac{2s_0}{s_1}} \sqrt{1 - |\dot{\eta}^* \eta^{(t-1)}|} + \sqrt{10}\widetilde{\Delta} \\ &\leq \tilde{\rho} \sqrt{1 - |\dot{\eta}^* \eta^{(t-1)}|} + \sqrt{10}\widetilde{\Delta}. \end{aligned}$$

Therefore, $\{|\dot{\eta}^* \eta^{(\tau)}|\}_{\tau=0}^t$ indeed monotonically increases unless $\sqrt{1 - |\dot{\eta}^* \eta^{(\tau)}|}$ reaches $\sqrt{10}\widetilde{\Delta}/(1 - \tilde{\rho})$ for some τ. The proof by induction is complete.

It follows that

$$\sqrt{1 - |\dot{\eta}^* \eta^{(t)}|} \leq \tilde{\rho}^t \sqrt{1 - |\dot{\eta}^* \eta^{(0)}|} + \frac{\sqrt{10}\widetilde{\Delta}}{1 - \tilde{\rho}},$$

or equivalently

$$d(\dot{\eta}, \eta^{(t)}) \leq \tilde{\rho}^t d(\dot{\eta}, \eta^{(0)}) + \frac{2\sqrt{5}\widetilde{\Delta}}{1 - \tilde{\rho}}.$$

■

D. Proof of the Guarantee for Algorithm 3

Proof of Theorem 14: We first show that, under the conditions in Theorem 14, the support T_j in Algorithm 3 contains $T_j' \subset \mathrm{supp}(x_{\cdot j})$ in Assumption 7. To this end, we prove that the norms of the rows of $D^* E$ indexed by T_j' are larger than those outside $\mathrm{supp}(x_{\cdot j})$. For a fixed $j \in [N]$, the j-th block of

D^*E is indexed by the set $(j-1)m+[m]$. Therefore, the goal is to show that

$$\min_{\ell\in T_j'} \|d^*_{\cdot((j-1)m+\ell)}E\|_2^2 > \max_{\ell'\in[m]\setminus\mathrm{supp}(x_{\cdot j})} \|d^*_{\cdot((j-1)m+\ell')}E\|_2^2,$$

or equivalently,

$$\min_{\ell\in T_j'} \sum_{k\in[n]} |\overline{a_{k\ell}}y_{kj}|^2 > \max_{\ell'\in[m]\setminus\mathrm{supp}(x_{\cdot j})} \sum_{k\in[n]} |\overline{a_{k\ell'}}y_{kj}|^2.$$

Since

$$\mathbb{E}|\overline{a_{k\ell}}y_{kj}|^2 = \frac{1}{n^2}|\lambda_k|^2(\|x_{\cdot j}\|_2^2 + |x_{\ell j}|^2) + \frac{1}{n}|w_{kj}|^2,$$

it suffices to show that for all $\ell \in T_j'$ and $\ell'' \in [m]$,

$$\frac{1}{n^2}\sum_{k\in[n]} |\lambda_k|^2|x_{\ell j}|^2 > 2\left|\sum_{k\in[n]}\left(|\overline{a_{k\ell''}}y_{kj}|^2 - \mathbb{E}|\overline{a_{k\ell''}}y_{kj}|^2\right)\right|. \tag{31}$$

Recall that

$$y_{kj} = \lambda_k a_{k\cdot}^\top x_{\cdot j} + w_{kj}.$$

By the triangle inequality and Lemma 26, for all $j \in [N]$ and $\ell \in [m]$,

$$\begin{aligned}
&\left|\sum_{k\in[n]}\left(|\overline{a_{k\ell}}y_{kj}|^2 - \mathbb{E}||\overline{a_{k\ell}}y_{kj}|^2\right)\right| \\
&\quad\le \left|\sum_{k\in[n]}\left(|\lambda_k\overline{a_{k\ell}}a_{k\cdot}^\top x_{\cdot j}|^2 - \mathbb{E}|\lambda_k\overline{a_{k\ell}}a_{k\cdot}^\top x_{\cdot j}|^2\right)\right| \\
&\qquad + 2\left|\sum_{k\in[n]}\mathrm{Re}\left(\lambda_k a_{k\ell}\overline{a_{k\ell}}a_{k\cdot}^\top x_{\cdot j}\overline{w_{kj}}\right)\right| \\
&\qquad + \left|\sum_{k\in[n]}\left(|\overline{a_{k\ell}}w_{kj}|^2 - \mathbb{E}|\overline{a_{k\ell}}w_{kj}|^2\right)\right| \\
&\quad\le C_6\left(1+\frac{C_W}{\sqrt{1-\theta}}\right)^2\frac{\|x_{\cdot j}\|_2^2\log^3(nmN)}{n^{3/2}},
\end{aligned}$$

with probability at least $1-n^{-c_6}$.

By Assumptions 4 and 7, if we plug the above result into (31), then the following sample complexity is sufficient for Algorithm 3 to correctly identify the subsets T_j' ($j\in[N]$) with probability at least $1-n^{-c_6}$:

$$n^{1/2} > \frac{2C_6}{\omega(1-\delta)}\left(1+\frac{C_W}{\sqrt{1-\theta}}\right)^2 s_0\log^3(nmN).$$

Thus the first half of Theorem 14 is proved.

Given that the support T_j covers the large entries indexed by T_j',

$$\begin{aligned}
\|\mathbb{E}\Pi_{T_x}D^*E - \frac{1}{n}x\lambda^\top\| &= \|\frac{1}{n}\Pi_{T_x}x\lambda^\top - \frac{1}{n}x\lambda^\top\| \\
&\le \sqrt{\frac{1+\delta}{n}\sum_{j\in[N],\ell'\in[m]\setminus T_j'}|x_{\ell' j}|^2} \\
&\le \sqrt{\frac{(1+\delta)\delta_X}{n}}.
\end{aligned} \tag{32}$$

We also have

$$\begin{aligned}
&\|\Pi_{T_x}D^*E - \mathbb{E}\Pi_{T_x}D^*E\| \\
&\quad\le \|\Omega_{T_x}D^*E_{\mathrm{s}} - \mathbb{E}\Omega_{T_x}D^*E_{\mathrm{s}}\| + \|\Omega_{T_x}D^*E_{\mathrm{n}}\| \\
&\quad\le \frac{1}{\alpha}(\|\Omega_{T_\eta}B_{\mathrm{s}}\Omega_{T_\eta}^* - \Omega_{T_\eta}\mathbb{E}B_{\mathrm{s}}\Omega_{T_\eta}^*\| + \|\Omega_{T_\eta}B_{\mathrm{n}}\Omega_{T_\eta}^*\|) \\
&\quad\le \frac{1}{\sqrt{n}}(\delta_B+\delta_W),
\end{aligned} \tag{33}$$

where the last inequality follows from Lemmas 16 and 21, given that the conditions of Theorem 12 are satisfied. By the triangle inequality, and (32) and (33),

$$\|\Pi_{T_x}D^*E - \frac{1}{n}x\lambda^\top\| \le \frac{1}{\sqrt{n}}(\delta_B+\delta_W+\sqrt{(1+\delta)\delta_X}),$$

where δ_B can be made arbitrarily small by a sufficiently large C in (13), δ_W can be made arbitrarily small by a sufficiently small C_W in Assumption 6, and the last term can be made arbitrarily small by a sufficiently small δ_X in Assumption 7. Therefore, the first left and right singular vectors u and v can become arbitrarily close to x and to $\lambda/\|\lambda\|_2$ (up to a global phase factor, i.e., a constant of unit modulus), respectively, and $|\dot\eta^*\eta^{(0)}|$ approaches

$$\frac{n^{3/2}+\|\lambda\|_2\|\gamma\|_2^2}{\sqrt{n^2+\|\lambda\|_2^2\|\gamma\|_2^2}\sqrt{n+\|\gamma\|_2^2}} > 1-2\delta.$$

The inequality follows from Assumption 4, i.e., $\sqrt{1-\delta}\le|\lambda_k|\le\sqrt{1+\delta}$, and $1/\sqrt{1+\delta}\le|\gamma_k| = 1/|\lambda_k|\le 1/\sqrt{1-\delta}$.■

VI. Numerical Experiments

In this section, we test the empirical performance of Algorithm 1 and Algorithm 2.

A. Subspace Case: Power Iteration vs. Least Squares

In Algorithm 1, we choose $\alpha = \sqrt{n}$, and $\beta = \|B\|$ (computed using another power iteration on B). We compare Algorithm 1 with the least squares approach in [7, Sec. 3.3], where $\gamma_1 = 1$ is used to avoid the trivial solution.

We generate $A \in \mathbb{C}^{n\times m}$ as a complex Gaussian random matrix, whose entries are drawn independently from $\mathcal{CN}(0,\frac{1}{n})$, i.e., the real and imaginary part are drawn independently from $\mathcal{N}(0,\frac{1}{2n})$. The unknown gains and phases λ_k are generated as follows:

$$\lambda_k = e^{\sqrt{-1}\varphi_k}\left(1+(\sqrt{1+\delta}-1)e^{\sqrt{-1}\varphi_k'}\right),\quad \forall k\in[n], \tag{34}$$

such that λ_k is on a small circle of radius $\sqrt{1+\delta}-1$ centered at a point on the unit circle, and φ_k and φ_k' are drawn independently from a uniform distribution on $[0,2\pi)$. Figure 1 visualizes one such synthesized λ_k in the complex plane. We set $\delta = 0.1$ in all the numerical experiments. The entries of $X\in\mathbb{C}^{m\times N}$ are drawn independently from $\mathcal{CN}(0,\frac{1}{Nm})$, so that the Frobenius norm of X is approximately 1. In the noisy setting, we generate complex white Gaussian noise $W\in\mathbb{C}^{n\times N}$, whose entries are drawn from $\mathcal{CN}(0,\frac{\sigma_W^2}{Nn})$.

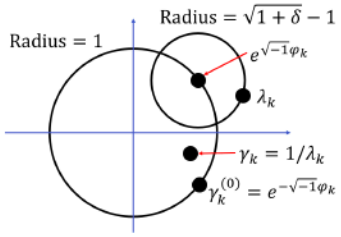

Fig. 1. Illustration of λ_k in the complex plane.

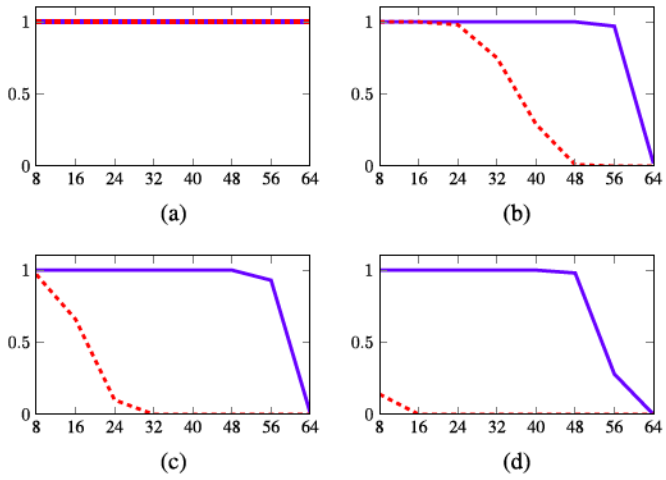

Fig. 2. Subspace case: The empirical success rates of power iteration (blue solid line) and least squares (red dashed line). The x-axis represents m, and the y-axis represents the empirical success rate. (a) – (d) are the results with $\sigma_W = 0$, 0.1, 0.2, and 0.5, respectively.

We define measurement signal-to-noise ratio (MSNR) and recovery signal-to-noise ratio (RSNR) as:

$$\text{MSNR} := 20\log_{10}\frac{\|\text{diag}(\lambda)AX\|_{\text{F}}}{\|W\|_{\text{F}}},$$
$$\text{RSNR} := -10\log_{10}(2-2|\dot{\eta}^*\eta^{(t)}|).$$

We test the two approaches at four noise levels: $\sigma_W = 0$, 0.1, 0.2, and 0.5, which roughly correspond to MSNR of ∞, 20 dB, 14 dB, and 6 dB. At these noise levels, we say the recovery is successful if the RSNR exceeds 30 dB, 20 dB, 14 dB, 6 dB, respectively. The success rates do not change dramatically as functions of these thresholds. In the experiments, we set $n = 128$, $N = 16$, and $m = 8, 16, 24, \ldots, 64$. For each m, we repeat the experiments 100 times and compute the empirical success rates, which are shown in Figure 2.

As seen in Figure 2(a), both power iteration and least squares achieve perfect recovery in the noiseless setting. However, as seen in Figures 2(b) – 2(d), power iteration is clearly more robust against noise than least squares, whose performance degrades more severely in the noisy settings.

