

A Deterministic Theory for Exact Non-Convex Phase Retrieval

Bariscan Yonel , *Student Member, IEEE*, and Birsen Yazici , *Senior Member, IEEE*

Abstract—In this paper, we analyze the non-convex framework of Wirtinger Flow (WF) for phase retrieval and identify a novel sufficient condition for universal exact recovery through the lens of low rank matrix recovery theory. Via a perspective in the lifted domain, we show that the WF iterates converge to a true solution with fully deterministic arguments under a single condition on the lifted forward model. To this end, a geometric relationship between the accuracy of spectral initialization and the validity of the *regularity condition* is derived. In particular, we determine that a certain concentration property on the spectral matrix must hold uniformly with a sufficiently tight constant. This culminates into a sufficient condition that is equivalent to a restricted isometry-type property over rank-1, positive semi-definite matrices, and amounts to a less stringent requirement on the lifted forward model than those of prominent low-rank-matrix-recovery methods in the literature. We characterize the performance limits of our framework in terms of the tightness of the concentration property via novel bounds on the convergence rate and on the signal-to-noise ratio such that the theoretical guarantees are valid using the spectral initialization at the proper sample complexity.

Index Terms—Wirtinger Flow, non-convex optimization, low rank matrix recovery, phase retrieval, lifting, exact recovery.

I. INTRODUCTION

A. Phase Retrieval

GENERALIZED phase retrieval (GPR) is a ubiquitous problem in science and engineering. The problem consists of the recovery of an object of interest $\mathbf{x} \in \mathbb{C}^N$ given the intensity only measurements of the form:

$$y_m = |\langle \mathbf{a}_m, \mathbf{x} \rangle|^2, \quad m = 1, 2, \dots, M, \quad (1)$$

where $\mathbf{a}_m \in \mathbb{C}^N$ denotes the m^{th} sampling vector. In literature, $\{\mathbf{a}_m\}_{m=1}^M$ most prominently corresponds to models such as Gaussian sampling [1], coded diffraction patterns [2], or the rows of a known linear transformation, such as the short time Fourier transform [3], or a particular imaging operator [8]. These models arise in problems such as X-Ray crystallography,

Manuscript received January 22, 2020; revised May 24, 2020; accepted June 24, 2020. Date of publication July 8, 2020; date of current version August 20, 2020. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. An Liu. This work was supported in part by the Air Force Office of Scientific Research (AFOSR) under the agreement FA9550-19-1-0284, in part by Office of Naval Research (ONR) under the agreement N0001418-1-2068 and in part by the National Science Foundation (NSF) under Grant ECCS-1809234. (Corresponding author: Birsen Yazici.)

The authors are with the Department of Electrical, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: yonelb@rpi.edu; yazici@ecse.rpi.edu).

Digital Object Identifier 10.1109/TP.2020.3007967

coded diffraction imaging, optical astronomy, quantum state tomography, array imaging, or blind channel estimation.

One approach to address GPR is via a least-squares formulation in which the ℓ_2 loss over intensity measurements in (1) is minimized as follows:

$$\underset{\mathbf{z}}{\text{minimize}} \quad f(\mathbf{z}) := \frac{1}{2M} \sum_{m=1}^M (y_m - |\langle \mathbf{a}_m, \mathbf{z} \rangle|^2)^2. \quad (2)$$

Alternative forms of f have also been popular in practice for optical imaging applications, where the ℓ_2 loss is computed as a mismatch of amplitudes instead of intensities [4]. In solving GPR by (2) or using the amplitudes as measurements, the objective function f is non-holomorphic and non-convex due to its invariance to global phase factors on the complex valued variable \mathbf{z} . Conventional methods from optical imaging literature reformulate (2) as a bilinear inverse problem by inserting the missing phase component as a variable, which is then solved by alternating minimization [5], [6], or non-convex analogs of feasibility problems [7]. However, these methods are not equipped with practical recovery guarantees, and carry the risk of getting stuck in local minima due to the non-convexity of the problem.

Despite the ill-posed nature of the problem, there has been a significant progress in the development of provably good GPR algorithms in the last decade. Such methods are characterized by either one or both of the following two principles: convexification of the equality constraints and the solution set, which include *lifting* based approaches [2], [8]–[10], or a provably accurate initialization, followed by an algorithmic map that refines the initial estimate on the original signal domain [1], [11], [12]. Notably, lifting-based approaches reformulate inversion from the quadratic equations of the form (1) into a convex semi-definite program while squaring the dimension of the inverse problem. As a result, these solvers have demanding implementation costs due to computational complexity and memory requirements, which limit their applicability for large scale sensing problems. Essentially, methods that operate on the original signal domain evade such practical bottlenecks arising from the increased dimensionality of the inverse problem.

B. Wirtinger Flow

The latter two-step approach for exact phase retrieval on the original signal domain was most prominently popularized by the seminal Wirtinger Flow (WF) framework [1]. In contrast to other state-of-the art exact phase retrieval methods, i.e., lifting

or linear programming based approaches [12]–[14], WF solves the original non-convex problem in (2) directly.

Given an initial estimate \mathbf{z}_0 , WF performs steepest descent iterations by means of Wirtinger derivatives of f as follows:

$$\mathbf{z}_{k+1} = \mathbf{z}_k - \frac{\mu_k}{\|\mathbf{z}_0\|^2} \nabla f(\mathbf{z}_k), \quad (3)$$

where ∇f is defined as the complex gradient operator, and μ_k is the step size. The premise of WF is that if \mathbf{z}_0 is *sufficiently accurate*, the iterates formed by (3) provably converge with a geometric rate to an element in the *global solution set* which is defined as follows:

Definition I.1: Global Solution Set. Let

$$P := \{e^{j\phi} \mathbf{x} : \phi \in [0, 2\pi]\}, \quad (4)$$

where $\mathbf{x} \in \mathbb{C}^N$ is the ground truth of intensity measurements (1). The set P is said to be the global solution set of (2).

In general for any $\mathbf{z} \in \mathbb{C}^N$, the non-convex set of the form $\{e^{j\Phi} \mathbf{z} : \Phi \in [0, 2\pi]\}$ represents an equivalence under the mapping of intensity only measurements. The convergence of algorithm iterates is governed by the following distance metric:

Definition I.2: Let $\mathbf{x} \in \mathbb{C}^N$ be an element of the solution set P . The distance of an element $\mathbf{z} \in \mathbb{C}^N$ to \mathbf{x} is defined as [1]:

$$\text{dist}(\mathbf{z}, \mathbf{x}) = \min_{\phi \in [0, 2\pi]} \|\mathbf{z} - \mathbf{x}e^{j\phi}\|. \quad (5)$$

The angle $\hat{\phi}$ where the minimum is achieved for a given $\mathbf{z} \in \mathbb{C}^N$ is denoted as $\Phi(\mathbf{z})$.

In literal terms, (5) quantifies the distance of an estimate to the closest point in P , eliminating the effect of non-uniqueness caused by the global phase factors. As such, the exact phase retrieval refers to the iterates converging to any of the elements in the global solution set.

Having to solve a non-convex problem, exact recovery guarantees of WF framework depend on the accuracy of the initial estimate \mathbf{z}_0 which is computed by the *spectral method* [11] as follows:

$$\mathbf{Y} = \frac{1}{M} \sum_{m=1}^M y_m \mathbf{a}_m \mathbf{a}_m^H. \quad (6)$$

The leading eigenvector of \mathbf{Y} , denoted as \mathbf{v}_0 , is scaled by the square root of the normalized ℓ_1 -norm of the data, i.e., $\lambda_0 = M^{-1} \|\mathbf{y}\|_1$ to yield the initial estimate $\mathbf{z}_0 = \sqrt{\lambda_0} \mathbf{v}_0$. Denoting $\text{dist}(\mathbf{z}_0, \mathbf{x}) = \epsilon \|\mathbf{x}\|$, the initial estimate determines an ϵ -neighborhood of P as follows:

Definition I.3: ϵ -Neighborhood of P . Let

$$E(\epsilon) = \{\mathbf{z} \in \mathbb{C}^N : \text{dist}(\mathbf{z}, P) \leq \epsilon \|\mathbf{x}\|\}, \quad (7)$$

where P is the global solution set as defined in (4). The set $E(\epsilon)$ is said to be the ϵ -neighborhood of P .

The main result of WF framework is that for Gaussian sampling and coded diffraction patterns, the initial estimate computed by the spectral method yields a small enough relative distance- ϵ , such that the following *regularity condition* holds with high probability for $M = \mathcal{O}(N \log N)$.

Condition I.1: Regularity Condition. The objective function f in (2) satisfies the regularity condition if, for all $\mathbf{z} \in E(\epsilon)$ the

following inequality holds

$$\text{Re} \left(\langle \nabla f(\mathbf{z}), (\mathbf{z} - \mathbf{x}e^{j\Phi(\mathbf{z})}) \rangle \right) \geq \frac{1}{\alpha} \text{dist}^2(\mathbf{z}, \mathbf{x}) + \frac{1}{\beta} \|\nabla f(\mathbf{z})\|^2 \quad (8)$$

with fixed $\alpha > 0$ and $\beta > 0$ such that $\alpha\beta > 4$.

Lemma 7.10 in [1] establishes that if the regularity condition is satisfied, the WF iterations are contractions with respect to the distance metric in (5) and all the algorithm iterates remain in $E(\epsilon)$. Essentially, the validity of (8) ensures that there exists no first order optimal point $\mathbf{z} \in E(\epsilon)$ other than the elements of P .

C. Related Work and Our Contributions

WF inspired several variants [15]–[20], which improve on its performance guarantees with respect to computational and sample complexity, as well as robustness to noise and outliers. The original WF framework has a $\mathcal{O}(1/N)$ specification on the algorithm step-size, yielding $\mathcal{O}(MN^2 \log 1/\epsilon_0)$ computational complexity for an ϵ_0 -relative accuracy on the final estimate of the optimization. Furthermore, the sample complexity of the method is $M = \mathcal{O}(N \log N)$ in order to control the heavy tailed distribution of the spectral matrix, which also relates to the local curvature of the objective function within $E(\epsilon)$. These issues are largely mitigated by the Truncated WF (TWF) framework [15], which guarantees convergence with a $\mathcal{O}(1)$ step-size for a linear computational complexity, and uniformly controls the tails of the spectral matrix when $M = \mathcal{O}(N)$ by devising a particular sample truncation scheme. Similar outcomes are achieved by alternative approaches that *reshape* the non-convex objective function to a loss over amplitudes [4], [18]–[20] in the form of reduced sample complexity, and faster convergence. Truncation methods were further studied in relation to increasing the robustness of these methods through median truncation [17] or noise estimation [16]. Spectral initialization schemes were also subject to further studies such as those involving the design of optimal pre-processing functions [21]–[25], generalizations [26], [27], or alternative formulations including highest correlation estimators [19], and orthogonality promoting methods [20].

The aforementioned WF-inspired works offer exact recovery guarantees for phase retrieval based on a wide range of theoretical arguments which are prominently probabilistic in nature, derived through the properties of statistical models assumed for the underlying measurement maps. Namely, the exact recovery analysis of state-of-the-art frameworks focus on establishing the regularity condition in (8) at the proper sample complexity, given that $E(\epsilon)$ is constructed by a specific initialization method. We instead take a low-rank matrix recovery based approach to phase retrieval, and conduct a geometric analysis of the optimization problem in (2) for arbitrary measurement models.

To this end, we develop a theoretical framework that unifies the key arguments that contribute to the exact recovery guarantees of WF and its intensity loss based variants under a single sufficient condition. Specifically, we show that one arrives at the restricted strong convexity property of the objective function around a global solution directly through a concentration bound of the spectral matrix due to the special structure of the set

of rank-1, positive semi-definite (PSD) matrices. We reach this conclusion by interpreting WF in the *lifted domain*. As a result, our framework establishes that the two steps of the non-convex optimization framework blueprinted by the seminal work in [1], i.e., the accuracy of spectral initialization, and the regularity condition, are geometric outcomes of a less restrictive sufficient condition related to the following concentration bound:

$$\|\mathbf{Y} - (\mathbf{x}\mathbf{x}^H + \|\mathbf{x}\|^2\mathbf{I})\| \leq \delta\|\mathbf{x}\|^2, \quad \text{for any } \mathbf{x} \in \mathbb{C}^N. \quad (9)$$

(9) is by no means an unexpected property. In fact, (9) is known to hold true with high probability for Gaussian sampling and coded diffraction patterns when $M = \mathcal{O}(N \log N)$ through the concentration of the Hessian of f around its expectation when evaluated at a global solution. Typically, it is used within the probabilistic analysis conducted for statistical models in relating the distance of the spectral initialization to the ground truth. Its uniformity over all $\mathbf{x} \in \mathbb{C}^N$ is also established when $M = \mathcal{O}(N)$ under the sample truncation scheme of TWF [15]. In our work, we show that this is a much stronger property than it is given credit for. Namely, we prove that if the concentration bound in (9) holds uniformly over all \mathbf{x} where δ is sufficiently tight with $\delta \leq 0.184$, then the regularity condition is redundant for the exact recovery guarantees of WF starting from the spectral initialization. In other words, there surely exists positive α, β with $\alpha\beta > 4$ in Condition I.1 such that (8) is satisfied *deterministically* via the restricted strong convexity of the objective function in (2).

The resulting deterministic convergence framework amounts to two key results in this paper. First, under the validity of our sufficient condition, we identify the best achievable convergence rate that guarantees the exact recovery of any unknown from its intensity-only measurements. The upper bound on the convergence rate is determined solely by the concentration bound parameter δ , which facilitates the derivation of an optimal fixed step-size $\mu_k = \mu(\delta)$ for the algorithm that is an $\mathcal{O}(1)$ constant. Secondly, in the presence of additive noise on the received intensity-only measurements, we determine a δ -dependent lower bound for the signal-to-noise ratio such that the algorithm is guaranteed to yield an estimate that is within a bounded perturbation from the ground truth on expectation. Essentially, these results characterize the relationship of value of δ to the trade-offs between model parameters and performance of the algorithm with respect to convergence and stability for practical purposes, provided that $\delta \leq 0.184$.

Related to our theory of WF as a non-convex optimization framework, it is observed in [15] and [28] that the regularity condition can be enforced by the restricted strong convexity condition due to the local Lipschitz differentiability of the objective function. Several key insights are built in [28] for exactly solving (2), in which a benign geometry for the objective function f is realized with high probability when the number of samples are sufficiently large as $\mathcal{O}(N \log^3 N)$ for the Gaussian sampling model. In particular, it is observed that all local minimizers of f are the elements of the global solution set P , and all its saddle points have a directional negative curvature with high probability, which allow vanilla gradient descent to converge to the exact solution even if the algorithm is initialized randomly.

