

A Geometric Analysis of Phase Retrieval

Ju Sun, Qing Qu, and John Wright

{js4038, qq2105, jw2966}@columbia.edu

Department of Electrical Engineering, Columbia University, New York, USA

February 22, 2016 Revised: February 29, 2016

Abstract

Can we recover a complex signal from its Fourier magnitudes? More generally, given a set of m measurements, $y_k = |\mathbf{a}_k^* \mathbf{x}|$ for $k = 1, \dots, m$, is it possible to recover $\mathbf{x} \in \mathbb{C}^n$ (i.e., length- n complex vector)? This *generalized phase retrieval* (GPR) problem is a fundamental task in various disciplines, and has been the subject of much recent investigation. Natural nonconvex heuristics often work remarkably well for GPR in practice, but lack clear theoretical explanations. In this paper, we take a step towards bridging this gap. We prove that when the measurement vectors \mathbf{a}_k 's are generic (i.i.d. complex Gaussian) and the number of measurements is large enough ($m \geq Cn \log^3 n$), with high probability, a natural least-squares formulation for GPR has the following benign geometric structure: (1) there are no spurious local minimizers, and all global minimizers are equal to the target signal \mathbf{x} , up to a global phase; and (2) the objective function has a negative curvature around each saddle point. This structure allows a number of iterative optimization methods to efficiently find a global minimizer, without special initialization. To corroborate the claim, we describe and analyze a second-order trust-region algorithm.

Keywords. Phase retrieval, Nonconvex optimization, Function landscape, Second-order geometry, Ridable saddles, Trust-region method, Inverse problems, Mathematical imaging

1 Introduction

1.1 Generalized Phase Retrieval and a Nonconvex Formulation

This paper concerns the problem of recovering an n -dimensional complex vector \mathbf{x} from the magnitude $y_k = |\mathbf{a}_k^* \mathbf{x}|$ of its projections onto a collection of known complex vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{C}^n$. Obviously, one can only hope to recover \mathbf{x} up to a global phase, as $\mathbf{x}e^{i\phi}$ for all $\phi \in [0, 2\pi)$ gives exactly the same set of measurements. The *generalized phase retrieval* problem asks whether it is possible to recover \mathbf{x} , up to this fundamental ambiguity:

Generalized Phase Retrieval Problem: Is it possible to *efficiently* recover an unknown \mathbf{x} from $y_k = |\mathbf{a}_k^* \mathbf{x}|$ ($k = 1, \dots, m$), up to a global phase factor $e^{i\phi}$?

This problem has attracted substantial recent interest, due to its connections to fields such as crystallography, optical imaging and astronomy. In these areas, one often has access only to the Fourier magnitudes of a complex signal \mathbf{x} , i.e., $|\mathcal{F}(\mathbf{x})|$ [Mil90, Rob93, Wal63, DF87]. The phase information is hard or infeasible to record due to physical constraints. The problem of recovering the signal \mathbf{x} from its Fourier magnitudes $|\mathcal{F}(\mathbf{x})|$ is naturally termed (Fourier) phase retrieval (PR).

It is easy to see PR as a special version of GPR, with the \mathbf{a}_k 's the Fourier basis vectors. GPR also sees applications in electron microscopy [MIJ⁺02], diffraction and array imaging [BDP⁺07, CMP11], acoustics [BCE06, Bal10], quantum mechanics [Cor06, Rei65] and quantum information [HMW13]. We refer the reader to survey papers [SEC⁺15, JEH15] for accounts of recent developments in the theory, algorithms, and applications of GPR.

For GPR, heuristic methods based on nonconvex optimization often work surprisingly well in practice (e.g., [Fie82, GS72], and many more cited in [SEC⁺15, JEH15]). However, investigation into provable recovery methods, particularly based on nonconvex optimization, has started only relatively recently [NJS13, CESV13, CSV13, CL14, CLS15a, WdM15, VX14, ABFM14, CLS15b, CC15, WWS15]. The surprising effectiveness of nonconvex heuristics on GPR remains largely mysterious. In this paper, we take a step towards bridging this gap.

We focus on a natural least-squares formulation¹ – discussed systematically in [SEC⁺15, JEH15] and studied theoretically in [CLS15b, WWS15],

$$\text{minimize}_{\mathbf{z} \in \mathbb{C}^n} f(\mathbf{z}) \doteq \frac{1}{2m} \sum_{k=1}^m \left(y_k^2 - |\mathbf{a}_k^* \mathbf{z}|^2 \right)^2. \quad (1.1)$$

We assume the \mathbf{a}_k 's are independent identically distributed (i.i.d.) complex Gaussian:

$$\mathbf{a}_k = \frac{1}{\sqrt{2}} (X_k + iY_k), \text{ with } X_k, Y_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n) \text{ independent.} \quad (1.2)$$

$f(\mathbf{z})$ is a fourth-order polynomial in \mathbf{z} , and is nonconvex. A-priori, there is little reason to believe that simple iterative methods can solve this problem without special initialization. Typical local convergence (i.e., convergence to a local minimizer) guarantees in optimization require an initialization near the target minimizer [Ber99]. Moreover, existing results on provable recovery using (1.1) and related formulations rely on careful initialization in the vicinity of the ground truth [NJS13, CLS15b, CC15, WWS15].

1.2 A Curious Experiment

We apply gradient descent to $f(\mathbf{z})$, starting from a *random initialization* $\mathbf{z}^{(0)}$:

$$\mathbf{z}^{(r+1)} = \mathbf{z}^{(r)} - \mu \nabla_{\mathbf{z}} f(\mathbf{z}^{(r)}),$$

where the step size μ is fixed for simplicity². The result is quite striking (Figure 1): for a fixed problem instance (fixed set of random measurements and fixed target \mathbf{x}), gradient descent seems to always return a *global minimizer* (i.e., the target \mathbf{x} up to a global phase shift), across many independent random initializations! This contrasts with the typical “mental picture” of nonconvex objectives as possessing many spurious local minimizers.

¹Another least-squares formulation, $\text{minimize}_{\mathbf{z}} \frac{1}{2m} \sum_{k=1}^m (y_k - |\mathbf{a}_k^* \mathbf{z}|)^2$, was studied in the seminal works [Fie82, GS72]. An obvious advantage of the $f(\mathbf{z})$ studied here is that it is differentiable.

²Mathematically, $f(\mathbf{z})$ is not complex differentiable; here the gradient is defined based on the Wirtinger derivatives [KD09]; see also [CLS15b]. This notion of gradient is a natural choice when optimizing real-valued functions of complex variables.

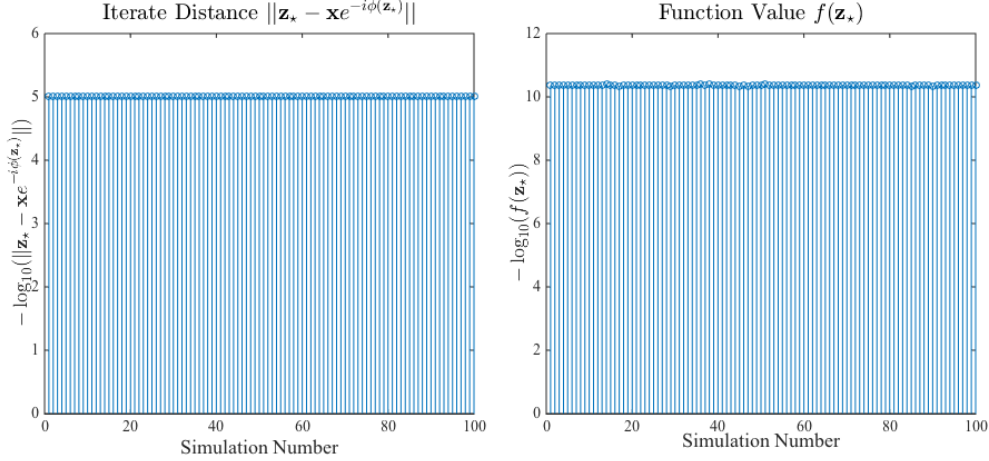


Figure 1: Gradient descent with random initialization seems to always return a global solution for (1.1)! Here $n = 100$, $m = 5n \log n$, step size $\mu = 0.05$, and stopping criterion is $\|\nabla_z f(z)\| \leq 10^{-5}$. We fix the set of random measurements and the ground-truth signal x . The experiments are repeated for 100 times with independent random initializations. z_* denotes the final iterate at convergence. (Left) Final distance to the target; (Right) Final function value (0 if globally optimized). Both vertical axes are on $-\log_{10}(\cdot)$ scale.

1.3 A Geometric Analysis

The numerical surprise described above is not completely isolated. Simple heuristic methods have been observed to work surprisingly well for practical PR [Fie82, GS72, SEC+15, JEH15]. In this paper, we take a step towards explaining this phenomenon. We show that *although the function (1.1) is nonconvex, when m is reasonably large, it actually has benign global geometry which allows it to be globally optimized by efficient iterative methods, regardless of the initialization.*

This geometric structure is evident for GPR in \mathbb{R}^2 . Figure 2 plots the function landscape of $f(z)$ for this case with large m (i.e., $\mathbb{E}_a[f(z)]$ approximately). Notice that (i) the only local minimizers are exactly $\pm x$ – they are also global minimizers;³ (ii) there are saddle points (and a local maximizer), but around them there is negative curvature in the $\pm x$ direction. Intuitively, any algorithm that can successfully escape this kind of saddle point (and local maximizer) can in fact find a global minimizer, i.e., recover the target signal x .

We prove that an analogous geometric structure exists, with high probability (w.h.p.)⁴, for GPR in \mathbb{C}^n , when m is reasonably large (Theorem 2.2). In particular, we show that when $m \geq Cn \log^3 n$, w.h.p., (i) the only local and also global minimizers to (1.1) are the target $xe^{i\phi}$ for $\phi \in [0, 2\pi)$; (ii) at any point in \mathbb{C}^n , either the gradient is large, or the curvature is negative in a certain direction, or it is near a minimizer. Moreover, in the vicinity of the minimizers, on the orthogonal complement of a single flat direction (which occurs because $f(ze^{i\phi}) = f(z)$ for every z, ϕ), the objective function is strongly convex.

Because of this global geometry, a wide range of efficient iterative methods can obtain a global minimizer to $f(z)$, regardless of initialization. Examples include the noisy gradient and stochastic gradient methods [GHJY15] (see also [LSJR16]), curvilinear search [Gol80] and trust-region methods

³Note that the global sign cannot be recovered.

⁴The probability is with respect to drawing of a_k 's.




Figure 2: Function landscape of (1.1) for $x = [1; 0]$ and $m \rightarrow \infty$. The only local and also global minimizers are $\pm x$. There are two saddle points near $\pm[0; 1/\sqrt{2}]$, around each there is a negative curvature direction along $\pm x$. (Left) The function graph; (Right) The same function visualized as a color image. The measurement vectors a_k 's are taken as i.i.d. standard real Gaussian in this version.

[CGT00, NP06, SQW15b]. The key property that the methods must possess is the ability to escape saddle points at which the Hessian has a strictly negative eigenvalue.

We corroborate this claim by developing a second-order trust-region method for this problem, and prove that (Theorem 3.10) (i) from any initialization, it efficiently obtains a close approximation (i.e., up to numerical precision) of the target x (up to a global phase) and (ii) it exhibits quadratic convergence in the vicinity of the global minimizers.

In sum, our geometric analysis produces the following result.

Informal Statement of Our Main Results; See Theorem 2.2 and Theorem 3.10. *When $m \geq Cn \log^3 n$, with probability at least $1 - cm^{-1}$, the function $f(z)$ has no spurious local minimizers. The only global minimizers are the target x and its equivalent copies, and at all saddle points the function has directional negative curvature. Moreover, with at least the same probability, the trust-region method with properly set step size parameter find a global minimizer of $f(z)$ in polynomial time, from an arbitrary initialization in the zero-centered complex ball with radius $R_0 \doteq 3(\frac{1}{m} \sum_{k=1}^m y_k^2)^{1/2}$. Here C and c are absolute positive constants.*

The choice of R_0 above allows us to state a result with a concise bound on the number of iterations required to converge. However, under our probability model, w.h.p., the trust-region method succeeds from any initialization. There are two caveats to this claim. First, one must choose the parameters of the method appropriately. Second, the number of iterations depends on how far away from the truth the method starts.

Our results asserts that when the a_k 's are *numerous* and *generic* enough, GPR can be solved efficiently. Similar conclusions have been obtained in [NJS13, CLS15b, CC15, WWS15]. One salient feature of our result is that the optimization method is “initialization free” - any initialization in the prescribed ball works, while all prior methods [NJS13, CLS15b, CC15, WWS15] require careful initializations that are already near the unknown target $xe^{i\phi}$. This distinctive property follows directly from the benign global geometry of $f(z)$. We believe that this sheds light on mechanism of the above numerical surprise.

1.4 Prior Arts and Connections

The survey papers [SEC⁺15, JEH15] provide accounts of recent progress on GPR. In this section, we focus on provable efficient (nonconvex) methods for GPR, and draw connections to other work on provable nonconvex heuristics for practical problems.

Provable methods for GPR. Although heuristic methods for GPR have been used effectively in practice [GS72, Fie82, SEC⁺15, JEH15], only recently have researchers begun to develop methods with provable performance guarantees. The first results of this nature were obtained using semidefinite programming (SDP) relaxations [CESV13, CSV13, CL14, CLS15a, WdM15, VX14]. While this represented a substantial advance in theory, the computational complexity of semidefinite programming limits the practicality of this approach.⁵

Recently, several provable *nonconvex* methods have been proposed for GPR. [NJS13] augmented the seminal error-reduction method [GS72] with spectral initialization and resampling to obtain the first provable nonconvex method for GPR. [CLS15b] studied the nonconvex formulation (1.1) under the same hypotheses as this paper, and showed that a combination of spectral initialization and local gradient descent recovers the true signal with near-optimal sample complexity. [CC15] worked with a different nonconvex formulation, and refined the spectral initialization and the local gradient descent with a step-adaptive truncation. With the modifications, they reduced the sample requirement to the optimal order.⁶ Compared to the SDP-based methods, these methods are more scalable and closer to methods used in practice. All three analyses are local in nature, and depend on the spectral initializer being sufficiently close to the target signal.

In contrast, we explicitly characterize the global function landscape of (1.1). Its benign geometric structure allows several algorithmic choices that need *no special initialization*. In fact, the spectral initialization used in [CLS15b] lands the iterate sequence in the region in which the objective is (restrictedly) strongly convex (\mathcal{R}_3 in Theorem 2.2). The analysis of [CLS15b] is based on a property that ensures the gradient descent method is locally contractive near the target set, which is closely linked to local convexity. [Sol14] and [WWS15] explicitly established local strong convexity near the target set for GPR in \mathbb{R}^n .

Geometric analysis of other nonconvex problems. The approach taken here is similar in spirit to our recent geometric analysis of a nonconvex formulation for complete dictionary learning [SQW15a]. For that problem, we also identified a similar geometric structure that allows efficient global optimization without special initialization. There, by analyzing the geometry of a nonconvex formulation, we derived a provable efficient algorithm for recovering square invertible dictionaries when the coefficient matrix has a constant fraction of nonzero entries. Previous results required the dictionary matrix to have far fewer nonzero entries. [SQW15b] provides a high-level overview of the common geometric structure that arises in dictionary learning, GPR and several other problems. Despite these similarities, GPR raises several novel technical challenges: the objective is heavy-tailed, and minimizing the number of measurements is important.

⁵Another line of research [BCE06, BBCE09, ABFM14] seeks to co-design the measurements and recovery algorithms based on frame- or graph-theoretic tools.

⁶In addition, [CC15] shows that the measurements can be non-adaptive, in the sense that a single, randomly chosen collection of vectors \mathbf{a}_i can simultaneously recover every $\mathbf{x} \in \mathbb{C}^n$. Results in [NJS13, CLS15b] and this paper pertain only to adaptive measurements that recover any fixed signal \mathbf{x} with high probability.

Our work sits amid the recent surge of work on provable nonconvex heuristics for practical problems. Besides GPR studied here, this line of work includes low-rank matrix recovery [KMO10, JNS13, Har14, HW14, NNS⁺14, JN14, SL14, WCCL15, SRO15, ZL15, TBSR15, CW15], tensor recovery [JO14, AGJ14a, AGJ14b, AJSN15, GHJY15], structured element pursuit [QSW14, HSSS15], dictionary learning [AAJ⁺13, AGM13, AAN13, ABGM14, AGMM15, SQW15a], mixed regression [YCS13, SA14], blind deconvolution [LWB13, LJ15, LLJB15], super resolution [EW15], phase synchronization [Bou16], numerical linear algebra [JJKN15], and so forth. Most of the methods adopt the strategy of initialization plus local refinement we alluded to above. In contrast, our geometric analysis allows flexible algorithm design (i.e., separation of geometry and algorithms) and gives some clues as to the behavior of nonconvex heuristics used in practice, which often succeed without clever initialization.

Recovering low-rank positive semidefinite matrices. The phase retrieval problem has a natural generalization to recovering low-rank positive semidefinite matrices. Consider the problem of recovering an unknown rank- r matrix $M \succeq \mathbf{0}$ in $\mathbb{R}^{n \times n}$ from linear measurement of the form $z_k = \text{tr}(\mathbf{A}_k M)$ with symmetric \mathbf{A}_k for $k = 1, \dots, m$. One can solve the problem by considering the “factorized” version: recovering $\mathbf{X} \in \mathbb{R}^{n \times r}$ (up to right invertible transform) from measurements $z_k = \text{tr}(\mathbf{X}^* \mathbf{A}_k \mathbf{X})$. This is a natural generalization of GPR, as one can write the GPR measurements as $y_k^2 = |\mathbf{a}_k^* \mathbf{x}|^2 = \mathbf{x}^* (\mathbf{a}_k \mathbf{a}_k^*) \mathbf{x}$. This generalization and related problems have recently been studied in [SRO15, ZL15, TBSR15, CW15].

1.5 Notations, Organization, and Reproducible Research

Basic notations and facts. Throughout the paper, we define complex inner product as: $\langle \mathbf{a}, \mathbf{b} \rangle \doteq \mathbf{a}^* \mathbf{b}$ for any $\mathbf{a}, \mathbf{b} \in \mathbb{C}^n$. We use $\mathbb{C}\mathbb{S}^{n-1}$ for the complex unit sphere in \mathbb{C}^n . $\mathbb{C}\mathbb{S}^{n-1}(\lambda)$ with $\lambda > 0$ denotes the centered complex sphere with radius λ in \mathbb{C}^n . Similarly, we use $\mathbb{C}\mathbb{B}^n(\lambda)$ to denote the centered complex ball of radius λ . We use $\mathcal{CN}(k)$ for a standard complex Gaussian vector of length k defined in (1.2). We reserve C and c , and their indexed versions to denote absolute constants. Their value vary with the context.

Let $\Re(z) \in \mathbb{R}^n$ and $\Im(z) \in \mathbb{R}^n$ denote the real and imaginary part of a complex vector $z \in \mathbb{C}^n$. It is easy to see that two complex vectors \mathbf{a} and \mathbf{b} are orthogonal in the geometric (real) sense if and only if $\Re(\mathbf{w}^* \mathbf{z}) = 0$.

For any z , obviously $f(z) = f(ze^{i\phi})$ for all ϕ , and the set $\{ze^{i\phi} : \phi \in [0, 2\pi)\}$ forms a one-dimensional (in the real sense) circle in \mathbb{C}^n . Throughout the paper, we reserve \mathbf{x} for the unknown target signal, and define the target set as $\mathcal{X} \doteq \{\mathbf{x}e^{i\phi} : \phi \in [0, 2\pi)\}$. Moreover, we define

$$\phi(z) \doteq \arg \min_{\phi \in [0, 2\pi)} \left\| z - \mathbf{x}e^{i\phi} \right\|, \quad \mathbf{h}(z) \doteq z - \mathbf{x}e^{i\phi(z)}, \quad \text{dist}(z, \mathcal{X}) \doteq \|\mathbf{h}(z)\|. \quad (1.3)$$

for any $z \in \mathbb{C}^n$. It is not difficult to see that $\Im(z^* \mathbf{x}e^{i\phi(z)}) = 0$ and also $\Re(z^* \mathbf{x}e^{i\phi(z)}) = |\mathbf{x}^* z|$. Moreover, $\mathbf{z}_T \doteq z / \|z\|$ and $-\mathbf{z}_T$ are the unit vectors tangent to the circle $\{ze^{i\phi} : \phi \in [0, 2\pi)\}$ at point z .

Wirtinger calculus. Consider a real-valued function $g(z) : \mathbb{C}^n \mapsto \mathbb{R}$. Unless g is constant, it is not complex differentiable. However, if one identifies \mathbb{C}^n with \mathbb{R}^{2n} and treats g as a function in the real domain, g may still be differentiable in the real sense. Doing calculus for g directly in

the real domain tends to produce cumbersome expressions. A more elegant way is adopting the Wirtinger calculus, which can be thought of a neat way of organizing the real partial derivatives. Here we only provide a minimal exposition of Wirtinger calculus; similar exposition is also given in [CLS15b]. A systematic development with emphasis on applications in optimization is provided in the article [KD09].

Let $z = \mathbf{x} + i\mathbf{y}$ where $\mathbf{x} = \Re(z)$ and $\mathbf{y} = \Im(z)$. For a complex-valued function $g(z) = u(\mathbf{x}, \mathbf{y}) + iv(\mathbf{x}, \mathbf{y})$, the Wirtinger derivative is well defined so long as the real-valued functions u and v are differentiable with respect to (w.r.t.) \mathbf{x} and \mathbf{y} . Under these conditions, the Wirtinger derivatives can be defined *formally* as

$$\begin{aligned}\frac{\partial g}{\partial z} &\doteq \left. \frac{\partial g(z, \bar{z})}{\partial z} \right|_{\bar{z} \text{ constant}} = \left[\frac{\partial g(z, \bar{z})}{\partial z_1}, \dots, \frac{\partial g(z, \bar{z})}{\partial z_n} \right] \Big|_{\bar{z} \text{ constant}} \\ \frac{\partial g}{\partial \bar{z}} &\doteq \left. \frac{\partial g(z, \bar{z})}{\partial \bar{z}} \right|_{z \text{ constant}} = \left[\frac{\partial g(z, \bar{z})}{\partial \bar{z}_1}, \dots, \frac{\partial g(z, \bar{z})}{\partial \bar{z}_n} \right] \Big|_{z \text{ constant}}.\end{aligned}$$

The notation above should only be taken at a formal level. Basically it says when evaluating $\partial g / \partial z$, one just treats \bar{z} as if it was a constant, and vice versa. To evaluate the individual partial derivatives, such as $\frac{\partial g(z, \bar{z})}{\partial z_i}$, all the usual rules of calculus apply.⁷

Note that above the partial derivatives $\frac{\partial g}{\partial z}$ and $\frac{\partial g}{\partial \bar{z}}$ are row vectors. The Wirtinger gradient and Hessian are defined as

$$\nabla g(z) = \left[\frac{\partial g}{\partial z}, \frac{\partial g}{\partial \bar{z}} \right]^* \quad \nabla^2 g(z) = \begin{bmatrix} \frac{\partial}{\partial z} \left(\frac{\partial g}{\partial z} \right)^* & \frac{\partial}{\partial \bar{z}} \left(\frac{\partial g}{\partial z} \right)^* \\ \frac{\partial}{\partial z} \left(\frac{\partial g}{\partial \bar{z}} \right)^* & \frac{\partial}{\partial \bar{z}} \left(\frac{\partial g}{\partial \bar{z}} \right)^* \end{bmatrix}, \quad (1.4)$$

where we sometimes write $\nabla_z g \doteq \left(\frac{\partial g}{\partial z} \right)^*$ and naturally $\nabla_{\bar{z}} g \doteq \left(\frac{\partial g}{\partial \bar{z}} \right)^*$. With gradient and Hessian, the second-order Taylor expansion of $g(z)$ at a point z_0 is defined as

$$\widehat{g}(\boldsymbol{\delta}; z_0) = g(z_0) + (\nabla g(z_0))^* \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 g(z_0) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}.$$

For numerical optimization, we are most interested in real-valued g . A real-valued g is stationary at a point z if and only if

$$\nabla_z g(z) = \mathbf{0}.$$

This is equivalent to the condition $\nabla_{\bar{z}} g = \mathbf{0}$, as $\nabla_z g = \overline{\nabla_{\bar{z}} g}$ when g is real-valued. The curvature of g at a stationary point z is dictated by the Wirtinger Hessian $\nabla^2 g(z)$. An important technical point is that the Hessian quadratic form involves left and right multiplication with a $2n$ -dimensional vector consisting of a conjugate pair $(\boldsymbol{\delta}, \bar{\boldsymbol{\delta}})$.

For our particular function $f(z) : \mathbb{C}^n \mapsto \mathbb{R}$ defined in (1.1), direct calculation gives

$$\nabla f(z) = \frac{1}{m} \sum_{k=1}^m \begin{bmatrix} (|\mathbf{a}_k^* z|^2 - y_k^2) (\mathbf{a}_k \mathbf{a}_k^*) z \\ (|\mathbf{a}_k^* z|^2 - y_k^2) (\mathbf{a}_k \mathbf{a}_k^*)^\top \bar{z} \end{bmatrix}, \quad (1.5)$$

⁷The precise definition is as follows: write $z = \mathbf{u} + i\mathbf{v}$. Then $\frac{\partial g}{\partial z} \doteq \frac{1}{2} \left(\frac{\partial g}{\partial \mathbf{u}} - i \frac{\partial g}{\partial \mathbf{v}} \right)$. Similarly, $\frac{\partial g}{\partial \bar{z}} \doteq \frac{1}{2} \left(\frac{\partial g}{\partial \mathbf{u}} + i \frac{\partial g}{\partial \mathbf{v}} \right)$.

$$\nabla^2 f(\mathbf{z}) = \frac{1}{m} \sum_{k=1}^m \begin{bmatrix} (2|\mathbf{a}_k^* \mathbf{z}|^2 - y_k^2) \mathbf{a}_k \mathbf{a}_k^* & (\mathbf{a}_k^* \mathbf{z})^2 \mathbf{a}_k \mathbf{a}_k^\top \\ (\mathbf{z}^* \mathbf{a}_k)^2 \overline{\mathbf{a}_k} \mathbf{a}_k^* & (2|\mathbf{a}_k^* \mathbf{z}|^2 - y_k^2) \overline{\mathbf{a}_k} \mathbf{a}_k^\top \end{bmatrix}. \quad (1.6)$$

Following the above notation, we write $\nabla_{\mathbf{z}} f(\mathbf{z})$ and $\nabla_{\overline{\mathbf{z}}} f(\mathbf{z})$ for denoting the first and second half of $\nabla f(\mathbf{z})$, respectively.

Organization. The remainder of this paper is organized as follows. In Section 2, we provide a quantitative characterization of the global geometry for GPR and highlight main technical challenges in establishing the results. Based on this characterization, in Section 3 we present a modified trust-region method for solving GPR from an arbitrary initialization, which leads to our main computational guarantee. In Section 4 we study the empirical performance of our method for GPR. Section 5 concludes the main body with a discussion of open problems. Section 6 and Section 7 collect the detailed proofs to technical results for the geometric analysis and algorithmic analysis, respectively.

Reproducible research. The code to reproduce all the figures and the experimental results can be found online:

https://github.com/sunju/pr_plain

2 The Geometry of the Objective Function

The low-dimensional example described in the introduction (Figure 2) provides some clues about the high-dimensional geometry of the objective function $f(\mathbf{z})$. Its properties can be seen most clearly through the population objective function $\mathbb{E}_a[f(\mathbf{z})]$, which can be thought of as a “large sample” version in which $m \rightarrow \infty$. We characterize this large-sample geometry in Section 2.1. In Section 2.2, we show that the most important characteristics of this large-sample geometry are present even when the number of observations m is close to the number of degrees of freedom n in the target \mathbf{x} . Section 2.3 describes several technical problems that arise in the finite sample analysis, and states a number of key intermediate results, which are proved in the appendix.

