

Solving Systems of Random Quadratic Equations via Truncated Amplitude Flow

Gang Wang¹, Student Member, IEEE, Georgios B. Giannakis, Fellow, IEEE, and Yonina C. Eldar, Fellow, IEEE

Abstract—This paper presents a new algorithm, termed *truncated amplitude flow* (TAF), to recover an unknown vector x from a system of quadratic equations of the form $y_i = |\langle a_i, x \rangle|^2$, where a_i 's are given random measurement vectors. This problem is known to be *NP-hard* in general. We prove that as soon as the number of equations is on the order of the number of unknowns, TAF recovers the solution exactly (up to a global unimodular constant) with high probability and complexity growing linearly with both the number of unknowns and the number of equations. Our TAF approach adapts the *amplitude-based* empirical loss function and proceeds in two stages. In the first stage, we introduce an *orthogonality-promoting* initialization that can be obtained with a few power iterations. Stage two refines the initial estimate by successive updates of scalable *truncated generalized gradient iterations*, which are able to handle the rather challenging nonconvex and nonsmooth amplitude-based objective function. In particular, when vectors x and a_i 's are real valued, our gradient truncation rule provably eliminates erroneously estimated signs with high probability to markedly improve upon its untruncated version. Numerical tests using synthetic data and real images demonstrate that our initialization returns more accurate and robust estimates relative to spectral initializations. Furthermore, even under the same initialization, the proposed amplitude-based refinement outperforms existing Wirtinger flow variants, corroborating the superior performance of TAF over state-of-the-art algorithms.

Index Terms—Nonconvex optimization, phase retrieval, amplitude-based cost function, orthogonality-promoting initialization, truncated gradient, linear convergence to global minimum.

I. INTRODUCTION

CONSIDER a system of m quadratic equations

$$y_i = |\langle a_i, x \rangle|^2, \quad 1 \leq i \leq m \quad (1)$$

where the data vector $y := [y_1 \ \dots \ y_m]^T$ and feature vectors $a_i \in \mathbb{R}^n$ or \mathbb{C}^n are known, whereas the vector $x \in \mathbb{R}^n$ or \mathbb{C}^n

is the wanted unknown. When $\{a_i\}_{i=1}^m$ and/or x are complex, the magnitudes of their inner-products $\{|\langle a_i, x \rangle|\}_{i=1}^m$ are given but phase information is lacking; in the real case only the signs of $\{\langle a_i, x \rangle\}_{i=1}^m$ are unknown. Assuming that the system of quadratic equations in (1) admits a unique solution x (up to a global unimodular constant), our objective is to reconstruct x from m phaseless quadratic equations, or equivalently, to recover the missing signs or phases of $\{\langle a_i, x \rangle\}_{i=1}^m$ under real- or complex-valued settings. It has been established that $m \geq 2n - 1$ or $m \geq 4n - 4$ generic measurements $\{a_i; y_i\}_{i=1}^m$ as in (1) suffice for uniquely determining an n -dimensional real-valued or complex-valued vector x [1], [2], respectively, while the former with $m = 2n - 1$ has also been shown to be necessary [1], [3].

The problem in (1) constitutes an instance of nonconvex quadratic programming, that is generally known to be *NP-hard* [4]. Specifically for real-valued vectors $\{a_i\}$ and x , problem (1) can be understood as a combinatorial optimization since one seeks a series of signs $\{s_i = \pm 1\}_{i=1}^m$, such that the solution to the system of linear equations $\langle a_i, x \rangle = s_i \psi_i$, where $\psi_i := \sqrt{y_i}$, obeys the given quadratic system. Evidently, there are a total of 2^m different combinations of $\{s_i\}_{i=1}^m$, among which only two lead to x up to a global sign. The complex case becomes even more complicated, where instead of a set of signs $\{s_i\}_{i=1}^m$, one must determine a collection of unimodular complex scalars $\{\sigma_i \in \mathbb{C}\}_{i=1}^m$. Special cases with $a_i > \mathbf{0}$ (entry-wise inequality), $x_i^2 = 1$, and $y_i = 0$, $1 \leq i \leq m$ correspond to the so-called *stone problem* [5, Section 3.4.1], [6].

In many fields of physical sciences and engineering, the problem of recovering the phase from intensity- or magnitude-only measurements is commonly referred to as *phase retrieval* [7]–[9]. Relevant application domains include X-ray crystallography [10], optics [11], array and high-power coherent diffractive imaging [12], [13], astronomy [14], and microscopy [15]. In these settings, due to physical limitations, optical sensors and detectors such as charge-coupled device (CCD) cameras, photosensitive films, and human eyes can record only the (squared) modulus of the Fresnel or Fraunhofer diffraction pattern, while losing the phase of the incident light striking the object. It has been shown that reconstructing a discrete, finite-duration signal from its Fourier transform magnitudes is generally *NP-complete* [16]. Even checking quadratic feasibility (i.e., whether a solution to a given quadratic system exists or not) is itself an *NP-hard* problem [17, Theorem 2.6]. Thus, despite its simple form and practical relevance across various fields, tackling the quadratic system in (1) is challenging and *NP-hard* in general.

Manuscript received July 3, 2016; revised April 1, 2017; accepted July 25, 2017. Date of publication September 26, 2017; date of current version January 18, 2018. G. Wang and G. B. Giannakis were supported by NSF under Grant 1500713 and Grant 1514056. This paper was presented at the 2016 Neural Information Processing Systems Conference.

G. Wang is with the Digital Technology Center and the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA, and also with the State Key Laboratory of Intelligent Control and Decision of Complex Systems, Beijing Institute of Technology, Beijing 100081, China (e-mail: gangwang@umn.edu).

G. B. Giannakis is with the Digital Technology Center and the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: georgios@umn.edu).

Y. C. Eldar is with the Department of Electrical Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel (e-mail: yonina@ee.technion.ac.il).

Communicated by C. Caramanis, Associate Editor for Machine Learning.

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

Digital Object Identifier 10.1109/TIT.2017.2756858

A. Prior Art

Adopting the least-squares criterion, the task of recovering \mathbf{x} from data y_i observed in additive white Gaussian noise (AWGN) can be recast as that of minimizing the *intensity-based* empirical loss [18]

$$\underset{\mathbf{z} \in \mathbb{R}^n / \mathbb{C}^n}{\text{minimize}} \quad f(\mathbf{z}) := \frac{1}{2m} \sum_{i=1}^m (y_i - |\langle \mathbf{a}_i, \mathbf{z} \rangle|^2)^2. \quad (2)$$

An alternative is to consider an *amplitude-based* loss, in which ψ_i is observed instead of y_i in AWGN [7], [19]

$$\underset{\mathbf{z} \in \mathbb{R}^n / \mathbb{C}^n}{\text{minimize}} \quad h(\mathbf{z}) := \frac{1}{2m} \sum_{i=1}^m (\psi_i - |\langle \mathbf{a}_i, \mathbf{z} \rangle|)^2. \quad (3)$$

Unfortunately, the presence of quadratic terms in (2) or the modulus in (3) renders the corresponding objective function nonconvex. Minimizing nonconvex objectives, which may exhibit many stationary points, is in general *NP-hard* [20]. In fact, even checking whether a given point is a local minimum or establishing convergence to a local minimum turns out to be *NP-complete* [20].

In the classical discretized one-dimensional (1D) phase retrieval, the amplitude vector $\boldsymbol{\psi}$ corresponds to the n -point Fourier transform of the n -dimensional signal \mathbf{x} [21]. It has been shown based on spectral factorization that in general there is no unique solution to 1D phase retrieval, even if we disregard trivial ambiguities [22]. To overcome this ill-posedness, several approaches have been suggested. One possibility is to assume additional constraints on the unknown signal such as sparsity [13], [23]–[25]. Other approaches rely on introducing redundancy into the measurements using for example, the short-time Fourier transform, or masks [26], [27]. Finally, recent works assume random measurements (e.g., Gaussian $\{\mathbf{a}_i\}$ designs) [6], [8], [18], [25], [28]–[30]. Henceforth, this paper focuses on random measurements $\{\psi_i\}$ obtained from independently and identically distributed (i.i.d.) Gaussian $\{\mathbf{a}_i\}$ designs.

Existing approaches to solving (2) (or related ones using the Poisson likelihood; see, e.g., [6]) or (3) fall under two categories: nonconvex and convex ones. Popular nonconvex solvers include alternating projection such as Gerchberg-Saxton [31] and Fineup [7], AltMinPhase [28], (Truncated) Wirtinger flow (WF/TWF) [6], [18], [32], and Karzmarz variants [33] as well as trust-region methods [34]. Inspired by WF, other relevant judiciously initialized counterparts have also been developed for faster semidefinite optimization [35], [36], blind deconvolution [37], and matrix completion [38]. Convex counterparts on the other hand rely on the so-called *matrix-lifting* technique or *Shor's* semidefinite relaxation to obtain the solvers abbreviated as PhaseLift [29], PhaseCut [39], and CoRK [40]. Further approaches dealing with noisy or sparse phase retrieval are discussed in [24], [31], and [42]–[45].

In terms of sample complexity, it has been proven that¹ $\mathcal{O}(n)$ noise-free random measurements suffice for uniquely determining a general signal [25]. It is also self-evident

¹The notation $\phi(n) = \mathcal{O}(g(n))$ means that there is a constant $c > 0$ such that $|\phi(n)| \leq c|g(n)|$.

that recovering a general n -dimensional \mathbf{x} requires at least $\mathcal{O}(n)$ measurements. Convex approaches enable exact recovery from the optimal bound $\mathcal{O}(n)$ of noiseless Gaussian measurements [45]; they are based on solving a semidefinite program with a matrix variable of size $n \times n$, thus incurring worst-case computational complexity on the order of $\mathcal{O}(n^{4.5})$ [39] that does not scale well with the dimension n . Upon exploiting the underlying problem structure, $\mathcal{O}(n^{4.5})$ can be reduced to $\mathcal{O}(n^3)$ [39]. Solving for vector variables, nonconvex approaches achieve significantly improved computational performance. Using formulation (3) and adopting a spectral initialization commonly employed in matrix completion [46], AltMinPhase establishes exact recovery with sample complexity $\mathcal{O}(n \log^3 n)$ under i.i.d. Gaussian $\{\mathbf{a}_i\}$ designs with resampling [28].

Concerning formulation (2), WF iteratively refines the spectral initial estimate by means of a gradient-like update, which can be approximately interpreted as a stochastic gradient descent variant [18], [32]. The follow-up TWF improves upon WF through a truncation procedure to separate gradient components of excessively extreme (large or small) sizes. Likewise, due to the heavy tails present in the initialization stage, data $\{y_i\}_{i=1}^m$ are pre-screened to yield improved initial estimates in the so-termed truncated spectral initialization method [6]. WF allows exact recovery from $\mathcal{O}(n \log n)$ measurements in $\mathcal{O}(mn^2 \log(1/\epsilon))$ time or flops to yield an ϵ -accurate solution for any given $\epsilon > 0$ [18], while TWF advances these to $\mathcal{O}(n)$ measurements and $\mathcal{O}(mn \log(1/\epsilon))$ time [6]. Interestingly, the truncation procedure in the gradient stage turns out to be useful in avoiding spurious stationary points in the context of nonconvex optimization, as will be justified in Section IV by the numerical comparison between our amplitude flow (AF) algorithms with or without the judiciously designed truncation rule. It is also worth mentioning that when $m \geq Cn \log^3 n$ for some sufficiently large positive constant C , the objective function in (3) is shown to admit benign geometric structure that allows certain iterative algorithms (e.g., trust-region methods) to efficiently find a global minimizer with random initializations [34]. Hence, the challenge of solving systems of random quadratic equations lies in the case where a near-optimal number of equations are involved, e.g., $m = 2n - 1$ in the real-valued setting.

Although achieving a linear (in the number of unknowns n) sample and computational complexity, the state-of-the-art TWF approach still requires at least $4n \sim 5n$ equations to yield stable empirical success rate (e.g., $\geq 99\%$) under the noiseless real-valued Gaussian model [6, Section 3], which is more than twice the known information-limit of $m = 2n - 1$ [1]. Similar though less obvious results hold in the complex-valued scenario. While the truncated spectral initialization in [6] improves upon the “plain-vanilla” spectral initialization, its performance still suffers when the number of measurements is relatively small and its advantage (over the untruncated one) diminishes as the number of measurements grows; see more details in Figure 4 and Section II. Furthermore, extensive numerical and experimental validation confirms that the *amplitude-based* cost function performs significantly better

than the *intensity-based* one [47]; that is, formulation (3) is superior to (2). Hence, besides enhancing initialization, markedly improved performance in the gradient stage can be expected by re-examining the amplitude-based cost function and incorporating judiciously designed gradient regularization rules.

B. This Paper

Along the lines of suitably initialized nonconvex schemes [6], [18] and inspired by [47], the present paper develops a linear-time (i.e., the computational time linearly in both dimensions m and n) algorithm to minimize the amplitude-based cost function, referred to as *truncated amplitude flow* (TAF). Our approach provably recovers an n -dimensional unknown signal \mathbf{x} exactly from a near-optimal number of noiseless random measurements, while also featuring near-perfect statistical performance in the noisy setting. TAF operates in two stages: In the first stage, we introduce an orthogonality-promoting initialization that is computable using a few power iterations. Stage two refines the initial estimate by successive updates of truncated generalized gradient iterations.

Our initialization is built upon the hidden orthogonality characteristics of high-dimensional random vectors [48], which is in contrast to spectral alternatives originating from the strong law of large numbers (SLLN) [6], [13], [18]. Furthermore, one challenge of phase retrieval lies in reconstructing the signs/phases of $\langle \mathbf{a}_i, \mathbf{x} \rangle$ in the real- or complex-valued settings. Our TAF's refinement stage leverages a simple yet effective regularization rule to eliminate the erroneously estimated phases in the generalized gradient components with high probability. Simulated tests corroborate that the proposed initialization returns more accurate and robust initial estimates than its spectral counterparts in the noiseless and noisy settings. In addition, our TAF (with gradient truncation) markedly improves upon its "plain-vanilla" version AF. Empirical results demonstrate the advantage of TAF over its competing alternatives.

Focusing on the same amplitude-based cost function, an independent work develops the so-termed reshaped Wirtinger flow (RWF) algorithm [49], which coincides with amplitude flow (AF). A slightly modified variant of spectral initialization [18] is used to obtain an initial guess, followed by a sequence of non-truncated generalized gradient iterations [49]. Numerical comparisons show that the proposed TAF method performs better than RWF especially when the number of equations approaches the information-theoretic limit ($2n - 1$ in the real case).

The remainder of this paper is organized as follows. The amplitude-based cost function, as well as the two algorithmic stages is described and analyzed in Section II. Section III summarizes the TAF algorithm and establishes its theoretical performance. Extensive simulated tests comparing TAF with Wirtinger-based approaches are presented in Section IV. Finally, main proofs are given in Section V, while technical details are deferred to the Appendix.

II. TRUNCATED AMPLITUDE FLOW

In this section, the two stages of our TAF algorithm are detailed. First, the challenge of handling the nonconvex and nonsmooth amplitude-based cost function is analyzed, and addressed by a carefully designed gradient regularization rule. Limitations of (truncated) spectral initializations are then pointed out, followed by a simple motivating example to inspire our orthogonality-promoting initialization method. For concreteness, the analysis will focus on the real-valued Gaussian model with $\mathbf{x} \in \mathbb{R}^n$ and i.i.d. design vectors $\mathbf{a}_i \in \mathbb{R}^n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. Numerical experiments using the complex-valued Gaussian model with $\mathbf{x} \in \mathbb{C}^n$ and i.i.d. $\mathbf{a}_i \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n) := \mathcal{N}(\mathbf{0}, \mathbf{I}_n/2) + j\mathcal{N}(\mathbf{0}, \mathbf{I}_n/2)$ will be discussed briefly.

To start, let us define the Euclidean distance of any estimate \mathbf{z} to the solution set: $\text{dist}(\mathbf{z}, \mathbf{x}) := \min \{\|\mathbf{z} + \mathbf{x}\|, \|\mathbf{z} - \mathbf{x}\|\}$ for real-valued signals, and $\text{dist}(\mathbf{z}, \mathbf{x}) := \min_{\phi \in [0, 2\pi)} \|\mathbf{z} - \mathbf{x}e^{j\phi}\|$ for complex-valued ones [18], where $\|\cdot\|$ denotes the Euclidean norm. Define also the indistinguishable global phase constant in the real-valued setting as

$$\phi(\mathbf{z}) := \begin{cases} 0, & \|\mathbf{z} - \mathbf{x}\| \leq \|\mathbf{z} + \mathbf{x}\|, \\ \pi, & \text{otherwise.} \end{cases} \quad (4)$$

Henceforth, fixing \mathbf{x} to be any solution of the given quadratic system (1), we always assume that $\phi(\mathbf{z}) = 0$; otherwise, \mathbf{z} is replaced by $e^{-j\phi(\mathbf{z})}\mathbf{z}$, but for simplicity of presentation, the constant phase adaptation term $e^{-j\phi(\mathbf{z})}$ will be dropped whenever it is clear from the context.

A. Truncated Generalized Gradient Stage

For brevity, collect all vectors $\{\mathbf{a}_i\}_{i=1}^m$ in the $m \times n$ matrix $\mathbf{A} := [\mathbf{a}_1 \cdots \mathbf{a}_m]^T$, and all amplitudes $\{\psi_i\}_{i=1}^m$ to form the vector $\boldsymbol{\psi} := [\psi_1 \cdots \psi_m]^T$. One can rewrite the amplitude-based cost function in matrix-vector representation as

$$\underset{\mathbf{z} \in \mathbb{R}^n}{\text{minimize}} \quad \ell(\mathbf{z}) := \frac{1}{m} \sum_{i=1}^m \ell_i(\mathbf{z}) = \frac{1}{2m} \|\boldsymbol{\psi} - |\mathbf{A}\mathbf{z}|\|^2 \quad (5)$$

where $\ell_i(\mathbf{z}) := \frac{1}{2}(\psi_i - |\mathbf{a}_i^T \mathbf{z}|)^2$ with the superscript T (\mathcal{H}) denoting (Hermitian) transpose; and with a slight abuse of notation, $|\mathbf{A}\mathbf{z}| := [|\mathbf{a}_1^T \mathbf{z}| \cdots |\mathbf{a}_m^T \mathbf{z}|]^T$. Apart from being nonconvex, $\ell(\mathbf{z})$ is also nondifferentiable, hence challenging the algorithmic design and analysis. In the presence of smoothness or convexity, convergence analysis of iterative algorithms relies either on continuity of the gradient (ordinary gradient methods) [50], or, on the convexity of the objective functional (subgradient methods) [51]. Although subgradient methods have found widespread applicability in nonsmooth optimization, they are limited to the class of convex functions [52, Page 4]. In nonconvex nonsmooth optimization settings, the so-termed *generalized gradient* broadens the scope of the (sub)gradient to the class of *almost everywhere* differentiable functions [53].

Consider a continuous but not necessarily differentiable function $h(\mathbf{z}) \in \mathbb{R}$ defined over an open region $\mathcal{S} \subseteq \mathbb{R}^n$. We then have the following definition.

Definition 1 [54, Definition 1.1]: The generalized gradient of a function h at z , denoted by ∂h , is the convex hull of the set of limits of the form $\lim_{k \rightarrow +\infty} \nabla h(z_k)$, where $z_k \rightarrow z$ as $k \rightarrow +\infty$, i.e.,

$$\partial h(z) := \text{conv} \left\{ \lim_{k \rightarrow +\infty} \nabla h(z_k) : z_k \rightarrow z, z_k \notin \mathcal{G}_\ell \right\}$$

where the symbol ‘conv’ signifies the convex hull of a set, and \mathcal{G}_ℓ denotes the set of points in \mathcal{S} at which h fails to be differentiable.

Having introduced the notion of a generalized gradient, and with t denoting the iteration count, our approach to solving (5) amounts to iteratively refining the initial guess z_0 (returned by the orthogonality-promoting initialization method to be detailed shortly) by means of the ensuing *truncated* generalized gradient iterations

$$z_{t+1} = z_t - \mu_t \partial \ell_{\text{tr}}(z_t). \quad (6)$$

Here, $\mu_t > 0$ is the step size, and the (truncated) generalized gradient $\partial \ell_{\text{tr}}(z_t)$ is given by

$$\partial \ell_{\text{tr}}(z_t) := \frac{1}{m} \sum_{i \in \mathcal{I}_{t+1}} \left(a_i^T z_t - \psi_i \frac{a_i^T z_t}{|a_i^T z_t|} \right) a_i \quad (7)$$

for some index set $\mathcal{I}_{t+1} \subseteq [m] := \{1, 2, \dots, m\}$ to be designed next. The convention $\frac{a_i^T z_t}{|a_i^T z_t|} := 0$ is adopted, if $a_i^T z_t = 0$. It is easy to verify that the update in (6) with a full generalized gradient in (7) monotonically decreases the objective function value in (5).

