

Constructing confidence intervals for the signals in sparse phase retrieval*

Yisha Yao

Department of Biostatistics, Yale University, New Haven, CT 06511
e-mail: yisha.yao@yale.edu

Abstract: In this paper, we provide a general methodology to draw statistical inferences on individual signal coordinates or linear combinations of them in sparse phase retrieval. Given an initial estimator of the targeting parameter, which is generated by some existing algorithm, we can modify it in a way that the modified version is asymptotically normal and unbiased. Then confidence intervals and hypothesis testing can be constructed based on this asymptotic normality. For conciseness, we focus on confidence intervals in this work, while a similar procedure can be adopted for hypothesis testing. Under some mild assumptions on the signal and sample size, we establish theoretical guarantees for the proposed method. These assumptions are generally weak in the sense that the dimension could exceed the sample size and many non-zero small coordinates are allowed. Furthermore, theoretical analysis reveals that the modified estimators for individual coordinates have uniformly bounded variance, and hence simultaneous inference is possible. Numerical simulations in a wide range of settings are supportive of our theoretical results.

MSC2020 subject classifications: Primary 60K35, 60K35; secondary 60K35.

Keywords and phrases: Sparse phase retrieval, statistical inference, confidence interval, high dimension.

Received February 2021.

1. Introduction

The problem of recovering a signal from its transformed measurements, referred to as phase retrieval, is fundamental in various applications, including optical imaging, X-ray crystallography, speech recognition, and so on [28]. It can be formulated into model (1.1).

$$y_j = |\mathbf{x}_j^* \boldsymbol{\beta}|^2 + \varepsilon_j, \quad j = 1, \dots, n, \quad (1.1)$$

where ε_j is a random noise with mean zero, $\mathbf{x}_j, \boldsymbol{\beta} \in \mathbb{C}^p$ or \mathbb{R}^p , and \mathbf{x}_j^* denotes the conjugate transpose of \mathbf{x}_j . Given the noise-contaminated magnitudes y_j 's and the design vectors \mathbf{x}_j 's, we need to recover the signal $\boldsymbol{\beta}$. The \mathbf{x}_j 's could be Fourier basis, Gaussian vectors, or other sensing vectors, depending on the specific scenario. Phase retrieval is difficult because the phase information is

*Y. Yao is supported partially by NSF grants DMS-1721495 and IIS-1741390.

totally lost in the data-acquisition process. Extensive literature is available on the theory and algorithms for estimating β .

The early-stage algorithms pioneered by Gerchberg and Saxton [15] and extended by Fienup [14] start with an arbitrary guess, then refine it by transforming back and forth between the signal domain and Fourier domain until all the constraints are satisfied. Since the violation between the iterate and the *a priori* knowledge is monotonically non-increasing, this type of algorithms get the name error reduction algorithms [13]. Such scheme is equivalent to alternating projections onto nonconvex sets [21, 1], but its convergence nature remains unknown. Besides, error reduction algorithms rely heavily on the prior information about the signal. Following the spirit of Gerchberg-Saxton algorithm, the alternating minimization is recently proposed [24]. It divides the data into a number of independent parts and uses a new part in each minimization step. Nevertheless, this strategy is of little practical value.

In most literature, phase retrieval is translated into a nonconvex minimization problem with various objective functions, for example:

$$\underset{\mathbf{b} \in \mathbb{C}^p / \mathbb{R}^p}{\text{minimize}} f(\mathbf{b}) = \frac{1}{4n} \sum_{j=1}^n (|\mathbf{x}_j^* \mathbf{b}|^2 - y_j)^2, \quad (1.2)$$

$$\underset{\mathbf{b} \in \mathbb{C}^p / \mathbb{R}^p}{\text{minimize}} f(\mathbf{b}) = \frac{1}{4n} \sum_{j=1}^n (|\mathbf{x}_j^* \mathbf{b}| - \phi_j)^2, \quad \phi_j = |\mathbf{x}_j^* \beta| + \varepsilon_j, \quad (1.3)$$

$$\underset{\mathbf{b} \in \mathbb{C}^p / \mathbb{R}^p}{\text{minimize}} f(\mathbf{b}) = \frac{1}{4n} \sum_{j=1}^n \left| |\mathbf{x}_j^* \mathbf{b}|^2 - y_j \right|. \quad (1.4)$$

Existing methods for solving (1.2)-(1.4) can be categorized into convex optimization type and gradient descent type. The former is based on Shor's convex relaxation [3]. It relaxes (1.2) to a convex minimization problem and solves this convex problem via semidefinite programming (SDP) [8, 7, 30]:

$$\begin{aligned} & \underset{\mathbf{B}}{\text{minimize}} && \text{trace}(\mathbf{B}), \\ & \text{subject to} && \mathbf{B} \succeq 0, \\ & && y_j = \text{trace}(\mathbf{x}_j \mathbf{x}_j^* \mathbf{B}), \quad j = 1, \dots, n, \end{aligned}$$

where $\mathbf{B} = \mathbf{b} \mathbf{b}^*$. Under noiseless Gaussian designs, SDP achieves exact recovery with sample size $O(p)$ [7]. Later, a modified version of SDP is proposed with the trace norm replaced by a reweighted trace norm, which is equivalent to minimizing a log-det function [5, 12]. Despite its reasonable performance and theoretical guarantees, SDP is computationally expensive because it optimizes over p^2 variables.

The second category are various types of gradient descent methods. The "Wirtinger flow" algorithm [6] targets the objective function (1.2). It obtains the starting point by spectral initialization, and refines the iterates via Wirtinger derivatives. A sample size $O(p \log p)$ is claimed to guarantee reasonable accuracy. The "truncated Wirtinger flow" [9] eliminates those abnormal data points

generated during the process to obtain a more reliable starting point as well as control the search direction. It exhibits more stable performance than the plain Wirtinger flow algorithm while advances the sample complexity to $O(p)$. The “truncated amplitude flow” algorithm [31] is targeting the objective function (1.3). It is also a two-stage procedure with orthogonality-promoting initialization followed by regularized gradient descent [31]. During its gradient descent stage, the signs of the components $\mathbf{x}_j^* \mathbf{b}$ are scrutinized to ensure a correct search direction. In a recent paper [11], the objective function (1.4) is transformed into the composition of a convex function and a smooth function, and the smooth function is further approximated by a linear function. The resulting objective function is convex and amenable to gradient descent. This method has slightly broader applications than other methods because it works on certain complex design vectors besides Gaussian designs. We do not give a complete bibliography here due to the vast amount of literature on this topic.

In many real-world applications, the signal has few nonzero coordinates, and far less measurements than the dimension of the signal are available. Phase retrieval in such context is referred to as sparse phase retrieval. The community has showed extensive interest on sparse phase retrieval during the past two decades [10]. Many of the algorithms for sparse phase retrieval are obtained by modifying the existing algorithms for non-sparse case. For example, certain norm regularization is added to the trace function in SDPs to promote sparsity [25, 22, 26], and a thresholding step is incorporated into each iteration of the gradient-descent-type algorithms [4, 32]. These modifications do not work for Fourier phase retrieval due to the ambiguity of translation and conjugate reflection. A novel method is proposed in [18] and demonstrated good performance in sparse Fourier phase retrieval. It first estimates the support via autocorrelation functions, and then solves an SDP over the support. Another recent algorithm is based on greedy local search [27]. It updates the signal support by interchanging the coordinate on support with the smallest gradient value with the coordinate off support with the largest gradient value. The objective function is also updated accordingly in each iteration. There are some other established methods for (sparse) phase retrieval [19, 16, 17], which we will not elaborate here to avoid unnecessary details.

Despite such intensive study on the algorithms for solving phase retrieval and sparse phase retrieval, statistical inferences about the signal is rarely touched. All the foregoing methods merely generate a point estimator for $\boldsymbol{\beta}$ and establish its convergence rate. No statistical inferences can be drawn on $\boldsymbol{\beta}$ or a function of $\boldsymbol{\beta}$ based on these point estimators. While in many real life applications, statistical inferences on the sparse signal are very much desired. For example, researchers might seek the 95% confidence interval of a certain coordinate β_k in order to adjust the receiver bandwidth. One major obstacle that thwarts statistical inferences on sparse phase retrieval is that the estimators generated by these algorithms cannot be written as an explicit function of the data. Thus, its sampling distribution or asymptotic distribution is in general not tractable.

In this paper, we propose a general method to construct confidence intervals for some simple function of the signals $\theta(\boldsymbol{\beta})$, *e.g.*, β_k . We also show that

the resulting confidence interval approximately attains the preassigned coverage probability when the sample size satisfies $n \gg (\log p)^2/s$ with s being the sparsity level of the true signal. Suppose we have an estimator for $\theta(\boldsymbol{\beta})$ available, which is asymptotically normal with mean $\theta(\boldsymbol{\beta})$. Then confidence intervals can be built based on this estimator and its asymptotic normality. Therefore, the key is to construct such an estimator, which we shall obtain as follows. First we pick an initial estimator output from some existing phase retrieval algorithm; second we modify the initial estimator in a way that the resulting estimator possesses all desired properties (asymptotically normal and centered at the true $\theta(\boldsymbol{\beta})$). The choice for the initial estimator will be discussed in Section 2. This method is inspired by the “debiased LASSO” in the literature of high-dimensional linear regressions [34, 29, 20, 2]. The LASSO is a shrinkage/thresholded estimator and hence biased. By adding a bias-correction term to the initial LASSO estimator, the authors obtain an asymptotically unbiased and normally distributed estimator, the debiased LASSO. Similarly in the case of sparse phase retrieval, existing algorithms always generate biased estimators because they are designed to promote sparsity. We will try to reduce the bias of the chosen initial estimator by adding a correction term that is distinct from the one in debiased LASSO. This correction term is the key to our method, and will be derived in Section 2. To reiterate the goal of this paper, given the measurements \mathbf{y} and the design matrix \mathbf{X} , we aim to construct confidence intervals with approximately preassigned coverage probabilities for a one-dimensional parameter, *e.g.*, $\theta = \theta(\boldsymbol{\beta}) = \beta_k$, a particular coordinate of the signal. To the best of our knowledge, this is the first study on statistical inference in the realm of sparse phase retrieval.

We organize the rest of this paper as follows. Section 2 elaborates our methodology and explains the rationale behind it. Section 3 presents the main theoretical guarantees for our method. Section 4 displays its empirical performance. Further discussions and perspectives are left to Section 5. And we leave all the proofs and technicality to Section 6.

Notation: Throughout the rest of the paper, we carry out the discussion based on model (1.1) and objective function (1.2). Wherever $f(\cdot)$ appears, it refers to the function in (1.2). However, the idea extends naturally to other phase retrieval models and objective functions. For conciseness, we just focus on the real case and consider model (1.1) with $\boldsymbol{\beta} \in \mathbb{R}^p$, $\|\boldsymbol{\beta}\|_0 \leq s$, $\mathbf{x}_1, \dots, \mathbf{x}_n$ *i.i.d.* $N(\mathbf{0}, I_p)$, and $\varepsilon_1, \dots, \varepsilon_n$ *i.i.d.* Gaussian noise $N(0, \sigma^2)$.

2. Methodology

This work is motivated by the celebrated debiasing techniques [33, 34, 29, 20, 2] for the high-dimensional linear regression problem and the emerging algorithms targeting sparse phase retrieval. Provided an initial estimator for β_k , which is biased and whose sampling distribution is intractable, we modify it so that the resulting estimator has approximately normal distribution centered at β_k in asymptotics. And confidence intervals can be constructed based on this asymptotic normality. Although we restrict our discussion to the sparse phase retrieval problem, the method is applicable to any M-estimating problem where

1. the Hessian matrix exists in a big enough neighborhood of the global maximizer and is invertible;
2. $\theta(\boldsymbol{\beta})$ is differentiable almost everywhere;
3. a good enough initial estimator is available.

