

Solving Random Quadratic Systems of Equations Is Nearly as Easy as Solving Linear Systems

YUXIN CHEN

Department of Statistics, Stanford University

EMMANUEL J. CANDÈS

Departments of Mathematics and of Statistics, Stanford University

Abstract

We consider the fundamental problem of solving quadratic systems of equations in n variables, where $y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2$, $i = 1, \dots, m$, and $\mathbf{x} \in \mathbb{R}^n$ is unknown. We propose a novel method, which starts with an initial guess computed by means of a spectral method and proceeds by minimizing a nonconvex functional as in the Wirtinger flow approach [13]. There are several key distinguishing features, most notably a distinct objective functional and novel update rules, which operate in an adaptive fashion and drop terms bearing too much influence on the search direction. These careful selection rules provide a tighter initial guess, better descent directions, and thus enhanced practical performance. On the theoretical side, we prove that for certain unstructured models of quadratic systems, our algorithms return the correct solution in linear time, i.e., in time proportional to reading the data $\{\mathbf{a}_i\}$ and $\{y_i\}$ as soon as the ratio m/n between the number of equations and unknowns exceeds a fixed numerical constant. We extend the theory to deal with noisy systems in which we only have $y_i \approx |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2$ and prove that our algorithms achieve a statistical accuracy, which is nearly unimprovable. We complement our theoretical study with numerical examples showing that solving random quadratic systems is both computationally and statistically not much harder than solving linear systems of the same size—hence the title of this paper. For instance, we demonstrate empirically that the computational cost of our algorithm is about four times that of solving a least squares problem of the same size. © 2016 Wiley Periodicals, Inc.

1 Introduction

1.1 Problem Formulation

Imagine we are given a set of m quadratic equations taking the form

$$(1.1) \quad y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2, \quad i = 1, \dots, m,$$

where the data $\mathbf{y} = [y_i]_{1 \leq i \leq m}$ and design vectors $\mathbf{a}_i \in \mathbb{R}^n / \mathbb{C}^n$ are known whereas $\mathbf{x} \in \mathbb{R}^n / \mathbb{C}^n$ is unknown. Having information about $|\langle \mathbf{a}_i, \mathbf{x} \rangle|^2$ —or, equivalently, $|\langle \mathbf{a}_i, \mathbf{x} \rangle|$ —means that we a priori know nothing about the phases or signs of the

linear products $\langle \mathbf{a}_i, \mathbf{x} \rangle$. The problem is this: can we hope to identify a solution, if any, compatible with this nonlinear system of equations?

This problem is combinatorial in nature as one can alternatively pose it as recovering the missing signs of $\langle \mathbf{a}_i, \mathbf{x} \rangle$ from magnitude-only observations. As is well-known, many classical combinatorial problems with Boolean variables may be cast as special instances of (1.1). As an example, consider the NP-hard *stone problem* [8], in which we have n stones, each of weight $w_i > 0$ ($1 \leq i \leq n$), which we would like to divide into two groups of equal sum weight. Letting $x_i \in \{-1, 1\}$ indicate which of the two groups the i^{th} stone belongs to, one can formulate this problem as solving the following quadratic system:

$$(1.2) \quad \begin{cases} x_i^2 = 1, \\ (w_1 x_1 + \cdots + w_n x_n)^2 = 0. \end{cases} \quad i = 1, \dots, n,$$

However simple this formulation may seem, even checking whether a solution to (1.2) exists is known to be NP-hard.

Moving from combinatorial optimization to the physical sciences, one application of paramount importance is the *phase retrieval* [23, 24] problem, which permeates through a wide spectrum of techniques, including X-ray crystallography, diffraction imaging, microscopy, and even quantum mechanics. In a nutshell, the problem of phase retrieval arises due to the physical limitation of optical sensors, which are often only able to record the intensities of the diffracted waves scattered by an object under study. Notably, upon illuminating an object \mathbf{x} , the diffraction pattern is of the form of $\mathbf{A}\mathbf{x}$; however, it is only possible to obtain intensity measurements $\mathbf{y} = |\mathbf{A}\mathbf{x}|^2$ leading to the quadratic system (1.1).¹ In the Fraunhofer regime where data is collected in the far-field zone, \mathbf{A} is given by the spatial Fourier transform. We refer to [43] for in-depth reviews of this subject.