The empirical phase transitions of power iteration are shown in Figure 3. We fix $N = 16$ and plot the phase transition with respect to n and m (Figure 3(a)); we then fix $n = 2m$ and plot the phase transition with respect to N and m (Figure 3(b)). Clearly, to achieve successful recovery, n must scale linearly with m, but N can be small compared to m and n. This confirms the sample complexity in Theorem 11, of $n \gtrsim m$ and $N \gtrsim 1$. Careful readers may notice in Figure 3(b) that for $N = 5$ the success rates at $m < 16$ are worse than those at $m \geq 16$. This seemingly peculiar phenomenon is caused by a small $n = 2m$, which does not belong to the large number regime associated with a high probability.

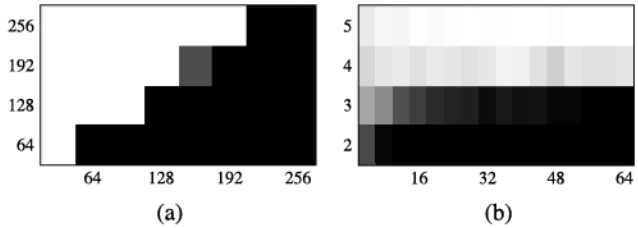

Fig. 3. The empirical phase transition of power iteration. Grayscale represents success rates, where white equals 1, and black equals 0. (a) The x-axis represents m, and the y-axis represents n. (b) The x-axis represents m, and the y-axis represents N.

B. Sparsity Case: Truncated Power Iteration vs. ℓ_1 Minimization

In the sparsity case, we use the same setup described in the previous section, except for the signal X. The supports of the s_0-sparse columns of X are chosen uniformly at random, and the nonzero entries follow $\mathcal{CN}(0, \frac{1}{Ns_0})$. This unstructured sparsity case is more challenging than the joint sparsity case in Theorem 12.

In Algorithm 2, we choose $\alpha = \sqrt{n}$, and $\beta = \|B\|$. In all the experiments, we assume that the sparsity level s_0 is known, and set $s_1 = 2s_0$ for convenience. A more sophisticated scheme that decreases s_1 as the iteration number increases may lead to better empirical performance [11].

For the experiment we suppose that the phases $\{\varphi_k\}_{k=1}^n$ in (34) are available, and let

$$\gamma^{(0)} := [e^{-\sqrt{-1}\varphi_1}, \ldots, e^{-\sqrt{-1}\varphi_n}]^\top \tag{35}$$

denote the initial estimate of γ, which is close to but different from the true γ, i.e., the entrywise inverse of λ in (34). See Figure 1 for an illustration of λ_k, γ_k, and $\gamma_k^{(0)}$. Then we initialize Algorithm 2 with $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, \gamma^{(0)\top}]^\top$.

We compare Algorithm 2 with an ℓ_1 minimization approach. Wang and Chi [8] adopted an approach tailored for the case where A is the DFT matrix and $\lambda_k \approx 1$. They use a linear constraint $\sum_{k\in[n]}\gamma_k = n$ to avoid the trivial solution of all zeros. For fair comparison, we revise their approach to accommodate arbitrary A and λ. The revised approach uses the alternating direction method of multipliers (ADMM) [61] to solve the following convex optimization problem[4]:

$$\begin{aligned}&\min_{\gamma,X}\ \|\text{vec}(X)\|_1\\ &\text{s.t. } \text{diag}(\gamma)Y = AX,\\ &\quad \gamma^{(0)*}\gamma = n.\end{aligned}$$

[4]In the noisy setting, one could replace the linear constraint $\text{diag}(\gamma)Y = AX$ with an ellipsoid constraint $\|\text{diag}(\gamma)Y - AX\|_{\text{F}} \leq \epsilon$. However, the parameter ϵ needs to be adjusted with noise levels. For fair comparison of robustness to noise, we use the linear constrained ℓ_1 minimization in the noisy setting (similar to [8]).

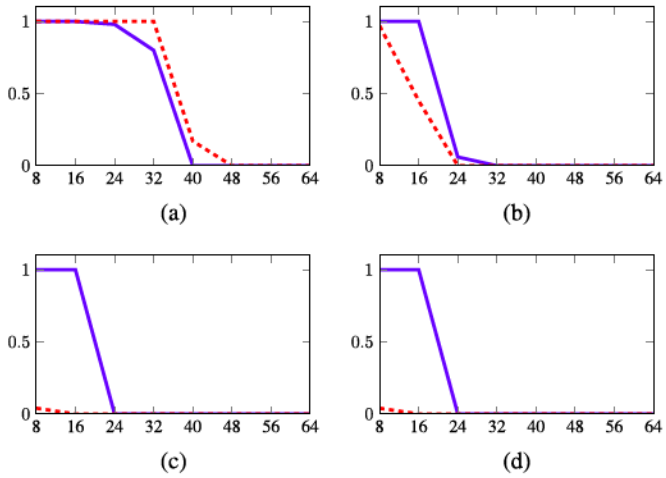

Fig. 4. Sparsity case: The empirical success rates of truncated power iteration (blue solid line) and ℓ_1 minimization (red dashed line). The x-axis represents s_0, and the y-axis represents the empirical success rate. (a) – (d) are the results with $\sigma_W = 0$, 0.1, 0.2, and 0.5, respectively.

Here, $\gamma^{(0)}$ is the initial estimate of γ defined in (35), and used as initialization in our Algorithm 2 in this comparison.

We conduct numerical experiments with the same four noise levels and criterion for successful recovery as in Section VI-A. In the experiments, we set $n = 128$, $m = 256$, $N = 16$, and $s_0 = 8, 16, 24, \ldots, 64$. For each s_0, we repeat the experiments 100 times and compute the empirical success rates, which are shown in Figure 4. In the noiseless case (Figure 4(a)), ℓ_1 minimization achieves a slightly higher success rate near the phase transition. However, truncated power iteration is more robust against noise than ℓ_1 minimization, which breaks down completely at the higher noise levels (Figures. 4(b) – 4(d)).

Figure 4(a) clearly shows that truncated power iteration recovers η successfully when $n = 128$, $N = 16$, and $s_0 = 32$. This suggests that truncated power iteration may succeed when n and N are (up to log factors) on the order of s_0 and 1, respectively. However, while the scaling with the number of sensors n agrees with Theorem 12, success with such small number of snapshots N is not guaranteed by our current theoretical analysis.

Next, we assume that only a subset of the phases $\{\varphi_k\}_{k=1}^n$ are available, and examine to what extent Algorithm 2 and ℓ_1 minimization depend on a good initial estimate of γ. In the numerical results shown in Figure 5, we consider only the noiseless setting of BGPC with sparsity, and set $s_0 = 4, 8, 12, \ldots, 32$. In Figures 5(a) and 5(b), we replace 1/2 and 3/4 of $\{\varphi_k\}_{k=1}^n$ with random phases, respectively, and use the resulting bad estimate $\gamma^{(0)}$ in Algorithm 2 and ℓ_1 minimization. As seen in Figure 5, truncated power iteration is less dependent on accurate initial estimate of γ.

We repeat the above experiments for the joint sparsity case, where we replace $\widetilde{\Pi}_{s_1}$ in Algorithm 2 with $\widetilde{\Pi}'_{s_1}$. We also replace the ℓ_1 norm $\|\mathrm{vec}(X)\|_1$ in the competing approach with a mixed norm:

$$\|X\|_{2,1} = \sum_{\ell\in[m]} \Big(\sum_{j\in[N]} |x_{\ell j}|^2 \Big)^{1/2},$$

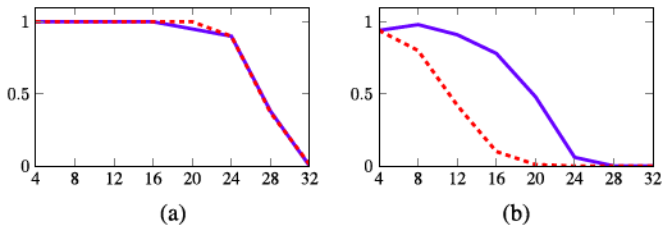

Fig. 5. Sparsity case: The empirical success rates of truncated power iteration (blue solid line) and ℓ_1 minimization (red dashed line), with bad initial estimate of the phases. The x-axis represents s_0, and the y-axis represents the empirical success rate. (a) and (b) are the results for which 1/2 and 3/4 of $\{\varphi_k\}_{k=1}^n$ are initialized with random phases.

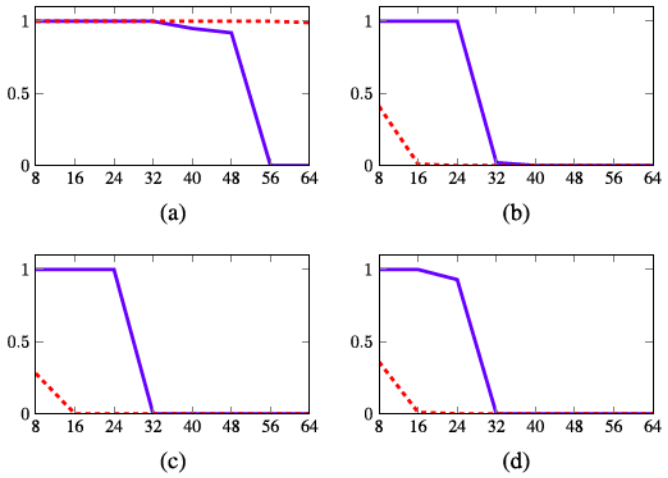

Fig. 6. Joint sparsity case: The empirical success rates of truncated power iteration (blue solid line) and mixed minimization (red dashed line). The x-axis represents s_0, and the y-axis represents the empirical success rate. (a) – (d) are the results with $\sigma_W = 0$, 0.1, 0.2, and 0.5, respectively.

which is a well-known convex method for the recovery of jointly sparse signals. The results for different noise levels and for inaccurate $\gamma^{(0)}$ are shown in Figures 6 and 7, respectively. In the joint sparsity case, truncated power iteration is robust against noise, but seems less robust against errors in the initial phase estimate. We conjecture that the failure of Algorithm 2 in the joint sparsity case is due to the restriction of $\widetilde{\Pi}'_{s_1}$. By projecting onto jointly sparse supports, the algorithm is likely to converge prematurely to an incorrect support. When compared to the results in Figures 7(a) and 7(b), Figures 7(c) and 7(d) show that using $\widetilde{\Pi}_{s_1}$ instead of $\widetilde{\Pi}'_{s_1}$ in the first half of the iterations indeed improves the performance of Algorithm 2 in the joint sparsity case. In the rest of the experiments, we use $\widetilde{\Pi}_{s_1}$ during the first half of the iterations in Algorithm 2 for the joint sparsity case.

Next, we plot the phase transitions for truncated power iteration. We fix $N = 16$ and $m = 2n$ and plot the empirical phase transition with respect to n and s_0 (sparsity case in Figure 8(a), and joint sparsity case in Figure 8(c)); we then fix $n = 4s_0$ and $m = 2n$ and plot the empirical phase transition with respect to N and s_0 (sparsity case in Figure 8(b), and joint sparsity case in Figure 8(d)). It is seen that, to achieve successful recovery, n must scale linearly with s_0, but N can be small compared to s_0 and n. On the one hand, the scaling

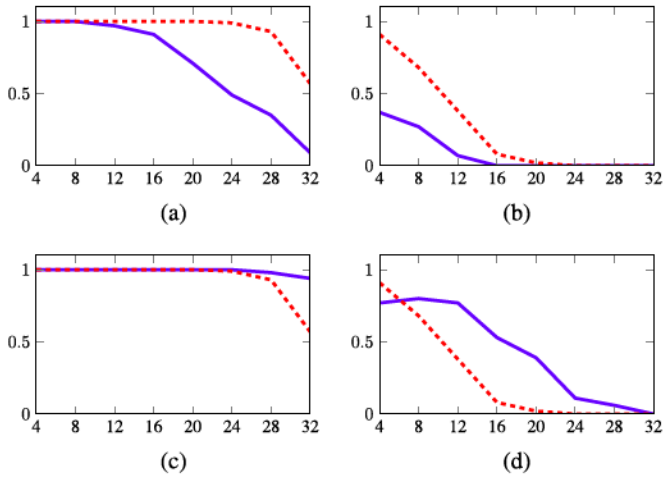

Fig. 7. Joint sparsity case: The empirical success rates of truncated power iteration with $\widetilde{\Pi}'_{s_1}$ (blue solid line) and mixed minimization (red dashed line), with bad initial estimate of the phases. The x-axis represents s_0, and the y-axis represents the empirical success rate. (a) and (b) are the results for which 1/2 and 3/4 of $\{\varphi_k\}_{k=1}^n$ are initialized with random phases. In (c) and (d), we repeat the experiments, but use $\widetilde{\Pi}_{s_1}$ instead of $\widetilde{\Pi}'_{s_1}$ in the first half of the iterations.