Through the implicit regularization properties of gradient descent and the incoherence property in the Gaussian sampling model, [29] improves the step size of WF to $\mathcal{O}(1/\log N)$ starting from spectral initialization when $M = \mathcal{O}(N \log N)$, thereby attaining a near linear computational complexity of $\mathcal{O}(MN \log N \log 1/\epsilon_0)$. [29] also identifies a $\mathcal{O}(N \log^3 N)$ sample complexity for the convergence of WF with a constant step-size. In contrast, [30] establishes the restricted strong convexity of the objective function (2) around the solution set at $\mathcal{O}(N \log^2 N)$ sample complexity for the Gaussian sampling model, with spectral initialization proven to fall within the strongly convex region with high probability. In this paper, we improve upon the required sample complexity identified in [30] to $\mathcal{O}(N \log N)$ for the restricted strong convexity of (2) in a sufficiently wide neighborhood around the solution set that is guaranteed to include the initial spectral estimate. Furthermore, our approach facilitates reducing the computational complexity of the algorithm to $\mathcal{O}(MN \log 1/\epsilon_0)$ for ϵ_0 -relative accuracy due to a constant $\mathcal{O}(1)$ step-size, which is an improvement over [1] and [29] at the identical sampling complexity.

A notable work in the literature in relation to our results is [26]. When coupled with the result in [30], [26] guarantees the exact recovery of a low rank- r matrix deterministically via restricted strong convexity with a properly set step-size that relates to the spectra of the ground truth. Another work that closely relates to our framework is [31], in which a non-convex approach based on a local restricted isometry property (RIP) over rank-2 matrices is considered as sufficient and established for the Gaussian sampling model with high probability for exactly solving the blind deconvolution problem. In essence, our framework stands in agreement with [26], [31]. Specifically, the sufficient condition we identify on the lifted forward model for exact non-convex phase retrieval *i*) deterministically yields a local RIP-2 type condition reminiscent to that of [31] when specified to hold only over the difference of two rank-1, PSD matrices, and *ii*) directly implies restricted strong convexity due to the duality between lifted domain and the underlying signal domain, after which results of [26] become applicable. Furthermore, under the validity of the concentration bound over all $\mathbf{x} \in \mathbb{C}^N$, the sufficient condition for exact recovery is effectively converted to a RIP on the lifted forward model of the problem over the set of rank-1, positive semi-definite (PSD) matrices. In this manner, the developed framework realizes the full characterization of an observation made in [26] by deriving the best uniform convergence rate and the optimal step-size for exact recovery only with respect to the restricted isometry constant (RIC). Thereby, our framework facilitates a bridge between the optimization based low rank matrix recovery methods, and the prominent statistical frameworks for phase retrieval with a unifying sufficient condition.

In recent years, RIP-based conditions on the lifted measurement map have been established as sufficient for global optimality in recovering a rank- r matrix from quadratic or bilinear equations via first-order non-convex optimization methods. RIP over rank- $6r$ with a RIC less than or equal to 1/10 is shown to be sufficient for exact recovery via Procrustes Flow [32]. RIP over rank- $4r$ matrices with a RIC less than or equal to

1/5 is shown to guarantee the strict saddle point condition and the absence of any spurious local minima for the ℓ_2 loss function [33]. RIP over rank- $2r$ matrices with a RIC less than or equal to 1/5 is shown to be sufficient for having no spurious local minima [34] whereas at most a RIC of 1/2 over rank- $2r$ matrices is postulated as necessary in [35]. Notably, the sufficient condition we establish in our framework corresponds to a less restrictive RIP-type condition than those of the non-convex low rank matrix recovery methods in [32]–[35], as it suffices that the property is satisfied only over the set of rank-1, PSD matrices. The major difference of our result stems from the fact that we are merely interested in restricted strong convexity within the ϵ -neighborhood obtained from the spectral initialization, and not the global properties of the optimization landscape. As a result, to the best of our knowledge, with this paper we establish the most minimal RIP based framework that fully characterizes the performance guarantees of non-convex phase retrieval via WF. Other key works for the non-convex rank- r matrix recovery problem include [36], in which the regularity condition of WF is considered and shown to hold for the Gaussian sampling model, and [26], [37] in which exact recovery, and stability are studied under restricted strong convexity and smoothness of the objective function, respectively. For further discussion on advances in non-convex low rank matrix recovery, we refer the reader to [38].

D. Notation and Organization of the Paper

The rest of the paper is organized as follows. In Section II, we provide a preliminary discussion on the interpretation of WF in the lifted domain. Section III contains our main results, and remarks. Section IV evaluates the robustness of WF in the presence of additive noise. In Section V, we present the proofs of our results. Section VI concludes the paper.

We denote the elements of finite dimensional vector spaces with lower case bold letters. Upper case bold and italic letters are allocated for matrices and sets, respectively. Caligraphic letters are allocated for operators that act on the lifted domain in $\mathbb{C}^{N \times N}$. In denoting the norms of elements in different domains, we use the following notation: $\|\cdot\|$ denotes the ℓ_2 norm when acting on a vector, and the spectral norm when acting on an operator. $\|\cdot\|_F$ and $\|\cdot\|_*$ denote the Frobenius and nuclear norms, respectively. \mathbf{I} and \mathcal{I} denote the identity operators on the vector space of \mathbb{C}^N , and the lifted domain in $\mathbb{C}^{N \times N}$, respectively.

II. WF IN THE LIFTED DOMAIN

We start by interpreting WF as a solver in the *lifted domain*, and adopt the concepts of the seminal work of PhaseLift in [2]. Lifting based approaches provide a profound perspective to the phase retrieval problem. In principle, these methods target the core issue of non-injectivity of phaseless measurement maps, which is a key step in formulating methods that guarantee exact recovery in phase retrieval literature [12], [39]. Notably, one can consider the measurement model in (1) as a mapping from a rank-1, positive semi-definite matrix $\mathbf{x}\mathbf{x}^H \in \mathbb{C}^{N \times N}$ instead of a quadratic map from the signal domain in \mathbb{C}^N . Lifting conceptualizes this observation:

Definition II.1: Lifting. Each measurement in (1) can be expressed in the form of an inner product of two rank-1 operators, $\tilde{\mathbf{X}} = \mathbf{x}\mathbf{x}^H$ and $\mathbf{A}_m = \mathbf{a}_m\mathbf{a}_m^H$ such that

$$y_m = \langle \mathbf{A}_m, \tilde{\mathbf{X}} \rangle_F \quad m = 1, \dots, M \quad (10)$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product. The process of transforming the signal recovery over \mathbb{C}^N to the recovery of the rank-1 unknown $\tilde{\mathbf{X}} \in \mathbb{C}^{N \times N}$ is known as lifting.

The lifting technique introduces a new, linear measurement map $\mathcal{A} : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^M$, which we refer to as *the lifted forward model*. Specifically, for the phaseless measurement model in (1), the domain of \mathcal{A} is constrained on the set of rank-1, positive semi-definite (PSD) matrices $R_1^+ = \{\mathbf{z}\mathbf{z}^H : \mathbf{z} \in \mathbb{C}^N\}$ as follows:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}\mathbf{x}^H) \quad (11)$$

where $\mathbf{y} = [y_1, y_2, \dots, y_M] \in \mathbb{R}^M$. As a result, each non-convex set of equivalent points under the mapping from the signal domain in \mathbb{C}^N to the phaseless measurements, i.e., $\{\mathbf{z}e^{i\Phi} : \Phi \in [0, 2\pi]\}$ for $\mathbf{z} \in \mathbb{C}^N$, is compressed into a single element $\mathbf{z}\mathbf{z}^H \in R_1^+$. Thereby, quadratic equality constraints over the signal domain are transformed to affine equality constraints in the lifted domain in $\mathbb{C}^{N \times N}$, which define a convex manifold.

In typical inference problems, \mathcal{A} has a non-trivial null space as the system of linear equations in (10) is severely underdetermined with $M \ll N^2$. Various studies approach the phase retrieval problem over the lifted domain, leveraging the low rank structure of the lifted ground truth $\tilde{\mathbf{X}} = \mathbf{x}\mathbf{x}^H$ and the subsequent LRMR theory from compressed sensing and matrix completion literature [40], [41]. The sufficient conditions on \mathcal{A} for exact recovery of $\mathbf{x}\mathbf{x}^H$ are primarily characterized by its null space [42]–[44] or restricted isometry properties on low rank [34], [41], [45] or PSD [2] matrices.

Knowing that (2) corresponds to the minimization of an ℓ_2 loss objective, WF exclusively iterates on the set of rank-1, PSD matrices by solving the following:

$$\underset{\mathbf{X}}{\text{minimize:}} \quad \frac{1}{2M} \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|^2 \quad \text{s.t.} \quad \mathbf{X} \in R_1^+, \quad (12)$$

where \mathbf{X} denotes the optimization variable in the lifted domain. The functional constraint on \mathbf{X} as rank-1, PSD matrix casts this minimization equivalent to minimizing over the signal domain variable \mathbf{z} , resulting with dimensionality reduction of the search space. This is practically enforced by a spectral projection within the gradient term ∇f , which can be expressed as

$$\nabla f(\mathbf{z}) = \frac{1}{M} \mathcal{A}^H \mathcal{A}(\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H) \mathbf{z}. \quad (13)$$

Beyond the immediate gains in practicality, the formulation in (12) reveals a theoretical advantage offered by the non-convex framework of WF, using which deterministic arguments for exact recovery similar to those of lifting-based methods can also be attained [27]. Moving from the convex relaxations of rank-minimization, WF corresponds to solving a non-convex feasibility problem, reminiscent of the optimizationless PhaseLift method in [46], and Uzawa's iterations in [40]. This yields an iterative scheme for the unrelaxed, non-convex form of the lifted problem, and enforces the rank-1, PSD structure on the

iterates. Thereby, the constraint set of WF is considerably less restrictive than those of the LRMR methods that solve for the lifted unknown $\mathbf{x}\mathbf{x}^H$.

Namely, the problem in (12) has a unique solution if there exists no element \mathbf{H} in the null space of \mathcal{A} , such that $\mathbf{x}\mathbf{x}^H + \mathbf{H}$ is a rank-1, PSD matrix. Observe that the validity of Condition I.1 and non-uniqueness of \mathcal{A} over the set of rank-1, PSD matrices result in a contradiction. Essentially, existence of a $\varphi \notin P$ that satisfies (1) ascertains either that the spectral initialization \mathbf{z}_0 is at a stationary point itself, or that the ϵ -neighborhood around the φ or P contains a stationary point, both of which violate the regularity condition. Hence, the degree to which \mathbf{Y} approximates $\mathbf{x}\mathbf{x}^H$ in the spectral norm sense as (9), also dictates the feasibility of violating the uniqueness condition of \mathcal{A} over the set of rank-1, PSD matrices. This is made further clear under the following observation.

Remark: The spectral matrix \mathbf{Y} in (6) is the *backprojection* estimate of the lifted ground truth, $\tilde{\mathbf{X}} = \mathbf{x}\mathbf{x}^H$, i.e.,

$$\mathbf{Y} = \frac{1}{M} \mathcal{A}^H(\mathbf{y}), \quad (14)$$

which, in the noise free case is $\mathbf{Y} = \frac{1}{M} \mathcal{A}^H \mathcal{A}(\mathbf{x}\mathbf{x}^H)$.

In practical terms, (9) becomes a condition on the *lifted normal operator*, $\mathcal{A}^H \mathcal{A}$. Thereby, the main result of this paper is that, if the following *concentration bound*

$$\left\| \frac{1}{M} \mathcal{A}^H \mathcal{A}(\mathbf{x}\mathbf{x}^H) - (\mathbf{x}\mathbf{x}^H + \|\mathbf{x}\|^2 \mathbf{I}) \right\| \leq \delta \|\mathbf{x}\|^2, \quad (15)$$

holds for all $\mathbf{x} \in \mathbb{C}^N$ with a δ that is sufficiently small, then via the special structure of the set of rank-1, PSD matrices, the iterations in (3) are guaranteed to converge to a solution in P via the restricted strong convexity of \mathcal{J} in $E(\epsilon)$. As δ gets smaller, the spectral initialization yields more accurate estimates due to favorable properties of the lifted normal operator over the set of rank-1, PSD matrices. Due to the fact that shrinkage on the value of δ is related to increasing the number of measurements M , the existence of sub-optimal minima accordingly vanishes. Hence, there exists a phase transition with respect to the value of δ , below which the tightness of the concentration bound can deterministically guarantee exact recovery from (1) using WF.

III. MAIN RESULTS

In this section, we prove (15) as a sufficient condition for exact phase retrieval for an arbitrary measurement model by establishing a set of lemmas through fully geometric arguments, given that it holds for all $\mathbf{x} \in \mathbb{C}^N$ with $\delta < 0.184$. Thereby, we use the form in (14) as a crucial element of our approach. Note that the lifted forward model in (11) may be a realization from a statistical model, or a deterministic measurement map. The following lemma characterizes the normal operator of the lifted forward model over the set of rank-1, positive semi-definite matrices.

Lemma III.1: Assume that (15) holds for an arbitrary $\mathbf{x} \in \mathbb{C}^N$. Then, the normal operator of the lifted forward model can be expressed as follows over the set of rank-1, PSD matrices:

$$\frac{1}{M} \mathcal{A}^H \mathcal{A} = \mathcal{I} + \mathcal{R} + \Delta, \quad (16)$$

where for any $\mathbf{z}\mathbf{z}^H \in R_1^+$, $\mathcal{R}(\mathbf{z}\mathbf{z}^H) = \|\mathbf{z}\|^2 \mathbf{I}$, with

$$\mathcal{R}(\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H) = (\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2) \mathbf{I}, \quad (17)$$

and $\Delta : R_1^+ \rightarrow \mathbb{C}^{N \times N}$ is a perturbation operator satisfying $\|\Delta(\mathbf{x}\mathbf{x}^H)\| \leq \delta \|\mathbf{x}\|^2$ for any $\mathbf{x} \in \mathbb{C}^N$ such that

$$\max_{\mathbf{v} \in \mathbb{C}^N \setminus \{\mathbf{0}\}} \frac{\|\Delta(\mathbf{v}\mathbf{v}^H)\|}{\|\mathbf{v}\mathbf{v}^H\|} \leq \delta, \quad (18)$$

for any deterministic model, or the Gaussian sampling model with high probability. \blacksquare

Proof: See Section V-A and Appendix A. \blacksquare

Specifically in the case of the Gaussian model, the operator $\mathcal{R} : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^{N \times N}$ characterizes the effect of the 4th moments of the sampling vectors. This term captures the diagonal bias of the spectral matrix in estimating the lifted signal, using the fact that the expectation of the lifted normal operator is linear on $\mathbb{C}^{N \times N}$.