Continuing this motivating line of thought, recorded intensities in any real-world application are always corrupted by at least a small amount of noise so that observed data are only about $|\langle \mathbf{a}_i, \mathbf{x} \rangle|^2$; i.e.,

$$(1.3) \quad y_i \approx |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2, \quad i = 1, \dots, m.$$

Although we present results for arbitrary noise distributions—even for nonstochastic noise—we shall pay particular attention to the Poisson data model, which assumes

$$(1.4) \quad y_i \stackrel{\text{ind.}}{\sim} \text{Poisson}(|\langle \mathbf{a}_i, \mathbf{x} \rangle|^2), \quad i = 1, \dots, m.$$

The reason that this statistical model is of special interest is that it naturally describes the variation in the number of photons detected by an optical sensor in various imaging applications.

¹ Here and below, for $\mathbf{z} \in \mathbb{C}^n$, $|\mathbf{z}|$ (respectively, $|\mathbf{z}|^2$) represents the vector of magnitudes $(|z_1|, \dots, |z_n|)^T$ (respectively, squared magnitudes $(|z_1|^2, \dots, |z_n|^2)^T$).

1.2 Nonconvex Optimization

Under a stochastic noise model with independent samples, a first impulse for solving (1.3) is to seek the maximum likelihood estimate (MLE), namely,

$$(1.5) \quad \text{minimize}_{\mathbf{z}} \quad - \sum_{i=1}^m \ell(\mathbf{z}; y_i),$$

where $\ell(\mathbf{z}; y_i)$ denotes the log-likelihood of a candidate solution \mathbf{z} given the outcome y_i . For instance, under the Poisson data model (1.4) one can write

$$(1.6) \quad \ell(\mathbf{z}; y_i) = y_i \log(|\mathbf{a}_i^* \mathbf{z}|^2) - |\mathbf{a}_i^* \mathbf{z}|^2$$

modulo some constant offset. Unfortunately, the log-likelihood is usually not concave, thus making the problem of computing the MLE NP-hard in general.

To alleviate this computational intractability, several convex surrogates have been proposed that work particularly well when the design vectors $\{\mathbf{a}_i\}$ are chosen at random [3, 9, 10, 14, 16, 18, 20, 25, 26, 29, 32, 33, 39, 44, 52]. The basic idea is to introduce a rank 1 matrix $\mathbf{X} = \mathbf{x} \mathbf{x}^*$ to linearize the quadratic constraints, and then relax the rank 1 constraint. Suppose we have Poisson data; then this strategy converts the problem into a convenient convex program:

$$\begin{aligned} & \text{minimize}_{\mathbf{X}} \quad \sum_{i=1}^m (\mu_i - y_i \log \mu_i) + \lambda \text{Tr}(\mathbf{X}) \\ & \text{subject to} \quad \mu_i = \mathbf{a}_i^\top \mathbf{X} \mathbf{a}_i, \quad 1 \leq i \leq m, \\ & \quad \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned}$$

Note that the log-likelihood function is augmented by the trace functional $\text{Tr}(\cdot)$, the role of which is to promote low-rank solutions. While such convex relaxation schemes enjoy intriguing performance guarantees in many aspects (e.g., they achieve minimal sample complexity and near-optimal error bounds for certain noise models), the computational cost typically far exceeds the order of n^3 . This limits applicability to large-dimensional data.

This paper follows another route: rather than lifting the problem into higher dimensions by introducing matrix variables, this paradigm maintains its iterates within the vector domain and optimizes the nonconvex objective directly (e.g., [13, 22, 24, 34, 36, 40–42, 54, 56]). One promising approach along this line is the recently proposed two-stage algorithm called *Wirtinger flow* (WF) [13]. Simply put, WF starts by computing a suitable initial guess $\mathbf{z}^{(0)}$ using a spectral method, and then successively refines the estimate via an update rule that bears a strong resemblance to a gradient descent scheme, namely,

$$\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{\mu_t}{m} \sum_{i=1}^m \nabla \ell(\mathbf{z}^{(t)}; y_i),$$

where $\mathbf{z}^{(t)}$ denotes the t^{th} iterate of the algorithm, and μ_t is the step size (or learning rate). Here, $\nabla \ell(\mathbf{z}; y_i)$ stands for the Wirtinger derivative with respect to \mathbf{z} , which in the real-valued case reduces to the ordinary gradient. The main results of [13] demonstrate that WF is surprisingly accurate for independent Gaussian design. Specifically, when $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ or $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) + j\mathcal{N}(\mathbf{0}, \mathbf{I})$:

- (1) WF achieves exact recovery from $m = O(n \log n)$ quadratic equations when there is no noise.²
- (2) WF attains ϵ -accuracy—in a relative sense—within $O(mn^2 \log(1/\epsilon))$ time (or flops).
- (3) In the presence of Gaussian noise, WF is stable and converges to the MLE as shown in [45].