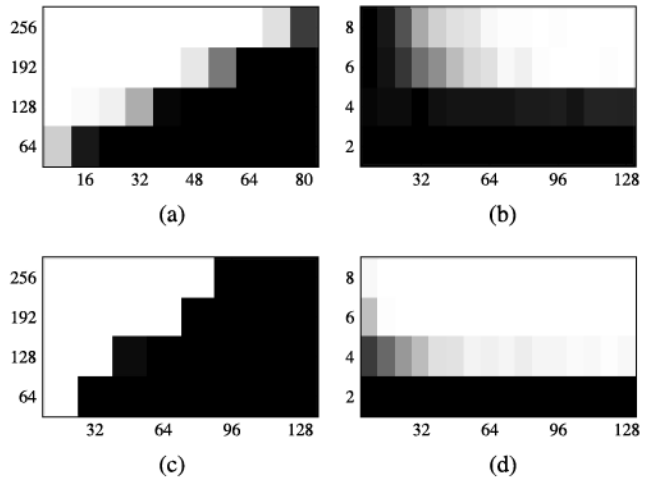

Fig. 8. The empirical phase transition of truncated power iteration. Grayscale represents success rates, where white equals 1, and black equals 0. (a) Sparsity case: The x-axis represents s_0, and the y-axis represents n. (b) Sparsity case: The x-axis represents s_0, and the y-axis represents N. (c) Joint sparsity case: The x-axis represents s_0, and the y-axis represents n. (d) Joint sparsity case: The x-axis represents s_0, and the y-axis represents N.

law $n \gtrsim s_0$ in Theorem 12 is confirmed by Figure 8; on the other hand, $N \gtrsim \sqrt{s_0}$ seems conservative and might be an artifact of our proof techniques. We have yet to come up with a theoretical guarantee that covers the more general sparsity case, or requires a less demanding sample complexity $N \gtrsim 1$. In Figures 8(b) and 8(d), the success rates at smaller s_0 are lower than those at a larger s_0, because the number of sensors $n = 4s_0$ is too small to yield a high probability.

C. Sparsity Case: Initialization

In this section, we examine the quality of the initialization produced by Algorithm 3 by comparing it with two different initializations: (i) the good initialization $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, \gamma^{(0)\top}]^\top$ aided by side information on the phase in Section VI-B; and (ii) a baseline initialization $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, \mathbf{1}_{n,1}^\top]^\top$. We use the same setting as in Section VI-B, except that $N = 32$. We let $\sigma_W = 0.1$, and claim the recovery is successful if the RSNR exceeds 20 dB. In the experiment for the joint sparsity case, for the reason mentioned in Section VI-B, we ignore the joint sparsity structure and estimate the support of different columns of X independently in the initialization and during the first half of the iterations. Only in the second half of the iterations, we use the projection $\widetilde{\Pi}'_{s_1}$ onto jointly sparse supports.

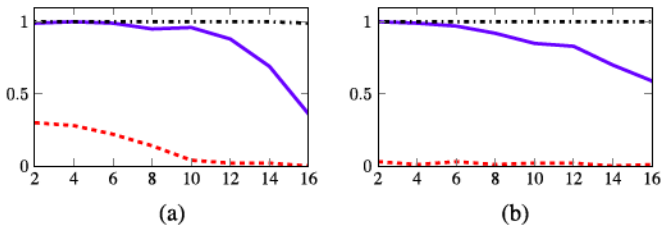

Fig. 9. The empirical success rates of truncated power iteration with the initialization in Algorithm 3 (blue solid line), with a baseline initialization $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, \mathbf{1}_{n,1}^\top]^\top$ (red dashed line), and with the accurate initialization $\eta^{(0)} = [\mathbf{0}_{Nm,1}^\top, \gamma^{(0)\top}]^\top$ with side information in Section VI-B (black dash-dot line). The x-axis represents s_0, and the y-axis represents the empirical success rate. (a) is the result for the sparsity case, and (b) is the result for the joint sparsity case.

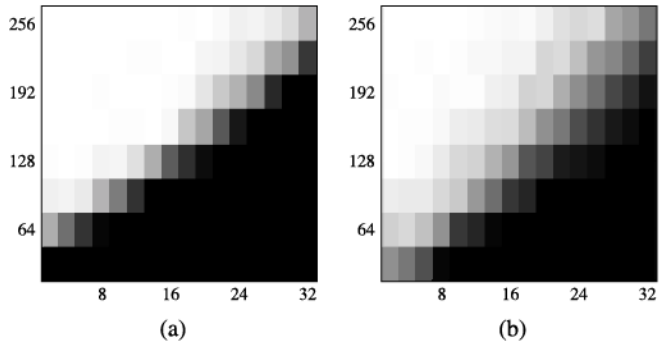

Fig. 10. The empirical phase transition of truncated power iteration with the initialization in Algorithm 3. The x-axis represents s_0, and the y-axis represents n. (a) is the result for the sparsity case, and (b) is the result for the joint sparsity case.

Figure 9 shows that, although the initialization provided by Algorithm 3 is not as good as the accurate initialization with side information, it is far better than the baseline. Figure 10 shows the empirical phase transition with respect to n and s_0, when Algorithm 3 is used to initialize truncated power iteration (sparsity case in Figure 10(a), and joint sparsity case in Figure 10(b)). The results suggest that when n scales linearly with s_0, Algorithm 3 can provide a sufficiently good initialization for truncated power iteration. For example, in 10(a), the success rate is 1 when $n = 256$ and $s_0 = 20$. Therefore, the sample complexity $n \gtrsim s_0^2$ in Theorem 14 could be overly conservative and an artifact of our analysis.

D. Dependence on Assumptions

In this section, we study how much the performance of Algorithms 1 – 3 depend on the Assumptions 3 – 7.

We first examine the importance of Assumption 3 by comparing the recovery success rate under four different models of $A \in \mathbb{C}^{n\times m}$:

- **Gaussian**: the entries are drawn independently from $\mathcal{CN}(0, \frac{1}{n})$.
- **Rademacher**: the entries are drawn independently from a two-point distribution on $\pm\frac{1}{\sqrt{n}}$, each with probability $\frac{1}{2}$ (a scaled version of the Rademacher distribution).
- **Random rotation**: When $n \geq m$, the columns of A are random orthonormal vectors (formed by computing the left singular vectors of a complex random Gaussian matrix). When $n < m$, the rows of A are random orthogonal vectors of ℓ_2 norm $\sqrt{\frac{m}{n}}$.
- **Partial Fourier**: A is a randomly subsampled DFT matrix. When $n \geq m$, the columns of A are a random subset of columns of the normalized $n \times n$ DFT matrix. When $n < m$, the rows of A are a random subset of row of the normalized $m \times m$ DFT matrix, scaled by $\sqrt{\frac{m}{n}}$.

In the subspace case, we compare the success rates of Algorithm 1 for the above four models, with $n = 128$, $m = 16$, and $N = 4$. In the joint sprasity case, we test Algorithm 2 with $n = 128$, $m = 256$, $s_0 = 16$, and $N = 4$. We generate λ and X following the same model as in Section VI-A, and for the joint sparsity case, assume that we have access to the same good initial estimate $\gamma^{(0)}$ as in Section VI-B. The measurement Y contains noise with $\sigma = 0.1$ (the MSNR $\approx$ 20 dB), and we declare successful recovery if the RSNR exceeds 20 dB. The success rates in the subspace case and the joint sparsity case are shown in Figures 11(a) and 11(b), respectively.

The empirical success rates in both the subspace and the joint sparsity cases show that, although we need the complex Gaussian random matrix model in Assumption 3 for the proof of our main results, in practice Algorithms 1 and 2 are just as successful for other models of A as they are for the Gaussian model. In other words, one does not need A to be a complex Gaussian random matrix for our algorithms to converge, and to be effective.

However, our initialization algorithm is not equally successful for all models of A. To demonstrate this, we test our algorithms for the joint sparsity case with $n = 128$, $m = 256$, $s_0 = 4$, and $N = 32$, and feed the initial estimates produced by Algorithm 3 to Algorithm 2. The success rates in Figure 11(c) clearly show that Algorithm 3 cannot produce sufficiently accurate estimates under the "Rademacher" model and the "partial Fourier" model of A. The condition on A under which Algorithm 3 is empirically successful requires further investigation, and is beyond the scope of this paper.

Next, we show that Assumptions 4 and 5 are important in terms of convergence rate and noise robustness of our algorithms. Despite the fact that, in the subspace case, the principal eigenvector of G corresponds to the ground truth regardless of the dynamic range in λ or the conditioning of X as long as the solution is unique in the subspace case, one can only expect the fast convergence and robust recovery in Theorem 11 under certain regularity conditions on λ and X. In the joint sparsity case, the success of truncated power iteration depends even more on the flatness of λ and the good conditioning of X. We demonstrate the importance of Assumptions 4 and 5 by relaxing them gradually, and observe the degradation in success rates.

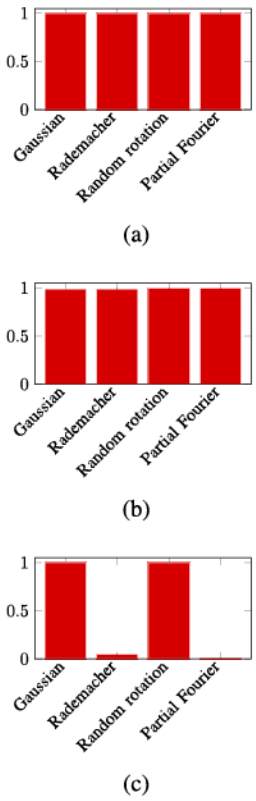

Fig. 11. The empirical success rates of our algorithms under four models of A. (a) is the result for the subspace case, (b) is the result for the joint sparsity case with a good approximate initial estimate, and (c) is the result for the joint sparsity case with initialization produced using Algorithm 3.

In particular, to compare with the success rates with an approximately flat λ that satisfies Assumption 4 (following the same model in Section VI-A), we multiply different entries of λ by additional gains – i.i.d. random variables following a uniform distribution on [0.5, 1] – which roughly increases the dynamic range of λ by a factor of 2. We also apply additional gains drawn from uniform distributions on [0.2, 1], [0.1, 1], or [0, 1], which makes λ progressively less flat by roughly increasing its dynamic range by a factor of 5, 10, or ∞. Similarly, we increase the condition number of X by a factor of about 2, 5, 10, or ∞, by multiplying the columns of X by i.i.d. random variables drawn from uniform distributions on [0.5, 1], [0.2, 1], [0.1, 1], or [0, 1]. According to Figure 12, the success rates, in both the subspace case and the joint sparsity case, decrease as the flatness of λ or the conditioning of X becomes worse. Such negative impact is more pronounced in the joint sparsity case.

The extensive numerical experiments in Sections VI-A and VI-B demonstrate that Algorithms 1 and 2 are successful under reasonable noise levels (MSNR = 20 dB, 14 dB, and 6 dB). The requirement on the noise level in the joint sparsity case of Assumption 6 is pessimistic when compared to empirical results, due to limitations of our theoretical analysis.

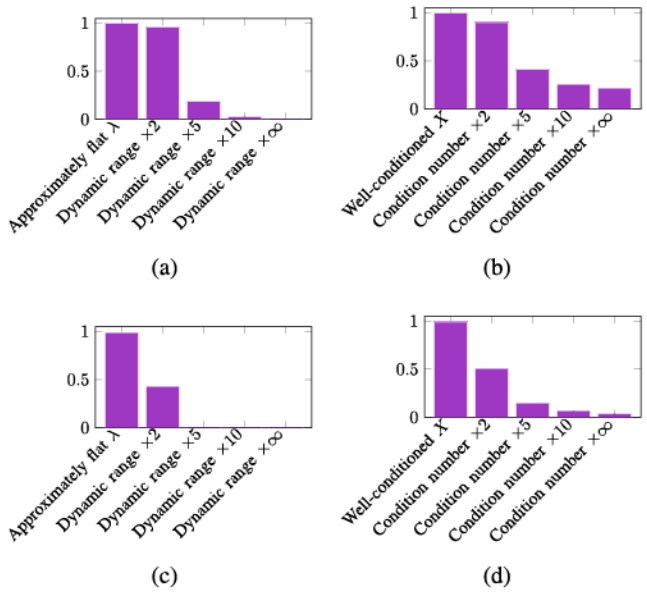

Fig. 12. The empirical success rates of our algorithms under different flatness of λ, and different conditioning of X. (a) and (b) are the results for the subspace case, (c) and (d) are the results for the joint sparsity case.

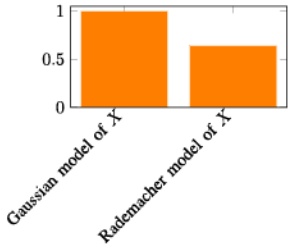

Fig. 13. The empirical success rates under two models of row sparse X: the Gaussian model and the Rademacher model.