2.1 A Glimpse of the Asymptotic Function Landscape

To characterize the geometry of $\mathbb{E}_a[f(\mathbf{z})]$ (written as $\mathbb{E}[f]$ henceforth), we simply calculate the expectation of the first and second derivatives of f at each point $\mathbf{z} \in \mathbb{C}^n$. We characterize the location of the critical points of the expectation, and use second derivative information to characterize their signatures. An important conclusion is that every local minimum of $\mathbb{E}[f]$ is of the form $\mathbf{x}e^{i\phi}$, and that all other critical points have a direction of strict negative curvature:

Theorem 2.1 *When $\mathbf{x} \neq \mathbf{0}$, the only critical points of $\mathbb{E}[f]$ are $\mathbf{0}$, \mathcal{X} and $\mathcal{S} \doteq \{\mathbf{z} : \mathbf{x}^* \mathbf{z} = 0, \|\mathbf{z}\| = \|\mathbf{x}\|/\sqrt{2}\}$, which are the local maximizer, the set of local/global minimizers, and the set of saddle points, respectively. Moreover, the saddle points and local maximizer have negative curvature in the $\mathbf{x}e^{i\phi(\mathbf{z})}$ direction.*

Proof We show the statement by partitioning the space \mathbb{C}^n into several regions and analyzing each region individually using the expected gradient and Hessian. These are calculated in Lemma 6.1,

and reproduced below:

$$\mathbb{E}[f] = \|\mathbf{x}\|^4 + \|z\|^4 - \|\mathbf{x}\|^2 \|z\|^2 - |\mathbf{x}^* z|^2, \quad (2.1)$$

$$\nabla \mathbb{E}[f] = \begin{bmatrix} \nabla_z \mathbb{E}[f] \\ \nabla_{\bar{z}} \mathbb{E}[f] \end{bmatrix} = \begin{bmatrix} \left(2\|z\|^2 \mathbf{I} - \|\mathbf{x}\|^2 \mathbf{I} - \mathbf{x}\mathbf{x}^* \right) z \\ \left(2\|z\|^2 \mathbf{I} - \|\mathbf{x}\|^2 \mathbf{I} - \mathbf{x}\mathbf{x}^* \right) \bar{z} \end{bmatrix}, \quad (2.2)$$

$$\nabla^2 \mathbb{E}[f] = \begin{bmatrix} 2zz^* - \mathbf{x}\mathbf{x}^* + \left(2\|z\|^2 - \|\mathbf{x}\|^2 \right) \mathbf{I} & 2zz^\top \\ 2\bar{z}z^* & 2\bar{z}z^\top - \bar{\mathbf{x}}\mathbf{x}^\top + \left(2\|z\|^2 - \|\mathbf{x}\|^2 \right) \mathbf{I} \end{bmatrix}. \quad (2.3)$$

Based on this, we observe:

- $z = \mathbf{0}$ is a critical point, and the Hessian

$$\nabla^2 \mathbb{E}[f(\mathbf{0})] = \text{diag} \left(-\mathbf{x}\mathbf{x}^* - \|\mathbf{x}\|^2 \mathbf{I}, -\bar{\mathbf{x}}\mathbf{x}^\top - \|\mathbf{x}\|^2 \mathbf{I} \right) \prec \mathbf{0}.$$

Hence, $z = \mathbf{0}$ is a local maximizer.

- In the region $\left\{ z : 0 < \|z\|^2 < \frac{1}{2} \|\mathbf{x}\|^2 \right\}$, we have

$$\begin{bmatrix} z \\ \bar{z} \end{bmatrix}^* \nabla \mathbb{E}[f] = 2 \left(2\|z\|^2 - \|\mathbf{x}\|^2 \right) \|z\|^2 - 2|\mathbf{x}^* z|^2 < 0.$$

So there is no critical point in this region.

- When $\|z\|^2 = \frac{1}{2} \|\mathbf{x}\|^2$, the gradient is $\nabla_z \mathbb{E}[f] = -\mathbf{x}\mathbf{x}^* z$. The gradient vanishes whenever $z \in \text{null}(\mathbf{x}\mathbf{x}^*)$, which is true if and only if $\mathbf{x}^* z = 0$. Thus, we can see that any $z \in \mathcal{S}$ is a critical point. Moreover, for any $z \in \mathcal{S}$,

$$\begin{bmatrix} \mathbf{x} e^{i\phi(z)} \\ \bar{\mathbf{x}} e^{i\phi(z)} \end{bmatrix}^* \nabla^2 \mathbb{E}[f] \begin{bmatrix} \mathbf{x} e^{i\phi(z)} \\ \bar{\mathbf{x}} e^{i\phi(z)} \end{bmatrix} = -2 \|\mathbf{x}\|^4.$$

Similarly, one can show that in z direction there is positive curvature. Hence, every $z \in \mathcal{S}$ is a saddle point.

- In the region $\left\{ z : \frac{1}{2} \|\mathbf{x}\|^2 < \|z\|^2 < \|\mathbf{x}\|^2 \right\}$, any potential critical point must satisfy

$$\left(2\|z\|^2 - \|\mathbf{x}\|^2 \right) z = \mathbf{x}\mathbf{x}^* z.$$

In other words, $2\|z\|^2 - \|\mathbf{x}\|^2$ is the positive eigenvalue of the rank-one PSD Hermitian matrix $\mathbf{x}\mathbf{x}^*$. Hence $2\|z\|^2 - \|\mathbf{x}\|^2 = \|\mathbf{x}\|^2$. This would imply that $\|z\| = \|\mathbf{x}\|$, which does not occur in this region.

- When $\|z\|^2 = \|\mathbf{x}\|^2$, critical points must satisfy

$$\left(\|\mathbf{x}\|^2 \mathbf{I} - \mathbf{x}\mathbf{x}^* \right) z = \mathbf{0},$$

and so $z \notin \text{null}(\mathbf{x}\mathbf{x}^*)$. Given that $\|z\| = \|\mathbf{x}\|$, we must have $z = \mathbf{x} e^{i\theta}$ for some $\theta \in [0, 2\pi)$. Since f is a nonnegative function, and $f(z) = 0$ for any $z \in \mathcal{X}$, \mathcal{X} is indeed also the global optimal set.

- For $\|z\| > \|x\|$, since the gradient $\left[\frac{z}{\bar{z}}\right]^* \nabla \mathbb{E}[f(z)] > 0$, there is no critical point present.

Summarizing the above observations completes the proof. \blacksquare

This result suggests that the same qualitative properties that we observed for $f(z)$ with $z \in \mathbb{R}^2$ also hold for higher-dimensional, complex z . The high-dimensional analysis is facilitated by the unitary invariance of the complex normal distribution – the properties of $\mathbb{E}[f]$ at a given point z depend only the norm of z and its inner product with the target vector x , i.e., x^*z . In the next section, we will show that the important qualitative aspects of this structure are preserved even when m is as small as $Cn \log^3 n$.

2.2 The Finite-Sample Landscape

The following theorem characterizes the geometry of the objective function $f(z)$, when the number of samples m is roughly on the order of the number n of degrees of freedom in the vector x . The main conclusion is that the space \mathbb{C}^n can be divided into three regions, in which the objective either exhibits negative curvature, strong gradient, or restricted strong convexity.

The result is not surprising in view of the above characterization of the “large-sample” landscape. The intuition is as follows: since the objective function is a sum of independent random variables, when m is sufficiently large, the function values, gradients and Hessians should be uniformly close to their expectations. Some care is required in making this intuition precise, however. Because the objective function contains fourth powers of Gaussian random variables, it is heavy tailed. Ensuring that f and its derivatives are uniformly close to their expectations requires $m \geq Cn^2$. This would be quite wasteful, since x has only n degrees of freedom.

Fortunately, when $m \geq Cn \text{ polylog}(n)$, w.h.p. f still has benign global geometry, even though its gradient is not uniformly close to its expectation. Perhaps surprisingly, the heavy tailed behavior of f only helps to prevent spurious local minimizers – away from the global minimizers and saddle points, the gradient can be sporadically large, but it cannot be sporadically small. This behavior will follow by expressing the decrease of the function along a certain carefully chosen descent direction as a sum of random variables which are heavy tailed, but are also *nonnegative*. Because they are nonnegative, their deviation below their expectation is bounded, and their lower-tail is well-behaved.

Our geometric characterization of the finite-sample objective function reflects these complexities. We prove that there is a partition of \mathbb{C}^n into regions of negative curvature, large gradient, and restricted strong convexity (near the optimizer x). The gradient region is further partitioned into two sub-regions, over which different canonical descent directions are studied. Our main geometric result is as follows:

Theorem 2.2 (Main Geometric Results) *There exist positive absolute constants c_a, c_b, c_c and C , such that when $m \geq Cn \log^3 n$, it holds with probability at least $1 - c_a \exp(-c_b m / \log m) - c_c m^{-1}$ that $f(z)$ has no spurious local minimizers and the only local/global minimizers are exactly the target set \mathcal{X} . More precisely, with the same probability,*

$$\frac{1}{\|x\|^2} \left[\frac{x e^{i\phi(z)}}{\bar{x} e^{-i\phi(z)}} \right]^* \nabla^2 f(z) \left[\frac{x e^{i\phi(z)}}{\bar{x} e^{-i\phi(z)}} \right] \leq -\frac{1}{100} \|x\|^2, \quad \forall z \in \mathcal{R}_1, \quad (\text{Negative Curvature})$$

$$\frac{z^* \nabla_z f(z)}{\|z\|} \geq \frac{1}{1000} \|x\|^2 \|z\|, \quad \forall z \in \mathcal{R}_2^z, \quad (\text{Large Gradient})$$

$$\frac{\Re(\mathbf{h}(\mathbf{z})^* \nabla_{\mathbf{z}} f(\mathbf{z}))}{\|\mathbf{h}(\mathbf{z})\|} \geq \frac{1}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\|, \quad \forall \mathbf{z} \in \mathcal{R}_2^{\mathbf{h}}, \quad (\text{Large Gradient})$$

$$\left[\frac{\mathbf{g}(\mathbf{z})}{\mathbf{g}(\mathbf{z})} \right]^* \nabla^2 f(\mathbf{z}) \left[\frac{\mathbf{g}(\mathbf{z})}{\mathbf{g}(\mathbf{z})} \right] \geq \frac{1}{4} \|\mathbf{x}\|^2, \quad \forall \mathbf{z} \in \mathcal{R}_3, \quad (\text{Restricted Strong Convexity})$$

where $\mathbf{h}(\mathbf{z})$ is defined in (1.3), and

$$\mathbf{g}(\mathbf{z}) \doteq \begin{cases} \mathbf{h}(\mathbf{z}) / \|\mathbf{h}(\mathbf{z})\| & \text{if } \text{dist}(\mathbf{z}, \mathcal{X}) \neq 0, \\ \mathbf{h} \in \mathcal{S} \doteq \{\mathbf{h} : \Im(\mathbf{h}^* \mathbf{z}) = 0, \|\mathbf{h}\| = 1\} & \text{if } \mathbf{z} \in \mathcal{X}. \end{cases}$$

Here the regions \mathcal{R}_1 , $\mathcal{R}_2^{\mathbf{z}}$, $\mathcal{R}_2^{\mathbf{h}}$ and \mathcal{R}_3 cover \mathbb{C}^n , and are defined as

$$\mathcal{R}_1 \doteq \left\{ \mathbf{z} : 8|\mathbf{x}^* \mathbf{z}|^2 + \frac{401}{100} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \leq \frac{398}{100} \|\mathbf{x}\|^4 \right\}, \quad (2.4)$$

$$\mathcal{R}_2^{\mathbf{z}} \doteq \left\{ \mathbf{z} : \Re(\langle \mathbf{z}, \nabla_{\mathbf{z}} \mathbb{E}[f] \rangle) \geq \frac{1}{100} \|\mathbf{z}\|^4 + \frac{1}{500} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \right\}, \quad (2.5)$$

$$\mathcal{R}_2^{\mathbf{h}} \doteq \left\{ \mathbf{z} : \Re(\langle \mathbf{h}(\mathbf{z}), \nabla_{\mathbf{z}} \mathbb{E}[f] \rangle) \geq \frac{1}{250} \|\mathbf{x}\|^2 \|\mathbf{z}\| \|\mathbf{h}(\mathbf{z})\|, \right. \\ \left. \frac{11}{20} \|\mathbf{x}\| \leq \|\mathbf{z}\| \leq \|\mathbf{x}\|, \text{dist}(\mathbf{z}, \mathcal{X}) \geq \frac{\|\mathbf{x}\|}{3} \right\}, \quad (2.6)$$

$$\mathcal{R}_3 \doteq \left\{ \mathbf{z} : \text{dist}(\mathbf{z}, \mathcal{X}) \leq \frac{1}{\sqrt{7}} \|\mathbf{x}\| \right\}. \quad (2.7)$$

Proof The quantitative statements are proved sequentially in Proposition 2.3, Proposition 2.4, Proposition 2.5, Proposition 2.6 and Proposition 2.7 in the next section. We next show \mathcal{X} are the only local/global minimizers. Obviously local minimizers will not occur in $\mathcal{R}_1 \cup \mathcal{R}_2^{\mathbf{z}} \cup \mathcal{R}_2^{\mathbf{h}}$, as at each such point either the gradient is nonzero, or there is a negative curvature direction. So local/global minimizers can occur only in \mathcal{R}_3 . From (1.5), it is easy to check that $\nabla_{\mathbf{z}} f(\mathbf{x}e^{i\phi}) = \mathbf{0}$ and $f(\mathbf{x}e^{i\phi}) = 0$ for any $\phi \in [0, 2\pi)$. Since $f(\mathbf{z}) \geq 0$ for all $\mathbf{z} \in \mathbb{C}^n$, all elements of \mathcal{X} are local/global minimizers. To see there is no other critical point in \mathcal{R}_3 , note that any point $\mathbf{z} \in \mathcal{R}_3 \setminus \mathcal{X}$ can be written as

$$\mathbf{z} = \mathbf{x}e^{i\phi(\mathbf{z})} + t\mathbf{g}, \quad \mathbf{g} \doteq \mathbf{h}(\mathbf{z}) / \|\mathbf{h}(\mathbf{z})\|, \quad t \doteq \text{dist}(\mathbf{z}, \mathcal{X}).$$

By the restricted strong convexity we have established, and the integral form of Taylor's theorem in Lemma A.2,

$$f(\mathbf{z}) = f(\mathbf{x}e^{i\phi(\mathbf{z})}) + t \left[\frac{\mathbf{g}}{\mathbf{g}} \right]^* \nabla f(\mathbf{x}e^{i\phi(\mathbf{z})}) + t^2 \int_0^1 (1-s) \left[\frac{\mathbf{g}}{\mathbf{g}} \right]^* \nabla^2 f(\mathbf{x}e^{i\phi(\mathbf{z})} + st\mathbf{g}) \left[\frac{\mathbf{g}}{\mathbf{g}} \right] ds \geq \frac{1}{8} \|\mathbf{x}\|^2 t^2.$$

similarly, we obtain

$$f(\mathbf{x}e^{i\phi(\mathbf{z})}) = 0 \geq f(\mathbf{z}) - t \left[\frac{\mathbf{g}}{\mathbf{g}} \right]^* \nabla f(\mathbf{z}) + t^2 \int_0^1 (1-s) \left[\frac{\mathbf{g}}{\mathbf{g}} \right]^* \nabla^2 f(\mathbf{z} - st\mathbf{g}) \left[\frac{\mathbf{g}}{\mathbf{g}} \right] ds \\ \geq f(\mathbf{z}) - \left[\frac{\mathbf{g}}{\mathbf{g}} \right]^* \nabla f(\mathbf{z}) + \frac{1}{8} \|\mathbf{x}\|^2 t^2.$$

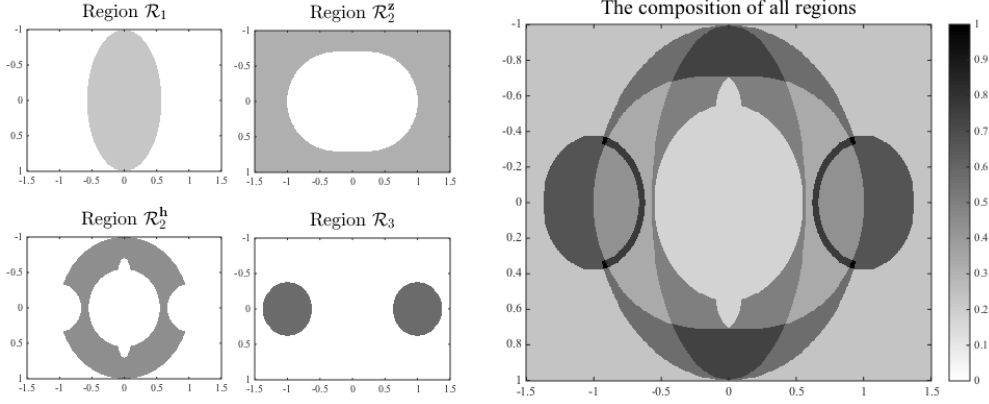


Figure 3: Schematic illustration of partitioning regions for Theorem 2.2. This plot corresponds to Figure 2, i.e., the target signal is $\mathbf{x} = [1; 0]$ and measurements are real Gaussians, such that the function is defined in \mathbb{R}^2 .

Summing up the above two inequalities, we obtain

$$t \begin{bmatrix} \mathbf{g} \\ \mathbf{g} \end{bmatrix}^* \nabla f(\mathbf{z}) \geq \frac{1}{4} \|\mathbf{x}\|^2 t^2 \implies \|\nabla f(\mathbf{z})\| \geq \frac{1}{4\sqrt{2}} \|\mathbf{x}\|^2 t,$$

as desired. ■

Figure 3 visualizes the different regions described in Theorem 2.2, and gives an idea of how they cover the space. For $f(\mathbf{z})$, a point $\mathbf{z} \in \mathbb{C}^n$ is either near a critical point such that the gradient $\nabla_{\mathbf{z}} f(\mathbf{z})$ is small (in magnitude), or far from a critical point such that the gradient is large. Any point in $\mathcal{R}_2^z \cup \mathcal{R}_2^h$ is far from a critical point, as the following is true:

$$\|\nabla_{\mathbf{z}} f(\mathbf{z})\| \geq \frac{\mathbf{z}^* \nabla_{\mathbf{z}} f(\mathbf{z})}{\|\mathbf{z}\|} \geq \frac{1}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\|, \text{ or } \|\nabla_{\mathbf{z}} f(\mathbf{z})\| \geq \frac{\Re(\mathbf{h}(\mathbf{z})^* \nabla_{\mathbf{z}} f(\mathbf{z}))}{\|\mathbf{h}(\mathbf{z})\|} \geq \frac{1}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\|.$$

The rest of the space consists of points near critical points. Since $\mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^h \cup \mathcal{R}_3$ cover the space, the rest points are included in $\mathcal{R}_1 \cup \mathcal{R}_3$. For any \mathbf{z} in \mathcal{R}_1 , the quantity

$$\frac{1}{\|\mathbf{x}\|^2} \begin{bmatrix} \mathbf{x} e^{i\phi(\mathbf{z})} \\ \overline{\mathbf{x}} e^{-i\phi(\mathbf{z})} \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \mathbf{x} e^{i\phi(\mathbf{z})} \\ \overline{\mathbf{x}} e^{-i\phi(\mathbf{z})} \end{bmatrix}$$

measures the local curvature of $f(\mathbf{z})$ in the $\mathbf{x} e^{i\phi(\mathbf{z})}$ direction. Strict negativity of this quantity implies that the neighboring critical point is either a local maximizer, or a saddle point. Moreover, $\mathbf{x} e^{i\phi(\mathbf{z})}$ is a local descent direction, even if $\nabla_{\mathbf{z}} f(\mathbf{z}) = \mathbf{0}$. For any $\mathbf{z} \in \mathcal{R}_3$, $\mathbf{g}(\mathbf{z})$ is the unit vector that points to $\mathbf{x} e^{i\phi(\mathbf{z})}$, and is also geometrically orthogonal to the $i\mathbf{x} e^{i\phi(\mathbf{z})}$ which is tangent the circle \mathcal{X} at $\mathbf{x} e^{i\phi(\mathbf{z})}$. The strict positivity of the quantity

$$\begin{bmatrix} \mathbf{g}(\mathbf{z}) \\ \mathbf{g}(\mathbf{z}) \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \mathbf{g}(\mathbf{z}) \\ \mathbf{g}(\mathbf{z}) \end{bmatrix}$$

implies that locally $f(\mathbf{z})$ is strongly convex in $\mathbf{g}(\mathbf{z})$ direction, although it is flat on the complex circle $\{\mathbf{z} e^{i\phi} : \phi \in [0, 2\pi)\}$. In particular, the result applied to $\mathbf{z} \in \mathcal{X}$ implies that on \mathcal{X} , $f(\mathbf{z})$ is strongly

convex in any direction orthogonal to \mathcal{X} . This observation, together with the fact that the Hessian is Lipschitz, implies that there is a neighborhood of \mathcal{X} on which $v^* \nabla^2 f(x) v > 0$ for *every* direction that is orthogonal to the trivial direction iz , not just the particular direction $g(z)$. This stronger property can be used to study the asymptotic convergence rate of algorithms; in particular, we will use it to obtain quadratic convergence for a certain variant of the trust-region method.

In the asymptotic version, we characterized only the critical points. In this finite-sample version, we characterize the whole space and particularly provide quantitative control for regions near critical points (i.e., $\mathcal{R}_1 \cup \mathcal{R}_3$). These concrete quantities are important for algorithm design and analysis (see Section 3).

In sum, our objective $f(z)$ has the benign geometry that each $z \in \mathbb{C}^n$ has either large gradient or directional negative curvature, or lies in the vicinity of a local minimizer around which the function is locally restrictedly strongly convex. Functions with this property lie in the rideable-saddle function class [GHJY15, SQW15b]. Functions in this class admit simple iterative methods (including the noisy gradient method, curvilinear search, and trust-region methods), which avoid being trapped near saddle points, and efficiently obtain a local minimizer. Theorem 2.2 shows that for our problem, every local minimizer is global, and so for our problem, these algorithms efficiently obtain a global minimizer.

2.3 Key Steps in the Geometric Analysis

Our proof strategy is fairly simple: we work out uniform bounds on the quantities for each of the four regions, and finally show the regions together cover the space. Since (1.1) and associated derivatives take the form of summation of m independent random variables, the proof involves concentration and covering arguments [Ver12]. The main challenge in our argument will be the heavy tailed nature of f and its gradient.

Proposition 2.3 (Negative Curvature) *When $m \geq Cn \log n$, it holds with probability at least $1 - c_a \exp(-c_b m / \log m) - c_c m^{-1}$ that*

$$\frac{1}{\|\mathbf{x}\|^2} \begin{bmatrix} \mathbf{x} e^{i\phi(z)} \\ \bar{\mathbf{x}} e^{-i\phi(z)} \end{bmatrix}^* \nabla^2 f(z) \begin{bmatrix} \mathbf{x} e^{i\phi(z)} \\ \bar{\mathbf{x}} e^{-i\phi(z)} \end{bmatrix} \leq -\frac{1}{100} \|\mathbf{x}\|^2$$

for all $z \in \mathcal{R}_1$ defined in (2.4). Here C , and c_a to c_c are positive absolute constants.

Proof See Section 6.2 on Page 30. ■

The expected gradient $\nabla_z \mathbb{E}[f(z)]$ is a linear combination of z and \mathbf{x} . We will divide \mathcal{R}_2 into two overlapped regions, \mathcal{R}_2^z and \mathcal{R}_2^h , roughly matching the case $\Re(z^* \nabla_z \mathbb{E}[f(z)]) > 0$ and the case $\Re((z - \mathbf{x} e^{i\phi(z)})^* \nabla_z \mathbb{E}[f(z)]) > 0$, respectively.

Proposition 2.4 *When $m \geq Cn \log n$, it holds with probability at least $1 - c_a \exp(-c_b m / \log m) - c_c m^{-1}$ that*

$$\frac{z^* \nabla_z f(z)}{\|z\|} \geq \frac{1}{1000} \|\mathbf{x}\|^2 \|z\|$$

for all $z \in \mathcal{R}_2^z$ defined in (2.5). Here C and c_a to c_c are positive absolute constants.

Proof See Section 6.3 on Page 30. ■

Proposition 2.5 When $m \geq Cn \log^3 n$, it holds with probability at least $1 - c_a \exp(-c_b m / \log^2 m) - c_c m^{-1}$ that

$$\Re(\mathbf{h}(\mathbf{z})^* \nabla_{\mathbf{z}} f(\mathbf{z})) \geq \frac{1}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\| \|\mathbf{h}(\mathbf{z})\|$$

for all $\mathbf{z} \in \mathcal{R}_2^h$ defined in (2.6). Here c_a to c_c and C are positive absolute constants.

Proof See Section 6.4 on Page 31. ■

Next, we show that for any $\mathbf{z} \in \mathbb{C}^n$ near \mathcal{X} , the objective f is strongly convex in the direction $\mathbf{z} - \mathbf{x}e^{i\phi(\mathbf{z})}$. This allows us to achieve a quadratic asymptotic rate of convergence with the modified trust-region algorithm we propose later.

Proposition 2.6 When $m \geq Cn \log n$ for a sufficiently large constant C , it holds with probability at least $1 - c_a m^{-1} - c_b \exp(-c_c m / \log m)$ that

$$\left[\frac{\mathbf{g}(\mathbf{z})}{\mathbf{g}(\mathbf{z})} \right]^* \nabla^2 f(\mathbf{z}) \left[\frac{\mathbf{g}(\mathbf{z})}{\mathbf{g}(\mathbf{z})} \right] \geq \frac{1}{4} \|\mathbf{x}\|^2$$

for all $\mathbf{z} \in \mathcal{R}_3$ defined in (2.7) and for all

$$\mathbf{g}(\mathbf{z}) \doteq \begin{cases} (\mathbf{z} - \mathbf{x}e^{i\phi(\mathbf{z})}) / \|\mathbf{z} - \mathbf{x}e^{i\phi(\mathbf{z})}\| & \text{if } \text{dist}(\mathbf{z}, \mathcal{X}) \neq 0, \\ \mathbf{h} \in \mathcal{S} \doteq \{\mathbf{h} : \Im(\mathbf{h}^* \mathbf{z}) = 0, \|\mathbf{h}\| = 1\} & \text{if } \mathbf{z} \in \mathcal{X}. \end{cases}$$

Here C , c_a to c_c are positive absolute constants.

Proof See Section 6.5 on Page 34. ■

Finally, we show that the regions we defined above cover the whole space. Formally,

Proposition 2.7 We have $\mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^h \cup \mathcal{R}_3 = \mathbb{C}^n$.

Proof See Section 6.6 on Page 35. ■

The main challenge is that the function (1.1) is fourth-order polynomial, and most quantities arising in the above propositions involve heavy-tailed random variables. For example, we need to control

$$\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^4 \quad \text{for all } \mathbf{z} \in \mathcal{R}_2^z \tag{2.8}$$

in proving Proposition 2.4,

$$\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \Re\left((\mathbf{z} - \mathbf{x}e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z}\right) \quad \text{for all } \mathbf{z} \in \mathcal{R}_2^h \tag{2.9}$$

in proving Proposition 2.5, and a quantity of the form

$$\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{w}|^2 |\mathbf{a}_k^* \mathbf{z}|^2 \quad \text{for all } \mathbf{w}, \mathbf{z} \quad (2.10)$$

in proving Proposition 2.6. With only $Cn \log^3 n$ samples, these quantities do not concentrate uniformly about their expectations. Fortunately, this heavy-tailed behavior does not prevent the objective function from being globally well-structured for optimization. Our bounds on the gradient and Hessian depend only on the *lower tails* of the above quantities. For (2.8) and (2.10) that are sum of independent nonnegative random variables, the lower tails concentrate uniformly as these lower-bounded variables are sub-Gaussian viewed from lower tails (see Lemma A.7 and Lemma 6.4). For (2.9), we carefully construct a proxy quantity that is summation of bounded random variables which uniformly bounds (2.9) from below.

