PHASE RETRIEVAL: GLOBAL CONVERGENCE OF GRADIENT DESCENT WITH OPTIMAL SAMPLE COMPLEXITY

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper addresses the phase retrieval problem, which aims to recover a signal vector $\boldsymbol{x}^\natural$ from m measurements $y_i = |\langle \boldsymbol{a}_i, \boldsymbol{x}^\natural \rangle|^2$, $i = 1, \ldots, m$. A standard approach is to solve a nonconvex least squares problem using gradient descent with random initialization, which is known to work efficiently given a sufficient number of measurements. However, whether $O(n)$ measurements suffice for gradient descent to recover the ground truth efficiently has remained an open question. Prior work has established that $O(n\,\mathrm{poly}(\log n))$ measurements are sufficient. In this paper, we resolve this open problem by proving that $m = O(n)$ Gaussian random measurements are sufficient to guarantee, with high probability, that the objective function has a benign global landscape. This sample complexity is optimal because at least $\Omega(n)$ measurements are required for exact recovery. The landscape result allows us to further show that gradient descent with a constant step size converges to the ground truth from almost any initial point.

1 INTRODUCTION

We study the problem of phase retrieval, which aims to recover a complex valued vector $\boldsymbol{x}^\natural \in \mathbb{C}^n$ from its intensity measurements

$$y_i = |\langle \boldsymbol{a}_i, \boldsymbol{x}^\natural \rangle|^2, \quad i = 1, \ldots, m, \tag{1}$$

where $\boldsymbol{a}_i \in \mathbb{C}^n$, $i = 1, \ldots, m$, are known complex vectors and m is the number of measurements. This problem has attracted high interest due to its broad applications in X-ray crystallography (Elser et al., 2018), microscopy (Miao et al., 2008), astronomy (Fienup & Dainty, 1987) and optical imaging (Shechtman et al., 2015).

The phase retrieval problem is NP-hard if only very few measurements, e.g. $m = n + 1$ (Fickus et al., 2014), are given. However, a wide range of algorithms can recover $\boldsymbol{x}^\natural$ up to a global phase shift provided enough measurements. Early methods with provable performance guarantees usually formulate it into a convex constrained optimization problem, such as a semidefinite programming problem (Candes et al., 2013; 2015a; Waldspurger et al., 2015) or basis pursuit problem (Goldstein & Studer, 2018). These methods are usually computationally challenging in high-dimensional cases. To address this issue, more recent works take nonconvex approaches, such as alternating minimization (Wen et al., 2012; Netrapalli et al., 2013; Waldspurger, 2018; Zhang, 2020); gradient descent type algorithms, including Wirtinger flow (Candes et al., 2015b; Chen & Candes, 2015; Ma et al., 2020), truncated amplitude flow (Wang et al., 2017), vanilla gradient descent (Chen et al., 2019), Riemannian gradient descent (Cai & Wei, 2024); and Newton type algorithms (Gao & Xu, 2017; Ma et al., 2018).

Convex methods mentioned above can achieve optimal sample complexity $m = O(n)$, but require an initialization close enough to the ground truth $\boldsymbol{x}^\natural$. For nonconvex algorithms, $O(n)$ sample complexity can be achieved with a careful initialization (Chen & Candes, 2015; Wang et al., 2017; Waldspurger, 2018; Cai & Wei, 2024). However, such initialization can be computationally inefficient when the dimension n is large. In practice, random initialization is more plausible. In the random initialization regime, it is known that $O(n \log^{13} n)$ samples are sufficient to guarantee a nearly linear convergence rate for vanilla gradient descent (Chen et al., 2019). To obtain a lower sample complexity in this regime, a benign global landscape result was analyzed for both the intensity measurement (1) (Sun et al., 2018b; Cai et al., 2023) and the amplitude measurement $y_i = |\langle \boldsymbol{a}_i, \boldsymbol{x}^\natural \rangle|$ (Cai et al., 2022b) for ℓ_2-loss function. Both the intensity and the amplitude model have no spurious local minima and saddle points are strict, provided $O(n\log n)$ and $O(n)$ samples respectively. Nevertheless, no global convergence guarantee of gradient descent algorithm can be easily deduced because the Lipschitz gradient assumption in the classical convergence results (Lee et al., 2016) does not hold, and the bounded iterates assumptions are hard to verify.

In this paper, we aim at showing the global convergence of vanilla gradient descent algorithm with arbitrary initialization for general phase retrieval problem with $O(n)$ intensity measurements. We propose a new tensor based criterion to show that $O(n)$ Gaussian samples of intensity measurements (1) are sufficient to guarantee a benign global landscape for phase retrieval problem. Furthermore, we show that such objective function has bounded gradient trajectories, which give us the certificate to apply the general global convergence result in (Josz, 2023). By combining the global benign landscape result and global convergence result, we conclude that given $O(n)$ Gaussian samples, with high probability, vanilla gradient descent initialized almost everywhere converge to the ground truth of phase retrieval problem.

1.1 Contributions

The main contributions of this paper can be summarized as follows.

1. We propose a new simple deterministic criterion and use it to verify the global benign landscape of intensity-based phase retrieval problem. In particular, with this simple criterion, we prove at once that the objective function admits restricted strong convexity near the ground truth, has Hessian with a negative eigenvalue near the origin, and has nonzero gradient elsewhere.
2. We utilize a new concentration inequality for random tensor to prove that $O(n)$ Gaussian samples are sufficient to satisfy our proposed criterion. This new technique circumvents the main challenge in literature that the objective and its gradient are heavy-tailed and hence reduce the logarithmic factor of the sample complexity.
3. We establish the boundedness of any gradient trajectories for intensity-based phase retrieval problem by revealing that its gradient trajectory only moves inside a linear subspace. This useful fact allows us to invoke a general global convergence result of gradient descent. Together with the global landscape results, our main result follows immediately.