In order to investigate how critical Assumption 7 is to the empirical performance of our initialization algorithm, we compare two models for the nonzero entries of X (in the joint sparsity case, with $n = 128$, $m = 256$, $s_0 = 4$, and $N = 32$):

- **Gaussian**: the nonzero entries follow $\mathcal{CN}(0, \frac{1}{Ns_0})$.
- **Rademacher**: the nonzero entries follow a two-point distribution on $\pm\frac{1}{\sqrt{Ns_0}}$, each with probability $\frac{1}{2}$ (a scaled version of the Rademacher distribution).

By our discussion in Section III-A, the Gaussian model satisfies Assumption 7 with $\omega = \frac{1}{4}$ and $\delta_X = \frac{1}{\sqrt{2\pi}}$, while the Rademacher model is the ideal case with $\omega = 1$ and $\delta_X = 0$. However, Figure 13 shows that Algorithms 2 and 3 have higher success rate under the Gaussian model. This suggests that, although Assumption 7 is important for our theoretical analysis of Algorithm 3, it cannot be used to predict the empirical performance of our algorithms, as better constants in Assumption 7 do not necessarily mean higher success rates.

E. Application: Inverse Rendering

In this section, we apply the power iteration algorithm to the inverse rendering problem in computational relighting – given images of an object under different but unknown lighting conditions (Figure 14(a)), and the surface normals of the object (Figure 14(b)), the goal is to recover both the albedo (also known as reflection coefficients) of the object surface and the lighting conditions. In this problem, the columns of $Y = \text{diag}(\lambda)AX \in \mathbb{R}^{n\times N}$ represent images under different lighting conditions, which are the products of the unknown albedo map $\lambda \in \mathbb{R}^n$ and the intensity maps of incident light under different conditions AX. For Lambertian surfaces, it is reasonable to assume that the intensity of incident light resides in a subspace spanned by the first nine spherical harmonics computed from the surface normals [2], which we denote by the columns of $A \in \mathbb{R}^{n\times 9}$. Then the columns of X are the coordinates of the spherical harmonic expansion, which parameterize the lighting conditions. We can solve for λ and X using Algorithm 1. Our approach is similar to that of Nguyen *et al.* [2], which also formulates inverse rendering as an eigenvector problem. Despite the fact that the two approaches solve for the eigenvectors of different matrices, they yield identical solutions in the ideal scenario where the model is exact and the solution is unique.

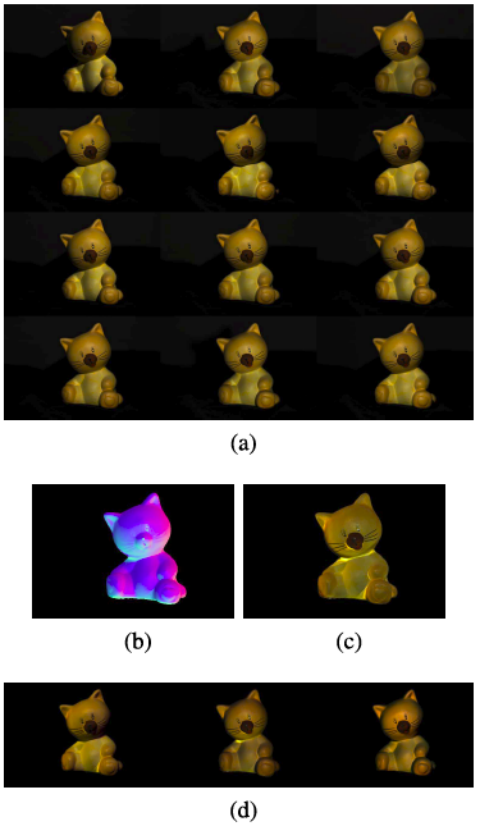

Fig. 14. Inverse rendering and relighting. (a) We use 12 images of the object under different lighting conditions. (b) The surface normals. The three dimensions of the normal vectors are represented by the RGB channels of the color image. (c) The recovered albedo map. (d) Computed images of the object under new lighting conditions.

In our experiment, we obtain $N = 12$ color images and the surface normals of an object under different lighting conditions,[5] and we compute the first $m = 9$ spherical harmonics.

[5]The images are downloaded from *https://courses.cs.washington.edu/courses/csep576/05wi/projects/project3/project3.htm* on September 16, 2017. The surface normals are computed using the method described in the same webpage.

We apply Algorithm 1 to each of the three color channels, and the albedo map recovered using 200 power iterations is shown in Figure 14(c). We also compute new images of the object under new lighting conditions (Figure 14(d)).

VII. Conclusion

We formulate the BGPC problem as an eigenvector problem, and propose to solve BGPC with power iteration, and solve BGPC with a sparsity structure with truncated power iteration. We give theoretical guarantees for the subspace case with a near optimal sample complexity, and for the joint sparsity case with a suboptimal sample complexity. Numerical experiments show that both power iteration and truncated power iteration can recover the unknown gain and phase, and the unknown signal, using a near optimal number of samples. It is an open problem to obtain theoretical guarantees with optimal sample complexities, for truncated power iteration that solves BGPC with joint sparsity or sparsity constraints.

Appendix

Proof of Lemma 15: We have

$$D^*D = I_N \otimes (A^*A), \tag{36}$$

$$D^*E_{\mathrm{s}} = \begin{bmatrix} \lambda_1\overline{a_{1\cdot}}a_{1\cdot}^\top x_{\cdot 1} & \cdots & \lambda_n\overline{a_{n\cdot}}a_{n\cdot}^\top x_{\cdot 1} \\ \vdots & \ddots & \vdots \\ \lambda_1\overline{a_{1\cdot}}a_{1\cdot}^\top x_{\cdot N} & \cdots & \lambda_n\overline{a_{n\cdot}}a_{n\cdot}^\top x_{\cdot N} \end{bmatrix}, \tag{37}$$

$$E_{\mathrm{s}}^*E_{\mathrm{s}} = \begin{bmatrix} |\lambda_1|^2\ a_{1\cdot}^\top XX^*\overline{a_{1\cdot}} & & \\ & \ddots & \\ & & |\lambda_n|^2\ a_{n\cdot}^\top XX^*\overline{a_{n\cdot}} \end{bmatrix}. \tag{38}$$

Under Assumptions 3 and 5, we have

$$\mathbb{E}D^*D = I_{Nm}, \tag{39}$$

$$\mathbb{E}D^*E_{\mathrm{s}} = \frac{1}{n}x\lambda^\top, \tag{40}$$

$$\begin{aligned}\mathbb{E}E_{\mathrm{s}}^*E_{\mathrm{s}} &= \frac{1}{n}\|X\|_{\mathrm{F}}^2\mathrm{diag}([|\lambda_1|^2,\ldots,|\lambda_n|^2]) \\ &= \frac{1}{n}\mathrm{diag}([|\lambda_1|^2,\ldots,|\lambda_n|^2]).\end{aligned} \tag{41}$$

Set $\alpha = \sqrt{n}$, we have

$$\mathbb{E}B_{\mathrm{s}} = \begin{bmatrix} I_{Nm} & \frac{1}{\sqrt{n}}x\lambda^\top \\ \frac{1}{\sqrt{n}}\overline{\lambda}x^* & \mathrm{diag}([|\lambda_1|^2,\ldots,|\lambda_n|^2]) \end{bmatrix},$$

and

$$\begin{aligned}&\mathbb{E}\Omega_{T_\eta}B_{\mathrm{s}}\Omega_{T_\eta}^* \\ &= \begin{bmatrix} I_{Ns} & \frac{1}{\sqrt{n}}\Omega_{T_x}x\lambda^\top \\ \frac{1}{\sqrt{n}}\overline{\lambda}x^*\Omega_{T_x}^* & \mathrm{diag}([|\lambda_1|^2,\ldots,|\lambda_n|^2]) \end{bmatrix} \\ &= P^*QP,\end{aligned}$$

where

$$\begin{aligned}P &= \mathrm{diag}([\mathbf{1}_{1,Ns}, \lambda^\top]), \\ Q &= \begin{bmatrix} I_{Ns} & \frac{1}{\sqrt{n}}\Omega_{T_x}x\mathbf{1}_{n,1}^\top \\ \frac{1}{\sqrt{n}}\mathbf{1}_{n,1}x^*\Omega_{T_x}^* & I_n \end{bmatrix}.\end{aligned}$$

The matrix Q has eigenvalues $0, 1, 1, \ldots, 1, 2$. The eigenvectors corresponding to 0 and 2 are $\mu = [(\Omega_{T_x}x)^\top, -\mathbf{1}_{n,1}^\top/\sqrt{n}]^\top/\sqrt{2}$ and $[(\Omega_{T_x}x)^\top, \mathbf{1}_{n,1}^\top/\sqrt{n}]^\top/\sqrt{2}$, respectively. Any vector orthogonal to these two vectors is an eigenvector of Q corresponding to 1. It follows that $Q + \mu\mu^* - I_{Ns+n}$ is positive semidefinite.

Since μ is a null vector of Q, we have $P^{-1}\mu$ is a null vector of P^*QP (note that $\Omega_{T_\eta}\eta = \sqrt{2}P^{-1}\mu$). Therefore, the smallest eigenvalue of the positive semidefinite matrix P^*QP is 0.

Next, we bound the largest eigenvalue of P^*QP, which satisfies

$$\begin{aligned}\max_{\|z\|_2\le 1}\|P^*QPz\|_2 &\le \sqrt{1+\delta}\max_{\|Pz\|_2\le\sqrt{1+\delta}}\|QPz\|_2 \\ &= (1+\delta)\max_{\|z\|_2\le 1}\|Qz\|_2 \\ &\le 2(1+\delta),\end{aligned} \tag{42}$$

where the first inequality follows from Assumption 4, and the second inequality follows from the largest eigenvalue of Q.

Next, we bound the second smallest eigenvalue of P^*QP, which satisfies

$$\begin{aligned}&\min_{z\perp P^{-1}\mu,\ \|z\|_2\ge 1}\|P^*QPz\|_2 \\ &\ge \sqrt{1-\delta}\min_{Pz\perp(PP^*)^{-1}\mu,\ \|Pz\|_2\ge\sqrt{1-\delta}}\|QPz\|_2 \\ &= (1-\delta)\min_{z\perp(PP^*)^{-1}\mu,\ \|z\|_2\ge 1}\|Qz\|_2 \\ &\ge (1-\delta)\min_{z\perp(PP^*)^{-1}\mu,\ \|z\|_2=1}\|(I_{Ns+n}-\mu\mu^*)z\|_2 \\ &= (1-\delta)\min_{z\perp(PP^*)^{-1}\mu,\ \|z\|_2=1}\sqrt{1-|\mu^*z|^2} \\ &= (1-\delta)\frac{|\mu^*(PP^*)^{-1}\mu|}{\|(PP^*)^{-1}\mu\|_2} \\ &\ge \frac{(1-\delta)^2}{1+\delta},\end{aligned} \tag{43}$$

where the first and third inequalities follow from Assumption 4, and the second inequality is due to the fact that $Q + \mu\mu^* - I_{Ns+n}$ is positive semidefinite.

By (42) and (43), all nonzero eigenvalues of $\mathbb{E}\Omega_{T_\eta}B_{\mathrm{s}}\Omega_{T_\eta}^*$ reside in the interval $[\frac{(1-\delta)^2}{1+\delta}, 2(1+\delta)]$. ∎

Proof of Lemma 17: We prove only the joint sparsity case. One can prove the subspace case by replacing s with m and getting rid of the union bound.