3 Optimization by Trust-Region Method (TRM)

Based on the geometric characterization in Section 2.2, we describe a second-order trust-region algorithm that produces close approximation (i.e., up to numerical precision) to the global minimizer of (1.1) in polynomial number of steps. One interesting aspect of f in the complex space is that each point has a “circle” of equivalent points that have the same function value. Thus, we constrain each step to move “orthogonal” to the trivial direction. This simple modification helps the algorithm to converge faster in practice, and proves important to the quadratic asymptotic convergence rate in theory.

3.1 A Modified Trust-Region Algorithm

The basic idea of the trust-region method is simple: we generate a sequence of iterates $\mathbf{z}^{(0)}, \mathbf{z}^{(1)}, \dots$, by repeatedly constructing quadratic approximations $\hat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}) \approx f(\mathbf{z}^{(r)} + \boldsymbol{\delta})$, minimizing \hat{f} to obtain a step $\boldsymbol{\delta}$, and setting $\mathbf{z}^{(r+1)} = \mathbf{z}^{(r)} + \boldsymbol{\delta}$. More precisely, we approximate $f(\mathbf{z})$ around $\mathbf{z}^{(r)}$ using the second-order Taylor expansion,

$$\hat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}) = f(\mathbf{z}^{(r)}) + \begin{bmatrix} \boldsymbol{\delta} \\ \boldsymbol{\delta} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r)}) + \frac{1}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \boldsymbol{\delta} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \boldsymbol{\delta} \\ \boldsymbol{\delta} \end{bmatrix},$$

and solve

$$\text{minimize}_{\boldsymbol{\delta} \in \mathbb{C}^n} \hat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}), \quad \text{subject to} \quad \Im(\boldsymbol{\delta}^* \mathbf{z}^{(r)}) = 0, \quad \|\boldsymbol{\delta}\| \leq \Delta, \quad (3.1)$$

to obtain the step $\boldsymbol{\delta}$. In (3.1), Δ controls the trust-region size. The first linear constraint further forces the movement $\boldsymbol{\delta}$ to be geometrically orthogonal to the $i\mathbf{z}$ direction, along which the possibility for reducing the function value is limited. Enforcing this linear constraint is a strategic modification to the classical trust-region subproblem.

Reduction to the standard trust-region subproblem. The modified trust-region subproblem is easily seen to be equivalent to a classical trust-region subproblem (with no constraint) over $2n - 1$ real variables. Notice that $\{\mathbf{w} \in \mathbb{C}^n : \Im(\mathbf{w}^* \mathbf{z}^{(r)}) = 0\}$ forms a subspace of dimension $2n - 1$ over \mathbb{R}^{2n} .

Take any matrix $\mathbf{U}(\mathbf{z}^{(r)}) \in \mathbb{C}^{n \times (2n-1)}$ whose columns form an orthonormal basis for the subspace, i.e., $\Re(\mathbf{U}_i^* \mathbf{U}_j) = \delta_{ij}$ for any columns \mathbf{U}_i and \mathbf{U}_j . The subproblem can then be reformulated as (\mathbf{U} short for $\mathbf{U}(\mathbf{z}^{(r)})$)

$$\text{minimize}_{\boldsymbol{\xi} \in \mathbb{R}^{2n-1}} \widehat{f}(\mathbf{U}\boldsymbol{\xi}; \mathbf{z}^{(r)}), \quad \text{subject to} \quad \|\boldsymbol{\xi}\| \leq \Delta. \quad (3.2)$$

Let us define

$$\mathbf{g}(\mathbf{z}^{(r)}) \doteq \begin{bmatrix} \mathbf{U} \\ \mathbf{U} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r)}), \quad \mathbf{H}(\mathbf{z}^{(r)}) \doteq \begin{bmatrix} \mathbf{U} \\ \mathbf{U} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U} \\ \mathbf{U} \end{bmatrix}. \quad (3.3)$$

Then, the quadratic approximation of $f(\mathbf{z})$ around $\mathbf{z}^{(r)}$ can be rewritten as

$$\widehat{f}(\boldsymbol{\xi}; \mathbf{z}^{(r)}) = f(\mathbf{z}^{(r)}) + \boldsymbol{\xi}^\top \mathbf{g}(\mathbf{z}^{(r)}) + \frac{1}{2} \boldsymbol{\xi}^\top \mathbf{H}(\mathbf{z}^{(r)}) \boldsymbol{\xi}. \quad (3.4)$$

By structure of the Wirtinger gradient $\nabla f(\mathbf{z}^{(r)})$ and Wirtinger Hessian $\nabla^2 f(\mathbf{z}^{(r)})$, $\mathbf{g}(\mathbf{z}^{(r)})$ and $\mathbf{H}(\mathbf{z}^{(r)})$ contain only real entries.

So, any method which can solve the classical trust-region subproblem can be directly applied to the modified problem (3.1). Although the resulting problem can be nonconvex, it can be solved in polynomial time, by root-finding or SDP relaxations. Our convergence guarantees assume an exact solution of this problem; we outline how to obtain such a solution via SDP relaxation. In practice, though, even very inexact solutions of the trust-region subproblem suffice.⁸ Inexact iterative solvers for the trust-region subproblem can be engineered to avoid the need to densely represent the Hessian; these methods have the attractive property that they attempt to optimize the amount of Hessian information that is used at each iteration, in order to balance rate of convergence and computation.

In the interest of theory, we describe briefly how to apply SDP relaxation to solve problem (3.2). This SDP relaxation has the important property that it is always exact, even if the Hessian is indefinite. By introducing

$$\widehat{\boldsymbol{\xi}} = \begin{bmatrix} \boldsymbol{\xi} \\ 1 \end{bmatrix}, \quad \boldsymbol{\Xi} = \widehat{\boldsymbol{\xi}} \widehat{\boldsymbol{\xi}}^\top, \quad \mathbf{M} = \begin{bmatrix} \mathbf{H}(\mathbf{z}^{(r)}) & \mathbf{g}(\mathbf{z}^{(r)}) \\ \mathbf{g}(\mathbf{z}^{(r)})^\top & 0 \end{bmatrix},$$

we can lift problem (3.2) as a semidefinite program (SDP):

$$\min_{\boldsymbol{\Xi}} \langle \boldsymbol{\Xi}, \mathbf{M} \rangle, \quad \text{s.t.} \quad \text{tr}(\boldsymbol{\Xi}) \leq \Delta^2 + 1, \quad \langle \mathbf{E}_{2n}, \boldsymbol{\Xi} \rangle = 1, \quad \boldsymbol{\Xi} \succeq \mathbf{0}, \quad (3.5)$$

where $\mathbf{E}_{2n} = \mathbf{e}_{2n} \mathbf{e}_{2n}^\top$, and $\langle \cdot, \cdot \rangle$ reduces to the usual real inner product of real-valued matrices.

Once the subproblem (3.5) is solved to optimal $\boldsymbol{\Xi}_*$, we can perform eigen-decomposition on $\boldsymbol{\Xi}_*$ as $\boldsymbol{\Xi}_* = \mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^\top$. Let \mathbf{v} be the principle eigenvector of \mathbf{V} , and let $\boldsymbol{\xi}_*$ be the first $2n - 1$ coordinate of \mathbf{v} , then the optimum of the original TRM subproblem (3.1) is recovered as $\boldsymbol{\delta}_* = \mathbf{U} \boldsymbol{\xi}_*$.

⁸This can also be proved, in a relatively straightforward way, using the geometry of the objective f . In the interest of brevity, we do not pursue this here.

3.2 Convergence Analysis

Our convergence proof proceeds as follows. Let δ^* denote the optimizer of the trust-region sub-problem at a point z . If $\|\nabla f(z)\|$ is bounded away from zero, or $\lambda_{\min}(\nabla^2 f(z))$ is bounded below zero, we can guarantee that $\widehat{f}(\delta^*, z) - f(z) < -\varepsilon$, for some ε which depends on our bounds on these quantities. Because $f(z + \delta^*) \approx \widehat{f}(\delta^*, z) < f(z) - \varepsilon$, we can guarantee (roughly) an ε decrease in the objective function at each iteration. Because this ε is uniformly bounded away from zero over the gradient and negative curvature regions, the algorithm can take at most finitely many steps in these regions. Once it enters the strong convexity region around the global minimizers, the algorithm behaves much like a typical Newton-style algorithm; in particular, it exhibits asymptotic quadratic convergence. Below, we prove quantitative versions of these statements. We begin by stating several basic facts that are useful for the convergence proof.

Norm of the target vector and initialization. In our problem formulation, $\|\mathbf{x}\|$ is not known ahead of time. However, it can be well estimated. When $\mathbf{a} \sim \mathcal{CN}(n)$, $\mathbb{E}|\mathbf{a}^* \mathbf{x}|^2 = \|\mathbf{x}\|^2$. By Bernstein's inequality, $\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 \geq \frac{1}{9} \|\mathbf{x}\|^2$ with probability at least $1 - \exp(-cm)$. Thus, with the same probability, the quantity $R_0 \doteq 3(\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2)^{1/2}$ is an upper bound for $\|\mathbf{x}\|$. For the sake of analysis, we will assume the initialization $\mathbf{z}^{(0)}$ is an *arbitrary point* over $\mathbb{CB}^n(R_0)$. Now consider a fixed $R_1 > R_0$. By Lemma 6.3, Lemma 6.4, and the fact that $\max_{k \in [m]} \|\mathbf{a}_k\|^4 \leq 10n^2 \log^2 m$ with probability at least $1 - c_a m^{-n}$, we have that the following estimate

$$\begin{aligned} & \inf_{\mathbf{z}, \mathbf{z}': \|\mathbf{z}\| \leq R_0, \|\mathbf{z}'\| \geq R_1} f(\mathbf{z}') - f(\mathbf{z}) \\ &= \inf_{\mathbf{z}, \mathbf{z}': \|\mathbf{z}\| \leq R_0, \|\mathbf{z}'\| \geq R_1} \frac{1}{m} \sum_{k=1}^m \left[|\mathbf{a}_k^* \mathbf{z}'|^4 - |\mathbf{a}_k^* \mathbf{z}|^4 - 2 |\mathbf{a}_k^* \mathbf{z}'|^2 |\mathbf{a}_k^* \mathbf{z}|^2 + 2 |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{x}|^2 \right] \\ &\geq \inf_{\mathbf{z}, \mathbf{z}': \|\mathbf{z}\| \leq R_0, \|\mathbf{z}'\| \geq R_1} \frac{199}{200} \|\mathbf{z}'\|^4 - 10n^2 \log^2 m \|\mathbf{z}\|^4 - \frac{201}{200} \left(\|\mathbf{z}'\|^2 \|\mathbf{x}\|^2 + |\mathbf{x}^* \mathbf{z}'|^2 \right) \\ &\geq \inf_{\mathbf{z}': \|\mathbf{z}'\| \geq R_1} \frac{199}{200} \|\mathbf{z}'\|^4 - 10n^2 \log^2 m R_0^4 - \frac{201}{100} \|\mathbf{z}'\|^2 R_0^2 \end{aligned}$$

holds with probability at least $1 - c_b m^{-1} - c_c \exp(-c_d m / \log m)$, provided $m \geq Cn \log n$ for a sufficiently large C . It can be checked that when

$$R_1 = 3\sqrt{n \log m} R_0, \tag{3.6}$$

we have

$$\inf_{\mathbf{z}': \|\mathbf{z}'\| \geq R_1} \frac{199}{200} \|\mathbf{z}'\|^4 - 10n^2 \log^2 m R_0^4 - \frac{201}{100} \|\mathbf{z}'\|^2 R_0^2 \geq 40n^2 \log^2 m R_0^4.$$

Thus, we conclude that when $m \geq Cn \log n$, w.h.p., the sublevel set $\{\mathbf{z} : f(\mathbf{z}) \leq f(\mathbf{z}^{(0)})\}$ is contained in the set

$$\Gamma \doteq \mathbb{CB}^n(R_1). \tag{3.7}$$

Lipschitz Properties We write $\mathbf{A} \doteq [\mathbf{a}_1, \dots, \mathbf{a}_m]$ so that $\|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2} = \max_{k \in [m]} \|\mathbf{a}_k\|$. We next provide estimates of Lipschitz constants of f and its derivatives, restricted to a slightly larger region than Γ :

Lemma 3.1 (Local Lipschitz Properties) *The Lipschitz constants for $f(\mathbf{z})$, $\nabla f(\mathbf{z})$, and $\nabla^2 f(\mathbf{z})$ over the set $\Gamma' \doteq \mathbb{C}\mathbb{B}^n(2R_1)$, denoted as L_f , L_g , and L_h respectively, can be taken as*

$$\begin{aligned} L_f &\doteq 7 \times 10^6 \cdot (n \log m)^{\frac{3}{2}} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^3, & L_g &\doteq 19000\sqrt{2}n \log m \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^2, \\ L_h &\doteq 480 \cdot (n \log m)^{\frac{1}{2}} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\| \end{aligned}$$

with probability at least $1 - c_a \exp(-c_b m)$, provided $m \geq Cn$ for a sufficiently large absolute constant C . Here c_a through c_e are positive absolute constants.

Proof See Section 7.2 on Page 38. ■

Property of Hessians near the Target Set \mathcal{X} . Define a region

$$\mathcal{R}'_3 \doteq \left\{ \mathbf{z} : \|\mathbf{h}(\mathbf{z})\| \leq \frac{1}{10L_h} \|\mathbf{x}\|^2 \right\}. \quad (3.8)$$

We will provide spectral upper and lower bounds for the (restricted) Hessian matrices $\mathbf{H}(\mathbf{z})$, where $\mathbf{H}(\mathbf{z})$ is as defined in (3.3). These bounds follow by bounding $\mathbf{H}(\mathbf{z})$ on \mathcal{X} , and then using the Lipschitz property of the Hessian to extend the bounds to a slightly larger region around \mathcal{X} .

Lemma 3.2 (Lower and Upper Bounds of Restricted Hessian in \mathcal{R}'_3) *When $m \geq Cn \log n$, it holds with probability at least $1 - c_a m^{-1} - c_b \exp(-c_c m / \log m)$ that*

$$m_H \mathbf{I} \preceq \mathbf{H}(\mathbf{z}) \preceq M_H \mathbf{I}$$

for all $\mathbf{z} \in \mathcal{R}'_3$ with $m_H = 22/25 \|\mathbf{x}\|^2$ and $M_H = 9/2 \|\mathbf{x}\|^2$. Here C , c_a to c_c are positive absolute constants.

Proof See Section 7.3 on Page 39. ■

3.3 Proof of TRM Convergence

We are now ready to prove the convergence of the TRM. Throughout, we will assume $m \geq Cn \log^3 n$ for a sufficiently large constant C , so that all the events of interest hold w.h.p..

Our initialization is an arbitrary point $\mathbf{z}^{(0)} \in \mathbb{C}\mathbb{B}^n(R_0) \subseteq \Gamma$. We will analyze effect of a trust-region step from any iterate $\mathbf{z}^{(r)} \in \Gamma$. Based on these arguments, we will show that whenever $\mathbf{z}^{(r)} \in \Gamma$, $\mathbf{z}^{(r+1)} \in \Gamma$, and so the entire iterate sequence remains in Γ . The analysis will use the fact that f and its derivatives are Lipschitz over the trust-region $\mathbf{z} + \mathbb{C}\mathbb{B}^n(\Delta)$. This follows from Proposition 3.1, provided

$$\Delta \leq R_1. \quad (3.9)$$

The next auxiliary lemma makes precise the intuition that whenever there exists a descent direction, the step size parameter Δ is sufficiently small, a trust-region step will decrease the objective.

Lemma 3.3 For any $z \in \Gamma$, suppose there exists a vector δ with $\|\delta\| \leq \Delta$ such that

$$\Im(\delta^* z) = 0 \quad \text{and} \quad f(z + \delta) \leq f(z) - d,$$

for a certain $d > 0$. Then the trust-region subproblem (3.1) returns a point δ_* with $\|\delta_*\| \leq \Delta$ and

$$f(z + \delta_*) \leq f(z) - d + \frac{2}{3}L_h\Delta^3.$$

Proof See Section 7.4 on Page 40. ■

The next proposition says when Δ is chosen properly, a trust-region step from a point with negative local curvature decreases the function value by a concrete amount.

Proposition 3.4 (Function Value Decrease in Negative Curvature Region \mathcal{R}_1) Suppose the current iterate $z^{(r)} \in \mathcal{R}_1 \cap \Gamma$, and our trust-region size satisfies

$$\Delta \leq \frac{1}{400L_h} \|\mathbf{x}\|^2. \quad (3.10)$$

Then an optimizer δ_* to (3.1) leads to $z^{(r+1)} = z^{(r)} + \delta_*$ that obeys

$$f(z^{(r+1)}) - f(z^{(r)}) \leq -d_1 \doteq -\frac{1}{400}\Delta^2 \|\mathbf{x}\|^2. \quad (3.11)$$

Proof See Section 7.5 on Page 40. ■

The next proposition shows that when Δ is chosen properly, a trust-region step from a point with strong gradient decreases the objective by a concrete amount.

Proposition 3.5 (Function Value Decrease in Large Gradient Region \mathcal{R}_2) Suppose our current iterate $z^{(r)} \in (\mathcal{R}_2^z \cup \mathcal{R}_2^h) \cap \mathcal{R}_1^c \cap \Gamma$, and our trust-region size satisfies

$$\Delta \leq \min \left\{ \frac{\|\mathbf{x}\|^3}{8000L_g}, \sqrt{\frac{3\|\mathbf{x}\|^3}{16000L_h}} \right\}. \quad (3.12)$$

Then an optimizer δ_* to (3.1) leads to $z^{(r+1)} = z^{(r)} + \delta_*$ that obeys

$$f(z^{(r+1)}) - f(z^{(r)}) \leq -d_2 \doteq -\frac{1}{4000}\Delta \|\mathbf{x}\|^3. \quad (3.13)$$

Proof See Section 7.6 on Page 41. ■

Now, we argue about \mathcal{R}_3 , in which the behavior of the algorithm is more complicated. For the region $\mathcal{R}_3 \setminus \mathcal{R}'_3$, the restricted strong convexity in radial directions around \mathcal{X} as established in Proposition 2.6 implies that the gradient at any point in $\mathcal{R}_3 \setminus \mathcal{R}'_3$ is nonzero. Thus, one can treat this as another strong gradient region, and carry out essentially the same argument as in Proposition 3.5.

Proposition 3.6 (Function Value Decrease in $\mathcal{R}_3 \setminus \mathcal{R}'_3$) Suppose our current iterate $\mathbf{z}^{(r)} \in \mathcal{R}_3 \setminus \mathcal{R}'_3$, and our trust-region size satisfies

$$\Delta \leq \min \left\{ \frac{\|\mathbf{x}\|^4}{160L_hL_g}, \sqrt{\frac{3}{320} \frac{\|\mathbf{x}\|^2}{L_h}} \right\}. \quad (3.14)$$

Then an optimizer δ_* to (3.1) leads to $\mathbf{z}^{(r+1)} = \mathbf{z}^{(r)} + \delta_*$ that obeys

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -d_3 \doteq -\frac{1}{80L_h} \Delta \|\mathbf{x}\|^4. \quad (3.15)$$

Proof See Section 7.7 on Page 42. ■

Our next several propositions show that when the iterate sequence finally moves into \mathcal{R}'_3 , it can be divided into two ordered phases, either of which can be absent: first, constrained steps in which the constraint of the trust-region subproblem is active, and second, unconstrained steps in which the trust-region constraint is inactive. The next proposition shows that when Δ is chosen properly, a constrained step in \mathcal{R}'_3 decreases the objective by a concrete amount.

Proposition 3.7 Suppose our current iterate $\mathbf{z}^{(r)} \in \mathcal{R}'_3$, and the trust-region subproblem takes a constrained step, i.e., the optimizer to (3.1) satisfies $\|\delta_*\| = \Delta$. We have the δ_* leads to

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -d_4 \doteq -\frac{m_H^2 \Delta^2}{4M_H}. \quad (3.16)$$

provided that

$$\Delta \leq m_H^2 / (4M_H L_h). \quad (3.17)$$

Here m_H and M_H are as defined in Lemma 3.2.

Proof See Section 7.8 on Page 43. ■

The next proposition shows that when Δ is properly tuned, an unconstrained step in \mathcal{R}'_3 dramatically reduces the norm of the gradient.

Proposition 3.8 (Quadratic Convergence of the Norm of the Gradient) Suppose our current iterate $\mathbf{z}^{(r)} \in \mathcal{R}'_3$, and the trust-region subproblem takes an unconstrained step, i.e., the unique optimizer to (3.1) satisfies $\|\delta_*\| < \Delta$. We have the δ_* leads to $\mathbf{z}^{(r+1)} = \mathbf{z}^{(r)} + \delta_*$ that obeys

$$\|\nabla f(\mathbf{z}^{(r+1)})\| \leq \frac{1}{m_H^2} (L_h + \frac{32}{\|\mathbf{x}\|} M_H) \|\nabla f(\mathbf{z}^{(r)})\|^2, \quad (3.18)$$

provided

$$\Delta \leq \|\mathbf{x}\| / 10. \quad (3.19)$$

Here M_H and m_H are as defined in Lemma 3.2.

Proof See Section 7.9 on Page 44. ■

The next proposition shows that when Δ is properly tuned, as soon as an unconstrained \mathcal{R}'_3 step is taken, all future iterations take unconstrained \mathcal{R}'_3 steps. Moreover, the sequence converges quadratically to the target set \mathcal{X} .

Proposition 3.9 (Quadratic Convergence of the Iterates in \mathcal{R}'_3) *Suppose the trust-region algorithm starts to take an unconstrained step in \mathcal{R}'_3 at $\mathbf{z}^{(r)}$ for a certain $r \in \mathbb{N}$. Then all future steps will be unconstrained steps in \mathcal{R}'_3 , and*

$$\left\| \mathbf{h}(\mathbf{z}^{(r+r')}) \right\| \leq \frac{4\sqrt{2}m_H^2}{\|\mathbf{x}\|^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right)^{-1} 2^{-2r'} \quad (3.20)$$

for all integers $r' \geq 1$, provided that

$$\Delta \leq \min \left\{ \frac{\|\mathbf{x}\|}{10}, \frac{m_H \|\mathbf{x}\|^2}{M_H \sqrt{40\sqrt{2}L_h(L_h + 32M_H/\|\mathbf{x}\|)}}, \frac{m_H^3}{\sqrt{2}M_H^2(L_h + 32M_H/\|\mathbf{x}\|)} \right\}. \quad (3.21)$$

Proof See Section 7.10 on Page 46. ■

Now we are ready to piece together the above technical propositions to prove our main algorithmic theorem.

Theorem 3.10 (TRM Convergence) *Suppose $m \geq Cn \log^3 n$ for a sufficiently large constant C . Then with probability at least $1 - c_a m^{-1}$, the trust-region algorithm with an arbitrary initialization $\mathbf{z}^{(0)} \in \mathbb{C}\mathbb{B}^n(R_0)$, with $R_0 = 3(\frac{1}{m} \sum_{k=1}^m y_k^2)^{1/2}$, will return a solution that is ε -close to the target set \mathcal{X} in*

$$\frac{c_b}{\Delta^2 \|\mathbf{x}\|^2} f(\mathbf{z}^{(0)}) + \log \log \left(\frac{c_c \|\mathbf{x}\|}{\varepsilon} \right) \quad (3.22)$$

steps, provided that

$$\Delta \leq c_d (n^{7/2} \log^{7/2} m)^{-1} \|\mathbf{x}\|. \quad (3.23)$$

Here c_a through c_d are positive absolute constants.

Proof When $m \geq C_1 n \log^3 n$ for a sufficiently large constant C_1 , the assumption of Theorem 2.2 is satisfied. Moreover, with probability at least $1 - c_2 m^{-1}$, the following estimates hold:

$$L_f = C_3 n^{5/2} \log^{5/2} m \|\mathbf{x}\|^3, \quad L_g = C_3 n^2 \log^2 m \|\mathbf{x}\|^2, \quad L_h = C_3 n^{3/2} \log^{3/2} m \|\mathbf{x}\|, \\ m_H = 22/25 \|\mathbf{x}\|^2, \quad M_H = 9/2 \|\mathbf{x}\|^2$$

for a certain positive absolute constant C_3 . From the technical lemmas and propositions in Section 3.3, it can be verified that when

$$\Delta \leq c_4 (n^{7/2} \log^{7/2} m)^{-1} \|\mathbf{x}\|,$$

for a positive absolute constant c_4 , all requirements on Δ are satisfied.

Write $\mathcal{R}_A \doteq \Gamma \setminus \mathcal{R}'_3$, where $\Gamma \doteq \mathbb{C}\mathbb{B}^n(R_1)$ with $R_1 = 3\sqrt{n \log m} R_0$. Then a step in Γ is either a \mathcal{R}_A or constrained \mathcal{R}'_3 step that reduces the objective value by a concrete amount, or an unconstrained \mathcal{R}'_3 step with all subsequent steps being unconstrained \mathcal{R}'_3 . From discussion in Section 3.2, for an arbitrary initialization $\mathbf{z}^{(0)} \in \Gamma$, our choice of R_1 ensures that w.h.p. the sublevel set $\Pi \doteq \{\mathbf{z} : f(\mathbf{z}) \leq f(\mathbf{z}^{(0)})\}$ is contained in Γ . Moreover, \mathcal{R}'_3 is also contained in Γ . \mathcal{R}_A and constrained \mathcal{R}'_3 steps reduce the objective function, and therefore cannot cause the iterate sequence to leave Π . So

\mathcal{R}_A and constrained \mathcal{R}'_3 steps stay within Γ . Since unconstrained \mathcal{R}'_3 steps stay within \mathcal{R}'_3 , they also stay within Γ , and the iterate sequence as a whole does not leave Γ .

In fact, the previous argument implies a generic iterate sequence consists of two phases: the first phase that takes consecutive \mathcal{R}_A or constrained \mathcal{R}'_3 steps, and the second phase that takes consecutive unconstrained \mathcal{R}'_3 steps. Either of the two can be absent depending on the initialization and parameter setting for the TRM algorithm.

By Proposition 3.4, 3.5, 3.6, and 3.7, from $z^{(0)}$ it takes at most

$$f(z^{(0)}) / \min(d_1, d_2, d_3, d_4)$$

for the iterate sequence to enter \mathcal{R}'_3 . Let r_0 denote the first iteration for which $z^{(r_0)} \in \mathcal{R}'_3$. Once the sequence enters \mathcal{R}'_3 , there are three possibilities:

- The sequence always takes constrained steps in \mathcal{R}'_3 and since the function $f(z)$ is lower bounded (≥ 0), it reaches the target set \mathcal{X} in finitely many steps.
- The sequence takes constrained steps until reaching certain point $z' \in \mathcal{R}'_3$ such that $f(z') \leq f(x) + d_4 = d_4$, where d_4 is defined in Proposition 3.7. Since a constrained step in \mathcal{R}'_3 must decrease the function value by at least d_4 , all future steps must be unconstrained. Proposition 3.9 suggests that the sequence will converge quadratically to the target set \mathcal{X} .
- The sequence starts to take unconstrained steps at a certain point $z'' \in \mathcal{R}'_3$ such that $f(z'') \geq f(x) + d_4$. Again Proposition 3.9 implies that the sequence will converge quadratically to the target set \mathcal{X} .