While these results formalize the advantages of WF, the computational complexity of WF is still much larger than the best one can hope for. Moreover, the statistical guarantee in terms of the sample complexity is weaker than that achievable by convex relaxations.³

1.3 Truncated Wirtinger Flow

This paper develops an efficient linear-time algorithm, which also enjoys near-optimal statistical guarantees. Following the spirit of WF, we propose a novel procedure called *truncated Wirtinger flow* (TWF) adopting a more adaptive gradient flow. Informally, TWF proceeds in two stages:

- (1) INITIALIZATION: compute an initial guess $\mathbf{z}^{(0)}$ by means of a spectral method applied to a subset \mathcal{T}_0 of the observations $\{y_i\}$;
- (2) LOOP: for $0 \leq t < T$,

$$(1.7) \quad \mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{\mu_t}{m} \sum_{i \in \mathcal{T}_{t+1}} \nabla \ell(\mathbf{z}^{(t)}; y_i)$$

for some index subset $\mathcal{T}_{t+1} \subseteq \{1, \dots, m\}$ determined by $\mathbf{z}^{(t)}$.

Three remarks are in order.

- First, we regularize both the initialization and the gradient flow in a data-dependent fashion by operating only upon some iteration-varying index subsets \mathcal{T}_t . This is a distinguishing feature of TWF in comparison to WF and other gradient descent variants. In words, \mathcal{T}_t corresponds to those data $\{y_i\}$ whose resulting spectral or gradient components are in some sense not excessively large; see Sections 2 and 3 for details. As we shall see

²The standard notation $f(n) = O(g(n))$ or $f(n) \lesssim g(n)$ (respectively, $f(n) = \Omega(g(n))$ or $f(n) \gtrsim g(n)$) means that there exists a constant $c > 0$ such that $|f(n)| \leq c|g(n)|$ (respectively, $|f(n)| \geq c|g(n)|$). $f(n) \asymp g(n)$ means that there exist constants $c_1, c_2 > 0$ such that $c_1|g(n)| \leq |f(n)| \leq c_2|g(n)|$.

³M. Soltanolkotabi recently informed us that the sample complexity of WF may be improved if one employs a better initialization procedure.

later, the main point is that this careful data-trimming procedure gives us a tighter initial guess and more stable search directions.

- Second, we recommend that the step size μ_t is either taken as some appropriate constant or determined by a backtracking line search. For instance, under appropriate conditions, we can take $\mu_t = 0.2$ for all t .
- Finally, the most expensive part of the gradient stage consists in computing $\nabla \ell(\mathbf{z}; y_i)$, $1 \leq i \leq m$, which can often be performed in an efficient manner. More concretely, under the *real-valued* Poisson data model (1.4) one has

$$\nabla \ell(\mathbf{z}; y_i) = 2 \left\{ \frac{y_i}{|\mathbf{a}_i^\top \mathbf{z}|^2} \mathbf{a}_i \mathbf{a}_i^\top \mathbf{z} - \mathbf{a}_i \mathbf{a}_i^\top \mathbf{z} \right\} = 2 \left(\frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \right) \mathbf{a}_i.$$

Thus, calculating $\{\nabla \ell(\mathbf{z}; y_i)\}$ essentially amounts to two matrix-vector products. Letting $\mathbf{A} := [\mathbf{a}_1, \dots, \mathbf{a}_m]^\top$ as before, we have

$$\sum_{i \in \mathcal{T}_{t+1}} \nabla \ell(\mathbf{z}^{(t)}; y_i) = \mathbf{A}^\top \mathbf{v}, \quad v_i = \begin{cases} 2 \frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}}, & i \in \mathcal{T}_{t+1}, \\ 0 & \text{otherwise.} \end{cases}$$

Hence, $\mathbf{A} \mathbf{z}$ gives \mathbf{v} and $\mathbf{A}^\top \mathbf{v}$ the desired regularized gradient.