2.1. Choice of the initial estimator

Thresholded Wirtinger Flow (TWF) is recently proposed for sparse phase retrieval [4]. It first generates a starting point by spectral initialization and then apply thresholded gradient descent to refine it. We choose the TWF output as our initial estimator because it is shown to achieve optimal minimax rate of convergence in the sparse phase retrieval setting [4]. More specifically, the authors proved that with high probability the TWF estimator of t -th iteration, $\tilde{\boldsymbol{\beta}}^{(t)}$, falls within a tiny ball centered at $\boldsymbol{\beta}$ (Theorem A.1). Written in mathematical formula,

$$\inf_{\|\boldsymbol{\beta}\|_0=s} \mathbb{P}(\mathbf{x}, \mathbf{y} | \boldsymbol{\beta}) \left\{ \min_{i=0,1} \|\tilde{\boldsymbol{\beta}}^{(t)} - (-1)^i \boldsymbol{\beta}\|_2 \leq \frac{1}{6} \left(1 - \frac{\mu}{16}\right)^t \|\boldsymbol{\beta}\|_2 + C \frac{\sigma}{\|\boldsymbol{\beta}\|_2} \sqrt{\frac{s \log p}{n}} \right\} > 1 - \frac{46}{n} - 10e^{-s} - \frac{t}{np^2} \tag{2.1}$$

for some absolute constant $C > 0$, provided the sample size satisfies $n \geq K(1 + \frac{\sigma}{\|\boldsymbol{\beta}\|_2})^2 s^2 \log(np)$ for some absolute constant $K > 0$ and the tuning parameters in the TWF algorithm are properly chosen. Here μ is the step size of gradient descent, which can be regarded as absolute constant once it is decided. If $t \asymp \log\left(\frac{\|\boldsymbol{\beta}\|_2^2 \sqrt{n}}{\sigma \sqrt{s \log p}}\right)$, one can obtain from the preceding result that with high probability

$$\min_{i=0,1} \|\tilde{\boldsymbol{\beta}}^{(t)} - (-1)^i \boldsymbol{\beta}\|_2 \lesssim \frac{\sigma}{\|\boldsymbol{\beta}\|_2} \sqrt{\frac{s \log p}{n}}. \tag{2.2}$$

This error rate is crucial for our bias-correction scheme to work, as will be revealed later in the proof part.

2.2. Bias-correction strategy

In this subsection, we will explain the bias-correction procedure in details. Suppose we are interested in a scalar parameter $\theta = \theta(\boldsymbol{\beta})$, which is a continuously differentiable function of $\boldsymbol{\beta}$. In our particular case, $\theta = \mathbf{e}_k^T \boldsymbol{\beta} = \beta_k$, $1 \leq k \leq p$. The TWF solution, denoted by $\tilde{\boldsymbol{\beta}}$, is biased due to the shrinkage nature of TWF. And so is $\theta(\tilde{\boldsymbol{\beta}})$. To correct the bias of $\theta(\tilde{\boldsymbol{\beta}})$, we adopt the idea of Low-Dimensional Projection Estimator (LDPE) proposed in [33]. The author considers a more general semi-low-dimensional (LD) approach where a high-dimensional (HD) model is decomposed as

$$\text{HD model} = \text{LD component} + \text{HD component}.$$

Such decomposition is to accommodate certain scenarios where people are only interested in some LD component instead of the whole model. The LD component can be one-dimensional, *e.g.*, a fixed direction in the model space. The HD component largely explains the model, but might not be related to or might be orthogonal to the LD component of interest. In such cases, it is natural to restrict the analysis on the LD component and leave aside the nuisance HD component.

In our problem, we have an HD model represented by the true parameter β . We are interested in a scalar parameter $\theta(\beta)$ instead of the whole β . Given a good initial estimator $\theta(\tilde{\beta})$ (computed based on $\tilde{\beta}$), we would search its vicinity for a $\hat{\theta}$ that is probably closer to the true θ . The key question is in which directions should we proceed to search for $\hat{\theta}$. Consider the one-dimensional sub-models $\{\tilde{\beta} + d \cdot \mathbf{u}, d \in \mathbb{R}, |d| \text{ is small}\}$, which start at $\tilde{\beta}$ and vary only in one direction $\mathbf{u} \in \mathbb{R}^p$. The direction \mathbf{u} is chosen to be the one that minimizes the Fisher information for estimating θ , and it is named the least favorable direction [33]. We pick the least favorable direction in order to guarantee reasonable performance in the worst case. The formula for $\mathbf{u}(\tilde{\beta})$ is given by

$$\mathbf{u}(\tilde{\beta}) = \arg \min_{\mathbf{u}} \{\mathbf{u}^\top \mathcal{I}(\tilde{\beta}) \mathbf{u} : \nabla \theta(\tilde{\beta})^T \mathbf{u}(\tilde{\beta}) = 1\} = \frac{\mathcal{I}(\tilde{\beta})^{-1} \nabla \theta(\tilde{\beta})}{\nabla \theta(\tilde{\beta})^T \mathcal{I}(\tilde{\beta})^{-1} \nabla \theta(\tilde{\beta})},$$

where

$$\mathcal{I}(\tilde{\beta}) = -\mathbb{E}_{\tilde{\beta}} \left[\frac{\partial^2 \ell(\mathbf{X}, \mathbf{y} | \mathbf{b})}{\partial \mathbf{b} \partial \mathbf{b}^T} \right]_{\mathbf{b}=\tilde{\beta}} = \frac{2}{\sigma^2} \sum_{j=1}^n 2(\mathbf{x}_j^\top \tilde{\beta})^2 \mathbf{x}_j \mathbf{x}_j^\top$$

is the Fisher information at $\tilde{\beta}$. We normalize \mathbf{u} with $\nabla \theta(\tilde{\beta})^T \mathbf{u}(\tilde{\beta}) = 1$ so that $\theta(\tilde{\beta} + d \cdot \mathbf{u}) - \theta(\tilde{\beta}) \approx d$. In view of the above, the decomposition for our specific problem can be written as

$$\beta - \tilde{\beta} = d \cdot \mathbf{u}(\tilde{\beta}) + \boldsymbol{\nu}, \quad (2.3)$$

where $d = \nabla \theta(\tilde{\beta})^\top (\beta - \tilde{\beta})$, and $\boldsymbol{\nu} = (\beta - \tilde{\beta}) - \nabla \theta(\tilde{\beta})^\top (\beta - \tilde{\beta}) \cdot \mathbf{u}(\tilde{\beta})$ is a nuisance component (with respect to θ) satisfying $\nabla \theta(\tilde{\beta})^\top \boldsymbol{\nu} = 0$. When both d and $\boldsymbol{\nu}$ are small, which is the case given $\beta \approx \tilde{\beta}$, the maximum likelihood estimator (MLE) in the least favorable direction is a natural candidate for d . An intuitive interpretation of (2.3) is as follows. The variation in β can be decomposed as the variation in direction \mathbf{u} and the variation in direction $\boldsymbol{\nu}$. The magnitude of the variation in \mathbf{u} equals that in θ . Variations in $\boldsymbol{\nu}$ would not change the value of θ . If β changes only in direction \mathbf{u} , it is difficult to detect the corresponding change in θ based on the likelihood function.

To simplify the notation, let $\tilde{\mathbf{u}} = \mathbf{u}(\tilde{\beta})$. The LDPE searches, in the least favorable direction $\tilde{\mathbf{u}}$, a parameter value that maximizes the likelihood of the occurring data sample, *i.e.*,

$$\hat{\theta} = \theta(\tilde{\beta}) + \arg \max_d \sum_{j=1}^n l(\mathbf{x}_j, y_j | \tilde{\beta} + \tilde{\mathbf{u}}d) = \theta(\tilde{\beta}) + \arg \min_d f(\tilde{\beta} + \tilde{\mathbf{u}}d). \quad (2.4)$$

The second equality above is because under model (1.1) and Gaussian noises,

$$\sum_{j=1}^n \ell(\mathbf{x}_j, y_j | \mathbf{b}) \propto -\frac{1}{2\sigma^2} \sum_{j=1}^n (y_j - |\mathbf{x}_j^T \mathbf{b}|^2)^2 = -\frac{2n}{\sigma^2} f(\mathbf{b}).$$

The score equation for the minimization problem in (2.4) is

$$\tilde{\mathbf{u}}^T \nabla f(\tilde{\boldsymbol{\beta}} + \tilde{\mathbf{u}}\hat{d}) = 0,$$

where $\hat{d} = \arg \min_d f(\tilde{\boldsymbol{\beta}} + \tilde{\mathbf{u}}d)$. By Taylor expansion,

$$\begin{aligned} \nabla f(\tilde{\boldsymbol{\beta}} + \tilde{\mathbf{u}}\hat{d}) - \nabla f(\tilde{\boldsymbol{\beta}}) &\approx \hat{d} \left[\frac{\partial^2 f(\mathbf{b})}{\partial \mathbf{b} \partial \mathbf{b}^T} \right]_{\mathbf{b}=\tilde{\boldsymbol{\beta}}} \tilde{\mathbf{u}} \\ &\approx \frac{\sigma^2}{2n} \hat{d} \hat{\mathcal{I}}(\tilde{\boldsymbol{\beta}}) \tilde{\mathbf{u}} \\ &= \frac{\sigma^2}{2n} \frac{\hat{d} \nabla \theta(\tilde{\boldsymbol{\beta}})}{\nabla \theta(\tilde{\boldsymbol{\beta}})^T \mathcal{I}(\tilde{\boldsymbol{\beta}})^{-1} \nabla \theta(\tilde{\boldsymbol{\beta}})}. \end{aligned}$$

Thus we have

$$0 = \tilde{\mathbf{u}}^T \nabla f(\tilde{\boldsymbol{\beta}} + \tilde{\mathbf{u}}\hat{d}) \approx \tilde{\mathbf{u}}^T \nabla f(\tilde{\boldsymbol{\beta}}) + \frac{\sigma^2}{2n} \frac{\hat{d}}{\nabla \theta(\tilde{\boldsymbol{\beta}})^T \mathcal{I}(\tilde{\boldsymbol{\beta}})^{-1} \nabla \theta(\tilde{\boldsymbol{\beta}})},$$

which solves

$$\hat{d} \approx -\frac{2}{\sigma^2} \nabla \theta(\tilde{\boldsymbol{\beta}})^T [\mathcal{I}(\tilde{\boldsymbol{\beta}})/n]^{-1} \nabla f(\tilde{\boldsymbol{\beta}}). \quad (2.5)$$

Plug (2.5) into (2.4), our corrected estimator is given by

$$\hat{\boldsymbol{\theta}} = \theta(\tilde{\boldsymbol{\beta}}) - \frac{2}{\sigma^2} \nabla \theta(\tilde{\boldsymbol{\beta}})^T [\mathcal{I}(\tilde{\boldsymbol{\beta}})/n]^{-1} \nabla f(\tilde{\boldsymbol{\beta}}). \quad (2.6)$$

Later we will implement the bias-corrected estimator (2.6) to the case where $\theta(\boldsymbol{\beta}) = \mathbf{e}_k^T \boldsymbol{\beta} = \beta_k$, and demonstrate some nice properties of this estimator. For consistency, hereinafter we will denote the corrected estimator as $\hat{\boldsymbol{\beta}}$ and the initial estimator as $\tilde{\boldsymbol{\beta}}$.

Before elaborating the analysis, let us acknowledge one fact. The exact sign of an individual β_k cannot be recovered because the measurements only provide magnitude information ($|\mathbf{x}_j^T \boldsymbol{\beta}|^2$ contaminated with noise). If $\tilde{\boldsymbol{\beta}}$ is a solution to the minimization problem (1.2), $-\tilde{\boldsymbol{\beta}}$ is also a solution. The signal $\boldsymbol{\beta}$ can be recovered only up to a global sign. Given an initial estimator $\tilde{\boldsymbol{\beta}}$, we define $\boldsymbol{\beta}^*$ to be whichever in $\{\boldsymbol{\beta}, -\boldsymbol{\beta}\}$ is closer to $\tilde{\boldsymbol{\beta}}$.

$$\boldsymbol{\beta}^* := \begin{cases} \boldsymbol{\beta} & \text{if } \|\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}}\|_2^2 \leq \|-\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}}\|_2^2, \\ -\boldsymbol{\beta} & \text{otherwise.} \end{cases}$$

We are only interested in gaining information about $\boldsymbol{\beta}^*$. This is the best we can do under model (1.1) in a sense that the global sign can never be retrieved.