Any stationary point z^* of $\ell(z)$ can be characterized by the following fixed-point equation [55], [56]

$$A^T \left(Az^* - \psi \odot \frac{Az^*}{|Az^*|} \right) = \mathbf{0} \quad (8)$$

for entry-wise product \odot , which may have many solutions. Clearly, if z^* is a solution, then so is $-z^*$. Furthermore, both solutions/global minimizers x and $-x$ satisfy (8) due to the fact that $Ax - \psi \odot \frac{Ax}{|Ax|} = \mathbf{0}$. Considering any stationary point $z^* \neq \pm x$ that has been adapted such that $\phi(z^*) = 0$, one can write

$$z^* = x + (A^T A)^{-1} A^T \left[\psi \odot \left(\frac{Az^*}{|Az^*|} - \frac{Ax}{|Ax|} \right) \right]. \quad (9)$$

Thus, a necessary condition for $z^* \neq x$ in (9) is $\frac{Az^*}{|Az^*|} \neq \frac{Ax}{|Ax|}$. Expressed differently, there must be sign differences between Az^* and Ax whenever one gets stuck with an undesirable stationary point z^* . Inspired by this observation, it is reasonable to devise algorithms that can detect and separate out the generalized gradient components corresponding to mistakenly estimated signs $\left\{ \frac{a_i^T z_t}{|a_i^T z_t|} \right\}$ along the iterates $\{z_t\}$.

Precisely, if z_t and x lie at different sides of the hyperplane $a_i^T z = 0$, then the sign of $a_i^T z_t$ will be different than that of $a_i^T x$; that is, $\frac{a_i^T x}{|a_i^T x|} \neq \frac{a_i^T z_t}{|a_i^T z_t|}$. Specifically, one can re-write the




Fig. 1. Geometric description of the proposed truncation rule on the i -th gradient component involving $a_i^T x = \psi_i$, where the red dot denotes the solution x and the black one is the origin. Hyperplanes $a_i^T z = \psi_i$ and $a_i^T z = 0$ (of $z \in \mathbb{R}^n$) passing through points $z = x$ and $z = \mathbf{0}$, respectively, are shown.

i -th generalized gradient component as

$$\begin{aligned} \partial \ell_i(z) &= \left(a_i^T z - \psi_i \frac{a_i^T z}{|a_i^T z|} \right) a_i \\ &= \left(a_i^T z - |a_i^T x| \frac{a_i^T x}{|a_i^T x|} \right) a_i + \left(\frac{a_i^T x}{|a_i^T x|} - \frac{a_i^T z}{|a_i^T z|} \right) \psi_i a_i \\ &= a_i a_i^T (z - x) + \left(\frac{a_i^T x}{|a_i^T x|} - \frac{a_i^T z}{|a_i^T z|} \right) \psi_i a_i \\ &= a_i a_i^T h + \underbrace{\left(\frac{a_i^T x}{|a_i^T x|} - \frac{a_i^T z}{|a_i^T z|} \right) \psi_i a_i}_{\triangleq r_i}, \end{aligned} \quad (10)$$

where $h := z - x$. Intuitively, the SLLN asserts that averaging the first term $a_i a_i^T h$ over m instances approaches h , which qualifies it as a desirable search direction. However, certain generalized gradient entries involve erroneously estimated signs of $a_i^T x$; hence, nonzero r_i terms exert a negative influence on the search direction h by dragging the iterate away from x , and they typically have sizable magnitudes as will be further elaborated in Remark 2 shortly.

Figure 1 demonstrates this from a geometric perspective, where the black dot denotes the origin, and the red dot the solution x ; here, $-x$ is omitted for ease of exposition. Assume without loss of generality that the i -th missing sign is positive, i.e., $a_i^T x = \psi_i$. As will be demonstrated in Theorem 1, with high probability, the initial estimate returned by our orthogonality-promoting method obeys $\|h\| \leq \rho \|x\|$ for some sufficiently small constant $\rho > 0$. Therefore, all points lying on or within the circle (or sphere in high-dimensional spaces) in Fig. 1 satisfy $\|h\| \leq \rho \|x\|$. If $a_i^T z = 0$ does not intersect with the circle, then all points within the circle satisfy $\frac{a_i^T z}{|a_i^T z|} = \frac{a_i^T x}{|a_i^T x|}$ qualifying the i -th generalized gradient as a desirable search (descent) direction in (10). If, on the other hand, $a_i^T z = 0$ intersects the circle, then points lying on the same side of $a_i^T z = 0$ with x in Fig. 1 admit correctly estimated signs, while points lying on different sides of $a_i^T z = 0$ with x would have $\frac{a_i^T z}{|a_i^T z|} \neq \frac{a_i^T x}{|a_i^T x|}$. This gives rise to a corrupted search direction in (10), implying that the corresponding generalized gradient component should be eliminated.

Nevertheless, it is difficult or even impossible to check whether the sign of $\mathbf{a}_i^T \mathbf{z}_t$ equals that of $\mathbf{a}_i^T \mathbf{x}$. Fortunately, as demonstrated in Fig. 1, most spurious generalized gradient components (those corrupted by nonzero \mathbf{r}_i terms) hover around the watershed hyperplane $\mathbf{a}_i^T \mathbf{z}_t = 0$. For this reason, TAF includes only those components having \mathbf{z}_t sufficiently away from its watershed, i.e.,

$$\mathcal{I}_{t+1} := \left\{ 1 \leq i \leq m \mid \frac{|\mathbf{a}_i^T \mathbf{z}_t|}{|\mathbf{a}_i^T \mathbf{x}|} \geq \frac{1}{1+\gamma} \right\}, \quad t \geq 0 \quad (11)$$

for an appropriately preselected threshold $\gamma > 0$. To be more specific, the light yellow color-coded area denoted by ξ_i^1 in Fig. 1 signifies the truncation region of \mathbf{z} : if $\mathbf{z} \in \xi_i^1$ satisfies the condition in (11), then the corresponding generalized gradient component $\partial \ell_i(\mathbf{z}; \psi_i)$ will be thrown out. However, the truncation rule may mis-reject certain ‘good’ gradients if \mathbf{z}_t lies in the upper part of ξ_i^1 ; ‘bad’ gradients may be missed as well if \mathbf{z}_t belongs to the spherical cap ξ_i^2 . Fortunately, as we will show in Lemmas 5 and 6, the probabilities of misses and mis-rejections are provably negligible, hence precluding a noticeable influence on the descent direction. Although not perfect, it turns out that such a regularization rule succeeds in detecting and eliminating most corrupted generalized gradient components with high probability, therefore maintaining a well-behaved search direction.

Regarding our gradient truncation or regularization rule in (11), two observations are in order.

Remark 1: The truncation rule in (11) includes only relatively sizable $\mathbf{a}_i^T \mathbf{z}_t$ ’s, hence enforcing the smoothness of the (truncated) objective function $\ell_{\text{tr}}(\mathbf{z}_t)$ at \mathbf{z}_t . Therefore, the truncated generalized gradient $\partial \ell_{\text{tr}}(\mathbf{z})$ employed in (6) and (7) boils down to the ordinary gradient or Wirtinger derivative $\nabla \ell_{\text{tr}}(\mathbf{z}_t)$ in the real- or complex-valued case.

Remark 2: As will be elaborated in (80) and (82), the quantities $(1/m) \sum_{i=1}^m \psi_i$ and $\max_{i \in [m]} \psi_i$ in (10) have magnitudes on the order of $\sqrt{\pi/2} \|\mathbf{x}\|$ and $\sqrt{m} \|\mathbf{x}\|$, respectively. In contrast, Proposition 1 asserts that the first term in (10) obeys $\|\mathbf{a}_i \mathbf{a}_i^T \mathbf{h}\| \approx \|\mathbf{h}\| \leq \rho \|\mathbf{x}\|$ for a sufficiently small $\rho \ll \sqrt{\pi/2}$. Thus, spurious generalized gradient components typically have large magnitudes. It turns out that our gradient regularization rule in (11) also throws out gradient components of large sizes. To see this, for all $\mathbf{z} \in \mathbb{R}^n$ such that $\|\mathbf{h}\| \leq \rho \|\mathbf{x}\|$ in (28), one can re-express

$$\sum_{i=1}^m \partial \ell_i(\mathbf{z}) = \sum_{i=1}^m \underbrace{\left(1 - \frac{|\mathbf{a}_i^T \mathbf{x}|}{|\mathbf{a}_i^T \mathbf{z}|} \right)}_{\triangleq \beta_i} \mathbf{a}_i \mathbf{a}_i^T \mathbf{z} \quad (12)$$

for some weight $\beta_i \in [-\infty, 1)$ assigned to the direction $\mathbf{a}_i \mathbf{a}_i^T \mathbf{z} \approx \mathbf{z}$ due to $\mathbb{E}[\mathbf{a}_i \mathbf{a}_i^T] = \mathbf{I}_n$. Then $\partial \ell_i(\mathbf{z})$ of an excessively large size corresponds to a large $|\mathbf{a}_i^T \mathbf{x}|/|\mathbf{a}_i^T \mathbf{z}|$ in (12), or equivalently a small $|\mathbf{a}_i^T \mathbf{z}|/|\mathbf{a}_i^T \mathbf{x}|$ in (11), thus causing the corresponding $\partial \ell_i(\mathbf{z})$ to be eliminated according to the truncation rule in (11).

Our truncation rule deviates from the intuition behind TWF, which throws away gradient components corresponding to

large-size $\{|\mathbf{a}_i^T \mathbf{z}_t|/|\mathbf{a}_i^T \mathbf{x}|\}$ in (11). As demonstrated by our analysis in Appendix VI-E, it rarely happens that a gradient component having large $|\mathbf{a}_i^T \mathbf{z}_t|/|\mathbf{a}_i^T \mathbf{x}|$ yields an incorrect sign of $\mathbf{a}_i^T \mathbf{x}$ under a sufficiently accurate initialization. Moreover, discarding too many samples (those for which $i \notin \mathcal{I}_{t+1}$ in TWF [6, Section 2.1]) introduces large bias into $(1/m) \sum_{i \in \mathcal{I}_{t+1}} \mathbf{a}_i \mathbf{a}_i^T \mathbf{h}$, so that TWF does not work well when m/n is close to the information-limit of $m/n \approx 2$. In sharp contrast, the motivation and objective of our truncation rule in (11) is to directly sense and eliminate gradient components that involve mistakenly estimated signs with high probability.

To demonstrate the power of TAF, numerical tests comparing all stages of (T)AF and (T)WF will be presented throughout our analysis. The basic test settings used in this paper are described next. For fairness, all pertinent algorithmic parameters involved in all compared schemes were set to their default values. Simulated estimates are averaged over 100 independent Monte Carlo (MC) realizations without mentioning this explicitly each time. Performance of different schemes is evaluated in terms of the relative root mean-square error, i.e.,

$$\text{Relative error} := \frac{\text{dist}(\mathbf{z}, \mathbf{x})}{\|\mathbf{x}\|}, \quad (13)$$

and the success rate among 100 trials, where a success is claimed for a trial if the returned estimate incurs a relative error less than 10^{-5} [6]. Simulated tests under both noiseless and noisy Gaussian models are performed, corresponding to $\psi_i = |\mathbf{a}_i^T \mathbf{x} + \eta_i|$ [28] with $\eta_i = 0$ and $\eta_i \sim \mathcal{N}(0, \sigma^2)$, respectively, with i.i.d. $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ or $\mathbf{a}_i \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n)$.

Numerical comparison depicted in Fig. 2 using the noiseless real-valued Gaussian model suggests that even when starting with the *same truncated spectral initialization*, TAF’s refinement outperforms those of TWF and WF, demonstrating the merits of our gradient update rule over TWF/WF. Furthermore, comparing TAF (gradient iterations in (6)-(7) with truncation in (11) initialized by the truncated spectral estimate) and AF (gradient iterations in (6)-(7) initialized by the truncated spectral estimate) corroborates the power of the truncation rule in (11).

B. Orthogonality-Promoting Initialization Stage

Leveraging the SLLN, spectral initialization methods estimate \mathbf{x} as the (appropriately scaled) leading eigenvector of $\mathbf{Y} := \frac{1}{m} \sum_{i \in \mathcal{T}_0} y_i \mathbf{a}_i \mathbf{a}_i^T$, where \mathcal{T}_0 is an index set accounting for possible data truncation. As asserted in [6], each summand $(\mathbf{a}_i^T \mathbf{x})^2 \mathbf{a}_i \mathbf{a}_i^T$ follows a heavy-tail probability density function lacking a moment generating function. This causes major performance degradation especially when the number of measurements is small. Instead of spectral initializations, we shall take another route to bypass this hurdle. To gain intuition into our initialization, a motivating example is presented first that reveals fundamental characteristics of high-dimensional random vectors.

Fixing any nonzero vector $\mathbf{x} \in \mathbb{R}^n$, generate data $\psi_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|$ using i.i.d. $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq i \leq m$. Evaluate the




Fig. 2. Empirical success rate for WF, TWF, AF, and TAF with the same truncated spectral initialization under the noiseless real-valued Gaussian model.




Fig. 3. Ordered squared normalized inner-product for pairs \mathbf{x} and \mathbf{a}_i , $\forall i \in [m]$ with m/n varying by 2 from 2 to 10, and $n = 1,000$.

following squared normalized inner-product

$$\cos^2 \theta_i := \frac{|\langle \mathbf{a}_i, \mathbf{x} \rangle|^2}{\|\mathbf{a}_i\|^2 \|\mathbf{x}\|^2} = \frac{\psi_i^2}{\|\mathbf{a}_i\|^2 \|\mathbf{x}\|^2}, \quad 1 \leq i \leq m \quad (14)$$

where θ_i is the angle between vectors \mathbf{a}_i and \mathbf{x} . Consider ordering all $\{\cos^2 \theta_i\}$ in an ascending fashion, and collectively denote them as $\boldsymbol{\xi} := [\cos^2 \theta_{[m]} \cdots \cos^2 \theta_{[1]}]^T$ with $\cos^2 \theta_{[1]} \geq \cdots \geq \cos^2 \theta_{[m]}$. Figure 3 plots the ordered entries in $\boldsymbol{\xi}$ for m/n varying by 2 from 2 to 10 with $n = 1,000$. Observe that almost all $\{\mathbf{a}_i\}$ vectors have a squared normalized inner-product with \mathbf{x} smaller than 10^{-2} , while half of the inner-products are less than 10^{-3} , which implies that \mathbf{x} is nearly orthogonal to a large number of \mathbf{a}_i 's.

This example corroborates the folklore that random vectors in high-dimensional spaces are almost always nearly orthogonal to each other [48]. This inspired us to pursue an *orthogonality-promoting initialization method*. Our key idea is to approximate \mathbf{x} by a vector that is most orthogonal to a subset of vectors $\{\mathbf{a}_i\}_{i \in \mathcal{I}_0}$, where \mathcal{I}_0 is an index set with cardinality $|\mathcal{I}_0| < m$ that includes indices of the smallest

squared normalized inner-products $\{\cos^2 \theta_i\}$. Since $\|\mathbf{x}\|$ appears in all inner-products, its exact value does not influence their ordering. Henceforth, we assume with no loss of generality that $\|\mathbf{x}\| = 1$.

Using data $\{(\mathbf{a}_i; \psi_i)\}$, evaluate $\cos^2 \theta_i$ according to (14) for each pair \mathbf{x} and \mathbf{a}_i . Instrumental for the ensuing derivations is noticing from the inherent near-orthogonal property of high-dimensional random vectors that the summation of $\cos^2 \theta_i$ over all indices $i \in \mathcal{I}_0$ should be very small; rigorous justification is deferred to Section V. Therefore, the sum $\sum_{i \in \mathcal{I}_0} \cos^2 \theta_i$ is also small, or according to (14), equivalently,

$$\sum_{i \in \mathcal{I}_0} \frac{|\langle \mathbf{a}_i, \mathbf{x} \rangle|^2}{\|\mathbf{a}_i\|^2 \|\mathbf{x}\|^2} = \frac{\mathbf{x}}{\|\mathbf{x}\|} \left(\sum_{i \in \mathcal{I}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} \right) \frac{\mathbf{x}}{\|\mathbf{x}\|} \quad (15)$$

is small. Therefore, a meaningful approximation of \mathbf{x} can be obtained by minimizing the former with \mathbf{x} replaced by the optimization variable \mathbf{z} , namely

$$\underset{\|\mathbf{z}\|=1}{\text{minimize}} \quad \mathbf{z}^T \left(\frac{1}{|\mathcal{I}_0|} \sum_{i \in \mathcal{I}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} \right) \mathbf{z}. \quad (16)$$

This amounts to finding the smallest eigenvalue and the associated eigenvector of $\mathbf{Y}_0 := \frac{1}{|\mathcal{I}_0|} \sum_{i \in \mathcal{I}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} \succeq \mathbf{0}$ (the symbol \succeq means positive semidefinite). Finding the smallest eigenvalue calls for eigen-decomposition or matrix inversion, each typically requiring computational complexity on the order of $\mathcal{O}(n^3)$. Such a computational burden may be intractable when n grows large. Applying a standard concentration result, we show how the computation can be significantly reduced.

Since $\mathbf{a}_i / \|\mathbf{a}_i\|$ has unit norm and is uniformly distributed on the unit sphere, it is uniformly spherically distributed.² Spherical symmetry implies that $\mathbf{a}_i / \|\mathbf{a}_i\|$ has zero mean and covariance matrix \mathbf{I}_n / n [57]. Appealing again to the SLLN, the sample covariance matrix $\frac{1}{m} \sum_{i=1}^m \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2}$ approaches \mathbf{I}_n / n as m grows. Simple derivations lead to

$$\sum_{i \in \mathcal{I}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} = \sum_{i=1}^m \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} - \sum_{i \in \bar{\mathcal{I}}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} \approx \frac{m}{n} \mathbf{I}_n - \sum_{i \in \bar{\mathcal{I}}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2} \quad (17)$$

where $\bar{\mathcal{I}}_0$ is the complement of \mathcal{I}_0 in the set $[m]$. Define $\mathbf{S} := [\mathbf{a}_1 / \|\mathbf{a}_1\| \cdots \mathbf{a}_m / \|\mathbf{a}_m\|]^T \in \mathbb{R}^{m \times n}$, and form $\bar{\mathbf{S}}_0$ by removing the rows of \mathbf{S} whose indices belong to \mathcal{I}_0 . Seeking the smallest eigenvalue of $\mathbf{Y}_0 = \frac{1}{|\mathcal{I}_0|} \mathbf{S}_0^T \mathbf{S}_0$ then reduces to computing the largest eigenvalue of the matrix

$$\bar{\mathbf{Y}}_0 := \frac{1}{|\bar{\mathcal{I}}_0|} \bar{\mathbf{S}}_0^T \bar{\mathbf{S}}_0, \quad (18)$$

namely,

$$\tilde{\mathbf{z}}_0 := \arg \max_{\|\mathbf{z}\|=1} \mathbf{z}^T \bar{\mathbf{Y}}_0 \mathbf{z} \quad (19)$$

which can be efficiently solved via simple power iterations.

²A random vector $\mathbf{z} \in \mathbb{R}^n$ is said to be spherical (or spherically symmetric) if its distribution does not change under rotations of the coordinate system; that is, the distribution of $\mathbf{P}\mathbf{z}$ coincides with that of \mathbf{z} for any given orthogonal $n \times n$ matrix \mathbf{P} .