A detailed specification of the algorithm is deferred to Section 2.

1.4 Numerical Surprises

To give the readers a sense of the practical power of TWF, we present here three illustrative numerical examples. Since it is impossible to recover the global sign—i.e., we cannot distinguish \mathbf{x} from $-\mathbf{x}$ —we will evaluate our solutions to the quadratic equations through the distance measure put forth in [13] representing the euclidean distance modulo a global sign: for complex-valued signals,

$$(1.8) \quad \text{dist}(\mathbf{z}, \mathbf{x}) := \min_{\varphi \in [0, 2\pi)} \|e^{-j\varphi} \mathbf{z} - \mathbf{x}\|,$$

while it is simply $\min \|\mathbf{z} \pm \mathbf{x}\|$ in the real-valued case. We shall use $\text{dist}(\hat{\mathbf{x}}, \mathbf{x})/\|\mathbf{x}\|$ throughout to denote the relative error of an estimate $\hat{\mathbf{x}}$. In the following, TWF proceeds by attempting to maximize the Poisson log-likelihood (1.6). Standalone MATLAB implementations of TWF are available at <http://statweb.stanford.edu/~candes/publications.html> (see [12] for straight WF implementations).

The first numerical example concerns the following two problems under noiseless real-valued data:

- (a) find $\mathbf{x} \in \mathbb{R}^n$ s.t. $b_i = \mathbf{a}_i^\top \mathbf{x}$, $1 \leq i \leq m$;
- (b) find $\mathbf{x} \in \mathbb{R}^n$ s.t. $b_i = |\mathbf{a}_i^\top \mathbf{x}|$, $1 \leq i \leq m$.

Apparently, (a) involves solving a linear system of equations (or a linear least squares problem), while (b) is tantamount to solving a quadratic system. Arguably the most popular method for solving large-scale least squares problems is the conjugate gradient (CG) method [38] applied to the normal equations. We are going to compare the computational efficiency between CG (for solving least




FIGURE 1.1. Relative errors of CG and TWF vs. iteration count. Here, $n = 1000$, $m = 8n$, and TWF is seeded using just 10 power iterations.

squares) and TWF with a step size $\mu_t \equiv 0.2$ (for solving a quadratic system). Set $m = 8n$ and generate $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $1 \leq i \leq m$, independently. This gives a matrix $\mathbf{A}^\top \mathbf{A}$ with a low condition number equal to about $(1 + \sqrt{1/8})^2 / (1 - \sqrt{1/8})^2 \approx 4.38$ by the Marchenko-Pastur law. Therefore, this is an ideal setting for CG as it converges extremely rapidly [49, theorem 38.5]. Figure 1.1 shows the relative estimation error of each method as a function of the iteration count, where TWF is seeded through 10 power iterations. For ease of comparison, we illustrate the iteration counts in different scales so that 4 TWF iterations are equivalent to 1 CG iteration.

Recognizing that each iteration of CG and TWF involves two matrix-vector products $\mathbf{A}\mathbf{z}$ and $\mathbf{A}^\top \mathbf{v}$, for such a design we reach a surprising observation:

Even when all phase information is missing, TWF is capable of solving a quadratic system of equations only about 4 times slower than solving a least squares problem of the same size!

To illustrate the applicability of TWF on real images, we turn to testing our algorithm on a digital photograph of the Stanford main quad containing 320×1280 pixels. We consider a type of measurement that falls under the category of coded diffraction patterns (CDP) [11] and set

$$(1.9) \quad \mathbf{y}^{(l)} = |\mathbf{F} \mathbf{D}^{(l)} \mathbf{x}|^2, \quad 1 \leq l \leq L.$$

Here \mathbf{F} stands for a discrete Fourier transform (DFT) matrix, and $\mathbf{D}^{(l)}$ is a diagonal matrix whose diagonal entries are independently and uniformly drawn from $\{1, -1, j, -j\}$ (phase delays). In phase retrieval, each $\mathbf{D}^{(l)}$ represents a random mask placed after the object so as to modulate the illumination patterns. When L masks are employed, the total number of quadratic measurements is $m = nL$. In this example, $L = 12$ random coded patterns are generated to measure each



(a)



(b)