In sum, by Proposition 3.4, Proposition 3.5, Proposition 3.6, Proposition 3.7, and Proposition 3.9, the number of iterations to obtain an ε -close solution to the target set \mathcal{X} can be grossly bounded by

$$\#\text{Iter} \leq \frac{f(z^{(0)})}{\min\{d_1, d_2, d_3, d_4\}} + \log \log \left(\frac{4\sqrt{2}m_H^2}{(L_h + 32M_H / \|\mathbf{x}\|) \|\mathbf{x}\|^2 \varepsilon} \right).$$

Using our previous estimates of m_H , M_H , and L_H , and taking $\min\{d_1, d_2, d_3, d_4\} = c_5 \Delta^2 \|\mathbf{x}\|^2$, we arrive at the claimed result. ■

4 Numerical Simulations

Our convergence analysis for the TRM is based on two idealizations: (i) the trust-region subproblem is solved exactly; and (ii) the step-size is fixed to be sufficiently small. These simplifications ease the analysis, but also render the TRM algorithm impractical. In practice, the trust-region subproblem is never exactly solved, and the trust-region step size is adjusted to the local geometry, by backtracking. It is relatively straightforward to modify our analysis to account for inexact subproblem solvers; for sake of brevity, we do not pursue this here.

In this section, we investigate experimentally the number of measurements m required to ensure that $f(z)$ is well-structured, in the sense of our theorems. This entails solving large instances of $f(z)$. To this end, we deploy the Manopt toolbox [BMAS14]⁹. Manopt is a user-friendly Matlab toolbox

⁹Available online: <http://www.manopt.org>.

that implements several sophisticated solvers for tackling optimization problems on Riemannian manifolds. The most developed solver is based on the TRM. This solver uses the truncated conjugate gradient (tCG; see, e.g., Section 7.5.4 of [CGT00]) method to (approximately) solve the trust-region subproblem (vs. the exact solver in our analysis). It also dynamically adjusts the step size. However, the original implementation (Manopt 2.0) is not adequate for our purposes. Their tCG solver uses the gradient as the initial search direction, which does not ensure that the TRM solver can escape from saddle points [ABG07, AMS09]. We modify the tCG solver, such that when the current gradient is small and there is a negative curvature direction (i.e., the current point is near a saddle point or a local maximizer for $f(z)$), the tCG solver explicitly uses the negative curvature direction¹⁰ as the initial search direction. This modification ensures the TRM solver always escapes saddle points/local maximizers with directional negative curvature. Hence, the modified TRM algorithm based on Manopt is expected to have the same qualitative behavior as the idealized version we analyzed.

We fix $n = 1,000$ and vary the ratio m/n from 4 to 10. For each m , we generate a fixed instance: a

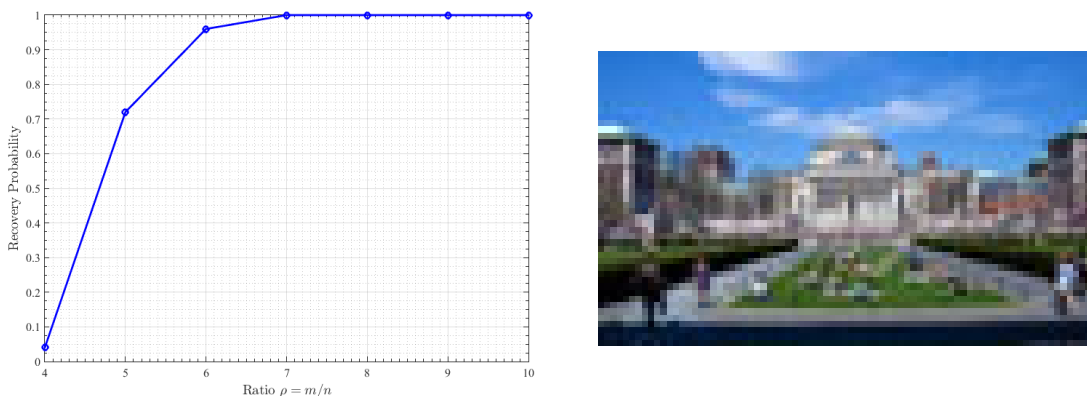


Figure 4: (Left) Recovery performance for GPR when optimizing (1.1) with the TRM. With $n = 1000$ and m varying, we consider a fixed problem instance for each m , and run the TRM algorithm 25 times from independently random initializations. The empirical recovery probability is a test of whether the benign geometric structure holds. (Right) A small “artistic” Columbia University campus image we use for comparing TRM and gradient descent.

fixed signal \mathbf{x} , and a fixed set of complex Gaussian vectors. We run the TRM algorithm 25 times for each problem instance, with independent random initializations. Successfully recovery is declared if at termination the optimization variable \mathbf{z}_∞ satisfies

$$\varepsilon_{\text{Rel}} \doteq \|\mathbf{z}_\infty - \mathbf{x}e^{i\phi(\mathbf{z}_\infty)}\| / \|\mathbf{x}\| \leq 10^{-3}.$$

The recovery probability is empirically estimated from the 25 repetitions for each m . Intuitively, when the recovery probability is below one, there are spurious local minimizers. In this case, the number of samples m is not large enough to ensure the finite-sample function landscape $f(z)$ to be qualitatively the same as the asymptotic version $\mathbb{E}_a[f(z)]$. Figure 4 shows the recovery performance. It seems that $m = 7n$ samples may be sufficient to ensure the geometric property holds.¹¹ On

¹⁰...adjusted in sign to ensure positive correlation with the gradient – if it does not vanish.

¹¹This prescription should be taken with a grain of salt, as here we have only tested a single fixed n .

the other hand, $m = 6n$ is not sufficient, whereas in theory it is known $4n$ samples are enough to guarantee measurement injectivity for complex signals [BCE06].¹²

We now briefly compare TRM and gradient descent in terms of running time. We take a small ($n = 80 \times 47$) image of Columbia University campus (Figure 4 (Right)), and make $m = 5n \log n$ complex Gaussian measurements. The TRM solver is the same as above, and the gradient descent solver is one with backtracking line search. We repeat the experiment 10 times, with independently generated random measurements and initializations each time. On average, the TRM solver returns a solution with $\varepsilon_{\text{Rel}} \leq 10^{-4}$ in about 2600 seconds, while the gradient descent solver produces a solution with $\varepsilon_{\text{Rel}} \sim 10^{-2}$ in about 6400 seconds. The point here is not to exhaustively benchmark the two – they both involve many implementation details and tuning parameters and they have very different memory requirements. It is just to suggest that second-order methods can be implemented in a practical manner for large-scale GPR problems.¹³

5 Discussion

In this work, we provide a complete geometric characterization of the nonconvex formulation (1.1) for the GPR problem. The benign geometric structure allows us to design a second-order trust-region algorithm that efficiently finds a global minimizer of (1.1), without special initialization. We close this paper by discussing possible extensions and relevant open problems.

Sample complexity and measurement schemes. Our result (Theorem 2.2 and Theorem 3.10) indicates that $m \geq C_1 n \log^3(n)$ samples are sufficient to guarantee the favorable geometric property and efficient recovery, while our simulations suggested that $C_2 n \log(n)$ or even $C_3 n$ is enough. For efficient recovery only, $m \geq C_4 n$ are known to be sufficient [CC15] (and for all signals; see also [CLS15b]). It is interesting to see if the gaps can be closed. Our current analysis pertains to Gaussian measurements only which are not practical, it is important to extend the geometric analysis to more practical measurement schemes, such as t-designs [GKK13] and masked Fourier transform measurements [CLS15a]. A preliminary study of the low-dimensional function landscape for the latter scheme produces very positive result; see Figure 5.

Sparse phase retrieval. A special case of GPR is when the underlying signal x is known to be sparse, which can be considered as a quadratic compressed sensing problem [OYVS13, OYDS13, OYDS12, LV13, JOH13, SBE14]. Since x is sparse, the lifted matrix $X = xx^*$ is sparse and has rank one. Thus, existing convex relaxation methods [OYVS13, OYDS13, LV13, JOH13] formulated it as a simultaneously low-rank and sparse recovery problem. For the latter problem, however, known convex relaxations are suboptimal [OJF⁺12, MHWG14]. Let k be the number of nonzeros in the target signal. [LV13, JOH13] showed that natural convex relaxations require $C_5 k^2 \log n$ samples for correct recovery, instead of the optimal order $O(k \log(n/k))$. A similar gap is also observed with

¹²Numerics in [CC15] suggest that under the same measurement model, $m = 5n$ is sufficient for efficient recovery. Our requirement on control of the whole function landscape and hence “initialization-free” algorithm may need the additional complexity.

¹³The main limitation in this experiment was not the TRM solver, but the need to store the vectors a_1, \dots, a_m . For other measurement models, such as the coded diffraction model [CLS15a], “matrix-free” calculation is possible, and storage is no longer a bottleneck.




Figure 5: Function landscape of (1.1) for $\mathbf{x} = [1; 0]$ and $m \rightarrow \infty$ for the masked Fourier transform measurements (coded diffraction model [CLS15a]). The landscape is qualitatively similar to that for the Gaussian model (Figure 2).

certain nonconvex methods [CLM15]. It is tempting to ask whether novel nonconvex formulations and analogous geometric analysis as taken here could shed light on this problem.

Other structured nonconvex problems. We have mentioned recent surge of works on provable nonconvex heuristics [JNS13, Har14, HW14, NNS⁺14, JN14, SL14, JO14, WCCL15, SRO15, ZL15, TBSR15, CW15, AGJ14a, AGJ14b, AJSN15, GHJY15, QSW14, HSSS15, AAJ⁺13, AGM13, AAN13, ABGM14, AGMM15, SQW15a, YCS13, SA14, LWB13, LJ15, LLJB15, EW15, Bou16, JJKN15]. While the initialization plus local refinement analyses generally produce interesting theoretical results, they do not explain certain empirical successes that do not rely on special initializations. The geometric structure and analysis we work with in our recent work [SQW15a, SQW15b] (see also [GHJY15] and [AG16]) seem promising in this regard. It is interesting to consider whether analogous geometric structure exists for other practical problems.

6 Proofs of Technical Results for Function Landscape

6.1 Auxiliary Lemmas

Lemma 6.1 For the function $f(\mathbf{z}) : \mathbb{C}^n \mapsto \mathbb{R}$ defined in (1.1), we have

$$\mathbb{E}[f(\mathbf{z})] = \|\mathbf{x}\|^4 + \|\mathbf{z}\|^4 - \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 - |\mathbf{x}^* \mathbf{z}|^2, \quad (6.1)$$

$$\nabla \mathbb{E}[f(\mathbf{z})] = \begin{bmatrix} \nabla_{\mathbf{z}} \mathbb{E}[f(\mathbf{z})] \\ \nabla_{\bar{\mathbf{z}}} \mathbb{E}[f(\mathbf{z})] \end{bmatrix} = \begin{bmatrix} \left(2\|\mathbf{z}\|^2 \mathbf{I} - \|\mathbf{x}\|^2 \mathbf{I} - \mathbf{x}\mathbf{x}^* \right) \mathbf{z} \\ \left(2\|\mathbf{z}\|^2 \mathbf{I} - \|\mathbf{x}\|^2 \mathbf{I} - \mathbf{x}\mathbf{x}^* \right) \bar{\mathbf{z}} \end{bmatrix}, \quad (6.2)$$

$$\nabla^2 \mathbb{E}[f(\mathbf{z})] = \begin{bmatrix} 2\mathbf{z}\mathbf{z}^* - \mathbf{x}\mathbf{x}^* + \left(2\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2 \right) \mathbf{I} & 2\mathbf{z}\mathbf{z}^\top \\ 2\bar{\mathbf{z}}\mathbf{z}^* & 2\bar{\mathbf{z}}\mathbf{z}^\top - \bar{\mathbf{x}}\mathbf{x}^\top + \left(2\|\mathbf{z}\|^2 - \|\mathbf{x}\|^2 \right) \mathbf{I} \end{bmatrix}. \quad (6.3)$$

Proof By definition (1.1), notice that

$$\mathbb{E}[f(\mathbf{z})] = \frac{1}{2} \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[\left(|\langle \mathbf{a}, \mathbf{x} \rangle|^2 - |\langle \mathbf{a}, \mathbf{z} \rangle|^2 \right)^2 \right]$$

$$= \frac{1}{2} \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{x} \rangle|^4 \right] + \frac{1}{2} \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{z} \rangle|^4 \right] - \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{x} \rangle|^2 |\langle \mathbf{a}, \mathbf{z} \rangle|^2 \right].$$

We now evaluate the three terms separately. Note that the law \mathcal{CN} is invariant to unitary transform. Thus,

$$\mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{x} \rangle|^4 \right] = \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{e}_1 \rangle|^4 \right] \|\mathbf{x}\|^4 = \mathbb{E}_{a \sim \mathcal{N}(0,1/2) + i \mathcal{N}(0,1/2)} \left[|a|^4 \right] \|\mathbf{x}\|^4 = 2 \|\mathbf{x}\|^4.$$

Similarly, we also obtain $\mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{z} \rangle|^4 \right] = 2 \|\mathbf{z}\|^4$. Now for the cross term,

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{x} \rangle|^2 |\langle \mathbf{a}, \mathbf{z} \rangle|^2 \right] \\ &= \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|\langle \mathbf{a}, \mathbf{e}_1 \rangle|^2 \left| \left\langle \mathbf{a}, s_1 e^{i\phi_1} \mathbf{e}_1 + s_2 e^{i\phi_2} \mathbf{e}_2 \right\rangle \right|^2 \right] \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \quad [\text{where } s_1^2 + s_2^2 = 1] \\ &= \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|a_1|^2 \left| s_1 \bar{a}_1 e^{i\phi_1} + s_2 \bar{a}_2 e^{i\phi_2} \right|^2 \right] \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \\ &= \mathbb{E}_{\mathbf{a} \sim \mathcal{CN}} \left[|a_1|^2 \left(s_1^2 |a_1|^2 + s_2^2 |a_2|^2 \right) \right] \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \\ &= (1 + s_1^2) \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 = \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 + |\mathbf{x}^* \mathbf{z}|^2. \end{aligned}$$

Gathering the above results, we obtain (6.1). By taking Wirtinger derivative (1.4) with respect to (6.1), we obtain the Wirtinger gradient and Hessian in (6.2), (6.3) as desired. \blacksquare

Lemma 6.2 For $\mathbf{a} \sim \mathcal{CN}(n)$ and any fixed vector $\mathbf{v} \in \mathbb{C}^n$, it holds that

$$\mathbb{E} \left[|\mathbf{a}^* \mathbf{v}|^2 \mathbf{a} \mathbf{a}^* \right] = \mathbf{v} \mathbf{v}^* + \|\mathbf{v}\|^2 \mathbf{I}, \quad \text{and} \quad \mathbb{E} \left[(\mathbf{a}^* \mathbf{v})^2 \mathbf{a} \mathbf{a}^\top \right] = 2 \mathbf{v} \mathbf{v}^\top.$$

Proof Observe that for $i \neq j$,

$$\mathbf{e}_i^* \mathbb{E} \left[|\mathbf{a}^* \mathbf{v}|^2 \mathbf{a} \mathbf{a}^* \right] \mathbf{e}_j = \sum_{q,\ell} \mathbb{E} \left[\overline{a(q)} a(\ell) v(q) \overline{v(\ell)} a(i) \overline{a(j)} \right] = \mathbb{E} \left[|a(i)|^2 |a(j)|^2 \right] v(i) \overline{v(j)} = v(i) \overline{v(j)}.$$

Similarly,

$$\begin{aligned} \mathbf{e}_i^* \mathbb{E} \left[|\mathbf{a}^* \mathbf{v}|^2 \mathbf{a} \mathbf{a}^* \right] \mathbf{e}_i &= \sum_{q,\ell} \mathbb{E} \left[\overline{a(q)} a(\ell) v(q) \overline{v(\ell)} |a(i)|^2 \right] \\ &= \mathbb{E} \left[|a(i)|^4 |v(i)|^2 \right] + \sum_{q \neq i} \mathbb{E} \left[|a(q)|^2 |v(q)|^2 |a(i)|^2 \right] = |v(i)|^2 + \|\mathbf{v}\|^2. \end{aligned}$$

Similar calculation yields the second expectation. \blacksquare

Lemma 6.3 Let $\mathbf{a}_1, \dots, \mathbf{a}_m$ be i.i.d. copies of $\mathbf{a} \sim \mathcal{CN}(n)$. For any $\delta \in (0, 1)$ and any $\mathbf{v} \in \mathbb{C}^n$, when $m \geq C(\delta)n \log n$, we have that with probability at least $1 - c_a \delta^{-2} m^{-1} - c_b \exp(-c_c \delta^2 m / \log m)$

$$\begin{aligned} \left\| \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{v}|^2 \mathbf{a}_k \mathbf{a}_k^* - (\mathbf{v} \mathbf{v}^* + \|\mathbf{v}\|^2 \mathbf{I}) \right\| &\leq \delta \|\mathbf{v}\|^2, \\ \left\| \frac{1}{m} \sum_{k=1}^m (\mathbf{a}_k^* \mathbf{v})^2 \mathbf{a}_k \mathbf{a}_k^\top - 2 \mathbf{v} \mathbf{v}^\top \right\| &\leq \delta \|\mathbf{v}\|^2. \end{aligned}$$

Here $C(\delta)$ is a constant depending on δ and c_a, c_b and c_c are positive absolute constants.

Proof We work out the results on $\frac{1}{m} \sum_{k=1}^m |a_k^* v|^2 a_k a_k^*$ first. By the unitary invariance of the Gaussian measure and rescaling, it is enough to consider $v = e_1$. We partition each vector a_k as $a_k = [a_k(1); \tilde{a}_k]$ and upper bound the target quantity as:

$$\begin{aligned} & \left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 \begin{bmatrix} |a_k(1)|^2 & a_k(1) \tilde{a}_k^* \\ a_k(1) \tilde{a}_k & \tilde{a}_k \tilde{a}_k^* \end{bmatrix} - (e_1 e_1^* + \mathbf{I}) \right\| \\ & \leq \left| \frac{1}{m} \sum_{k=1}^m (|a_k(1)|^4 - 2) \right| + \left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 \begin{bmatrix} 0 & a_k(1) \tilde{a}_k^* \\ a_k(1) \tilde{a}_k & \mathbf{0} \end{bmatrix} \right\| \\ & \quad + \left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (\tilde{a}_k \tilde{a}_k^* - \mathbf{I}_{n-1}) \right\| + \left| \frac{1}{m} \sum_{k=1}^m (|a_k(1)|^2 - 1) \right|. \end{aligned}$$

By Chebyshev's inequality, we have with probability at least $1 - c_1 \delta^{-2} m^{-1}$,

$$\left| \frac{1}{m} \sum_{k=1}^m (|a_k(1)|^4 - 2) \right| \leq \frac{\delta}{4} \quad \text{and} \quad \left| \frac{1}{m} \sum_{k=1}^m (|a_k(1)|^2 - 1) \right| \leq \frac{\delta}{4}.$$

To bound the second term, we note that

$$\begin{aligned} \left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 \begin{bmatrix} 0 & a_k(1) \tilde{a}_k^* \\ a_k(1) \tilde{a}_k & \mathbf{0} \end{bmatrix} \right\| &= \left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 a_k(1) \tilde{a}_k^* \right\| \\ &= \sup_{\mathbf{w} \in \mathbb{C}^{n-1}: \|\mathbf{w}\|=1} \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 a_k(1) \tilde{a}_k^* \mathbf{w}. \end{aligned}$$

For all \mathbf{w} and all $k \in [m]$, $\tilde{a}_k^* \mathbf{w}$ is distributed as $\mathcal{CN}(1)$ that is independent of the $\{a_k(1)\}$ sequence. So for one realization of $\{a_k(1)\}$, the Hoeffding-type inequality of Lemma A.5 implies

$$\mathbb{P} \left[\frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 a_k(1) \tilde{a}_k^* \mathbf{w} > t \right] \leq e \exp \left(- \frac{c_2 m^2 t^2}{\sum_{k=1}^m |a_k(1)|^6} \right),$$

for any \mathbf{w} with $\|\mathbf{w}\| = 1$ and any $t > 0$. Taking $t = \delta/8$, together with a union bound on a $1/2$ -net on the sphere, we obtain

$$\mathbb{P} \left[\left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 a_k(1) \tilde{a}_k^* \right\| > \delta/4 \right] \leq e \exp \left(- \frac{c_2 m^2 \delta^2}{64 \sum_{k=1}^m |a_k(1)|^6} + 12(n-1) \right).$$

Now an application of Chebyshev's inequality gives that $\sum_{k=1}^m |a_k(1)|^6 \leq 20m$ with probability at least $1 - c_3 m^{-1}$. Substituting this into the above, we conclude that whenever $m \geq C_4 \delta^{-2} n$ for some sufficiently large C_4 ,

$$\left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 a_k(1) \tilde{a}_k^* \right\| \leq \delta/4$$

with probability at least $1 - c_3 m^{-1} - \exp(-c_5 \delta^2 m)$.

To bound the third term, we note that

$$\left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (\tilde{\mathbf{a}}_k \tilde{\mathbf{a}}_k^* - \mathbf{I}_{n-1}) \right\| = \sup_{\mathbf{w} \in \mathbb{C}^{n-1}: \|\mathbf{w}\|=1} \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (|\tilde{\mathbf{a}}_k^* \mathbf{w}|^2 - 1).$$

For all fixed \mathbf{w} and all $k \in [m]$, $\tilde{\mathbf{a}}_k^* \mathbf{w} \sim \mathcal{CN}(1)$. Thus, $|\tilde{\mathbf{a}}_k^* \mathbf{w}|^2 - 1$ is centered sub-exponential. So for one realization of $\{a_k(1)\}$, Bernstein's inequality (Lemma A.6) implies

$$\mathbb{P} \left[\frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (|\tilde{\mathbf{a}}_k^* \mathbf{w}|^2 - 1) > t \right] \leq 2 \exp \left(-c_6 \min \left(\frac{t^2}{c_7^2 \sum_{k=1}^m |a_k(1)|^4}, \frac{t}{c_7 \max_{i \in [m]} |a_k(1)|^2} \right) \right)$$

for any fixed \mathbf{w} with $\|\mathbf{w}\| = 1$ and any $t > 0$. Taking $t = \delta/8$, together with a union bound on a $1/2$ -net on the sphere, we obtain

$$\begin{aligned} \mathbb{P} \left[\left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (\tilde{\mathbf{a}}_k \tilde{\mathbf{a}}_k^* - \mathbf{I}_{n-1}) \right\| > \frac{\delta}{4} \right] \\ \leq 2 \exp \left(-c_6 \min \left(\frac{m^2 \delta^2 / 64}{c_7^2 \sum_{k=1}^m |a_k(1)|^4}, \frac{m \delta / 8}{c_7 \max_{i \in [m]} |a_k(1)|^2} \right) + 12(n-1) \right). \end{aligned}$$

Chebyshev's inequality and the union bound give that

$$\sum_{k=1}^m |a_k(1)|^4 \leq 10m, \quad \text{and} \quad \max_{i \in [m]} |a_k(1)|^2 \leq 10 \log m$$

hold with probability at least $1 - c_8 m^{-1} - m^{-4}$. To conclude, when $m \geq C_9(\delta) \delta^{-2} n \log n$ for some sufficiently large constant $C_9(\delta)$,

$$\left\| \frac{1}{m} \sum_{k=1}^m |a_k(1)|^2 (\tilde{\mathbf{a}}_k \tilde{\mathbf{a}}_k^* - \mathbf{I}_{n-1}) \right\| \leq \frac{\delta}{4}$$

with probability at least $1 - c_8 m^{-1} - m^{-4} - 2 \exp(-c_{10} \delta^2 m / \log m)$.

Collecting the above bounds and probabilities yields the claimed results. Similar arguments prove the claim on $\frac{1}{m} \sum_{k=1}^m (\mathbf{a}_k^* \mathbf{v}) \mathbf{a}_k \mathbf{a}_k^\top$ also, completing the proof. \blacksquare

Lemma 6.4 *Let $\mathbf{a}_1, \dots, \mathbf{a}_m$ be i.i.d. copies of $\mathbf{a} \sim \mathcal{CN}(n)$. For any $\delta \in (0, 1)$, when $m \geq C(\delta) n \log n$, it holds with probability at least $1 - c' \exp(-c(\delta)m) - c'' m^{-n}$ that*

$$\begin{aligned} \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{w}|^2 &\geq (1 - \delta) \left(\|\mathbf{w}\|^2 \|\mathbf{z}\|^2 + |\mathbf{w}^* \mathbf{z}|^2 \right) \quad \text{for all } \mathbf{z}, \mathbf{w} \in \mathbb{C}^n, \\ \frac{1}{m} \sum_{k=1}^m [\Re(\mathbf{a}_k^* \mathbf{z})(\mathbf{w}^* \mathbf{a}_k)]^2 &\geq (1 - \delta) \left(\frac{1}{2} \|\mathbf{z}\|^2 \|\mathbf{w}\|^2 + \frac{3}{2} [\Re \mathbf{z}^* \mathbf{w}]^2 - \frac{1}{2} [\Im \mathbf{z}^* \mathbf{w}]^2 \right) \quad \text{for all } \mathbf{z}, \mathbf{w} \in \mathbb{C}^n. \end{aligned}$$

Here $C(\delta)$ and $c(\delta)$ are constants depending on δ and c' and c'' are positive absolute constants.

Proof By Lemma 6.2, $\mathbb{E} \left[|\mathbf{a}^* \mathbf{w}|^2 |\mathbf{a}^* \mathbf{z}|^2 \right] = \|\mathbf{w}\|^2 \|\mathbf{z}\|^2 + |\mathbf{w}^* \mathbf{z}|^2$. By homogeneity, it is enough to prove the result for all $\mathbf{w}, \mathbf{z} \in \mathbb{CS}^{n-1}$. For a pair of fixed $\mathbf{w}, \mathbf{z} \in \mathbb{CS}^{n-1}$, Lemma A.7 implies that for any $\delta \in (0, 1)$,

$$\sum_{k=1}^m |\mathbf{a}_k^* \mathbf{w}|^2 |\mathbf{a}_k^* \mathbf{z}|^2 \geq \left(1 - \frac{\delta}{2}\right) m \left(1 + |\mathbf{w}^* \mathbf{z}|^2\right)$$

with probability at least $1 - \exp(-c_1 \delta^2 m)$. For a certain $\varepsilon \in (0, 1)$ to be fixed later and an ε -net $N_\varepsilon^1 \times N_\varepsilon^2$ that covers $\mathbb{CS}^{n-1} \times \mathbb{CS}^{n-1}$, we have that the event

$$\mathcal{E}_0 \doteq \left\{ \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{w}|^2 |\mathbf{a}_k^* \mathbf{z}|^2 \geq \left(1 - \frac{\delta}{2}\right) m \left(1 + |\mathbf{w}^* \mathbf{z}|^2\right) \quad \forall \mathbf{w}, \mathbf{z} \in N_\varepsilon^1 \times N_\varepsilon^2 \right\}$$

holds with probability at least $1 - \exp(-c_1 \delta^2 m + 4n \log(3/\varepsilon))$ by a simple union bound. Now conditioned on \mathcal{E}_0 , we have for every $\mathbf{z} \in \mathbb{CS}^{n-1}$ can be written as $\mathbf{z} = \mathbf{z}_0 + \mathbf{e}$ for certain $\mathbf{z}_0 \in N_\varepsilon^1$ and \mathbf{e} with $\|\mathbf{e}\| \leq \varepsilon$; similarly $\mathbf{w} = \mathbf{w}_0 + \boldsymbol{\zeta}$ for $\mathbf{w}_0 \in N_\varepsilon^2$ and $\|\boldsymbol{\zeta}\| \leq \varepsilon$. For the function $g(\mathbf{w}, \mathbf{z}) \doteq \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{w}|^2$, with high probability,

$$\begin{aligned} \left\| \frac{\partial g}{\partial \mathbf{w}} \right\| &= \left\| \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \mathbf{w}^* \mathbf{a}_k \bar{\mathbf{a}}_k \right\| \leq \|\mathbf{z}\|^2 \|\mathbf{w}\| \left\| \sum_{k=1}^m \|\mathbf{a}_k\|^2 \mathbf{a}_k \mathbf{a}_k^* \right\| \leq 10mn \sqrt{\log m}, \\ \left\| \frac{\partial g}{\partial \mathbf{z}} \right\| &= \left\| \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{w}|^2 \mathbf{z}^* \mathbf{a}_k \bar{\mathbf{a}}_k \right\| \leq \|\mathbf{w}\|^2 \|\mathbf{z}\| \left\| \sum_{k=1}^m \|\mathbf{a}_k\|^2 \mathbf{a}_k \mathbf{a}_k^* \right\| \leq 10mn \sqrt{\log m}, \end{aligned}$$

as $\max_{k \in [m]} \|\mathbf{a}_k\|^2 \leq 5n \log m$ with probability at least $1 - c_2 m^{-n}$, and $\|\sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^*\| \leq 2m$ with probability at least $1 - \exp(-c_3 m)$. Thus,

$$\sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{w}|^2 \geq \left(1 - \frac{\delta}{3}\right) m - 40\varepsilon mn \log m + \left(1 - \frac{\delta}{3}\right) m \left(|\mathbf{w}_0^* \mathbf{z}_0|^2 - 4\varepsilon\right).$$

Taking $\varepsilon = c_4(\delta)/(n \log m)$ for a sufficiently small $c_4(\delta) > 0$, we obtain that with probability at least $1 - \exp(-c_1 \delta^2 m + 4n \log(3n \log m / c_4(\delta))) - c_5 m^{-n}$,

$$\sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{w}|^2 \geq \left(1 - \frac{2}{3}\delta\right) m \left(1 + |\mathbf{w}_0^* \mathbf{z}_0|^2\right).$$

which, together with continuity of the function $(\mathbf{w}, \mathbf{z}) \mapsto |\mathbf{w}^* \mathbf{z}|^2$, implies

$$\sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 |\mathbf{a}_k^* \mathbf{w}|^2 \geq (1 - \delta) m \left(1 + |\mathbf{w}^* \mathbf{z}|^2\right).$$

It is enough to take $m \geq C_6 \delta^{-2} n \log n$ to ensure the desired event happens with high probability.