Notation: Denote $\mathbb{R}$ and $\mathbb{C}$ as the real and complex fields respectively. Denote $\|\boldsymbol{x}\|$ as the ℓ_2-norm of vector $\boldsymbol{x}\in\mathbb{C}^n$. For $\boldsymbol{a},\boldsymbol{b}\in\mathbb{C}^n$, denote $\langle\boldsymbol{a},\boldsymbol{b}\rangle=\sum_{i=1}^n a_i\overline{b_i}$, where $\overline{b_i}$ is the complex conjugate of b_i. We denote $\boldsymbol{a}\otimes\boldsymbol{b}$ as the tensor product of $\boldsymbol{a}$ and $\boldsymbol{b}$ and $\boldsymbol{a}^{\otimes r}$ as the tensor product of r vectors $\boldsymbol{a}$. Denote $\boldsymbol{I}_n$ as the $n\times n$ identity matrix and $\boldsymbol{O}_n$ as the $n\times n$ zero matrix. We call $\mathbf{T}=\boldsymbol{u}_1\otimes\cdots\otimes\boldsymbol{u}_p$ a tensor of rank 1 and order p. Denote $\|\mathbf{T}\|=\|\boldsymbol{u}_1\|\cdots\|\boldsymbol{u}_p\|$ and $\|\mathbf{T}\|_{\mathrm{op}}=\sup_{\|\boldsymbol{x}_1\otimes\cdots\otimes\boldsymbol{x}_p\|=1}\langle\mathbf{T},\boldsymbol{x}_1\otimes\cdots\otimes\boldsymbol{x}_p\rangle$ as its operator norm. Denote $\boldsymbol{H}_f(\boldsymbol{x})$ the Hessian of a function f at a point $\boldsymbol{x}$. We write $f(n)=O(g(n))$ if $f(n)\le C_1g(n)$, $f(n)=\Omega(g(n))$ if $f(n)\ge C_2g(n)$, and $f(n)=\Theta(g(n))$ if $C_2g(n)\le f(n)\le C_1g(n)$ for some constants $C_1,C_2>0$.

2 Problem formulation and related work

We first present the problem formulation for phase retrieval and the algorithm to solve it, and then review related work and provide a comparative overview of state-of-the-art convergence results and our new result.

2.1 Problem formulation

We denote $\boldsymbol{x}^\natural\in\mathbb{C}^n$ as the ground truth vector we want to recover. Intensity-based phase retrieval problem aims at minimizing the empirical ℓ_2-loss of m intensity measurements $y_i=|\langle\boldsymbol{a}_i,\boldsymbol{x}^\natural\rangle|^2$,

$$\min_{\boldsymbol{x}\in\mathbb{C}^n}\quad f(\boldsymbol{x}):=\sum_{i=1}^m\left(|\langle\boldsymbol{a}_i,\boldsymbol{x}\rangle|^2-y_i\right)^2, \tag{2}$$

Since all vectors of the form $e^{i\theta}\boldsymbol{x}^\natural$ are global minimizers of f, we can only expect to recover $\boldsymbol{x}^\natural$ up to a phase shift by solving (2). For the purpose of comparison, we also recall the amplitude-based formulation mentioned in the introduction,

$$\min_{\boldsymbol{x}\in\mathbb{C}^n} \quad \sum_{i=1}^{m} \left(|\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle| - \sqrt{y_i}\right)^2.$$

It is easy to see that the intensity-based formulation has a smooth objective function while the amplitude-based formulation has a nonsmooth one.

In order to apply vanilla gradient descent to solve (2) as it is applied in the real case (Chen et al., 2019), we introduce the following mapping for any $\boldsymbol{v} \in \mathbb{C}^n$:

$$\boldsymbol{v} \mapsto \boldsymbol{v}^+ := \begin{bmatrix} \mathrm{Re}(\boldsymbol{v}) \\ \mathrm{Im}(\boldsymbol{v}) \end{bmatrix}, \quad \boldsymbol{v} \mapsto \boldsymbol{v}^- := \boldsymbol{M}\boldsymbol{v}^+ = \begin{bmatrix} -\mathrm{Im}(\boldsymbol{v}) \\ \mathrm{Re}(\boldsymbol{v}) \end{bmatrix} \text{ where } \boldsymbol{M} := \begin{bmatrix} \boldsymbol{O}_n & -\boldsymbol{I}_n \\ \boldsymbol{I}_n & \boldsymbol{O}_n \end{bmatrix}. \tag{3}$$

By using (3), we can define $\boldsymbol{a}_i^+$, $\boldsymbol{a}_i^-$, $\boldsymbol{x}^+$, $\boldsymbol{x}^-$, $\boldsymbol{x}^{\natural+}$ and $\boldsymbol{x}^{\natural-}$. Then, with a little abuse of notation over f, we can obtain a equivalent form of f in terms of $\boldsymbol{x}^+$ as

$$f(\boldsymbol{x}) = \sum_{i=1}^{m} \left(|\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle|^2 - y_i\right)^2 = \sum_{i=1}^{m} \left(\langle \boldsymbol{A}_i\boldsymbol{x}^+, \boldsymbol{x}^+\rangle - y_i\right)^2 := f(\boldsymbol{x}^+), \tag{4}$$

where $\boldsymbol{A}_i = \boldsymbol{a}_i^+(\boldsymbol{a}_i^+)^{\mathrm{T}} + \boldsymbol{a}_i^-(\boldsymbol{a}_i^-)^{\mathrm{T}}$. Now we can treat f as a function from $\mathbb{R}^{2n}$ to $\mathbb{R}$ w.r.t. $\boldsymbol{x}^+$.

The gradient descent algorithm to minimize f is given by

$$\boldsymbol{x}_{k+1}^+ = \boldsymbol{x}_k^+ - \alpha_k \nabla_{\boldsymbol{x}^+} f(\boldsymbol{x}_k^+), \quad \forall\, k \in \mathbb{N}, \tag{5}$$

where $\boldsymbol{x}_0^+ \in \mathbb{R}^{2n}$ is any given initial point and $\alpha_k > 0$ is any step size. The goal of this paper is to study the global convergence and sample complexity of gradient descent algorithm (5) for f defined in (4). As we did in (5), all gradient and Hessian in the rest of this paper will be taken w.r.t. $\boldsymbol{x}^+$.