To prepare the proofs in Section 6, we derive formulae to simplify a few notations here. The average Fisher information matrices at β^* and $\tilde{\beta}$ can be approximated by

$$\begin{aligned}\mathcal{I}(\beta^*)/n &\approx \frac{2}{\sigma^2} \mathbb{E}_{\mathbf{X}, \epsilon} [2(\mathbf{x}_j^T \beta^*)^2 \mathbf{x}_j \mathbf{x}_j^T] = \frac{4}{\sigma^2} (\|\beta^*\|_2^2 I + 2\beta^* \beta^{*T}), \\ \mathcal{I}(\tilde{\beta})/n &\approx \frac{2}{\sigma^2} \mathbb{E}_{\mathbf{X}, \epsilon} [2(\mathbf{x}_j^T \tilde{\beta})^2 \mathbf{x}_j \mathbf{x}_j^T] = \frac{4}{\sigma^2} (\|\tilde{\beta}\|_2^2 I + 2\tilde{\beta} \tilde{\beta}^T).\end{aligned}\quad (2.7)$$

In our specific context, $\theta(\beta) = \mathbf{e}_k^T \beta$ and $\nabla \theta(\beta) = \mathbf{e}_k$. If we define

$$\begin{aligned}\mathbf{w}_k &= -\frac{2}{\sigma^2} [\mathcal{I}(\beta)/n]^{-1} \nabla \theta(\beta) = -\frac{1}{2} (\|\beta^*\|_2^2 I + 2\beta^* \beta^{*T})^{-1} \mathbf{e}_k \\ &= -\frac{1}{2\|\beta^*\|_2^2} \left(I - \frac{2\beta^* \beta^{*T}}{3\|\beta^*\|_2^2} \right) \mathbf{e}_k\end{aligned}\quad (2.8)$$

and

$$\begin{aligned}\tilde{\mathbf{w}}_k &= -\frac{2}{\sigma^2} [\mathcal{I}(\tilde{\beta})/n]^{-1} \nabla \theta(\tilde{\beta}) = -\frac{1}{2} (\|\tilde{\beta}\|_2^2 I + 2\tilde{\beta} \tilde{\beta}^T)^{-1} \mathbf{e}_k \\ &= -\frac{1}{2\|\tilde{\beta}\|_2^2} \left(I - \frac{2\tilde{\beta} \tilde{\beta}^T}{3\|\tilde{\beta}\|_2^2} \right) \mathbf{e}_k,\end{aligned}\quad (2.9)$$

the estimator (2.6) can adopt a simpler formula

$$\hat{\beta}_k = \tilde{\beta}_k + \tilde{\mathbf{w}}_k^T \nabla f(\tilde{\beta}). \quad (2.10)$$

Remark 2.1. Equation (2.7) is derived under the assumption that $\tilde{\beta}$ is independent from the design vectors $\{\mathbf{x}_j\}_{j=1, \dots, n}$. This assumption is realistic in certain specially designed procedures. For example, we can randomly split the i.i.d. data $(\mathbf{x}_j, y_j)_{j=1, \dots, n}$ into two parts. We use the first part to obtain $\tilde{\beta}$ and the second part to construct the bias-correction term. The detailed procedure will be addressed in Section 2.3.

2.3. Data split and swap

In the analysis so far, we implicitly assumed that $\tilde{\beta}$ is independent from the design vectors \mathbf{x}_j 's and the noises ϵ_j 's, which especially ease the derivation of (2.7). However, this assumption fails if $\tilde{\beta}$ is obtained from the same set of design vectors \mathbf{x}_j 's and measurements y_j 's by TWF iterations. To ensure the independence between $\tilde{\beta}$ and $(\{\mathbf{x}_j\}_j, \{\epsilon_j\}_j)$, we design the three-step procedure. First, splitting the whole data set randomly into two parts, using the first part of data $(\mathbf{X}_1, \mathbf{y}_1)$ to obtain $\tilde{\beta}_1$ via TWF and using the second part of data $(\mathbf{X}_2, \mathbf{y}_2)$ to debias $\tilde{\beta}_1$; second, swapping the two parts of data, using the second part to generate $\tilde{\beta}_2$ and the first part to debias $\tilde{\beta}_2$; finally, combining the two debiased estimators. Theoretically, the required sample size to obtain a good initial estimator $\tilde{\beta}$ is different from the required sample size for efficient bias-correction. But we set the two parts to be of equal size n since the two parts will be swapped for both purposes.

3. Theoretical result

In this section, we will establish the asymptotic normality of the bias-corrected estimator (2.10). And the confidence interval based on this asymptotic distribution is shown to achieve roughly the preassigned coverage probability. Furthermore, theoretical analysis reveals that the confidence interval has sharp width and attains approximately the coverage probability simultaneously for all β_k . Note we assume σ is known because in many applications people have *a priori* knowledge about the noise level. Even if the noise level is unknown, there are efficient ways to estimate it, for example, the method proposed in [4]. In the following, Theorem 3.1 and Theorem 3.2 present some theoretical guarantees for the data-splitting scheme and the data-swapping scheme, respectively.

Before stating the theorems, we introduce some *global assumptions*. Suppose the design matrix and the true signal satisfy (3.1)-(3.3).

$$n \geq K \left(1 + \frac{\sigma}{\|\beta^*\|_2}\right)^2 s^2 \log(np), \quad K \text{ is an absolute constant.} \quad (3.1)$$

$$\frac{\log p}{\sqrt{ns}} = o(1). \quad (3.2)$$

$$\|\beta^*\|_2 = O(\sqrt{s}). \quad (3.3)$$

If $s \ll \log p$, assumptions (3.2) and (3.3) imply assumption (3.1). However, we do not impose here any restrictions on s . Our method could work in cases where the signal sparsity is not strong, *i.e.*, there could be many non-zero small coordinates.

Throughout the rest of the paper, let

$$\epsilon'_n = C_1 \frac{(\log p)^{\frac{3}{2}}}{sn} + C_2 \frac{(\log p)^2}{(sn)^{\frac{3}{2}}} + C_3 \frac{\log p}{(sn)^{\frac{1}{2}}} + C_4 \frac{n^{\frac{1}{2}}}{p^2} \quad (3.4)$$

for some absolute constants C_1, C_2, C_3, C_4 , and let

$$\epsilon''_n = \frac{50}{n} + 10e^{-s} + \frac{M \log n}{np^2} \quad (3.5)$$

for some absolute constant M . It is easy to see that $\epsilon'_n, \epsilon''_n \rightarrow 0$ under assumption (3.2).

Theorem 3.1 essentially states that under global assumptions (3.1)-(3.2) the estimator $\hat{\beta}_k$ generated by the data-splitting scheme is approximately asymptotic normal centered at the true parameter. The *i.i.d.* data generated from model (1.1) are randomly split into two halves, $(\mathbf{X}_1, \mathbf{y}_1)$ and $(\mathbf{X}_2, \mathbf{y}_2)$. Each half has a sample size n . We use $(\mathbf{X}_1, \mathbf{y}_1)$ to obtain the initial estimator and $(\mathbf{X}_2, \mathbf{y}_2)$ to correct the bias.

Theorem 3.1. *Under the data-splitting scheme and assumption (3.3), if the initial estimator $\tilde{\beta}$ satisfies the two conditions in (3.6) with probability at least*

$1 - \epsilon_0$ for some $\epsilon_0 = o(1)$,

$$\text{supp}(\tilde{\boldsymbol{\beta}}) \subset S := \text{supp}(\boldsymbol{\beta}^*), \quad \|\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 \leq \frac{C_0 \sigma}{\|\boldsymbol{\beta}^*\|_2} \sqrt{\frac{s \log p}{n}}, \quad (3.6)$$

we have

$$\mathbb{P} \left\{ \left| \sqrt{n}(\hat{\beta}_k - \beta_k^*) - Z_k \right| \leq \epsilon'_n \right\} \geq 1 - \epsilon_0 - \frac{4}{n},$$

where

$$Z_k = -\frac{1}{\sqrt{n}} \sum_{j=1}^n \varepsilon_j (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k), \quad (3.7)$$

and $\hat{\beta}_k$, $\tilde{\mathbf{w}}_k$, ϵ'_n are given by (2.10), (2.9), (3.4), respectively.

Remark 3.1. In particular, it has been established in [4] that under assumption (3.1) and with iteration number $t \asymp \log \left(\frac{\|\boldsymbol{\beta}^*\|_2^2 \sqrt{n}}{\sigma \sqrt{s \log p}} \right)$, the TWF estimator satisfies (3.6) with $\epsilon_0 = \frac{46}{n} + 10e^{-s} + \frac{M \log n}{np^2}$. If we choose the TWF estimator as $\tilde{\boldsymbol{\beta}}$, under assumptions (3.1) and (3.3), $\hat{\beta}_k$ would achieve

$$\mathbb{P} \left\{ \left| \sqrt{n}(\hat{\beta}_k - \beta_k^*) - Z_k \right| \leq \epsilon'_n \right\} \geq 1 - \epsilon''_n. \quad (3.8)$$

The assumption (3.2) further implies that both ϵ'_n and ϵ''_n go to zero as $n \rightarrow \infty$. Therefore, $\hat{\beta}_k$ is approximately normal centered at β_k^* in asymptotics. It is clear that $\tilde{\boldsymbol{\beta}}$ is independent from \mathbf{x}_j 's in our data-splitting regime. Z_k has limiting distribution $N(0, \sigma^2 \tau_k^2)$, where

$$\begin{aligned} \tau_k^2 &= \|\tilde{\boldsymbol{\beta}}\|_2^2 \|\tilde{\mathbf{w}}_k\|_2^2 + 2(\tilde{\boldsymbol{\beta}}^T \tilde{\mathbf{w}}_k)^2 = \|\tilde{\boldsymbol{\beta}}\|_2^2 \|\tilde{\mathbf{w}}_k\|_2^2 + 2 \left(-\frac{\tilde{\beta}_k}{2\|\tilde{\boldsymbol{\beta}}\|_2^2} + \frac{\tilde{\beta}_k}{3\|\tilde{\boldsymbol{\beta}}\|_2^2} \right)^2 \\ &= \|\tilde{\boldsymbol{\beta}}\|_2^2 \|\tilde{\mathbf{w}}_k\|_2^2 + \frac{\tilde{\beta}_k^2}{18\|\tilde{\boldsymbol{\beta}}\|_2^4} \\ &\leq \left(\|\boldsymbol{\beta}^*\|_2 + \|\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 \right)^2 \left[\frac{5}{6\|\boldsymbol{\beta}^*\|_2^2} + O\left(\frac{\|\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2}{\|\boldsymbol{\beta}^*\|_2^3} \right) \right]^2 + \frac{1}{18\|\tilde{\boldsymbol{\beta}}\|_2^2} \\ &\asymp \frac{3}{4s} + O\left(\sqrt{\frac{\log p}{ns^3}} \right) \end{aligned}$$

by (3.3), (2.9), (2.2). Such sharp width is uniformly achievable for all coordinates, which allows the possibility of constructing simultaneous confidence intervals. More importantly, the bias term ϵ'_n is of smaller order compared with τ_k .