To show the second inequality, first notice that $\mathbb{E} [\Re(\mathbf{a}_k^* \mathbf{z})(\mathbf{w}^* \mathbf{a}_k)]^2 = \frac{1}{2} \|\mathbf{z}\|^2 \|\mathbf{w}\|^2 + \frac{3}{2} [\Re \mathbf{z}^* \mathbf{w}]^2 - \frac{1}{2} [\Im \mathbf{z}^* \mathbf{w}]^2$. The argument then proceeds to apply the discretization trick as above. \blacksquare

6.2 Proof of Proposition 2.3

Proof Direct calculation shows that

$$\begin{aligned}
& \left[\frac{\mathbf{x}e^{i\phi(z)}}{\bar{\mathbf{x}}e^{-i\phi(z)}} \right]^* \nabla^2 f(z) \left[\frac{\mathbf{x}e^{i\phi(z)}}{\bar{\mathbf{x}}e^{-i\phi(z)}} \right] \\
&= \frac{1}{m} \sum_{k=1}^m \left(4 |\mathbf{a}_k^* z|^2 |\mathbf{a}_k^* \mathbf{x}|^2 - 2 |\mathbf{a}_k^* \mathbf{x}|^4 + 2\Re \left[(\mathbf{a}_k^* z)^2 (\mathbf{x}^* \mathbf{a}_k)^2 e^{-2i\phi(z)} \right] \right) \\
&= \frac{1}{m} \sum_{k=1}^m \left(2 |\mathbf{a}_k^* z|^2 |\mathbf{a}_k^* \mathbf{x}|^2 - 2 |\mathbf{a}_k^* \mathbf{x}|^4 \right) \\
&\quad + \frac{1}{m} \sum_{k=1}^m \left(2 |\mathbf{a}_k^* z|^2 |\mathbf{a}_k^* \mathbf{x}|^2 + 2\Re \left[(\mathbf{a}_k^* z)^2 (\mathbf{x}^* \mathbf{a}_k)^2 e^{-2i\phi(z)} \right] \right).
\end{aligned}$$

Lemma 6.3 implies that when $m \geq C_1 n \log n$, with high probability,

$$\frac{2}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 |\mathbf{a}_k^* z|^2 \leq 2 |\mathbf{x}^* z|^2 + \frac{401}{200} \|\mathbf{x}\|^2 \|z\|^2.$$

On the other hand, by Lemma A.7, we have that

$$\frac{2}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^4 \geq \frac{399}{100} \|\mathbf{x}\|^4$$

holds with probability at least $1 - \exp(-c_2 m)$. For the second summation, we have

$$\begin{aligned}
& \frac{1}{m} \sum_{k=1}^m \left(2 |\mathbf{a}_k^* z|^2 |\mathbf{a}_k^* \mathbf{x}|^2 + 2\Re \left[(\mathbf{a}_k^* z)^2 (\mathbf{x}^* \mathbf{a}_k)^2 e^{-2i\phi(z)} \right] \right) \\
&= \left[\frac{z}{\bar{z}} \right]^* \nabla^2 f(\mathbf{x}e^{i\phi(z)}) \left[\frac{z}{\bar{z}} \right] \leq \left[\frac{z}{\bar{z}} \right]^* \nabla^2 \mathbb{E} \left[f(\mathbf{x}e^{i\phi(z)}) \right] \left[\frac{z}{\bar{z}} \right] + \frac{1}{200} \|\mathbf{x}\|^2 \|z\|^2 \leq 6 |\mathbf{x}^* z|^2 + \frac{401}{200} \|\mathbf{x}\|^2 \|z\|^2,
\end{aligned}$$

with high probability, provided $m \geq C_3 n \log n$, according to Lemma 6.3.

Collecting the above estimates, we have that when $m \geq C_4 n \log n$ for a sufficiently large constant C_4 , with high probability,

$$\left[\frac{\mathbf{x}e^{i\phi(z)}}{\bar{\mathbf{x}}e^{-i\phi(z)}} \right]^* \nabla^2 f(z) \left[\frac{\mathbf{x}e^{i\phi(z)}}{\bar{\mathbf{x}}e^{-i\phi(z)}} \right] \leq \frac{401}{100} \|\mathbf{x}\|^2 \|z\|^2 + 8 |\mathbf{x}^* z|^2 - \frac{399}{100} \|\mathbf{x}\|^4 \leq -\frac{1}{100} \|\mathbf{x}\|^4$$

for all $z \in \mathcal{R}_1$. Dividing both sides of the above by $\|\mathbf{x}\|^2$ gives the claimed results. \blacksquare

6.3 Proof of Proposition 2.4

Proof Note that

$$z^* \nabla_z f(z) = \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* z|^4 - \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 |\mathbf{a}_k^* z|^2.$$

By Lemma 6.4, when $m \geq C_1 n \log n$ for some sufficiently large C_1 , with high probability,

$$\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^4 \geq \frac{199}{100} \|\mathbf{z}\|^4$$

for all $\mathbf{z} \in \mathbb{C}^n$. On the other hand, Lemma 6.3 implies that when $m \geq C_2 n \log n$ for some sufficiently large C_2 , with high probability,

$$\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 |\mathbf{a}_k^* \mathbf{z}|^2 \leq |\mathbf{x}^* \mathbf{z}|^2 + \frac{1001}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2.$$

for all $\mathbf{z} \in \mathbb{C}^n$. Combining the above estimates, we have that when $m \geq \max(C_1, C_2) n \log n$, with high probability,

$$\mathbf{z}^* \nabla_{\mathbf{z}} f(\mathbf{z}) \geq \frac{199}{100} \|\mathbf{z}\|^4 - \frac{1001}{1000} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 - |\mathbf{x}^* \mathbf{z}|^2 \geq \frac{1}{500} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2$$

for all $\mathbf{z} \in \mathcal{R}_2^z$, as desired. ■

6.4 Proof of Proposition 2.5

Proof We abbreviate $\phi(\mathbf{z})$ as ϕ below. Note that

$$(\mathbf{z} - \mathbf{x} e^{i\phi})^* \nabla_{\mathbf{z}} f(\mathbf{z}) = \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 (\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} - \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 (\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z}.$$

We first bound the second term. By Lemma 6.3, when $m \geq C_1 n \log n$ for a sufficiently large constant C_1 , with high probability, for all $\mathbf{z} \in \mathbb{C}^n$,

$$\begin{aligned} & \Re \left(\frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 (\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right) \\ &= \Re \left((\mathbf{z} - \mathbf{x} e^{i\phi})^* (\|\mathbf{x}\|^2 \mathbf{I} + \mathbf{x} \mathbf{x}^*) \mathbf{z} \right) + \Re \left((\mathbf{z} - \mathbf{x} e^{i\phi})^* \Delta \mathbf{z} \right) \quad (\text{for a certain } \Delta \text{ with } \|\Delta\| \leq \|\mathbf{x}\|^2 / 100) \\ &\leq \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 + |\mathbf{x}^* \mathbf{z}|^2 - 2 \|\mathbf{x}\|^2 |\mathbf{x}^* \mathbf{z}| + \frac{1}{1000} \|\mathbf{x}\|^2 \|\mathbf{z} - \mathbf{x} e^{i\phi}\| \|\mathbf{z}\|. \end{aligned}$$

To bound the first term, for a fixed τ to be determined later, define:

$$\begin{aligned} S(\mathbf{z}) &\doteq \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \Re \left((\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right), \\ S_1(\mathbf{z}) &\doteq \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \Re \left((\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right) \mathbb{1}_{|\mathbf{a}_k^* \mathbf{x}| \leq \tau} \\ S_2(\mathbf{z}) &\doteq \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \Re \left((\mathbf{z} - \mathbf{x} e^{i\phi})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right) \mathbb{1}_{|\mathbf{a}_k^* \mathbf{x}| \leq \tau} \mathbb{1}_{|\mathbf{a}_k^* \mathbf{z}| \leq \tau}. \end{aligned}$$

Obviously $S_1(\mathbf{z}) \geq S_2(\mathbf{z})$ for all \mathbf{z} as

$$\begin{aligned} S_1(\mathbf{z}) - S_2(\mathbf{z}) &= \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \Re \left((\mathbf{z} - \mathbf{x} e^{i\theta})^* \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right) \mathbb{1}_{|\mathbf{a}_k^* \mathbf{x}| \leq \tau} \mathbb{1}_{|\mathbf{a}_k^* \mathbf{z}| > \tau} \\ &\geq \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^2 \left(|\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{x}| |\mathbf{a}_k^* \mathbf{z}| \right) \mathbb{1}_{|\mathbf{a}_k^* \mathbf{x}| \leq \tau} \mathbb{1}_{|\mathbf{a}_k^* \mathbf{z}| > \tau} \geq 0. \end{aligned}$$

Now for an $\varepsilon \in (0, \|\mathbf{x}\|)$ to be fixed later, consider an ε -net N_ε for the ball $\mathbb{CB}^n(\|\mathbf{x}\|)$, with $|N_\varepsilon| \leq (3\|\mathbf{x}\|/\varepsilon)^{2n}$. On the complement of the event $\{\max_{k \in [m]} |\mathbf{a}_k^* \mathbf{x}| > \tau\}$, we have for any $t > 0$ that

$$\begin{aligned} &\mathbb{P}[S(\mathbf{z}) - \mathbb{E}[S(\mathbf{z})] < -t, \forall \mathbf{z} \in N_\varepsilon] \\ &\leq |N_\varepsilon| \mathbb{P}[S(\mathbf{z}) - \mathbb{E}[S(\mathbf{z})] < -t] \\ &\leq |N_\varepsilon| \mathbb{P}[S_1(\mathbf{z}) - \mathbb{E}[S_1(\mathbf{z})] < -t + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S(\mathbf{z})]|]. \end{aligned}$$

Because $S_1(\mathbf{z}) \geq S_2(\mathbf{z})$ as shown above,

$$\begin{aligned} &\mathbb{P}[S_1(\mathbf{z}) - \mathbb{E}[S_1(\mathbf{z})] < -t + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S(\mathbf{z})]|] \\ &\leq \mathbb{P}[S_2(\mathbf{z}) - \mathbb{E}[S_2(\mathbf{z})] < -t + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S(\mathbf{z})]| + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S_2(\mathbf{z})]|]. \end{aligned}$$

Thus, the unconditional probability can be bounded as

$$\begin{aligned} &\mathbb{P}[S(\mathbf{z}) - \mathbb{E}[S(\mathbf{z})] < -t, \forall \mathbf{z} \in N_\varepsilon] \\ &\leq |N_\varepsilon| \mathbb{P}[S_2(\mathbf{z}) - \mathbb{E}[S_2(\mathbf{z})] < -t + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S(\mathbf{z})]| + |\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S_2(\mathbf{z})]|] \\ &\quad + \mathbb{P}\left[\max_{k \in [m]} |\mathbf{a}_k^* \mathbf{x}| > \tau\right]. \end{aligned}$$

Taking $\tau = \sqrt{10 \log m} \|\mathbf{x}\|$, we obtain

$$\begin{aligned} &\mathbb{P}\left[\max_{k \in [m]} |\mathbf{a}_k^* \mathbf{x}| > \tau\right] \leq m \exp\left(-\frac{10 \log m}{2}\right) \leq m^{-4}, \\ &|\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S(\mathbf{z})]| \leq \sqrt{\mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^6 |\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})|^2\right]} \sqrt{\mathbb{P}[|\mathbf{a}^* \mathbf{x}| > \tau]} \leq 4\sqrt{3} m^{-5/2} \|\mathbf{z}\|^3 \|\mathbf{z} - \mathbf{x} e^{i\phi}\|, \\ &|\mathbb{E}[S_1(\mathbf{z})] - \mathbb{E}[S_2(\mathbf{z})]| \leq \sqrt{\mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^6 |\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})|^2 \mathbb{1}_{|\mathbf{a}^* \mathbf{x}| \leq \tau}\right]} \sqrt{\mathbb{P}[|\mathbf{a}^* \mathbf{z}| > \tau]} \\ &\leq 4\sqrt{3} m^{-5/2} \|\mathbf{z}\|^3 \|\mathbf{z} - \mathbf{x} e^{i\phi}\|, \end{aligned}$$

where we have used $\|\mathbf{z}\| \leq \|\mathbf{x}\|$ to simplify the last inequality. Now we used the moment-control Bernstein's inequality (Lemma A.8) to get a bound for probability on deviation of $S_2(\mathbf{z})$. To this end, we have

$$\begin{aligned} &\mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^6 \left|\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})\right|^2 \mathbb{1}_{|\mathbf{a}^* \mathbf{x}| \leq \tau} \mathbb{1}_{|\mathbf{a}^* \mathbf{z}| \leq \tau}\right] \leq \tau^2 \mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^4 \left|\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})\right|^2\right] \\ &\leq 240 \log m \|\mathbf{x}\|^2 \|\mathbf{z}\|^4 \|\mathbf{z} - \mathbf{x} e^{i\phi}\|^2 \\ &\mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^{3p} \left|\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})\right|^p \mathbb{1}_{|\mathbf{a}^* \mathbf{x}| \leq \tau} \mathbb{1}_{|\mathbf{a}^* \mathbf{z}| \leq \tau}\right] \leq \tau^{2p} \mathbb{E}\left[|\mathbf{a}^* \mathbf{z}|^p \left|\mathbf{a}^*(\mathbf{z} - \mathbf{x} e^{i\phi})\right|^p\right] \end{aligned}$$

$$\leq \left(10 \log m \|\mathbf{x}\|^2\right)^p p! \|z\|^p \left\|z - \mathbf{x}e^{i\phi}\right\|^p,$$

where the second inequality holds for any integer $p \geq 3$. Hence one can take

$$\begin{aligned}\sigma^2 &= 240 \log^2 m \|\mathbf{x}\|^4 \|z\|^2 \left\|z - \mathbf{x}e^{i\phi}\right\|^2, \\ R &= 10 \log m \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\|\end{aligned}$$

in Lemma A.8, and

$$t = \frac{1}{1000} \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\|$$

in the deviation inequality of $S_2(z)$ to obtain

$$\mathbb{P} \left[S_2(z) - \mathbb{E}[S_2(z)] < -\frac{1}{200} \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\| \right] \leq \exp \left(-\frac{c_2 m}{\log^2 m} \right),$$

where we have used the fact $\|z\| \leq \|\mathbf{x}\|$ and assumed $4\sqrt{3}m^{-5/2} \leq 1/200$ to simplify the probability. Thus, with probability at least $1 - m^{-4} - \exp(-c_2 m / \log^2 m + 2n \log(3 \|\mathbf{x}\| / \varepsilon))$, it holds that

$$S(z) \geq 2 \|z\|^4 - 2 \|z\|^2 |\mathbf{x}^* z| - \frac{1}{1000} \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\| \quad \forall z \in N_\varepsilon. \quad (6.4)$$

Moreover, for any $z, z' \in \mathcal{R}_2^h$, we have

$$\begin{aligned}& |S(z) - S(z')| \\ & \leq \frac{1}{m} \sum_{k=1}^m \left| |\mathbf{a}_k^* z|^2 - |\mathbf{a}_k^* z'|^2 \right| |\mathbf{h}^*(z) \mathbf{a}_k \mathbf{a}_k^* z| + \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* z'|^2 |\mathbf{h}^*(z) \mathbf{a}_k \mathbf{a}_k^* z - \mathbf{h}^*(z') \mathbf{a}_k \mathbf{a}_k^* z'| \\ & \leq 4 \max_{k \in [m]} \|\mathbf{a}_k\|^4 \|\mathbf{x}\|^3 \|z - z'\| + 5 \max_{k \in [m]} \|\mathbf{a}_k\|^4 \|\mathbf{x}\|^3 \|z - z'\| \\ & \leq 90n^2 \log^2 m \|\mathbf{x}\|^3 \|z - z'\|,\end{aligned}$$

as $\max_{k \in [m]} \|\mathbf{a}_k\|^4 \leq 10n^2 \log^2 m$ with probability at least $1 - c_3 m^{-n}$, and $11 \|\mathbf{x}\| / 20 \leq \|z\| \leq \|\mathbf{x}\|$, and also $\|\mathbf{x}e^{i\phi(z)} - \mathbf{x}e^{i\phi(z')}\| \leq 2 \|z - z'\|$ for $z, z' \in \mathcal{R}_2^h$. Every $z \in \mathcal{R}_2^h$ can be written as $z = z' + e$, with $z' \in N_\varepsilon$ and $\|e\| \leq \varepsilon$. Thus,

$$\begin{aligned}S(z) &\geq S(z') - 90n^2 \log^2 m \|\mathbf{x}\|^3 \varepsilon \\ &\geq 2 \|z'\|^4 - 2 \|z'\|^2 |\mathbf{x}^* z'| - \frac{1}{1000} \|\mathbf{x}\|^2 \|z'\| \left\|z' - \mathbf{x}e^{i\phi}\right\| - 90n^2 \log^2 m \|\mathbf{x}\|^3 \varepsilon \\ &\geq 2 \|z\|^4 - 2 \|z\|^2 |\mathbf{x}^* z| - \frac{1}{1000} \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\| - 11\varepsilon \|\mathbf{x}\|^3 - 90n^2 \log^2 m \|\mathbf{x}\|^3 \varepsilon,\end{aligned}$$

where the additional $11\varepsilon \|\mathbf{x}\|^3$ term in the third line is to account for the change from z' to z , which has been simplified by assumptions that $11/20 \cdot \|\mathbf{x}\| \leq \|z\| \leq \|\mathbf{x}\|$ and that $\varepsilon \leq \|\mathbf{x}\|$. Choosing $\varepsilon = \|\mathbf{x}\| / (c_5 n^2 \log^2 m)$ for a sufficiently large $c_5 > 0$ and additionally using $\text{dist}(z, \mathcal{X}) \geq \|\mathbf{x}\| / 3$, we obtain that

$$S(z) \geq 2 \|z\|^4 - 2 \|z\|^2 \Re \left(\mathbf{x}^* z e^{-i\phi} \right) - \frac{1}{500} \|\mathbf{x}\|^2 \|z\| \left\|z - \mathbf{x}e^{i\phi}\right\| \quad (6.5)$$

for all $z \in \mathcal{R}_2^h$, with probability at least $1 - c_6 m^{-1} - c_7 \exp(-c_2 m / \log^2 m + c_9 n \log(C_8 n \log m))$.

Combining the above estimates, when $m \geq C_{10} n \log^3 n$ for sufficiently large constant C_{10} , with high probability,

$$\Re \left((z - x e^{i\phi})^* \nabla_z f(z) \right) \geq \frac{1}{1000} \|x\|^2 \|z\| \left\| z - x e^{i\phi} \right\|$$

for all $z \in \mathcal{R}_2^h$, as claimed. \blacksquare

6.5 Proof of Proposition 2.6

Proof It is enough to prove that for all unit vectors g that are geometrically orthogonal to ix , i.e., $g \in \mathcal{T} \doteq \{z : \Im(z^* x) = 0, \|z\| = 1\}$ and all $t \in [0, \|x\| / \sqrt{7}]$, the following holds:

$$\begin{bmatrix} g \\ \bar{g} \end{bmatrix}^* \nabla^2 f(x + tg) \begin{bmatrix} g \\ \bar{g} \end{bmatrix} \geq \frac{1}{4} \|x\|^2.$$

Direct calculation shows

$$\begin{aligned} & \begin{bmatrix} g \\ \bar{g} \end{bmatrix}^* \nabla^2 f(x + tg) \begin{bmatrix} g \\ \bar{g} \end{bmatrix} \\ &= \frac{1}{m} \sum_{k=1}^m 4 |\mathbf{a}_k^*(x + tg)|^2 |\mathbf{a}_k^* g|^2 - 2 |\mathbf{a}_k^* x|^2 |\mathbf{a}_k^* g|^2 + 2 \Re [(t \mathbf{a}_k^* g + \mathbf{a}_k^* x)^2 (g^* \mathbf{a}_k)^2] \\ &\geq \frac{1}{m} \sum_{k=1}^m 4 |\mathbf{a}_k^*(x + tg)|^2 |\mathbf{a}_k^* g|^2 - 2 |\mathbf{a}_k^* x|^2 |\mathbf{a}_k^* g|^2 + 4 [\Re(t \mathbf{a}_k^* g + \mathbf{a}_k^* x)(g^* \mathbf{a}_k)]^2 - 2 |(t \mathbf{a}_k^* g + \mathbf{a}_k^* x)(g^* \mathbf{a}_k)|^2 \\ &\geq \frac{1}{m} \sum_{k=1}^m 2 |\mathbf{a}_k^*(x + tg)|^2 |\mathbf{a}_k^* g|^2 - 2 |\mathbf{a}_k^* x|^2 |\mathbf{a}_k^* g|^2 + 4 [\Re(t \mathbf{a}_k^* g + \mathbf{a}_k^* x)(g^* \mathbf{a}_k)]^2. \end{aligned}$$

Lemma 6.4 implies when $m \geq C_1 n \log n$ for sufficiently large constant C_1 , with high probability,

$$\frac{1}{m} \sum_{k=1}^m 2 |\mathbf{a}_k^*(x + tg)|^2 |\mathbf{a}_k^* g|^2 \geq \frac{199}{100} |(x + tg)^* g|^2 + \frac{199}{100} \|x + tg\|^2 \|g\|^2 \quad (6.6)$$

for all $g \in \mathbb{C}^n$ and all $t \in [0, \|x\| / \sqrt{7}]$. Lemma 6.3 implies that when $m \geq C_2 n \log n$ for sufficiently large constant C_2 , with high probability,

$$\frac{1}{m} \sum_{k=1}^m 2 |\mathbf{a}_k^* x|^2 |\mathbf{a}_k^* g|^2 \leq \frac{201}{100} |x^* g|^2 + \frac{201}{100} \|x\|^2 \|g\|^2 \quad (6.7)$$

for all $g \in \mathbb{C}^n$. Moreover, Lemma 6.4 implies when $m \geq C_3 n \log n$ for sufficiently large constant C_3 , with high probability,

$$\frac{4}{m} \sum_{k=1}^m [\Re(t \mathbf{a}_k^* g + \mathbf{a}_k^* x)(g^* \mathbf{a}_k)]^2 \geq 2 \|x + tg\|^2 \|g\|^2 + 6 |(x + g)^* g|^2 - \frac{1}{400} \|x\|^2 \|g\|^2$$

for all $g \in \mathcal{T}$, where we have used that $\Im(g^* x) = 0 \implies \Im(x + g)^* g = 0$ to simplify the results.

Collecting the above estimates, we obtain that when $m \geq C_4 n \log n$, with high probability,

$$\begin{aligned} & \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla^2 f(\mathbf{x} + t\mathbf{g}) \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right] \\ & \geq \left(\frac{399}{100} \|\mathbf{x} + t\mathbf{g}\|^2 - \frac{161}{80} \|\mathbf{x}\|^2 \right) + \left(\frac{799}{100} |(\mathbf{x} + t\mathbf{g})^* \mathbf{g}|^2 - \frac{201}{100} |\mathbf{x}^* \mathbf{g}|^2 \right) \\ & = \frac{791}{400} \|\mathbf{x}\|^2 + \frac{598}{100} |\mathbf{x}^* \mathbf{g}|^2 + \frac{1198}{100} t^2 + \frac{2396}{100} t \Re(\mathbf{x}^* \mathbf{g}). \end{aligned}$$

To provide a lower bound for the above, we let $\Re(\mathbf{x}^* \mathbf{g}) = \mathbf{x}^* \mathbf{g} = \lambda \|\mathbf{x}\|$ with $\lambda \in [-1, 1]$ and $t = \eta \|\mathbf{x}\|$ with $\eta \in [0, 1/\sqrt{7}]$. Then

$$\frac{598}{100} |\mathbf{x}^* \mathbf{g}|^2 + \frac{1198}{100} t^2 + \frac{2396}{100} t \Re(\mathbf{x}^* \mathbf{g}) = \|\mathbf{x}\|^2 \left(\frac{598}{100} \lambda^2 + \frac{1198}{100} \eta^2 + \frac{2396}{100} \lambda \eta \right) \doteq \|\mathbf{x}\|^2 \phi(\lambda, \eta).$$

For any fixed η , it is easy to see that minimizer occurs when $\lambda = -\frac{599}{299} \eta$. Plugging this into $\phi(\lambda, \eta)$, one obtains $\phi(\lambda, \eta) \geq -\frac{241}{20} \eta^2 \geq -\frac{241}{140}$. Thus,

$$\left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla^2 f(\mathbf{x} + t\mathbf{g}) \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right] \geq \left(\frac{791}{400} - \frac{241}{140} \right) \|\mathbf{x}\|^2 \geq \frac{1}{4} \|\mathbf{x}\|^2,$$

as claimed. ■

6.6 Proof of Proposition 2.7

Proof For convenience, we will define a relaxed \mathcal{R}_2^h region

$$\mathcal{R}_2^{h'} \doteq \left\{ \mathbf{z} : \Re(\langle \mathbf{h}(\mathbf{z}), \nabla_{\mathbf{z}} \mathbb{E}[f] \rangle) \geq \frac{1}{250} \|\mathbf{x}\|^2 \|\mathbf{z}\| \|\mathbf{h}(\mathbf{z})\|, \|\mathbf{z}\| \leq \|\mathbf{x}\| \right\} \supset \mathcal{R}_2^h$$

and try to show that $\mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^{h'} \cup \mathcal{R}_3 = \mathbb{C}^n$. In the end, we will discuss how this implies the claimed result.

We will first divide \mathbb{C}^n into several (overlapping) regions, and show that each such region is a subset of $\mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^{h'} \cup \mathcal{R}_3$.

Cover $\mathcal{R}_a \doteq \{ \mathbf{z} : |\mathbf{x}^* \mathbf{z}| \leq \frac{1}{2} \|\mathbf{x}\| \|\mathbf{z}\| \}$: In this case, when $\|\mathbf{z}\|^2 \leq \frac{398}{601} \|\mathbf{x}\|^2$,

$$8 |\mathbf{x}^* \mathbf{z}|^2 + \frac{401}{100} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \leq \frac{601}{100} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \leq \frac{398}{100} \|\mathbf{x}\|^4.$$

On the other hand, when $\|\mathbf{z}\|^2 \geq \frac{626}{995} \|\mathbf{x}\|^2$,

$$\frac{501}{500} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 + |\mathbf{x}^* \mathbf{z}|^2 \leq \frac{313}{250} \|\mathbf{x}\|^2 \|\mathbf{z}\|^2 \leq \frac{199}{100} \|\mathbf{z}\|^4.$$

Since $\frac{398}{601} > \frac{626}{995}$, we conclude that $\mathcal{R}_a \subset \mathcal{R}_1 \cup \mathcal{R}_2^z$.