We begin by considering the spectral initialization scheme. Namely, through Lemma III.1, the concentration bound in (15) indicates a proper scaling factor for the unit-norm eigenvector of \mathbf{Y} . This is derived in the following corollary which implies that the lifted forward model \mathcal{A} is a *tight frame*.

Corollary III.1: Assume that the assumptions of Lemma III.1 hold. Then, \mathcal{A} satisfies the following identity:

$$(2 - \delta) \|\mathbf{x}\mathbf{x}^H\|_F^2 \leq \frac{1}{M} \|\mathcal{A}(\mathbf{x}\mathbf{x}^H)\|^2 \leq (2 + \delta) \|\mathbf{x}\mathbf{x}^H\|_F^2. \quad (19)$$

Furthermore, if (19) holds for all $\mathbf{x} \in \mathbb{C}^N$, then the concentration bound in (15) equivalently holds uniformly over all $\mathbf{x} \in \mathbb{C}^N$. \blacksquare

Proof: See Section V-A. \blacksquare

As a result, $\lambda_0 = (2M)^{-1/2} \|\mathbf{y}\|$ is an estimator for the energy of the signal $\|\mathbf{x}\|^2$. Using the norm estimate obtained by Corollary III.1 for λ_0 , the distance of the spectral initialization yields the following ϵ -neighborhood as a function of the concentration parameter δ .

Lemma III.2: Assume that the assumptions of Lemma III.1 hold. Let \mathbf{z}_0 be the estimate $\mathbf{z}_0 = \sqrt{\lambda_0} \mathbf{v}_0$ where \mathbf{v}_0 is the eigenvector corresponding to the leading eigenvalue of the spectral matrix \mathbf{Y} in (14), and λ_0 is the signal energy estimate obtained as

$$\lambda_0 = \frac{1}{\sqrt{2M}} \|\mathbf{y}\|. \quad (20)$$

Then, the initial estimate \mathbf{z}_0 satisfies $\text{dist}^2(\mathbf{z}_0, \mathbf{x}) \leq \epsilon^2 \|\mathbf{x}\|^2$, where

$$\epsilon^2 = 1 + \sqrt{1 + \frac{\delta}{2}} - 2\sqrt{(1 - 2\delta) \left(1 + \frac{\delta}{2}\right)^{1/2}}. \quad (21)$$

Proof: See Section V-B. \blacksquare

Note that $E(\epsilon)$ is formed using (21) under the assumption that $\delta < 0.5$, which is required to have a valid estimate via spectral initialization. Next, we introduce the following lemma to characterize the relation between the distance metric introduced in (5), and the distance in the lifted domain.

Lemma III.3: Let $\mathbf{z} \in E(\epsilon)$ of \mathbf{x} , with $\epsilon = \epsilon_0 \|\mathbf{x}\|$ satisfying (21). Then, we have

$$h_1(\delta) \text{dist}(\mathbf{z}, \mathbf{x}) \|\mathbf{x}\| \leq \|\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H\|_F \leq h_2(\delta) \text{dist}(\mathbf{z}, \mathbf{x}) \|\mathbf{x}\|, \quad (22)$$

where $h_1 = \sqrt{(1-\epsilon)(2-\epsilon)}$, and $h_2 = (2+\epsilon)$. \blacksquare

Proof: See Section V-C.

Lemma III.3 states that the distance between the lifted signals $\mathbf{z}\mathbf{z}^H$ and $\mathbf{x}\mathbf{x}^H$ is of the rate of the distance of the signals $\mathbf{z}, \mathbf{x} \in \mathbb{C}^N$, when $\mathbf{z} \in E(\epsilon)$ of \mathbf{x} . Essentially, the distance metric in (5) locally tracks the error on the constraint set of rank-1, PSD matrices in the lifted domain. The outcome of this result, together with Lemma III.1, is the following *local restricted isometry*-type property.

Lemma III.4: Under the assumptions of Lemmas III.1 and III.3, for any $\mathbf{x} \in \mathbb{C}^N$ and $\mathbf{z} \in E(\epsilon)$, the lifted forward model satisfies

$$\begin{aligned} (1 - \hat{\delta})\|\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H\|_F^2 &\leq \frac{1}{M}\|\mathcal{A}(\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H)\|^2 \\ &\leq (2 + \hat{\delta})\|\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H\|_F^2, \end{aligned} \quad (23)$$

where $\hat{\delta} = \frac{\sqrt{2}(2+\epsilon)}{\sqrt{(1-\epsilon)(2-\epsilon)}}\delta$.

Proof: See Section V-D. \blacksquare

We refer to $\hat{\delta}$ as the local restricted isometry constant (RIC) of the lifted forward model over rank-2 matrices. However, note that this property holds over a very particular subset of rank-2 matrices even beyond the locality with respect to $\mathbf{x} \in \mathbb{C}^N$. Namely, (23) states that the distance between two elements $\mathbf{z}\mathbf{z}^H, \mathbf{x}\mathbf{x}^H \in R_1^+$ is approximately preserved under the mapping of \mathcal{A} , if $\mathbf{z} \in E(\epsilon)$ of \mathbf{x} . The significance of Lemma III.4 is the fact that the restricted isometry property in (23) is derived as a deterministic consequence of the concentration bound of the spectral matrix. The Lemmas III.3 and III.4 culminate to yield the Lipschitz differentiability of the objective function, stated in the following lemma:

Lemma III.5: Assume that the assumptions of Lemmas III.3 and III.4 hold. Then, for any $\mathbf{z} \in E(\epsilon)$, the objective function f in (2) is local Lipschitz differentiable at $\mathbf{x} \in P$ with

$$\|\nabla f(\mathbf{z})\| \leq c(\delta) \cdot \text{dist}(\mathbf{z}, \mathbf{x})\|\mathbf{x}\|^2 \quad (24)$$

where $c(\delta) = (1 + \epsilon)(2 + \epsilon)(2 + \delta)$ is the local Lipschitz constant. \blacksquare

Proof: See Section V-E. \blacksquare

Invoking the result of Lemma III.5, to establish the regularity condition for f , it is sufficient to show that for any $\mathbf{z} \in E(\epsilon)$

$$\text{Re}\left(\langle \nabla f(\mathbf{z}), (\mathbf{z} - \mathbf{x}e^{i\phi(\mathbf{z})}) \rangle\right) \geq \left(\frac{1}{\alpha} + \frac{c^2(\delta)\|\mathbf{x}\|^4}{\beta}\right) \text{dist}^2(\mathbf{z}, \mathbf{x}) \quad (25)$$

which is equivalent to the *restricted strong convexity* of the objective function in $E(\epsilon)$. By the definition of strong convexity around the closest solution $\hat{\mathbf{x}} = e^{i\Phi(\mathbf{z})}\mathbf{x}$ to an estimate $\mathbf{z} \in E(\epsilon)$, this condition, is implied if the objective function satisfies

$$f(\mathbf{z}) \geq f(\hat{\mathbf{x}}) + \text{Re}(\nabla f(\hat{\mathbf{x}})^H(\mathbf{z} - \hat{\mathbf{x}})) + \frac{L}{2}\|\mathbf{z} - \hat{\mathbf{x}}\|^2, \quad (26)$$

where L equals to the multiplier of the distance term in (25). Having $f(\hat{\mathbf{x}})$ and $\nabla f(\hat{\mathbf{x}})$ equal 0 by definition, the restricted strong convexity in $E(\epsilon)$ is simply reduced to the following

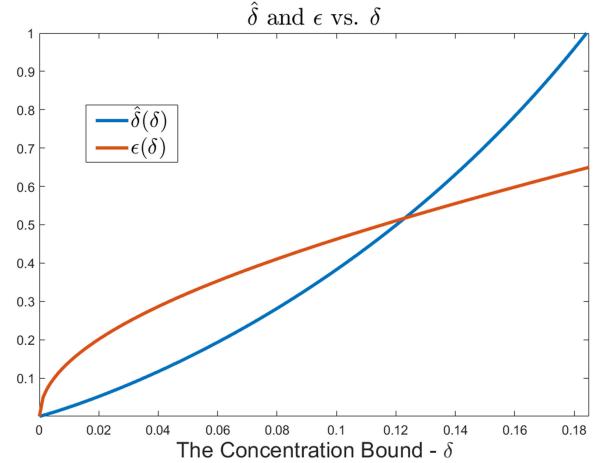


Fig. 1. The $\hat{\delta}$ and ϵ values in the region that the regularity condition is guaranteed to hold. Observe that the limiting value is the local restricted isometry constant, $\hat{\delta}$, which controls the uniqueness property in the lifted problem locally.

condition:

$$f(\mathbf{z}) \geq \frac{1}{2} \left(\frac{1}{\alpha} + \frac{c^2(\delta)\|\mathbf{x}\|^4}{\beta} \right) \text{dist}^2(\mathbf{z}, \mathbf{x}), \quad (27)$$

for any $\mathbf{z} \in E(\epsilon)$. Writing f explicitly in terms of the lifted signals as $f(\mathbf{z}) = \|\mathcal{A}(\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H)\|^2/2M$, and applying the lower bounds from Lemmas III.3 and III.4, we have

$$f(\mathbf{z}) \geq \frac{(1 - \hat{\delta})h_1^2(\delta)}{2} \text{dist}^2(\mathbf{z}, \mathbf{x})\|\mathbf{x}\|^2. \quad (28)$$

Thus, the regularity condition is satisfied by setting α and β such that $\alpha\beta > 4$, and

$$\frac{1}{\alpha\|\mathbf{x}\|^2} + \frac{c^2(\delta)\|\mathbf{x}\|^2}{\beta} \leq (1 - \hat{\delta})h_1^2(\delta) := h(\delta). \quad (29)$$

The final form we derive in (29) results in a number of notable outcomes regarding the non-convex optimization theory of the WF framework:

1) We show that there exists a regime in which the regularity condition holds by default. This regime is characterized by the concentration bound of the spectral matrix in (15), as $\hat{\delta}$ is solely a function of δ . Observe that the validity of this regime depends on attaining $\hat{\delta} < 1$ which constrains the tightness of the concentration property in (15). This numerically yields an upper bound of $\delta \leq 0.184$ as shown in Fig. 1.

2) (29) provides an interpretation of the algorithm parameters consistent with the original work of [1]: Fig. 2 demonstrates the range of values the constants c and h can attain in the valid region of δ . Notably, the values of these $\mathcal{O}(1)$ constants characterize the convergence rate of the algorithm, as α and β are required to be sufficiently large constants for (29) to hold. Observe that (29) implies setting $\alpha = \mathcal{O}(1/\|\mathbf{x}\|^2)$, and $\beta = \mathcal{O}(\|\mathbf{x}\|^2)$, hence $\alpha\beta = \mathcal{O}(1)$. Since we clearly have $h < 2$, and $c > 4$, the condition in (29) holds with $\alpha\beta > 4$ by definition. Hence, the regularity condition is satisfied, and a step size $\mu' \leq 2/\beta$ can be chosen to yield a convergence rate of $2\mu'/\alpha$ via [1]. This step size μ' is then $\mathcal{O}(1/\|\mathbf{x}\|^2)$. Hence, the definition of the updates

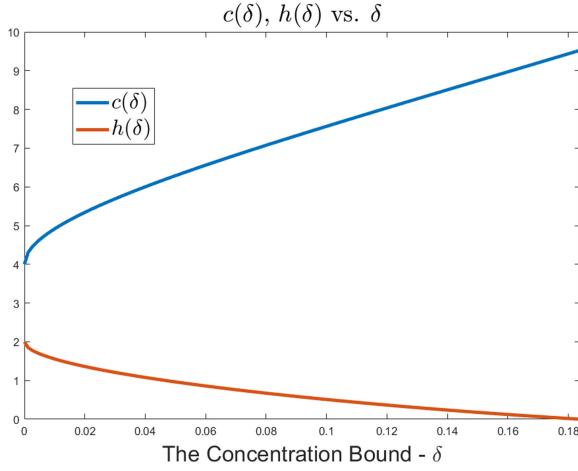


Fig. 2. The c and h values in the region that the regularity condition is guaranteed to hold. The two constants characterize the values of α and β parameters for the convergence of the iterates generated by WF updates.

in (3) requires an approximate normalization term¹ of $\|\mathbf{z}_0\|^2$ on a scalar entity $\mu_k = \mathcal{O}(1)$.

3) The condition in (29) effectively places an upper bound on the convergence rate of the algorithm under which exact recovery of any $\mathbf{x} \in \mathbb{C}^N$ is guaranteed. Simply fixing $\kappa = \alpha\beta$ and re-organizing (29), we obtain the following:

$$\frac{1}{\alpha\|\mathbf{x}\|^2} \left(1 + \frac{c^2(\delta)\|\mathbf{x}\|^4}{\kappa} \alpha^2 - h(\delta)\|\mathbf{x}\|^2\alpha \right) \leq 0. \quad (30)$$

Since we have $\alpha > 0$ by definition and $h(\delta) > 0$ by constraining δ , it suffices to consider the non-negativity of the discriminant of the quadratic equation with respect to α in (30) for the overall condition to hold, which yields

$$h^2(\delta)\|\mathbf{x}\|^4 - \frac{4}{\kappa}c^2(\delta)\|\mathbf{x}\|^4 \geq 0. \quad (31)$$

As a result, knowing that $4/\kappa$ is an upper bound on the rate of convergence of the algorithm via [1], we obtain the best achievable geometric convergence rate as a function of the concentration bound parameter as

$$\frac{4}{\alpha\beta} \leq \frac{h^2(\delta)}{c^2(\delta)} := r(\delta) = \left(\frac{(1 - \hat{\delta})(1 - \epsilon)(2 - \epsilon)}{(2 + \delta)(1 + \epsilon)(2 + \epsilon)} \right)^2. \quad (32)$$

Thereby, beyond directly guaranteeing the existence of a pair of (α, β) to satisfy the regularity condition when sufficiently small, the δ -value fully characterizes the practicality and iteration complexity of the algorithm via $r(\delta)$ in (32). Note that $r(\delta)$ is the best convergence rate the algorithm can achieve uniformly over all $\mathbf{x} \in \mathbb{C}^N$.