When $\|\mathbf{x}\| \neq 1$, the estimate $\tilde{\mathbf{z}}_0$ from (19) is scaled so that its norm matches approximately that of \mathbf{x} , which is estimated as $\sqrt{\frac{1}{m} \sum_{i=1}^m y_i}$, or more accurately $\sqrt{\frac{n \sum_{i=1}^m y_i}{\sum_{i=1}^m \|\mathbf{a}_i\|^2}}$. To motivate these estimates, using the rotational invariance property of normal distributions, it suffices to consider the case where $\mathbf{x} = \|\mathbf{x}\| \mathbf{e}_1$, with \mathbf{e}_1 denoting the first canonical vector of \mathbb{R}^n . Indeed,

$$\begin{aligned} \left\langle \mathbf{a}_i, \frac{\mathbf{x}}{\|\mathbf{x}\|} \right\rangle^2 &= |\langle \mathbf{a}_i, \mathbf{U} \mathbf{e}_1 \rangle|^2 \\ &= \left| \langle \mathbf{U}^T \mathbf{a}_i, \mathbf{e}_1 \rangle \right|^2 \stackrel{d}{=} |\langle \mathbf{a}_i, \mathbf{e}_1 \rangle|^2 \end{aligned} \quad (20)$$

where $\mathbf{U} \in \mathbb{R}^{n \times n}$ is some unitary matrix, and $\stackrel{d}{=}$ means that terms on both sides of the equality have the same distribution. It is then easily verified that

$$\frac{1}{m} \sum_{i=1}^m y_i = \frac{1}{m} \sum_{i=1}^m \mathbf{a}_{i,1}^2 \|\mathbf{x}\|^2 \approx \|\mathbf{x}\|^2, \quad (21)$$

where the last approximation arises from the following concentration result $(1/m) \sum_{i=1}^m \mathbf{a}_{i,1}^2 \approx \mathbb{E}[\mathbf{a}_{i,1}^2] = 1$ using again the SLLN. Regarding the second estimate, one can rewrite its square as

$$\frac{n \sum_{i=1}^m y_i}{\sum_{i=1}^m \|\mathbf{a}_i\|^2} = \frac{1}{m} \sum_{i=1}^m y_i \cdot \frac{n}{(1/m) \cdot \sum_{i=1}^m \|\mathbf{a}_i\|^2}. \quad (22)$$

It is clear from (21) that the first term on the right hand side of (22) approximates $\|\mathbf{x}\|^2$. The second term approaches 1 because the denominator $(1/m) \cdot \sum_{i=1}^m \|\mathbf{a}_i\|^2 \approx n$ appealing to the SLLN again and the fact that $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. For simplicity, we choose to work with the first norm estimate

$$\mathbf{z}_0 = \sqrt{\frac{\sum_{i=1}^m y_i}{m}} \tilde{\mathbf{z}}_0. \quad (23)$$

It is worth highlighting that, compared to the matrix $\mathbf{Y} := \frac{1}{m} \sum_{i \in \mathcal{I}_0} y_i \mathbf{a}_i \mathbf{a}_i^T$ used in spectral methods, our constructed matrix $\tilde{\mathbf{Y}}_0$ in (18) does not depend on the observed data $\{y_i\}$ explicitly; the dependence is only through the choice of the index set \mathcal{I}_0 . The novel orthogonality-promoting initialization thus enjoys two advantages over its spectral alternatives: a1) it does not suffer from heavy-tails of the fourth-order moments of Gaussian $\{\mathbf{a}_i\}$ vectors common in spectral initialization schemes; and, a2) it is less sensitive to noisy and corrupted data.

Figure 4 compares three different initialization schemes including spectral initialization [18], [28], truncated spectral initialization [6], and the proposed orthogonality-promoting initialization. The relative error of their returned initial estimates versus the measurement/unknown ratio m/n is depicted under the noiseless and noisy real-valued Gaussian models, where $\mathbf{x} \in \mathbb{R}^{1,000}$ was randomly generated and m/n increases by 2 from 2 to 20. Clearly, all schemes enjoy improved performance as m/n increases in both noiseless and noisy settings. The orthogonality-promoting initialization achieves consistently superior performance over its competing spectral alternatives under both noiseless and noisy Gaussian data. Interestingly, the spectral and truncated spectral schemes

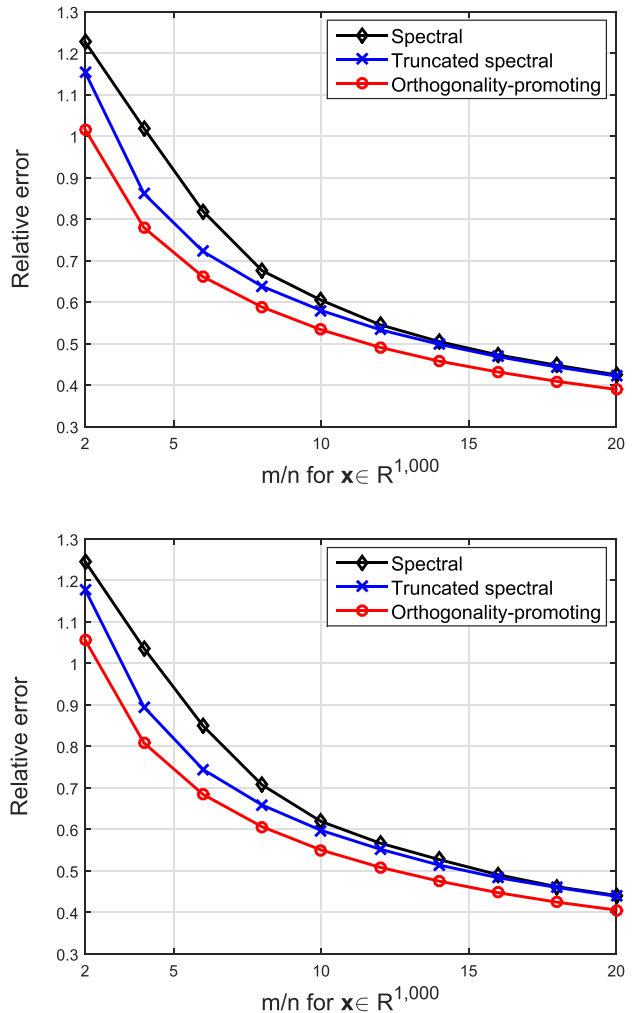


Fig. 4. Relative error of initial estimates versus m/n for: i) the spectral method [18]; ii) the truncated spectral method [6]; and iii) our orthogonality-promoting method with $n = 1,000$, and m/n varying by 2 from 2 to 20. Top: Noiseless real-valued Gaussian model with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and $\eta_i = 0$. Bottom: Noisy real-valued Gaussian model with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and $\sigma^2 = 0.2^2 \|\mathbf{x}\|^2$.

exhibit similar performance when m/n becomes sufficiently large (e.g., $m/n \geq 14$ in the noiseless setup or $m/n \geq 16$ in the noisy one). This confirms that the truncation helps only if m/n is relatively small. Indeed, the truncation discards measurements of excessively large or small sizes emerging from the heavy tails of the data distribution. Hence, its advantage over the non-truncated spectral initialization diminishes as the number of measurements increases, which gradually straightens out the heavy tails.

III. MAIN RESULTS

The TAF algorithm is summarized in Algorithm 1. Default values are set for pertinent algorithmic parameters. Assuming independent data samples $\{(\mathbf{a}_i; \psi_i)\}$ drawn from the noiseless real-valued Gaussian model, the following result establishes the theoretical performance of TAF.

Theorem 1 (Exact recovery): Let $\mathbf{x} \in \mathbb{R}^n$ be an arbitrary signal vector, and consider (noise-free) measurements $\psi_i = |\mathbf{a}_i^T \mathbf{x}|$, in which $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq i \leq m$. Then with

Algorithm 1 Truncated Amplitude Flow (TAF)

- 1: **Input:** Amplitude data $\{\psi_i := |\langle \mathbf{a}_i, \mathbf{x} \rangle|\}_{i=1}^m$ and design vectors $\{\mathbf{a}_i\}_{i=1}^m$; the maximum number of iterations T ; by default, take constant step sizes $\mu = 0.6$ or 1 for the real- or complex-valued models, truncation thresholds $|\bar{\mathcal{I}}_0| = \lceil \frac{1}{6}m \rceil$, and $\gamma = 0.7$.
- 2: **Set** $\bar{\mathcal{I}}_0$ as the set of indices corresponding to the $|\bar{\mathcal{I}}_0|$ largest values of $\{\psi_i / \|\mathbf{a}_i\|\}$.
- 3: **Initialize** \mathbf{z}_0 to $\sqrt{\frac{\sum_{i=1}^m \psi_i^2}{m}} \bar{\mathbf{z}}_0$, where $\bar{\mathbf{z}}_0$ is the normalized leading eigenvector of $\bar{\mathbf{Y}}_0 := \frac{1}{|\bar{\mathcal{I}}_0|} \sum_{i \in \bar{\mathcal{I}}_0} \frac{\mathbf{a}_i \mathbf{a}_i^T}{\|\mathbf{a}_i\|^2}$.
- 4: **Loop:** for $t = 0$ to $T - 1$

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \frac{\mu}{m} \sum_{i \in \mathcal{I}_{t+1}} \left(\mathbf{a}_i^T \mathbf{z}_t - \psi_i \frac{\mathbf{a}_i^T \mathbf{z}_t}{|\mathbf{a}_i^T \mathbf{z}_t|} \right) \mathbf{a}_i$$

$$\text{where } \mathcal{I}_{t+1} := \left\{ 1 \leq i \leq m \mid |\mathbf{a}_i^T \mathbf{z}_t| \geq \frac{1}{1+\gamma} \psi_i \right\}.$$

- 5: **Output:** \mathbf{z}_T .

probability at least $1 - (m+5)e^{-n/2} - e^{-c_0 m} - 1/n^2$ for some universal constant $c_0 > 0$, the initialization \mathbf{z}_0 returned by the orthogonality-promoting method in Algorithm 1 satisfies

$$\text{dist}(\mathbf{z}_0, \mathbf{x}) \leq \rho \|\mathbf{x}\| \quad (24)$$

with $\rho = 1/10$ (or any sufficiently small positive constant), provided that $m \geq c_1 |\bar{\mathcal{I}}_0| \geq c_2 n$ for some numerical constants $c_1, c_2 > 0$, and sufficiently large n . Furthermore, choosing a constant step size $\mu \leq \mu_0$ along with a truncation level $\gamma \geq 1/2$, and starting from any initial guess \mathbf{z}_0 satisfying (24), successive estimates of the TAF solver (tabulated in Algorithm 1) obey

$$\text{dist}(\mathbf{z}_t, \mathbf{x}) \leq \rho (1-v)^t \|\mathbf{x}\|, \quad t = 0, 1, 2, \dots \quad (25)$$

for some $0 < v < 1$, which holds with probability exceeding $1 - (m+5)e^{-n/2} - 8e^{-c_0 m} - 1/n^2$.

Typical parameter values for TAF in Algorithm 1 are $\mu = 0.6$, and $\gamma = 0.7$. The proof of Theorem 1 is relegated to Section V. Theorem 1 asserts that: i) TAF reconstructs the solution \mathbf{x} exactly as soon as the number of equations is about the number of unknowns, which is theoretically order optimal. Our numerical tests demonstrate that for the real-valued Gaussian model, TAF achieves a success rate of 100% when m/n is as small as 3, which is slightly larger than the information limit of $m/n = 2$ (Recall that $m \geq 2n - 1$ is necessary for the uniqueness.) This is a significant reduction in the sample complexity ratio, which is 5 for TWF and 7 for WF. Surprisingly, TAF also enjoys a success rate of over 50% when m/n is the information limit 2, which has not yet been presented for any existing algorithms. See further discussion in Section IV; and, ii) TAF converges exponentially fast with convergence rate independent of the dimension n . Specifically, TAF requires at most $\mathcal{O}(\log(1/\epsilon))$ iterations to achieve any given solution accuracy $\epsilon > 0$ (a.k.a., $\text{dist}(\mathbf{z}_t, \mathbf{x}) \leq \epsilon \|\mathbf{x}\|$), with iteration cost $\mathcal{O}(mn)$. Since the truncation takes time

²The symbol $\lceil \cdot \rceil$ is the ceiling operation returning the smallest integer greater than or equal to the given number.

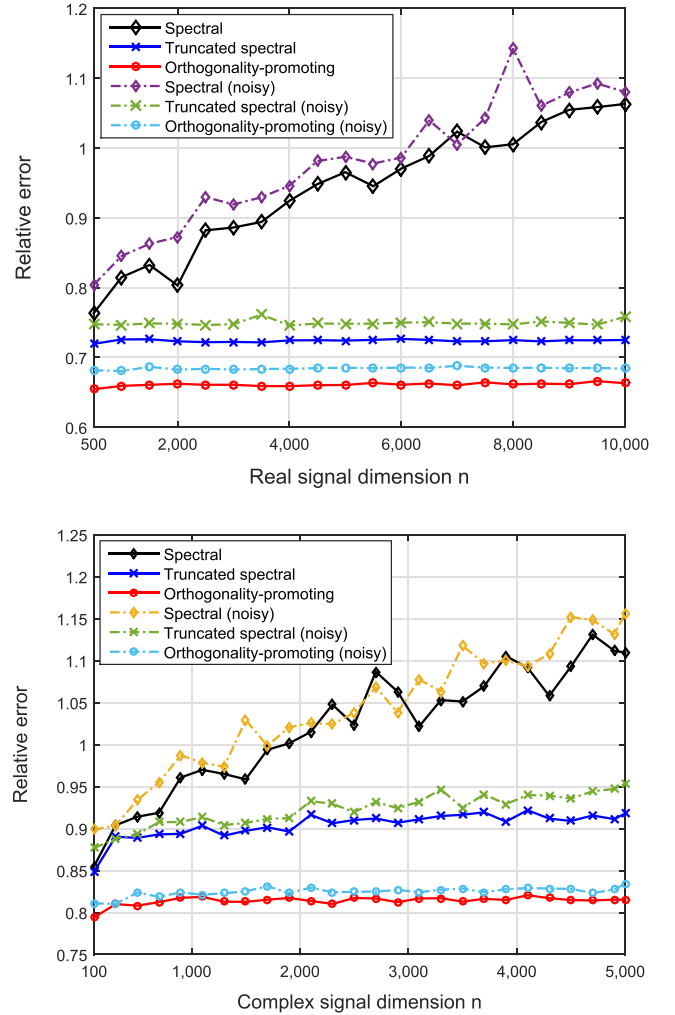


Fig. 5. Average relative error of estimates obtained from 100 MC trials using: i) the spectral method [18], [28]; ii) the truncated spectral method [6]; and iii) the proposed orthogonality-promoting method on noise-free (solid lines) and noisy (dotted lines) instances with $m/n = 6$, and n varying from 500/100 to 10,000/5,000 for real-/complex-valued vectors. Top: Real-valued Gaussian model with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and $\sigma^2 = 0.2^2 \|\mathbf{x}\|^2$. Bottom: Complex-valued Gaussian model with $\mathbf{x} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n)$, $\mathbf{a}_i \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n)$, and $\sigma^2 = 0.2^2 \|\mathbf{x}\|^2$.

on the order of $\mathcal{O}(m)$, the computational burden of TAF per iteration is dominated by the evaluation of the gradient components. The latter involves two matrix-vector multiplications that are computable in $\mathcal{O}(mn)$ flops, namely, $\mathbf{A}\mathbf{z}_t$ yields \mathbf{u}_t , and $\mathbf{A}^T \mathbf{v}_t$ the gradient, where $\mathbf{v}_t := \mathbf{u}_t - \psi \odot \frac{\mathbf{u}_t}{|\mathbf{u}_t|}$. Hence, the total running time of TAF is $\mathcal{O}(mn \log(1/\epsilon))$, which is proportional to the time taken to read the data $\mathcal{O}(mn)$.

In the noisy setting, TAF is stable under additive noise. To be specific, consider the amplitude-based noisy data model $\psi_i = |\mathbf{a}_i^T \mathbf{x}| + \eta_i$. It can be shown that the truncated amplitude flow estimates in Algorithm 1 satisfy

$$\text{dist}(\mathbf{z}_t, \mathbf{x}) \lesssim (1-v)^t \|\mathbf{x}\| + \frac{1}{\sqrt{m}} \|\boldsymbol{\eta}\|, \quad t = 0, 1, \dots \quad (26)$$

with high probability for all $\mathbf{x} \in \mathbb{R}^n$, provided that $m \geq c_1 |\bar{\mathcal{I}}_0| \geq c_2 n$ for sufficiently large n and the noise is bounded $\|\boldsymbol{\eta}\|_\infty \leq c_3 \|\mathbf{x}\|$ with $\boldsymbol{\eta} := [\eta_1 \dots \eta_n]^T$, where $0 < v < 1$,




Fig. 6. Relative initialization error of the initialization solving the minimum eigenvalue problem in (16) via the Lanczos method and by solving the maximum eigenvalue problem in (19).

and $c_1, c_2, c_3 > 0$ are some universal constants. The proof can be directly adapted from those of Theorem 1 above and [6, Theorem 2].

IV. SIMULATED TESTS

In this section, we provide additional numerical tests evaluating performance of the proposed scheme relative to (T)WF³ and AF. The initial estimate was found based on 50 power iterations, and was subsequently refined by $T = 1,000$ gradient-type iterations in each scheme. The Matlab implementations of TAF are available at <https://gangwg.github.io/TAF/> for reproducibility.

Top panel in Fig. 5 presents the average relative error of three initialization methods on a series of noiseless/noisy real-valued Gaussian problems with $m/n = 6$ fixed, and n varying from 500 to 10^4 , while those for the corresponding complex-valued Gaussian instances are shown in the bottom panel. Clearly, the proposed initialization method returns more accurate and robust estimates than the spectral ones. Under the same condition for the real-valued Gaussian model, Fig. 6 compares the initialization implemented in Algorithm 1 obtained by solving the maximum eigenvalue problem in (19) with the one obtained by tackling the minimum eigenvalue problem in (16) via the Lanczos method [58]. When the number of equations is relatively small (less than about $3n$), the former performs better than the latter. Interestingly though, the latter works remarkably well and almost halves the error incurred by the implemented initialization of Algorithm 1 as soon as the number of equations becomes larger than 4.

To demonstrate the power of TAF, Fig. 7 plots the relative error of recovering a real-valued signal in logarithmic scale versus the iteration count under the information-limit of $m = 2n - 1$ noiseless i.i.d. Gaussian measurements [1]. In this case, since the returned initial estimate is relatively far from the




Fig. 7. Relative error versus iteration for TAF for a noiseless real-valued Gaussian model under the information-limit of $m = 2n - 1$.

optimal solution (see Fig. 4), TAF converges slowly for the first 200 iterations or so due to elimination of a significant amount of ‘bad’ generalized gradient components (corrupted by mistakenly estimated signs). As the iterate gets more accurate and lands within a small-size neighborhood of \mathbf{x} , TAF converges exponentially fast to the globally optimal solution. It is worth emphasizing that no existing method succeeds in this case. Figure 8 compares the empirical success rate of three schemes under both real-valued and complex-valued Gaussian models with $n = 10^3$ and m/n varying by 0.1 from 1 to 7, where a success is claimed if the estimate has a relative error less than 10^{-5} . For real-valued vectors, TAF achieves a success rate of over 50% when $m/n = 2$, and guarantees perfect recovery from about $3n$ measurements; while for complex-valued ones, TAF enjoys a success rate of 95% when $m/n = 3.4$, and ensures perfect recovery from about $4.5n$ measurements.

To demonstrate the stability of TAF, the relative mean-squared error (MSE)

$$\text{Relative MSE} := \frac{\text{dist}^2(\mathbf{z}_T, \mathbf{x})}{\|\mathbf{x}\|^2}$$

as a function of the signal-to-noise ratio (SNR) is plotted for different m/n values. We consider the noisy model $\psi_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle| + \eta_i$ with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{1,000})$ and real-valued independent Gaussian sensing vectors $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{1,000})$, in which m/n takes values $\{6, 8, 10\}$, and the SNR in dB, given by

$$\text{SNR} := 10 \log_{10} \frac{\sum_{i=1}^m |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2}{\sum_{i=1}^m \eta_i^2}$$

is varied from 10 dB to 50 dB. Averaging over 100 independent trials, Fig. 9 demonstrates that the relative MSE for all m/n values scales inversely proportional to SNR, hence justifying the stability of TAF under bounded additive noise.

The next experiment evaluates the efficacy of the proposed initialization method, simulating all schemes initialized by the truncated spectral initial estimate [6] and the orthogonality-promoting initial estimate. Evidently, all algorithms except

³Matlab codes directly downloaded from the authors’ websites: <http://statweb.stanford.edu/~candes/TWF/algorithm.html>; <http://www-bcf.usc.edu/~soltanol/WFcode.html>.