Cover $\mathcal{R}_b \doteq \{z : |\mathbf{x}^*z| \geq \frac{1}{2} \|\mathbf{x}\| \|z\|, \|z\| \leq \frac{57}{100} \|\mathbf{x}\|\}$: In this case,

$$8 |\mathbf{x}^*z|^2 + \frac{401}{100} \|\mathbf{x}\|^2 \|z\|^2 \leq \frac{1201}{100} \|\mathbf{x}\|^2 \|z\|^2 \leq \frac{398}{100} \|\mathbf{x}\|^4.$$

So \mathcal{R}_b is covered by \mathcal{R}_1 .

Cover $\mathcal{R}_c \doteq \{z : \frac{1}{2} \|\mathbf{x}\| \|z\| \leq |\mathbf{x}^*z| \leq \frac{99}{100} \|\mathbf{x}\| \|z\|, \|z\| \geq \frac{11}{20} \|\mathbf{x}\|\}$: We show this region is covered by \mathcal{R}_2^z and $\mathcal{R}_2^{h'}$. First, for any $z \in \mathcal{R}_c$, when $\|z\| \geq \sqrt{\frac{1983}{1990}} \|\mathbf{x}\|$,

$$\frac{501}{500} \|\mathbf{x}\|^2 \|z\|^2 + |\mathbf{x}^*z|^2 \leq \frac{1983}{1000} \|\mathbf{x}\|^2 \|z\|^2 \leq \frac{199}{100} \|z\|^4,$$

implying that $\mathcal{R}_c \cap \{z : \|z\| \geq \sqrt{\frac{1983}{1990}} \|\mathbf{x}\|\} \subset \mathcal{R}_2^z$. Next we suppose $\|z\| = \lambda \|\mathbf{x}\|$ and $|\mathbf{x}^*z| = \eta \|\mathbf{x}\| \|z\|$, where $\lambda \in [\frac{11}{20}, \sqrt{\frac{1984}{1990}}]$ and $\eta \in [\frac{1}{2}, \frac{99}{100}]$, and show the rest of \mathcal{R}_c is covered by $\mathcal{R}_2^{h'}$. To this end, it is enough to verify that

$$2 \left(\|\mathbf{x}\|^2 - \|z\|^2 \right) |\mathbf{x}^*z| + 2 \|z\|^4 - \|\mathbf{x}\|^2 \|z\|^2 - |\mathbf{x}^*z|^2 - \frac{1}{250} \|\mathbf{x}\|^2 \|z\| \sqrt{\|\mathbf{x}\|^2 + \|z\|^2 - 2 |\mathbf{x}^*z|} \geq 0$$

over this subregion. Writing the left as a function of λ, η and eliminating $\|\mathbf{x}\|$ and $\|z\|$, it is enough to show

$$h(\lambda, \eta) \doteq 2(1 - \lambda^2)\eta + 2\lambda^3 - \lambda - \eta^2\lambda - \frac{1}{250} \sqrt{1 + \lambda^2 - 2\eta\lambda} \geq 0,$$

which is implied by

$$p(\lambda, \eta) \doteq 2(1 - \lambda^2)\eta + 2\lambda^3 - \lambda - \eta^2\lambda \geq \frac{49}{10000},$$

as $\frac{1}{250} \sqrt{1 + \lambda^2 - 2\eta\lambda} < 49/10000$. Let \mathbf{H}_p be the Hessian matrix of this bivariate function, it is easy to verify that $\det(\mathbf{H}_p) = -4(\eta + \lambda)^2 - 36\lambda^2 < 0$ for all valid (λ, η) . Thus, the minimizer must occur on the boundary. For any fixed λ , $2(1 - \lambda^2)\eta - \eta^2\lambda$ is minimized at either $\eta = 99/100$ or $\eta = 1/2$. When $\eta = 99/100$, p is minimized at $\lambda = (4 \cdot 0.99 + \sqrt{40 \cdot 0.99^2 + 24})/12 < \sqrt{1984/1990}$, giving $p \geq 0.019$; when $\eta = 1/2$, p is minimized when $\lambda = (4 \cdot 0.5 + \sqrt{40 \cdot 0.5^2 + 24})/12 = (2 + \sqrt{34})/12$, giving $p \geq 0.3$. Overall, $p \geq 0.019 > 49/10000$, as desired.

Cover $\mathcal{R}_d \doteq \{z : \frac{99}{100} \|\mathbf{x}\| \|z\| \leq |\mathbf{x}^*z| \leq \|\mathbf{x}\| \|z\|, \|z\| \geq \frac{11}{20} \|\mathbf{x}\|\}$: We show that this region is covered by \mathcal{R}_2^z , \mathcal{R}_3 , and $\mathcal{R}_2^{h'}$ together. First, for any $z \in \mathcal{R}_d$, when $\|z\| \geq \sqrt{\frac{1001}{995}} \|\mathbf{x}\|$,

$$\frac{501}{500} \|\mathbf{x}\|^2 \|z\|^2 + |\mathbf{x}^*z|^2 \leq \frac{1001}{500} \|\mathbf{x}\|^2 \|z\|^2 \leq \frac{199}{100} \|z\|^4.$$

So $\mathcal{R}_d \cap \{z : \|z\| \geq \sqrt{\frac{1001}{995}} \|\mathbf{x}\|\} \subset \mathcal{R}_2^z$. Next, we show that any $z \in \mathcal{R}_d$ with $\|z\| \leq 24/25 \cdot \|\mathbf{x}\|$ is contained in $\mathcal{R}_2^{h'}$. Similar to the above argument for \mathcal{R}_c , it is enough to show

$$p(\lambda, \eta) \doteq 2(1 - \lambda^2)\eta + 2\lambda^3 - \lambda - \eta^2\lambda \geq 0.00185,$$

as $\frac{1}{250} \sqrt{1 + \lambda^2 - 2\eta\lambda} < 0.00185$ in this case. Since the Hessian is again always indefinite, we check the optimal value for $\eta = 99/100$ and $\eta = 1$ and do the comparison. It can be verified $p \geq 0.00627 > 0.00185$ in this case. So $\mathcal{R}_d \cap \{z : \|z\| \leq \frac{24}{25} \|\mathbf{x}\|\} \subset \mathcal{R}_2^{h'}$. Finally, we consider the case $\frac{23}{25} \|\mathbf{x}\| \leq \|z\| \leq \sqrt{\frac{1005}{995}} \|\mathbf{x}\|$. A λ, η argument as above leads to

$$\|\mathbf{h}(z)\|^2 = \|\mathbf{x}\|^2 + \|z\|^2 - 2|\mathbf{x}^*z| < \frac{1}{7} \|\mathbf{x}\|^2,$$

implying that $\mathcal{R}_d \cap \left\{z : \frac{23}{25} \|\mathbf{x}\| \leq \|z\| \leq \sqrt{\frac{1005}{995}} \|\mathbf{x}\|\right\} \subset \mathcal{R}_3$.

In summary, now we obtain that $\mathbb{C}^n = \mathcal{R}_a \cup \mathcal{R}_b \cup \mathcal{R}_c \cup \mathcal{R}_d \subset \mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^{h'} \cup \mathcal{R}_3$. Observe that $\mathcal{R}_{h'}$ is only used to cover $\mathcal{R}_c \cup \mathcal{R}_d$, which is in turn a subset of $\{z : \|z\| \geq 11 \|\mathbf{x}\| / 20\}$. Thus, $\mathbb{C}^n = \mathcal{R}_1 \cup \mathcal{R}_2^z \cup (\mathcal{R}_2^{h'} \cap \{z : \|z\| \geq 11 \|\mathbf{x}\| / 20\}) \cup \mathcal{R}_3$. Moreover, by the definition of \mathcal{R}_3 ,

$$\begin{aligned} & \mathcal{R}_1 \cup \mathcal{R}_2^z \cup (\mathcal{R}_2^{h'} \cap \{z : \|z\| \geq 11 \|\mathbf{x}\| / 20\}) \cup \mathcal{R}_3 \\ & \subset \mathcal{R}_1 \cup \mathcal{R}_2^z \cup (\mathcal{R}_2^{h'} \cap \{z : \|z\| \geq 11 \|\mathbf{x}\| / 20\} \cap \mathcal{R}_3^c) \cup \mathcal{R}_3 \\ & \subset \mathcal{R}_1 \cup \mathcal{R}_2^z \cup \mathcal{R}_2^h \cup \mathcal{R}_3 \subset \mathbb{C}^n, \end{aligned}$$

implying the claimed coverage. ■

7 Proofs of Technical Results for Trust-Region Algorithm

7.1 Auxiliary Lemmas

Lemma 7.1 *When $m \geq Cn$ for a sufficiently large C , it holds with probability at least $1 - c_a \exp(-c_b m)$ that*

$$\frac{1}{m} \sum_{k=1}^m \left| |\mathbf{a}_k^* z|^2 - |\mathbf{a}_k^* w|^2 \right| \leq \frac{3}{2} \|z - w\| (\|z\| + \|w\|)$$

for all $z, w \in \mathbb{C}^n$. Here C, c_a, c_b are positive absolute constants.

Proof Lemma 3.1 in [CSV13] has shown that when $m \geq C_1 n$, it holds with probability at least $1 - c_2 \exp(-c_3 m)$ that

$$\frac{1}{m} \sum_{k=1}^m \left| |\mathbf{a}_k^* z|^2 - |\mathbf{a}_k^* w|^2 \right| \leq \frac{3}{2\sqrt{2}} \|zz^* - ww^*\|_*$$

for all z and w , where $\|\cdot\|_*$ is the nuclear norm that sums up singular values. The claims follows from

$$\|zz^* - ww^*\|_* \leq \sqrt{2} \|zz^* - ww^*\| \leq \sqrt{2} \|z - w\| (\|z\| + \|w\|),$$

completing the proof. ■

Lemma 7.2 When $m \geq Cn \log n$, with probability at least $1 - c_a m^{-1} - c_b \exp(-c_c m / \log m)$,

$$\left\| \nabla^2 f(\mathbf{x} e^{i\psi}) - \mathbb{E} \left[\nabla^2 f(\mathbf{x} e^{i\psi}) \right] \right\| \leq \frac{1}{100} \|\mathbf{x}\|^2$$

for all $\psi \in [0, 2\pi)$. Here C, c_a to c_c are positive absolute constants.

Proof By Lemma 6.3, we have that

$$\begin{aligned} & \left\| \nabla^2 f(\mathbf{x} e^{i\psi}) - \mathbb{E} \left[\nabla^2 f(\mathbf{x} e^{i\psi}) \right] \right\| \\ & \leq \left\| \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 \mathbf{a}_k \mathbf{a}_k - \left(\|\mathbf{x}\|^2 \mathbf{I} + \mathbf{x} \mathbf{x}^* \right) \right\| + \left\| \frac{1}{m} \sum_{k=1}^m (\mathbf{a}_k^* \mathbf{x})^2 \mathbf{a}_k \mathbf{a}_k^\top e^{i2\psi} - 2\mathbf{x} \mathbf{x}^\top e^{i2\psi} \right\| \\ & \leq \frac{1}{200} \|\mathbf{x}\|^2 + \frac{1}{200} \|\mathbf{x}\|^2 \leq \frac{1}{100} \|\mathbf{x}\|^2 \end{aligned}$$

holds with high probability when $m \geq C_1 n \log n$ for a sufficiently large C_1 . ■

7.2 Proof of Lemma 3.1

Proof For any $\mathbf{z}, \mathbf{z}' \in \Gamma'$, we have

$$\begin{aligned} |f(\mathbf{z}) - f(\mathbf{z}')| &= \frac{1}{2m} \left| \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{z}|^4 - |\mathbf{a}_k^* \mathbf{z}'|^4 - 2 \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 (|\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{z}'|^2) \right| \\ &\leq \frac{1}{2m} \sum_{k=1}^m (|\mathbf{a}_k^* \mathbf{z}|^2 + |\mathbf{a}_k^* \mathbf{z}'|^2) \left| |\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{z}'|^2 \right| + \frac{1}{m} \sum_{k=1}^m |\mathbf{a}_k^* \mathbf{x}|^2 \left| |\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{z}'|^2 \right| \\ &\leq 4R_1^2 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \cdot \frac{3}{2} \cdot 4R_1 \|\mathbf{z} - \mathbf{z}'\| + 2 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^2 \cdot \frac{3}{2} \cdot 4R_1 \|\mathbf{z} - \mathbf{z}'\| \\ &\leq (24R_1^3 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 + 12 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^2 R_1) \|\mathbf{z} - \mathbf{z}'\|, \end{aligned}$$

where in the third line we invoked results of Lemma 7.1, and hence the derived inequality holds with high probability when $m \geq C_1 n$. Similarly, for the gradient,

$$\begin{aligned} & \|\nabla f(\mathbf{z}) - \nabla f(\mathbf{z}')\| \\ &= \frac{\sqrt{2}}{m} \left\| \sum_{k=1}^m \left(|\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{x}|^2 \right) \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} - \sum_{k=1}^m \left(|\mathbf{a}_k^* \mathbf{z}'|^2 - |\mathbf{a}_k^* \mathbf{x}|^2 \right) \mathbf{a}_k \mathbf{a}_k^* \mathbf{z}' \right\| \\ &\leq \frac{\sqrt{2}}{m} \sum_{k=1}^m \left\| \left(|\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{z}'|^2 \right) \mathbf{a}_k \mathbf{a}_k^* \mathbf{z} \right\| + \sqrt{2} \left\| \frac{1}{m} \sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^* |\mathbf{a}_k^* \mathbf{z}'|^2 \right\| \|\mathbf{z} - \mathbf{z}'\| \\ &\quad + \sqrt{2} \left\| \frac{1}{m} \sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^* |\mathbf{a}_k^* \mathbf{x}|^2 \right\| \|\mathbf{z} - \mathbf{z}'\| \\ &\leq \sqrt{2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \cdot 2R_1 \cdot \frac{3}{2} \cdot 4R_1 \|\mathbf{z} - \mathbf{z}'\| + (8\sqrt{2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 R_1^2 + 2\sqrt{2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^2) \|\mathbf{z} - \mathbf{z}'\| \\ &\leq (20\sqrt{2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 R_1^2 + 2\sqrt{2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\|^2) \|\mathbf{z} - \mathbf{z}'\|, \end{aligned}$$

where from the second to the third inequality we used the fact $\left\| \frac{1}{m} \sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^* \right\| \leq 2$ with probability at least $1 - \exp(-c_2 m)$. Similarly for the Hessian,

$$\begin{aligned}
& \left\| \nabla^2 f(\mathbf{z}) - \nabla^2 f(\mathbf{z}') \right\| \\
&= \sup_{\|\mathbf{w}\|=1} \left\| \frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* (\nabla^2 f(\mathbf{z}) - \nabla^2 f(\mathbf{z}')) \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} \right\| \\
&\leq \sup_{\|\mathbf{w}\|=1} 2 \left\| \frac{1}{m} \sum_{k=1}^m (|\mathbf{a}_k^* \mathbf{z}|^2 - |\mathbf{a}_k^* \mathbf{z}'|^2) |\mathbf{a}_k^* \mathbf{w}|^2 \right\| + \left\| \frac{1}{m} \sum_{k=1}^m \Re((\mathbf{a}_k^* \mathbf{z})^2 - (\mathbf{a}_k^* \mathbf{z}')^2) (\mathbf{w}^* \mathbf{a}_k)^2 \right\| \\
&\leq 2 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \cdot \frac{3}{2} \cdot 4R_1 \|\mathbf{z} - \mathbf{z}'\| + \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \cdot 4R_1 \cdot \|\mathbf{z} - \mathbf{z}'\| \cdot 2 \\
&\leq 16 \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 R_1 \|\mathbf{z} - \mathbf{z}'\|,
\end{aligned}$$

where to obtain the third inequality we used that $\frac{1}{m} \|\mathbf{A}^*\|^2 \leq 2$ with probability at least $1 - \exp(-c_3 m)$ when $m \geq C_4 n$ for a sufficiently large constant C_4 .

Since $R_0 \leq 10 \|\mathbf{x}\|$ with probability at least $1 - \exp(-c_5 m)$ when $m \geq C_6 n$, by definition of R_1 , we have $R_1 \leq 30(n \log m)^{1/2} \|\mathbf{x}\|$ with high probability. Substituting this estimate into the above bounds yields the claimed results. \blacksquare

7.3 Proof of Lemma 3.2

Proof For the upper bound, we have that for all $\mathbf{z} \in \mathcal{R}'_3$,

$$\begin{aligned}
\|\mathbf{H}(\mathbf{z})\| &\leq \|\nabla^2 f(\mathbf{z})\| \leq \left\| \nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) \right\| + L_h \|\mathbf{h}(\mathbf{z})\| \\
&\leq \left\| \nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) - \mathbb{E} \left[\nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) \right] \right\| + \left\| \mathbb{E} \left[\nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) \right] \right\| + \frac{1}{10} \|\mathbf{x}\|^2 \\
&\leq \frac{1}{100} \|\mathbf{x}\|^2 + 4 \|\mathbf{x}\|^2 + \frac{1}{10} \|\mathbf{x}\|^2 \leq \frac{9}{2} \|\mathbf{x}\|^2,
\end{aligned}$$

where to obtain the third line we applied Lemma 7.2. To show the lower bound for all $\mathbf{z} \in \mathcal{R}'_3$, it is equivalent to show that

$$\frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} \geq m_H, \quad \forall \|\mathbf{w}\| = 1 \text{ with } \Im(\mathbf{w}^* \mathbf{z}) = 0, \text{ and } \forall \mathbf{z} \in \mathcal{R}'_3.$$

By Lemma 3.1 and Lemma 7.2, with high probability, we have

$$\begin{aligned}
\frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} &\geq \frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* \nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} - L_h \|\mathbf{h}(\mathbf{z})\| \|\mathbf{w}\|^2 \\
&\geq \frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* \mathbb{E} \left[\nabla^2 f(\mathbf{x} e^{i\phi(\mathbf{z})}) \right] \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} - \left(\frac{1}{10} + \frac{1}{100} \right) \|\mathbf{x}\|^2 \\
&= \left(1 - \frac{1}{100} - \frac{1}{10} \right) \|\mathbf{x}\|^2 + |\mathbf{w}^* \mathbf{x}|^2 + 2\Re \left((\mathbf{w}^* \mathbf{x} e^{i\phi(\mathbf{z})})^2 \right) \\
&\geq \frac{89}{100} \|\mathbf{x}\|^2 + \Re \left((\mathbf{w}^* \mathbf{x} e^{i\phi(\mathbf{z})})^2 \right).
\end{aligned}$$

Since $\Im(\mathbf{w}^* \mathbf{z}) = 0$, we have $\Re((\mathbf{w}^* \mathbf{z})^2) = |\mathbf{w}^* \mathbf{z}|^2$. Thus,

$$\begin{aligned}
\Re\left((\mathbf{w}^* \mathbf{x} e^{i\phi(\mathbf{z})})^2\right) &= \Re\left((\mathbf{w}^* \mathbf{z} - \mathbf{w}^* \mathbf{h}(\mathbf{z}))^2\right) \\
&= |\mathbf{w}^* \mathbf{z}|^2 + \Re\left((\mathbf{w}^* \mathbf{h}(\mathbf{z}))^2\right) - 2\Re\left((\mathbf{w}^* \mathbf{h}(\mathbf{z}))(\mathbf{w}^* \mathbf{z})\right) \\
&\geq |\mathbf{w}^* \mathbf{z}|^2 - \|\mathbf{w}\|^2 \|\mathbf{h}(\mathbf{z})\|^2 - 2\|\mathbf{w}\|^2 \|\mathbf{h}(\mathbf{z})\| \|\mathbf{z}\| \\
&\geq -\frac{1}{100L_h^2} \|\mathbf{x}\|^4 - \frac{2}{10L_h} \|\mathbf{x}\|^2 \left(\|\mathbf{x}\| + \frac{1}{10L_h} \|\mathbf{x}\|^2\right) \\
&\geq -\frac{1}{100} \|\mathbf{x}\|^2,
\end{aligned}$$

where we obtained the last inequality based on the fact that $L_h \doteq 480(n \log m)^{1/2} \|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \|\mathbf{x}\| \geq 150 \|\mathbf{x}\|$ whenever $\|\mathbf{A}\|_{\ell^1 \rightarrow \ell^2}^2 \geq 1$; this holds with high probability when $m \geq C_1 n$ for large enough constant C_1 . Together we obtain

$$\frac{1}{2} \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \mathbf{w} \\ \overline{\mathbf{w}} \end{bmatrix} \geq \frac{89}{100} \|\mathbf{x}\|^2 - \frac{1}{100} \|\mathbf{x}\|^2 \geq \frac{22}{25} \|\mathbf{x}\|^2,$$

as desired. ■

7.4 Proof of Lemma 3.3

Proof In view of Lemma A.3, we have

$$\begin{aligned}
f(\mathbf{z} + \boldsymbol{\delta}_*) &\leq \widehat{f}(\boldsymbol{\delta}_*; \mathbf{z}) + \frac{1}{3} L_h \Delta^3 \\
&\leq \widehat{f}(\boldsymbol{\delta}; \mathbf{z}) + \frac{1}{3} L_h \Delta^3 \\
&\leq f(\mathbf{z} + \boldsymbol{\delta}) + \frac{2}{3} L_h \Delta^3 \\
&\leq f(\mathbf{z}) - d + \frac{2}{3} L_h \Delta^3,
\end{aligned}$$

as desired. ■

7.5 Proof of Proposition 3.4

Proof In view of Proposition 2.3, consider direction $\boldsymbol{\delta} \doteq \mathbf{x} e^{i\phi(\mathbf{z})} / \|\mathbf{x}\|$. Obviously, vectors of the form $t\sigma\boldsymbol{\delta}$ are feasible for (3.1) for any $t \in [0, \Delta]$ and $\sigma \doteq -\text{sign}([\boldsymbol{\delta}^*, \overline{\boldsymbol{\delta}}^*] \nabla f(\mathbf{z}^{(r)}))$. By Lemma A.2, we obtain

$$\begin{aligned}
f(\mathbf{z}^{(r)} + t\sigma\boldsymbol{\delta}) &= f(\mathbf{z}^{(r)}) + t\sigma \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r)}) + t^2 \int_0^1 (1-s) \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)} + \sigma st\boldsymbol{\delta}) \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix} ds \\
&\leq f(\mathbf{z}^{(r)}) + \frac{t^2}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix} \\
&\quad + t^2 \int_0^1 (1-s) \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix}^* \left[\nabla^2 f(\mathbf{z}^{(r)} + \sigma st\boldsymbol{\delta}) - \nabla^2 f(\mathbf{z}^{(r)}) \right] \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix} ds \\
&\leq f(\mathbf{z}^{(r)}) + \frac{t^2}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \boldsymbol{\delta} \\ \overline{\boldsymbol{\delta}} \end{bmatrix} + \frac{L_h}{3} t^3.
\end{aligned}$$

Thus, we have

$$f(\mathbf{z}^{(r)} + t\sigma\boldsymbol{\delta}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{200}t^2\|\mathbf{x}\|^2 + \frac{L_h}{3}t^3.$$

Taking $t = \Delta$ and applying Lemma 3.3, we have

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{200}\Delta^2\|\mathbf{x}\|^2 + \frac{L_h}{3}\Delta^3 + \frac{2}{3}L_h\Delta^3 \leq -\frac{1}{200}\Delta^2\|\mathbf{x}\|^2 + L_h\Delta^3 \leq -\frac{1}{400}\|\mathbf{x}\|^2\Delta^2,$$

where we obtain the very last inequality using the assumption that $\Delta \leq \|\mathbf{x}\|^2 / (400L_h)$, completing the proof. \blacksquare

7.6 Proof of Proposition 3.5

Proof We take

$$\boldsymbol{\delta} = \begin{cases} -\mathbf{z}^{(r)} / \|\mathbf{z}^{(r)}\| & \mathbf{z}^{(r)} \in \mathcal{R}_2^z \\ -\mathbf{h}(\mathbf{z}^{(r)}) / \|\mathbf{h}(\mathbf{z}^{(r)})\| & \mathbf{z}^{(r)} \in \mathcal{R}_2^h \end{cases}.$$

Obviously vectors of the form $t\boldsymbol{\delta}$ is feasible for (3.1) for any $t \in [0, \Delta]$. By Lemma A.2, we have

$$\begin{aligned} f(\mathbf{z}^{(r)} + t\boldsymbol{\delta}) &= f(\mathbf{z}^{(r)}) + t \int_0^1 \left[\frac{\boldsymbol{\delta}}{\delta} \right]^* \nabla f(\mathbf{z}^{(r)} + s t \boldsymbol{\delta}) ds \\ &= f(\mathbf{z}^{(r)}) + t \left[\frac{\boldsymbol{\delta}}{\delta} \right]^* \nabla f(\mathbf{z}^{(r)}) + t \int_0^1 \left[\frac{\boldsymbol{\delta}}{\delta} \right]^* \left[\nabla f(\mathbf{z}^{(r)} + s t \boldsymbol{\delta}) - \nabla f(\mathbf{z}^{(r)}) \right] ds \\ &\leq f(\mathbf{z}^{(r)}) + t \left[\frac{\boldsymbol{\delta}}{\delta} \right]^* \nabla f(\mathbf{z}^{(r)}) + t^2 L_g. \end{aligned}$$

By Proposition 2.4 and Proposition 2.5, we have

$$f(\mathbf{z}^{(r)} + t\boldsymbol{\delta}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{1000}t\|\mathbf{x}\|^2\|\mathbf{z}^{(r)}\| + t^2 L_g.$$

Since $\{\mathbf{z} : \|\mathbf{z}\| \leq \|\mathbf{x}\|/2\} \subset \mathcal{R}_1$, $\mathbf{z}^{(r)}$ of interest here satisfies $\|\mathbf{z}^{(r)}\| \geq \|\mathbf{x}\|/2$. Thus,

$$f(\mathbf{z}^{(r)} + t\boldsymbol{\delta}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{2000}t\|\mathbf{x}\|^3 + t^2 L_g.$$

Combining the above with Lemma 3.3, we obtain

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{2000}\Delta\|\mathbf{x}\|^3 + \Delta^2 L_g + \frac{2}{3}L_h\Delta^3 \leq -\frac{1}{4000}\Delta\|\mathbf{x}\|^3,$$

provided

$$\Delta \leq \min \left\{ \frac{\|\mathbf{x}\|^3}{8000L_g}, \sqrt{\frac{3\|\mathbf{x}\|^3}{16000L_h}} \right\},$$

as desired. \blacksquare

7.7 Proof of Proposition 3.6

Proof By Proposition 2.6 and the integral form of Taylor's theorem in Lemma A.2, we have that for any \mathbf{g} satisfying $\Im(\mathbf{g}^* \mathbf{x}) = 0$ and $\|\mathbf{g}\| = 1$ and any $t \in [0, \|\mathbf{x}\| / \sqrt{7}]$,

$$\begin{aligned} f(\mathbf{x} + t\mathbf{g}) &= f(\mathbf{x}) + t \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla f(\mathbf{x}) + t^2 \int_0^1 (1-s) \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla^2 f(\mathbf{x} + st\mathbf{g}) \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right] ds \\ &\geq f(\mathbf{x}) + t \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla f(\mathbf{x}) + \frac{1}{8} \|\mathbf{x}\|^2 t^2. \end{aligned}$$

Similarly, we have

$$f(\mathbf{x}) \geq f(\mathbf{x} + t\mathbf{g}) - t \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla f(\mathbf{x} + t\mathbf{g}) + \frac{1}{8} \|\mathbf{x}\|^2 t^2.$$