Organizing the arguments developed in this section, we state the following for exact phase retrieval via Wirtinger Flow:

Theorem III.1: Assume that the concentration property

$$\left\| \frac{1}{M} \mathcal{A}^H \mathcal{A}(\mathbf{x}\mathbf{x}^H) - (\mathbf{x}\mathbf{x}^H + \|\mathbf{x}\|^2 \mathbf{I}) \right\| \leq \delta \|\mathbf{x}\|^2.$$

¹The approximation argument can be followed from [1] in the proof of Lemma 7.10.

holds for all $\mathbf{x} \in \mathbb{C}^N$ with $\delta \leq 0.184$. Then, the initial estimate \mathbf{z}_0 obtained from the spectral matrix in (6) using the scaling factor in (20) satisfies

$$\text{dist}^2(\mathbf{z}_0, \mathbf{x}) \leq \epsilon^2 \|\mathbf{x}\|^2,$$

for all \mathbf{x} , where $\epsilon^2 \leq 1 + \sqrt{1 + \frac{\delta}{2}} - 2\sqrt{(1 - 2\delta)(1 + \frac{\delta}{2})^{1/2}}$, and for the iterates generated by (3) with a fixed step size of $\mu/\|\mathbf{z}_0\|^2 \leq 2/\beta$ we have that

$$\text{dist}^2(\mathbf{z}_k, \mathbf{x}) \leq \epsilon^2 \left(1 - \frac{2\mu}{\alpha} \right)^k \|\mathbf{x}\|^2,$$

with the best achievable convergence rate of

$$\frac{2\mu}{\alpha} \leq \frac{h^2(\delta)}{c^2(\delta)} := \left(\frac{(1 - \hat{\delta})(1 - \epsilon)(2 - \epsilon)}{(2 + \delta)(1 + \epsilon)(2 + \epsilon)} \right)^2,$$

for the recovery of any $\mathbf{x} \in \mathbb{C}^N$, where $\hat{\delta} = \frac{\sqrt{2}(2+\epsilon)}{\sqrt{(1-\epsilon)(2-\epsilon)}}\delta$.

Remark: For a deterministic model, (15) to hold for an unspecified, arbitrary $\mathbf{x} \in \mathbb{C}^N$ is equivalent to holding over all $\mathbf{x} \in \mathbb{C}^N$ [47]. However in the probabilistic setting, the probability that (15) holds uniformly over all $\mathbf{x} \in \mathbb{C}^N$ is naturally more stringent than it is for an arbitrary \mathbf{x} . In Appendix A, we provide the uniformity of (15) for the Gaussian sampling model when $M = \mathcal{O}(N \log N)$.

Complementing our performance guarantees in Theorem III.1, we can further derive an optimal step-size that achieves our upper bound on the rate of convergence in (32). Knowing that $\beta = \beta'/\|\mathbf{x}\|^2$, $\alpha = \alpha'/\|\mathbf{x}\|^2$ with $\mu_k = \mu \leq 2/\beta'$ defined as the largest step-size allowed, and using (29) we have that:

$$\frac{2(h(\delta) - 1/\alpha')}{c^2(\delta)} = \frac{r(\delta)\alpha'}{2}, \quad (33)$$

when $2\mu/\alpha$ attains $r(\delta)$ with equality. Solving for α' yields $\alpha' = 2/h(\delta)$, which translates to a constant step-size of $\mu(\delta) = h(\delta)/c^2(\delta)$. Hence the WF iterations provably converge to a true solution at a constant step-size $\mu_k = \mu = \mathcal{O}(1)$ with $\mu > 0$ as long as $\delta \leq 0.184$ in (15).

The evolution of the optimal convergence rate $r(\delta)$ and the step-size $\mu(\delta)$ within our exact recovery regime is provided in Fig. 3. The characterization of the convergence rate and the optimal step-size with respect to δ highlights a key contribution of our deterministic framework in Theorem III.1. Specifically considering the Gaussian model at a fixed number of samples at the proper complexity, picking a smaller δ is equivalent to relaxing the probability of success for exact recovery. This in turn indicates moving up on the $\mu(\delta)$ curve such that the algorithm is ran with a larger step-size corresponding to the smaller δ value. Therefore in practical terms, the behavior in Fig. 3 quantifies the trade-off between rate of convergence of the algorithm, and the probability of exactly recovering the ground truth $\mathbf{x} \in \mathbb{C}^N$. Alternatively an analogous phenomenon characterizes the trade-off between the number of samples and the algorithm performance, for a desired fixed probability of success. As a result, we establish an explicit relationship between the step-size chosen in practice, and the underlying model parameters such as the number of

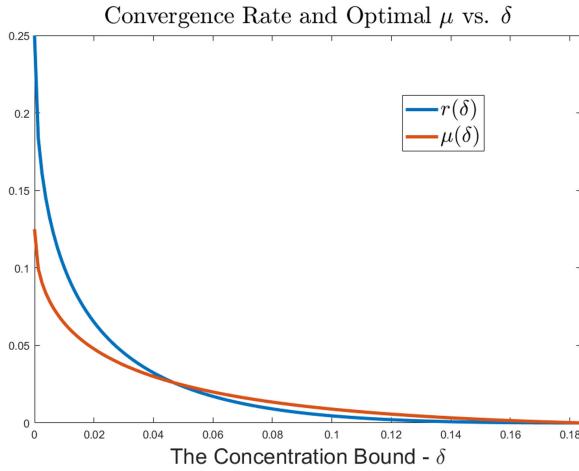


Fig. 3. The $r(\delta)$ upper bound on rate of converge and the optimal step size $\mu(\delta)$ with respect to the concentration bound parameter δ in the exact recovery regime of our framework. The optimal step-size is defined as the μ that satisfies the upper bound on the rate-of convergence.

samples, and the probability of success for non-convex phase retrieval from intensity measurements via WF.

Note that the uniformity of (15) in the complex Gaussian sampling model was shown when $M = \mathcal{O}(N)$ in [15] via sample truncation, however not as a sufficient condition for exact recovery in the manner we established in Theorem III.1. Our framework essentially captures the impact of such a sample truncation scheme through its relation to the properties of the lifted forward model. In this sense, one could consider Theorem III.1 as an abstraction of the theoretical guarantees in [15], where we have identified a novel, more minimal sufficient condition for arbitrary lifted forward models. As a result of such an abstraction, we extend the favorable properties attained in [15] such as exact recovery with an $\mathcal{O}(1)$ step-size and the linear computational complexity to a more general problem setting which includes that of the original WF. It can also be observed that the numerical cases evaluated in [1] for the Gaussian model are well within the range of values we identify for the validity of our Theorem III.1. Hence, the theoretical means developed for Theorem III.1 are consistent with the convergence behavior demonstrated in [1].

In establishing Theorem III.1 for exact phase retrieval, we necessarily use the specific structure of the diagonal bias term in the expectation of the spectral matrix. This is in contrast to our work in [27], in which the spectral matrix is an unbiased estimator of the lifted signal. Nonetheless, Corollary III.1 highlights a key advantage of the non-convex framework of WF. Via the removal of convex relaxations and solving the perturbed problem in the lifted domain over the set of rank-1, PSD matrices, the RIP-type properties required by semi-definite programming and lifting-based approaches are relaxed to smaller, more specific domain of matrices. Under the lens of LRMR theory, WF not only offers computational advantages, but also less stringent theoretical means to achieve exact recovery if the step-size is properly controlled. The deterministic and less stringent nature of our recovery guarantees also opens up promising possibilities

for the study of more structured models for problems such as wave-based imaging, where estimates on δ value would relate to parameters such as bandwidth, central frequency, or resolution [8], [47].

Overall, we further stress a few notable outcomes of Theorem III.1, which include the following observations:

- Via the established deterministic convergence framework given the concentration bound, our result proves that the restricted strong convexity property of the objective function is achieved with $\mathcal{O}(N \log N)$ samples for the complex Gaussian sampling model. This is a $\log N$ factor less than the sample complexity reported in [30].
- Our sufficient condition has to hold only over the rank-1, PSD matrices, which is less stringent than those studied in the non-convex LRMR literature. Additionally, a universal upper bound on the relative distance- ϵ via Fig. 1 is attained within our exact recovery regime. Hence, the concept of sufficient accuracy of the spectral initialization is captured by a quantitative measure.
- Another observation is that the upper bound on the concentration property of the spectral matrix in the phase retrieval problem requires a smaller constant than the one in the *interferometric inversion problem* we studied in [27] (0.184 as opposed to 0.214), in which the relative phase information of a pair of measurements is retained. This is indeed an intuitive outcome, as more information is lost when measurements are phaseless, compared to the interferometric case. A similar outcome is observed in the upper bound obtained for the geometric convergence rate of the algorithm, which approaches to 0 for the case of interferometric inversion as $\delta \rightarrow 0$. As a result, the impact of the additional loss of phase information is directly captured in the sufficient conditions and the performance guarantees of the algorithm in solving the different types of quadratic systems of equations.

IV. ROBUSTNESS

In this section, we assess the robustness of the WF algorithm in the presence of additive noise in the measurements. We show that for the general problem setting of

$$\mathbf{y} = \mathcal{A}(\mathbf{x}\mathbf{x}^H) + \boldsymbol{\eta}, \quad (34)$$

the results presented in Theorem III.1 for the noise-free case in Section III are attained upto a bounded perturbation for $\mathbb{E}[\boldsymbol{\eta}] = 0$. It should be noted that, the ℓ_2 mismatch function minimized in the problem formulation fits the data for i.i.d. additive white Gaussian noise model $\{\eta_m\}_{m=1}^M$ in the maximum likelihood sense. Despite this, the results presented in this section have no specification on the distribution of the noise term $\boldsymbol{\eta}$, similar to those of [15], which were derived for the Poisson loss function.

Our first goal is to establish the validity of the spectral initialization for our exact recovery guarantees with respect to the SNR of measurements in (34) by utilizing the arguments developed over the lifted domain. Namely, for our subsequent convergence theory to hold, we derived numerical constraints on both the concentration bound (i.e., δ), and the distance of

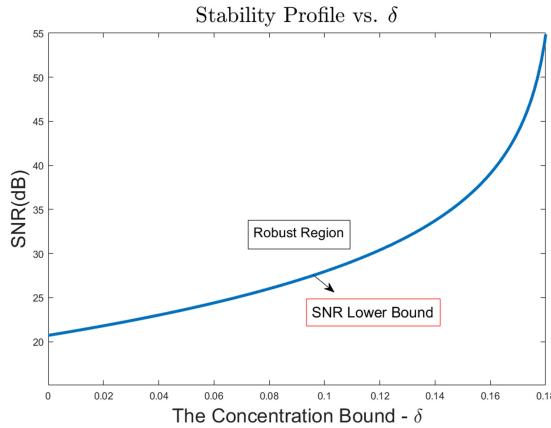


Fig. 4. The SNR lower bound required for our convergence framework to hold. Below the lower bound spectral initialization violates the sufficient conditions for Theorem IV.1. In this case, pre-processing techniques or other initialization schemes must be pursued.

the initial estimate obtained from the spectral method (i.e., ϵ). These constraints characterize the amount of perturbation the algorithm can tolerate, which is stated in the following lemma and shown in Fig. 4.

Lemma IV.1: Consider the spectral matrix formed by (6), using the noisy measurements in (34). Moreover, let the concentration bound in (15) hold as stated in the setup of Theorem III.1. Then, the spectral matrix \mathbf{Y} satisfies,

$$\mathbb{E}_{\boldsymbol{\eta}} [\|\mathbf{Y} - (\mathbf{x}\mathbf{x}^H + \|\mathbf{x}\|^2\mathbf{I})\|] \leq \left(\delta + \frac{(2 + \delta)}{\sqrt{\text{SNR}}} \right) \|\mathbf{x}\|^2, \quad (35)$$

where SNR stands for signal-to-noise-ratio, defined as $\text{SNR} = \|\mathcal{A}(\mathbf{x}\mathbf{x}^H)\|^2 / \mathbb{E}[\|\boldsymbol{\eta}\|^2]$.

Proof: See Section V-F. \blacksquare

Analogous to the one-to-one relationship of the ϵ -distance of the spectral initialization and the concentration bound δ in the noise-free phase retrieval problem, the noisy scenario has the additional dependence on the SNR of the measurements through $\tilde{\delta} := \delta + (2 + \delta)/\sqrt{\text{SNR}}$. With the presence of the SNR term, there exists a level of noise as a function of δ beyond which the concentration bound in (35) is insufficient to guarantee an effective spectral initialization. This restriction is directly determined by two constraints: $\epsilon < 1$ and $\tilde{\delta} < 1$, in order to retain a valid regime where convergence arguments from Theorem III.1 hold true for the noise-free component of the gradients. Thereby, we obtain a region over the (δ, SNR) domain such that the spectral method produces a valid estimator under $\tilde{\delta}$.

A direct manner to determine this region is by enforcing $\tilde{\delta} \leq 0.184$, through which accuracy of spectral initialization and subsequent arguments within the $E(\epsilon)$ are preserved, yielding an SNR lower bound of

$$\text{SNR(dB)} \geq 20 \log \frac{2 + \delta}{0.184 - \delta}. \quad (36)$$

Fig. 4 depicts that the lowest SNR value of 20.7dB is attained at $\delta = 0$. Although it is derived in a straight-forward manner, (36) is the best lower bound that can be characterized by our framework. This is rather surprising at first glance, since the $\tilde{\delta}$ term is only

affected by the perturbation resulting from noise through the ϵ parameter, as δ and $\tilde{\delta}$ are properties of the underlying lifted forward model \mathcal{A} , which is independent of noise. However, these still prove to be consequential for the stability of the algorithm under additive noise because of constraints that arise from the convergence arguments, beyond those related to the validity of the spectral initialization.

In particular, under additive noise and the assumptions of Lemma IV.1, within the valid region for the spectral initialization defined by (36), the convergence guarantees of WF are perturbed by a constant factor that is a function of SNR.