It is well-known that, for sufficiently large n, a Gaussian random matrix satisfies RIP [55]. Here, we use a bound for real Gaussian random matrices [62], and present its extension to complex Gaussian random matrices. Let $T \subset [m]$ denote an index set of cardinality s, i.e., $|T| = s < n$. Let $\widehat{A} \coloneqq [\mathrm{Re}(A)\Omega_T^*, \mathrm{Im}(A)\Omega_T^*]$. By [62, Th. 2.13],

$$\mathbb{P}\Big[\|2\widehat{A}^*\widehat{A} - I_{2s}\| \le 3\Big(\sqrt{\frac{2s}{n}} + \varepsilon\Big)\Big] \ge 1 - 2\exp\Big(-\frac{n\varepsilon^2}{2}\Big).$$

Note also that

$$\begin{aligned}\Omega_T A^*A\Omega_T^* &= \Omega_T \mathrm{Re}(A)^\top \mathrm{Re}(A)\Omega_T^* \\ &\quad + \sqrt{-1}\Omega_T \mathrm{Re}(A)^\top \mathrm{Im}(A)\Omega_T^* \\ &\quad - \sqrt{-1}\Omega_T \mathrm{Im}(A)^\top \mathrm{Re}(A)\Omega_T^* \\ &\quad + \Omega_T \mathrm{Im}(A)^\top \mathrm{Im}(A)\Omega_T^*. \\ \|\Omega_T A^*A\Omega_T^* - I_s\| &\le \|\Omega_T \mathrm{Re}(A)^\top \mathrm{Re}(A)\Omega_T^* - I_s/2\| \\ &\quad + \|\Omega_T \mathrm{Re}(A)^\top \mathrm{Im}(A)\Omega_T^*\| \\ &\quad + \|\Omega_T \mathrm{Im}(A)^\top \mathrm{Re}(A)\Omega_T^*\| \\ &\quad + \|\Omega_T \mathrm{Im}(A)^\top \mathrm{Im}(A)\Omega_T^* - I_s/2\| \\ &\le 4\|\widehat{A}^*\widehat{A} - I_{2s}/2\|.\end{aligned}$$

It follows that

$$\begin{aligned}\mathbb{P}\Big[\|\Omega_T A^*A\Omega_T^* - I_s\| \le 6\Big(\sqrt{\frac{s}{n}} + \varepsilon\Big)\Big] \\ \ge 1 - 2\exp\Big(-\frac{n\varepsilon^2}{2}\Big).\end{aligned}$$

Therefore, there exist constants $C_1, c_1 > 0$, such that

$$\begin{aligned}\mathbb{P}\Big[\|\Omega_T A^*A\Omega_T^* - I_s\| &\le C_1\sqrt{\frac{s}{n}\log m},\ \forall T \text{ s.t. } |T| = s\Big] \\ &\ge 1 - 2\binom{m}{s}\exp\Big(-\Big(\frac{C_1}{6} - 1\Big)^2\frac{s}{2}\log m\Big) \\ &\ge 1 - m^{-c_1 s},\end{aligned}$$

where the first inequality follows from a union bound, and setting $\varepsilon = (\frac{C_1}{6} - 1)\sqrt{\frac{s}{n}\log m}$; the second inequality follows from Stirling's approximation $\binom{m}{s} \le (\frac{em}{s})^s$.

We obtain Lemma 17 by applying the above bound to every diagonal block of the block diagonal matrix $\Omega_{T_x} D^* D \Omega_{T_x}^*$. ■

Proof of Lemma 18: By a consequence of the Hanson-Wright inequality (see [63, Th. 2.1], and its complexification in [63, Sec. 3.1]), there exists an absolute constant c_2' such that

$$\mathbb{P}\Big[\big|\sqrt{n}\|X^\top a_{k\cdot}\|_2 - 1\big| \le \varepsilon\Big] \ge 1 - 2\exp\Big(-\frac{c_2'\varepsilon^2}{\|X\|^2}\Big). \quad (44)$$

Set $\varepsilon = C_2'\|X\|\sqrt{\log n}$ for some $C_2' > 0$, then by a union bound, there exists an absolute constant $c_2 > 0$ such that

$$\mathbb{P}\Big[\big|\sqrt{n}\|X^\top a_{k\cdot}\|_2 - 1\big| \le C_2'\|X\|\sqrt{\log n},\ \forall k \in [n]\Big] \ge 1 - n^{-c_2}. \quad (45)$$

By Assumption 4,

$$\begin{aligned}\mathbb{P}\Big[|\lambda_k|^2\big|a_{k\cdot}^\top X X^* \overline{a_{k\cdot}} - \frac{1}{n}\big| &\le \frac{(2C_2' + C_2'^2)(1+\delta)}{n} \\ &\cdot \max\Big\{\|X\|\sqrt{\log n}, \|X\|^2\log n\Big\},\ \forall k \in [n]\Big] \\ &\ge 1 - n^{-c_2}. \quad (46)\end{aligned}$$

The spectral norm $\|X\|$ is bounded in Assumption 5:

Subspace case: $\|X\|^2 \le (1+\theta)\max\{\frac{1}{N}, \frac{1}{m}\}$,

Joint sparsity case: $\|X\|^2 \le (1+\theta)\max\{\frac{1}{N}, \frac{1}{s_0}\}$.

Therefore, Lemma 18 follows from (38), (41), and (46). ■

Proof of Lemma 19: By (37), the columns of D^*E_s are independent random vectors. Define

$$\phi_k := \begin{bmatrix} \overline{a_{k\cdot}} a_{k\cdot}^\top x_{\cdot 1} \\ \overline{a_{k\cdot}} a_{k\cdot}^\top x_{\cdot 2} \\ \vdots \\ \overline{a_{k\cdot}} a_{k\cdot}^\top x_{\cdot N} \end{bmatrix}.$$

Then $D^*E_s = [\phi_1, \phi_2, \ldots, \phi_n]\mathrm{diag}(\lambda)$. Next, we bound the spectral norm of the random matrix $\Phi - \mathbb{E}\Phi$, where $\Phi := [\phi_1, \phi_2, \ldots, \phi_n]$, using matrix Bernstein inequality [57, Th. 1.6]. We need the following bounds to proceed:

1) A bound on $\|\phi_k - \mathbb{E}\phi_k\|_2$.
First, by [63, Sec. 3.1, Th. 2.1], there exists a constant c_3'

$$\mathbb{P}\Big[\big|\sqrt{n}\|a_{k\cdot}\|_2 - \sqrt{m}\big| \le \varepsilon\Big] \ge 1 - 2\exp(-c_3'\varepsilon^2).$$

By a union bound over all $k \in [n]$, there exists a constant C_3' such that

$$\begin{aligned}\mathbb{P}\Big[\big|\sqrt{n}\|a_{k\cdot}\|_2 - \sqrt{m}\big| &\le C_3'\sqrt{\log n},\ \forall k \in [n]\Big] \\ &\ge 1 - 2n\exp\big(-c_3'C_3'^2\log n\big) \\ &\ge 1 - n^{-c_2}. \quad (47)\end{aligned}$$

Note that

$$\|\mathbb{E}\phi_k\|_2 = 7\frac{1}{n}\|X\|_{\mathrm{F}} = \frac{1}{n}, \quad \|\phi_k\|_2 \le \|a_{k\cdot}\|_2\|X^\top a_{k\cdot}\|_2.$$

By (45) and (47), there exists a constant C_3'', such that with probability at least $1 - 2n^{-c_2}$,

$$\begin{aligned}&\|\phi_k - \mathbb{E}\phi_k\|_2 \\ &\le \frac{C_3''}{n}\max\Big\{\sqrt{m}, \sqrt{\log n}\Big\}\max\Big\{1, \sqrt{\frac{\log n}{N}}, \sqrt{\frac{\log n}{m}}\Big\} \\ &\le \frac{C_3''\sqrt{m}}{n},\end{aligned}$$

for all $k \in [n]$, where the second inequality uses the assumption that $\min\{N, m\} > \log n$.

2) A bound on $\|\mathbb{E}[(\Phi - \mathbb{E}\Phi)^*(\Phi - \mathbb{E}\Phi)]\|$.
One should observe that

$$\begin{aligned}\mathbb{E}[(\phi_k - \mathbb{E}\phi_k)^*(\phi_k - \mathbb{E}\phi_k)] &= \frac{m}{n^2}, \\ \mathbb{E}[(\phi_k - \mathbb{E}\phi_k)^*(\phi_{k'} - \mathbb{E}\phi_{k'})] &= 0,\end{aligned}$$

for $k \ne k'$. Therefore,

$$\begin{aligned}\mathbb{E}[(\Phi - \mathbb{E}\Phi)^*(\Phi - \mathbb{E}\Phi)] &= \frac{m}{n^2}I_n, \\ \|\mathbb{E}[(\Phi - \mathbb{E}\Phi)^*(\Phi - \mathbb{E}\Phi)]\| &= \frac{m}{n^2}.\end{aligned}$$

3) A bound on $\|\mathbb{E}[(\Phi - \mathbb{E}\Phi)(\Phi - \mathbb{E}\Phi)^*]\|$.
Since $\{\phi_k\}_{k=1}^n$ are i.i.d. random vectors,

$$\begin{aligned}&\mathbb{E}[(\Phi - \mathbb{E}\Phi)(\Phi - \mathbb{E}\Phi)^*] \\ &= \sum_{k=1}^n \mathbb{E}[(\phi_k - \mathbb{E}\phi_k)(\phi_k - \mathbb{E}\phi_k)^*] \\ &= n\mathbb{E}[(\phi_1 - \mathbb{E}\phi_1)(\phi_1 - \mathbb{E}\phi_1)^*] \\ &= n[\mathbb{E}(\phi_1\phi_1^*) - (\mathbb{E}\phi_1)(\mathbb{E}\phi_1)^*] \\ &= \frac{1}{n}(X^\top\overline{X} \otimes I_m).\end{aligned}$$

By Assumption 5, in the subspace case,

$$\|\mathbb{E}[(\Phi-\mathbb{E}\Phi)(\Phi-\mathbb{E}\Phi)^*]\| = \frac{1}{n}\|X^\top \overline{X}\| \le \frac{1+\theta}{n}\max\{\frac{1}{N},\frac{1}{m}\}.$$

Given the above bounds, we apply the matrix Bernstein inequality [57, Th. 1.6] as follows:

$$\mathbb{P}\Big[\|\Phi-\mathbb{E}\Phi\|\le\varepsilon \,\Big|\, \|\phi_k-\mathbb{E}\phi_k\|_2\le R,\ \forall k\in[n]\Big] \ge 1-(Nm+n)\exp\Big(-\frac{\varepsilon^2/2}{\sigma^2+R\varepsilon/3}\Big),$$

where

$$\sigma^2=\max\Big\{\frac{m}{n^2},\frac{1+\theta}{nN},\frac{1+\theta}{nm}\Big\},\qquad R=\frac{C_3''\sqrt{m}}{n}.$$

It follows that

$$\mathbb{P}\Big[\|\Phi-\mathbb{E}\Phi\|\le\varepsilon\Big] \ge 1-(Nm+n)\exp\Big(-\frac{\varepsilon^2/2}{\sigma^2+R\varepsilon/3}\Big)-2n^{-c_2},$$

where the last term $2n^{-c_2}$ bounds the probability that $\|\phi_k-\mathbb{E}\phi_k\|_2 > R$ for some k. Hence there exist constants $C_3, c_3 > 0$ such that

$$\mathbb{P}\Big[\|\Phi-\mathbb{E}\Phi\|\le\frac{C_3}{\sqrt{1+\delta}}\max\Big\{\sqrt{\frac{\log(Nm+n)}{nN}},\sqrt{\frac{\log(Nm+n)}{nm}},\frac{\sqrt{m}\log(Nm+n)}{n}\Big\}\Big]\ge 1-n^{-c_3}.$$

Lemma 18 follows from the above bound, and

$$\|\Omega_{T_x}D^*E_s-\mathbb{E}\Omega_{T_x}D^*E_s\| = \|\Phi-\mathbb{E}\Phi\|\|\mathrm{diag}(\lambda)\| \le \sqrt{1+\delta}\|\Phi-\mathbb{E}\Phi\|.\ \blacksquare$$

Proof of Lemma 20: We introduce some notations for this proof. We use B_p^n and $B_{S_p}^{m,n}$ to denote unit balls in $\mathbb{C}^n$ with ℓ_p norm, and in $\mathbb{C}^{m\times n}$ with Schatten p norm, respectively. The projection on the support set T is denoted by Π_T. For a set $\mathcal{A}$ of matrices, $d_{\mathrm{F}}(\mathcal{A})$ and $d_{\mathrm{op}}(\mathcal{A})$ denote the radii of $\mathcal{A}$ in the Frobenius norm and in the spectral norm, respectively. We use $\gamma_2(\mathcal{A},\|\cdot\|)$ the γ_2 functional of $\mathcal{A}$, which is another way to quantify the size of $\mathcal{A}$ [58, Sec. 2.2]. These are key quantities in the upper bound of the supremum of an asymmetric second-order process [58, Th. 2.3], which we use to prove Lemma 20.

Note that

$$\max_{\substack{T\subset[m]\\|T|=s}}\|\Omega_{T_x}D^*E_s-\mathbb{E}\Omega_{T_x}D^*E_s\| = \max_{\substack{T\subset[m]\\|T|=s}}\ \max_{\substack{v\in B_2^{mN}\\(I_N\otimes\Pi_T)v=v}}\ \max_{u\in B_2^n}|v^*\Phi u-\mathbb{E}v^*\Phi u|, \quad (48)$$

where $\Phi = D^*E_s$. Let $z=\sqrt{n}[a_{1\cdot}^*,\dots,a_{n\cdot}^*]^\top$. Then z follows $\mathcal{CN}(\mathbf{0}_{mn,1}, I_{mn})$ and $v^*\Phi u$ is written as a quadratic form in z as follows:

$$v^*\Phi u = \sum_{k=1}^n\sum_{j=1}^N u_k a_{k\cdot}^\top x_{\cdot j}v_{\cdot j}^*\overline{a_{k\cdot}} = z^*(\mathrm{diag}(u)\otimes\Pi_{T_0})\Big(\frac{1}{n}I_n\otimes XV^*\Big)z, \quad (49)$$

where $u=[u_1,\dots,u_n]^\top$, $v=[v_{\cdot 1}^\top,\dots,v_{\cdot N}^\top]^\top$, $V=[v_{\cdot 1},\dots,v_{\cdot N}]$, and $T_0=\{i\in[m]|\|e_i^\top X\|_2>0\}$ denotes the row support of $X=[x_{\cdot 1},\dots,x_{\cdot N}]$.