Combining the above two inequalities, we obtain

$$t \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* (\nabla f(\mathbf{x} + t\mathbf{g}) - \nabla f(\mathbf{x})) \geq \frac{1}{4} \|\mathbf{x}\|^2 t^2 \implies \left[\frac{\mathbf{g}}{\bar{\mathbf{g}}} \right]^* \nabla f(\mathbf{x} + t\mathbf{g}) \geq \frac{1}{4} \|\mathbf{x}\|^2 t \geq \frac{1}{40L_h} \|\mathbf{x}\|^4,$$

where to obtain the very last bound we have used the fact $\min_{\mathbf{z} \in \mathcal{R}_3 \setminus \mathcal{R}'_3} \|\mathbf{h}(\mathbf{z})\| \geq \|\mathbf{x}\|^2 / (10L_h)$ due to (3.8). This implies that for all $\mathbf{z} \in \mathcal{R}_3 \setminus \mathcal{R}'_3$,

$$\left[\frac{\mathbf{h}(\mathbf{z})}{\bar{\mathbf{h}}(\mathbf{z})} \right]^* \nabla f(\mathbf{z}) \geq \frac{1}{40L_h} \|\mathbf{x}\|^4. \quad (7.1)$$

The rest arguments are very similar to that of Proposition 3.5. Take $\boldsymbol{\delta} = -\mathbf{h}(\mathbf{z}^{(r)}) / \|\mathbf{h}(\mathbf{z}^{(r)})\|$ and it can checked vectors of the form $t\boldsymbol{\delta}$ for $t \in [0, \Delta]$ are feasible for (3.1). By Lemma A.2, we have

$$\begin{aligned} f(\mathbf{z}^{(r)} + t\boldsymbol{\delta}) &= f(\mathbf{z}^{(r)}) + t \int_0^1 \left[\frac{\boldsymbol{\delta}}{\bar{\boldsymbol{\delta}}} \right]^* \nabla f(\mathbf{z}^{(r)} + st\boldsymbol{\delta}) ds \\ &= f(\mathbf{z}^{(r)}) + t \left[\frac{\boldsymbol{\delta}}{\bar{\boldsymbol{\delta}}} \right]^* \nabla f(\mathbf{z}^{(r)}) + t \int_0^1 \left[\frac{\boldsymbol{\delta}}{\bar{\boldsymbol{\delta}}} \right]^* \left[\nabla f(\mathbf{z}^{(r)} + st\boldsymbol{\delta}) - \nabla f(\mathbf{z}^{(r)}) \right] ds \\ &\leq f(\mathbf{z}^{(r)}) + t \left[\frac{\boldsymbol{\delta}}{\bar{\boldsymbol{\delta}}} \right]^* \nabla f(\mathbf{z}^{(r)}) + t^2 L_g \\ &\leq f(\mathbf{z}^{(r)}) - \frac{1}{40L_h} t \|\mathbf{x}\|^4 + t^2 L_g, \end{aligned}$$

where to obtain the last line we have used (7.1). Combining the above with Lemma 3.3, we obtain

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -\frac{1}{40L_h} \Delta \|\mathbf{x}\|^4 + \Delta^2 L_g + \frac{2}{3} L_h \Delta^3 \leq -\frac{1}{80L_h} \Delta \|\mathbf{x}\|^4,$$

provided

$$\Delta \leq \min \left\{ \frac{\|\mathbf{x}\|^4}{160L_h L_g}, \sqrt{\frac{3}{320}} \frac{\|\mathbf{x}\|^2}{L_h} \right\},$$

as desired. ■

7.8 Proof of Proposition 3.7

Proof If we identify \mathbb{C}^n with \mathbb{R}^{2n} , it can be easily verified that the orthoprojectors of a vector \mathbf{w} onto \mathbf{z} and its orthogonal complement are

$$\mathcal{P}_{\mathbf{z}}(\mathbf{w}) = \frac{\Re(\mathbf{z}^* \mathbf{w}) \mathbf{z}}{\|\mathbf{z}\|^2}, \quad \text{and} \quad \mathcal{P}_{\mathbf{z}^\perp}(\mathbf{w}) = \mathbf{w} - \frac{\Re(\mathbf{z}^* \mathbf{w}) \mathbf{z}}{\|\mathbf{z}\|^2}.$$

Now at any point $\mathbf{z}^{(r)} \in \mathcal{R}'_3$, consider a feasible direction of the form $\boldsymbol{\delta} \doteq -t \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)})$ ($0 \leq t \leq \Delta / \|\mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)})\|$) to the trust-region subproblem (3.1). The local quadratic approximation obeys

$$\begin{aligned} \widehat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}) &= f(\mathbf{z}^{(r)}) + \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r)}) + \frac{1}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix} \\ &\leq f(\mathbf{z}^{(r)}) - 2t \left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\|^2 + t^2 M_H \left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\|^2 \\ &= f(\mathbf{z}^{(r)}) - 2t \left(1 - \frac{M_H}{2} t \right) \left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\|^2, \end{aligned}$$

where M_H is as defined in Lemma 3.2. Taking $t = \min\{M_H^{-1}, \Delta / \|\mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)})\|\}$, we have

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}) - f(\mathbf{z}^{(r)}) \leq -\min\{M_H^{-1}, \Delta / \|\mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)})\|\} \left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\|^2.$$

Let \mathbf{U} be an orthogonal (in geometric sense) basis for the space $\{\mathbf{w} : \Im(\mathbf{w}^* \mathbf{z}^{(r)}) = 0\}$. In view of the transformed gradient and Hessian in (3.3), it is easy to see

$$\left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\| = \frac{1}{\sqrt{2}} \left\| \mathbf{g}(\mathbf{z}^{(r)}) \right\|,$$

where $\mathbf{g}(\mathbf{z}^{(r)})$ is the transformed gradient. To lower bound $\left\| \mathcal{P}_{(\mathbf{iz}^{(r)})^\perp} \nabla_{\mathbf{z}^{(r)}} f(\mathbf{z}^{(r)}) \right\|$, recall the step is constrained, we have

$$\Delta \leq \left\| \mathbf{H}^{-1}(\mathbf{z}^{(r)}) \mathbf{g}(\mathbf{z}^{(r)}) \right\| \leq \left\| \mathbf{H}^{-1}(\mathbf{z}^{(r)}) \right\| \left\| \mathbf{g}(\mathbf{z}^{(r)}) \right\| \leq \frac{1}{\lambda_{\min}(\mathbf{H}(\mathbf{z}^{(r)}))} \left\| \mathbf{g}(\mathbf{z}^{(r)}) \right\|.$$

By Lemma 3.2, $\lambda_{\min}(\mathbf{H}(\mathbf{z}^{(r)})) \geq m_H$. Thus,

$$\left\| \mathbf{g}(\mathbf{z}^{(r)}) \right\| \geq m_H \Delta.$$

Hence we have

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{z}^{(r)}) - f(\mathbf{z}^{(r)}) \leq -\min \left\{ \frac{m_H^2 \Delta^2}{2M_H}, \frac{\Delta^2 m_H}{\sqrt{2}} \right\} \leq -\frac{m_H^2 \Delta^2}{2M_H},$$

where the last simplification is due to that $M_H \geq m_H$. By Lemma A.3, we have

$$f(\mathbf{z}^{(r)} + \boldsymbol{\delta}) - f(\mathbf{z}^{(r)}) \leq -\frac{m_H^2 \Delta^2}{2M_H} + \frac{L_h}{3} \Delta^3.$$

Therefore, for $\mathbf{z}^{(r+1)} = \mathbf{z}^{(r)} + \boldsymbol{\delta}_*$, Lemma 3.3 implies that

$$f(\mathbf{z}^{(r+1)}) - f(\mathbf{z}^{(r)}) \leq -\frac{m_H^2 \Delta^2}{2M_H} + L_h \Delta^3.$$

The claimed result follows provided $\Delta \leq m_H^2 / (4M_H L_h)$, completing the proof. \blacksquare

7.9 Proof of Proposition 3.8

Before proceeding, we note one important fact that is useful below. For any \mathbf{z} , we have

$$\mathcal{P}_{i\mathbf{z}} \nabla_{\mathbf{z}} f(\mathbf{z}) = \frac{\Re((i\mathbf{z})^* \nabla_{\mathbf{z}} f(\mathbf{z}))}{\|\mathbf{z}\|^2} i\mathbf{z} = \mathbf{0}.$$

Thus, if $\mathbf{U}(\mathbf{z})$ is an (geometrically) orthonormal basis constructed for the space $\{\mathbf{w} : \Im(\mathbf{w}^* \mathbf{z}) = 0\}$ (as defined around (3.3)), it is easy to verify that

$$\begin{bmatrix} \mathbf{U} \\ \overline{\mathbf{U}} \end{bmatrix} \begin{bmatrix} \mathbf{U} \\ \overline{\mathbf{U}} \end{bmatrix}^* \nabla f(\mathbf{z}) = 2\nabla f(\mathbf{z}). \quad (7.2)$$

We next prove Proposition 3.8.

Proof Throughout the proof, we write $\mathbf{g}^{(r)}$, $\mathbf{H}^{(r)}$ and $\mathbf{U}^{(r)}$ short for $\mathbf{g}(\mathbf{z}^{(r)})$, $\mathbf{H}(\mathbf{z}^{(r)})$ and $\mathbf{U}(\mathbf{z}^{(r)})$, respectively. Given an orthonormal basis $\mathbf{U}^{(r)}$ for $\{\mathbf{w} : \Im(\mathbf{w}^* \mathbf{z}^{(r)}) = 0\}$, the unconstrained optimality condition of the trust region method implies that

$$\mathbf{H}^{(r)} \boldsymbol{\xi}_* + \mathbf{g}^{(r)} = \mathbf{0} \iff \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* + \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r)}) = \mathbf{0}.$$

Thus, we have

$$\begin{aligned} & \|\nabla f(\mathbf{z}^{(r+1)})\| \\ &= \frac{1}{2} \left\| \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r+1)}) \right\| \\ &= \frac{1}{2} \left\| \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix}^* \nabla f(\mathbf{z}^{(r+1)}) - \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix}^* \left(\nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* + \nabla f(\mathbf{z}^{(r)}) \right) \right\| \\ &\leq \frac{1}{2} \left\| \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix}^* \left[\nabla f(\mathbf{z}^{(r+1)}) - \nabla f(\mathbf{z}^{(r)}) - \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* \right] \right\| \\ &\quad + \frac{1}{2} \left\| \left(\begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix}^* - \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix}^* \right) \left(\nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* + \nabla f(\mathbf{z}^{(r)}) \right) \right\| \\ &\leq \left\| \nabla f(\mathbf{z}^{(r+1)}) - \nabla f(\mathbf{z}^{(r)}) - \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* \right\| \\ &\quad + \frac{1}{2} \left\| \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r+1)} \\ \overline{\mathbf{U}^{(r+1)}} \end{bmatrix}^* - \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix}^* \right\| \left\| \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* + \nabla f(\mathbf{z}^{(r)}) \right\|. \end{aligned}$$

By Taylor's theorem and Lipschitz property in Lemma 3.1, we have

$$\begin{aligned} & \left\| \nabla f(\mathbf{z}^{(r+1)}) - \nabla f(\mathbf{z}^{(r)}) - \nabla^2 f(\mathbf{z}^{(r)}) \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* \right\| \\ &= \left\| \int_0^1 \left[\nabla^2 f(\mathbf{z}^{(r)} + t \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_*) - \nabla^2 f(\mathbf{z}^{(r)}) \right] \begin{bmatrix} \mathbf{U}^{(r)} \\ \overline{\mathbf{U}^{(r)}} \end{bmatrix} \boldsymbol{\xi}_* dt \right\| \end{aligned}$$

$$\begin{aligned}
&\leq \|\boldsymbol{\xi}_\star\| \int_0^1 \left\| \nabla^2 f(\mathbf{z}^{(r)}) + t \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right] \boldsymbol{\xi}_\star - \nabla^2 f(\mathbf{z}^{(r)}) \right\| dt \\
&\leq \frac{1}{2} L_h \|\boldsymbol{\xi}_\star\|^2.
\end{aligned} \tag{7.3}$$

Moreover,

$$\begin{aligned}
\left\| \nabla f(\mathbf{z}^{(r)}) \right\| &= \frac{1}{\sqrt{2}} \left\| \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right]^* \nabla f(\mathbf{z}^{(r)}) \right\| \\
&= \frac{1}{\sqrt{2}} \left\| - \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right]^* \nabla^2 f(\mathbf{z}^{(r)}) \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right] \boldsymbol{\xi}_\star \right\| \leq \sqrt{2} \left\| \nabla^2 f(\mathbf{z}^{(r)}) \right\| \|\boldsymbol{\xi}_\star\|,
\end{aligned}$$

where to obtain the second equality we have used the optimality condition discussed at start of the proof. Thus, using the result above, we obtain

$$\left\| \nabla^2 f(\mathbf{z}^{(r)}) \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right] \boldsymbol{\xi}_\star + \nabla f(\mathbf{z}^{(r)}) \right\| \leq 2\sqrt{2} \left\| \nabla^2 f(\mathbf{z}^{(r)}) \right\| \|\boldsymbol{\xi}_\star\|. \tag{7.4}$$

On the other hand,

$$\begin{aligned}
&\left\| \left[\frac{\mathbf{U}^{(r+1)}}{\mathbf{U}^{(r+1)}} \right] \left[\frac{\mathbf{U}^{(r+1)}}{\mathbf{U}^{(r+1)}} \right]^* - \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right] \left[\frac{\mathbf{U}^{(r)}}{\mathbf{U}^{(r)}} \right]^* \right\| \\
&\leq \left\| \mathbf{U}^{(r+1)}(\mathbf{U}^{(r+1)})^* - \mathbf{U}^{(r)}(\mathbf{U}^{(r)})^* \right\| + \left\| \mathbf{U}^{(r+1)}(\mathbf{U}^{(r+1)})^\top - \mathbf{U}^{(r)}(\mathbf{U}^{(r)})^\top \right\|.
\end{aligned}$$

Write $\mathbf{U}^{(r+1)} = \mathbf{U}_{\Re}^{(r+1)} + i\mathbf{U}_{\Im}^{(r+1)}$, where $\mathbf{U}_{\Re}^{(r+1)}$ and $\mathbf{U}_{\Im}^{(r+1)}$ collect respectively entrywise real and imaginary parts of $\mathbf{U}^{(r+1)}$. It is not difficult to verify that $\mathbf{V}^{(r+1)} \doteq [\mathbf{U}_{\Re}^{(r+1)}; \mathbf{U}_{\Im}^{(r+1)}] \in \mathbb{R}^{2n \times (2n-1)}$ is an orthonormal matrix. We also define $\mathbf{V}^{(r)}$ accordingly. Thus,

$$\begin{aligned}
\left\| \mathbf{U}^{(r+1)}(\mathbf{U}^{(r+1)})^* - \mathbf{U}^{(r)}(\mathbf{U}^{(r)})^* \right\| &= \left\| [\mathbf{I}, i\mathbf{I}] \left(\mathbf{V}^{(r+1)}(\mathbf{V}^{(r+1)})^\top - \mathbf{V}^{(r)}(\mathbf{V}^{(r)})^\top \right) [\mathbf{I}, -i\mathbf{I}]^\top \right\| \\
&\leq 2 \left\| \mathbf{V}^{(r+1)}(\mathbf{V}^{(r+1)})^\top - \mathbf{V}^{(r)}(\mathbf{V}^{(r)})^\top \right\| \\
&\leq 2\sqrt{2} \left\| \mathbf{V}^{(r+1)}(\mathbf{V}^{(r+1)})^\top - \mathbf{V}^{(r)}(\mathbf{V}^{(r)})^\top \right\|_R,
\end{aligned}$$

where from the second to the third line we translate the complex operator norm to the real operator norm. Similarly, we also get

$$\left\| \mathbf{U}^{(r+1)}(\mathbf{U}^{(r+1)})^\top - \mathbf{U}^{(r)}(\mathbf{U}^{(r)})^\top \right\| \leq 2\sqrt{2} \left\| \mathbf{V}^{(r+1)}(\mathbf{V}^{(r+1)})^\top - \mathbf{V}^{(r)}(\mathbf{V}^{(r)})^\top \right\|_R.$$

Since $iz^{(r)}$ is the normal vector of the space generated by $\mathbf{U}^{(r)}$, $[-z_{\Im}^{(r)}; z_{\Re}^{(r)}]$ is the corresponding normal vector of $\mathbf{V}^{(r)}$. By Lemma A.9, the largest principal angle θ_1 between the subspaces designated by $\mathbf{V}^{(r+1)}$ and $\mathbf{V}^{(r)}$ are the angle between their normal vectors $\mathbf{a} \doteq [-z_{\Im}^{(r)}; z_{\Re}^{(r)}]$ and $\mathbf{b} \doteq [-z_{\Im}^{(r+1)}; z_{\Re}^{(r+1)}]$. Here we have decomposed $\mathbf{z}^{(r+1)}$ and $\mathbf{z}^{(r)}$ into real and imaginary parts. Similarly we define $\mathbf{c} \doteq [-(\boldsymbol{\delta}_\star)_{\Im}; (\boldsymbol{\delta}_\star)_{\Re}]$. By the law of cosines,

$$\cos \theta_1 = \frac{\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - \|\mathbf{c}\|^2}{2\|\mathbf{a}\|\|\mathbf{b}\|} \geq 1 - \frac{\|\mathbf{c}\|^2}{2\|\mathbf{a}\|\|\mathbf{b}\|} = 1 - \frac{\|\boldsymbol{\xi}_\star\|^2}{2\|\mathbf{z}^{(r)}\|\|\mathbf{z}^{(r+1)}\|}.$$

Since $\|\mathbf{z}^{(r)}\| \geq \min_{\mathbf{z} \in \mathcal{R}_3} \|\mathbf{z}\| \geq (1 - 1/\sqrt{7}) \|\mathbf{x}\| \geq 3\|\mathbf{x}\|/5$, and $\|\mathbf{z}^{(r+1)}\| \geq \|\mathbf{z}^{(r)}\| - \Delta \geq \|\mathbf{x}\|/2$ provided

$$\Delta \leq \|\mathbf{x}\|/10,$$

we obtain that

$$\cos \theta_1 \geq 1 - \frac{5}{3\|\mathbf{x}\|^2} \|\boldsymbol{\xi}_*\|^2.$$

Thus, by Lemma A.9 again,

$$\begin{aligned} \left\| \mathbf{V}^{(r+1)}(\mathbf{V}^{(r+1)})^\top - \mathbf{V}^{(r)}(\mathbf{V}^{(r)})^\top \right\|_R &= \sqrt{1 - \cos^2 \theta_1} \\ &\leq \sqrt{\frac{10}{3\|\mathbf{x}\|^2} \|\boldsymbol{\delta}_*\|^2 + \frac{25}{9\|\mathbf{x}\|^4} \|\boldsymbol{\xi}_*\|^4} \leq \frac{2}{\|\mathbf{x}\|} \|\boldsymbol{\xi}_*\|, \end{aligned} \quad (7.5)$$

where we used the assumption $\Delta \leq \|\mathbf{x}\|/10$ again to obtain the last inequality.

Collecting the above results, we obtain

$$\left\| \nabla f(\mathbf{z}^{(r+1)}) \right\| \leq \left(\frac{1}{2} L_h + \frac{16}{\|\mathbf{x}\|} M_H \right) \|\boldsymbol{\xi}_*\|^2. \quad (7.6)$$

Invoking the optimality condition again, we obtain

$$\|\boldsymbol{\xi}_*\|^2 = \left\| (\mathbf{H}^{(r)})^{-1} \mathbf{g}^{(r)} \right\|^2 \leq \frac{1}{m_H^2} \left\| \mathbf{g}^{(r)} \right\|^2 = \frac{2}{m_H^2} \left\| \nabla f(\mathbf{z}^{(r)}) \right\|^2. \quad (7.7)$$

Here $(\mathbf{H}^{(r)})^{-1}$ is well-defined because Lemma 3.2 shows that $\|\mathbf{H}^{(r)}\| \geq m_H$ for all $\mathbf{z}^{(r)} \in \mathcal{R}'_3$. Combining the last two estimates, we complete the proof. \blacksquare

7.10 Proof of Proposition 3.9

Proof Throughout the proof, we write $\mathbf{g}^{(r)}$, $\mathbf{H}^{(r)}$ and $\mathbf{U}^{(r)}$ short for $\mathbf{g}(\mathbf{z}^{(r)})$, $\mathbf{H}(\mathbf{z}^{(r)})$ and $\mathbf{U}(\mathbf{z}^{(r)})$, respectively.

We first show $\mathbf{z}^{(r+1)}$ stays in \mathcal{R}'_3 . From proof of Proposition 3.6, we know that for all $\mathbf{z} \in \mathcal{R}_3$, the following estimate holds:

$$\|\nabla f(\mathbf{z})\| \geq \frac{1}{4\sqrt{2}} \|\mathbf{x}\|^2 \|\mathbf{h}(\mathbf{z})\|.$$

From Proposition 3.8, we know that

$$\|\nabla f(\mathbf{z}^{(r+1)})\| \leq \frac{1}{m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|\nabla f(\mathbf{z}^{(r)})\|^2$$

provided $\Delta \leq \|\mathbf{x}\|/10$. Moreover,

$$\|\nabla f(\mathbf{z}^{(r)})\|^2 = \frac{1}{2} \left\| \mathbf{g}^{(r)} \right\|^2 \leq M_H^2 \left\| (\mathbf{H}^{(r)})^{-1} \mathbf{g}^{(r)} \right\|^2 \leq M_H^2 \Delta^2,$$

where the last inequality followed because step r is unconstrained. Combining the above estimates, we obtain that

$$\|\nabla f(\mathbf{z}^{(r+1)})\| \leq \frac{1}{m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) M_H^2 \Delta^2.$$

Thus,

$$\|\mathbf{h}(\mathbf{z}^{(r+1)})\| \leq \frac{4\sqrt{2}}{\|\mathbf{x}\|^2} \|\nabla f(\mathbf{z}^{(r+1)})\| \leq \frac{4\sqrt{2}}{m_H^2 \|\mathbf{x}\|^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) M_H^2 \Delta^2.$$

So, provided

$$\frac{4\sqrt{2}}{m_H^2 \|\mathbf{x}\|^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) M_H^2 \Delta^2 \leq \frac{1}{10L_h} \|\mathbf{x}\|^2,$$

$\mathbf{z}^{(r+1)}$ stays in \mathcal{R}'_3 .

Next we show the next step will also be an unconstrained step when Δ is sufficiently small. We have

$$\begin{aligned} & \|(\mathbf{H}^{(r+1)})^{-1} \mathbf{g}^{(r+1)}\| \\ & \leq \frac{1}{m_H} \|\mathbf{g}^{(r+1)}\| = \frac{\sqrt{2}}{m_H} \|\nabla f(\mathbf{z}^{(r+1)})\| \\ & \leq \frac{\sqrt{2}}{m_H^3} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|\nabla f(\mathbf{z}^{(r)})\|^2 = \frac{1}{\sqrt{2}m_H^3} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|\mathbf{g}^{(r)}\|^2 \\ & \leq \frac{M_H^2}{\sqrt{2}m_H^3} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|(\mathbf{H}^{(r)})^{-1} \mathbf{g}^{(r)}\|^2 \leq \frac{M_H^2}{\sqrt{2}m_H^3} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \Delta^2, \end{aligned}$$

where we again applied results of Proposition 3.8 to obtain the third line, and applied the optimality condition to obtain the fourth line. Thus, whenever

$$\frac{M_H^2}{\sqrt{2}m_H^3} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \Delta < 1,$$

the transformed trust-region subproblem has its minimizer $\boldsymbol{\xi}^{(r+1)}$ with $\|\boldsymbol{\xi}^{(r+1)}\| < \Delta$. This implies the minimizer $\boldsymbol{\delta}^{(r+1)}$ to the original trust-region subproblem satisfies $\boldsymbol{\delta}^{(r+1)} < \Delta$, as $\|\boldsymbol{\delta}^{(r+1)}\| = \|\boldsymbol{\xi}^{(r+1)}\|$. Thus, under the above condition the $(r+1)$ -th step is also unconstrained.

Repeating the above arguments for all future steps implies that all future steps will be constrained within \mathcal{R}'_3 .

We next provide an explicit estimate for the rate of convergence in terms of distance of the iterate to the target set \mathcal{X} . Again by Proposition 3.8,

$$\begin{aligned} \|\nabla f(\mathbf{z}^{(r+r')})\| & \leq m_H^2 \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right)^{-1} \left(\frac{1}{m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|\nabla f(\mathbf{z}^{(r)})\| \right)^{2^{r'}} \\ & \leq m_H^2 \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right)^{-1} \left(\frac{1}{\sqrt{2}m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \|\mathbf{g}^{(r)}\| \right)^{2^{r'}} \end{aligned}$$

$$\leq m_H^2 \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right)^{-1} \left(\frac{M_H}{\sqrt{2}m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \Delta \right)^{2r'}.$$

Thus, provided

$$\frac{M_H}{\sqrt{2}m_H^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right) \Delta \leq \frac{1}{2},$$

we have

$$\left\| \mathbf{h}(\mathbf{z}^{(r+r')}) \right\| \leq \frac{4\sqrt{2}}{\|\mathbf{x}\|^2} \|\nabla f(\mathbf{z}^{(r+r')})\| \leq \frac{4\sqrt{2}m_H^2}{\|\mathbf{x}\|^2} \left(L_h + \frac{32}{\|\mathbf{x}\|} M_H \right)^{-1} 2^{-2r'},$$

as claimed. ■

Acknowledgement. This work was partially supported by funding from the Gordon and Betty Moore Foundation, the Alfred P. Sloan Foundation, and the grants ONR N00014-13-1-0492, NSF CCF 1527809, and NSF IIS 1546411. We thank Nicolas Boumal for helpful discussion related to the Manopt package. We thank Mahdi Soltanolkotabi for pointing us to his early result on the local convexity around the target set for GPR in \mathbb{R}^n . We also thank Yonina Eldar, Kishore Jaganathan, Xiaodong Li for helpful feedback on a prior version of this paper.