Theorem IV.1: Assume that the assumptions of Lemma IV.1 hold. Then, for the identical procedure and values of constants α, β in the setup of Theorem III.1, we have

$$\mathbb{E}_{\boldsymbol{\eta}} [\text{dist}(\mathbf{z}_k, \mathbf{x})] \lesssim \epsilon \left(1 - \frac{2\mu}{\alpha} \right)^{\frac{k}{2}} \|\mathbf{x}\| + \alpha' \frac{(2 + \delta)}{\sqrt{\text{SNR}}} \|\mathbf{x}\|, \quad (37)$$

where $\mu = \mu_k / \|\mathbf{z}_0\|^2 \leq 2/\beta$, $\alpha' = \mathcal{O}(1)$ such that $\alpha = \alpha' / \|\mathbf{z}_0\|^2$, and $\epsilon^2 \leq 1 + \sqrt{1 + \frac{\tilde{\delta}}{2}} - 2\sqrt{(1 - 2\tilde{\delta})(1 + \frac{\tilde{\delta}}{2})^{1/2}}$, with $\tilde{\delta} = \delta + \frac{(2 + \delta)}{\sqrt{\text{SNR}}}$.

Proof: See Section V-G. \blacksquare

As a result of Theorem IV.1, we observe a crucial element for determining the trade-off between the α and β parameters. The stability guarantees directly incentivize allocating a small value for the parameter α , which is inversely related to the magnitude of the β parameter. Since $\mu \leq 2/\beta$ by definition, a tighter stability bound requires a trade-off from the step size of the algorithm. This outcome is indeed expected, as lower SNR in the measurements means more variance for the gradient estimates at the iterative refinement stage, hence one should take less confident steps to counter inaccurate update terms. Our framework perfectly captures this phenomenon, and requires small step sizes for improved stability in the algorithm performance while optimizing the noisy landscape of the objective function over the signal domain.

Furthermore, to guarantee that the iterates formed using the noisy measurements remain in the $E(\epsilon)$, there is an effective upper bound on the α parameter such that

$$\alpha \|\mathbf{x}\|^2 \leq \frac{\epsilon \sqrt{\text{SNR}}}{(1 + \epsilon)(2 + \delta)}. \quad (38)$$

Equivalently, this is a lower bound on the value of $1/(\alpha \|\mathbf{x}\|^2)$ in (29), which requires

$$\frac{(1 + \epsilon)(2 + \delta)}{\epsilon \sqrt{\text{SNR}}} \leq \frac{1}{\alpha \|\mathbf{x}\|^2} < (1 - \hat{\delta})(1 - \epsilon)(2 - \epsilon) \quad (39)$$

to be satisfied for a finite β to exist to attain a practical step-size for the algorithm. Thereby, on expectation, iterative updates that are contractions with respect to the distance metric can be achieved, if SNR is sufficiently high to satisfy

$$\frac{(1 + \epsilon)(2 + \delta)}{\epsilon \sqrt{\text{SNR}}} < (1 - \hat{\delta})(1 - \epsilon)(2 - \epsilon), \quad (40)$$

where both ϵ and $\hat{\delta}$ are SNR dependent. Since the left- and right-hand-sides of the inequality monotonically decrease and

increase, respectively, with increasing SNR, there exists a transition point at any fixed $\delta < 0.184$ value beyond which the inequality holds true. This numerical characterization of the SNR requirements for the convergence of WF precisely corresponds to the lower bound in (36).

Beyond our characterization of the (δ, SNR) domain, pre-processing techniques developed in the phase retrieval literature [17], [22], [25] can be adopted to improve the SNR bound by relaxing our requirements on δ through the value of ϵ . As a result, alternative techniques that improve the accuracy of the spectral estimation can be pursued for practical purposes in order to sustain the convergence guarantees given our sufficient condition below the provided SNR lower bound.

V. PROOFS

In this section we present the proofs of our arguments used in establishing Theorem III.1, and its corollaries. Notably, our results are derived in a deterministic framework based on geometric arguments unlike the probabilistic theory of original WF theory [1]. The probabilistic nature of the convergence analysis prominent in phase retrieval literature is thereby compressed into a single condition on the lifted forward model. From this point on, we frequently use $\tilde{\mathbf{V}} = \mathbf{v}\mathbf{v}^H$ to denote the lifted signal corresponding to an element $\mathbf{v} \in \mathbb{C}^N$ for notational brevity.

A. Proofs of Lemma III.1 and Corollary III.1

1) *Lemma III.1:* Reprising (15), from the definition of the lifted forward model \mathcal{A} , and spectral matrix \mathbf{Y} we have for any $\mathbf{x} \in \mathbb{C}^N$:

$$\left\| \frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{X}}) - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I}) \right\| \leq \delta \|\mathbf{x}\|^2. \quad (41)$$

Over the set of rank-1 matrices $R_1 = \{\mathbf{u}\mathbf{v}^H : \mathbf{u}, \mathbf{v} \in \mathbb{C}^N\}$, we define the operator $\mathcal{R} : R_1 \rightarrow \text{span}(\mathbf{I}) \subset \mathbb{C}^{N \times N}$, such that $\mathcal{R}(\mathbf{u}\mathbf{v}^H) = (\mathbf{v}^H \mathbf{u})\mathbf{I}$. Then, we define $(\frac{1}{M} \mathcal{A}^H \mathcal{A} - \mathcal{I} - \mathcal{R})(\tilde{\mathbf{X}}) = \Delta(\tilde{\mathbf{X}})$, and by (41) we have that $\|\Delta(\tilde{\mathbf{X}})\| \leq \delta \|\mathbf{x}\|^2$.

Next, we represent the rank-1, PSD matrix $\tilde{\mathbf{Z}} = \mathbf{z}\mathbf{z}^H$ as a linear combination of rank-1 elements in R_1 . Letting $\mathbf{e} = \mathbf{z} - \mathbf{x}$, we have

$$\tilde{\mathbf{E}} = \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} - \mathbf{e}\mathbf{e}^H - \mathbf{x}\mathbf{e}^H, \quad (42)$$

hence $\tilde{\mathbf{Z}} = \tilde{\mathbf{E}} + \mathbf{e}\mathbf{e}^H + \mathbf{x}\mathbf{e}^H + \tilde{\mathbf{X}}$. In the range of \mathcal{R} , distributing over the terms on the right hand side we obtain

$$\mathcal{R}(\tilde{\mathbf{Z}}) = (\|\mathbf{e}\|^2 + \mathbf{x}^H \mathbf{e} + \mathbf{e}^H \mathbf{x} + \|\mathbf{x}\|^2) \mathbf{I}, \quad (43)$$

which, from the definition of \mathbf{e} , precisely equals to $\|\mathbf{z}\|^2 \mathbf{I}$. Thus having $R_1^+ \subset R_1$, by moving \mathbf{x} to the left-hand-side \mathcal{R} satisfies,

$$\mathcal{R}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}) = \|\mathbf{z}\|^2 - \|\mathbf{x}\|^2. \quad (44)$$

Since $\mathcal{I} + \mathcal{R}$ is linear over elements in R_1^+ , Δ is necessarily linear as well knowing that $\mathcal{A}^H \mathcal{A}$ is linear over all $\mathbb{C}^{N \times N}$. In addition, since $\langle \mathcal{R}(\tilde{\mathbf{X}}), \tilde{\mathbf{X}} \rangle_F = \|\tilde{\mathbf{X}}\|_F^2$, \mathcal{R} is self-adjoint, hence Δ is a Hermitian operator by definition. Finally, since

$\|\Delta(\tilde{\mathbf{X}})\| \leq \delta \|\mathbf{x}\|^2$, we have

$$\max_{\mathbf{v} \in \mathbb{C}^N \setminus \{0\}} \frac{|\mathbf{v}^H \Delta(\tilde{\mathbf{X}}) \mathbf{v}|}{\|\mathbf{v}\|^2} \leq \delta \|\mathbf{x}\|^2. \quad (45)$$

Equivalently, utilizing the view in the lifted domain, we have, $|\mathbf{v}^H \Delta(\tilde{\mathbf{X}}) \mathbf{v}| = |\langle \Delta(\tilde{\mathbf{X}}), \tilde{\mathbf{V}} \rangle_F| = |\langle \Delta(\tilde{\mathbf{V}}), \tilde{\mathbf{X}} \rangle_F|$. Hence the concentration property directly implies uniformity of the bound over the set of the domain of operator Δ , i.e., the set of rank-1, PSD matrices as

$$\max_{\mathbf{v} \in \mathbb{C}^N \setminus \{0\}} \frac{|\mathbf{v}^H \Delta(\mathbf{v}\mathbf{v}^H) \mathbf{v}|}{\|\mathbf{x}\|^2 \|\mathbf{v}\|^2} \leq \delta, \quad (46)$$

again, for any fixed $\mathbf{x} \in \mathbb{C}^N$. For an arbitrary deterministic map that already satisfies (41), this completes the proof as the condition also holds for $\hat{\mathbf{v}}$ that maximizes $|\langle \Delta(\tilde{\mathbf{V}}), \tilde{\mathbf{X}} \rangle_F| / \|\mathbf{x}\|^2$ as $|\langle \Delta(\tilde{\mathbf{V}}), \hat{\mathbf{v}}\hat{\mathbf{v}}^H \rangle_F| / \|\hat{\mathbf{v}}\|^2$, which, by definition, is the spectral norm of $\|\Delta(\mathbf{v}\mathbf{v}^H)\|$. For uniformity in the complex Gaussian model, we refer the reader to Appendix A.

2) *Corollary III.1:* The proof of Corollary III.1 then directly follows from (15) and the definition of the spectral matrix. From (41) and the definition of the spectral norm, we have

$$\frac{|\mathbf{v}^H \left(\frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{X}}) - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I}) \right) \mathbf{v}|}{\|\mathbf{v}\|^2} \leq \delta \|\mathbf{x}\|^2, \quad (47)$$

for any $\mathbf{v} \in \mathbb{C}^N \setminus \{0\}$. Hence, for $\mathbf{v} = \mathbf{x}$, using the representation in the lifted domain via the definition of the Frobenius inner product, we have

$$\begin{aligned} |\langle \frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{X}}) - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I}), \mathbf{x}\mathbf{x}^H \rangle_F| &\leq \delta \|\mathbf{x}\|^4, \\ |\frac{1}{M} \|\mathcal{A}(\tilde{\mathbf{X}})\|^2 - \|\tilde{\mathbf{X}}\|_F^2 - \langle \|\mathbf{x}\|^2 \mathbf{I}, \tilde{\mathbf{X}} \rangle_F| &\leq \delta \|\tilde{\mathbf{X}}\|_F^2. \end{aligned} \quad (48)$$

From the definition of the Frobenius inner product with the identity matrix \mathbf{I} in $\mathbb{C}^{N \times N}$, only the diagonal elements are multiplied and summed with $\tilde{\mathbf{X}} = \mathbf{x}\mathbf{x}^H$, which corresponds to $\|\mathbf{x}\|^2$. Hence, we get, for any $\mathbf{x} \in \mathbb{C}^N$,

$$\left| \frac{1}{M} \|\mathcal{A}(\tilde{\mathbf{X}})\|^2 - 2\|\tilde{\mathbf{X}}\|_F^2 \right| \leq \delta \|\tilde{\mathbf{X}}\|_F^2. \quad (49)$$

Thus, the condition in Corollary III.1 is directly implied by (41). Furthermore, again using the fact that Δ is Hermitian, consider the maxima over the unit sphere $\|\mathbf{x}\| = 1$ as

$$\max_{\|\mathbf{x}\|=1} \|\Delta(\mathbf{x}\mathbf{x}^H)\| = \max_{\|\mathbf{x}\|=1} \max_{\|\mathbf{v}\|=1} |\mathbf{v}^H \Delta(\mathbf{x}\mathbf{x}^H) \mathbf{v}| \quad (50)$$

$$= \max_{\|\mathbf{x}\|=1} \max_{\|\mathbf{v}\|=1} |\langle \sqrt{\Delta}(\mathbf{x}\mathbf{x}^H), \sqrt{\Delta}(\mathbf{v}\mathbf{v}^H) \rangle_F|, \quad (51)$$

where $\sqrt{\Delta} : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^{N \times N}$ is a (non-unique) square root of the operator Δ . Then from Cauchy-Schwartz we have that

$$\max_{\|\mathbf{x}\|=1} \|\Delta(\mathbf{x}\mathbf{x}^H)\| \leq \max_{\|\mathbf{x}\|=1} \max_{\|\mathbf{v}\|=1} \|\sqrt{\Delta}(\mathbf{x}\mathbf{x}^H)\|_F \|\sqrt{\Delta}(\mathbf{v}\mathbf{v}^H)\|_F \quad (52)$$

in which we clearly have a maximization that is fully split such that

$$\max_{\|\mathbf{x}\|=1} \|\Delta(\mathbf{x}\mathbf{x}^H)\| \leq \max_{\|\mathbf{x}\|=1} \|\sqrt{\Delta}(\mathbf{x}\mathbf{x}^H)\|_F^2. \quad (53)$$

From our construction of $\sqrt{\Delta}$, the upper bound in (53) yields

$$\max_{\|\mathbf{x}\|=1} \|\Delta(\mathbf{x}\mathbf{x}^H)\| \leq \max_{\|\mathbf{x}\|=1} |\mathbf{x}^H \Delta(\mathbf{x}\mathbf{x}^H) \mathbf{x}|. \quad (54)$$

Since $\|\Delta(\mathbf{x}\mathbf{x}^H)\| = \max_{\|\mathbf{v}\|=1} |\mathbf{v}^H \Delta(\mathbf{x}\mathbf{x}^H) \mathbf{v}| \geq |\mathbf{x}^H \Delta(\mathbf{x}\mathbf{x}^H) \mathbf{x}|$ by definition, (54) implies that the maxima necessarily occurs at $\mathbf{v} = \mathbf{x}/\|\mathbf{x}\|$, as

$$\max_{\|\mathbf{x}\|=1} \|\Delta(\mathbf{x}\mathbf{x}^H)\| = \max_{\|\mathbf{x}\|=1} |\langle \Delta(\mathbf{x}\mathbf{x}^H), \mathbf{x}\mathbf{x}^H \rangle_F|. \quad (55)$$

Hence, the concentration bound holds uniformly with a constant δ that is equivalent to the restricted isometry constant of the lifted forward model over the set of rank-1, PSD matrices.