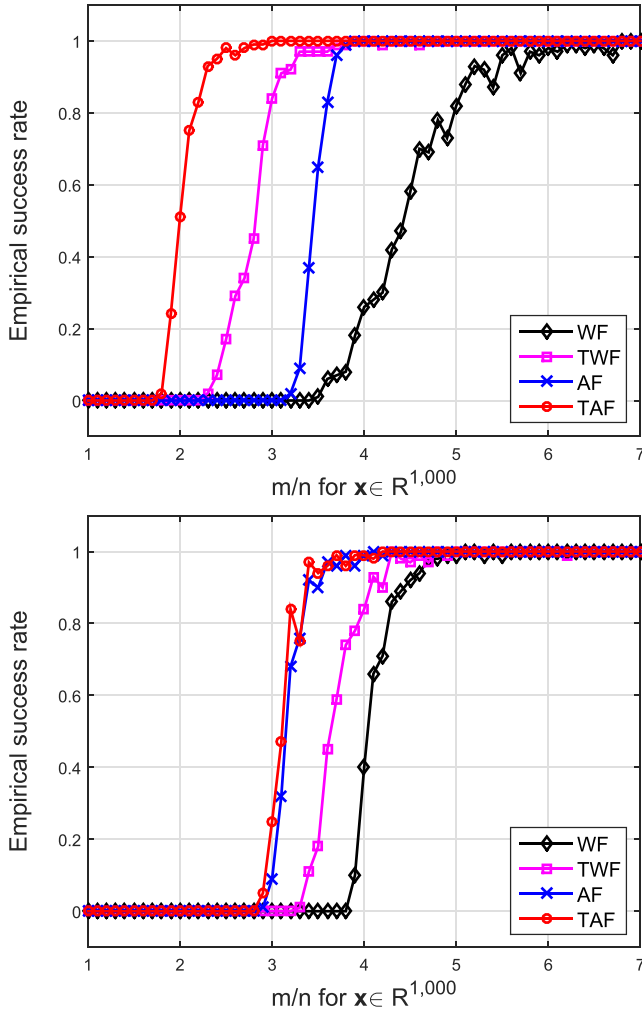


Fig. 8. Empirical success rate for WF, TWF, AF, and TAF with $n = 1,000$ and m/n varying by 0.1 from 1 to 7. Top: Noiseless real-valued Gaussian model with $x \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$; Bottom: Noiseless complex-valued Gaussian model with $x \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n)$ and $\mathbf{a}_i \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_n)$.




Fig. 9. Relative MSE versus SNR for TAF when ψ_i 's follow the amplitude-based noisy data model.

WF admit a significant performance improvement when initialized by the proposed orthogonality-promoting initialization relative to the truncated spectral initialization. Nevertheless,




Fig. 10. Empirical success rate for WF, TWF, AF, and TAF initialized by the truncated spectral and the orthogonality-promoting initializations with $n = 1,000$ and m/n varying by 0.1 from 1 to 7.

TAF with our developed orthogonality-promoting initialization enjoys superior performance over all simulated approaches.

Finally, to examine the effectiveness and scalability of TAF in real-world conditions, we simulate recovery of the Milky Way Galaxy image⁴ $X \in \mathbb{R}^{1080 \times 1920 \times 3}$ shown in Fig. 11. The first two indices encode the pixel locations, and the third the RGB (red, green, blue) color bands. Consider the coded diffraction pattern (CDP) measurements with random masks [6], [18], [27]. Letting $x \in \mathbb{R}^n$ be a vectorization of a certain band of X and postulating a number K of random masks, one can further write

$$\psi^{(k)} = |FD^{(k)}x|, \quad 1 \leq k \leq K, \quad (27)$$

where F denotes the $n \times n$ discrete Fourier transform matrix, and $D^{(k)}$ is a diagonal matrix holding entries sampled uniformly at random from $\{1, -1, j, -j\}$ (phase delays) on its diagonal, with j denoting the imaginary unit. Each $D^{(k)}$ represents a random mask placed after the object [27]. With $K = 6$ masks implemented in our experiment, the total number of quadratic measurements is $m = 6n$. Every algorithm was run independently on each of the three bands. A number 100 of power iterations were used to obtain an initialization, which was refined by 100 gradient-type iterations. The relative errors after our orthogonality-promoting initialization and after 100 TAF iterations are 0.6807 and 9.8631×10^{-5} , respectively, and the recovered images are displayed in Fig. 11. In sharp contrast, TWF reconstructs images of corresponding relative errors 1.3801 and 1.3409, which are far away from the ground truth.

Regarding running times in all performed experiments, TAF converges slightly faster than TWF, while both are markedly faster than WF. All experiments were implemented using MATLAB on an Intel CPU @ 3.4 GHz (32 GB RAM) computer.

⁴Downloaded from <http://pics-about-space.com/milky-way-galaxy>.



Fig. 11. The recovered Milky Way Galaxy images after i) truncated spectral initialization (top); ii) orthogonality-promoting initialization (middle); and iii) 100 TAF gradient iterations refining the orthogonality-promoting initialization (bottom).

V. PROOFS

This section presents the main ideas behind the proof of Theorem 1, and establishes a few necessary lemmas. Technical details are deferred to the Appendix. Relative to WF and TWF, our objective function involves nonsmoothness and nonconvexity, rendering the proof of exact recovery of TAF nontrivial. In addition, our initialization method starts from a rather different perspective than spectral alternatives, so that the tools involved in proving performance of our initialization deviate from those of spectral methods [6], [18], [28]. Part of our proof is adapted from [6], [18], and [55].

The proof of Theorem 1 consists of two parts: Section V-A justifies the performance of the proposed orthogonality-promoting initialization, which essentially achieves any given constant relative error as soon as the number of equations is on the order of the number of unknowns, namely, $m \asymp n$.⁵

⁵The notations $\phi(n) = \mathcal{O}(g(n))$ or $\phi(n) \gtrsim g(n)$ (respectively, $\phi(n) \lesssim g(n)$) means there exists a numerical constant $c > 0$ such that $\phi(n) \leq cg(n)$, while $\phi(n) \asymp g(n)$ means $\phi(n)$ and $g(n)$ are orderwise equivalent.

Section V-B demonstrates theoretical convergence of TAF to the solution of the quadratic system in (1) at a geometric rate provided that the initial estimate has a sufficiently small constant relative error as in (24). The two stages of TAF can be performed independently, meaning that better initialization methods, if available, could be adopted to initialize our truncated generalized gradient iterations; likewise, our initialization may be applied to other iterative optimization algorithms.

A. Constant Relative Error by Orthogonality-Promoting Initialization

This section concentrates on proving guaranteed performance of the proposed orthogonality-promoting initialization method, as asserted in the following proposition. An alternative approach may be found in [59].

Proposition 1: Fix $\mathbf{x} \in \mathbb{R}^n$ arbitrarily, and consider the noiseless case $\psi_i = |\mathbf{a}_i^\top \mathbf{x}|$, where $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq i \leq m$. Then with probability at least $1 - (m+5)e^{-n/2} - e^{-c_0 m}$

$1/n^2$ for some universal constant $c_0 > 0$, the initialization \mathbf{z}_0 returned by the orthogonality-promoting method satisfies

$$\text{dist}(\mathbf{z}_0, \mathbf{x}) \leq \rho \|\mathbf{x}\| \quad (28)$$

for $\rho = 1/10$ or any positive constant, with the proviso that $m \geq c_1 |\bar{\mathcal{I}}_0| \geq c_2 n$ for some numerical constants $c_1, c_2 > 0$ and sufficiently large n .

Due to homogeneity in (28), it suffices to consider the case $\|\mathbf{x}\| = 1$. Assume for the moment that $\|\mathbf{x}\| = 1$ is known and \mathbf{z}_0 has been scaled such that $\|\mathbf{z}_0\| = 1$ in (23). The error between the employed \mathbf{x} 's norm estimate $\sqrt{\frac{1}{m} \sum_{i=1}^m y_i}$ and the unknown norm $\|\mathbf{x}\| = 1$ will be accounted for at the end of this section. Instrumental in proving Proposition 1 is the following result, whose proof is provided in Appendix VI-A.

Lemma 1: Consider the noiseless data $\psi_i = |\mathbf{a}_i^T \mathbf{x}|$, where $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq i \leq m$. For any unit vector $\mathbf{x} \in \mathbb{R}^n$, there exists a vector $\mathbf{u} \in \mathbb{R}^n$ with $\mathbf{u}^T \mathbf{x} = 0$ and $\|\mathbf{u}\| = 1$ such that

$$\frac{1}{2} \left\| \mathbf{x} \mathbf{x}^T - \mathbf{z}_0 \mathbf{z}_0^T \right\|_F^2 \leq \frac{\|\bar{\mathbf{S}}_0 \mathbf{u}\|^2}{\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2} \quad (29)$$

for $\mathbf{z}_0 = \tilde{\mathbf{z}}_0$, where the unit vector $\tilde{\mathbf{z}}_0$ is given in (19), and $\bar{\mathbf{S}}_0$ is formed by removing the rows of $\mathbf{S} := [\mathbf{a}_1 / \|\mathbf{a}_1\| \cdots \mathbf{a}_m / \|\mathbf{a}_m\|]^T \in \mathbb{R}^{m \times n}$ if their indices do not belong to the set $\bar{\mathcal{I}}_0$ specified in Algorithm 1.

We now turn to prove Proposition 1. The first step consists in upper-bounding the term on the right-hand-side of (29). Specifically, its numerator is upper bounded, and the denominator lower bounded, as summarized in Lemma 2 and Lemma 3 next; their proofs are provided in Appendix VI-B and Appendix VI-C, respectively.

Lemma 2: In the setup of Lemma 1, if $|\bar{\mathcal{I}}_0| \geq c'_1 n$, then

$$\|\bar{\mathbf{S}}_0 \mathbf{u}\|^2 \leq 1.01 |\bar{\mathcal{I}}_0| / n \quad (30)$$

holds with probability at least $1 - 2e^{-c_K n}$, where c'_2 and c_K are some universal constants.

Lemma 3: In the setup of Lemma 1, the following holds with probability at least $1 - (m+1)e^{-n/2} - e^{-c_0 m} - 1/n^2$,

$$\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 \geq \frac{0.99 |\bar{\mathcal{I}}_0|}{2.3n} [1 + \log(m/|\bar{\mathcal{I}}_0|)] \quad (31)$$

provided that $|\bar{\mathcal{I}}_0| \geq c'_1 n$, $m \geq c'_2 |\bar{\mathcal{I}}_0|$, and $m \geq c'_3 n$ for some absolute constants $c'_1, c'_2, c'_3 > 0$, and sufficiently large n .

Leveraging the upper and lower bounds in (30) and (31), one arrives at

$$\frac{\|\bar{\mathbf{S}}_0 \mathbf{u}\|^2}{\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2} \leq \frac{2.4}{1 + \log(m/|\bar{\mathcal{I}}_0|)} \triangleq \kappa \quad (32)$$

which holds with probability at least $1 - (m+3)e^{-n/2} - e^{-c_0 m} - 1/n^2$, assuming that $m \geq c'_1 |\bar{\mathcal{I}}_0|$, and $m \geq c'_2 n$, $|\bar{\mathcal{I}}_0| \geq c'_3 n$ for some absolute constants $c'_1, c'_2, c'_3 > 0$, and sufficiently large n .

The bound κ in (32) is meaningful only when the ratio $\log(m/|\bar{\mathcal{I}}_0|) > 1.4$, i.e., $m/|\bar{\mathcal{I}}_0| > 4$, because the left hand side is expressible in terms of $\sin^2 \theta$, and therefore, enjoys a trivial upper bound of 1. Henceforth, we will assume $m/|\bar{\mathcal{I}}_0| > 4$.

Empirically, $\lfloor m/|\bar{\mathcal{I}}_0| \rfloor = 6$, or equivalently $|\bar{\mathcal{I}}_0| = \lfloor \frac{1}{6} m \rfloor$ in Algorithm 1 works well when m/n is relatively small. Note further that the bound κ can be made arbitrarily small by letting $m/|\bar{\mathcal{I}}_0|$ be large enough. Without any loss of generality, let us take $\kappa := 0.001$. An additional step leads to the wanted bound on the distance between $\tilde{\mathbf{z}}_0$ and \mathbf{x} ; similar arguments are found in [18, Section 7.8]. Recall that

$$|\mathbf{x}^T \tilde{\mathbf{z}}_0|^2 = \cos^2 \theta = 1 - \sin^2 \theta \geq 1 - \kappa. \quad (33)$$

Therefore,

$$\begin{aligned} \text{dist}^2(\tilde{\mathbf{z}}_0, \mathbf{x}) &\leq \|\tilde{\mathbf{z}}_0\|^2 + \|\mathbf{x}\|^2 - 2|\mathbf{x}^T \tilde{\mathbf{z}}_0| \\ &\leq (2 - 2\sqrt{1 - \kappa}) \|\mathbf{x}\|^2 \\ &\approx \kappa \|\mathbf{x}\|^2. \end{aligned} \quad (34)$$

Coming back to the case in which $\|\mathbf{x}\|$ is unknown stated prior to Lemma 1, the unit eigenvector $\tilde{\mathbf{z}}_0$ is scaled by an estimate of $\|\mathbf{x}\|$ to yield the initial guess $\mathbf{z}_0 = \sqrt{\frac{1}{m} \sum_{i=1}^m y_i} \tilde{\mathbf{z}}_0$. Using the results in [19, Lemma 7.8], the following holds with high probability

$$\|\mathbf{z}_0 - \tilde{\mathbf{z}}_0\| = \|\|\mathbf{z}_0\| - 1\| \leq (1/20) \|\mathbf{x}\|. \quad (35)$$

Summarizing the two inequalities, we conclude that

$$\text{dist}(\mathbf{z}_0, \mathbf{x}) \leq \|\mathbf{z}_0 - \tilde{\mathbf{z}}_0\| + \text{dist}(\tilde{\mathbf{z}}_0, \mathbf{x}) \leq (1/10) \|\mathbf{x}\|. \quad (36)$$

The initialization thus obeys $\text{dist}(\mathbf{z}_0, \mathbf{x})/\|\mathbf{x}\| \leq 1/10$ for any $\mathbf{x} \in \mathbb{R}^n$ with high probability provided that $m \geq c_1 |\bar{\mathcal{I}}_0| \geq c_2 n$ holds for some universal constants $c_1, c_2 > 0$ and a sufficiently large n .

B. Exact Recovery From Noiseless Data

We now prove that with accurate enough initial estimates, TAF converges at a geometric rate to \mathbf{x} with high probability (i.e., the second part of Theorem 1). To be specific, with initialization obeying (28) in Proposition 1, TAF reconstructs the solution \mathbf{x} exactly in linear time. To start, it suffices to demonstrate that the TAF's update rule (i.e., Step 4 in Algorithm 1) is locally contractive within a sufficiently small neighborhood of \mathbf{x} , as asserted in the following proposition.

Proposition 2 (Local error contraction): Consider the noise-free measurements $\psi_i = |\mathbf{a}_i^T \mathbf{x}|$ with i.i.d. Gaussian design vectors $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq i \leq m$, and fix any $\gamma \geq 1/2$. There exist universal constants $c_0, c_1 > 0$ and $0 < \nu < 1$ such that with probability at least $1 - 7e^{-c_0 m}$, the following holds

$$\text{dist}^2\left(\mathbf{z} - \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{x}\right) \leq (1 - \nu) \text{dist}^2(\mathbf{z}, \mathbf{x}) \quad (37)$$

for all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ obeying (28) with the proviso that $m \geq c_1 n$ and that the constant step size μ satisfies $0 < \mu \leq \mu_0$ for some $\mu_0 > 0$.

Proposition 2 demonstrates that the distance of TAF's successive iterates to \mathbf{x} is monotonically decreasing once the algorithm enters a small-size neighborhood around \mathbf{x} . This neighborhood is commonly referred to as the *basin of attraction*; see further discussions in [6], [19], [33], [37], and [39]. In other words, as soon as one lands within the basin

of attraction, TAF's iterates remain in this region and will be attracted to \mathbf{x} exponentially fast. To substantiate Proposition 2, recall the *local regularity condition*, which was first developed in [18] and plays a fundamental role in establishing linear convergence to global optimum of nonconvex optimization approaches such as WF and TWF [6], [18], [30], [32].

Consider the update rule of TAF

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}_t), \quad t = 0, 1, 2, \dots \quad (38)$$

where the truncated gradient $\nabla \ell_{\text{tr}}(\mathbf{z}_t)$ (as carefully justified in Remark 1) evaluated at some point $\mathbf{z}_t \in \mathbb{R}^n$ is given by

$$\frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}_t) \triangleq \frac{1}{m} \sum_{i \in \mathcal{I}} \left(\mathbf{a}_i^T \mathbf{z}_t - \psi_i \frac{\mathbf{a}_i^T \mathbf{z}_t}{|\mathbf{a}_i^T \mathbf{z}_t|} \right) \mathbf{a}_i.$$

The truncated gradient $\nabla \ell_{\text{tr}}(\mathbf{z})$ is said to satisfy the local regularity condition, or LRC(μ, λ, ϵ) for some constant $\lambda > 0$, provided that

$$\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle \geq \frac{\mu}{2} \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2 + \frac{\lambda}{2} \|\mathbf{h}\|^2 \quad (39)$$

holds for all $\mathbf{z} \in \mathbb{R}^n$ such that $\|\mathbf{h}\| = \|\mathbf{z} - \mathbf{x}\| \leq \epsilon \|\mathbf{x}\|$ for some constant $0 < \epsilon < 1$, where the ball $\|\mathbf{z} - \mathbf{x}\| \leq \epsilon \|\mathbf{x}\|$ is the so-called *basin of attraction*. Simple linear algebra along with the regularity condition in (39) leads to

$$\begin{aligned} \text{dist}^2 \left(\mathbf{z} - \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{x} \right) &= \left\| \mathbf{z} - \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) - \mathbf{x} \right\|^2 \\ &= \|\mathbf{h}\|^2 - 2\mu \left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle + \left\| \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2 \\ &\leq \|\mathbf{h}\|^2 - 2\mu \left(\frac{\mu}{2} \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2 + \frac{\lambda}{2} \|\mathbf{h}\|^2 \right) \\ &\quad + \left\| \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2 \\ &= (1 - \lambda\mu) \|\mathbf{h}\|^2 = (1 - \lambda\mu) \text{dist}^2(\mathbf{z}, \mathbf{x}) \end{aligned} \quad (40)$$

for all \mathbf{z} obeying $\|\mathbf{h}\| \leq \epsilon \|\mathbf{x}\|$. Evidently, if the LRC(μ, λ, ϵ) is proved for TAF, then (37) follows upon letting $\nu := \lambda\mu$.

1) *Proof of the Local Regularity Condition in (39)*: By definition, justifying the local regularity condition in (39) entails controlling the norm of the truncated gradient $\frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z})$, i.e., bounding the last term in (40). Recall that

$$\frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) = \frac{1}{m} \sum_{i \in \mathcal{I}} \left(\mathbf{a}_i^T \mathbf{z} - \psi_i \frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} \right) \mathbf{a}_i \triangleq \frac{1}{m} \mathbf{A} \mathbf{v} \quad (42)$$

where $\mathcal{I} := \{1 \leq i \leq m \mid |\mathbf{a}_i^T \mathbf{z}| \geq |\mathbf{a}_i^T \mathbf{x}| / (1 + \gamma)\}$, and $\mathbf{v} := [v_1 \ \dots \ v_m]^T \in \mathbb{R}^m$ with $v_i := \frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} (|\mathbf{a}_i^T \mathbf{z}| - \psi_i) \mathbb{1}_{\{|\mathbf{a}_i^T \mathbf{z}| \geq |\mathbf{a}_i^T \mathbf{x}| / (1 + \gamma)\}}$. Now, consider

$$\begin{aligned} |v_i|^2 &= \left| \left(|\mathbf{a}_i^T \mathbf{z}| - |\mathbf{a}_i^T \mathbf{x}| \right) \mathbb{1}_{\{|\mathbf{a}_i^T \mathbf{z}| \geq |\mathbf{a}_i^T \mathbf{x}| / (1 + \gamma)\}} \right|^2 \\ &\leq \left| |\mathbf{a}_i^T \mathbf{z}| - |\mathbf{a}_i^T \mathbf{x}| \right|^2 \leq \left| \mathbf{a}_i^T \mathbf{h} \right|^2 \end{aligned} \quad (43)$$

where $\mathbf{h} = \mathbf{z} - \mathbf{x}$. Appealing to [29, Lemma 3.1], fixing any $\delta' > 0$, the following holds for any $\mathbf{h} \in \mathbb{R}^n$ with probability at least $1 - e^{-m\delta'^2/2}$:

$$\|\mathbf{v}\|^2 = \sum_{i=1}^m v_i^2 \leq \sum_{i=1}^m \left| \mathbf{a}_i^T \mathbf{h} \right|^2 \leq (1 + \delta') m \|\mathbf{h}\|^2. \quad (44)$$

On the other hand, standard matrix concentration results confirm that the largest singular value of $\mathbf{A} = [\mathbf{a}_1 \ \dots \ \mathbf{a}_m]^T$ with i.i.d. Gaussian $\{\mathbf{a}_i\}$ satisfies $\sigma_1 := \|\mathbf{A}\| \leq (1 + \delta'') \sqrt{m}$ for some $\delta'' > 0$ with probability exceeding $1 - 2e^{-c_0 m}$ as soon as $m \geq c_1 n$ for sufficiently large $c_1 > 0$, where $c_1 > 0$ is a universal constant depending on δ'' [57, Remark 5.25]. Combining (42), (43), and (44) yields

$$\begin{aligned} \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\| &\leq \frac{1}{m} \|\mathbf{A}\| \cdot \|\mathbf{v}\| \\ &\leq (1 + \delta') (1 + \delta'') \|\mathbf{h}\| \\ &\leq (1 + \delta)^2 \|\mathbf{h}\|, \quad \delta := \max\{\delta', \delta''\} \end{aligned} \quad (45)$$

which holds with high probability. This condition essentially asserts that the truncated gradient of the objective function $\ell(\mathbf{z})$ or the search direction is well behaved (the function value does not vary too much).