To fully extract information contained in the data, we further apply the data-swapping scheme and present the corresponding theoretical results in Theorem 3.2. After randomly splitting the data into two halves, we go through

two rounds of constructions. In the first round, we obtain the initial estimator $\hat{\beta}_1$ from $(\mathbf{X}_1, \mathbf{y}_1)$, and correct its bias using $(\mathbf{X}_2, \mathbf{y}_2)$, resulting in $\hat{\beta}_{1k} = \tilde{\beta}_{1k} + \tilde{\mathbf{w}}_{1k}^T \nabla f_2(\tilde{\beta}_1)$; in the second round, we obtain the initial estimator $\hat{\beta}_2$ from $(\mathbf{X}_2, \mathbf{y}_2)$ and correct its bias using $(\mathbf{X}_1, \mathbf{y}_1)$, resulting in $\hat{\beta}_{2k} = \tilde{\beta}_{2k} + \tilde{\mathbf{w}}_{2k}^T \nabla f_1(\tilde{\beta}_2)$. Finally we combine the two estimators in a way that the resulting estimator has the smallest asymptotic variance. Here

$$\begin{aligned} \tilde{\mathbf{w}}_{1k} &:= w(\tilde{\beta}_1, \theta = \beta_k) = -\frac{1}{2} \left(\|\tilde{\beta}_1\|_2^2 I + 2\tilde{\beta}_1 \tilde{\beta}_1^T \right)^{-1} \mathbf{e}_k, \\ \tilde{\mathbf{w}}_{2k} &:= w(\tilde{\beta}_2, \theta = \beta_k) = -\frac{1}{2} \left(\|\tilde{\beta}_2\|_2^2 I + 2\tilde{\beta}_2 \tilde{\beta}_2^T \right)^{-1} \mathbf{e}_k, \\ \nabla f_2(\tilde{\beta}_1) &= \frac{1}{n} \sum_{j=1}^n [(\mathbf{x}_{2j}^T \tilde{\beta}_1)^2 - y_{2j}] (\mathbf{x}_{2j}^T \tilde{\beta}_1) \mathbf{x}_{2j}, \\ \nabla f_1(\tilde{\beta}_2) &= \frac{1}{n} \sum_{j=1}^n [(\mathbf{x}_{1j}^T \tilde{\beta}_2)^2 - y_{1j}] (\mathbf{x}_{1j}^T \tilde{\beta}_2) \mathbf{x}_{1j}. \end{aligned}$$

Finally we linearly combine $\hat{\beta}_{1k}$ and $\hat{\beta}_{2k}$ in a way that the resulting estimator $\hat{\beta}_k^{swap}$ has the smallest asymptotic variance. $\hat{\beta}_k^{swap}$ contains more information of the data than $\hat{\beta}_{1k}$ or $\hat{\beta}_{2k}$ alone.

Theorem 3.2. (i) Under the global assumptions (3.1)-(3.3), the estimator

$$\hat{\beta}_k^{swap} = \frac{\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2} \cdot \left(\tilde{\beta}_{1k} + \tilde{\mathbf{w}}_{1k}^T \nabla f_2(\tilde{\beta}_1) \right) + \frac{\tau_{1k}^2}{\tau_{1k}^2 + \tau_{2k}^2} \cdot \left(\tilde{\beta}_{2k} + \tilde{\mathbf{w}}_{2k}^T \nabla f_1(\tilde{\beta}_2) \right) \tag{3.9}$$

achieves the smallest asymptotic variance among all asymptotically unbiased estimators that are convex combinations of $\hat{\beta}_{1k}$ and $\hat{\beta}_{2k}$. Furthermore,

$$\mathbb{P} \left\{ \left| \sqrt{n}(\hat{\beta}_k^{swap} - \beta_k^*) - \left(\frac{\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2} Z_{1k} + \frac{\tau_{1k}^2}{\tau_{1k}^2 + \tau_{2k}^2} Z_{2k} \right) \right| \leq \epsilon'_n \right\} \geq 1 - 2\epsilon''_n, \tag{3.10}$$

where

$$\begin{aligned} Z_{1k} &= \frac{1}{\sqrt{n}} \sum_{j=1}^n \varepsilon_{2j} (\mathbf{x}_{2j}^T \tilde{\beta}_1) (\mathbf{x}_{2j}^T \tilde{\mathbf{w}}_{1k}), \\ Z_{2k} &= \frac{1}{\sqrt{n}} \sum_{j=1}^n \varepsilon_{1j} (\mathbf{x}_{1j}^T \tilde{\beta}_2) (\mathbf{x}_{1j}^T \tilde{\mathbf{w}}_{2k}), \\ \tau_{1k}^2 &= \|\tilde{\beta}_1\|_2^2 \|\tilde{\mathbf{w}}_{1k}\|_2^2 + 2(\tilde{\beta}_1^T \tilde{\mathbf{w}}_{1k})^2, \\ \tau_{2k}^2 &= \|\tilde{\beta}_2\|_2^2 \|\tilde{\mathbf{w}}_{2k}\|_2^2 + 2(\tilde{\beta}_2^T \tilde{\mathbf{w}}_{2k})^2. \end{aligned}$$

(ii)

$$\mathbb{P}\left\{\left|\hat{\beta}_k^{swap} - \beta_k^*\right| \leq \frac{\sigma r}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}} + O\left(\frac{1}{p^2}\right) + \frac{\epsilon'_n}{\sqrt{n}}\right\} \geq 2\Phi(r) - 1 - 2\epsilon''_n - \frac{2}{n}. \tag{3.11}$$

(iii)

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left\{\max_{k \in [p]} \left|\sqrt{n}(\hat{\beta}_k^{swap} - \beta_k^*)\right| \leq \sqrt{\frac{3}{8s}} \sigma \Phi^{-1}\left(1 - \frac{\alpha}{2}\right)\right\} \geq 1 - \alpha. \tag{3.12}$$

Remark 3.2. Note, asymptotically, $\tau_{1k}^2, \tau_{2k}^2 \asymp \frac{3}{4s}$ for all $k = 1, \dots, p$ by Remark 3.1. Hence $\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2} \leq \frac{1}{2} \max\{\tau_{1k}^2, \tau_{2k}^2\} \leq \frac{3}{8s}$. The asymptotic variance of $\hat{\beta}_k^{swap}$ is shrunk by a factor of 2 compared to that of $\hat{\beta}_{1k}$ or $\hat{\beta}_{2k}$. The uniformly bounded variance for all $\sqrt{n}(\hat{\beta}_k^{swap} - \beta_k^*)$ in (3.12) allows Bonferroni adjustment to control familywise error rate in simultaneous interval estimation.

We are also able to establish the theoretical guarantee for constructing Scheffe’s simultaneous confidence intervals using $\hat{\beta}_k^{swap}$, which is stated in the following corollary.

Corollary 3.1. Under the global assumptions (3.1)-(3.3), for $\forall \mathbf{h} \neq 0, \mathbf{h} \in \mathbb{R}^p$, the estimator in (3.9) enjoys the property below.

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left\{\left|\mathbf{h}^T(\hat{\beta}^{swap} - \beta^*)\right| \leq \sqrt{\frac{\sigma^2}{n} \chi_{p,\alpha}^2 \mathbf{h}^T \mathbf{V} \mathbf{h}}\right\} \geq 1 - \alpha, \tag{3.13}$$

where

$$V_{kl} = a_k a_l \left[\|\tilde{\beta}_1\|_2^2 (\tilde{\mathbf{w}}_{1k}^T \tilde{\mathbf{w}}_{1l}) + 2(\tilde{\beta}_1^T \tilde{\mathbf{w}}_{1k})(\tilde{\beta}_1^T \tilde{\mathbf{w}}_{1l}) \right] + (1 - a_k)(1 - a_l) \left[\|\tilde{\beta}_2\|_2^2 (\tilde{\mathbf{w}}_{2k}^T \tilde{\mathbf{w}}_{2l}) + 2(\tilde{\beta}_2^T \tilde{\mathbf{w}}_{2k})(\tilde{\beta}_2^T \tilde{\mathbf{w}}_{2l}) \right]$$

with $a_k = \frac{\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}$ and $a_l = \frac{\tau_{2l}^2}{\tau_{1l}^2 + \tau_{2l}^2}$.

A key point to the success of $\hat{\beta}_k^{swap}$ is that its bias is of smaller order than the corresponding standard deviation. And hence $\hat{\beta}_k^{swap}$ would center at β_k^* with an approximately normal distribution in asymptotics. To avoid redundancy, we will not go through the case where the targeting parameter is a linear combination of the signal coordinates, $\theta(\beta) = \mathbf{a}^T \beta$. The construction procedure is almost identical to what aforementioned except that \mathbf{e}_k is replaced by $\mathbf{a} \in \mathbb{R}^P$.

4. Numerical simulation

In this section, we implement our method in a variety of settings to assess its empirical performance. Moreover, by comparing the simulation results in


different combinations of sparsity (s), sample size ($2n$), and noise-to-signal ratio (NSR) $\sigma^2/\|\beta\|_2^2$, we get a general idea about how on those factors affect the performance.

Throughout our simulation, we set the signal dimension $p = 1000$. All the tuning parameters in the TWF algorithm are chosen according to [4] and kept fixed. In each choice of (n, s, NSR), we generate the signal β by randomly picking the support and assigning *i.i.d.* $N(0, 1)$ to the coordinates on support. Among these nonzero coordinates, there are a few large coordinates ($\beta_k \approx 3$), some median coordinates ($\beta_k \approx 1$), and some small coordinates ($\beta_k \approx 0.1$). Later we will differentiate the performances of our method in cases of large, median, and small coordinates, respectively. Given this β , the following procedure is repeated independently for 100 times: first, generating $2n$ random vectors \mathbf{x}_j 's *i.i.d.* $N(0, I)$, $2n$ noises ε_j 's *i.i.d.* $N(0, \sigma^2)$, and the measurements \mathbf{y} by (1.1); second, obtaining the TWF estimator using the whole dataset (\mathbf{X}, \mathbf{y}) and record the errors $\tilde{\beta}_k - \beta_k^*$ of four large coordinates ($|\beta_k^*| \approx 3$), four median coordinates ($|\beta_k^*| \approx 1$), and four small coordinates ($|\beta_k^*| \approx 0.1$); third, implementing the data-swap scheme, obtaining the debiased TWF $\hat{\beta}^{swap}$, and recording the errors $\hat{\beta}_k^{swap} - \beta_k^*$ of the four large/median/small coordinates, respectively. Therefore, every summary statistic in Table 1 comes from a pool of 400 errors (100 times \times 4 coordinates) and each histogram in Figure 1, Figure 2, Figure 3 represents the distribution of a pool of 400 errors. While the average coverage probability in Table 2 is calculated based on all coordinates.


The performance of our method is assessed in terms of several aspects, including biasness (Table 1), variance (Table 1), asymptotic normality (Figure 1, Figure 2, Figure 3), and coverage probability (Table 2).

Judging from Table 1 (“mae” stands for “median absolute error”), debiased TWF achieves close-to-zero average bias at the cost of slightly larger variance than TWF in all settings except when sparsity s is too large ($s=200$) or sample size $2n$ is too small ($n=2000$). The reason behind such phenomenon could be as follows. The bias-correction term with non-zero mean is supposed to neutralize the bias of TWF. Though it brings extra variance, the amount of this extra variance is negligible as long as n is large enough. Nevertheless, when n is small or s is large, the bias-correction term does not concentrate tightly around its mean. Instead of neutralizing the bias of TWF, it adds extra bias and much larger variance. Table 1 also exhibits a parallel trend between TWF and debiased TWF that their bias and variance get larger as s/NSR increases or n decreases. This is because the performance of debiased TWF depends on the quality of TWF while TWF gets worse as s/NSR increases or n decreases, which has been demonstrated in the original paper [4]. Although there are two abnormal cases ($s=200$ and $n=2000$) in Table 1 where the debiased TWF fails, our main theory (Theorem 3.2) is not violated because the (s, n, p) in these cases are largely deviated from assumption (3.1) and (3.2).

Figure 1, Figure 2, and Figure 3 demonstrate the unbiasedness and approximate normality of the debiased TWF in all settings. Figure 1 explores the relationship between the sparsity s and the quality of debiased TWF by fix-




(a) Histograms of TWF errors for large β_k (1st row), median β_k (2nd row), small β_k (3rd row).