B. Proof of Lemma III.2

As shown in [1] we know that (15) implies $|\mathbf{v}_0^H \mathbf{x}|^2 \geq (1-2\delta)\|\mathbf{x}\|^2$, where $\|\mathbf{v}_0\| = 1$ is the leading eigenvector of the spectral matrix \mathbf{Y} . Using Corollary III.1, we know that $\lambda_0 = \|\mathbf{y}\|/\sqrt{2M}$ is an estimate for the energy of the unknown signal \mathbf{x} such that $(\sqrt{1-\delta/2})\|\mathbf{x}\|^2 \leq \lambda_0 \leq (\sqrt{1+\delta/2})\|\mathbf{x}\|^2$. Using the definition of the distance metric, and the lower bound from [1], for $\mathbf{z}_0 = \sqrt{\lambda_0}\mathbf{v}_0$, we have

$$\text{dist}^2(\mathbf{z}_0, \mathbf{x}) \leq \left(\frac{\lambda_0}{\|\mathbf{x}\|^2} + 1 - 2\sqrt{\frac{\lambda_0}{\|\mathbf{x}\|^2}\|\mathbf{x}\|\sqrt{1-2\delta}} \right) \|\mathbf{x}\|^2. \quad (56)$$

Since the right-hand-side is a convex quadratic function of $\sqrt{\lambda_0}$, its maximum value is reached at the boundary values of $\sqrt{\lambda_0}$. Setting $\lambda_0 = (\sqrt{1+\delta/2})$, the upper bound is monotonically greater than at $\lambda_0 = (\sqrt{1-\delta/2})$ for all valid values for δ , hence we conclude that

$$\text{dist}^2(\mathbf{z}_0, \mathbf{x}) \leq \left(\sqrt{1 + \frac{\delta}{2}} + 1 - 2\sqrt{(1-2\delta)\sqrt{1 + \frac{\delta}{2}}} \right) \|\mathbf{x}\|^2. \quad (57)$$

C. Proof of Lemma III.3

1) *Proof of the Upper Bound:* Let $\hat{\mathbf{x}}$ be the closest solution in P to an arbitrary $\mathbf{z} \in E(\epsilon)$. From reverse triangle inequality we have $(1-\epsilon)\|\mathbf{x}\| \leq \|\mathbf{z}\| \leq (1+\epsilon)\|\mathbf{x}\|$. Setting $\mathbf{e} = \mathbf{z} - \hat{\mathbf{x}}$, by (42) we have

$$\tilde{\mathbf{Z}} - \tilde{\mathbf{X}} = \tilde{\mathbf{E}} + \mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H. \quad (58)$$

Then, for the Frobenius norm of the error in the lifted domain, we have

$$\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \leq \|\tilde{\mathbf{E}}\|_F + \|\mathbf{e}\hat{\mathbf{x}}^H\|_F + \|\hat{\mathbf{x}}\mathbf{e}^H\|_F. \quad (59)$$

Since all the elements on the right-hand-side are rank-1, and $\mathbf{z} \in E(\epsilon)$, by definition, we have $\|\cdot\| = \|\cdot\|_F$, and

$$\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \leq \|\mathbf{e}\|^2 + 2\|\mathbf{x}\|\|\mathbf{e}\| \leq (2+\epsilon)\|\mathbf{e}\|\|\mathbf{x}\|, \quad (60)$$

which yields the upper bound as $\|\mathbf{e}\| = \text{dist}(\mathbf{z}, \mathbf{x})$.

2) *Proof of the Lower Bound in Lemma III.3:* Expanding $\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F$ by the definition of the Frobenius inner product, we have

$$\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2 = \|\tilde{\mathbf{Z}}\|_F^2 + \|\tilde{\mathbf{X}}\|_F^2 - 2\text{Re}\langle \tilde{\mathbf{Z}}, \tilde{\mathbf{X}} \rangle_F. \quad (61)$$

Due to rank-1 property, the Frobenius inner product reduces to $2\text{Re}\langle \tilde{\mathbf{Z}}, \tilde{\mathbf{X}} \rangle_F = 2|\langle \mathbf{z}, \mathbf{x} \rangle|^2$, and $\|\tilde{\mathbf{Z}}\|_F^2 = \|\mathbf{z}\|^4$, $\|\tilde{\mathbf{X}}\|_F^2 = \|\mathbf{x}\|^4$. Hence, (61) equals to $\|\mathbf{z}\|^4 + \|\mathbf{x}\|^4 - 2|\langle \mathbf{z}, \mathbf{x} \rangle|^2$, and

$$\begin{aligned} (\|\mathbf{z}\|^4 - |\langle \mathbf{z}, \mathbf{x} \rangle|^2) + (\|\mathbf{x}\|^4 - |\langle \mathbf{z}, \mathbf{x} \rangle|^2) = \\ (\|\mathbf{z}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|)(\|\mathbf{z}\|^2 - |\langle \mathbf{z}, \mathbf{x} \rangle|) + (62) \\ (\|\mathbf{x}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|)(\|\mathbf{x}\|^2 - |\langle \mathbf{z}, \mathbf{x} \rangle|). \end{aligned}$$

Since $\text{dist}^2(\mathbf{z}, \mathbf{x}) = \|\mathbf{z}\|^2 + \|\mathbf{x}\|^2 - 2|\langle \mathbf{z}, \mathbf{x} \rangle| = \|\mathbf{z} - \hat{\mathbf{x}}\|^2 \geq 0$, we can lower bound (61) using (62) as

$$\begin{aligned} \|\tilde{\mathbf{z}} - \tilde{\mathbf{x}}\|_F^2 \geq \min((\|\mathbf{z}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|), (\|\mathbf{x}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|)) \\ \times (\|\mathbf{z}\|^2 + \|\mathbf{x}\|^2 - 2|\langle \mathbf{z}, \mathbf{x} \rangle|). \end{aligned} \quad (63)$$

Knowing that $\text{dist}^2(\mathbf{z}, \mathbf{x}) \leq \epsilon^2\|\mathbf{x}\|^2$ and the result from the reverse triangle inequality on $\|\mathbf{z}\|$, the terms within the minimization are further lower bounded using

$$\begin{aligned} 2|\langle \mathbf{z}, \mathbf{x} \rangle| &\geq \|\mathbf{z}\|^2 + \|\mathbf{x}\|^2 - \epsilon^2\|\mathbf{x}\|^2 \\ |\langle \mathbf{z}, \mathbf{x} \rangle| &\geq (1-\epsilon)\|\mathbf{x}\|^2. \end{aligned} \quad (64)$$

We then get the bound on the scalar multiplying $\text{dist}^2(\mathbf{z}, \mathbf{x})$ as

$$\begin{aligned} \min((\|\mathbf{z}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|), (\|\mathbf{x}\|^2 + |\langle \mathbf{z}, \mathbf{x} \rangle|)) \\ \geq ((1-\epsilon)^2 + (1-\epsilon))\|\mathbf{x}\|^2, \end{aligned} \quad (65)$$

which yields the lower bound of Lemma III.3

$$\|\tilde{\mathbf{z}} - \tilde{\mathbf{x}}\|_F \geq \sqrt{(1-\epsilon)(2-\epsilon)} \text{dist}(\mathbf{z}, \mathbf{x})\|\mathbf{x}\|, \quad (66)$$

and the proof is complete.

D. Proof of Lemma III.4

From Lemma III.1, we have

$$\begin{aligned} \frac{1}{M} \|\mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\|^2 &= \langle \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}) \rangle_F = \\ \langle \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} + (\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)\mathbf{I} + \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F. \end{aligned} \quad (67)$$

From the linearity of the inner product, (67) becomes $\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2 + (\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)(\langle \mathbf{I}, \tilde{\mathbf{Z}} \rangle_F - \langle \mathbf{I}, \tilde{\mathbf{X}} \rangle_F) + \langle \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F$. Since the Frobenius inner product of a matrix with the identity matrix \mathbf{I} is simply the sum of its diagonal terms, and the lifted signals have the auto-correlation of their entries at diagonals, the second term reduces to $(\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)^2 = \|\mathbf{z}\|^4 + \|\mathbf{x}\|^4 - 2\|\mathbf{z}\|^2\|\mathbf{x}\|^2$. From Cauchy-Schwartz, $(\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)^2 \leq \|\mathbf{z}\|^4 + \|\mathbf{x}\|^4 - 2|\langle \mathbf{z}, \mathbf{x} \rangle|^2 = \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2$, by definition. Hence, we obtain

$$\frac{1}{M} \|\mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\|^2 \leq 2\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2 + \langle \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F, \quad (68)$$

and in the other direction, since $(\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)^2$ is the square of a real valued quantity, it is lower bounded by 0, which yields

$$\frac{1}{M} \|\mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\|^2 \geq \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2 + \langle \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F. \quad (69)$$

It remains to upper bound the quantity $|\langle \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F|$. Using the definition in (58), and the linearity of Δ , from Cauchy

Schwartz inequality, we have, for any $\mathbf{z} \in E(\epsilon)$,

$$\begin{aligned} & |\langle \Delta(\tilde{\mathbf{E}} + \mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H), \tilde{\mathbf{Z}} - \tilde{\mathbf{X}} \rangle_F| \\ & \leq \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_* \left(\|\Delta(\tilde{\mathbf{E}})\| + \|\Delta(\mathbf{e}\hat{\mathbf{x}}^H)\| + \|\Delta(\hat{\mathbf{x}}\mathbf{e}^H)\| \right), \\ & \leq \sqrt{2}\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \delta (\|\tilde{\mathbf{E}}\| + \|\mathbf{e}\hat{\mathbf{x}}^H\| + \|\hat{\mathbf{x}}\mathbf{e}^H\|) \quad (70) \\ & = \sqrt{2}\|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \delta (\|\mathbf{e}\|^2 + 2\|\mathbf{x}\|\|\mathbf{e}\|) \\ & \leq \sqrt{2}(2 + \epsilon)\delta \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \text{dist}(\mathbf{z}, \mathbf{x})\|\mathbf{x}\|. \end{aligned}$$

Finally, using the lower bound from Lemma III.3 where $\epsilon < 1$, we obtain

$$\begin{aligned} & \sqrt{2}(2 + \epsilon)\delta \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F \text{dist}(\mathbf{z}, \mathbf{x})\|\mathbf{x}\| \\ & \leq \frac{\sqrt{2}(2 + \epsilon)\delta}{\sqrt{(1 - \epsilon)(2 - \epsilon)}} \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F^2. \quad (71) \end{aligned}$$

Thereby, setting $\hat{\delta} = \frac{\sqrt{2}(2 + \epsilon)}{\sqrt{(1 - \epsilon)(2 - \epsilon)}}\delta$, and combining the bounds (68), (69) and (71), the proof is complete.

E. Proof of Lemma III.5

Recall the definition of the gradient in (13). By Lemma III.4 a RIP-type property is satisfied locally for $\mathbf{z}\mathbf{z}^H - \mathbf{x}\mathbf{x}^H$ if $\mathbf{z} \in E(\epsilon)$. As a result we can express $\nabla f(\mathbf{z}) = \tilde{\mathbf{Z}}\mathbf{z} - \tilde{\mathbf{X}}\mathbf{z} + (\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)\mathbf{z} + \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\mathbf{z} = \|\mathbf{z}\|^2\mathbf{z} - (\hat{\mathbf{x}}^H\mathbf{z})\hat{\mathbf{x}} + (\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2)\mathbf{z} + \Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\mathbf{z}$, with $\hat{\mathbf{x}}$ again denoting the closest solution in P to a given $\mathbf{z} \in \mathbb{C}^N$. Then, we upper bound $\nabla f(\mathbf{z})$ as follows:

$$\begin{aligned} \|\nabla f(\mathbf{z})\| & \leq \|\|\mathbf{z}\|^2\mathbf{z} - (\hat{\mathbf{x}}^H\mathbf{z})\hat{\mathbf{x}}\| + \|\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2\|\mathbf{z}\| \\ & \quad + \|\Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\|\|\mathbf{z}\| \quad (72) \end{aligned}$$

from which, knowing that $\|\mathbf{z}\|^2\mathbf{z} - (\hat{\mathbf{x}}^H\mathbf{z})\hat{\mathbf{x}} = (\|\mathbf{z}\|^2 - \hat{\mathbf{x}}^H\mathbf{z})\mathbf{z} + (\hat{\mathbf{x}}^H\mathbf{z})(\mathbf{z} - \hat{\mathbf{x}})$, where $\mathbf{e} = \mathbf{z} - \hat{\mathbf{x}}$ with $\hat{\mathbf{x}}$ again denoting the closest solution to \mathbf{z} in P , and $\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2 \leq \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F$, we have

$$\begin{aligned} \|\nabla f(\mathbf{z})\| & \leq \|\mathbf{z}\|(\text{dist}(\mathbf{z}, \mathbf{x})(\|\mathbf{z}\| + \|\mathbf{x}\|) \\ & \quad + \|\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}\|_F + \|\Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\|). \quad (73) \end{aligned}$$

Again considering the expression $\tilde{\mathbf{Z}} - \tilde{\mathbf{X}} = \tilde{\mathbf{E}} + \mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H$, we have $\|\Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\| \leq \|\Delta(\tilde{\mathbf{E}})\| + \|\Delta(\mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H)\|$. Since $\mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H$ is at most of rank-2 by definition, let $\hat{\mathbf{x}}\mathbf{e}^H + \hat{\mathbf{x}}\mathbf{e}^H = \sum_{i=1}^2 \lambda_i \mathbf{v}_i \mathbf{v}_i^H$, by which we obtain $\|\Delta(\mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H)\| \leq |\lambda_1| \|\Delta(\mathbf{v}_1 \mathbf{v}_1^H)\| + |\lambda_2| \|\Delta(\mathbf{v}_2 \mathbf{v}_2^H)\|$. Thereby, using Lemma III.1, we have $\|\Delta(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\| \leq \delta(\|\mathbf{e}\|^2 + |\lambda_1| + |\lambda_2|)$. Furthermore, $|\lambda_1| + |\lambda_2|$ precisely corresponds to the nuclear norm of $\mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H$, which is upper bounded as $\|\mathbf{e}\hat{\mathbf{x}}^H + \hat{\mathbf{x}}\mathbf{e}^H\|_* \leq \|\mathbf{e}\hat{\mathbf{x}}^H\|_* + \|\hat{\mathbf{x}}\mathbf{e}^H\|_* \leq 2\|\mathbf{e}\|\|\mathbf{x}\|$. As a result, invoking the upper bound on the error in the lifted domain via Lemma III.3, we obtain

$$\begin{aligned} \|\nabla f(\mathbf{z})\| & \leq \|\mathbf{z}\| ((\|\mathbf{z}\| + \|\mathbf{x}\|) + (1 + \delta)(2 + \epsilon)\|\mathbf{x}\|) \text{dist}(\mathbf{z}, \mathbf{x}) \\ & \leq (1 + \epsilon)(2 + \epsilon)(2 + \delta)\|\mathbf{x}\|^2 \text{dist}(\mathbf{z}, \mathbf{x}). \quad (74) \end{aligned}$$

In (74) we've used the fact that, for $\mathbf{z} \in E(\epsilon)$, $\|\mathbf{z}\| \leq (1 + \epsilon)\|\mathbf{x}\|$. Hence, the proof of local Lipschitz differentiability is complete, with a constant $c(\delta) = (1 + \epsilon)(2 + \epsilon)(2 + \delta)$.