2.2 Related work

Global benign landscape results usually refer to properties that all local minima of a function are global minima and all saddle points are strict, i.e., Hessian matrix having a negative eigenvalue at the saddle point. For phase retrieval problem, existing results obtain $O(n)$ sample complexity when the objective function is quadratic-like far from the origin, including amplitude model (piecewise quadratic) (Cai et al., 2022b), quotient intensity model (quartic divided by quadratic) (Cai et al., 2022a), and perturbed amplitude model (truncated quartic) (Cai et al., 2021). For purely quartic objective function, like intensity model, people usually require $O(n\,\mathrm{poly}(\log n))$ samples to obtain the global benign landscape result (Sun et al., 2018b; Cai et al., 2023). Table 1 summarizes the above results.

Work	Objective function	Sample complexity	Probability
Sun et al. (2018a)	quartic	$O(n\log^3 n)$	$1-\Theta(1/m)$
Cai et al. (2021)	truncated quartic	$O(n)$	$1-O(1/m^2)$
Cai et al. (2022a)	quartic over quadratic	$O(n)$	$1-\exp(-\Theta(m))$
Cai et al. (2022b)	piecewise quadratic	$O(n)$	$1-\exp(-\Theta(m))$
Cai et al. (2023)	quartic	$O(n\log n)$	$1-\Theta(1/m)$
Our work	**quartic**	$\boldsymbol{O(n)}$	$\boldsymbol{1-\exp(-\Theta(m))}$

Table 1: Global benign landscape results of different objective functions and sample complexity in literature. Here "quartic" and "quadratic" means 4th order and 2nd order polynomial respectively.

Convergence results for gradient descent type algorithms to solve phase retrieval problem either requires special initialization that are usually already close a ground truth (Candes et al., 2015b; Chen & Candes, 2015; Wang et al., 2017; Cai & Wei, 2024) or suboptimal sample complexity $O(n\log^{13} n)$ (Chen et al., 2019), as summarized in Table 2. In the former case, the iterates start nearby a ground truth, where a desirable local landscape occurs and is enough to guarantee a contraction property of the distance of iterates towards ground truth. In the latter case, a careful analysis of some nonlinear approximate dynamics and complicated concentration results are required to show that after a short period of time, the iterates will enter the neighborhood with contraction property. Contrary to the existing proof idea, we develop a new and concise proof of global convergence results for phase retrieval problem by connecting global benign landscape results and continuous time gradient dynamics results.

Work	Algorithm	Initialization	Sample complexity
Candes et al. (2015b)	WF	spectral	$O(n \log n)$
Chen & Candes (2015)	TWF	spectral	$O(n)$
Wang et al. (2017)	TAF	orthogonality-promoting	$O(n)$
Chen et al. (2019)	GD	$\mathcal{N}(\mathbf{0}, n^{-1}\boldsymbol{I}_n)$	$O(n \log^{13} n)$
Cai & Wei (2024)	TRGrad	close to ground truth	$O(n)$
Our work	**GD**	**a.e. on** $\mathbb{R}^{2n}$	$\boldsymbol{O(n)}$

Table 2: Global convergence results of gradient descent type algorithms with different initialization and sample complexity in literature. Here "GD" stands for vanilla gradient descent, "WF" for Wirtinger flow, "TWF" for truncated Wirtinger flow, "TAF" for truncated amplitude flow, "TRGrad" for truncated Riemannian gradient descent.

3 Main results

We introduce tensors $\mathsf{T}, \mathsf{S} \in \left(\mathbb{R}^{2n}\right)^{\otimes 4}$ that will be useful to formulate our main results. Define T as

$$\mathsf{T} := \frac{1}{c} \sum_i (\boldsymbol{a}_i^+)^{\otimes 4},$$

where $c = m\sigma^4$ and $\sigma^2 = \mathrm{Var}((\boldsymbol{a}_i^+)_1)$. We also define S as

$$\mathsf{S}_{i_1,i_2,i_3,i_4} := \mathbf{1}_{i_1=i_2,i_3=i_4} + \mathbf{1}_{i_1=i_3,i_2=i_4} + \mathbf{1}_{i_1=i_4,i_2=i_3},$$

where $(\mathbf{1}_{i_1=i_2,i_3=i_4})_{i_1,i_2,i_3,i_4} = 1$ if $i_1 = i_2$ and $i_3 = i_4$ and 0 else (and similarly for other tensors).

The main results of this paper are given in Theorem 1 and Theorem 2. Theorem 1 gives a tensor based criterion for global benign landscape of f.

Theorem 1. *Assume $\|\mathsf{T} - \mathsf{S}\|_{\mathrm{op}} \leq \delta_0$ for some constant $\delta_0 > 0$ small enough. The only local minima of f defined in (4) are global minima $\boldsymbol{x}^\natural e^{i\theta}$ and all saddle points of f are strict*[1].

The proof of Theorem 1 is available in Section 3.2. The main idea is to show that when $\|\mathsf{T} - \mathsf{S}\|_{\mathrm{op}}$ is small, the objective function f is restricted strongly convex near the ground truth, has indefinite Hessian near the origin, and has nonzero gradient elsewhere.

With Theorem 1 at hand, we can derive our global convergence result in Theorem 2.

Theorem 2. *Assume $\|\mathsf{T} - \mathsf{S}\|_{\mathrm{op}} \leq \delta_0$ for some constant $\delta_0 > 0$ small enough. For almost every initial point $\boldsymbol{x}_0^+ \in \mathbb{R}^{2n}$, there exists $\bar{\alpha} > 0$ such that for any step sizes $(\alpha_k)_{k\in\mathbb{N}}$ satisfying $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\alpha_k \in (0, \bar{\alpha}]$, $\forall k \in \mathbb{N}$, the gradient descent algorithm (5) converges to a global minimizer of f defined in (4).*

The proof of Theorem 2 is available in Section 3.3. The main idea is to show that f has bounded gradient trajectories, and then use the general convergence result (Josz, 2023) together with Theorem 1 to conclude the global convergence of gradient descent with random initialization for phase retrieval problem.