(b) Histograms of debiased TWF errors for large β_k (4th row), median β_k (5th row), small β_k (6th row).

FIG 1. From left to right, plots correspond to simulation settings $(n=3000, s=50, NSR=0.3)$, $(n=3000, s=100, NSR=0.3)$, $(n=3000, s=150, NSR=0.3)$, $(n=3000, s=200, NSR=0.3)$ with $p=1000$ fixed.




(a) Histograms of TWF errors for large β_k (1st row), median β_k (2nd row), small β_k (3rd row).




(b) Histograms of debiased TWF errors for large β_k (4th row), median β_k (5th row), small β_k (6th row).

FIG 2. From left to right, plots correspond to simulation settings $(n=2000, s=100, NSR=0.3)$, $(n=3000, s=100, NSR=0.3)$, $(n=4000, s=100, NSR=0.3)$, $(n=5000, s=100, NSR=0.3)$ with $p=1000$ fixed.



(a) Histograms of TWF errors for large β_k (1st row), median β_k (2nd row), small β_k (3rd row).



(b) Histograms of debiased TWF errors for large β_k (4th row), median β_k (5th row), small β_k (6th row).

FIG 3. From left to right, plots correspond to simulation settings ($n=3000, s=100, NSR=0.2$), ($n=3000, s=100, NSR=0.3$), ($n=3000, s=100, NSR=0.4$), ($n=3000, s=100, NSR=0.5$) with $p=1000$ fixed.

TABLE 1
 Summary statistics of TWF errors and debiased TWF errors for large β_k , median β_k , and small β_k under various simulation settings.

n	s	NSR		large coordinates		median coordinates		small coordinates	
				TWF	de-TWF	TWF	de-TWF	TWF	de-TWF
3000	50	0.3	bias	0.0108	0.0016	0.0421	0.0017	-0.0607	0.0020
			sd	0.0174	0.0203	0.0186	0.0216	0.0188	0.0226
			mae	0.0144	0.0135	0.0415	0.0155	0.0604	0.0151
3000	100	0.3	bias	0.0367	-0.0061	0.0573	-0.0054	-0.0872	0.0054
			sd	0.0295	0.0838	0.0291	0.0540	0.0173	0.0513
			mae	0.0375	0.0316	0.0567	0.0311	0.0987	0.0353
3000	150	0.3	bias	-0.0370	-0.0359	-0.0928	0.0475	0.0975	0.0025
			sd	0.0483	0.1495	0.0498	0.1287	0.0084	0.1346
			mae	0.0442	0.0838	0.0921	0.0992	0.1000	0.0939
3000	200	0.3	bias	0.1007	0.2577	0.1197	-0.0474	-0.0989	0.0153
			sd	0.0679	0.3521	0.0681	0.2335	0.0068	0.2175
			mae	0.0976	0.2857	0.1209	0.1630	0.1000	0.1523
2000	100	0.3	bias	0.0481	-0.0271	0.0844	-0.0128	-0.0972	-0.0143
			sd	0.0457	0.4064	0.0475	0.2029	0.0091	0.1421
			mae	0.0503	0.1264	0.0840	0.1179	0.1000	0.0986
4000	100	0.3	bias	-0.0216	-0.0021	-0.0499	0.0027	0.0780	-0.0044
			sd	0.0247	0.0311	0.0252	0.0331	0.0205	0.0350
			mae	0.0255	0.0223	0.0498	0.0241	0.0798	0.0248
5000	100	0.3	bias	0.0288	0.0018	0.0362	-0.0017	-0.0625	0.0037
			sd	0.0193	0.0227	0.0192	0.0233	0.0190	0.0236
			mae	0.0282	0.0156	0.0362	0.0154	0.0624	0.0159
3000	100	0.2	bias	-0.0229	0.0078	-0.0410	-0.0014	0.0662	-0.0050
			sd	0.0196	0.0848	0.0210	0.0521	0.0209	0.0380
			mae	0.0242	0.0221	0.0408	0.0251	0.0678	0.0243
3000	100	0.4	bias	-0.0346	-0.0083	-0.0827	0.0140	0.0953	-0.0081
			sd	0.0425	0.0594	0.0418	0.0585	0.0137	0.0634
			mae	0.0380	0.0420	0.0841	0.0425	0.1000	0.0412
3000	100	0.5	bias	-0.0462	-0.0076	-0.0980	0.0141	0.0972	-0.0082
			sd	0.0480	0.0630	0.0507	0.0685	0.0097	0.0707
			mae	0.0494	0.0445	0.0965	0.0438	0.1000	0.0487

ing $p=1000$, $n=3000$, $NSR=0.3$ while varying $s=50, 100, 150, 200$ from left to right. The normality and unbiasedness hold better in small s . Besides, the distributions spread wider as s increases in every row. Such phenomenon is not surprising because the asymptotic unbiasedness and normality of $\hat{\beta}_k^{swap}$ rely on that $\frac{\sqrt{n}\|\Delta\|_2^2}{\|\beta^*\|_2}$ is asymptotically negligible. When s increases to an extent that the assumption is violated, the quality of $\tilde{\beta}$ drops and $\|\Delta\|_2 = \|\tilde{\beta} - \beta^*\|_2$ cannot be controlled. Similarly, the errors of debiased TWF distribute more Gaussian and spread narrower as n increases in Figure 2 or as NSR decreases in Figure 3.

When checking the accuracy of Theorem 3.2 (ii), we use the confidence interval $\left(\hat{\beta}_k^{swap} - \frac{\sigma\Phi^{-1}(0.98)}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}}, \hat{\beta}_k^{swap} + \frac{\sigma\Phi^{-1}(0.98)}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}} \right)$. For each setting, a new β is generated, and we repeat constructing the confidence interval for 200 times with independently generated $(\mathbf{X}, \varepsilon, \mathbf{y})$. Note, the theoretic coverage probability is $2 \times 0.98 - 1 = \frac{102}{n}$, in which we count the term ϵ_n'' . In asymptotics the term ϵ_n'' is negligible, yet in our simulation $\frac{102}{n} \in (1.36\%, 2\%)$. Therefore, the

TABLE 2

Mean coverage probabilities of the confidence interval in Theorem 3.2 under various settings.

n	s	σ	all coor	large coor	median coor	small coor
5000	40	5	92.86	92.50	94.75	93.25
6000	40	5	93.29	93.63	92.63	94.38
6000	50	5	92.78	93.00	92.63	94.25
7500	50	5	93.42	93.88	92.88	93.25
5000	40	10	92.91	92.38	94.13	92.38
6000	40	10	93.66	92.75	94.75	93.25
6000	50	10	92.59	92.25	92.25	91.63
7500	50	10	93.42	93.88	93.88	95.13

theoretic coverage probabilities are ranging from 94% to 94.64% depending on the variation of n . Table 2 shows that most average coverage probabilities are slightly below their theoretical values. The reason for this phenomenon lies in the extra term $\frac{\epsilon'_n}{\sqrt{n}}$ in Theorem 3.2 (ii). Again, in asymptotics $\frac{\epsilon'_n}{\sqrt{n}}$ is negligible compared to $\frac{\sigma r}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}}$, but in our simulation these two terms are of the same order. Thus, leaving the term $\frac{\epsilon'_n}{\sqrt{n}}$ out has diminished the coverage probability to certain extent. Despite this flaw, the results in Table 2 imply that Theorem 3.2 (ii) is informative of the actual performance of the debiased TWF. We can see the general trends: the average coverage probabilities get better as n increases or s decreases, and not sensitive to σ .

5. Discussion

We propose in this work a general approach for drawing statistical inferences on the sparse signal in phase retrieval. A new estimator $\hat{\beta}_k^{swap}$ for the individual signal coordinate β_k has been constructed by adding a bias-correction term to the TWF estimator. Under mild assumptions on \mathbf{X} and β and sample size requirement (3.2), $\hat{\beta}_k^{swap}$ has asymptotic Gaussian distribution centered at β_k^* . This property allows construction of confidence intervals with approximately preassigned probabilities as well as hypothesis testing on the signal of interest. Our new estimator can also be used as a point estimator. Compared with the plain TWF estimator, $\hat{\beta}_k^{swap}$ achieves asymptotic unbiasedness at the cost of a slightly larger variance.

There remain some open problems. For instance, can we draw statistical inferences on more complicated functions of β , such as a group of coordinates (similar to that in [23]) or a multidimensional-valued function? Can this method be extended to Fourier designs, which are more applicable? The former is difficult since it involves non-convex optimization over matrices. The later is even more challenging since the Fourier phase retrieval problem is generally considered not solved. And hence we do not have an initial estimator available yet. In summary, there is still long way to go before the Fourier phase retrieval problem is solved and related statistical inferences can be drawn.

6. Proofs

PROOF OF THEOREM 3.1.

For the objective function (1.2), we have

$$\nabla f(\tilde{\boldsymbol{\beta}}) = \frac{1}{n} \sum_{j=1}^n [(\mathbf{x}_j^T \tilde{\boldsymbol{\beta}})^2 - y_j] (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) \mathbf{x}_j.$$

Let $\Delta = \tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$, and $\hat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^*$ can be decomposed as

$$\begin{aligned} \hat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^* &= \tilde{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^* + \tilde{\mathbf{w}}_k^T \nabla f(\tilde{\boldsymbol{\beta}}) \\ &= \mathbf{e}_k^T \Delta + \frac{1}{n} \sum_{j=1}^n [(\mathbf{x}_j^T \tilde{\boldsymbol{\beta}})^2 - y_j] (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) \\ &= \mathbf{e}_k^T \Delta + \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^3 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) + \frac{3}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^2 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*) \\ &\quad + \frac{2}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 - \frac{1}{n} \sum_{j=1}^n \varepsilon_j (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) \\ &= \mathbf{e}_k^T \Delta + \frac{2}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) (\mathbf{x}_j^T \mathbf{w}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 + \frac{2}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) [\mathbf{x}_j^T (\tilde{\mathbf{w}}_k - \mathbf{w}_k)] (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \\ &\quad + \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^3 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) + \frac{3}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^2 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*) \\ &\quad - \frac{1}{n} \sum_{j=1}^n \varepsilon_j (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k). \end{aligned} \tag{6.1}$$

We will bound these terms separately.

By (3.6), we have $\text{supp}(\Delta) \subseteq S$ with high probability. Since

$$\begin{aligned} \mathbf{w}_k &= -\frac{1}{2} (\|\boldsymbol{\beta}^*\|_2^2 I + 2\boldsymbol{\beta}^* \boldsymbol{\beta}^{*T})^{-1} \mathbf{e}_k \\ &= -\frac{1}{2\|\boldsymbol{\beta}^*\|_2^2} \left(I - \frac{2}{3} \frac{\boldsymbol{\beta}^*}{\|\boldsymbol{\beta}^*\|_2} \frac{\boldsymbol{\beta}^{*T}}{\|\boldsymbol{\beta}^*\|_2} \right) \mathbf{e}_k \\ &= -\frac{1}{2\|\boldsymbol{\beta}^*\|_2^2} \mathbf{e}_k + \frac{\beta_k^*}{3\|\boldsymbol{\beta}^*\|_2^4} \boldsymbol{\beta}^*, \end{aligned}$$

$\text{supp}(\mathbf{w}_k) \subseteq S \cup \{k\}$. Similarly, $\text{supp}(\tilde{\mathbf{w}}_k) \subseteq S \cup \{k\}$. Let $\bar{S} = S \cup \{k\}$.