F. Proof of Lemma IV.1

$$\begin{aligned} \text{Using the definition of the spectral matrix as } \mathbf{Y} = \frac{1}{M} \mathcal{A}^H(\mathbf{y}), \\ \text{we have, for any } \mathbf{x} \in \mathbb{C}^N, \text{ and any realization of } \boldsymbol{\eta} \in \mathbb{R}^M \\ \|\mathbf{Y} - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I})\| \leq \|\frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{X}}) - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I})\| \\ + \|\frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta})\|, \quad (75) \end{aligned}$$

where $\mathbf{y} = \mathcal{A}(\tilde{\mathbf{X}}) + \boldsymbol{\eta}$, and the inequality simply follows from the triangle inequality. Under the assumption that our sufficient condition in (15) holds, we get

$$\|\mathbf{Y} - (\tilde{\mathbf{X}} + \|\mathbf{x}\|^2 \mathbf{I})\| \leq \delta \|\mathbf{x}\|^2 + \|\frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta})\|. \quad (76)$$

Then, using the definition of the spectral norm, we have

$$\begin{aligned} \|\frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta})\| & = \max_{\mathbf{z} \in \mathbb{C}^N, \|\mathbf{z}\|=1} |\langle \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}), \mathbf{z}\mathbf{z}^H \rangle_F| \\ & = \max_{\mathbf{z} \in \mathbb{C}^N, \|\mathbf{z}\|=1} |\langle \frac{1}{M} \boldsymbol{\eta}, \mathcal{A}(\mathbf{z}\mathbf{z}^H) \rangle| \quad (77) \\ & \leq \frac{1}{M} \|\boldsymbol{\eta}\| \|\mathcal{A}(\mathbf{z}\mathbf{z}^H)\| \leq \frac{\sqrt{2 + \delta}}{\sqrt{M}} \|\boldsymbol{\eta}\| \end{aligned}$$

where we used the adjoint definition over the Frobenius inner product for \mathcal{A} , followed by Cauchy-Schwartz inequality, and the upper bound from Corollary III.1 $\|\mathcal{A}(\mathbf{z}\mathbf{z}^H)\|^2/M \leq (2 + \delta)\|\mathbf{z}\mathbf{z}^H\|_F^2$, with $\|\mathbf{z}\| = 1$. Reorganizing the last inequality in (77), we obtain

$$\|\frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta})\| \leq \sqrt{2 + \delta} \left(\frac{1}{\sqrt{M}} \|\mathcal{A}(\tilde{\mathbf{X}})\| \right) \frac{\|\boldsymbol{\eta}\|}{\|\mathcal{A}(\tilde{\mathbf{X}})\|}. \quad (78)$$

Invoking the upper bound from Corollary III.1 once again, along with the definition of SNR, and Jensen's inequality, we obtain

$$\mathbb{E}_{\boldsymbol{\eta}} \left[\sqrt{2 + \delta} \left(\frac{1}{\sqrt{M}} \|\mathcal{A}(\tilde{\mathbf{X}})\| \right) \frac{\|\boldsymbol{\eta}\|}{\|\mathcal{A}(\tilde{\mathbf{X}})\|} \right] \leq \frac{2 + \delta}{\sqrt{\text{SNR}}} \|\tilde{\mathbf{X}}\|_F, \quad (79)$$

and plugging the bound in (79) into (76) the proof is complete.

G. Proof of Theorem IV.1

Given the lifted domain definition of the clean gradient term in (13), using the noisy measurements $\mathbf{y} = \mathcal{A}(\mathbf{x}\mathbf{x}^H) + \boldsymbol{\eta}$ with $\boldsymbol{\eta} \in \mathbb{R}^M$, we define

$$\nabla \tilde{f}(\mathbf{z}) = \left(\frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}}) + \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}) \right) \mathbf{z}. \quad (80)$$

Having $\nabla f(\mathbf{z}) = \frac{1}{M} \mathcal{A}^H \mathcal{A}(\tilde{\mathbf{Z}} - \tilde{\mathbf{X}})\mathbf{z}$ as the ideal gradient estimate and setting $\mu_{k+1} = \mu'$, we analyze the following updates:

$$\begin{aligned} \mathbf{z}_{k+1} & = (\mathbf{z}_k - \frac{\mu'}{\|\mathbf{x}\|^2} \nabla f(\mathbf{z}_k)) + \frac{\mu'}{\|\mathbf{x}\|^2} \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}) \mathbf{z}_k \\ & = \hat{\mathbf{z}}_{k+1} + \frac{\mu}{\|\mathbf{x}\|^2} \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}) \mathbf{z}_k, \quad (81) \end{aligned}$$

where $\hat{\mathbf{z}}_{k+1}$ denotes the iterate obtained from the ideal update given the current estimate \mathbf{z}_k . We now approach the proof by induction. Starting from the first iteration $k = 0$, we have

$$\mathbf{z}_1 = \hat{\mathbf{z}}_1 + \frac{\mu}{\|\mathbf{x}\|^2} \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}) \mathbf{z}_0, \quad (82)$$

$$\begin{aligned} \text{dist}(\mathbf{z}_1, \mathbf{x}) &\leq \|\mathbf{z}_1 - \hat{\mathbf{x}}_1\| \leq \|\mathbf{z}_1 - \hat{\mathbf{z}}_1\| + \|\hat{\mathbf{z}}_1 - \hat{\mathbf{x}}_1\| \\ &= \text{dist}(\hat{\mathbf{z}}_1, \mathbf{x}) + \|\mathbf{z}_1 - \hat{\mathbf{z}}_1\|. \end{aligned} \quad (83)$$

Furthermore given (38), the iterates are guaranteed to stay in the ϵ -neighborhood determined by the spectral initialization, which via Theorem III.1 guarantees

$$\text{dist}(\hat{\mathbf{z}}_1, \mathbf{x}) \leq \epsilon \|\mathbf{x}\| \left(1 - \frac{2\mu'}{\alpha'}\right)^{1/2}, \quad (84)$$

under the validity of (36), where $\mu'/\alpha' = \mu/\alpha$. Repeating for the next iteration, under the validity of (36) and (38), we have that

$$\begin{aligned} \text{dist}(\mathbf{z}_2, \mathbf{x}) &\leq \|\mathbf{z}_2 - \hat{\mathbf{z}}_2\| + \text{dist}(\hat{\mathbf{z}}_2, \mathbf{x}) \\ &\leq \left(1 - \frac{2\mu}{\alpha}\right) \epsilon \|\mathbf{x}\| + \sum_{l=1}^2 \left(1 - \frac{2\mu}{\alpha}\right)^{\frac{l-1}{2}} \|\mathbf{z}_l - \hat{\mathbf{z}}_l\|. \end{aligned} \quad (85)$$

and by induction, we obtain

$$\begin{aligned} \text{dist}(\mathbf{z}_k, \mathbf{x}) &\leq \epsilon \left(1 - \frac{2\mu}{\alpha}\right)^{k/2} \|\mathbf{x}\| \\ &\quad + \sum_{l=1}^k \left(1 - \frac{2\mu}{\alpha}\right)^{\frac{l-1}{2}} \|\mathbf{z}_l - \hat{\mathbf{z}}_l\|. \end{aligned} \quad (86)$$

Recalling that $\|\mathbf{z}_l - \hat{\mathbf{z}}_l\| = \|\frac{\mu'}{\|\mathbf{x}\|^2} \frac{1}{M} \mathcal{A}^H(\boldsymbol{\eta}) \mathbf{z}_{l-1}\|$, we can bound the term within the summation as follows:

$$\|\mathbf{z}_l - \hat{\mathbf{z}}_l\| \leq \frac{\mu'}{\|\mathbf{x}\|^2} \|\mathbf{z}_{l-1}\| \frac{(2+\delta)}{\sqrt{\text{SNR}}} \|\mathbf{x}\|^2 := U_l. \quad (87)$$

After summing both sides in (87), we approximately obtain

$$\sum_{l=1}^k U_l \approx \frac{(2+\delta)}{\sqrt{\text{SNR}}} \sum_{l=1}^k \mu' \left(1 - \frac{2\mu'}{\alpha'}\right)^{\frac{l-1}{2}} \|\mathbf{x}\| \quad (88)$$

where we have used that the norms of the iterates $\|\mathbf{z}_{l-1}\| \approx \|\mathbf{x}\|$. Then, applying the geometric sum formula, we obtain

$$\begin{aligned} \sum_{l=1}^k U_l &\approx \frac{(2+\delta)}{\sqrt{\text{SNR}}} \mu' \|\mathbf{x}\| \frac{1 - (1 - \frac{2\mu'}{\alpha'})^{\frac{k-1}{2}}}{1 - (1 - \frac{2\mu'}{\alpha'})^{\frac{1}{2}}} \\ &\leq \frac{(2+\delta)}{\sqrt{\text{SNR}}} \mu' \|\mathbf{x}\| \frac{1 + \sqrt{1 - \frac{2\mu'}{\alpha'}}}{(1 - \sqrt{1 - \frac{2\mu'}{\alpha'}})(1 + \sqrt{1 - \frac{2\mu'}{\alpha'}})} \\ &\leq \frac{(2+\delta)}{\sqrt{\text{SNR}}} \|\mathbf{x}\| \frac{2\mu' \alpha'}{2\mu'} \end{aligned} \quad (89)$$

which completes the proof. \blacksquare

VI. CONCLUSION

This paper analyzes the exact recovery guarantees of the non-convex phase retrieval framework of Wirtinger Flow through a novel perspective in the lifted domain. Our approach quantifies a regime in which the concentration bound of the spectral matrix geometrically implies the validity of the regularity condition. As a result, we identify a sufficient condition under which the convergence to the true solution is guaranteed deterministically via Wirtinger Flow, starting from the estimate obtained from the spectral initialization. Notably, our results address a theoretical gap that exists in phase retrieval literature, in which convergence arguments are predominantly probabilistic in nature. Furthermore, the deterministic convergence arguments developed in this paper rely on a less stringent restricted isometry type property than those of state-of-the art low rank matrix recovery methods. Although numerical simulations on specific problem domains are beyond the scope of this paper, our results culminate into a framework that is highly relevant to applications such as wave-based imaging, in which the underlying scattering phenomenon is typically a deterministic map. In future work, we will study the impact of regularization on our framework, and investigate improvements on our recovery guarantees via alternative initialization schemes to improve its recovery guarantees towards practical measurement models, such as coded diffraction patterns and 2D-Fourier slices. Further directions of our interest also include the study of amplitude-based loss functions to obtain a more inclusive framework beyond our specified loss function over intensity measurements. Namely, the superior sample complexity and convergence rates these methods were shown to attain motivates the study of a RIP-based approach for the analysis of amplitude-based loss functions or for possible analogies between the two formulations.

APPENDIX

A. Uniformity for the Gaussian Model

We begin from the concentration bound which is known to hold for the complex Gaussian sampling model when $M \geq \mathcal{O}(N \log N)$ via [1]:

$$\left\| \frac{1}{M} \mathcal{A}^H \mathcal{A}(\mathbf{x} \mathbf{x}^H) - (\|\mathbf{x}\|^2 \mathbf{I} + \mathbf{x} \mathbf{x}^H) \right\| \leq \delta \|\mathbf{x}\|^2, \quad (90)$$

at a fixed $\delta > 0$ with probability $1 - 5e^{-\gamma N}$ for an appropriately chosen $\gamma(\delta) > 0$, for any fixed $\mathbf{x} \in \mathbb{C}^N$ via unitary invariance on an event² E_0 which holds with probability $1 - 4/N^2$. We need to show (90) holds for all $\mathbf{x} \in \mathbb{C}^N$ uniformly in the event that E_0 holds. Note that due to the Hermitian property of Δ , we established in Lemma III.1 that the condition holds uniformly over \mathbf{x} if \mathbf{v} is fixed and vice-versa. We also know from Corollary III.1 that we can equivalently show that $|\langle \Delta(\mathbf{x} \mathbf{x}^H), \mathbf{x} \mathbf{x}^H \rangle| \leq \delta$ holds uniformly over the unit sphere in \mathbb{C}^N to establish uniformity of (90), since the maxima over \mathbf{x} was proven to be identical for the two conditions.