Let

$$\mathcal{A}=\{A_u|u\in B_2^n\},$$

and

$$\mathcal{B}=\{B_v|v\in B_2^{mN},\ (I_N\otimes\Pi_T)v=v\},$$

where A_u and B_v are left and right factors in the quadratic form in (49), i.e.,

$$A_u=\mathrm{diag}(u)\otimes\Pi_{T_0},$$

and

$$B_v=\frac{1}{n}I_n\otimes XV^*.$$

Then (48) is equivalent to

$$\sup_{A_u\in\mathcal{A}}\sup_{B_v\in\mathcal{B}}|z^*A_uB_vz-\mathbb{E}z^*A_uB_vz|,$$

which is a supremum of an asymmetric second-order process. We use the result on suprema of asymmetric second-order chaos processes by Lee and Junge [58, Th. 2.3], which extends the original result by Krahmer *et al.* [64] to asymmetric cases.

Next, we compute the key quantities, given as functions of $\mathcal{A}$ and $\mathcal{B}$, which we need to apply [58, Th. 2.3]. Let $A_u\in\mathcal{A}$. Since $|T_0|\le s_0$, we have

$$\|A_u\|_{\mathrm{F}}=\sqrt{s_0}\|u\|_2\le\sqrt{s_0}$$

and the radius of $\mathcal{A}$ in the Frobenius norm satisfies

$$d_{\mathrm{F}}(\mathcal{A})\le\sqrt{s_0}.$$

On the other hand,

$$\|A_u\|=\|u\|_\infty\le 1,$$

which implies that the radius of $\mathcal{A}$ in the spectral norm satisfies

$$d_{\mathrm{op}}(\mathcal{A})\le 1.$$

Moreover, for $A_u, A_u'\in\mathcal{A}$, we have

$$\|A_u-A_{u'}\|=\|u-u'\|_\infty.$$

Therefore, by the Dudley's inequality [65],

$$\begin{aligned}\gamma_2(\mathcal{A},\|\cdot\|)&\lesssim\int_0^\infty\sqrt{\log N(\mathcal{A},\|\cdot\|;t)}dt\\&\le\int_0^\infty\sqrt{\log N(B_2^n,\|\cdot\|_\infty;t)}dt\\&\lesssim\int_0^\infty\sqrt{\log N(B_1^n,\|\cdot\|_2;t)}dt\\&\lesssim\log^{3/2}n,\end{aligned}$$

where the third step follows from the entropy duality result by Artstein *et al.* [66] and the last step follows from Maurey's empirical method [67] (also see [68, Lemma 3.1]). Collecting the above estimates shows that the relevant quantities are given by

$$\begin{aligned}
&\gamma_2(\mathcal{A}, \|\cdot\|)(d_{\mathrm{F}}(\mathcal{A}) + \gamma_2(\mathcal{A}, \|\cdot\|)) + d_{\mathrm{F}}(\mathcal{A})d_{\mathrm{op}}(\mathcal{A}) \\
&\qquad \lesssim \max\{\sqrt{s_0}\log^{3/2} n,\ \log^3\ n\}, \\
&d_{\mathrm{op}}(\mathcal{A})(\gamma_2(\mathcal{A}, \|\cdot\|) + d_{\mathrm{F}}(\mathcal{A})) \\
&\qquad \lesssim \max\{\sqrt{s_0},\ \log^{3/2} n\}, \\
&d_{\mathrm{op}}(\mathcal{A})^2 \le 1.
\end{aligned}$$

Next we consider the other set $\mathcal{B}$. Let $B_v \in \mathcal{B}$. Then

$$\|B_v\|_{\mathrm{F}} = \frac{1}{\sqrt{n}}\|XV^*\|_{\mathrm{F}} \le \frac{1}{\sqrt{n}}\|X\|\|V\|_{\mathrm{F}} = \frac{1}{\sqrt{n}}\|X\|.$$

Therefore

$$d_{\mathrm{F}}(\mathcal{B}) \le \frac{1}{\sqrt{n}}\|X\|.$$

On the other hand,

$$\|B_v\| = \frac{1}{n}\|XV^*\| \le \frac{1}{n}\|X\|\|V\|,$$

which implies

$$d_{\mathrm{op}}(\mathcal{B}) \le \frac{1}{n}\|X\|.$$

Moreover, for $B_v, B_{v'} \in \mathcal{B}$, we have

$$\|B_v - B_{v'}\| \le \frac{1}{n}\|X\|\|V - V'\|,$$

where $V' = [v'_{\cdot 1}, \ldots, v'_{\cdot N}]$ and $v' = [v'^{\top}_{\cdot 1}, \ldots, v'^{\top}_{\cdot N}]^{\top}$. Therefore,

$$\begin{aligned}
&\gamma_2(\mathcal{B}, \|\cdot\|) \\
&\quad \lesssim \frac{1}{n}\|X\| \int_0^{\infty} \sqrt{\log N(\cup_{|T|=s} \Pi_T B_{S_2^{m,N}}, \|\cdot\|_{S_\infty^{m,N}}; t)}dt \\
&\quad \le \frac{1}{n}\|X\| \int_0^{1} \sqrt{\log N(\cup_{|T|=s} \Pi_T B_{S_2^{m,N}}, \|\cdot\|_{S_\infty^{m,N}}; t)}dt \\
&\quad \le \frac{1}{n}\|X\| \int_0^{1} \sqrt{\log \sum_{|T|=s} N(\Pi_T B_{S_2^{m,N}}, \|\cdot\|_{S_\infty^{m,N}}; t)}dt \\
&\quad \le \frac{1}{n}\|X\| \int_0^{1} \sqrt{s\log m + \log N(B_{S_2^{s,N}}, \|\cdot\|_{S_\infty^{s,N}}; t)}dt \\
&\quad \le \frac{1}{n}\|X\| \Big(\sqrt{s\log m} \\
&\qquad + \int_0^{1} \sqrt{\log N(B_{S_2^{s,N}}, \|\cdot\|_{S_\infty^{s,N}}; t)}dt\Big) \\
&\quad \lesssim \frac{1}{n}\|X\|\sqrt{s+N}\log(sN+m),
\end{aligned}$$

where the last step follows from Lemma 27. Therefore, the parameters for $\mathcal{B}$ are estimated as

$$\begin{aligned}
&\gamma_2(\mathcal{B}, \|\cdot\|)(d_{\mathrm{F}}(\mathcal{B}) + \gamma_2(\mathcal{B}, \|\cdot\|)) + d_{\mathrm{F}}(\mathcal{B})d_{\mathrm{op}}(\mathcal{B}) \\
&\qquad \lesssim \frac{1}{n^2}\|X\|^2((s+N)\log^2(sN+m) \\
&\qquad\quad + \sqrt{s+N}\sqrt{n}\log(sN+m)), \\
&d_{\mathrm{op}}(\mathcal{B})(\gamma_2(\mathcal{B}, \|\cdot\|) + d_{\mathrm{F}}(\mathcal{B})) \\
&\qquad \lesssim \frac{1}{n^2}\|X\|^2(\sqrt{s+N}\log(sN+m) + \sqrt{n}), \\
&d_{\mathrm{op}}(\mathcal{B})^2 \le \frac{1}{n^2}\|X\|^2.
\end{aligned}$$

According to [58, Th. 2.3], the optimal upper bound is obtained as the geometric mean of the dominant parameters for the two sets. More precisely, the suprema is (up to an absolute constant) no larger than

$$\frac{s_0^{1/4}(s+N)^{1/4}(\sqrt{n}+\sqrt{s+N})^{1/2}}{n} \cdot \|X\|\log^3\ n\log(sN+m)$$

with probability $1 - n^{-c_3}$. By Assumptions 4 and 5,

$$\begin{aligned}
&|\lambda_k| \le \sqrt{1+\delta}, \\
&\|X\| \le \max\left\{\sqrt{\frac{1+\theta}{N}},\ \sqrt{\frac{1+\theta}{s_0}}\right\},
\end{aligned}$$

which completes the proof. ∎

Lemma 27:

$$\int_0^{\infty} \sqrt{\log N(B_{S_2^{m,N}}, tB_{S_\infty^{m,N}})}dt \lesssim \sqrt{m+N}\log(mN).$$

Proof of Lemma 27: First, by the dual entropy result by Artstein *et al.* [66], we have

$$\log N(B_{S_2^{m,N}}, tB_{S_\infty^{m,N}}) \lesssim \log N(B_{S_1^{m,N}}, tB_{S_2^{m,N}}).$$

Then we approximate the S_1 ball as a polytope using a trick proposed by Junge and Lee [68]. Let R be the set of all rank-1 matrices in the unit sphere of $S_2^{m,N}$. Then $B_{S_1^{m,N}}$ is the absolute convex hull of R. We construct an ϵ-net Δ_m of the sphere S^{m-1}. Then

$$|\Delta_m| \le \left(1 + \frac{2}{\epsilon}\right)^m.$$

For an arbitrary $f \in S^{m-1}$, we have a sequence $\{f_l\}_{l=1}^{\infty} \subset \Delta_m$ such that

$$f = \sum_{l=1}^{\infty} \alpha_l f_l,$$

and

$$\sum_{l=1}^{\infty} |\alpha_l| \le \frac{1}{1-\epsilon}.$$

The existence of such a sequence follows from the optimality of the construction of the net. Similarly we construct an ϵ-net $\Delta_N \subset S^{N-1}$ of S^{N-1}. Then

$$|\Delta_N| \le \left(1 + \frac{2}{\epsilon}\right)^N.$$

For an arbitrary $g \in S^{N-1}$, we have a sequence $\{g_k\}_{k=1}^{\infty} \subset \Delta_N$ such that

$$g = \sum_{k=1}^{\infty} \beta_k g_k$$

and

$$\sum_{k=1}^{\infty} |\beta_k| \leq \frac{1}{1-\epsilon}.$$

Therefore,

$$fg^* = \sum_{l,k=1}^{\infty} \alpha_l \beta_k f_l g_k^*$$

and

$$\sum_{l,k=1}^{\infty} |\alpha_l||\beta_k| \leq \left(\frac{1}{1-\epsilon}\right)^2.$$

We can choose ϵ so that

$$\left(\frac{1}{1-\epsilon}\right)^2 \leq 2$$

and

$$1 + \frac{2}{\epsilon} \leq 8.$$

Let $\Delta_{m,N} = \Delta_m \times \Delta_N$. Then

$$\log(|\Delta_{m,N}|) \leq (m+N)\log 8$$

and

$$B_{S_1^{m,N}} \subset 2\text{absconv}(\Delta_{m,N}).$$

Now, it suffices to compute

$$\int_0^{\infty} \sqrt{\log N(2\text{absconv}(\Delta_{m,N}), tB_{S_2^{m,N}})}dt.$$

Then use a change of variable and get

$$\int_0^{\infty} \sqrt{\log N(2\text{absconv}(\Delta_{m,N}), tB_{S_2^{m,N}})}dt = 2\int_0^{\infty} \sqrt{\log N(\text{absconv}(\Delta_{m,N}), tB_{S_2^{m,N}})}dt.$$

Let $\Delta_{m,N} = \{q_1, \ldots, q_M\}$, where $M = |\Delta_{m,N}|$. Define linear mapping $Q : \ell_1^M \to \ell_2^{mN}$ by $Q(e_i) = \text{vec}(q_i)$ for $i = 1, \ldots, M$. Since $\|\text{vec}(q_i)\|_2 = \|q_i\|_{S_2} = 1$ for all i, we have

$$\|Q : \ell_1^M \to \ell_2^{mN}\| = 1.$$

Note

$$\int_0^{\infty} \sqrt{\log N(\text{absconv}(\Delta_{m,N}), tB_{S_2^{m,N}})}dt = \int_0^{\infty} \sqrt{\log N(Q(B_1^M), tB_{\ell_2^{mN}})}dt.$$

By a version of Maurey's empirical method (see for example [68, Proposition 3.2]), we have

$$\int_0^{\infty} \sqrt{\log N(Q(B_1^M), tB_{\ell_2^{mN}})}dt \lesssim \sqrt{\log M}\log(mN) \lesssim \sqrt{m+N}\log(mN).$$

This completes the proof. ∎

Proof of Lemma 22: Bear in mind that the columns of $\Psi := D^*E_{\text{n}}$, which we denote by $\{\psi_k\}_{k=1}^n$, are independent random vectors with zero mean:

$$\psi_k := \begin{bmatrix} \overline{a_{k\cdot}} w_{k1} \\ \overline{a_{k\cdot}} w_{k2} \\ \vdots \\ \overline{a_{k\cdot}} w_{kN} \end{bmatrix}.$$

We bound $\|D^*E_{\text{n}}\|$ using the matrix Bernstein inequality [57, Th. 1.6]. We need the following bounds:

1) A bound on $\|\psi_k\|_2$.
Since

$$\|\psi_k\|_2 \leq \|a_{k\cdot}\|_2 \|w_{k\cdot}\|_2$$

By (47), and $m > \log n$,

$$\|\psi_k\|_2 \leq (C_3' + 1)\sqrt{\frac{m}{n}} \times \sqrt{N} \max_{k\in[n], j\in[N]} |w_{kj}|,$$

with probability at least $1 - n^{-c_2}$.