The fact that the supports of Δ , \mathbf{w}_k and $\tilde{\mathbf{w}}_k$ are inside \bar{S} grants the applicability of Lemma A.1 here, which implies

$$\left| \mathbf{w}_k^T \left(\frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \mathbf{x}_j \mathbf{x}_j^T - (\|\boldsymbol{\beta}^*\|_2^2 I + 2\boldsymbol{\beta}^* \boldsymbol{\beta}^{*T}) \right) \Delta \right|$$

$$= \left| \mathbf{w}_k^T \left(\frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \mathbf{x}_{j\bar{S}} \mathbf{x}_{j\bar{S}}^T - (\|\boldsymbol{\beta}^*\|_2^2 I_{\bar{S}} + 2\boldsymbol{\beta}^* \boldsymbol{\beta}^{*T}) \right) \Delta \right| \leq \delta \|\boldsymbol{\beta}^*\|_2^2 \|\Delta\|_2 \|\mathbf{w}_k\|_2 \quad (6.2)$$

with probability at least $1 - \frac{1}{n}$, provided $n \geq C(\delta)(s+1) \log(s+1)$. Here $C(\delta)$ is constant only depending on δ , and $I_{\bar{S}}$ is a diagonal matrix with the diagonal elements in \bar{S} equal to 1 and others equal to 0.

By (2.8) and simple algebra, we have

$$\mathbf{e}_k^T \Delta + 2\mathbf{w}_k^T (\|\boldsymbol{\beta}^*\|_2^2 I_{\bar{S}} + 2\boldsymbol{\beta}^* \boldsymbol{\beta}^{*T}) \Delta = 0.$$

Combining the above two formulae, the first two terms in (6.1) can be bounded by

$$\left| \mathbf{e}_k^T \Delta + \frac{2}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) (\mathbf{x}_j^T \mathbf{w}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \right| \leq \delta \|\boldsymbol{\beta}^*\|_2^2 \|\Delta\|_2 \|\mathbf{w}_k\|_2 \quad (6.3)$$

with high probability if $n \geq C(\delta)(s+1) \log(s+1)$.

By similar argument, we have

$$\left| \tilde{\mathbf{w}}_k^T \left(\frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^2 \mathbf{x}_j \mathbf{x}_j^T - (\|\Delta\|_2^2 I + 2\Delta \Delta^T) \right) \Delta \right| \leq \delta \|\Delta\|_2^3 \|\tilde{\mathbf{w}}_k\|_2$$

with probability at least $1 - \frac{1}{n}$ if $n \geq C(\delta)(s+1) \log(s+1)$. Thus, the fourth term in (6.1) is bounded by

$$\begin{aligned} \left| \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^3 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) \right| &\leq 3\|\Delta\|_2^2 |\Delta^T \tilde{\mathbf{w}}_k| + \delta \|\Delta\|_2^3 \|\tilde{\mathbf{w}}_k\|_2 \\ &\leq 3\|\Delta\|_2^3 \|\tilde{\mathbf{w}}_k\|_2 + \delta \|\Delta\|_2^3 \|\tilde{\mathbf{w}}_k\|_2 \end{aligned} \quad (6.4)$$

with high probability.

Again via similar reasoning, the fifth term in (6.1) falls within $3\|\Delta\|_2^2 (\tilde{\mathbf{w}}_k^T \boldsymbol{\beta}^*) + 6(\Delta^T \tilde{\mathbf{w}}_k)(\Delta^T \boldsymbol{\beta}^*) \pm 3\delta \|\Delta\|_2^2 \|\tilde{\mathbf{w}}_k\|_2 \|\boldsymbol{\beta}^*\|_2$ with probability at least $1 - \frac{1}{n}$ provided $n \geq C(\delta)(s+1) \log(s+1)$. And we bound it by

$$\begin{aligned} &\left| \frac{3}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta)^2 (\mathbf{x}_j^T \tilde{\mathbf{w}}_k) (\mathbf{x}_j^T \boldsymbol{\beta}^*) \right| \\ &\leq 3\|\Delta\|_2^2 |\tilde{\mathbf{w}}_k^T \boldsymbol{\beta}^*| + 6|\Delta^T \tilde{\mathbf{w}}_k| |\Delta^T \boldsymbol{\beta}^*| + 3\delta \|\Delta\|_2^2 \|\tilde{\mathbf{w}}_k\|_2 \|\boldsymbol{\beta}^*\|_2 \\ &\leq 9\|\Delta\|_2^2 \|\tilde{\mathbf{w}}_k\|_2 \|\boldsymbol{\beta}^*\|_2 + 3\delta \|\Delta\|_2^2 \|\tilde{\mathbf{w}}_k\|_2 \|\boldsymbol{\beta}^*\|_2. \end{aligned} \quad (6.5)$$

The derivations of (6.3)-(6.5) require a common condition $n \geq C(\delta)(s+1) \log(s+1)$ with exactly the same $C(\delta)$.

Next, we deal with the third term in (6.1).

$$\begin{aligned} & \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) [\mathbf{x}_j^T (\tilde{\mathbf{w}}_k - \mathbf{w}_k)] (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \\ &= \Delta^T \left(\frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \mathbf{x}_j \mathbf{x}_j^T - (\|\boldsymbol{\beta}^*\|_2^2 I + 2\boldsymbol{\beta}^* \boldsymbol{\beta}^{*T}) \right) (\tilde{\mathbf{w}}_k - \mathbf{w}_k) \\ &+ \|\boldsymbol{\beta}^*\|_2^2 \Delta^T (\tilde{\mathbf{w}}_k - \mathbf{w}_k) + 2(\boldsymbol{\beta}^{*T} \Delta) \left(\boldsymbol{\beta}^{*T} (\tilde{\mathbf{w}}_k - \mathbf{w}_k) \right). \end{aligned}$$

Applying Lemma A.1 one more time, the third term falls within

$$\|\boldsymbol{\beta}^*\|_2^2 \Delta^T (\tilde{\mathbf{w}}_k - \mathbf{w}_k) + 2(\boldsymbol{\beta}^{*T} \Delta) \left(\boldsymbol{\beta}^{*T} (\tilde{\mathbf{w}}_k - \mathbf{w}_k) \right) \pm \delta \|\boldsymbol{\beta}^*\|_2^2 \|\Delta\|_2 \|\tilde{\mathbf{w}}_k - \mathbf{w}_k\|_2$$

with probability at least $1 - \frac{1}{n}$, given $n \geq C(\delta)(s+1) \log(s+1)$. The magnitude of $\|\tilde{\mathbf{w}}_k - \mathbf{w}_k\|_2$ in terms of $\|\boldsymbol{\beta}^*\|_2$ and $\|\Delta\|_2$ can be estimated by

$$\begin{aligned} \tilde{\mathbf{w}}_k - \mathbf{w}_k &= \frac{1}{2} \left(\frac{1}{\|\boldsymbol{\beta}^*\|_2^2} - \frac{1}{\|\tilde{\boldsymbol{\beta}}\|_2^2} \right) \mathbf{e}_k + \frac{\tilde{\beta}_k}{3\|\tilde{\boldsymbol{\beta}}\|_2^4} \Delta + \frac{1}{3} \left(\frac{\tilde{\beta}_k}{\|\tilde{\boldsymbol{\beta}}\|_2^4} - \frac{\beta_k^*}{\|\boldsymbol{\beta}^*\|_2^4} \right) \boldsymbol{\beta}^*, \\ \|\tilde{\mathbf{w}}_k - \mathbf{w}_k\|_2 &\leq \frac{1}{2} \left| \frac{1}{\|\boldsymbol{\beta}^*\|_2^2} - \frac{1}{\|\tilde{\boldsymbol{\beta}}\|_2^2} \right| + \frac{|\tilde{\beta}_k|}{3\|\tilde{\boldsymbol{\beta}}\|_2^4} \|\Delta\|_2 + \frac{1}{3} \left| \frac{\tilde{\beta}_k}{\|\tilde{\boldsymbol{\beta}}\|_2^4} - \frac{\beta_k^*}{\|\boldsymbol{\beta}^*\|_2^4} \right| \|\boldsymbol{\beta}^*\|_2 \\ &= O\left(\frac{\|\Delta\|_2}{\|\boldsymbol{\beta}^*\|_2^3}\right). \end{aligned} \tag{6.6}$$

The last equality is because

$$\left| \frac{1}{\|\boldsymbol{\beta}^*\|_2^2} - \frac{1}{\|\tilde{\boldsymbol{\beta}}\|_2^2} \right| = \frac{\|\boldsymbol{\beta}^* + \Delta\|_2^2 - \|\boldsymbol{\beta}^*\|_2^2}{\|\boldsymbol{\beta}^*\|_2^2 \|\tilde{\boldsymbol{\beta}}\|_2^2} \leq \frac{2\|\boldsymbol{\beta}^*\|_2 \|\Delta\|_2 + \|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2^2 \|\tilde{\boldsymbol{\beta}}\|_2^2} = O\left(\frac{\|\Delta\|_2}{\|\boldsymbol{\beta}^*\|_2^3}\right),$$

and

$$\begin{aligned} \left| \frac{\tilde{\beta}_k}{\|\tilde{\boldsymbol{\beta}}\|_2^4} - \frac{\beta_k^*}{\|\boldsymbol{\beta}^*\|_2^4} \right| &= \frac{1}{\|\boldsymbol{\beta}^*\|_2^4 \|\tilde{\boldsymbol{\beta}}\|_2^4} \left| \|\boldsymbol{\beta}^*\|_2^4 \tilde{\beta}_k - \|\tilde{\boldsymbol{\beta}}\|_2^4 \beta_k^* \right| \\ &\leq \frac{\|\boldsymbol{\beta}^*\|_2^4 |\tilde{\beta}_k - \beta_k^*|}{\|\boldsymbol{\beta}^*\|_2^4 \|\tilde{\boldsymbol{\beta}}\|_2^4} + \frac{\left| \|\tilde{\boldsymbol{\beta}}\|_2^4 - \|\boldsymbol{\beta}^*\|_2^4 \right| |\beta_k^*|}{\|\boldsymbol{\beta}^*\|_2^4 \|\tilde{\boldsymbol{\beta}}\|_2^4} \\ &\leq \frac{\|\Delta\|_2}{\|\tilde{\boldsymbol{\beta}}\|_2^4} + \frac{\left(\|\boldsymbol{\beta}^*\|_2 + \|\Delta\|_2 \right)^4 - \|\boldsymbol{\beta}^*\|_2^4}{\|\boldsymbol{\beta}^*\|_2^3 \|\tilde{\boldsymbol{\beta}}\|_2^4} \\ &\leq O\left(\frac{\|\Delta\|_2}{\|\boldsymbol{\beta}^*\|_2^3}\right). \end{aligned}$$

Therefore, we bound the third term in (6.1) by

$$\begin{aligned} & \left| \frac{2}{n} \sum_{j=1}^n (\mathbf{x}_j^T \Delta) [\mathbf{x}_j^T (\tilde{\mathbf{w}}_k - \mathbf{w}_k)] (\mathbf{x}_j^T \boldsymbol{\beta}^*)^2 \right| \\ & \leq 3 \|\boldsymbol{\beta}^*\|_2^2 \|\Delta\|_2 \|\tilde{\mathbf{w}}_k - \mathbf{w}_k\|_2 + \delta \|\boldsymbol{\beta}^*\|_2^2 \|\Delta\|_2 \|\tilde{\mathbf{w}}_k - \mathbf{w}_k\|_2 \\ & = 3 \|\boldsymbol{\beta}^*\|_2^2 \cdot O\left(\frac{\|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2^3}\right) + \delta \|\boldsymbol{\beta}^*\|_2^2 \cdot O\left(\frac{\|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2^3}\right). \end{aligned} \quad (6.7)$$