Finally, by using a concentration result, we verify the criterion in Theorem 1 and Theorem 2 when $\{\boldsymbol{a}_i\}_{i=1}^m$ are i.i.d. random vectors distributed as $\mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{I}_{2n})$, which yields the following application

Corollary 1. *Let $\{\boldsymbol{a}_i^+\}_{i=1}^m$ be i.i.d. random vectors distributed as $\mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{I}_{2n})$. There exist $K, \beta > 0$ such that if $m \geq Kn$, then with at least $1 - e^{-\beta m}$ probability, the conclusions in Theorem 1 and Theorem 2 hold.*

[1] A strict saddle point is a critical point at which the Hessian has a strictly negative eigenvalue.

The rest of this section will be organized as follows. In Section 3.1, we introduce an equivalent formulation of f as an inner product of some tensors. In Section 3.2, we provide necessary ingredients for proving Theorem 1. In Section 3.3, we show that f has bounded gradient trajectories and deduce Theorem 2. In Section 3.4, we apply a tensor concentration result and Theorems 1 and 2 to obtain Corollary 1.

3.1 Equivalent formulation of the problem

In this section, we provide an equivalent tensor based formulation for f. We write Hermitian product as a function of real vectors $\boldsymbol{x}^+$ and $\boldsymbol{x}^-$:

$$|\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle|^2 = \langle \boldsymbol{a}_i^+, \boldsymbol{x}^+\rangle^2 + \langle \boldsymbol{a}_i^+, \boldsymbol{x}^-\rangle^2. \tag{6}$$

To make use of tensors, for $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^{2n}$, we develop the following product of scalar products as:

$$\begin{aligned}\langle \boldsymbol{a}_i^+, \boldsymbol{u}\rangle^2 \langle \boldsymbol{a}_i^+, \boldsymbol{v}\rangle^2 &= \sum_{i_1,i_2,i_3,i_4} (\boldsymbol{a}_i^+)_{i_1} (\boldsymbol{a}_i^+)_{i_2} (\boldsymbol{a}_i^+)_{i_3} (\boldsymbol{a}_i^+)_{i_4} u_{i_1} u_{i_2} v_{i_3} v_{i_4} \\ &= \langle (\boldsymbol{a}_i^+)^{\otimes 4}, \boldsymbol{u}^{\otimes 2} \otimes \boldsymbol{v}^{\otimes 2}\rangle. \end{aligned}\tag{7}$$

Combining (6) and (7), we obtain an alternative expression of f as in Proposition 1.

Proposition 1. *f can be written as*

$$f(\boldsymbol{x}^+) = \left\langle \sum_i (\boldsymbol{a}_i^+)^{\otimes 4}, \mathbf{U}(\boldsymbol{x}^+) \right\rangle = c\langle \mathbf{T}, \mathbf{U}(\boldsymbol{x}^+)\rangle,$$

where $\mathbf{U}(\boldsymbol{x}^+)$ is a tensor defined by

$$\mathbf{U}(\boldsymbol{x}^+) := \sum_k \varepsilon_k \boldsymbol{x}_{1,k}^{\otimes 2} \otimes \boldsymbol{x}_{2,k}^{\otimes 2}$$

with $\boldsymbol{x}_{1,k}, \boldsymbol{x}_{2,k} \in \{\boldsymbol{x}^+, \boldsymbol{x}^-, \boldsymbol{x}^{\natural +}, \boldsymbol{x}^{\natural -}\}$ and $\varepsilon_k = \pm 1$.

The proof of Proposition 1 is in Appendix A.1.1. We regard $\mathbf{U}$ as a function of $\boldsymbol{x}^+$ even if it involves $\boldsymbol{x}^-$, because $\boldsymbol{x}^- = \boldsymbol{M}\boldsymbol{x}^+$ is also a function of $\boldsymbol{x}^+$. With the above alternative form of f, the assumption $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} \leq \delta_0$ allows us to the landscape of f by studying the landscape of its approximation $c\langle \mathbf{S}, \mathbf{U}(\boldsymbol{x}^+)\rangle$. A detailed analysis of this approximation is given in Section 3.2.

3.2 Landscape results

In this subsection, we prove that all critical points of f are either global minimizers or strict saddles. Given the fact that $\mathbf{T}$ is close to $\mathbf{S}$, we can expect that $f = c\langle \mathbf{T}, \mathbf{U}\rangle$ is close to $c\langle \mathbf{S}, \mathbf{U}\rangle$ in some sense. Therefore, our landscape analysis of f is inspired by the landscape information of this approximation. Proposition 2 gives an equivalent formula for this approximation function.

Proposition 2. *We have*

$$g(\boldsymbol{x}^+) := c\langle \mathbf{S}, \mathbf{U}(\boldsymbol{x}^+)\rangle = 8c\left(\|\boldsymbol{x}\|^4 + \|\boldsymbol{x}^\natural\|^4 - |\langle \boldsymbol{x}, \boldsymbol{x}^\natural\rangle|^2 - \|\boldsymbol{x}\|^2\|\boldsymbol{x}^\natural\|^2\right).$$

The proof of Proposition 2 is in Appendix A.2.1. Local minimizers of g are global minimizers and other critical points of g are strict saddles (see the analysis in Appendix A.2.7). Motivated by this, we prove that similar geometrical properties also hold for f. We define the following 3 regions:

$$\mathcal{R}^1_{\delta_0} := \left\{ \boldsymbol{x} \;\middle|\; (\|\boldsymbol{x}\| + \|\boldsymbol{x}^\natural\|)^3 \leq \frac{\|\nabla g(\boldsymbol{x}^+)\|}{C_1 \delta_0 c} \right\}, \tag{8}$$

$$\mathcal{R}^2_{\delta_0} := \left\{ \boldsymbol{x} \;\middle|\; 8\frac{|\langle \boldsymbol{x}, \boldsymbol{x}^\natural\rangle|^2}{\|\boldsymbol{x}^\natural\|^4} + \left(4 + \frac{1}{4}\delta_0 C_2\right)\frac{\|\boldsymbol{x}\|^2}{\|\boldsymbol{x}^\natural\|^2} \leq \left(4 - \frac{1}{4}\delta_0 C_2\right) \right\}, \tag{9}$$

$$\mathcal{R}^3_{\delta_0} := \left\{ \boldsymbol{x} \;\middle|\; d(\boldsymbol{x}, \mathcal{G}) \leq \frac{16 - 5C_2\delta_0}{192 + 4C_2\delta_0} \right\}, \tag{10}$$

where c, δ_0, C_1, C_2 are constants and $\mathcal{G}$ is the set of global minimizers of f.