Appendices

A Basic Tools and Results

Lemma A.1 (Even Moments of Complex Gaussian) *For $a \sim \mathcal{CN}(1)$, it holds that*

$$\mathbb{E} \left[|a|^{2p} \right] = p! \quad \forall p \in \mathbb{N}.$$

Proof Write $a = x + iy$, then $x, y \sim_{i.i.d.} \mathcal{N}(0, 1/2)$. Thus,

$$\mathbb{E} \left[|a|^{2p} \right] = \mathbb{E}_{x,y} \left[(x^2 + y^2)^p \right] = \frac{1}{2^p} \mathbb{E}_{z \sim \chi^2(2)} \left[z^p \right] = \frac{1}{2^p} 2^p p! = p!,$$

as claimed. ■

Lemma A.2 (Integral Form of Taylor's Theorem) *Consider any continuous function $f(\mathbf{z}) : \mathbb{C}^n \mapsto \mathbb{R}$ with continuous first- and second-order Wirtinger derivatives. For any $\delta \in \mathbb{C}^n$ and scalar $t \in \mathbb{R}$, we have*

$$f(\mathbf{z} + t\delta) = f(\mathbf{z}) + t \int_0^1 \begin{bmatrix} \delta \\ \bar{\delta} \end{bmatrix}^* \nabla f(\mathbf{z} + s\delta) ds,$$

$$f(\mathbf{z} + t\delta) = f(\mathbf{z}) + t \begin{bmatrix} \delta \\ \bar{\delta} \end{bmatrix}^* \nabla f(\mathbf{z}) + t^2 \int_0^1 (1-s) \begin{bmatrix} \delta \\ \bar{\delta} \end{bmatrix}^* \nabla^2 f(\mathbf{z} + s\delta) \begin{bmatrix} \delta \\ \bar{\delta} \end{bmatrix} ds.$$

Proof Since f is continuous differentiable, by the fundamental theorem of calculus,

$$f(\mathbf{z} + t\boldsymbol{\delta}) = f(\mathbf{z}) + \int_0^t \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z} + \tau\boldsymbol{\delta}) d\tau.$$

Moreover, by integral by part, we obtain

$$\begin{aligned} f(\mathbf{z} + t\boldsymbol{\delta}) &= f(\mathbf{z}) + \left[(\tau - t) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z} + \tau\boldsymbol{\delta}) \right] \Big|_0^t - \int_0^t (\tau - t) d \left[\begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z} + \tau\boldsymbol{\delta}) \right] \\ &= f(\mathbf{z}) + t \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z}) + \int_0^t (t - \tau) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z} + \tau\boldsymbol{\delta}) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix} d\tau. \end{aligned}$$

Change of variable $\tau = st$ ($0 \leq s \leq 1$) gives the claimed result. ■

Lemma A.3 (Error of Quadratic Approximation) Consider any continuous function $f(\mathbf{z}) : \mathbb{C}^n \mapsto \mathbb{R}$ with continuous first- and second-order Wirtinger derivatives. Suppose its Hessian $\nabla^2 f(\mathbf{z})$ is L_h -Lipschitz. Then the second-order approximation

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{z}) = f(\mathbf{z}) + \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla f(\mathbf{z}) + \frac{1}{2} \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* \nabla^2 f(\mathbf{z}) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}$$

around each point \mathbf{z} obeys

$$\left| f(\mathbf{z} + \boldsymbol{\delta}) - \widehat{f}(\boldsymbol{\delta}; \mathbf{z}) \right| \leq \frac{1}{3} L_h \|\boldsymbol{\delta}\|^3.$$

Proof By integral form of Taylor's theorem in Lemma A.2,

$$\begin{aligned} \left| f(\mathbf{z} + \boldsymbol{\delta}) - \widehat{f}(\boldsymbol{\delta}; \mathbf{z}) \right| &= \left| \int_0^1 (1 - \tau) \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix}^* [\nabla^2 f(\mathbf{z} + \tau\boldsymbol{\delta}) - \nabla^2 f(\mathbf{z})] \begin{bmatrix} \boldsymbol{\delta} \\ \bar{\boldsymbol{\delta}} \end{bmatrix} d\tau \right| \\ &\leq 2 \|\boldsymbol{\delta}\|^2 \int_0^1 (1 - \tau) \|\nabla^2 f(\mathbf{z} + \tau\boldsymbol{\delta}) - \nabla^2 f(\mathbf{z})\| d\tau \\ &\leq 2L_h \|\boldsymbol{\delta}\|^3 \int_0^1 (1 - \tau)\tau d\tau = \frac{L_h}{3} \|\boldsymbol{\delta}\|^3, \end{aligned}$$

as desired. ■

Lemma A.4 (Spectrum of Complex Gaussian Matrices) Let \mathbf{X} be an $n_1 \times n_2$ ($n_1 > n_2$) matrices with i.i.d. \mathcal{CN} entries. Then,

$$\sqrt{n_1} - \sqrt{n_2} \leq \mathbb{E}[\sigma_{\min}(\mathbf{X})] \leq \mathbb{E}[\sigma_{\max}(\mathbf{X})] \leq \sqrt{n_1} + \sqrt{n_2}.$$

Moreover, for each $t \geq 0$, it holds with probability at least $1 - 2 \exp(-t^2)$ that

$$\sqrt{n_1} - \sqrt{n_2} - t \leq \sigma_{\min}(\mathbf{X}) \leq \sigma_{\max}(\mathbf{X}) \leq \sqrt{n_1} + \sqrt{n_2} + t.$$

Lemma A.5 (Hoeffding-type Inequality, Proposition 5.10 of [Ver12]) Let X_1, \dots, X_N be independent centered sub-Gaussian random variables, and let $K = \max_i \|X_i\|_{\psi_2}$, where the sub-Gaussian norm

$$\|X_i\|_{\psi_2} \doteq \sup_{p \geq 1} p^{-1/2} (\mathbb{E} [|X|^p])^{1/p}. \quad (\text{A.1})$$

Then for every $\mathbf{b} = [b_1; \dots; b_N] \in \mathbb{C}^N$ and every $t \geq 0$, we have

$$\mathbb{P} \left(\left| \sum_{k=1}^N b_k X_k \right| \geq t \right) \leq e \exp \left(-\frac{ct^2}{K^2 \|\mathbf{b}\|_2^2} \right). \quad (\text{A.2})$$

Here c is a universal constant.

Lemma A.6 (Bernstein-type Inequality, Proposition 5.17 of [Ver12]) Let X_1, \dots, X_N be independent centered sub-exponential random variables, and let $K = \max_i \|X_i\|_{\psi_1}$, where the sub-exponential norm

$$\|X_i\|_{\psi_1} \doteq \sup_{p \geq 1} p^{-1} (\mathbb{E} [|X|^p])^{1/p}. \quad (\text{A.3})$$

Then for every $\mathbf{b} = [b_1; \dots; b_N] \in \mathbb{C}^N$ and every $t \geq 0$, we have

$$\mathbb{P} \left(\left| \sum_{k=1}^N b_k X_k \right| \geq t \right) \leq 2 \exp \left(-c \min \left(\frac{t^2}{K^2 \|\mathbf{b}\|_2^2}, \frac{t}{K \|\mathbf{b}\|_\infty} \right) \right). \quad (\text{A.4})$$

Here c is a universal constant.

Lemma A.7 (Subgaussian Lower Tail for Nonnegative RV's, Problem 2.9 of [BLM13]) Let X_1, \dots, X_N be i.i.d. copies of the nonnegative random variable X with finite second moment. Then it holds that

$$\mathbb{P} \left[\frac{1}{N} \sum_{i=1}^N (X_i - \mathbb{E}[X_i]) < -t \right] \leq \exp \left(-\frac{Nt^2}{2\sigma^2} \right)$$

for any $t > 0$, where $\sigma^2 = \mathbb{E}[X^2]$.

Proof For any $\lambda > 0$, we have

$$\log \mathbb{E} \left[e^{-\lambda(X - \mathbb{E}[X])} \right] = \lambda \mathbb{E}[X] + \log \mathbb{E} \left[e^{-\lambda X} \right] \leq \lambda \mathbb{E}[X] + \mathbb{E} \left[e^{-\lambda X} \right] - 1,$$

where the last inequality holds thanks to $\log u \leq u - 1$ for all $u > 0$. Moreover, using the fact $e^u \leq 1 + u + u^2/2$ for all $u \leq 0$, we obtain

$$\log \mathbb{E} \left[e^{-\lambda(X - \mathbb{E}[X])} \right] \leq \frac{1}{2} \lambda^2 \mathbb{E}[X^2] \iff \mathbb{E} \left[e^{-\lambda(X - \mathbb{E}[X])} \right] \leq \exp \left(\frac{1}{2} \lambda^2 \mathbb{E}[X^2] \right).$$

Thus, by the usual exponential transform trick, we obtain that for any $t > 0$,

$$\mathbb{P} \left[\sum_{i=1}^N (X_i - \mathbb{E}[X_i]) < -t \right] \leq \exp \left(-\lambda t + N \lambda^2 \mathbb{E}[X^2] / 2 \right).$$

Taking $\lambda = t/(N\sigma^2)$ and making change of variable for t give the claimed result. \blacksquare

Lemma A.8 (Moment-Control Bernstein's Inequality for Random Variables) Let X_1, \dots, X_p be i.i.d. copies of a real-valued random variable X . Suppose that there exist some positive number R and σ_X^2 such that

$$\mathbb{E}[X^2] \leq \sigma_X^2, \quad \text{and} \quad \mathbb{E}[|X|^m] \leq \frac{m!}{2} \sigma_X^2 R^{m-2}, \quad \text{for all integers } m \geq 3.$$

Let $S \doteq \frac{1}{p} \sum_{k=1}^p X_k$, then for ... , it holds that

$$\mathbb{P}[|S - \mathbb{E}[S]| \geq t] \leq 2 \exp\left(-\frac{pt^2}{2\sigma_X^2 + 2Rt}\right).$$

Lemma A.9 (Angles Between Two Subspaces) Consider two linear subspaces \mathcal{U}, \mathcal{V} of dimension k in \mathbb{R}^n ($k \in [n]$) spanned by orthonormal bases \mathbf{U} and \mathbf{V} , respectively. Suppose $\pi/2 \geq \theta_1 \geq \theta_2 \geq \dots \geq \theta_k \geq 0$ are the principal angles between \mathcal{U} and \mathcal{V} . Then it holds that

i) $\min_{\mathbf{Q} \in O_k} \|\mathbf{U} - \mathbf{V}\mathbf{Q}\| \leq \sqrt{2 - 2\cos\theta_1}$;

ii) $\sin\theta_1 = \|\mathbf{U}\mathbf{U}^* - \mathbf{V}\mathbf{V}^*\|$;

iii) Let \mathcal{U}^\perp and \mathcal{V}^\perp be the orthogonal complement of \mathcal{U} and \mathcal{V} , respectively. Then $\theta_1(\mathcal{U}, \mathcal{V}) = \theta_1(\mathcal{U}^\perp, \mathcal{V}^\perp)$.

Proof Proof to i) is similar to that of II. Theorem 4.11 in [SS90]. For $2k \leq n$, w.l.o.g., we can assume \mathbf{U} and \mathbf{V} are the canonical bases for \mathcal{U} and \mathcal{V} , respectively. Then

$$\min_{\mathbf{Q} \in O_k} \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma}\mathbf{Q} \\ -\mathbf{\Sigma}\mathbf{Q} \\ \mathbf{0} \end{bmatrix} \right\| \leq \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma} \\ -\mathbf{\Sigma} \\ \mathbf{0} \end{bmatrix} \right\| \leq \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma} \\ -\mathbf{\Sigma} \end{bmatrix} \right\|.$$

Now by definition

$$\begin{aligned} \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma} \\ -\mathbf{\Sigma} \end{bmatrix} \right\|^2 &= \max_{\|\mathbf{x}\|=1} \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma} \\ -\mathbf{\Sigma} \end{bmatrix} \mathbf{x} \right\|^2 = \max_{\|\mathbf{x}\|=1} \sum_{i=1}^k (1 - \cos\theta_i)^2 x_i^2 + \sin^2\theta_i x_i^2 \\ &= \max_{\|\mathbf{x}\|=1} \sum_{i=1}^k (2 - 2\cos\theta_i) x_i^2 \leq 2 - 2\cos\theta_1. \end{aligned}$$

Note that the upper bound is achieved by taking $\mathbf{x} = \mathbf{e}_1$. When $2k > n$, by the results from CS decomposition (see, e.g., I Theorem 5.2 of [SS90]).

$$\min_{\mathbf{Q} \in O_k} \left\| \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{\Gamma} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \\ \mathbf{\Sigma} & \mathbf{0} \end{bmatrix} \right\| \leq \left\| \begin{bmatrix} \mathbf{I} - \mathbf{\Gamma} \\ -\mathbf{\Sigma} \end{bmatrix} \right\|,$$

and the same argument then carries through. To prove ii), note the fact that $\sin\theta_1 = \|\mathbf{U}\mathbf{U}^* - \mathbf{V}\mathbf{V}^*\|$ (see, e.g., Theorem 4.5 and Corollary 4.6 of [SS90]). Obviously one also has

$$\sin\theta_1 = \|\mathbf{U}\mathbf{U}^* - \mathbf{V}\mathbf{V}^*\| = \|(\mathbf{I} - \mathbf{U}\mathbf{U}^*) - (\mathbf{I} - \mathbf{V}\mathbf{V}^*)\|,$$

while $\mathbf{I} - \mathbf{U}\mathbf{U}^*$ and $\mathbf{I} - \mathbf{V}\mathbf{V}^*$ are projectors onto \mathcal{U}^\perp and \mathcal{V}^\perp , respectively. This completes the proof. ■

References

- [AAJ⁺13] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, Praneeth Netrapalli, and Rashish Tandon, *Learning sparsely used overcomplete dictionaries via alternating minimization*, arXiv preprint arXiv:1310.7991 (2013).
- [AAN13] Alekh Agarwal, Animashree Anandkumar, and Praneeth Netrapalli, *Exact recovery of sparsely used overcomplete dictionaries*, arXiv preprint arXiv:1309.1952 (2013).
- [ABFM14] Boris Alexeev, Afonso S. Bandeira, Matthew Fickus, and Dustin G. Mixon, *Phase retrieval with polarization*, SIAM Journal on Imaging Sciences 7 (2014), no. 1, 35–66.
- [ABG07] Pierre-Antoine. Absil, Christopher G. Baker, and Kyle A. Gallivan, *Trust-region methods on Riemannian manifolds*, Foundations of Computational Mathematics 7 (2007), no. 3, 303–330.
- [ABGM14] Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma, *More algorithms for provable dictionary learning*, arXiv preprint arXiv:1401.0579 (2014).
- [AG16] Anima Anandkumar and Rong Ge, *Efficient approaches for escaping higher order saddle points in non-convex optimization*, arXiv preprint arXiv:1602.05908 (2016).
- [AGJ14a] Animashree Anandkumar, Rong Ge, and Majid Janzamin, *Analyzing tensor power method dynamics: Applications to learning overcomplete latent variable models*, arXiv preprint arXiv:1411.1488 (2014).
- [AGJ14b] ———, *Guaranteed non-orthogonal tensor decomposition via alternating rank-1 updates*, arXiv preprint arXiv:1402.5180 (2014).
- [AGM13] Sanjeev Arora, Rong Ge, and Ankur Moitra, *New algorithms for learning incoherent and overcomplete dictionaries*, arXiv preprint arXiv:1308.6273 (2013).
- [AGMM15] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra, *Simple, efficient, and neural algorithms for sparse coding*, arXiv preprint arXiv:1503.00778 (2015).
- [AJSN15] Animashree Anandkumar, Prateek Jain, Yang Shi, and Uma Naresh Niranjan, *Tensor vs matrix methods: Robust tensor decomposition under block sparse perturbations*, arXiv preprint arXiv:1510.04747 (2015).
- [AMS09] Pierre-Antoine. Absil, Robert Mahoney, and Rodolphe Sepulchre, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2009.
- [Bal10] Radu V. Balan, *On signal reconstruction from its spectrogram*, Information Sciences and Systems (CISS), 44th Annual Conference on, IEEE, 2010, pp. 1–4.
- [BBCE09] Radu Balan, Bernhard G. Bodmann, Peter G. Casazza, and Dan Edidin, *Painless reconstruction from magnitudes of frame coefficients*, Journal of Fourier Analysis and Applications 15 (2009), no. 4, 488–501.
- [BCE06] Radu Balana, Pete Casazzab, and Dan Edidin, *On signal reconstruction without phase*, Applied and Computational Harmonic Analysis 20 (2006), no. 3, 345 – 356.
- [BDP⁺07] Oliver Bunk, Ana Diaz, Franz Pfeiffer, Christian David, Bernd Schmitt, Dillip K. Satapathy, and J. Friso van der Veen, *Diffraction imaging for periodic samples: retrieving one-dimensional concentration profiles across microfluidic channels*, Acta Crystallographica Section A 63 (2007), no. 4, 306–314.
- [Ber99] Dimitri P. Bertsekas, *Nonlinear programming*.
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford University Press, 2013.
- [BMAS14] Nicolas Boumal, Bamdev Mishra, P.-A. Absil, and Rodolphe Sepulchre, *Manopt, a Matlab toolbox for optimization on manifolds*, Journal of Machine Learning Research 15 (2014), 1455–1459.

- [Bou16] Nicolas Boumal, *Nonconvex phase synchronization*, arXiv preprint arXiv:1601.06114 (2016).
- [CC15] Yuxin Chen and Emmanuel J. Candès, *Solving random quadratic systems of equations is nearly as easy as solving linear systems*, arXiv preprint arXiv:1505.05114 (2015).
- [CESV13] Emmanuel J. Candès, Yonina C. Eldar, Thomas Strohmer, and Vladislav Voroninski, *Phase retrieval via matrix completion*, *SIAM Journal on Imaging Sciences* **6** (2013), no. 1.
- [CGT00] Andrew R. Conn, Nicholas I.M. Gould, and Philippe L. Toint, *Trust region methods*, vol. 1, SIAM, 2000.
- [CL14] Emmanuel J. Candès and Xiaodong Li, *Solving quadratic equations via phaselift when there are about as many equations as unknowns*, *Foundations of Computational Mathematics* **14** (2014), no. 5, 1017–1026.
- [CLM15] T. Tony Cai, Xiaodong Li, and Zongming Ma, *Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow*, arXiv preprint arXiv:1506.03382 (2015).
- [CLS15a] Emmanuel J. Candès, Xiaodong Li, and Mahdi Soltanolkotabi, *Phase retrieval from coded diffraction patterns*, *Applied and Computational Harmonic Analysis* **39** (2015), no. 2, 277–299.
- [CLS15b] ———, *Phase retrieval via wirtinger flow: Theory and algorithms*, *Information Theory, IEEE Transactions on* **61** (2015), no. 4, 1985–2007.
- [CMP11] Anwei Chai, Miguel Moscoso, and George Papanicolaou, *Array imaging using intensity-only measurements*, *Inverse Problems* **27** (2011), no. 1, 015005.
- [Cor06] John V. Corbett, *The pauli problem, state reconstruction and quantum-real numbers*, *Reports on Mathematical Physics* **57** (2006), no. 1, 53–68.
- [CSV13] Emmanuel J. Candès, Thomas Strohmer, and Vladislav Voroninski, *Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming*, *Communications on Pure and Applied Mathematics* **66** (2013), no. 8, 1241–1274.
- [CW15] Yudong Chen and Martin J. Wainwright, *Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees*, arXiv preprint arXiv:1509.03025 (2015).
- [DF87] Chris Dainty and James R. Fienup, *Phase retrieval and image reconstruction for astronomy*, *Image Recovery: Theory and Application* (1987), 231–275.
- [EW15] Armin Eftekhari and Michael B. Wakin, *Greed is super: A fast algorithm for super-resolution*, arXiv preprint arXiv:1511.03385 (2015).
- [Fie82] James R. Fienup, *Phase retrieval algorithms: a comparison*, *Applied Optics* **21** (1982), no. 15, 2758–2769.
- [GHJY15] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan, *Escaping from saddle points—online stochastic gradient for tensor decomposition*, *Proceedings of The 28th Conference on Learning Theory*, 2015, pp. 797–842.
- [GKK13] David Gross, Felix Kraemer, and Richard Kueng, *A partial derandomization of phaselift using spherical designs*, arXiv preprint arXiv:1310.2267 (2013).
- [Gol80] Donald Goldfarb, *Curvilinear path steplength algorithms for minimization which use directions of negative curvature*, *Mathematical programming* **18** (1980), no. 1, 31–40.
- [GS72] R. W. Gerchberg and W. Owen Saxton, *A practical algorithm for the determination of the phase from image and diffraction plane pictures*, *Optik* **35** (1972), 237–246.
- [Har14] Moritz Hardt, *Understanding alternating minimization for matrix completion*, *Foundations of Computer Science (FOCS)*, 2014 IEEE 55th Annual Symposium on, IEEE, 2014, pp. 651–660.

- [HMW13] Teiko Heinosaari, Luca Mazzarella, and Michael M. Wolf, *Quantum tomography under prior information*, *Communications in Mathematical Physics* **318** (2013), no. 2, 355–374.
- [HSS15] Samuel B. Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer, *Speeding up sum-of-squares for tensor decomposition and planted sparse vectors*, arXiv preprint arXiv:1512.02337 (2015).
- [HW14] Moritz Hardt and Mary Wootters, *Fast matrix completion without the condition number*, *Proceedings of The 27th Conference on Learning Theory*, 2014, pp. 638–678.
- [JEH15] Kishore Jaganathan, Yonina C. Eldar, and Babak Hassibi, *Phase retrieval: An overview of recent developments*, arXiv preprint arXiv:1510.07713 (2015).
- [JKN15] Prateek Jain, Chi Jin, Sham M. Kakade, and Praneeth Netrapalli, *Computing matrix squareroot via non convex local search*, arXiv preprint arXiv:1507.05854 (2015).
- [JN14] Prateek Jain and Praneeth Netrapalli, *Fast exact matrix completion with finite samples*, arXiv preprint arXiv:1411.1087 (2014).
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi, *Low-rank matrix completion using alternating minimization*, *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*, ACM, 2013, pp. 665–674.
- [JO14] Prateek Jain and Sewoong Oh, *Provable tensor factorization with missing data*, *Advances in Neural Information Processing Systems*, 2014, pp. 1431–1439.
- [JOH13] Kishore Jaganathan, Samet Oymak, and Babak Hassibi, *Sparse phase retrieval: Convex algorithms and limitations*, *Proceedings of IEEE International Symposium on Information Theory*, IEEE, 2013, pp. 1022–1026.
- [KD09] Ken Kreutz-Delgado, *The complex gradient operator and the $\mathbb{C}\mathbb{R}$ -calculus*, arXiv preprint arXiv:0906.4835 (2009).
- [KMO10] Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh, *Matrix completion from a few entries*, *Information Theory, IEEE Transactions on* **56** (2010), no. 6, 2980–2998.
- [LJ15] Kiryung Lee and Marius Junge, *RIP-like properties in subsampled blind deconvolution*, arXiv preprint arXiv:1511.06146 (2015).
- [LLJB15] Kiryung Lee, Yanjun Li, Marius Junge, and Yoram Bresler, *Blind recovery of sparse signals from subsampled convolution*, arXiv preprint arXiv:1511.06149 (2015).
- [LSJR16] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht, *Gradient descent converges to minimizers*, arXiv preprint arXiv:1602.04915 (2016).
- [LV13] Xiaodong Li and Vladislav Voroninski, *Sparse signal recovery from quadratic measurements via convex programming*, *SIAM Journal on Mathematical Analysis* **45** (2013), no. 5, 3019–3033.
- [LWB13] Kiryung Lee, Yihong Wu, and Yoram Bresler, *Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization*, arXiv preprint arXiv:1312.0525 (2013).
- [MHWG14] Cun Mu, Bo Huang, John Wright, and Donald Goldfarb, *Square deal: Lower bounds and improved convex relaxations for tensor recovery.*, *Journal of Machine Learning Research* **1** (2014), 1–48.
- [MIJ⁺02] Jianwei Miao, Tetsuya Ishikawa, Bart Johnson, Erik H. Anderson, Barry Lai, and Keith O. Hodgson, *High resolution 3D X-Ray diffraction microscopy*, *Phys. Rev. Lett.* **89** (2002), no. 8, 088303.
- [Mil90] R. P. Millane, *Phase retrieval in crystallography and optics*, *Journal of the Optical Society of America A* **7** (1990), no. 3, 394–411.
- [NJS13] Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi, *Phase retrieval using alternating minimization*, *Advances in Neural Information Processing Systems*, 2013, pp. 2796–2804.

- [NNS⁺14] Praneeth Netrapalli, Uma Naresh. Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain, *Non-convex robust PCA*, Advances in Neural Information Processing Systems, 2014, pp. 1107–1115.
- [NP06] Yurii Nesterov and Boris T. Polyak, *Cubic regularization of newton method and its global performance*, Mathematical Programming **108** (2006), no. 1, 177–205.
- [OJF⁺12] Samet Oymak, Amin Jalali, Maryam Fazel, Yonina C. Eldar, and Babak Hassibi, *Simultaneously structured models with application to sparse and low-rank matrices*, arXiv preprint arXiv:1212.3753 (2012).
- [OYDS12] Henrik Ohlsson, Allen Y. Yang, Roy Dong, and S. Shankar Sastry, *CPRL – An extension of compressive sensing to the phase retrieval problem*, Advances in Neural Information Processing Systems, 2012.
- [OYDS13] ———, *Compressive phase retrieval from squared output measurements via semidefinite programming*, arXiv preprint arXiv:1111.6323 (2013).
- [OYVS13] Henrik Ohlsson, Allen Y. Yang, Michel Verhaegen, and S. Shankar Sastry, *Quadratic basis pursuit*, arXiv preprint arXiv:1301.7002 (2013).
- [QSW14] Qing Qu, Ju Sun, and John Wright, *Finding a sparse vector in a subspace: Linear sparsity using alternating directions*, Advances in Neural Information Processing Systems, 2014, pp. 3401–3409.
- [Rei65] H. Reichenbach, *Philosophic foundations of quantum mechanics*, University of California Press, 1965.
- [Rob93] W. Harrison Robert, *Phase problem in crystallography*, Journal of the Optical Society of America A **10** (1993), no. 5, 1046–1055.
- [SA14] Hanie Sedghi and Animashree Anandkumar, *Provable tensor methods for learning mixtures of classifiers*, arXiv preprint arXiv:1412.3046 (2014).
- [SBE14] Yoav Shechtman, Amir Beck, and Yonina C. Eldar, *GESPAR: Efficient phase retrieval of sparse signals*, Signal Processing, IEEE Transactions on **62** (2014), no. 4, 928–938.
- [SEC⁺15] Yoav Shechtman, Yonina C. Eldar, Oren Cohen, Henry N. Chapman, Jianwei Miao, and Mordechai Segev, *Phase retrieval with application to optical imaging: A contemporary overview*, Signal Processing Magazine, IEEE **32** (2015), no. 3, 87–109.
- [SL14] Ruoyu Sun and Zhi-Quan Luo, *Guaranteed matrix completion via non-convex factorization*, arXiv preprint arXiv:1411.8003 (2014).
- [Sol14] Mahdi Soltanolkotabi, *Algorithms and theory for clustering and nonconvex quadratic programming*, Ph.D. thesis, Stanford University, 2014.
- [SQW15a] Ju Sun, Qing Qu, and John Wright, *Complete dictionary recovery over the sphere*, arXiv preprint arXiv:1504.06785 (2015).
- [SQW15b] ———, *When are nonconvex problems not scary?*, arXiv preprint arXiv:1510.06096 (2015).
- [SRO15] Christopher De Sa, Christopher Re, and Kunle Olukotun, *Global convergence of stochastic gradient descent for some non-convex matrix problems*, The 32nd International Conference on Machine Learning, vol. 37, 2015, pp. 2332–2341.
- [SS90] Gilbert W. Stewart and Ji-guang Sun, *Matrix perturbation theory*, Academic press, 1990.
- [TBSR15] Stephen Tu, Ross Boczar, Mahdi Soltanolkotabi, and Benjamin Recht, *Low-rank solutions of linear matrix equations via procrustes flow*, arXiv preprint arXiv:1507.03566 (2015).

- [Ver12] Roman Vershynin, *Introduction to the non-asymptotic analysis of random matrices*, Compressed Sensing (Yonina C. Eldar and Gitta Kutyniok, eds.), Cambridge University Press, 2012, pp. 210–268.
- [VX14] Vladislav Voroninski and Zhiqiang Xu, *A strong restricted isometry property, with an application to phaseless compressed sensing*, arXiv preprint arXiv:1404.3811 (2014).
- [Wal63] Adriaan Walther, *The question of phase retrieval in optics*, Journal of Modern Optics **10** (1963), no. 1, 41–49.
- [WCCL15] Ke Wei, Jian-Feng Cai, Tony F. Chan, and Shingyu Leung, *Guarantees of Riemannian optimization for low rank matrix recovery*, arXiv preprint arXiv:1511.01562 (2015).
- [WdM15] Irène Waldspurger, Alexandre d`Aspremont, and Stéphane Mallat, *Phase recovery, maxcut and complex semidefinite programming*, Mathematical Programming **149** (2015), no. 1-2, 47–81.
- [WWS15] Chris D. White, Rachel Ward, and Sujay Sanghavi, *The local convexity of solving quadratic equations*, arXiv preprint arXiv:1506.07868 (2015).
- [YCS13] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi, *Alternating minimization for mixed linear regression*, arXiv preprint arXiv:1310.3745 (2013).
- [ZL15] Qinqing Zheng and John Lafferty, *A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements*, arXiv preprint arXiv:1506.06081 (2015).