The term $\|\tilde{\mathbf{w}}_k\|_2$ appears in (6.4) and (6.5), and needs to be bounded. By (2.8) and (6.6),

$$\begin{aligned} \|\mathbf{w}_k\|_2 & \leq \frac{1}{2\|\boldsymbol{\beta}^*\|_2^2} + \frac{\beta_k^*}{3\|\boldsymbol{\beta}^*\|_2^3} \leq \frac{1}{2\|\boldsymbol{\beta}^*\|_2^2} + \frac{1}{3\|\boldsymbol{\beta}^*\|_2^2} \\ & \leq \frac{5}{6\|\boldsymbol{\beta}^*\|_2^2}, \\ \|\tilde{\mathbf{w}}_k\|_2 & \leq \|\mathbf{w}_k\|_2 + O\left(\frac{\|\Delta\|_2}{\|\boldsymbol{\beta}^*\|_2^3}\right) \\ & \leq \frac{5}{6\|\boldsymbol{\beta}^*\|_2^2} + O\left(\frac{\|\Delta\|_2}{\|\boldsymbol{\beta}^*\|_2^3}\right). \end{aligned}$$

So far, the terms that differ $\hat{\beta}_k - \beta_k^*$ from $-\frac{1}{n} \sum_{j=1}^n \varepsilon_j (\mathbf{x}_j^T \tilde{\boldsymbol{\beta}}) (\mathbf{x}_j^T \tilde{\mathbf{w}}_k)$ (an asymptotically normal random variable) have been concentrated around their means with the concentration errors associated with δ . We can set $\delta = \frac{1}{p^2}$ so that all concentration errors vanish in an order faster than $\frac{\sqrt{n}}{p^2}$. And we can just ignore these terms. The only thing left to check is whether these mean terms are negligible after multiplying by \sqrt{n} , *i.e.* to show that $\sqrt{n}(\hat{\beta}_k - \beta_k^*)$ is approximately normal as $n \rightarrow \infty$. The goal is to show

$$\left| \sqrt{n}(\hat{\beta}_k - \beta_k^*) - Z_k \right| = o_p(1).$$

Under assumption (3.1), $n \geq C(\frac{1}{p^2})(s+1) \log(s+1)$ holds. Thus, by (6.1), (6.3)-(6.5), (6.7), and (3.7), we have with probability at least $1 - \frac{4}{n}$,

$$\begin{aligned} & \left| \sqrt{n}(\hat{\beta}_k - \beta_k^*) - Z_k \right| \\ & \leq \sqrt{n} \left(3 \|\Delta\|_2^3 \|\tilde{\mathbf{w}}_k\|_2 + 9 \|\Delta\|_2^2 \|\tilde{\mathbf{w}}_k\|_2 \|\boldsymbol{\beta}^*\|_2 + 3 \|\boldsymbol{\beta}^*\|_2^2 \cdot O\left(\frac{\|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2^3}\right) \right) + O\left(\frac{\sqrt{n}}{p^2}\right) \\ & \leq \sqrt{n} \left(\frac{5}{2} \frac{\|\Delta\|_2^3}{\|\boldsymbol{\beta}^*\|_2^2} + O\left(\frac{\|\Delta\|_2^4}{\|\boldsymbol{\beta}^*\|_2^3}\right) + \frac{15}{2} \frac{\|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2} + O\left(\frac{\|\Delta\|_2^3}{\|\boldsymbol{\beta}^*\|_2^2}\right) + O\left(\frac{\|\Delta\|_2^2}{\|\boldsymbol{\beta}^*\|_2}\right) \right) + O\left(\frac{\sqrt{n}}{p^2}\right). \end{aligned} \quad (6.8)$$

Here $C(\frac{1}{p^2})$ is a constant only depending on $\frac{1}{p^2}$. Together with (3.6), we obtain

$$\left| \sqrt{n}(\hat{\beta}_k - \beta_k^*) - Z_k \right|$$

$$\leq C_1 \frac{\sigma^3}{\|\beta^*\|_2^5} \frac{s \log p \sqrt{s \log p}}{n} + C_2 \frac{\sigma^4}{\|\beta^*\|_2^7} \frac{s^2 (\log p)^2}{n \sqrt{n}} + C_3 \frac{\sigma^2}{\|\beta^*\|_2^3} \frac{s \log p}{\sqrt{n}} + C_4 \frac{\sqrt{n}}{p^2} \tag{6.9}$$

with probability at least $1 - \epsilon_0 - \frac{4}{n}$. Plugging in $\|\beta^*\|_2 = O(\sqrt{s})$, the right hand side of (6.9) becomes $C'_1 \frac{(\log p)^{\frac{3}{2}}}{sn} + C'_2 \frac{(\log p)^2}{(sn)^{\frac{3}{2}}} + C'_3 \frac{\log p}{(sn)^{\frac{1}{2}}} + C'_4 \frac{n^{\frac{1}{2}}}{p^2}$, which is equal to ϵ'_n . \square

PROOF OF THEOREM 3.2.

(i) It is intuitive that a reasonable estimator would combine $\hat{\beta}_{1k}$ and $\hat{\beta}_{2k}$ so as to integrate both pieces of information. We know that Z_{1k} and Z_{2k} have asymptotic variances $\sigma^2 \tau_{1k}^2$ and $\sigma^2 \tau_{2k}^2$, respectively. By Theorem 3.1 and Remark 3.1,

$$\begin{aligned} \hat{\beta}_{1k} &:= \beta_k^* + \frac{Z_{1k}}{\sqrt{n}} + op\left(\frac{\sigma \tau_{1k}}{\sqrt{n}}\right) = \beta_k^* + N\left(0, \frac{\sigma^2 \tau_{1k}^2}{n}\right) + op\left(\frac{\sigma \tau_{1k}}{\sqrt{n}}\right), \\ \hat{\beta}_{2k} &:= \beta_k^* + \frac{Z_{2k}}{\sqrt{n}} + op\left(\frac{\sigma \tau_{2k}}{\sqrt{n}}\right) = \beta_k^* + N\left(0, \frac{\sigma^2 \tau_{2k}^2}{n}\right) + op\left(\frac{\sigma \tau_{2k}}{\sqrt{n}}\right). \end{aligned}$$

Judging from these two formulae, any convex combination of $\hat{\beta}_{1k}$ and $\hat{\beta}_{2k}$ remains asymptotically unbiased (has asymptotic mean equal to β_k^*) and possibly attain a smaller asymptotic variance. Suppose we have the final estimator given by

$$\hat{\beta}_k^{swap} = a \hat{\beta}_{1k} + (1 - a) \hat{\beta}_{2k} \quad \text{for } a \in (0, 1),$$

then

$$\sqrt{n}(\hat{\beta}_k^{swap} - \beta_k^*) = a \cdot Z_{1k} + (1 - a) \cdot Z_{2k} + op(1).$$

It is easy to verify that $cov(Z_{1k}, Z_{2k}) = 0$, and hence the asymptotic variance of $\hat{\beta}_k^{swap}$ is $\sigma^2 \left(a^2 \tau_{1k}^2 + (1 - a)^2 \tau_{2k}^2 \right)$, which is a quadratic form in a . When $a = \frac{\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}$, the asymptotic variance attains minimum, in which case $\hat{\beta}_k^{swap}$ has approximately asymptotic distribution $N\left(0, \sigma^2 \frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}\right)$.

(ii) Let $a = \frac{\tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}$. By Lemma A.1, for n satisfying (3.1),

$$\begin{aligned} \mathbb{P} \left\{ a^2 \left| \frac{Var(Z_{1k})}{\sigma^2} - \tau_{1k}^2 \right| \leq \frac{a^2}{p^2} \|\tilde{\beta}_1\|_2^2 \|\tilde{\mathbf{w}}_{1k}\|_2^2 \right\} &\geq 1 - \frac{1}{n}, \\ \mathbb{P} \left\{ (1 - a)^2 \left| \frac{Var(Z_{2k})}{\sigma^2} - \tau_{2k}^2 \right| \leq \frac{(1 - a)^2}{p^2} \|\tilde{\beta}_2\|_2^2 \|\tilde{\mathbf{w}}_{2k}\|_2^2 \right\} &\geq 1 - \frac{1}{n}. \end{aligned}$$

Let $Z_k^{swap} = a \cdot Z_{1k} + (1 - a) \cdot Z_{2k}$, then

$$\mathbb{P} \left\{ \left| \frac{Var(Z_k^{swap})}{\sigma^2} - a^2 \tau_{1k}^2 - (1 - a)^2 \tau_{2k}^2 \right| \leq \mathbf{U}(a) \right\} \geq 1 - \frac{2}{n}, \tag{6.10}$$

where

$$\mathbf{U}(a) = \frac{1}{p^2} \left(a^2 \|\tilde{\beta}_1\|_2^2 \|\tilde{\mathbf{w}}_{1k}\|_2^2 + (1-a)^2 \|\tilde{\beta}_2\|_2^2 \|\tilde{\mathbf{w}}_{2k}\|_2^2 \right).$$

Similar to the argument in Theorem 3.1,

$$\begin{aligned} \|\tilde{\beta}_1\|_2^2 \|\tilde{\mathbf{w}}_{1k}\|_2^2 &\leq \left(\|\beta^*\|_2 + 2\|\Delta\|_2 \right)^2 \left(\frac{5}{6\|\beta^*\|_2^2} + O\left(\frac{\|\Delta\|_2}{\|\beta^*\|_2^3}\right) \right)^2 \\ &\asymp \frac{25}{36s} + O\left(\sqrt{\frac{\log p}{ns^3}}\right). \end{aligned} \quad (6.11)$$

Plugging (6.11) into (6.10), together with the fact $a^2 + (1-a)^2 \leq \frac{1}{2}$, we have

$$\mathbb{P} \left\{ \left| \frac{\text{Var}(Z_k^{swap})}{\sigma^2} - a^2\tau_{1k}^2 - (1-a)^2\tau_{2k}^2 \right| \leq \frac{1}{p^2} \left(\frac{25}{72s} + O\left(\sqrt{\frac{\log p}{ns^3}}\right) \right) \right\} \geq 1 - \frac{2}{n}.$$

Further plugging in the value of a ,

$$\mathbb{P} \left\{ \text{Var}(Z_k^{swap}) \leq \sigma^2 \frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2} + O\left(\frac{1}{p^2}\right) \right\} \geq 1 - \frac{2}{n}.$$

By (3.9), we get

$$\mathbb{P} \left\{ \left| \hat{\beta}_k^{swap} - \beta_k^* \right| \leq \frac{1}{\sqrt{n}} |Z_k^{swap}| + \frac{\epsilon'_n}{\sqrt{n}} \right\} \geq 1 - 2\epsilon''_n.$$

Let $\Phi(\cdot)$ be the cumulative distribution function of standard normal distribution, for $\forall r > 0$,

$$\mathbb{P} \left\{ \frac{1}{\sqrt{n}} |Z_k^{swap}| \leq \frac{r \sqrt{\text{Var}(Z_k^{swap})}}{\sqrt{n}} \right\} = 2\Phi(r) - 1.$$

The above three formulae imply (3.11), which means the interval $\left(\hat{\beta}_k^{swap} - \frac{\sigma r}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}} + \frac{\epsilon'_n}{\sqrt{n}}, \hat{\beta}_k^{swap} + \frac{\sigma r}{\sqrt{n}} \sqrt{\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2}} + \frac{\epsilon'_n}{\sqrt{n}} \right)$ has asymptotic coverage probability at least $2\Phi(r) - 1$.