We have related $\|\nabla \ell_{\text{tr}}(\mathbf{z})\|^2$ to $\|\mathbf{h}\|^2$ through (45). Therefore, a more conservative lower bound for $\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle$ in LRC can be given in terms of $\|\mathbf{h}\|^2$. It is equivalent to show that the truncated gradient $\frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z})$ ensures sufficient descent [38], i.e., it obeys a uniform lower bound along the search direction \mathbf{h} taking the form

$$\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle \gtrsim \|\mathbf{h}\|^2 \quad (46)$$

which occupies the remaining of this section. Formally, this can be stated as follows.

Proposition 3: Consider the noiseless measurements $\psi_i = |\mathbf{a}_i^T \mathbf{x}|$, and fix any sufficiently small constant $\epsilon > 0$. There exist universal constants $c_0, c_1 > 0$ such that if $m > c_1 n$, then the following holds with probability exceeding $1 - 4e^{-c_0 m}$:

$$\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle \geq 2(1 - \zeta_1 - \zeta_2 - 2\epsilon) \|\mathbf{h}\|^2 \quad (47)$$

for all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ such that $\|\mathbf{h}\| / \|\mathbf{x}\| \leq \rho$ for $0 < \rho \leq 1/10$ and any fixed $\gamma \geq 1/2$.

Before justifying Proposition 3, we introduce the following events.

Lemma 4: Fix any $\gamma > 0$. For each $i \in [m]$, define

$$\mathcal{E}_i := \left\{ \frac{|\mathbf{a}_i^T \mathbf{z}|}{|\mathbf{a}_i^T \mathbf{x}|} \geq \frac{1}{1 + \gamma} \right\}, \quad (48)$$

$$\mathcal{D}_i := \left\{ \frac{|\mathbf{a}_i^T \mathbf{h}|}{|\mathbf{a}_i^T \mathbf{x}|} \geq \frac{2 + \gamma}{1 + \gamma} \right\}, \quad (49)$$

$$\text{and } \mathcal{K}_i := \left\{ \frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} \neq \frac{\mathbf{a}_i^T \mathbf{x}}{|\mathbf{a}_i^T \mathbf{x}|} \right\} \quad (50)$$

where $\mathbf{h} = \mathbf{z} - \mathbf{x}$. Under the condition $\|\mathbf{h}\| / \|\mathbf{x}\| \leq \rho$, the following inclusion holds for all nonzero $\mathbf{z}, \mathbf{h} \in \mathbb{R}^n$

$$\mathcal{E}_i \cap \mathcal{K}_i \subseteq \mathcal{D}_i \cap \mathcal{K}_i. \quad (51)$$

Proof: From Fig. 1, it is clear that if $\mathbf{z} \in \xi_i^2$, then the sign of $\mathbf{a}_i^T \mathbf{z}$ will be different than that of $\mathbf{a}_i^T \mathbf{x}$. The region ξ_i^2 can be readily specified by the conditions that

$$\frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} \neq \frac{\mathbf{a}_i^T \mathbf{x}}{|\mathbf{a}_i^T \mathbf{x}|}$$

and

$$\frac{|\mathbf{a}_i^T \mathbf{h}|}{|\mathbf{a}_i^T \mathbf{x}|} \geq 1 + \frac{1}{1+\gamma} = \frac{2+\gamma}{1+\gamma}.$$

Under our initialization condition $\|\mathbf{h}\|/\|\mathbf{x}\| \leq \rho$, it is self-evident that \mathcal{D}_i describes two symmetric spherical caps over $\mathbf{a}_i^T \mathbf{x} = \psi_i$ with one being ξ_i^2 . Hence, it holds that $\mathcal{E}_i \cap \mathcal{K}_i = \xi_i^2 \subseteq \mathcal{D}_i \cap \mathcal{K}_i$. \square

To prove (47), consider rewriting the truncated gradient in terms of the events defined in Lemma 4:

$$\begin{aligned} & \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \\ &= \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{z} - \left| \mathbf{a}_i^T \mathbf{x} \right| \frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} \right) \mathbf{a}_i \mathbb{1}_{\mathcal{E}_i} \\ &= \frac{1}{m} \sum_{i=1}^m \mathbf{a}_i \mathbf{a}_i^T \mathbf{h} \mathbb{1}_{\mathcal{E}_i} - \frac{1}{m} \sum_{i=1}^m \left(\frac{\mathbf{a}_i^T \mathbf{z}}{|\mathbf{a}_i^T \mathbf{z}|} - \frac{\mathbf{a}_i^T \mathbf{x}}{|\mathbf{a}_i^T \mathbf{x}|} \right) \left| \mathbf{a}_i^T \mathbf{x} \right| \mathbf{a}_i \mathbb{1}_{\mathcal{E}_i}. \end{aligned} \quad (52)$$

Using the definitions and properties in Lemma 4, one further arrives at

$$\begin{aligned} & \left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle \\ & \geq \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{E}_i} - \frac{1}{m} \sum_{i=1}^m \left| \mathbf{a}_i^T \mathbf{x} \right| \left| \mathbf{a}_i^T \mathbf{h} \right| \mathbb{1}_{\mathcal{E}_i \cap \mathcal{K}_i} \\ & \geq \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{E}_i} - \frac{2}{m} \sum_{i=1}^m \left| \mathbf{a}_i^T \mathbf{x} \right| \left| \mathbf{a}_i^T \mathbf{h} \right| \mathbb{1}_{\mathcal{D}_i \cap \mathcal{K}_i} \\ & \geq \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{E}_i} - \frac{1+\gamma}{2+\gamma} \cdot \frac{2}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{D}_i \cap \mathcal{K}_i} \end{aligned} \quad (53)$$

where the last inequality arises from the property $|\mathbf{a}_i^T \mathbf{x}| \leq \frac{1+\gamma}{2+\gamma} |\mathbf{a}_i^T \mathbf{h}|$ by the definition of \mathcal{D}_i .

Establishing the regularity condition or Proposition 3, boils down to lower bounding the right-hand side of (53), namely, to lower bounding the first term and to upper bounding the second one. By the SLLN, the first term in (53) approximately gives $\|\mathbf{h}\|^2$ as long as our truncation procedure does not eliminate too many generalized gradient components (i.e., summands in the first term). Regarding the second, one would expect its contribution to be small under our initialization condition in (28) and as the relative error $\|\mathbf{h}\|/\|\mathbf{x}\|$ decreases. Specifically, under our initialization, \mathcal{D}_i is provably a rare event, thus eliminating the possibility of the second term exerting a noticeable influence on the first term. Rigorous analyses concerning the two terms are elaborated in Lemma 5 and Lemma 6, whose proofs are provided in Appendix VI-D and Appendix VI-E, respectively.

Lemma 5: Fix $\gamma \geq 1/2$ and $\rho \leq 1/10$, and let \mathcal{E}_i be defined in (48). For independent random variables $W \sim \mathcal{N}(0, 1)$ and $Z \sim \mathcal{N}(0, 1)$, set

$$\zeta_1 := 1 - \min \left\{ \mathbb{E} \left[\mathbb{1} \left\{ \left| \frac{1-\rho+W}{\rho} \right| \geq \frac{\sqrt{1.01}}{\rho(1+\gamma)} \right\} \right], \mathbb{E} \left[Z^2 \mathbb{1} \left\{ \left| \frac{1-\rho+W}{\rho} \right| \geq \frac{\sqrt{1.01}}{\rho(1+\gamma)} \right\} \right] \right\}. \quad (54)$$

Then for any $\epsilon > 0$ and any vector \mathbf{h} obeying $\|\mathbf{h}\|/\|\mathbf{x}\| \leq \rho$, the following holds with probability exceeding $1 - 2e^{-c_5 \epsilon^2 m}$:

$$\frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{E}_i} \geq (1 - \zeta_1 - \epsilon) \|\mathbf{h}\|^2 \quad (55)$$

provided that $m > (c_6 \cdot \epsilon^{-2} \log \epsilon^{-1})n$ for some universal constants $c_5, c_6 > 0$.

To have a sense of how large the quantities involved in Lemma 5 are,

when $\gamma = 0.7$ and $\rho = 1/10$, it holds that

$$\mathbb{E} \left[\mathbb{1} \left\{ \left| \frac{1-\rho+W}{\rho} \right| \geq \frac{\sqrt{1.01}}{\rho(1+\gamma)} \right\} \right] \approx 0.92$$

and

$$\mathbb{E} \left[Z^2 \mathbb{1} \left\{ \left| \frac{1-\rho+W}{\rho} \right| \geq \frac{\sqrt{1.01}}{\rho(1+\gamma)} \right\} \right] \approx 0.99$$

hence leading to $\zeta_1 \approx 0.08$.

Having derived a lower bound for the first term in the right-hand side of (53), it remains to deal with the second one.

Lemma 6: Fix $\gamma > 0$ and $\rho \leq 1/10$, and let $\mathcal{D}_i, \mathcal{K}_i$ be defined in (49), (50), respectively. For any constant $\epsilon > 0$, there exists some universal constants $c_5, c_6 > 0$ such that

$$\frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{D}_i \cap \mathcal{K}_i} \leq (\zeta_2' + \epsilon) \|\mathbf{h}\|^2 \quad (56)$$

holds with probability at least $1 - 2e^{-c_5 \epsilon^2 m}$ provided that $m/n > c_6 \cdot \epsilon^{-2} \log \epsilon^{-1}$ for some universal constants $c_5, c_6 > 0$, where $\zeta_2' = 0.9748 \sqrt{\rho \tau / (0.99 \tau^2 - \rho^2)}$ with $\tau = (2 + \gamma)/(1 + \gamma)$.

With our TAF default parameters $\rho = 1/10$ and $\gamma = 0.7$, we have $\zeta_2' \approx 0.2463$. Using (53), (55), and (56), choosing m/n exceeding some sufficiently large constant such that $c_0 \leq c_5 \epsilon^2$, and denoting $\zeta_2 := 2\zeta_2'(1 + \gamma)/(2 + \gamma)$, the following holds with probability exceeding $1 - 4e^{-c_0 m}$

$$\left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle \geq (1 - \zeta_1 - \zeta_2 - 2\epsilon) \|\mathbf{h}\|^2 \quad (57)$$

for all \mathbf{x} and \mathbf{z} such that $\|\mathbf{h}\|/\|\mathbf{x}\| \leq \rho$ for $0 < \rho \leq 1/10$ and any fixed $\gamma \geq 1/2$. This combined with (39) and (41) proves Proposition 2 for appropriately chosen $\mu > 0$ and $\lambda > 0$.

To conclude this section, an estimate for the working step size is provided next. Plugging the results of (45) and (47) into (40) suggests that

$$\begin{aligned} & \text{dist}^2 \left(\mathbf{z} - \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{x} \right) \\ &= \|\mathbf{h}\|^2 - 2\mu \left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle + \left\| \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2 \\ &\leq \left\{ 1 - \mu \left[2(1 - \zeta_1 - \zeta_2 - 2\epsilon) - \mu(1 + \delta)^4 \right] \right\} \|\mathbf{h}\|^2 \\ &\triangleq (1 - \nu) \|\mathbf{h}\|^2, \end{aligned} \quad (59)$$

and also that $\lambda = 2(1 - \zeta_1 - \zeta_2 - 2\epsilon) - \mu(1 + \delta)^4 \triangleq \lambda_0$ in the local regularity condition in (39). Clearly, it holds that $0 < \lambda < 2(1 - \zeta_1 - \zeta_2)$. Taking ϵ and δ to be sufficiently small, one obtains the feasible range of the step size for TAF

$$\mu \leq \frac{2(0.99 - \zeta_1 - \zeta_2) \triangleq \mu_0}{1.05^4}. \quad (60)$$

In particular, under default parameters in Algorithm 1, $\mu_0 = 0.8388$ and $\lambda_0 = 1.22$, thus concluding the proof of Theorem 1.

VI. CONCLUSION

This paper developed a linear-time algorithm termed TAF for solving generally unstructured systems of random quadratic equations. Our TAF algorithm builds on three key ingredients: an orthogonality-promoting initialization, along with a simple yet effective gradient truncation rule, as well as scalable gradient-like iterations. Numerical tests using synthetic data and real images corroborate the superior performance of TAF over state-of-the-art solvers of the same type.

A few timely and pertinent future research directions are worth pointing out. First, in parallel with spectral initialization methods, the proposed orthogonality-promoting initialization can be applied for semidefinite optimization [36], matrix completion [38], [46], as well as blind deconvolution [37]. It is also interesting to investigate suitable gradient regularization rules in more general nonconvex optimization settings. Extending the theory to the more challenging case where \mathbf{a}_i 's are generated from the coded diffraction pattern model [27] constitutes another meaningful direction.

APPENDIX A PROOFS FOR SECTION V

A. Proof of Lemma 1

By homogeneity of (28), it suffices to work with the case where $\|\mathbf{x}\| = 1$. It is easy to check that

$$\begin{aligned} \frac{1}{2} \|\mathbf{x}\mathbf{x}^T - \tilde{\mathbf{z}}_0\tilde{\mathbf{z}}_0^T\|_F^2 &= \frac{1}{2}\|\mathbf{x}\|^4 + \frac{1}{2}\|\tilde{\mathbf{z}}_0\|^4 - |\mathbf{x}^T\tilde{\mathbf{z}}_0|^2 \\ &= 1 - |\mathbf{x}^T\tilde{\mathbf{z}}_0|^2 \\ &= 1 - \cos^2\theta \end{aligned} \quad (61)$$

where $0 \leq \theta \leq \pi/2$ is the angle between the spaces spanned by \mathbf{x} and $\tilde{\mathbf{z}}_0$. Then one can write

$$\mathbf{x} = \cos\theta \tilde{\mathbf{z}}_0 + \sin\theta \tilde{\mathbf{z}}_0^\perp, \quad (62)$$

where $\tilde{\mathbf{z}}_0^\perp \in \mathbb{R}^n$ is a unit vector that is orthogonal to $\tilde{\mathbf{z}}_0$ and has a nonnegative inner product with \mathbf{x} . Likewise,

$$\mathbf{x}^\perp := -\sin\theta \tilde{\mathbf{z}}_0 + \cos\theta \tilde{\mathbf{z}}_0^\perp, \quad (63)$$

in which $\mathbf{x}^\perp \in \mathbb{R}^n$ is a unit vector orthogonal to \mathbf{x} .

Since $\tilde{\mathbf{z}}_0$ is the solution to the maximum eigenvalue problem

$$\tilde{\mathbf{z}}_0 := \arg \max_{\|z\|=1} z^T \bar{\mathbf{Y}}_0 z \quad (64)$$

for $\bar{\mathbf{Y}}_0 := \frac{1}{|\mathcal{I}_0|} \bar{\mathbf{S}}_0^T \bar{\mathbf{S}}_0$, it is the leading eigenvector of $\bar{\mathbf{Y}}_0$, i.e., $\bar{\mathbf{Y}}_0 \tilde{\mathbf{z}}_0 = \lambda_1 \tilde{\mathbf{z}}_0$, where $\lambda_1 > 0$ is the largest eigenvalue of $\bar{\mathbf{Y}}_0$. Premultiplying (62) and (63) by $\bar{\mathbf{S}}_0$ yields

$$\bar{\mathbf{S}}_0 \mathbf{x} = \cos\theta \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0 + \sin\theta \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp, \quad (65a)$$

$$\bar{\mathbf{S}}_0 \mathbf{x}^\perp = -\sin\theta \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0 + \cos\theta \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp. \quad (65b)$$

Pythagoras' relationship now gives

$$\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 = \cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 + \sin^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2, \quad (66a)$$

$$\|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2 = \sin^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 + \cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2, \quad (66b)$$

where the cross-terms vanish because $\tilde{\mathbf{z}}_0^T \bar{\mathbf{S}}_0^T \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp = |\mathcal{I}_0| \tilde{\mathbf{z}}_0^T \bar{\mathbf{Y}}_0 \tilde{\mathbf{z}}_0^\perp = \lambda_1 |\mathcal{I}_0| \tilde{\mathbf{z}}_0^T \tilde{\mathbf{z}}_0^\perp = 0$ following from the definition of $\tilde{\mathbf{z}}_0^\perp$.

We next construct the following expression:

$$\begin{aligned} \sin^2\theta \|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 - \|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2 &= \sin^2\theta \left(\cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 + \sin^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \right) \\ &\quad - \left(\sin^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 + \cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \right) \\ &= \sin^2\theta \left(\cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 - \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 + \sin^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \right) - \\ &\quad \cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \\ &= \sin^4\theta \left(\|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 - \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 \right) - \cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \leq 0. \end{aligned} \quad (67)$$

Regarding the last inequality, since $\tilde{\mathbf{z}}_0$ maximizes the term $\tilde{\mathbf{z}}_0^T \bar{\mathbf{Y}}_0 \tilde{\mathbf{z}}_0 = \frac{1}{|\mathcal{I}_0|} \tilde{\mathbf{z}}_0^T \bar{\mathbf{S}}_0^T \bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0$ according to (64), then in (67) the first term $\|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 - \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0\|^2 \leq 0$ holds for any unit vector $\tilde{\mathbf{z}}_0^\perp \in \mathbb{R}^n$. In addition, the second term $-\cos^2\theta \|\bar{\mathbf{S}}_0 \tilde{\mathbf{z}}_0^\perp\|^2 \leq 0$, thus yielding $\sin^2\theta \|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 - \|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2 \leq 0$. For any nonzero $\mathbf{x} \in \mathbb{R}^n$, it holds that

$$\sin^2\theta = 1 - \cos^2\theta \leq \frac{\|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2}{\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2}. \quad (68)$$

Upon letting $\mathbf{u} = \mathbf{x}^\perp$, the last inequality taken together with (61) concludes the proof of (29).