2) A bound on $\|\mathbb{E}\Psi^*\Psi\|$.
Since

$$\mathbb{E}\Psi^*\Psi = \frac{m}{n}\text{diag}([\|w_{1\cdot}\|_2^2, \|w_{2\cdot}\|_2^2, \ldots, \|w_{k\cdot}\|_2^2]),$$

we have

$$\|\mathbb{E}\Psi^*\Psi\| = \frac{m}{n}\max_{k\in[n]} \|w_{k\cdot}\|_2^2 \leq \frac{mN}{n} \max_{k\in[n], j\in[N]} |w_{kj}|^2.$$

3) A bound on $\|\mathbb{E}\Psi\Psi^*\|$.
Since

$$\mathbb{E}\Psi\Psi^* = \sum_{k\in[n]} \frac{1}{n}\text{diag}([|w_{k1}|^2, |w_{k2}|^2, \ldots, |w_{kN}|^2]) \otimes I_m,$$

we have

$$\|\mathbb{E}\Psi\Psi^*\| = \frac{1}{n}\max_{j\in[N]} \sum_{k\in[n]} |w_{kj}|^2 \leq \max_{k\in[n], j\in[N]} |w_{kj}|^2.$$

Given the above bounds, we completes the proof using the matrix Bernstein inequality (similar to the proof of Lemma 19). There exist constants $C_4, c_4 > 0$ such that

$$\|D^*E_{\text{n}}\| = \|\Psi\| \leq C_4 \max\left\{\sqrt{\log(Nm+n)}, \sqrt{\frac{Nm}{n}}\log(Nm+n)\right\} \max_{k\in[n], j\in[N]} |w_{kj}|,$$

with probability at least $1 - n^{-c_4}$. ∎

Proof of Lemma 23: Note that

$$\max_{\substack{T\subset[m] \\ |T|=s}} \|\Omega_{T_x} D^*E_{\text{n}}\| = \max_{\substack{T\subset[m] \\ |T|=s}} \max_{\substack{v\in B_2^{mN} \\ (I_N\otimes\Pi_T)v=v}} \max_{u\in B_2^n} |v^*\Psi u|,$$

where

$$\Psi = D^*E_{\text{n}} = \left[I_N \otimes \overline{a_{1\cdot}} \ \ldots \ I_N \otimes \overline{a_{n\cdot}}\right] \begin{bmatrix} w_{1\cdot} & & \\ & \ddots & \\ & & w_{n\cdot} \end{bmatrix}.$$

Let $z = \sqrt{n}[a_{1\cdot}^*, \ldots, a_{n\cdot}^*]^\top$. Then z is a standard Gaussian vector, and

$$v^*\Psi u = \frac{1}{\sqrt{n}}(\mathbf{1}_{1,n} \otimes v^*)(E_n \otimes I_m)(\mathrm{diag}(u) \otimes I_m)z.$$

Let

$$q_{u,v} := \frac{1}{\sqrt{n}}(\mathrm{diag}(u)^* \otimes I_m)(E_n^* \otimes I_m)(\mathbf{1}_{n,1} \otimes v).$$

The L_2 metric is given by

$$\begin{aligned} d((u,v),(u',v')) &= \sqrt{\mathbb{E}(q_{u,v}^* z - q_{u',v'}^* z)^2} \\ &= \|q_{u,v} - q_{u',v'}\|_2. \end{aligned}$$

Indeed,

$$\begin{aligned} d((u,v),(u',v')) &\le d((u,v),(u,v')) + d((u,v'),(u',v')) \\ &\le \|\mathrm{diag}(u-u')\|_\infty \|E_n\| \|v\|_2 \\ &\quad + \|\mathrm{diag}(u')\|_\infty \|E_n\| \|v - v'\|_2 \\ &\le \|\mathrm{diag}(u-u')\|_\infty \|E_n\| \\ &\quad + \|E_n\| \|v - v'\|_2. \end{aligned}$$

Let $\Gamma_s = \{v \in B_2^{mN} : T \subset [m],\ |T| = s,\ (I_N \otimes \Pi_T)v = v\}$. By Dudley's theorem (see e.g., [65, Th. 11.17]), we have

$$\begin{aligned} &\mathbb{E} \sup_{\substack{T\subset[m]\\|T|=s}} \sup_{\substack{v\in B_2^{mN}\\(I_N\otimes\Pi_T)v=v}} \sup_{u\in B_2^n} v^*\Psi u \\ &\le 24 \int_0^\infty \sqrt{\log N(\Gamma_s \times B_2^n, d(\cdot); \epsilon)}\, d\epsilon \\ &\le 24\|E_n\| \Big(\int_0^\infty \sqrt{\log N(\Gamma_s, \|\cdot\|_2; \epsilon)}\, d\epsilon \\ &\quad + \int_0^\infty \sqrt{\log N(B_2^n, \|\cdot\|_\infty; \epsilon)}\, d\epsilon \Big) \\ &\le 24\|E_n\| \Big(\int_0^\infty \sqrt{\log N(\Gamma_s, \|\cdot\|_2; \epsilon)}\, d\epsilon \\ &\quad + \int_0^\infty \sqrt{\log N(B_1^n, \|\cdot\|_2; \epsilon)}\, d\epsilon \Big) \\ &\lesssim \|E_n\|(\sqrt{s\log m} + \sqrt{Ns} + \log^{3/2} n). \end{aligned}$$

By an extension of Dudley's inequality to moments [69, p. 263, Sec. 8.9],

$$\begin{aligned} &\Big(\mathbb{E} \sup_{\substack{T\subset[m]\\|T|=s}} \sup_{\substack{v\in B_2^{mN}\\(I_N\otimes\Pi_T)v=v}} \sup_{u\in B_2^n} |v^*\Psi u|^p\Big)^{1/p} \\ &\qquad \lesssim \|E_n\|(\sqrt{s\log m} + \sqrt{Ns} + \log^{3/2} n)\sqrt{p}. \end{aligned}$$

By a variation of Markov's inequality [69, Proposition 7.11], there exist absolute constants $C_4, c_4 > 0$ such that

$$\begin{aligned} &\sup_{\substack{T\subset[m]\\|T|=s}} \sup_{\substack{v\in B_2^{mN}\\(I_N\otimes\Pi_T)v=v}} \sup_{u\in B_2^n} |v^*\Psi u| \\ &\qquad \le C_4\|E_n\|(\sqrt{s\log m} + \sqrt{Ns} + \log^{3/2} n)\sqrt{\log n}, \end{aligned}$$

with probability at least $1 - n^{-c_4}$.

Therefore, Lemma 23 follows from

$$\|E_n\| = \max_{k\in[n]} \|w_{k\cdot}\|_2 \le \sqrt{N} \max_{k\in[n],j\in[N]} |w_{kj}|. \qquad \blacksquare$$

Proof of Lemma 24: If assumptions 3 – 5 are satisfied, then by (45),

$$\|y_{k\cdot}\|_2 \le \frac{(C_2'+1)\sqrt{1+\delta}}{\sqrt{n}} \max\Big\{1, \|X\|\sqrt{\log n}\Big\}$$

for all $k \in [n]$, with probability at least $1 - n^{-c_2}$.

Since

$$E_s^* E_n = \mathrm{diag}([y_{1\cdot}^* w_{1\cdot}, y_{2\cdot}^* w_{2\cdot}, \ldots, y_{n\cdot}^* w_{n\cdot}]),$$

there exist constants $C_5 = (C_2'+1)\sqrt{(1+\delta)(1+\theta)} > 0$ such that

$$\begin{aligned} \|E_s^* E_n\| &\le \max_k \|y_{k\cdot}\|_2 \times \sqrt{N} \max_{k\in[n],j\in[N]} |w_{kj}| \\ &\le \frac{C_5}{\sqrt{1+\theta}}\sqrt{\frac{N}{n}} \max\Big\{1, \|X\|\sqrt{\log n}\Big\} \max_{k\in[n],j\in[N]} |w_{kj}|, \end{aligned}$$

with probability at least $1 - n^{-c_2}$. Therefore, Lemma 24 follows from Assumption 5. $\blacksquare$

Proof of Lemma 25: Lemma 25 follows from

$$E_n^* E_n = \mathrm{diag}([\|w_{1\cdot}\|_2^2, \|w_{2\cdot}\|_2^2, \ldots, \|w_{n\cdot}\|_2^2]). \qquad \blacksquare$$

Proof of Lemma 26: We prove these inequalities using the Hoeffding's inequality.

For all $j \in [N]$, $\ell \in [m]$, and $k \in [n]$,

$$\begin{aligned} &\Big||\overline{a_{k\ell}} a_{k\cdot}^\top x_{\cdot j}|^2 - \mathbb{E}|\overline{a_{k\ell}} a_{k\cdot}^\top x_{\cdot j}|^2\Big| \\ &\le |a_{k\ell}|^2 |a_{k\cdot}^\top x_{\cdot j}|^2 + \frac{1}{n^2}(\|x_{\cdot j}\|_2^2 + |x_{\ell j}|^2) \\ &\le C_6' \frac{\log(nm)}{n} \cdot \frac{\|x_{\cdot j}\|_2^2 \log(nN)}{n} + \frac{2\|x_{\cdot j}\|_2^2}{n^2} \\ &\le \frac{(C_6'+2)\|x_{\cdot j}\|_2^2 \log^2(nmN)}{n^2}, \end{aligned}$$

where the third line is true with probability at least $1 - n^{-c_6'}$ for some absolute constant c_6'. We show this by applying a Chernoff bound and a union bound to $|a_{k\ell}|^2$, and applying the Hanson-Wright inequality (44) and a union bound to $|a_{k\cdot}^\top x_{\cdot j}|^2$. Then it follows from the Hoeffding's inequality and a union bound, that there exist absolute constants $C_6, c_6 > 0$ such that for all $j \in [N]$ and $\ell \in [m]$ we have (20).

Similarly, for all $j \in [N]$, $\ell \in [m]$, and $k \in [n]$,

$$|a_{k\ell}|^2 |a_{k\cdot}^\top x_{\cdot j}| \le C_6' \frac{\log(nm)}{n} \cdot \frac{\|x_{\cdot j}\|_2 \sqrt{\log(nN)}}{\sqrt{n}},$$

with probability at least $1 - n^{-c_6'}$. By the Hoeffding's inequality and a union bound, we have (21). Here we use the following facts: By Assumption 5, $\|x_{\cdot j}\| \ge \sqrt{\frac{1-\theta}{N}}$. By Assumption 6, $\max_{k\in[n],j\in[N]} |w_{kj}| \le \frac{C_W}{\sqrt{nN}}$.

For $\ell \in [m]$ and $k \in [n]$,

$$|a_{k\ell}|^2 - \mathbb{E}|a_{k\ell}|^2 \le C_6' \frac{\log(nm)}{n},$$

with probability at least $1 - n^{-c_6'}$. By the Hoeffding's inequality and a union bound, we have (22). $\blacksquare$

REFERENCES

[1] Y. Li, K. Lee, and Y. Bresler, "Blind gain and phase calibration for low-dimensional or sparse signal sensing via power iteration," in *Proc. Int. Conf. Sampling Theory Appl. (SampTA)*, Jul. 2017, pp. 119–123.

[2] H. Q. Nguyen, S. Liu, and M. N. Do, "Subspace methods for computational relighting," in *Computational Imaging XI*, C. A. Bouman, I. Pollak, and P. J. Wolfe, Eds. Bellingham, WA, USA: SPIE, Feb. 2013.

[3] R. L. Morrison, M. N. Do, and D. C. Munson, "MCA: A multichannel approach to SAR autofocus," *IEEE Trans. Image Process.*, vol. 18, no. 4, pp. 840–853, Apr. 2009.

[4] A. Paulraj and T. Kailath, "Direction of arrival estimation by eigenstructure methods with unknown sensor gain and phase," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, vol. 10, Apr. 1985, pp. 640–643.

[5] Y. Li, K. Lee, and Y. Bresler, "Identifiability in bilinear inverse problems with applications to subspace or sparsity-constrained blind gain and phase calibration," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 822–842, Feb. 2017.

[6] Y. Li, K. Lee, and Y. Bresler, "Optimal sample complexity for blind gain and phase calibration," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5549–5556, Nov. 2016.