Our goal is to show that

- $\mathcal{R}^1_{\delta_0}$ contains no critical points of f,
- $\mathcal{R}^2_{\delta_0}$ only contains strict saddles of f,
- $\mathcal{R}^3_{\delta_0}$ only contains global minimizers.

If we write $\boldsymbol{x}^+ = \boldsymbol{y} + \mu \boldsymbol{x}^{\natural+} + \nu \boldsymbol{x}^{\natural-}$ with $\boldsymbol{y} \in \mathbb{R}^{2n}$ orthogonal to $\boldsymbol{x}^{\natural+}, \boldsymbol{x}^{\natural-}$, we can represent regions $\mathcal{R}^2_{\delta_0}, \mathcal{R}^3_{\delta_0}$ on a 3-d plot with coordinates $(\|\boldsymbol{y}\|, \langle \boldsymbol{x}^+, \boldsymbol{x}^{\natural+}\rangle, \langle \boldsymbol{x}^+, \boldsymbol{x}^{\natural-}\rangle)$ as shown in Figure 1.

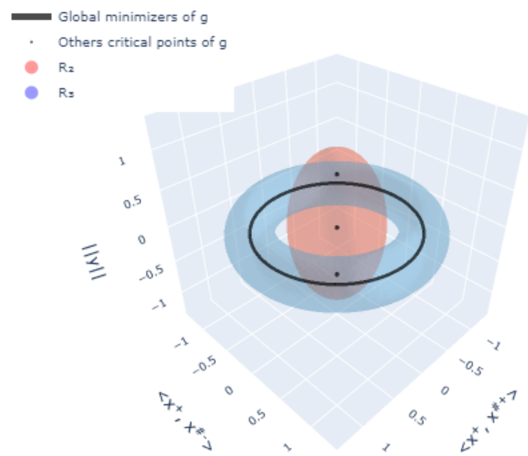

Figure 1: Overview of the regions considered in the proof

First we prove that $\mathcal{R}^1_{\delta_0}$ has no critical points. We bound the entries of the gradient of $\mathbf{U}$ and use the concentration result $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} \leq \delta_0$ to lower bound $\langle \nabla f(\boldsymbol{x}^+), \nabla g(\boldsymbol{x}^+)\rangle$ by $\|\nabla g(\boldsymbol{x}^+)\|^2 - C_1 \|\nabla g(\boldsymbol{x}^+)\| (\|\boldsymbol{x}\| + \|\boldsymbol{x}^\natural\|)^3$ and deduce Proposition 3.

Proposition 3. *Assuming $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} < \delta_0$, we have for some absolute constant C_1,*

$$\|\nabla g(\boldsymbol{x}^+)\| \geq C_1 \delta_0 c (\|\boldsymbol{x}\| + \|\boldsymbol{x}^\natural\|)^3 \implies \nabla f(\boldsymbol{x}^+) \neq 0.$$

The proof of Proposition 3 is in Appendix A.2.2. This means that as long as the gradient of g is large enough at $\boldsymbol{x}^+$, then $\boldsymbol{x}^+$ is not a critical point and of f and thus $\mathcal{R}^1_{\delta_0}$ does not contain any critical point of f.

To show that the critical points in $\mathcal{R}^2_{\delta_0}$ and $\mathcal{R}^3_{\delta_0}$ are either strict saddles or global minima respectively, we need to bound the distance between the Hessian of f and g. We control the Hessian of the entries of $\mathbf{U}$ and use the concentration $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} \leq \delta_0$ to bound $\boldsymbol{u}^T \boldsymbol{H}_f \boldsymbol{u} - \boldsymbol{u}^T \boldsymbol{H}_g \boldsymbol{u}$ in Proposition 4.

Proposition 4. *Assume $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} < \delta_0$. For all vectors u and for some absolute constant C_2,*

$$|\boldsymbol{u}^T \boldsymbol{H}_f(\boldsymbol{x}^+) \boldsymbol{u} - \boldsymbol{u}^T \boldsymbol{H}_g(\boldsymbol{x}^+) \boldsymbol{u}| < C_2 \delta_0 c \|\boldsymbol{u}\|^2 (\|\boldsymbol{x}\| + \|\boldsymbol{x}^\natural\|)^2. \tag{11}$$

The proof of Proposition 4 is in Appendix A.2.3. To derive geometrical properties for f, we need an explicit formula for the Hessian of g, which can be easily deduced from Proposition 2. The result is given in Proposition 5.

Proposition 5. *For all $\boldsymbol{x}^+$ we have*

$$\boldsymbol{H}_g(\boldsymbol{x}^+) = 8c(8\boldsymbol{x}^+(\boldsymbol{x}^+)^{\mathrm{T}} + 4\|\boldsymbol{x}^+\|^2 \boldsymbol{I}_n - 2\boldsymbol{x}^{\natural+}(\boldsymbol{x}^{\natural+})^{\mathrm{T}} - 2\boldsymbol{x}^{\natural-}(\boldsymbol{x}^{\natural-})^{\mathrm{T}} - 2\|\boldsymbol{x}^{\natural+}\|^2 \boldsymbol{I}_n). \tag{12}$$

The proof of Proposition 5 is in Appendix A.2.4. Combining Proposition 4 and Proposition 5, if $\boldsymbol{H}_g(\boldsymbol{x}^+)$ has a sufficiently negative eigenvalue along some direction, then $\boldsymbol{H}_f(\boldsymbol{x}^+)$ will also have a negative eigenvlaue along the same direction. In region $\mathcal{R}^2_{\delta_0}$, we find that $\boldsymbol{H}_g(\boldsymbol{x}^+)$ indeed has a sufficiently negative eigenvalue along $\boldsymbol{x}^{\natural+}$, which yields Proposition 6.