B. Proof of Lemma 2

Assume $\|\mathbf{x}\| = 1$. Let $\mathbf{s} \in \mathbb{R}^n$ be sampled uniformly at random on the unit sphere, which has zero mean and covariance matrix \mathbf{I}_n/n . Let also $\mathbf{U} \in \mathbb{R}^{n \times n}$ be a unitary matrix such that $\mathbf{U}\mathbf{x} = \mathbf{e}_1$, where \mathbf{e}_1 is the first canonical vector in \mathbb{R}^n . It is then easy to verify that the following holds for any fixed threshold $0 < \tau < 1$ [59]

$$\begin{aligned} \mathbb{E}[s s^T | (s^T \mathbf{x})^2 > \tau] &= \mathbf{U} \mathbb{E}[\mathbf{U}^T s s^T \mathbf{U} | (s^T \mathbf{U} \mathbf{U}^T \mathbf{x})^2 > \tau] \mathbf{U}^T \\ &\stackrel{(i)}{=} \mathbf{U} \mathbb{E}[\tilde{s} \tilde{s}^T | (\tilde{s}^T \mathbf{e}_1)^2 > \tau] \mathbf{U}^T \\ &= \mathbf{U} \mathbb{E}[\tilde{s} \tilde{s}^T | \tilde{s}_1^2 > \tau] \mathbf{U}^T \\ &= \mathbf{U} \left[\begin{array}{c} \mathbb{E}[\tilde{s}_1^2 | \tilde{s}_1^2 > \tau] \mathbb{E}[\tilde{s}_1 \tilde{s}_1^T | \tilde{s}_1^2 > \tau] \\ \mathbb{E}[\tilde{s}_1 \tilde{s}_1^T | \tilde{s}_1^2 > \tau] \end{array} \right] \mathbf{U}^T \\ &\stackrel{(ii)}{=} \mathbf{U} \left[\begin{array}{c} \mathbb{E}[\tilde{s}_1^2 | \tilde{s}_1^2 > \tau] \mathbf{0}^T \\ \mathbf{0} \end{array} \right] \mathbf{U}^T \\ &\stackrel{(iii)}{=} \mathbb{E}[\tilde{s}_2^2 | \tilde{s}_1^2 > \tau] \mathbf{I}_n + (\mathbb{E}[\tilde{s}_1^2 | \tilde{s}_1^2 > \tau] - \mathbb{E}[\tilde{s}_2^2 | \tilde{s}_1^2 > \tau]) \mathbf{x} \mathbf{x}^T \\ &\triangleq \mathbf{C}_1 \mathbf{I}_n + \mathbf{C}_2 \mathbf{x} \mathbf{x}^T \end{aligned} \quad (69)$$

with the constants $C_1 := \mathbb{E}[\bar{s}_1^2 | \bar{s}_1^2 > \tau] < \frac{1-\tau}{n-1}$, $C_2 := \mathbb{E}[\bar{s}_1^2 | \bar{s}_1^2 > \tau] - C_1 > 0$, and $s_{\setminus 1} \in \mathbb{R}^{n-1}$ denoting the subvector of $s \in \mathbb{R}^n$ after removing the first entry from s . Here, the result (i) follows upon defining $\tilde{s} := U^T s$, which obeys the uniformly spherical distribution too using the rotational invariance. The equality (ii) is due to the zero-mean and symmetrical properties of the uniformly spherical distribution. Finally, to derive (iii), we have used the fact $\mathbf{x} = U\mathbf{e}_1 = \mathbf{u}_1$, the first column of U , which arises from $U^T \mathbf{x} = \mathbf{e}_1$ and $UU^T = \mathbf{I}_n$.

By the argument above, assume without loss of generality that $\mathbf{x} = \mathbf{e}_1$. Consider now the truncated vector $s_{\setminus 1} | (s^T \mathbf{x})^2 > \tau$, or equivalently, $s_{\setminus 1} | s_1^2 > \tau$. It is then clear that $s_{\setminus 1} | s_1^2 > \tau$ is bounded, and thus subgaussian; furthermore, the next hold

$$\mathbb{E}[s_{\setminus 1} | s_1^2 > \tau] = \mathbf{0} \quad (70a)$$

$$\mathbb{E}[(s_{\setminus 1} | s_1^2 > \tau)(s_{\setminus 1} | s_1^2 > \tau)^T] = C_1 \mathbf{I}_{n-1} \quad (70b)$$

where (70b) is obtained as a submatrix of the first term in (69) since the second term $C_2 \mathbf{e}_1 \mathbf{e}_1^T$ is removed.

Considering a unit vector \mathbf{x}^\perp such that $\mathbf{x}^T \mathbf{x}^\perp = \mathbf{e}_1^T \mathbf{x}^\perp = 0$, there exists a unit vector $\mathbf{d} \in \mathbb{R}^{n-1}$ such that $\mathbf{x}^\perp = [0 \ \mathbf{d}^T]^T$. Thus, it holds that

$$\|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2 = \|\bar{\mathbf{S}}_0 [0 \ \mathbf{d}^T]^T\|^2 = \|\bar{\mathbf{S}}_{0, \setminus 1} \mathbf{d}\|^2 \quad (71)$$

where $\bar{\mathbf{S}}_{0, \setminus 1} \in \mathbb{R}^{|\bar{\mathcal{I}}_0| \times (n-1)}$ is obtained through deleting the first column in $\bar{\mathbf{S}}_0$, which is denoted by $\bar{\mathbf{S}}_{0,1}$; that is, $\bar{\mathbf{S}}_0 = [\bar{\mathbf{S}}_{0,1} \ \bar{\mathbf{S}}_{0, \setminus 1}]$.

The rows of $\bar{\mathbf{S}}_{0, \setminus 1}$ may therefore be viewed as independent realizations of the conditional random vector $s_{\setminus 1}^T | s_1^2 > \tau$, with the threshold τ being the $|\bar{\mathcal{I}}_0|$ -largest value in $\{y_i / \|\mathbf{a}_i\|^2\}_{i=1}^m$. Standard concentration inequalities on the sum of random positive semi-definite matrices composed of independent non-isotropic subgaussian rows [57, Remark 5.40] confirm that

$$\left\| \frac{1}{|\bar{\mathcal{I}}_0|} \bar{\mathbf{S}}_{0, \setminus 1}^T \bar{\mathbf{S}}_{0, \setminus 1} - C_1 \mathbf{I}_{n-1} \right\| \leq \sigma C_1 \leq \frac{(1-\tau)\sigma}{n-1} \quad (72)$$

holds with probability at least $1 - 2e^{-c_K n}$ as long as $|\bar{\mathcal{I}}_0|/n$ is sufficiently large, where σ is a numerical constant that can take arbitrarily small values, and $c_K > 0$ is a universal constant. Without loss of generality, let us work with $\sigma := 0.005$ in (72). Then for any unit vector $\mathbf{d} \in \mathbb{R}^{n-1}$, the following inequality holds with probability at least $1 - 2e^{-c_K n}$:

$$\left| \frac{1}{|\bar{\mathcal{I}}_0|} \mathbf{d}^T \bar{\mathbf{S}}_{0, \setminus 1}^T \bar{\mathbf{S}}_{0, \setminus 1} \mathbf{d} - C_1 \right| \leq \frac{0.01}{n} \quad (73)$$

for $n \geq 3$. Therefore, one readily concludes that

$$\|\bar{\mathbf{S}}_0 \mathbf{x}^\perp\|^2 = |(\mathbf{x}^\perp)^T \bar{\mathbf{S}}^T \mathbf{S} \mathbf{x}^\perp| \leq 1.01 |\bar{\mathcal{I}}_0|/n \quad (74)$$

holds with probability at least $1 - 2e^{-c_K n}$, provided that $|\bar{\mathcal{I}}_0|/n$ exceeds some constant. Note that c_K depends on the maximum subgaussian norm of rows of \mathbf{S} , and we assume without loss of generality $c_K \geq 1/2$. Hence, $\|\bar{\mathbf{S}}_0 \mathbf{u}\|^2$ in (29) is upper bounded simply by letting $\mathbf{u} = \mathbf{x}^\perp$ in (74).

C. Proof of Lemma 3

We next pursue a meaningful lower bound for $\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2$ in (31). When $\mathbf{x} = \mathbf{e}_1$, one has $\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 = \|\bar{\mathbf{S}}_0 \mathbf{e}_1\|^2 = \sum_{i=1}^{|\bar{\mathcal{I}}_0|} \bar{s}_{i,1}^2$, where $\{\bar{s}_{i,1}\}_{i=1}^{|\bar{\mathcal{I}}_0|}$ are entries of the first column of $\bar{\mathbf{S}}_0$. It is further worth mentioning that all squared entries of any spherical random vector obey the *Beta* distribution with parameters $\alpha = \frac{1}{2}$, and $\beta = \frac{n-1}{2}$, i.e., $\bar{s}_{i,j}^2 \sim \text{Beta}(\frac{1}{2}, \frac{n-1}{2})$ for all i, j , [60, Lemma 2]. Although they have closed-form probability density functions (pdfs) that may facilitate deriving a lower bound, we take another route detailed as follows. A simple yet useful inequality is established first.

Lemma 7: Given m fractions obeying $1 > \frac{p_1}{q_1} \geq \frac{p_2}{q_2} \geq \dots \geq \frac{p_m}{q_m} > 0$, in which $p_i, q_i > 0, \forall i \in [m]$, the following holds for all $1 \leq k \leq m$

$$\sum_{i=1}^k \frac{p_i}{q_i} \geq \sum_{i=1}^k \frac{p_{[i]}}{q_{[1]}} \quad (75)$$

where $p_{[i]}$ denotes the i -th largest one among $\{p_i\}_{i=1}^m$, and hence, $q_{[1]}$ is the maximum in $\{q_i\}_{i=1}^m$.

Proof: For any $k \in [m]$, according to the definition of $q_{[1]}$, it holds that $p_{[1]} \geq p_{[2]} \geq \dots \geq p_{[k]}$, so $\frac{p_{[1]}}{q_{[1]}} \geq \frac{p_{[2]}}{q_{[1]}} \geq \dots \geq \frac{p_{[k]}}{q_{[1]}}$. Considering $q_{[1]} \geq q_i, \forall i \in [m]$, and letting $j_i \in [m]$ be the index such that $p_{j_i} = p_{[i]}$, then $\frac{p_{j_i}}{q_{j_i}} = \frac{p_{[i]}}{q_{j_i}} \geq \frac{p_{[i]}}{q_{[1]}}$ holds for any $i \in [k]$. Therefore, $\sum_{i=1}^k \frac{p_{j_i}}{q_{j_i}} = \sum_{i=1}^k \frac{p_{[i]}}{q_{j_i}} \geq \sum_{i=1}^k \frac{p_{[i]}}{q_{[1]}}$. Note that $\left\{ \frac{p_{[i]}}{q_{j_i}} \right\}_{i=1}^k$ comprise a subset of terms in $\left\{ \frac{p_i}{q_i} \right\}_{i=1}^m$. On the other hand, according to our assumption, $\sum_{i=1}^k \frac{p_i}{q_i}$ is the largest among all sums of k summands; hence, $\sum_{i=1}^k \frac{p_i}{q_i} \geq \sum_{i=1}^k \frac{p_{[i]}}{q_{j_i}}$ yields $\sum_{i=1}^k \frac{p_i}{q_i} \geq \sum_{i=1}^k \frac{p_{[i]}}{q_{[1]}}$ concluding the proof. \square

Without loss of generality and for simplicity of exposition, let us assume that indices of \mathbf{a}_i 's have been re-ordered such that

$$\frac{a_{1,1}^2}{\|\mathbf{a}_1\|^2} \geq \frac{a_{2,1}^2}{\|\mathbf{a}_2\|^2} \geq \dots \geq \frac{a_{m,1}^2}{\|\mathbf{a}_m\|^2}, \quad (76)$$

where $a_{i,1}$ denotes the first element of \mathbf{a}_i . Therefore, writing $\|\bar{\mathbf{S}}_0 \mathbf{e}_1\|^2 = \sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2 / \|\mathbf{a}_i\|^2$, the next task amounts to finding the sum of the $|\bar{\mathcal{I}}_0|$ largest out of all m entities in (76). Applying the result (75) in Lemma 7 gives

$$\sum_{i=1}^{|\bar{\mathcal{I}}_0|} \frac{a_{i,1}^2}{\|\mathbf{a}_i\|^2} \geq \sum_{i=1}^{|\bar{\mathcal{I}}_0|} \frac{a_{[i],1}^2}{\max_{i \in [m]} \|\mathbf{a}_i\|^2}, \quad (77)$$

in which $a_{[i],1}^2$ stands for the i -th largest entity in $\{a_{i,1}^2\}_{i=1}^m$.

Observe that for i.i.d. random vectors $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, the property $\mathbb{P}(\|\mathbf{a}_i\|^2 \geq 2.3n) \leq e^{-n/2}$ holds for large enough n (e.g., $n \geq 20$), which can be understood upon substituting $\xi := n/2$ into the following standard result [61, Lemma 1]

$$\mathbb{P}\left(\|\mathbf{a}_i\|^2 - n \geq 2\sqrt{\xi} + 2\xi\right) \leq e^{-\xi}. \quad (78)$$

In addition, one readily deduces that $\mathbb{P}\left(\max_{i \in [m]} \|\mathbf{a}_i\| \leq \sqrt{2.3n}\right) \geq 1 - me^{-n/2}$. We will henceforth build our subsequent proofs on this event

without stating this explicitly each time encountering it. Therefore, (77) can be lower bounded by

$$\|\bar{\mathbf{S}}\mathbf{x}\|^2 = \sum_{i=1}^{|\bar{\mathcal{I}}_0|} \frac{a_{i,1}^2}{\|\mathbf{a}_i\|^2} \geq \sum_{i=1}^{|\bar{\mathcal{I}}_0|} \frac{a_{i,1}^2}{\max_{i \in [m]} \|\mathbf{a}_i\|^2} \geq \frac{1}{2.3n} \sum_{i=1}^{|\bar{\mathcal{I}}_0|} |a_{i,1}|^2 \quad (79)$$

which holds with probability at least $1 - me^{-n/2}$. The task left for bounding $\|\bar{\mathbf{S}}\mathbf{x}\|^2$ is to derive a meaningful lower bound for $\sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2$. Roughly speaking, because the ratio $|\bar{\mathcal{I}}_0|/m$ is small, e.g., $|\bar{\mathcal{I}}_0|/m \leq 1/5$, a trivial result consists of bounding $(1/|\bar{\mathcal{I}}_0|) \sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2$ by its sample average $(1/m) \sum_{i=1}^m a_{i,1}^2$. The latter can be bounded using its ensemble mean, i.e., $\mathbb{E}[a_{i,1}^2] = 1$, $\forall i \in [\bar{\mathcal{I}}_0]$, to yield $(1/m) \sum_{i=1}^m a_{i,1}^2 \geq (1 - \epsilon) \mathbb{E}[a_{i,1}^2] = 1 - \epsilon$, which holds with high probability for some numerical constant $\epsilon > 0$ [29, Lemma 3.1]. Therefore, one has a candidate lower bound $\sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2 \geq (1 - \epsilon)|\bar{\mathcal{I}}_0|$. Nonetheless, this lower bound is in general too loose, and it contributes to a relatively large upper bound on the wanted term in (29).

To obtain an alternative bound, let us examine first the typical size of the maximum in $\{a_{i,1}^2\}_{i=1}^m$. Observe obviously that the modulus $|a_{i,1}|$ follows the half-normal distribution having the pdf $p(r) = \sqrt{2/\pi} \cdot e^{-r^2/2}$, $r > 0$, and it is easy to verify that

$$\mathbb{E}[|a_{i,1}|] = \sqrt{2/\pi}. \quad (80)$$

Then integrating the pdf from 0 to $+\infty$ yields the corresponding accumulative distribution function (cdf) expressible in terms of the error function $\mathbb{P}(|a_{i,1}| > \xi) = 1 - \text{erf}(\xi/2)$, i.e., $\text{erf}(\xi) := 2/\sqrt{\pi} \cdot \int_0^\xi e^{-r^2} dr$. Appealing to a lower bound on the complimentary error function $\text{erfc}(\xi) := 1 - \text{erf}(\xi)$ from [62, Theorem 2], one establishes that $\mathbb{P}(|a_{i,1}| > \xi) = 1 - \text{erf}(\xi/2) \geq (3/5)e^{-\xi^2/2}$. Additionally, direct application of probability theory and Taylor expansion confirms that

$$\begin{aligned} \mathbb{P}\left(\max_{i \in [m]} |a_{i,1}| \geq \xi\right) &= 1 - \left[\mathbb{P}(|a_{i,1}| \leq \xi)\right]^m \\ &\geq 1 - \left(1 - 0.6e^{-\xi^2/2}\right)^m \\ &\geq 1 - e^{-0.6me^{-\xi^2/2}}. \end{aligned} \quad (81)$$

Choosing now $\xi := \sqrt{2 \log n}$ leads to

$$\mathbb{P}\left(\max_{i \in [m]} |a_{i,1}| \geq \sqrt{2 \log n}\right) \geq 1 - e^{-0.6m/n} \geq 1 - o(1) \quad (82)$$

which holds with the proviso that m/n is large enough, and the symbol $o(1)$ represents a small constant probability. Thus, provided that m/n exceeds some large constant, the event $\max_{i \in [m]} a_{i,1}^2 \geq 2 \log n$ occurs with high probability. Hence, one may expect a tighter lower bound than $(1 - \epsilon_0)|\bar{\mathcal{I}}_0|$, which is on the same order of m under the assumption that $|\bar{\mathcal{I}}_0|/m$ is about a constant.

Although $a_{i,1}^2$ obeys the *Chi-square* distribution with $k = 1$ degrees of freedom, its cdf is rather complicated and does not admit a nice closed-form expression. A small trick is hence taken in the sequel. Assume without loss of generality that

both m and $|\bar{\mathcal{I}}_0|$ are even. Grouping two consecutive $a_{i,1}^2$'s together, introduce a new variable $\vartheta[i] := a_{[2k-1],1}^2 + a_{[2k],1}^2$, $\forall k \in [m/2]$, hence yielding a sequence of ordered numbers, i.e., $\vartheta[1] \geq \vartheta[2] \geq \dots \geq \vartheta[m/2] > 0$. Then, one can equivalently write the wanted sum as

$$\sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2 = \sum_{i=1}^{|\bar{\mathcal{I}}_0|/2} \vartheta[i]. \quad (83)$$

On the other hand, for i.i.d. standard normal random variables $\{a_{i,1}\}_{i=1}^m$, let us consider grouping randomly two of them and denote the corresponding sum of their squares by $\chi_k := a_{k_i,1}^2 + a_{k_j,1}^2$, where $k_i \neq k_j \in [m]$, and $k \in [m/2]$. It is self-evident that the χ_k 's are identically distributed obeying the *Chi-square* distribution with $k = 2$ degrees of freedom, having the pdf

$$p(r) = \frac{1}{2} e^{-r/2}, \quad r \geq 0, \quad (84)$$

and the following complementary cdf (ccdf)

$$\mathbb{P}(\chi_k \geq \xi) := \int_\xi^\infty \frac{1}{2} e^{-r/2} dr = e^{-\xi/2}, \quad \forall \xi \geq 0. \quad (85)$$

Ordering all χ_k 's, summing the $|\bar{\mathcal{I}}_0|/2$ largest ones, and comparing the resultant sum with the one in (83) confirms that

$$\sum_{i=1}^{|\bar{\mathcal{I}}_0|/2} \chi_{[i]} \leq \sum_{i=1}^{|\bar{\mathcal{I}}_0|/2} \vartheta_{[i]} = \sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{i,1}^2, \quad \forall |\bar{\mathcal{I}}_0| \in [m]. \quad (86)$$

Upon setting $\mathbb{P}(\chi_k \geq \xi) = |\bar{\mathcal{I}}_0|/m$, one obtains an estimate of $\chi_{|\bar{\mathcal{I}}_0|/2}$, the $(|\bar{\mathcal{I}}_0|/2)$ -th largest value in $\{\chi_k\}_{k=1}^{m/2}$ as follows

$$\hat{\chi}_{|\bar{\mathcal{I}}_0|/2} := 2 \log(m/|\bar{\mathcal{I}}_0|). \quad (87)$$

Furthermore, applying the Hoeffding-type inequality [57, Proposition 5.10] and leveraging the convexity of the ccdf in (85), one readily establishes that

$$\mathbb{P}\left(\hat{\chi}_{|\bar{\mathcal{I}}_0|/2} - \chi_{|\bar{\mathcal{I}}_0|/2} > \xi\right) \leq e^{-\frac{1}{4}m\xi^2 e^{-\xi} (|\bar{\mathcal{I}}_0|/m)^2}, \quad \forall \xi > 0. \quad (88)$$

Taking without loss of generality $\xi := 0.05 \hat{\chi}_{|\bar{\mathcal{I}}_0|/2} = 0.1 \log(m/|\bar{\mathcal{I}}_0|)$ gives

$$\mathbb{P}\left(\chi_{|\bar{\mathcal{I}}_0|/2} < 0.95 \hat{\chi}_{|\bar{\mathcal{I}}_0|/2}\right) \leq e^{-c_0 m} \quad (89)$$

for some universal constants $c_0, c_\chi > 0$, and sufficiently large n such that $|\bar{\mathcal{I}}_0|/m \gtrsim c_\chi > 0$. The remaining part in this section assumes that this event occurs.