Now, on the event that E_0 holds, (41) by definition directly implies $|\langle \Delta(\mathbf{x} \mathbf{x}^H), \mathbf{x} \mathbf{x}^H \rangle_F| \leq \delta$, for any $\mathbf{x} \in \mathbb{C}^N$ with $\|\mathbf{x}\| = 1$

²See A.4.2 in [1] for details.

at the same probability when $M \geq \mathcal{O}(N \log N)$. It then remains to extend this result over all $\mathbf{x} \in \mathbb{C}^N$ via an ϵ -net argument. Having $|\langle \Delta(\mathbf{xx}^H), \mathbf{xx}^H \rangle_F| = |M^{-1} \sum_{m=1}^M |\mathbf{a}_m^H \mathbf{x}|^4 - 2| \leq \tilde{\delta}$, with $\mathbb{E}[|\mathbf{a}_m^H \mathbf{x}|^4] = 2$, we follow the methodology of [48] for bounding the a quantity around its expectation uniformly. Assume the bound we have at hand holds with $\tilde{\delta}$ at any \mathbf{z} on the unit sphere, such that γ is properly set to have

$$\Pr \left(\left| \frac{1}{2M} \|\mathcal{A}(\mathbf{zz}^H)\|^2 - 1 \right| \geq \frac{\tilde{\delta}}{2} \right) \leq 5e^{-\gamma N}. \quad (91)$$

We first need to perform the union bound over a properly defined ϵ -net of \mathcal{S}_ϵ on the unit sphere, and then control any perturbations when generalizing the result over to the whole domain. To determine the appropriate ϵ -net, we first have to consider how *tight* the perturbation must be bounded. From (91) we trivially have

$$(1 - \tilde{\delta}/2) \leq \frac{1}{\sqrt{2M}} \|\mathcal{A}(\mathbf{zz}^H)\| \leq (1 + \tilde{\delta}/2), \quad (92)$$

which will hold for all $\mathbf{z} \in \mathcal{S}_\epsilon$ once the union bound is performed. Next we define A as the smallest number such that

$$\frac{1}{\sqrt{2M}} \|\mathcal{A}(\mathbf{xx}^H)\| \leq (1 + A) \quad \text{for all } \|\mathbf{x}\| = 1. \quad (93)$$

Hence, by definition of A we have that

$$\|\mathcal{A}(\mathbf{xx}^H)\| \leq \|\mathcal{A}(\mathbf{zz}^H)\| + \|\mathcal{A}(\mathbf{xx}^H - \mathbf{zz}^H)\| \quad (94)$$

$$\frac{1}{\sqrt{2M}} \|\mathcal{A}(\mathbf{xx}^H)\| \leq 1 + \tilde{\delta}/2 + (1 + A) \|\mathbf{xx}^H - \mathbf{zz}^H\|_* \quad (95)$$

where we have used the fact that $\|\mathcal{A}(\mathbf{xx}^H - \mathbf{zz}^H)\| \leq |\lambda_1| \|\mathcal{A}(\mathbf{v}_1 \mathbf{v}_1^H)\| + |\lambda_2| \|\mathcal{A}(\mathbf{v}_2 \mathbf{v}_2^H)\|$, with $\mathbf{xx}^H - \mathbf{zz}^H = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^H + \lambda_2 \mathbf{v}_2 \mathbf{v}_2^H$. Since by definition A is the smallest value for which (93) holds, we have that

$$A \leq \tilde{\delta}/2 + (1 + A)(|\lambda_1| + |\lambda_2|). \quad (96)$$

Setting $\sigma = |\lambda_1| + |\lambda_2|$, we have

$$A \leq \frac{\tilde{\delta}/2 + \sigma}{1 - \sigma}, \quad (97)$$

where we wish to have an ϵ such that $A \leq \tilde{\delta}$, which indicates

$$\frac{\tilde{\delta}/2 + \sigma}{1 - \eta} \leq \tilde{\delta} \rightarrow \sigma \leq \frac{\tilde{\delta}}{2(1 + \tilde{\delta})} \quad (98)$$

which is satisfied for $\sigma \leq \tilde{\delta}/4$ by definition since $\tilde{\delta} < 1$. Now we know from Lemma III.3 that for any \mathbf{x} with $\|\mathbf{x}\| = 1$, and any \mathbf{z} where $\|\mathbf{x} - \mathbf{z}\| \leq \epsilon \|\mathbf{x}\|$ we have that

$$\sigma = \|\mathbf{xx}^H - \mathbf{zz}^H\|_* \leq (2 + \epsilon) \|\mathbf{x} - \mathbf{z}\| \|\mathbf{x}\| \leq (2 + \epsilon) \epsilon. \quad (99)$$

As a result, we can determine an ϵ -net on the unit sphere in the signal domain, such that a perturbation is controlled in the lifted domain. It essentially requires that $\tilde{\delta}/4 = 2\epsilon + \epsilon^2$. Since $\tilde{\delta}/4 + 1 = (1 + \epsilon)^2$, with an ϵ -net \mathcal{S}_ϵ with

$$\epsilon = \sqrt{\frac{\tilde{\delta}}{4} + 1} - 1 \approx \frac{\tilde{\delta}}{8} \quad (100)$$

we obtain the desired result of $A \leq \tilde{\delta}$. The lower inequality then follows as $1 - \tilde{\delta}/2 - (1 + \tilde{\delta})\tilde{\delta}/4 \geq 1 - \tilde{\delta}$ as noted in [48]. Thereby, generalizing a union bound over the set \mathcal{S}_ϵ with cardinality k , for all $\mathbf{z} \in \mathcal{S}_\epsilon$ we have that

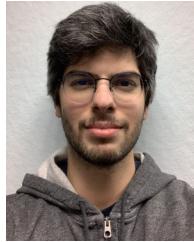
$$\Pr \left(\left| \frac{1}{2M} \|\mathcal{A}(\mathbf{zz}^H)\|^2 - 1 \right| \leq \frac{\tilde{\delta}}{2} \right) \leq 1 - 5(24/\tilde{\delta})^k e^{-\gamma N}, \quad (101)$$

where we used the covering number of \mathcal{S}_ϵ from [48] via reference [49]. Hence the concentration bound in (90) holds uniformly over all \mathbf{x} with $\delta = 2\tilde{\delta}$ with probability $1 - 5(24/\tilde{\delta})^k e^{-\gamma N} - 4N^{-2}$ when $M = \mathcal{O}(N \log N)$, and the proof is complete.

REFERENCES

- [1] E. J. Candes, 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.
- [2] E. J. Candes and T. Strohmer, "PhaseLift: Exact and stable recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, Aug. 2013.
- [3] T. Bendory, Y. C. Eldar, and N. Boumal, "Non-convex phase retrieval from stft measurements," *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 467–484, Jan. 2018.
- [4] H. Zhang and Y. Liang, "Reshaped Wirtinger flow for solving quadratic systems of equations," *Adv. Neural Inf. Process. Syst.*, pp. 2622–2630, 2016.
- [5] 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, 1972.
- [6] J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," *Opt. Lett.*, vol. 3, no. 1, pp. 27–29, 1978.
- [7] H. H. Bauschke, P. L. Combettes, and D. R. Luke, "Hybrid projection–reflection method for phase retrieval," *J. Opt. Soc. Am. A*, vol. 20, no. 6, pp. 1025–1034, 2003.
- [8] A. Chai, M. Moscoso, and G. Papanicolaou, "Array imaging using intensity-only measurements," *IOP Inverse Problems J.*, vol. 27, no. 1, pp. 1–16, Jan. 2011.
- [9] E. J. Candes, Y. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. Imag. Sci.*, vol. 6, no. 1, pp. 199–225, 2013.
- [10] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Math. Program.*, vol. 149, no. 1, pp. 47–81, 2015.
- [11] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2796–2804.
- [12] T. Goldstein and C. Studer, "PhaseMax: Convex phase retrieval via basis pursuit," *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 2675–2689, Apr. 2018.
- [13] P. Hand and V. Voroninski, "An elementary proof of convex phase retrieval in the natural parameter space via the linear program phasemax," *Commun. Math. Sci.*, vol. 16, no. 7, pp. 2047–2051, 2018.
- [14] S. Bahmani and R. Justin, "A flexible convex relaxation for phase retrieval," *Electron. J. Statist.*, vol. 11, no. 2, pp. 5254–5281, 2017.
- [15] Y. Chen and E. J. Candes, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," *Commun. Pure Appl. Math.*, vol. 70, pp. 0822–0883, 2017.
- [16] J. Chen, L. Wang, X. Zhang, and Q. Gu, "Robust Wirtinger flow for phase retrieval with arbitrary corruption," 2017.
- [17] H. Zhang, Y. Chi, and Y. Liang, "Median-truncated nonconvex approach for phase retrieval with outliers," *IEEE Trans. Inf. Theory*, vol. 64, no. 11, pp. 7287–7310, Nov. 2018.
- [18] G. Wang, G. Giannakis, Y. Saad, and J. Chen, "Solving most systems of random quadratic equations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 1867–1877.
- [19] G. Wang, G. B. Giannakis, Y. Saad, and J. Chen, "Phase retrieval via reweighted amplitude flow," *IEEE Trans. Signal Process.*, vol. 66, no. 11, pp. 2818–2833, Jun. 2018.
- [20] G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving systems of random quadratic equations via truncated amplitude flow," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 773–794, Feb. 2018.

- [21] Y. M. Lu and G. Li, "Phase transitions of spectral initialization for high-dimensional nonconvex estimation," *Inf. Inference: A J. IMA*, 2019.
- [22] M. Mondelli and A. Montanari, "Fundamental limits of weak recovery with applications to phase retrieval," *Found. Comput. Math.*, vol. 19, no. 3, pp. 703–773, 2019.
- [23] W. Luo, W. Alghamdi, and Y. M. Lu, "Optimal spectral initialization for signal recovery with applications to phase retrieval," *IEEE Trans. Signal Process.*, vol. 67, no. 9, pp. 2347–2356, May 2019.
- [24] R. Ghods, A. S. Lan, T. Goldstein, and C. Studer, "Linear spectral estimators and an application to phase retrieval," in *Proc. 35th Int. Conf. Mach. Learn.*, vol. 80, 2018, pp. 1734–1743.
- [25] B. Gao and Z. Xu, "Phaseless recovery using the Gauss–Newton method," *IEEE Trans. Signal Process.*, vol. 65, no. 22, pp. 5885–5896, Nov. 2017.
- [26] S. Bhojanapalli, A. Kyriolidis, and S. Sanghavi, "Dropping convexity for faster semi-definite optimization," in *Proc. Conf. Learn. Theory*, 2016, pp. 530–582.
- [27] B. Yonel and B. Yazici, "A generalization of Wirtinger flow for exact interferometric inversion," *SIAM J. Imag. Sci.*, vol. 12, no. 4, pp. 2119–2164, 2019.
- [28] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," *Found. Comput. Math.*, vol. 18, no. 5, pp. 1131–1198, 2018.
- [29] C. Ma, K. Wang, Y. Chi, and Y. Chen, "Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution," *Found. Comput. Math.*, vol. 20, no. 3, pp. 451–632, 2020.
- [30] S. Sanghavi, R. Ward, and C. D. White, "The local convexity of solving systems of quadratic equations," *Results Math.*, vol. 71, no. 3-4, pp. 569–608, 2017.
- [31] X. Li, S. Ling, T. Strohmer, and K. Wei, "Rapid, robust, and reliable blind deconvolution via nonconvex optimization," *Appl. Comput. Harmon. Anal.*, vol. 47, pp. 893–934, Nov. 2019.
- [32] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht, "Low-rank solutions of linear matrix equations via procrustes flow," in *Proc. 35th Int. Conf. Mach. Learn.*, vol. 48, pp. 964–973, 2016.
- [33] Z. Zhu, Q. Li, G. Tang, and M. B. Wakin, "Global optimality in low-rank matrix optimization," *IEEE Trans. Signal Process.*, vol. 66, no. 13, pp. 3614–3628, Jul. 2018.
- [34] S. Bhojanapalli, B. Neyshabur, and N. Srebro, "Global optimality of local search for low rank matrix recovery," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3873–3881.
- [35] R. Zhang, C. Josz, S. Sojoudi, and J. Lavaei, "How much restricted isometry is needed in nonconvex matrix recovery?" in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 5586–5597.
- [36] Q. Zheng and J. Lafferty, "A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 109–117.
- [37] L. Wang, X. Zhang, and Q. Gu, "A unified computational and statistical framework for nonconvex low-rank matrix estimation," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, vol. 54, pp. 981–990, 2017.
- [38] Y. Chi, Y. M. Lu, and Y. Chen, "Nonconvex optimization meets low-rank matrix factorization: An overview," *IEEE Trans. Sig. Process.*, vol. 67, no. 20, pp. 5239–5269, Aug. 2019.
- [39] 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.
- [40] J.-F. Cai, E. J. Candes, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [41] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.
- [42] B. Recht, W. Xu, and B. Hassibi, "Necessary and sufficient conditions for success of the nuclear norm heuristic for rank minimization," in *Proc. 47th IEEE Conf. Decis. Control*, 2008, pp. 3065–3070.
- [43] B. Recht, W. Xu, and B. Hassibi, "Null space conditions and thresholds for rank minimization," *Math. Program.*, vol. 127, no. 1, pp. 175–202, 2011.
- [44] S. Oymak, K. Mohan, M. Fazel, and B. Hassibi, "A simplified approach to recovery conditions for low rank matrices," in *Proc. IEEE Int. Symp. Inf. Theory Proc.*, 2011, pp. 2318–2322.
- [45] T. T. Cai, "Sharp RIP bound for sparse signal and low-rank matrix recovery," *Appl. Comput. Harmon. Anal.*, vol. 35, pp. 74–93, 2013.
- [46] L. Demanet and P. Hand, "Stable optimizationless recovery from phaseless linear measurements," *J. Fourier Anal. Appl.*, vol. 20, no. 1, pp. 199–221, 2014.
- [47] B. Yonel, I.-Y. Son, and B. Yazici, "Exact multistatic interferometric imaging via generalized Wirtinger flow," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 711–726, Jan. 2020.
- [48] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [49] R. A. DeVore and G. G. Lorentz, *Constructive Approximation*, vol. 303. Berlin, Germany: Springer, 1993.



Bariscan Yonel (Student Member, IEEE) received the B.Sc. degree in electrical engineering from Koc University in Istanbul, Turkey, in 2015. He subsequently joined Rensselaer Polytechnic Institute (RPI) in Troy, NY, where he is currently a Ph.D. candidate in electrical engineering. His research and thesis work focus on theoretical guarantees and practical limitations for solving quadratic equations in high dimensional inference and wave-based imaging problems, using low rank matrix recovery theory and computationally efficient non-convex algorithms. Further research interests include applications of machine learning, compressed sensing and optimization methods for signal processing and computational imaging.



Birsen Yazici (Senior Member, IEEE) received the B.S. degree in electrical engineering and mathematics, in 1988, from Bogazici University, Istanbul Turkey, and M.S. and Ph.D. degrees in mathematics and electrical engineering both from Purdue University, West Lafayette, IN, in 1990 and 1994, respectively. From September 1994 until 2000, she was a Research Engineer at General Electric Company Global Research Center, Schenectady NY. During her tenure in industry, she worked on radar, transportation, industrial and medical imaging systems. From 2001 to June 2003, she was an Assistant Professor at Drexel University, Electrical and Computer Engineering Department. In 2003, she joined Rensselaer Polytechnic Institute where she is currently a Full Professor in the Department of Electrical, Computer and Systems Engineering and in the Department of Biomedical Engineering. Prof. Yazici's research interests span the areas of statistical signal processing, inverse problems in imaging, image reconstruction, biomedical optics, radar and X-ray imaging. She served as an associate editor for the *IEEE TRANSACTIONS ON IMAGE PROCESSING* from 2008 to 2012, *IEEE TRANSACTIONS ON GEOSCIENCES AND REMOTE SENSING* from 2014 to 2018, and for *SIAM Journal on Imaging Science* from 2010 to 2014. She currently serves as an Associate Editor for *IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING* and as a Distinguished Lecturer of the IEEE Aerospace and Electronics Systems Society. She is the recipient of the Rensselaer Polytechnic Institute 2007 and 2013 School of Engineering Research Excellence awards. She holds 11 US patents.