Proposition 6. *Assuming $\|\mathbf{T}-\mathbf{S}\|_{\mathrm{op}} < \delta_0$, we have*

$$8\frac{|\langle \boldsymbol{x}, \boldsymbol{x}^\natural\rangle|^2}{\|\boldsymbol{x}^\natural\|^4} + \left(4+\frac{1}{4}\delta_0 C_2\right)\frac{\|\boldsymbol{x}\|^2}{\|\boldsymbol{x}^\natural\|^2} \le \left(4-\frac{1}{4}\delta_0 C_2\right) \implies (\boldsymbol{x}^{\natural+})^T \boldsymbol{H}_f(\boldsymbol{x}^+)\boldsymbol{x}^{\natural+} < 0.$$

The proof of Proposition 6 is in Appendix A.2.5. It shows that all possible critical points in the region $\mathcal{R}^2_{\delta_0}$ are strict saddles.

Finally we study the critical points in $\mathcal{R}^3_{\delta_0}$. We prove that all possible critical points in this region are global minimizers. To do so, we use the concept of restricted strong convexity similar to Sun et al. (2018b). The set of global minimizers is the circle $\mathcal{G} = \{\boldsymbol{x}^\natural e^{i\theta} : \theta \in [0, 2\pi)\}$, so we cannot expect f to be strongly convex in this region. However, we can expect to have strong convexity in the directions orthogonal to $T_{\boldsymbol{x}^+}(\mathcal{G})$ (the line tangent to the circle of global minimizers at $\boldsymbol{x}^+$ where $\boldsymbol{x} \in \mathcal{G}$). Again, we combine Proposition 4 and Proposition 5 to prove that we indeed have restricted convexity and deduce that there are no critical points near $\mathcal{G}$ in Proposition 7.

Proposition 7. *Assume $\|\mathbf{T}-\mathbf{S}\|_{\mathrm{op}} < \delta_0$. If*

$$d(\boldsymbol{x}, \mathcal{G}) \le \frac{16 - 5C_2\delta_0}{192 + 4C_2\delta_0}$$

then $\boldsymbol{x}$ is a critical point if and only if $\boldsymbol{x}$ is a global minimizer.

The proof of Proposition 7 is in Appendix A.2.6. To conclude the landscape result, the last step is prove that $\mathcal{R}^1_{\delta_0} \cup \mathcal{R}^2_{\delta_0} \cup \mathcal{R}^3_{\delta_0} = \mathbb{R}^{2n}$ for δ_0 small enough.

The key idea is that $\mathcal{R}^1_{\delta_0}$ is increasing in the sense of set inclusion as δ_0 gets smaller. The only points in the complementary of $\mathcal{R}^1_{\delta_0}$ are those close to critical points. As $\mathcal{R}^2_{\delta_0}$ and $\mathcal{R}^3_{\delta_0}$ contain an open neighbourhood of the critical points, the complementary of $\mathcal{R}^1_{\delta_0}$ is strictly included in $\mathcal{R}^2_{\delta_0} \cup \mathcal{R}^2_{\delta_0}$ for small enough δ_0. Therefore, we can finally conclude the landscape result for f in Proposition 8.

Proposition 8. *For δ_0 small enough, we have*

$$\mathcal{R}^1_{\delta_0} \cup \mathcal{R}^2_{\delta_0} \cup \mathcal{R}^3_{\delta_0} = \mathbb{R}^{2n}.$$

The proof of Proposition 8 is in Appendix A.2.7. Combining Propositions 3 and 6 to 8, the result in Theorem 1 follows immediately.

3.3 Convergence of gradient descent

In this subsection, we prove Theorem 2 by using the fact that f has bounded gradient trajectories and our landscape results in Theorem 1.

We say that f has bounded gradient trajectories if for every $\boldsymbol{x}^+_0 \in \mathbb{R}^{2n}$, the solution $\boldsymbol{x}^+(\cdot)$ to the following initial value problem

$$(\boldsymbol{x}^+)'(t) = -\nabla f(\boldsymbol{x}^+(t)), \quad \boldsymbol{x}^+(0) = \boldsymbol{x}^+_0 \tag{13}$$

satisfies that $\|\boldsymbol{x}^+(t)\| \le c_{\boldsymbol{x}^+_0}$ for all $t \ge 0$, where $c_{\boldsymbol{x}^+_0}$ is a constant dependent on $\boldsymbol{x}^+_0$.

The following Proposition 9 verifies that f has bounded gradient trajectories for general positive semidefinite matrices $\{\boldsymbol{A}_i\}_{i=1}^m$, not necessarily of the form $\boldsymbol{a}_i^+(\boldsymbol{a}_i^+)^{\mathrm{T}} + \boldsymbol{a}_i^-(\boldsymbol{a}_i^-)^{\mathrm{T}}$.

Proposition 9. *Let $\boldsymbol{A}_i \in \mathbb{R}^{2n\times 2n}$ be symmetric positive semidefinite and $y_i \in \mathbb{R}$ for all $i = 1, \ldots, m$. Then (4) has bounded subgradient trajectories.*

The proof of Proposition 9 is in Appendix A.3.1. Therefore, we are certified to apply results in Josz (2023, Corollary 1) and Theorem 1 to conclude Theorem 2.

3.4 Concentration results

The proof of Corollary 1 is mainly based on a concentration result of the tensor $\mathbf{T}$. We prove that $\mathbb{E}(\mathbf{T}) = \mathbf{S}$ and then use the concentration result from (Even & Massoulié, 2021) to control the deviation of $\mathbf{T}$ from its mean.

Noticing that $(\boldsymbol{a}_i^+)^{\otimes 4}_{i_1,i_2,i_3,i_4} = (\boldsymbol{a}_i^+)_{i_1}(\boldsymbol{a}_i^+)_{i_2}(\boldsymbol{a}_i^+)_{i_3}(\boldsymbol{a}_i^+)_{i_4}$, we can easily deduce Proposition 10.