Choosing $\xi := 4 \log n$ and substituting this into the ccdf in (85) leads to

$$\mathbb{P}(\chi \leq 4 \log n) = 1 - 1/n^2. \quad (90)$$

Notice that each summand in $\sum_{i=1}^{|\bar{\mathcal{I}}_0|/2} \chi_{[i]} \geq \sum_{i=1}^{m/2} \chi_i \mathbb{1}_{\bar{\mathcal{E}}_i}$ is Chi-square distributed, and hence could be unbounded, so we choose to work with the truncation $\sum_{i=1}^{m/2} \chi_i \mathbb{1}_{\bar{\mathcal{E}}_i}$, where the $\mathbb{1}_{\bar{\mathcal{E}}_i}$'s are independent copies of $\mathbb{1}_{\bar{\mathcal{E}}}$, and $\mathbb{1}_{\bar{\mathcal{E}}}$ denotes the indicator function for the ensuing events

$$\bar{\mathcal{E}} := \left\{ \chi \geq \hat{\chi}_{|\bar{\mathcal{I}}_0|/2} \right\} \cap \left\{ \chi \leq 4 \log n \right\}. \quad (91)$$

Apparently, it holds that $\sum_{i=1}^{|\bar{\mathcal{I}}_0|/2} \chi_{[i]} \geq \sum_{i=1}^{m/2} \chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i}$. One further establishes that

$$\begin{aligned} \mathbb{E} \left[\chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i} \right] &:= \int_{\hat{\chi}_{|\bar{\mathcal{I}}_0|/2}}^{4 \log n} \frac{1}{2} r e^{-r/2} dr \\ &= \left(\hat{\chi}_{|\bar{\mathcal{I}}_0|/2} + 2 \right) e^{-\hat{\chi}_{|\bar{\mathcal{I}}_0|/2/2} - (4 \log n + 2)} e^{-2 \log n} \\ &= \frac{2|\bar{\mathcal{I}}_0|}{m} \left[1 + \log(m/|\bar{\mathcal{I}}_0|) \right] - \frac{(4 \log n + 2)}{n^2}. \end{aligned} \quad (92)$$

The task of bounding $\sum_{i=1}^{|\bar{\mathcal{I}}_0|} a_{[i],1}^2$ in (86) now boils down to bounding $\sum_{i=1}^{m/2} \chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i}$ from its expectation in (92). A convenient way to accomplish this is using the Bernstein inequality [57, Proposition 5.16], that deals with bounded random variables. That also justifies introducing the upper-bound truncation on χ in (91). Specifically, define

$$\vartheta_i := \chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i} - \mathbb{E} \left[\chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i} \right], \quad 1 \leq i \leq m/2. \quad (93)$$

Thus, $\{\vartheta_i\}_{i=1}^{m/2}$ are i.i.d. centered and bounded random variables following from the mean-subtraction and the upper-bound truncation. Further, according to the cdf (85) and the definition of sub-exponential random variables [57, Definition 5.13], the terms $\{\vartheta_i\}_{i=1}^{m/2}$ are sub-exponential. Then, the following

$$\left| \sum_{i=1}^{m/2} \vartheta_i \right| \geq \tau \quad (94)$$

holds with probability at least $1 - 2e^{-c_s \min(\tau/K_s, \tau^2/K_s^2)}$, in which $c_s > 0$ is a universal constant, and $K_s := \max_{i \in [m/2]} \|\vartheta_i\|_{\psi_1}$ represents the maximum subexponential norm of the ϑ_i 's.

Indeed, K_s can be found as follows [57, Definition 5.13]:

$$\begin{aligned} K_s &:= \sup_{p \geq 1} p^{-1} \left(\mathbb{E} \left[|\vartheta_i|^p \right] \right)^{1/p} \\ &\leq (4 \log n - 2 \log(m/|\bar{\mathcal{I}}_0|)) \left[|\bar{\mathcal{I}}_0|/m - 1/n^2 \right] \\ &\leq \frac{2|\bar{\mathcal{I}}_0|}{m} \log(n^2 |\bar{\mathcal{I}}_0|/m) \\ &\leq \frac{4|\bar{\mathcal{I}}_0|}{m} \log n. \end{aligned} \quad (95)$$

Choosing $\tau := 8|\bar{\mathcal{I}}_0|/(c_s m) \cdot \log^2 n$ in (94) yields

$$\begin{aligned} \sum_{i=1}^{m/2} \chi_i \mathbb{1}_{\tilde{\mathcal{E}}_i} &\geq |\bar{\mathcal{I}}_0| \left[1 + \log(m/|\bar{\mathcal{I}}_0|) \right] - 8|\bar{\mathcal{I}}_0|/(c_s m) \cdot \log^2 n \\ &\quad - m(2 \log n + 1)/n^2 \\ &\geq (1 - \epsilon_s) |\bar{\mathcal{I}}_0| \left[1 + \log(m/|\bar{\mathcal{I}}_0|) \right] \end{aligned} \quad (96)$$

for some small constant $\epsilon_s > 0$, which holds with probability at least $1 - me^{-n/2} - e^{-c_0 m} - 1/n^2$ as long as m/n exceeds some numerical constant and n is sufficiently large. Therefore, combining (79), (86), and (96), one concludes that the following holds with high probability

$$\|\bar{\mathbf{S}}_0 \mathbf{x}\|^2 = \sum_{i=1}^{|\bar{\mathcal{I}}_0|} \frac{a_{i,1}^2}{\|\mathbf{a}_i\|^2} \geq (1 - \epsilon_s) \frac{|\bar{\mathcal{I}}_0|}{2.3n} \left[1 + \log(m/|\bar{\mathcal{I}}_0|) \right]. \quad (97)$$

Taking $\epsilon_s := 0.01$ without loss of generality concludes the proof of Lemma 3.

D. Proof of Lemma 5

Let us first prove the argument for a fixed pair \mathbf{h} and \mathbf{x} , such that \mathbf{h} and \mathbf{z} are independent of $\{\mathbf{a}_i\}_{i=1}^m$, and then apply a covering argument. To start, introduce a Lipschitz-continuous counterpart for the discontinuous indicator function [6, Appendix A.2]

$$\chi_E(\theta) := \begin{cases} 1, & |\theta| \geq \frac{\sqrt{1.01}}{1+\gamma}, \\ 100(1+\gamma)^2 \theta^2 - 100, & \frac{1}{1+\gamma} \leq |\theta| < \frac{\sqrt{1.01}}{1+\gamma}, \\ 0, & |\theta| < \frac{1}{1+\gamma} \end{cases} \quad (98)$$

with Lipschitz constant $\mathcal{O}(1)$. Recall that $\mathcal{E}_i = \left\{ \left| \frac{\mathbf{a}_i^T \mathbf{z}}{\mathbf{a}_i^T \mathbf{x}} \right| \geq \frac{1}{1+\gamma} \right\}$, so it holds that $0 \leq \chi_E \left(\frac{\mathbf{a}_i^T \mathbf{z}}{\mathbf{a}_i^T \mathbf{x}} \right) \leq \mathbb{1}_{\mathcal{E}_i}$ for any $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^n$, thus yielding

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{E}_i} &\geq \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \chi_E \left(\frac{\mathbf{a}_i^T \mathbf{z}}{\mathbf{a}_i^T \mathbf{x}} \right) \\ &= \frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \chi_E \left(1 + \frac{\mathbf{a}_i^T \mathbf{h}}{\mathbf{a}_i^T \mathbf{x}} \right). \end{aligned} \quad (99)$$

By homogeneity and rotational invariance of normal distributions, it suffices to prove the case where $\mathbf{x} = \mathbf{e}_1$ and $\|\mathbf{h}\|/\|\mathbf{x}\| = \|\mathbf{h}\| \leq \rho$. According to (99), lower bounding the first term in (53) can be achieved by lower bounding $\sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \chi_E \left(1 + \frac{\mathbf{a}_i^T \mathbf{h}}{\mathbf{a}_i^T \mathbf{x}} \right)$ instead. To that end, let us find the mean of $\left(\mathbf{a}_i^T \mathbf{h} \right)^2 \chi_E \left(1 + \frac{\mathbf{a}_i^T \mathbf{h}}{\mathbf{a}_i^T \mathbf{x}} \right)$. Note that $\left(\mathbf{a}_i^T \mathbf{h} \right)^2$ and $\chi_E \left(1 + \frac{\mathbf{a}_i^T \mathbf{h}}{\mathbf{a}_i^T \mathbf{x}} \right)$ are dependent. Introduce an orthonormal matrix \mathbf{U}_h that contains $\mathbf{h}^T/\|\mathbf{h}\|$ as its first row, i.e.,

$$\mathbf{U}_h := \begin{bmatrix} \mathbf{h}^T/\|\mathbf{h}\| \\ \tilde{\mathbf{U}}_h \end{bmatrix} \quad (100)$$

for some orthogonal matrix $\tilde{\mathbf{U}}_h \in \mathbb{R}^{(n-1) \times n}$ such that \mathbf{U}_h is orthonormal. Moreover, define $\tilde{\mathbf{h}} := \mathbf{U}_h \mathbf{h}$, and $\tilde{\mathbf{a}}_i := \mathbf{U}_h \mathbf{a}_i$; and let $\tilde{a}_{i,1}$ and $\tilde{\mathbf{a}}_{i,\setminus 1}$ denote the first entry and the remaining entries in the vector $\tilde{\mathbf{a}}_i$; likewise for $\tilde{\mathbf{h}}$. Then, for any \mathbf{h} such that $\|\mathbf{h}\| \leq \rho$, we have

$$\begin{aligned} &\mathbb{E} \left[\left(\mathbf{a}_i^T \mathbf{h} \right)^2 \chi_E \left(1 + \frac{\mathbf{a}_i^T \mathbf{h}}{\mathbf{a}_i^T \mathbf{x}} \right) \right] \\ &= \mathbb{E} \left[\left(\tilde{a}_{i,1} \tilde{h}_1 \right)^2 \chi_E \left(1 + \frac{\tilde{a}_{i,1} \tilde{h}_1}{\tilde{\mathbf{a}}_{i,1}^T \tilde{\mathbf{h}}_1} \right) \right] \\ &\quad + \mathbb{E} \left[\left(\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1} \right)^2 \chi_E \left(1 + \frac{\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1}}{\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1}} \right) \right] \\ &= \tilde{h}_1^2 \mathbb{E} \left[\tilde{a}_{i,1}^2 \chi_E \left(1 + \frac{\tilde{a}_{i,1} \tilde{h}_1}{\tilde{\mathbf{a}}_{i,1}^T \tilde{\mathbf{h}}_1} \right) \right] \\ &\quad + \mathbb{E} \left[\left(\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1} \right)^2 \right] \mathbb{E} \left[\chi_E \left(1 + \frac{\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1}}{\tilde{\mathbf{a}}_{i,\setminus 1}^T \tilde{\mathbf{h}}_{\setminus 1}} \right) \right] \end{aligned}$$

$$\begin{aligned}
&= \tilde{h}_1^2 \mathbb{E} \left[\tilde{a}_{i,1}^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \right] \\
&\quad + \|\tilde{\mathbf{h}}_{\setminus 1}\|^2 \mathbb{E} \left[\chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \right] \\
&\geq (\tilde{h}_1^2 + \|\tilde{\mathbf{h}}_{\setminus 1}\|^2) \min \left\{ \mathbb{E} \left[\tilde{a}_{i,1}^2 \chi_E \left(\left| 1 + h_1 + \frac{\mathbf{a}_{i,\setminus 1}^T \mathbf{h}_{\setminus 1}}{a_{i,1}} \right| \right) \right], \right. \\
&\quad \left. \mathbb{E} \left[\chi_E \left(\left| 1 + h_1 + \frac{\mathbf{a}_{i,\setminus 1}^T \mathbf{h}_{\setminus 1}}{a_{i,1}} \right| \right) \right] \right\} \\
&\geq \|\mathbf{h}\|^2 \min \left\{ \mathbb{E} \left[\tilde{a}_{i,1}^2 \chi_E \left(\left| 1 - \rho + \frac{a_{i,2} \rho}{a_{i,1}} \right| \right) \right], \right. \\
&\quad \left. \mathbb{E} \left[\chi_E \left(\left| 1 - \rho + \frac{a_{i,2} \rho}{a_{i,1}} \right| \right) \right] \right\} = (1 - \zeta_1) \|\mathbf{h}\|^2 \quad (101)
\end{aligned}$$

where the second equality follows from the independence between $\tilde{a}_{i,1}^T \tilde{\mathbf{h}}_{\setminus 1}$ and $\mathbf{a}_i^T \mathbf{h}$, the second inequality holds for $\rho \leq 1/10$ and $\gamma \geq 1/2$, and the last equality comes from the definition of ζ_1 in (93). Notice that $\varrho := (\mathbf{a}_i^T \mathbf{h})^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \leq (\mathbf{a}_i^T \mathbf{h})^2 \stackrel{d}{=} \|\mathbf{h}\|^2 a_{i,1}^2$ is a subexponential variable, and thus its subexponential norm $\|\varrho\|_{\psi_1} := \sup_{p \geq 1} [\mathbb{E}(|\varrho|^p)]^{1/p}$ is finite.

Direct application of the Bernstein-type inequality [57, Proposition 5.16] confirms that for any $\epsilon > 0$, the following

$$\begin{aligned}
&\frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{h})^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \\
&\quad \geq \mathbb{E} \left[(\mathbf{a}_i^T \mathbf{h})^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \right] - \epsilon \|\mathbf{h}\|^2 \\
&\quad \geq (1 - \zeta_1 - \epsilon) \|\mathbf{h}\|^2 \quad (102)
\end{aligned}$$

holds with probability at least $1 - e^{-c_5 m \epsilon^2}$ for some numerical constant $c_5 > 0$ provided that $\epsilon \leq \|\varrho\|_{\psi_1}$ by assumption.

To obtain uniform control over all vectors \mathbf{z} and \mathbf{x} such that $\|\mathbf{z} - \mathbf{x}\| \leq \rho$, the net covering argument is applied [57, Definition 5.1]. Let \mathcal{S}_ϵ be an ϵ -net of the unit sphere, \mathcal{L}_ϵ be an ϵ -net of $[0, \rho]$, and define

$$\mathcal{N}_\epsilon := \{(\mathbf{z}, \mathbf{h}, t) : (\mathbf{z}_0, \mathbf{h}_0, t_0) \in \mathcal{S}_\epsilon \times \mathcal{S}_\epsilon \times \mathcal{L}_\epsilon\}. \quad (103)$$

Since the cardinality $|\mathcal{S}_\epsilon| \leq (1 + 2/\epsilon)^n$ [57, Lemma 5.2], then

$$|\mathcal{N}_\epsilon| \leq (1 + 2/\epsilon)^{2n} \rho/\epsilon \leq (1 + 2/\epsilon)^{2n+1} \quad (104)$$

due to the fact that $\rho/\epsilon < 2/\epsilon < 1 + 2/\epsilon$ for $0 < \rho < 1$.

Consider now any $(\mathbf{z}, \mathbf{h}, t)$ obeying $\|\mathbf{h}\| = t \leq \rho$. There exists a pair $(\mathbf{z}_0, \mathbf{h}_0, t_0) \in \mathcal{N}_\epsilon$ such that $\|\mathbf{z} - \mathbf{z}_0\|, \|\mathbf{h} - \mathbf{h}_0\|$, and $|t - t_0|$ are each at most ϵ . Taking the union bound yields

$$\begin{aligned}
&\frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{h}_0)^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}_0}{a_{i,1}} \right| \right) \\
&\quad \geq \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{h}_0)^2 \chi_E \left(\left| 1 - t_0 + \frac{a_{i,2} t_0}{a_{i,1}} \right| \right) \\
&\quad \geq (1 - \zeta_1 - \epsilon) \|\mathbf{h}_0\|^2, \quad \forall (\mathbf{z}_0, \mathbf{h}_0, t_0) \in \mathcal{N}_\epsilon \quad (105)
\end{aligned}$$

with probability at least $1 - (1 + 2/\epsilon)^{2n+1} e^{-c_5 \epsilon^2 m} \geq 1 - e^{-c_0 m}$, which follows by choosing m such that $m \geq (c_6 \cdot \epsilon^{-2} \log \epsilon^{-1}) n$ for some constant $c_6 > 0$.

Recall that $\chi_E(\tau)$ is Lipschitz continuous, thus

$$\begin{aligned}
&\left| \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{h})^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \right. \\
&\quad \left. - (\mathbf{a}_i^T \mathbf{h}_0)^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}_0}{a_{i,1}} \right| \right) \right| \\
&\lesssim \frac{1}{m} \sum_{i=1}^m \left| (\mathbf{a}_i^T \mathbf{h})^2 - (\mathbf{a}_i^T \mathbf{h}_0)^2 \right| \\
&= \frac{1}{m} \sum_{i=1}^m \left| \mathbf{a}_i^T (\mathbf{h} \mathbf{h}^T - \mathbf{h}_0 \mathbf{h}_0^T) \mathbf{a}_i \right| \\
&\lesssim c_7 \sum_{i=1}^m \left| \mathbf{h} \mathbf{h}^T - \mathbf{h}_0 \mathbf{h}_0^T \right| \leq 2.5 c_7 \|\mathbf{h} - \mathbf{h}_0\| \|\mathbf{h}\| \\
&\leq 2.5 c_7 \rho \epsilon \quad (106)
\end{aligned}$$

for some numerical constant c_7 and provided that $\epsilon < 1/2$ and $m \geq (c_6 \cdot \epsilon^{-2} \log \epsilon^{-1}) n$, where the first inequality arises from the Lipschitz property of $\chi_E(\tau)$, the second uses the results in [6, Lemma 1], and the third from [6, Lemma 2].

Putting all results together confirms that with probability exceeding $1 - 2e^{-c_0 m}$, we have

$$\begin{aligned}
&\frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{h})^2 \chi_E \left(\left| 1 + \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \\
&\quad \geq [1 - \zeta_1 - (1 + 2.5 c_7 \rho) \epsilon] \|\mathbf{h}\|^2 \quad (107)
\end{aligned}$$

for all vectors $\|\mathbf{h}\| / \|\mathbf{x}\| \leq \rho$, concluding the proof.

E. Proof of Lemma 6

Similar to the proof in Section VI-D, it is convenient to work with the following auxiliary function instead of the discontinuous indicator function

$$\chi_D(\theta) := \begin{cases} 1, & |\theta| \geq \frac{2+\gamma}{1+\gamma} \\ -100 \left(\frac{1+\gamma}{2+\gamma} \right)^2 \theta^2 + 100, & \sqrt{0.99} \cdot \frac{2+\gamma}{1+\gamma} \leq |\theta| < \frac{2+\gamma}{1+\gamma} \\ 0, & |\theta| < \sqrt{0.99} \cdot \frac{2+\gamma}{1+\gamma} \end{cases} \quad (108)$$

which is Lipschitz continuous in θ with Lipschitz constant $\mathcal{O}(1)$. For $\mathcal{D}_i = \left\{ \left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \geq \frac{2+\gamma}{1+\gamma} \right\}$, it holds that $0 \leq \mathbb{1}_{\mathcal{D}_i} \leq$

$\chi_D \left(\left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right)$ for any $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^n$. Assume without loss of generality that $\mathbf{x} = \mathbf{e}_1$. Then for $\gamma > 0$ and $\rho \leq 1/10$, it holds that

$$\frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\left\{ \frac{|\mathbf{a}_i^T \mathbf{h}|}{|a_{i,1}^T \mathbf{x}|} \geq \frac{2+\gamma}{1+\gamma} \right\}} \leq \frac{1}{m} \sum_{i=1}^m \chi_D \left(\left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}^T \mathbf{x}} \right| \right)$$

$$\begin{aligned}
&= \frac{1}{m} \sum_{i=1}^m \chi_D \left(\left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \right) \\
&= \frac{1}{m} \sum_{i=1}^m \chi_D \left(\left| h_1 + \frac{\mathbf{a}_{i,\setminus 1}^T \mathbf{h}_{\setminus 1}}{a_{i,1}} \right| \right) \\
&= \frac{1}{m} \sum_{i=1}^m \chi_D \left(\left| h_1 + \frac{a_{i,2}}{a_{i,1}} \|\mathbf{h}_{\setminus 1}\| \right| \right) \\
&\stackrel{(i)}{\leq} \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\left\{ \left| h_1 + \frac{a_{i,2}}{a_{i,1}} \|\mathbf{h}_{\setminus 1}\| \right| \geq \sqrt{0.99} \frac{2+\gamma}{1+\gamma} \right\}} \quad (109)
\end{aligned}$$

where the last inequality arises from the definition of χ_D . Note that $a_{i,2}/a_{i,1}$ obeys the standard Cauchy distribution, i.e., $a_{i,2}/a_{i,1} \sim \text{Cauchy}(0, 1)$ [63]. Transformation properties of Cauchy distributions assert that $h_1 + \frac{a_{i,2}}{a_{i,1}} \|\mathbf{h}_{\setminus 1}\| \sim \text{Cauchy}(h_1, \|\mathbf{h}_{\setminus 1}\|)$ [64]. Recall that the cdf of a Cauchy distributed random variable $w \sim \text{Cauchy}(\mu_0, \alpha)$ is given by [63]

$$F(w; \mu_0, \alpha) = \frac{1}{\pi} \arctan \left(\frac{w - \mu_0}{\alpha} \right) + \frac{1}{2}. \quad (110)$$