[7] S. Ling and T. Strohmer. (2016). "Self-calibration and bilinear inverse problems via linear least squares." [Online]. Available: https://arxiv.org/abs/1611.04196

[8] L. Wang and Y. Chi, "Blind deconvolution from multiple sparse inputs," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1384–1388, Oct. 2016.

[9] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: Johns Hopkins Univ. Press, 1996.

[10] B. Moghaddam, Y. Weiss, and S. Avidan, "Spectral bounds for sparse PCA: Exact and greedy algorithms," in *Advances in Neural Information Processing Systems*, Y. Weiss, P. B. Schölkopf, and J. C. Platt, Eds. Cambridge, MA, USA: MIT Press, 2006, pp. 915–922.

[11] X.-T. Yuan and T. Zhang, "Truncated power method for sparse eigenvalue problems," *J. Mach. Learn. Res.*, vol. 14, no. 4, pp. 899–925, 2013.

[12] Y. Rockah, H. Messer, and P. M. Schultheiss, "Localization performance of arrays subject to phase errors," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-24, no. 4, pp. 402–410, Jul. 1988.

[13] A. J. Weiss and B. Friedlander, "Eigenstructure methods for direction finding with sensor gain and phase uncertainties," *Circuits, Syst. Signal Process.*, vol. 9, no. 3, pp. 271–300, 1990.

[14] L. Balzano and R. Nowak, "Blind calibration of sensor networks," in *Proc. 6th Int. Symp. Inf. Process. Sensor Netw. (IPSN)*, Apr. 2007, pp. 79–88.

[15] J. Lipor and L. Balzano, "Robust blind calibration via total least squares," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2014, pp. 4244–4248.

[16] Ç. Bilen, G. Puy, R. Gribonval, and L. Daudet, "Convex optimization approaches for blind sensor calibration using sparsity," *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4847–4856, Sep. 2014.

[17] S. Ling and T. Strohmer, "Self-calibration and bilinear inverse problems via linear least squares," *SIAM J. Imag. Sci.*, vol. 11, no. 1, pp. 252–292, Jan. 2018.

[18] Y. C. Eldar, W. Liao, and S. Tang. (2017). "Sensor calibration for off-the-grid spectral estimation." [Online]. Available: https://arxiv.org/abs/1707.03378

[19] L. Tong, G. Xu, and T. Kailath, "A new approach to blind identification and equalization of multipath channels," in *Proc. Conf. Rec. 25th Asilomar Conf. Signals, Syst. Comput.*, vol. 2, Nov. 1991, pp. 856–860.

[20] E. Moulines, P. Duhamel, J.-F. Cardoso, and S. Mayrargue, "Subspace methods for the blind identification of multichannel FIR filters," *IEEE Trans. Signal Process.*, vol. 43, no. 2, pp. 516–525, Feb. 1995.

[21] G. Harikumar and Y. Bresler, "FIR perfect signal reconstruction from multiple convolutions: Minimum deconvolver orders," *IEEE Trans. Signal Process.*, vol. 46, no. 1, pp. 215–218, Jan. 1998.

[22] G. Harikumar and Y. Bresler, "Perfect blind restoration of images blurred by multiple filters: Theory and efficient algorithms," *IEEE Trans. Image Process.*, vol. 8, no. 2, pp. 202–219, Feb. 1999.

[23] K. Lee, F. Krahmer, and J. Romberg. (2017). "Spectral methods for passive imaging: Non-asymptotic performance and robustness." [Online]. Available: https://arxiv.org/abs/1708.04343

[24] K. Lee, N. Tian, and J. Romberg, "Fast and guaranteed blind multichannel deconvolution under a bilinear system model," *IEEE Trans. Inf. Theory*, vol. 64, no. 7, pp. 4792–4818, Jul. 2018.

[25] M. A. Davenport and J. Romberg, "An overview of low-rank matrix recovery from incomplete observations," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 608–622, Jun. 2016.

[26] Y. Li, K. Lee, and Y. Bresler, "Optimal sample complexity for stable matrix recovery," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 81–85.

[27] M. Kech and F. Krahmer, "Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems," *SIAM J. Appl. Algebra Geometry*, vol. 1, no. 1, pp. 20–37, 2017.

[28] K. Lee, Y. Wu, and Y. Bresler, "Near-optimal compressed sensing of a class of sparse low-rank matrices via sparse power factorization," *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1666–1698, Mar. 2018.

[29] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *IEEE Trans. Inf. Theory*, vol. 60, no. 3, pp. 1711–1732, Mar. 2014.

[30] S. Ling and T. Strohmer, "Self-calibration and biconvex compressive sensing," *Inverse Problems*, vol. 31, no. 11, p. 115002, 2015.

[31] Y. Chi, "Guaranteed blind sparse spikes deconvolution via lifting and convex optimization," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 782–794, Jun. 2016.

[32] K. Lee, Y. Li, M. Junge, and Y. Bresler, "Blind recovery of sparse signals from subsampled convolution," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 802–821, Feb. 2017.

[33] X. Li, S. Ling, T. Strohmer, and K. Wei, "Rapid, robust, and reliable blind deconvolution via nonconvex optimization," *Appl. Comput. Harmon. Anal.*, to be published, doi: 10.1016/j.acha.2018.01.001.

[34] S. Bahmani and J. Romberg, "Lifting for blind deconvolution in random mask imaging: Identifiability and convex relaxation," *SIAM J. Imag. Sci.*, vol. 8, no. 4, pp. 2203–2238, 2015.

[35] G. Tang and B. Recht, "Convex blind deconvolution with random masks," in *Proc. Comput. Opt. Sens. Imag.*, 2014, p. CW4C-1.

[36] V. Cambareri and L. Jacques, "A non-convex blind calibration method for randomised sensing strategies," in *Proc. 4th Int. Workshop Compressed Sens. Theory Appl. Radar, Sonar Remote Sens. (CoSeRa)*, Sep. 2016, pp. 16–20.

[37] V. Cambareri, A. Moshtaghpour, and L. Jacques, "A greedy blind calibration method for compressed sensing with unknown sensor gains," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2017, pp. 1132–1136.

[38] A. Ahmed and L. Demanet. (2016). "Leveraging diversity and sparsity in blind deconvolution." [Online]. Available: https://arxiv.org/abs/1610.06098

[39] A. Cosse. (2017). "From blind deconvolution to blind super-resolution through convex programming." [Online]. Available: https://arxiv.org/abs/1709.09279

[40] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, Aug. 1982.

[41] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, Aug. 2013.

[42] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," in *Advances in Neural Information Processing Systems*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. New York, NY, USA: Curran Associates, 2013, pp. 2796–2804.

[43] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.

[44] T. T. Cai, X. Li, and Z. Ma, "Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow," *Ann. Statist.*, vol. 44, no. 5, pp. 2221–2251, Oct. 2016.

[45] S. Bahmani and J. Romberg, "Phase retrieval meets statistical learning theory: A flexible convex relaxation," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, in Proceedings of Machine Learning Research, vol. 54, A. Singh and J. Zhu, Eds. Fort Lauderdale, FL, USA: PMLR, Apr. 2017, pp. 252–260.

[46] T. Goldstein and C. Studer. (2016). "Phasemax: Convex phase retrieval via basis pursuit." [Online]. Available: https://arxiv.org/abs/1610.07531

[47] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: A contemporary overview," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, May 2015.

[48] J. Sun, Q. Qu, and J. Wright. (2015). "When are nonconvex problems not scary?" [Online]. Available: https://arxiv.org/abs/1510.06096

[49] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2980–2998, Jun. 2010.

[50] P. Jain, P. Netrapalli, and S. Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proc. 45th Annu. ACM Symp. Theory Comput. (STOC)*, 2013, pp. 665–674.

[51] R. Sun and Z.-Q. Luo, "Guaranteed matrix completion via non-convex factorization," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6535–6579, Nov. 2016.

[52] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere I: Overview and the geometric picture," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 853–884, Feb. 2017.

[53] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere II: Recovery by Riemannian trust-region method," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 885–914, Feb. 2017.

[54] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 2379–2383.

[55] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[56] C. Davis and W. M. Kahan, "The rotation of eigenvectors by a perturbation. III," *SIAM J. Numer. Anal.*, vol. 7, no. 1, pp. 1–46, Mar. 1970.

[57] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Found. Comput. Math.*, vol. 12, no. 4, pp. 389–434, Aug. 2011.

[58] K. Lee and M. Junge. (2015). "RIP-like properties in subsampled blind deconvolution." [Online]. Available: https://arxiv.org/abs/1511.06146

[59] Q. Berthet and P. Rigollet. (2013). "Computational lower bounds for sparse PCA." [Online]. Available: https://arxiv.org/abs/1304.0828

[60] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Uniqueness guarantees and recovery algorithms," *IEEE Trans. Signal Process.*, vol. 65, no. 9, pp. 2402–2410, May 2017.

[61] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.

[62] K. R. Davidson and S. J. Szarek, "Local operator theory, random matrices and banach spaces," in *Handbook Geometry Banach Spaces*. Amsterdam, The Netherlands: Elsevier BV, 2001, pp. 317–366.

[63] M. Rudelson and R. Vershynin, "Hanson-Wright inequality and sub-Gaussian concentration," *Electron. Commun. Probab.*, vol. 18, no. 82, p. 9, 2013.

[64] F. Krahmer, S. Mendelson, and H. Rauhut, "Suprema of chaos processes and the restricted isometry property," *Commun. Pure Appl. Math.*, vol. 67, no. 11, pp. 1877–1904, Nov. 2014.

[65] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes*. Berlin, Germany: Springer-Verlag, 2011.

[66] S. Artstein, V. Milman, and S. J. Szarek, "Duality of metric entropy," *Ann. Math.*, vol. 159, no. 3, pp. 1313–1328, 2004.

[67] B. Carl, "Inequalities of Bernstein-Jackson-type and the degree of compactness of operators in Banach spaces," *Annales de l'institut Fourier*, vol. 35, no. 3, pp. 79–118, 1985.

[68] M. Junge and K. Lee. (2017). "Generalized notions of sparsity and restricted isometry property. Part I: A unified framework." [Online]. Available: https://arxiv.org/abs/1706.09410

[69] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Basel, Switzerland: Springer, 2013.

Yanjun Li received the Bachelor of Engineering degree in Automation from Tsinghua University (Beijing, China) in 2012, with the highest honor. He received the M.S. degree and the Ph.D. degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2015 and 2018, respectively. His research interests center around low-dimensional and parsimonious structures in signal processing and machine learning. He is the recipient of a Best Student Paper Prize at SPARS 2015. He received the Robert T. Chien Memorial Award in 2016, the Hong, McCully, and Allen Fellowship, and a Scott Dissertation Completion Fellowship in 2017, from the University of Illinois at Urbana-Champaign.

Kiryung Lee received the Ph.D. degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2012. He joined the Department of Electrical and Computer Engineering at The Ohio State University as an Assistant Professor in 2018. Prior to this, he was a Research Scientist II in the School of Electrical and Computer Engineering at Georgia Institute of Technology in 2016-2018 and a Visiting Assistant Professor in the Department of Statistics at the University of Illinois at Urbana-Champaign in 2014-2015. In his earlier career, he worked as a research engineer at the Electronic Telecommunications Research Institute and the LG Electronics in 2002-2004 and 2004-2006, respectively. His research interests are in developing mathematical theory and computational methods for optimization problems in imaging, signal processing, statistics, and machine learning.

Yoram Bresler received the B.Sc. *(cum laude)* and M.Sc. degrees from the Technion, Israel Institute of Technology, in 1974 and 1981, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, USA, in 1986, all in electrical engineering. In 1987 he joined the University of Illinois at Urbana-Champaign, where he currently holds the position of a Founder Professor of Engineering with the Departments of Electrical and Computer Engineering and Bioengineering, and the Coordinated Science Laboratory. He is also President and Chief Technology Officer at InstaRecon, Inc., a startup company he cofounded to commercialize breakthrough technology for tomographic reconstruction developed in his academic research. His current research interests include machine learning and statistical signal processing and their applications to inverse problems in imaging, and in particular compressed sensing, computed tomography, and magnetic resonance imaging.

Dr. Bresler is an IEEE Fellow, and Fellow of the AIMBE. He received two Senior Paper Awards from the IEEE Signal Processing society, and two papers he coauthored with his students received the Young Author Award from the Signal Processing society in 2002 and 2016, respectively. He is the recipient of a 1991 NSF Presidential Young Investigator Award, the Technion Fellowship in 1995, and the Xerox Senior Award for Faculty Research in 1998. He was named a University of Illinois Scholar in 1999, appointed as an Associate at the Center for Advanced Study of the University in 2001-2002, and Faculty Fellow at the National Center for Super Computing Applications in 2006. In 2016 he was appointed an IEEE Signal Processing Society Distinguished Lecturer. He has served on the editorial board of a number of journals including the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING, *Machine Vision and Applications*, and the *SIAM Journal on Imaging Science*, and on various committees of the IEEE.