Proposition 10.

$$\mathbb{E}\left(\sum_i (\boldsymbol{a}_i^+)^{\otimes 4}\right) = m\sigma^4 \mathbf{S}.$$

The proof of Proposition 10 is in Appendix A.4.1. Using the result in (Even & Massoulié, 2021), if $\mathbf{T}_1, \dots \mathbf{T}_m$ are 4-th order Kronecker product of normally distributed random vectors, then

$$\mathbb{P}\left(\left\|\frac{1}{m}\sum_i \mathbf{T}_i - \frac{1}{m}\mathbb{E}\left(\sum_i \mathbf{T}_i\right)\right\|_{\mathrm{op}} \geq C\sqrt{\frac{n + \ln(n) + \beta m}{m}}\right) \leq e^{-\beta m}.$$

Applying this result to the normalized tensors $\mathbf{T}_i = \left(\frac{\boldsymbol{a}_i}{\sigma}\right)^{\otimes 4}$, we verify the concentration assumption of Theorems 1 and 2 in Proposition 11.

Proposition 11. *Let $\{\boldsymbol{a}_i^+\}_{i=1}^m$ be i.i.d. random vectors distributed as $\mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_{2n})$. There exists some absolute constants $K, \beta > 0$, such that for all $m \geq Kn$, with probability at least $1 - e^{-\beta m}$,*

$$\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} < \delta_0.$$

The proof of Proposition 11 can be found in Appendix A.4.2. Corollary 1 then follows naturally.

4 Experiments

In this section, we first show the concentration of tensor $\mathbf{T}$ and that $m = O(n)$ is enough to ensure $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}}$ small numerically. Then, we show the convergence of the loss function for gradient descent with a fixed n versus different m and that all trajectories do converge at a linear rate when m is large enough even for large initializations.

We generate a sample of $\ell = 5$ sets of m vectors $\{\boldsymbol{a}_i^+\}_{i=1}^m$ i.i.d and distributed as $\mathcal{N}(\mathbf{0}, \boldsymbol{I}_{2n})$ for various m, n and computed an approximation of $\sup_{\|\boldsymbol{u}_1\|=\|\boldsymbol{u}_2\|=\|\boldsymbol{u}_3\|=\|\boldsymbol{u}_4\|=1} \langle \mathbf{T}, \boldsymbol{u}_1 \otimes \boldsymbol{u}_2 \otimes \boldsymbol{u}_3 \otimes \boldsymbol{u}_4 \rangle$ to estimate $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}}$ and averaged the result over the ℓ samples. The result is displayed in Figure 2.

We can observe the concentration of $\mathbf{T}$ around $\mathbf{S}$ when m is large enough for fixed n. Moreover, the boundary of the region $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} \leq \delta_0$ for some δ_0 is approximately linear: for $\delta_0 = 0.4$, the values of m such that $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}} \leq \delta_0$ are approximately $m \geq 2000n$. This validates that $m = O(n)$ samples are sufficient to ensure $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}}$ is small enough. However, in practice, we don't need as much samples as $m = 2000n$. As we will see in the next experiment, in practice the number of samples needed to recover the ground truth is about $m \approx 10n$ for almost all initial points.

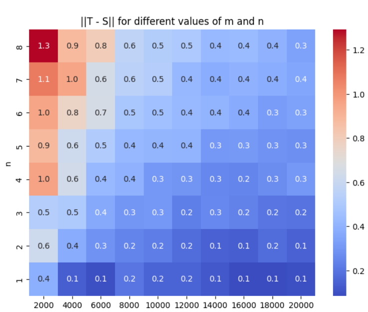

Figure 2: Values of $\|\mathbf{T} - \mathbf{S}\|_{\mathrm{op}}$ for different values of m and n

In Figure 3, we plot the values $\tilde{f} = \frac{f}{m}$ at each iteration of vanilla gradient descent for $n = 4$ and $m \in \{10, 20, 30, 40\}$. We normalize the function f by $\frac{1}{m}$ to have comparable magnitude for different values of m. For each value of m, we chose $\{\boldsymbol{a}_i^+\}_{i=1}^m, \boldsymbol{x}^{\natural+}$ i.i.d and distributed as

$\mathcal{N}(\mathbf{0}, \boldsymbol{I}_{2n})$. To study the influence of potentially large initializations, we select $\ell = 20$ initial points following a uniform distribution on $[-10, 10]^n$, a learning rate $\eta = 5 \cdot 10^{-5}$ and $N = 5000$ iterations.

As we may expect, some trajectories do not converge to a global minimum when m is relatively small: for $m = 10$ and $m = 20$, there are 5 and 1 trajectories not converging to a global minimum respectively. However, when m is large enough (for $m = 30$ and $m = 40$ in our experiment), all initializations converge to a global minimum of $\tilde{f}$ with a linear rate. Note that the rate of convergence also seems to be better when the number of samples increases.

This suggests that the effective K with which we have convergence of trajectories with high probability for $m \geq Kn$ should be around $K \approx 10$ as a rough estimate. Note that from Fickus et al. (2014) theoretically there are at least $m = 4n - 2$ measurements needed to recover the ground truth vector, and using $m \approx 10n$ measurements is close to optimal.

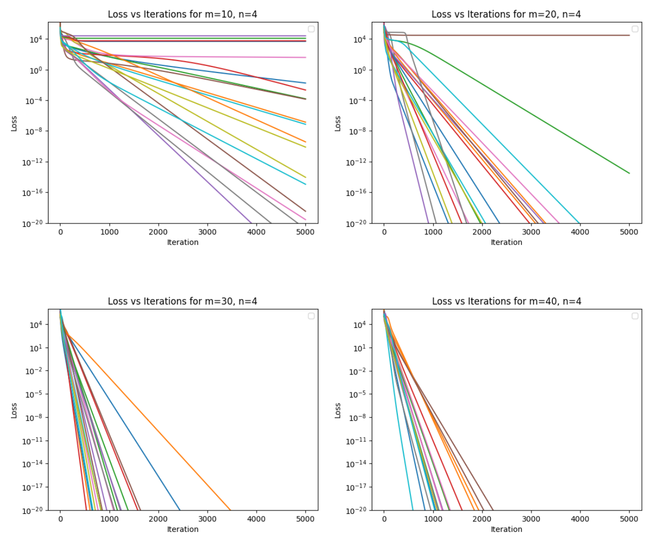

Figure 3: Loss of 20 trajectories for $n = 4$ and $m \in [10, 20, 30, 40]$

5 Conclusion and discussions

In this paper, we provided a tensor based criterion that guarantees global convergence of vanilla gradient descent with random initialization for phase retrieval problem. We first showed that the objective function has a benign global landscape. Then we proved that given the number of measurements $m \geq Kn$, the criterion is satisfied with high probability. We also showed that the objective function has bounded gradient trajectories, which allows us to utilize the proposed landscape results and a general convergence result in literature to conclude our main convergence result.