It is easy to check that when $\|\mathbf{h}_{\setminus 1}\| = 0$, the indicator function $\mathbb{1}_{\mathcal{D}_i} = 0$ due to $|h_1| \leq \rho < \sqrt{0.99}(2+\gamma)/(1+\gamma)$. Consider only $\|\mathbf{h}_{\setminus 1}\| \neq 0$ next. Define for notational brevity $w := a_{i,2}/a_{i,1}$, $\alpha := \|\mathbf{h}_{\setminus 1}\|$, as well as $\mu_0 := h_1/\alpha$ and $w_0 := \sqrt{0.99} \frac{2+\gamma}{\alpha(1+\gamma)}$. Then,

$$\begin{aligned}
&\mathbb{E}[\mathbb{1}_{\{\mu_0+w \geq w_0\}}] = 1 - [F(w_0; \mu_0, 1) - F(-w_0; \mu_0, 1)] \\
&= 1 - \frac{1}{\pi} [\arctan(w_0 - \mu_0) - \arctan(-w_0 - \mu_0)] \\
&\stackrel{(i)}{=} \frac{1}{\pi} \arctan \left(\frac{2w_0}{w_0^2 - \mu_0^2 - 1} \right) \\
&\stackrel{(ii)}{\leq} \frac{1}{\pi} \cdot \frac{2w_0}{w_0^2 - \mu_0^2 - 1} \\
&\stackrel{(iii)}{\leq} \frac{1}{\pi} \cdot \frac{2\sqrt{0.99}\rho(2+\gamma)/(1+\gamma)}{0.99(2+\gamma)^2/(1+\gamma)^2 - \rho^2} \\
&\leq 0.0646 \quad (111)
\end{aligned}$$

for all $\gamma > 0$ and $\rho \leq 1/10$. In deriving (i), we used the property $\arctan(u) + \arctan(v) = \arctan\left(\frac{u+v}{1-uv}\right) \pmod{\pi}$ for any $uv \neq 1$. Concerning (ii), the inequality $\arctan(x) \leq x$ for $x \geq 0$ is employed. Plugging given parameter values and using $\|\mathbf{h}_{\setminus 1}\| \leq \|\mathbf{h}\| \leq \rho$ confirms (iii). Next, $\mathbb{1}_{\{\mu_0+w \geq w_0\}}$ is bounded; and it is known that all bounded random variables are subexponential. Thus, upon applying the Bernstein-type inequality [57, Corollary 5.17], the next holds with probability at least $1 - e^{-c_5 m \epsilon^2}$ for some numerical constant $c_5 > 0$ and any sufficiently small $\epsilon > 0$:

$$\begin{aligned}
\frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\left\{ \left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \geq \frac{2+\gamma}{1+\gamma} \right\}} &\leq \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\left\{ \left| h_1 + \frac{a_{i,2}}{a_{i,1}} \|\mathbf{h}_{\setminus 1}\| \right| \geq \sqrt{0.99} \frac{2+\gamma}{1+\gamma} \right\}} \\
&\leq (1+\epsilon) \mathbb{E} \left[\mathbb{1}_{\left\{ \left| h_1 + \frac{a_{i,2}}{a_{i,1}} \|\mathbf{h}_{\setminus 1}\| \right| \geq \sqrt{0.99} \frac{2+\gamma}{1+\gamma} \right\}} \right] \\
&\leq \frac{1+\epsilon}{\pi} \cdot \frac{2\sqrt{0.99}\rho(2+\gamma)/(1+\gamma)}{0.99(2+\gamma)^2/(1+\gamma)^2 - \rho^2}. \quad (112)
\end{aligned}$$

On the other hand, it is easy to establish that the following holds true for any fixed $\mathbf{h} \in \mathbb{R}^n$:

$$\mathbb{E} \left[\left(\mathbf{a}_i^T \mathbf{h} \right)^4 \right] = \mathbb{E} \left[a_{i,1}^4 \right] \|\mathbf{h}\|^4 = 3 \|\mathbf{h}\|^4 \quad (113)$$

which has also been used in Lemma 1 [6] and Lemma 6.1 [34]. Furthermore, recalling our working assumption $\|\mathbf{a}_i\| \leq \sqrt{2.3n}$ and $\|\mathbf{h}\| \leq \rho \|\mathbf{x}\|$, the random variables $(\mathbf{a}_i^T \mathbf{h})^4$ are bounded, and thus they are subexponential [57]. Appealing again to the Bernstein-type inequality for subexponential random variables [57, Proposition 5.16] and provided that $m/n > c_6 \epsilon^{-2} \log \epsilon^{-1}$ for some numerical constant $c_6 > 0$, we have

$$\frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^4 \leq 3(1+\epsilon) \|\mathbf{h}\|^4 \quad (114)$$

which holds with probability exceeding $1 - e^{-c_5 m \epsilon^2}$ for some universal constant $c_5 > 0$ and any sufficiently small $\epsilon > 0$.

Combining results (112), (114), leveraging the Cauchy-Schwartz inequality, and considering $\mathcal{D}_i \cap \mathcal{K}_i$ only consisting of a spherical cap, the following holds for any $\rho \leq 1/10$ and $\gamma > 0$:

$$\begin{aligned}
&\frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^2 \mathbb{1}_{\mathcal{D}_i \cap \mathcal{K}_i} \\
&\leq \sqrt{\frac{1}{m} \sum_{i=1}^m \left(\mathbf{a}_i^T \mathbf{h} \right)^4} \sqrt{\frac{1}{2} \cdot \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\left\{ \left| \frac{\mathbf{a}_i^T \mathbf{h}}{a_{i,1}} \right| \geq \frac{2+\gamma}{1+\gamma} \right\}}} \\
&\leq \sqrt{3(1+\epsilon) \|\mathbf{h}\|^4} \sqrt{\frac{1+\epsilon}{\pi} \cdot \frac{\sqrt{0.99}\rho(2+\gamma)/(1+\gamma)}{0.99(2+\gamma)^2/(1+\gamma)^2 - \rho^2}} \\
&\triangleq (\zeta'_2 + \epsilon') \|\mathbf{h}\|^2 \quad (115)
\end{aligned}$$

where $\zeta'_2 := 0.9748 \sqrt{\rho\tau/(0.99\tau^2 - \rho^2)}$ with $\tau := (2+\gamma)/(1+\gamma)$, which holds with probability at least $1 - 2e^{-c_0 m}$. The latter arises upon choosing $c_0 \leq c_5 \epsilon^2$ in $1 - 2e^{-c_5 m \epsilon^2}$, which can be accomplished by taking the ratio m/n sufficiently large.

ACKNOWLEDGMENTS

The authors would like to thank Prof. John Duchi for pointing out an error in an initial draft of this paper. They also thank Profs. Soltanolkotabi, Yuxin Chen, and Drs. Kejun Huang and Ju Sun for helpful discussions.

REFERENCES

- [1] R. Balan, P. Casazza, and D. Edidin, "On signal reconstruction without phase," *Appl. Comput. Harmon. Anal.*, vol. 20, no. 3, pp. 345–356, May 2006.
- [2] A. Conca, D. Edidin, M. Hering, and C. Vinzant, "An algebraic characterization of injectivity in phase retrieval," *Appl. Comput. Harmon. Anal.*, vol. 38, no. 2, pp. 346–356, Mar. 2015.
- [3] A. S. Bandeira, J. Cahill, D. G. Mixon, and A. A. Nelson, "Saving phase: Injectivity and stability for phase retrieval," *Appl. Comput. Harmon. Anal.*, vol. 37, no. 1, pp. 106–125, 2014.
- [4] P. M. Pardalos and S. A. Vavasis, "Quadratic programming with one negative eigenvalue is NP-hard," *J. Global Optim.*, vol. 1, no. 1, pp. 15–22, 1991.
- [5] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, vol. 2. Philadelphia, PA, USA: SIAM, 2001.
- [6] Y. Chen and E. J. Candès, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," *Commun. Pure Appl. Math.*, vol. 70, no. 5, pp. 822–883, May 2017.

- [7] J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," *Opt. Lett.*, vol. 3, no. 1, pp. 27–29, Jul. 1978.
- [8] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM Rev.*, vol. 57, no. 2, pp. 225–251, May 2015.
- [9] K. Jaganathan, Y. C. Eldar, and B. Hassibi, "Phase retrieval: An overview of recent developments," in *Optical Compressive Imaging*. Boca Raton, FL, USA: CRC Press, 2015, pp. 263–296.
- [10] J. Miao, P. Charalambous, J. Kirz, and D. Sayre, "Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens," *Nature*, vol. 400, no. 6742, pp. 342–344, Jul. 1999.
- [11] R. P. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Amer. A, Opt. Image Sci.*, vol. 7, no. 3, pp. 394–411, 1990.
- [12] A. Chai, M. Moscoso, and G. Papanicolaou, "Array imaging using intensity-only measurements," *Inverse Problems*, vol. 27, no. 1, p. 015005, Jan. 2011.
- [13] S. Marchesini, Y.-C. Tu, and H.-T. Wu, "Alternating projection, ptychographic imaging and phase synchronization," *Appl. Comput. Harmon. Anal.*, vol. 41, no. 3, pp. 815–851, Jun. 2016.
- [14] J. C. Dainty and J. R. Fienup, "Phase retrieval and image reconstruction for astronomy," in *Image Recovery: Theory and Application*, H. Stark, Ed. New York, NY, USA: Academic, 1987, pp. 231–275.
- [15] J. Miao, I. Ishikawa, Q. Shen, and T. Earnest, "Extending X-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes," *Annu. Rev. Phys. Chem.*, vol. 59, no. 1, pp. 387–410, May 2008.
- [16] H. Sahinoglou and S. D. Cabrera, "On phase retrieval of finite-length sequences using the initial time sample," *IEEE Trans. Circuits Syst.*, vol. 38, no. 8, pp. 954–958, Aug. 1991.
- [17] C. J. Hillar and L.-H. Lim, "Most tensor problems are NP-hard," *J. ACM*, vol. 60, no. 6, p. 45, 2013.
- [18] 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.
- [19] G. Wang and G. B. Giannakis, "Solving random systems of quadratic equations via truncated generalized gradient flow," in *Proc. Adv. Neural Inf. Process. Syst.*, Barcelona, Spain, 2016, pp. 568–576.
- [20] K. G. Murty and S. N. Kabadi, "Some NP-complete problems in quadratic and nonlinear programming," *Math. Program.*, vol. 39, no. 2, pp. 117–129, 1987.
- [21] 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.
- [22] E. Hofstetter, "Construction of time-limited functions with specified autocorrelation functions," *IEEE Trans. Inf. Theory*, vol. 10, no. 2, pp. 119–126, Apr. 1964.
- [23] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient phase retrieval of sparse signals," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, Feb. 2014.
- [24] 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.
- [25] Y. C. Eldar and S. Mendelson, "Phase retrieval: Stability and recovery guarantees," *Appl. Comput. Harmon. Anal.*, vol. 36, no. 3, pp. 473–494, May 2014.
- [26] Y. C. Eldar, P. Sidorenko, D. G. Mixon, S. Barel, and O. Cohen, "Sparse phase retrieval from short-time Fourier measurements," *IEEE Signal Process. Lett.*, vol. 22, no. 5, pp. 638–642, May 2015.
- [27] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Appl. Comput. Harmon. Anal.*, vol. 39, no. 2, pp. 277–299, Sep. 2015.
- [28] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 18, pp. 4814–4826, Sep. 2015.
- [29] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Appl. Comput. Harmon. Anal.*, vol. 66, no. 8, pp. 1241–1274, Aug. 2013.
- [30] H. Zhang, Y. Chi, and Y. Liang, (2016). "Median-truncated nonconvex approach for phase retrieval with outliers." [Online]. Available: <https://arxiv.org/abs/1603.03805>
- [31] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237–246, Nov. 1972.
- [32] M. Soltanolkotabi, "Algorithms and theory for clustering and nonconvex quadratic programming," Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, USA, 2014.
- [33] K. Wei, "Solving systems of phaseless equations via Kaczmarz methods: A proof of concept study," *Inverse Problems*, vol. 31, no. 12, p. 125008, 2015.
- [34] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," *J. Found. Comput. Math.*, 2017. [Online]. Available: <https://doi.org/10.1007/s10208-017-9365-9>
- [35] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht. (2015). "Low-rank solutions of linear matrix equations via Procrustes flow." [Online]. Available: <https://arxiv.org/abs/1507.03566>
- [36] S. Sanghavi, R. Ward, and C. D. White, "The local convexity of solving systems of quadratic equations," *Results Math.*, vol. 71, pp. 569–608, Jun. 2017.
- [37] X. Li, S. Ling, T. Strohmer, and K. Wei. (2016). "Rapid, robust, and reliable blind deconvolution via nonconvex optimization." [Online]. Available: <https://arxiv.org/abs/1606.04933>
- [38] 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.
- [39] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, *MaxCut* and complex semidefinite programming," *Math. Program.*, vol. 149, nos. 1–2, pp. 47–81, 2015.
- [40] K. Huang, Y. C. Eldar, and N. D. Sidiropoulos, "Phase retrieval from 1D Fourier measurements: Convexity, uniqueness, and algorithms," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6105–6117, Dec. 2016.
- [41] C. Qian, N. D. Sidiropoulos, K. Huang, L. Huang, and H. C. So, "Phase retrieval using feasible point pursuit: Algorithms and Cramér–Rao bound," *IEEE Trans. Signal Process.*, vol. 64, no. 20, pp. 5282–5296, Oct. 2016.
- [42] G. Wang, G. B. Giannakis, and J. Chen, "Scalable solvers of random quadratic equations via stochastic truncated amplitude flow," *IEEE Trans. Signal Process.*, vol. 65, no. 8, pp. 1961–1974, Apr. 2017.
- [43] G. Wang, L. Zhang, G. B. Giannakis, J. Chen, and M. Akçakaya. (2016). "Sparse phase retrieval via truncated amplitude flow." [Online]. Available: <https://arxiv.org/abs/1611.07641>
- [44] G. Wang, G. B. Giannakis, Y. Saad, and J. Chen. (2017). "Solving most high-dimensional systems of random quadratic equations." [Online]. Available: <https://arxiv.org/abs/1705.10407>
- [45] E. J. Candès and X. Li, "Solving quadratic equations via PhaseLift when there are about as many equations as unknowns," *Found. Comput. Math.*, vol. 14, no. 5, pp. 1017–1026, 2014.
- [46] 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.
- [47] L.-H. Yeh *et al.*, "Experimental robustness of Fourier ptychography phase retrieval algorithms," *Opt. Exp.*, vol. 23, no. 26, pp. 33214–33240, Dec. 2015.
- [48] T. Cai, J. Fan, and T. Jiang, "Distributions of angles in random packing on spheres," *J. Mach. Learn. Res.*, vol. 14, no. 1, pp. 1837–1864, Jan. 2013.
- [49] H. Zhang, Y. Zhou, Y. Liang, and Y. Chi. (2016). "Reshaped Wirtinger flow and incremental algorithm for solving quadratic system of equations." [Online]. Available: <https://arxiv.org/abs/1605.07719>
- [50] N. Z. Shor, "A class of almost-differentiable functions and a minimization method for functions of this class," *Cybern. Syst. Anal.*, vol. 8, no. 4, pp. 599–606, Jul. 1972.
- [51] R. Rockafellar and R. J.-B. Wets, *Variational Analysis*. Berlin, Germany: Springer-Verlag, 1998.
- [52] N. Z. Shor, K. C. Kiwiel, and A. Ruszczyński, *Minimization Methods for Non-Differentiable Functions*. New York, NY, USA: Springer-Verlag, 1985.
- [53] F. H. Clarke, *Optimization and Nonsmooth Analysis*, vol. 5. Philadelphia, PA, USA: SIAM, 1990.
- [54] F. H. Clarke, "Generalized gradients and applications," *Trans. Amer. Math. Soc.*, vol. 205, pp. 247–262, Apr. 1975.
- [55] P. Chen, A. Fannjiang, and G.-R. Liu. (2015). "Phase retrieval with one or two diffraction patterns by alternating projection with null initialization." [Online]. Available: <https://arxiv.org/abs/1510.07379>
- [56] P. Chen and A. Fannjiang (2015). "Fourier phase retrieval with a single mask by Douglas–Rachford algorithm." [Online]. Available: <https://arxiv.org/abs/1509.00888>
- [57] R. Vershynin. (2010). "Introduction to the non-asymptotic analysis of random matrices." [Online]. Available: <https://arxiv.org/abs/1011.3027>

- [58] Y. Saad, *Iterative Methods for Sparse Linear Systems*. Philadelphia, PA, USA: SIAM, 2003.
- [59] J. C. Duchi and F. Ruan. (2017). "Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval." [Online]. Available: <https://arxiv.org/abs/1705.02356>
- [60] S. Cambanis, S. Huang, and G. Simons, "On the theory of elliptically contoured distributions," *J. Multivariate Anal.*, vol. 11, no. 3, pp. 368–385, Sep. 1981.
- [61] B. Laurent and P. Massart, "Adaptive estimation of a quadratic functional by model selection," *Ann. Stat.*, vol. 28, no. 5, pp. 1302–1338, Oct. 2000.
- [62] S.-H. Chang, P. C. Cosman, and L. B. Milstein, "Chernoff-type bounds for the Gaussian error function," *IEEE Trans. Commun.*, vol. 59, no. 11, pp. 2939–2944, Nov. 2011.
- [63] T. S. Ferguson, "A representation of the symmetric bivariate Cauchy distribution," *Ann. Math. Stat.*, vol. 33, no. 4, pp. 1256–1266, 1962.
- [64] H. Y. Lee, G. J. Parka, and H. M. Kim, "A clarification of the Cauchy distribution," *Commun. Stat. Appl. Methods*, vol. 21, no. 2, pp. 183–191, Mar. 2014.

Gang Wang (S'12) received the B.Eng. degree in Electrical Engineering and Automation from the Beijing Institute of Technology, Beijing, China, in 2011. He is currently working toward his Ph.D. degree in the Department of Electrical and Computer Engineering at the University of Minnesota, Minneapolis, MN, USA.

His research interests focus on the areas of high-dimensional statistical signal processing, stochastic and nonconvex optimization with applications to autonomous energy grids, and deep learning. He received a National Scholarship from China in 2014, the Student Travel Awards from the NSF (2016) and NIPS (2017), and a Best Student Paper Award at the 2017 European Signal Processing Conference.

Georgios B. Giannakis (F'97) received his Diploma in Electrical Engineering from the National Technical University of Athens, Greece, 1981. From 1982 to 1986 he was with the University of Southern California, where he received his MSc. in Electrical Engineering, 1983, MSc. in Mathematics, 1986, and Ph.D. in Electrical Engineering, 1986. He was with the University of Virginia from 1987 to 1998, and since 1999 he has been a professor with the University of Minnesota, where he holds an Endowed Chair in Wireless Telecommunications, a University of Minnesota McKnight Presidential Chair in Electrical and Computer Engineering, and serves as director of the Digital Technology Center.

His general interests span the areas of communications, networking and statistical signal processing - subjects on which he has published more than 400 journal papers, 700 conference papers, 25 book chapters, two edited books and two research monographs (h-index 127). Current research focuses on learning from Big Data, wireless cognitive radios, and network science with applications to social, brain, and power networks with renewables. He is the (co-) inventor of 30 patents issued, and the (co-) recipient of 8 best paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (2015). He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society.

Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in Physics in 1995 and the B.Sc. degree in Electrical Engineering in 1996 both from Tel-Aviv University (TAU), Tel-Aviv, Israel, and the Ph.D. degree in Electrical Engineering and Computer Science in 2002 from the Massachusetts Institute of Technology (MIT), Cambridge.

She is currently a Professor in the Department of Electrical Engineering at the Technion - Israel Institute of Technology, Haifa, Israel, where she holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT, an Adjunct Professor at Duke University, and was a Visiting Professor at Stanford University, Stanford, CA. She is a member of the Israel Academy of Sciences and Humanities (elected 2017), an IEEE Fellow and a EURASIP Fellow. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyo Tomiyasu Award (2016). She was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technions Award for Excellence in Teaching (two times). She received several best paper awards and best demo awards together with her research students and colleagues including the SIAM outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel.

She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She is the Editor in Chief of *Foundations and Trends in Signal Processing*, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an associate editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She was Co-Chair and Technical Co-Chair of several international conferences and workshops.

She is author of the book *Sampling Theory: Beyond Bandlimited Systems* and co-author of the books *Compressed Sensing* and *Convex Optimization Methods in Signal Processing and Communications*, all published by Cambridge University Press.