(iii) By Remark 3.1, $\frac{\tau_{1k}^2 \tau_{2k}^2}{\tau_{1k}^2 + \tau_{2k}^2} \leq \frac{1}{2} \max\{\tau_{1k}^2, \tau_{2k}^2\} \asymp \frac{3}{8s} + O\left(\sqrt{\frac{\log p}{ns^3}}\right)$ for all $k = 1, 2, \dots, p$. Plus (3.11) holds uniformly over k , we obtain

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{k \in [p]} \left| \sqrt{n}(\hat{\beta}_k^{swap} - \beta_k^*) \right| \leq \sqrt{\frac{3}{8s}} \sigma r \right\} \geq 2\phi(r) - 1,$$

and then replace r by $\Phi(1 - \frac{\alpha}{2})$. □

PROOF OF COROLLARY 3.1

It is easy to verify that the asymptotic distribution of $\hat{\beta}^{swap}$ is multinormal $N(\beta^*, \frac{\sigma^2}{n}\mathbf{V})$, where

$$\begin{aligned} V_{kl} &= \frac{1}{\sigma^2} Cov(Z_k^{swap}, \tilde{Z}_l) = \frac{1}{\sigma^2} Cov\left(a_k Z_{1k} + (1-a_k)Z_{2k}, a_l Z_{1l} + (1-a_l)Z_{2l}\right) \\ &= \frac{1}{\sigma^2} \left[a_k a_l Cov(Z_{1k}, Z_{1l}) + (1-a_k)(1-a_l)Cov(Z_{2k}, Z_{2l}) \right] \\ &= a_k a_l \left[\|\tilde{\beta}_1\|_2^2 (\tilde{\mathbf{w}}_{1k}^T \tilde{\mathbf{w}}_{1l}) + 2(\tilde{\beta}_1^T \tilde{\mathbf{w}}_{1k})(\tilde{\beta}_1^T \tilde{\mathbf{w}}_{1l}) \right] \\ &\quad + (1-a_k)(1-a_l) \left[\|\tilde{\beta}_2\|_2^2 (\tilde{\mathbf{w}}_{2k}^T \tilde{\mathbf{w}}_{2l}) + 2(\tilde{\beta}_2^T \tilde{\mathbf{w}}_{2k})(\tilde{\beta}_2^T \tilde{\mathbf{w}}_{2l}) \right]. \end{aligned}$$

By linear algebra,

$$\sup_{\mathbf{h} \neq 0, \mathbf{h} \in \mathbb{R}^p} \frac{|\mathbf{h}^T (\hat{\beta}^{swap} - \beta^*)|^2}{\mathbf{h}^T \mathbf{V} \mathbf{h}} = (\hat{\beta}^{swap} - \beta^*)^T \mathbf{V}^{-1} (\hat{\beta}^{swap} - \beta^*) \stackrel{d}{\rightarrow} \frac{\sigma^2}{n} \chi_p^2.$$

Therefore, the coverage probability in the worst direction is bounded below,

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \sup_{\mathbf{h} \neq 0, \mathbf{h} \in \mathbb{R}^p} \frac{|\mathbf{h}^T (\hat{\beta}^{swap} - \beta^*)|^2}{\mathbf{h}^T (\frac{\sigma^2}{n} \mathbf{V}) \mathbf{h}} \leq \chi_{p,\alpha}^2 \right\} \geq 1 - \alpha.$$

It further implies, for $\forall h \neq 0, h \in \mathbb{R}^p$,

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \left| \mathbf{h}^T (\hat{\beta}^{swap} - \beta^*) \right| \leq \sqrt{\frac{\sigma^2}{n} \chi_{p,\alpha}^2 \mathbf{h}^T \mathbf{V} \mathbf{h}} \right\} \geq 1 - \alpha. \quad \square$$

Appendix A: Appendix section

The theorem and lemma below are stated and proved in [4].

Theorem A.1. [4] Suppose the tuning parameters in the thresholded Wirtinger algorithm are suitably chosen, and the sample size $n \geq K(1 + \frac{\sigma}{\|\beta\|_2})^2 s^2 \log(np)$ for some absolute constant $K > 0$, Let $S = \text{support}(\beta)$, then

$$\begin{aligned} \inf_{\|\beta\|_0=s} \mathbb{P}_{(\mathbf{x}, \mathbf{y}|\beta)} \left\{ \text{supp}(\tilde{\beta}^{(t)}) \subseteq S \text{ and} \right. \\ \left. \min_{i=0,1} \|\tilde{\beta}^{(t)} - (-1)^i \beta\|_2 \leq \frac{1}{6} \left(1 - \frac{w}{16}\right)^t \|\beta^*\|_2 + C \frac{\sigma}{\|\beta^*\|_2} \sqrt{\frac{s \log p}{n}} \right\} \\ > 1 - \frac{46}{n} - 10e^{-s} - \frac{t}{np^2} \end{aligned}$$

for some absolute constant $C > 0$, where w is the gradient descent step size. When $\frac{\sigma}{\|\beta^*\|_2} = o(\sqrt{\frac{n}{\log n}})$ and is unknown, we can estimate $\|\beta^*\|_2^2$ by

$$\phi^2 := \widehat{\|\beta^*\|_2^2} = \frac{1}{n} \sum_{j=0}^n y_j$$

and define

$$\hat{\sigma} = \sqrt{\left(\frac{1}{n} \sum_{j=0}^n y_j^2\right) - 3\phi^4}.$$

Then with probability at least $1 - \frac{1}{n}$, there holds $\frac{\hat{\sigma}}{\phi^2} \asymp \frac{\sigma}{\|\beta^*\|_2}$. If the sample size $n \geq K(1 + \frac{\hat{\sigma}}{\phi^2})^2 s^2 \log(np)$, then the above claim holds with the first term on the right hand side $\frac{46}{n}$ replaced by $\frac{47}{n}$.

Lemma A.1. [4] Suppose \mathbf{x}_j are i.i.d. $N(0, I_{p \times p})$. Then on an event with probability at least $1 - \frac{1}{n}$, we have

$$\left\| \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j^T \mathbf{b})^2 \mathbf{x}_{jS} \mathbf{x}_{jS}^T - (\|\mathbf{b}\|_2^2 I_S + 2\mathbf{b}\mathbf{b}^T) \right\| \leq \delta \|\mathbf{b}\|_2^2$$

provided $n \geq C(\delta)s \log s$, where $C(\delta)$ is constant only depending on δ . Here, I_S is a diagonal matrix with the diagonal elements in S equal to 1, whereas others equal to 0. And $\text{supp}(\mathbf{b}) \subset S$.

Acknowledgments

The author would like to thank Cun-Hui Zhang and Pierre Bellec for several constructive advises and enlightening discussions. The author also would like to thank the anonymous reviewers for the valuable input.

References

- [1] BAUSCHKE, H. H., COMBETTES, P. L. and LUKE, D. R. (2003). Hybrid projection–reflection method for phase retrieval. *JOSA A* **20** 1025–1034. [MR1914365](#)
- [2] BELLEC, P. C. and ZHANG, C.-H. (2019). De-Biasing The Lasso With Degrees-of-Freedom Adjustment. *arXiv preprint arXiv:1902.08885*.
- [3] BEN-TAL, A. and NEMIROVSKI, A. (2001). *Lectures on modern convex optimization: analysis, algorithms, and engineering applications* **2**. Siam. [MR1857264](#)

- [4] CAI, T. T., LI, X. and MA, Z. (2016). Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow. *The Annals of Statistics* **44** 2221–2251. [MR3546449](#)
- [5] CANDES, E. J., ELДАР, Y. C., STROHMER, T. and VORONINSKI, V. (2015). Phase retrieval via matrix completion. *SIAM review* **57** 225–251. [MR3345342](#)
- [6] CANDES, E. J., LI, X. and SOLTANOLKOTABI, M. (2015). Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory* **61** 1985–2007. [MR3332993](#)
- [7] CANDES, E. J., STROHMER, T. and VORONINSKI, V. (2013). Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics* **66** 1241–1274. [MR3069958](#)
- [8] CHAI, A., MOSCOSO, M. and PAPANICOLAOU, G. (2010). Array imaging using intensity-only measurements. *Inverse Problems* **27** 015005. [MR2746408](#)
- [9] CHEN, Y. and CANDES, E. (2015). Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems* 739–747. [MR3628877](#)
- [10] DONOHO, D. L. (2006). Compressed sensing. *IEEE Transactions on information theory* **52** 1289–1306. [MR2241189](#)
- [11] DUCHI, J. C. and RUAN, F. (2017). Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *arXiv preprint arXiv:1705.02356*. [MR3994397](#)
- [12] FAZEL, M., HINDI, H. and BOYD, S. P. (2003). Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003*. **3** 2156–2162. IEEE.
- [13] FIENUP, J. R. (1982). Phase retrieval algorithms: a comparison. *Applied optics* **21** 2758–2769.
- [14] FIENUP, J. R. (1987). Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint. *JOSA A* **4** 118–123.
- [15] GERCHBERG, R. W. (1972). A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* **35** 237–246.
- [16] IWEN, M., VISWANATHAN, A. and WANG, Y. (2017). Robust sparse phase retrieval made easy. *Applied and Computational Harmonic Analysis* **42** 135–142. [MR3574564](#)
- [17] JAGANATHAN, K., ELДАР, Y. C. and HASSIBI, B. (2015). Phase retrieval: An overview of recent developments. *arXiv preprint arXiv:1510.07713*. [MR3674475](#)
- [18] JAGANATHAN, K., OYMAK, S. and HASSIBI, B. (2012). Recovery of sparse 1-D signals from the magnitudes of their Fourier transform. In *2012 IEEE International Symposium on Information Theory Proceedings* 1473–1477. IEEE.
- [19] JAGANATHAN, K., OYMAK, S. and HASSIBI, B. (2013). Sparse phase re-

- trieval: Convex algorithms and limitations. In *2013 IEEE International Symposium on Information Theory* 1022–1026. IEEE. [MR3620357](#)
- [20] JAVANMARD, A. and MONTANARI, A. (2014). Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory. *IEEE Transactions on Information Theory* **60** 6522–6554. [MR3265038](#)
- [21] LEVI, A. and STARK, H. (1984). Image restoration by the method of generalized projections with application to restoration from magnitude. *JOSA A* **1** 932–943. [MR0758183](#)
- [22] LI, X. and VORONINSKI, V. (2013). Sparse signal recovery from quadratic measurements via convex programming. *SIAM Journal on Mathematical Analysis* **45** 3019–3033. [MR3106479](#)
- [23] MITRA, R. and ZHANG, C.-H. (2016). The benefit of group sparsity in group inference with de-biased scaled group Lasso. *Electronic Journal of Statistics* **10** 1829–1873. [MR3522662](#)
- [24] NETRAPALLI, P., JAIN, P. and SANGHAVI, S. (2013). Phase retrieval using alternating minimization. In *Advances in Neural Information Processing Systems* 2796–2804. [MR3385838](#)
- [25] OHLSSON, H., YANG, A. Y., DONG, R. and SASTRY, S. S. (2011). Compressive phase retrieval from squared output measurements via semidefinite programming. *arXiv preprint arXiv:1111.6323* 1–27.
- [26] OYMAK, S., JALALI, A., FAZEL, M., ELDAR, Y. C. and HASSIBI, B. (2015). Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Transactions on Information Theory* **61** 2886–2908. [MR3342310](#)
- [27] SHECHTMAN, Y., BECK, A. and ELDAR, Y. C. (2014). GESPAR: Efficient phase retrieval of sparse signals. *IEEE transactions on signal processing* **62** 928–938. [MR3160324](#)
- [28] SHECHTMAN, Y., ELDAR, Y. C., COHEN, O., CHAPMAN, H. N., MIAO, J. and SEGEV, M. (2015). Phase retrieval with application to optical imaging: a contemporary overview. *IEEE signal processing magazine* **32** 87–109.
- [29] VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42** 1166–1202. [MR3224285](#)
- [30] WALDSPURGER, I., D’ASPREMONT, A. and MALLAT, S. (2015). Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming* **149** 47–81. [MR3300456](#)
- [31] WANG, G., GIANNAKIS, G. B. and ELDAR, Y. C. (2018). Solving systems of random quadratic equations via truncated amplitude flow. *IEEE Transactions on Information Theory* **64** 773–794. [MR3762591](#)
- [32] WANG, G., ZHANG, L., GIANNAKIS, G. B., AKÇAKAYA, M. and CHEN, J. (2017). Sparse phase retrieval via truncated amplitude flow. *IEEE Transactions on Signal Processing* **66** 479–491. [MR3756889](#)
- [33] ZHANG, C.-H. (2011). Statistical inference for high-dimensional data. *Mathematisches Forschungsinstitut Oberwolfach: Very High Dimensional Semiparametric Models, Report* **48** 28–31.

- [34] ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76** 217–242. [MR3153940](#)