Finally, we discuss the limitation of our paper, which also leads to a potential future direction. From our proof, only a local linear convergence rate can be obtained once the iterates enter the region $\mathcal{R}^3_{\delta_0}$, but there is no information on how long the gradient descent algorithm will take to enter $\mathcal{R}^3_{\delta_0}$ starting from almost every point in $\mathcal{R}^1_{\delta_0}$ or $\mathcal{R}^2_{\delta_0}$. A more detailed analysis is needed for how long the iterates will stay in $\mathcal{R}^1_{\delta_0}$ and $\mathcal{R}^2_{\delta_0}$. For $\mathcal{R}^1_{\delta_0}$, the analysis is relatively simple. Since the gradient norm is lower bounded, the objective value will drop by a constant for every iteration. The key challenge is to ensure the iterates will not revisit $\mathcal{R}^2_{\delta_0}$ once leaving it. With the global landscape result provided in this paper, we expect to obtain a nearly linear iteration complexity result with only $O(n)$ samples in a following paper.

REFERENCES

Jian-Feng Cai, Meng Huang, Dong Li, and Yang Wang. The global landscape of phase retrieval ii: Quotient intensity models. *Ann. Appl. Math*, 38(1):62–114, 2022a.

Jian-Feng Cai, Meng Huang, Dong Li, and Yang Wang. Solving phase retrieval with random initial guess is nearly as good as by spectral initialization. *Applied and Computational Harmonic Analysis*, 58:60–84, 2022b.

Jian-Feng Cai, Meng Huang, Dong Li, and Yang Wang. Nearly optimal bounds for the global geometric landscape of phase retrieval. *Inverse Problems*, 39(7):075011, 2023.

Jianfeng Cai and Ke Wei. Solving systems of phaseless equations via riemannian optimization with optimal sampling complexity. *Journal of Computational Mathematics*, 42(3):755–783, 2024.

Jianfeng Cai, Meng Huang, Dong Li, and Yang Wang. The global landscape of phase retrieval i: Perturbed amplitude models. *Annals of Applied Mathematics*, 37(4):437, 2021.

Emmanuel J Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

Emmanuel J Candes, Yonina C Eldar, Thomas Strohmer, and Vladislav Voroninski. Phase retrieval via matrix completion. *SIAM review*, 57(2):225–251, 2015a.

Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015b.

Yuxin Chen and Emmanuel Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Advances in Neural Information Processing Systems*, 28, 2015.

Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176:5–37, 2019.

Veit Elser, Ti-Yen Lan, and Tamir Bendory. Benchmark problems for phase retrieval. *SIAM Journal on Imaging Sciences*, 11(4):2429–2455, 2018.

Mathieu Even and Laurent Massoulié. Concentration of non-isotropic random tensors with applications to learning and empirical risk minimization. In *Conference on Learning Theory*, pp. 1847–1886. PMLR, 2021.

Matthew Fickus, Dustin G Mixon, Aaron A Nelson, and Yang Wang. Phase retrieval from very few measurements. *Linear Algebra and its Applications*, 449:475–499, 2014.

C Fienup and J Dainty. Phase retrieval and image reconstruction for astronomy. *Image recovery: theory and application*, 231:275, 1987.

Bing Gao and Zhiqiang Xu. Phaseless recovery using the gauss–newton method. *IEEE Transactions on Signal Processing*, 65(22):5885–5896, 2017.

Tom Goldstein and Christoph Studer. Phasemax: Convex phase retrieval via basis pursuit. *IEEE Transactions on Information Theory*, 64(4):2675–2689, 2018.

Cédric Josz. Global convergence of the gradient method for functions definable in o-minimal structures. *Mathematical Programming*, 202(1):355–383, 2023.

Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on learning theory*, pp. 1246–1257. PMLR, 2016.

Chao Ma, Xin Liu, and Zaiwen Wen. Globally convergent levenberg-marquardt method for phase retrieval. *IEEE Transactions on Information Theory*, 65(4):2343–2359, 2018.

Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution. *Foundations of Computational Mathematics*, 20:451–632, 2020.

Jianwei Miao, Tetsuya Ishikawa, Qun Shen, and Thomas Earnest. Extending x-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes. *Annu. Rev. Phys. Chem.*, 59(1):387–410, 2008.

Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi. Phase retrieval using alternating minimization. *Advances in Neural Information Processing Systems*, 26, 2013.

Yoav Shechtman, Yonina C Eldar, Oren Cohen, Henry Nicholas Chapman, Jianwei Miao, and Mordechai Segev. Phase retrieval with application to optical imaging: a contemporary overview. *IEEE signal processing magazine*, 32(3):87–109, 2015.

Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18(5):1131–1198, October 2018a. ISSN 1615-3375. doi: 10.1007/s10208-017-9365-9.

Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18:1131–1198, 2018b.

Irene Waldspurger. Phase retrieval with random gaussian sensing vectors by alternating projections. *IEEE Transactions on Information Theory*, 64(5):3301–3312, 2018.

Irene Waldspurger, Alexandre d'Aspremont, and Stéphane Mallat. Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming*, 149:47–81, 2015.

Gang Wang, Georgios B Giannakis, and Yonina C Eldar. Solving systems of random quadratic equations via truncated amplitude flow. *IEEE Transactions on Information Theory*, 64(2):773–794, 2017.

Zaiwen Wen, Chao Yang, Xin Liu, and Stefano Marchesini. Alternating direction methods for classical and ptychographic phase retrieval. *Inverse Problems*, 28(11):115010, 2012.

Teng Zhang. Phase retrieval by alternating minimization with random initialization. *IEEE Transactions on Information Theory*, 66(7):4563–4573, 2020.