(c)

FIGURE 1.2. The recovered images after (a) spectral initialization; (b) regularized spectral initialization; and (c) 50 TWF gradient iterations following the regularized initialization.

color band (i.e., red, green, or blue) separately. The experiment is carried out on a MacBook Pro equipped with a 3GHz Intel Core i7 and 16GB of memory. We run 50 iterations of the truncated power method for initialization, and 50 regularized gradient iterations, which in total costs 43.9 seconds or 2400 FFTs for each color band. The relative errors after regularized spectral initialization and after 50 TWF iterations are 0.4773 and 2.16×10^{-5} , respectively, with the recovered images displayed in Figure 1.2. In comparison, the spectral initialization using 50 untruncated power iterations returns an image of relative error 1.409 , which is almost like a random guess and extremely far from the truth.

While the above experiments concern noiseless data, the numerical surprise extends to the noisy realm. Suppose the data are drawn according to the Poisson noise model (1.4), with $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ independently generated. Figure 1.3 displays the empirical relative mean-square error (MSE) of TWF as a function of the




FIGURE 1.3. Relative MSE vs. SNR in dB. The curves are shown for two settings: TWF for solving quadratic equations (blue), and MLE had we observed additional phase information (green). The results are shown for $n = 100$, and each point is averaged over 50 Monte Carlo trials.

signal-to-noise ratio (SNR), where the relative MSE for an estimate $\hat{\mathbf{x}}$ and the SNR are defined as⁴

$$(1.10) \quad \text{MSE} := \frac{\text{dist}^2(\hat{\mathbf{x}}, \mathbf{x})}{\|\mathbf{x}\|^2} \quad \text{and} \quad \text{SNR} := 3\|\mathbf{x}\|^2.$$

Both SNR and MSE are displayed on a dB scale (i.e., the values of $10 \log_{10}(\text{SNR})$ and $10 \log_{10}(\text{rel. MSE})$ are plotted).

To evaluate the accuracy of the TWF solutions, we consider the performance achieved by MLE applied to an *ideal* problem in which the true phases are revealed. In this ideal scenario, in addition to the data $\{y_i\}$ we are further given exact phase information $\{\varphi_i = \text{sign}(\mathbf{a}_i^\top \mathbf{x})\}$. Such precious information gives away the phase retrieval problem and makes the MLE efficiently solvable since the MLE problem with side information

$$\begin{aligned} & \text{minimize}_{\mathbf{z} \in \mathbb{R}^n} && -\sum_{i=1}^m y_i \log(|\mathbf{a}_i^\top \mathbf{z}|^2) + (\mathbf{a}_i^\top \mathbf{z})^2 \\ & \text{subject to} && \varphi_i = \text{sign}(\mathbf{a}_i^\top \mathbf{z}) \end{aligned}$$

can be cast as a convex program

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^n} \quad -\sum_{i=1}^m 2y_i \log(\varphi_i \mathbf{a}_i^\top \mathbf{z}) + (\mathbf{a}_i^\top \mathbf{z})^2.$$

Figure 1.3 illustrates the empirical performance for this ideal problem. The plots demonstrate that even when all phases are erased, TWF yields a solution of nearly the best possible quality, since it only incurs an extra 1.5 dB loss compared to ideal

⁴To justify the definition of SNR, note that the signals and noise are captured by $\mu_i = (\mathbf{a}_i^\top \mathbf{x})^2$ and $y_i - \mu_i$, $1 \leq i \leq m$, respectively. The ratio of the signal power to the noise power is therefore

$$\frac{\sum_{i=1}^m \mu_i^2}{\sum_{i=1}^m \text{Var}[y_i]} = \frac{\sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{x}|^4}{\sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{x}|^2} \approx \frac{3m\|\mathbf{x}\|^4}{m\|\mathbf{x}\|^2} = 3\|\mathbf{x}\|^2.$$

MLE computed with all true phases revealed. This phenomenon arises regardless of the SNR!

1.5 Main Results

The preceding numerical discoveries unveil promising features of TWF in three aspects: (1) exponentially fast convergence, (2) exact recovery from noiseless data with sample complexity $O(n)$, and (3) nearly minimal mean-square loss in the presence of noise. This paper offers a formal substantiation of all these findings. To this end, we assume a tractable model in which the design vectors \mathbf{a}_i 's are independent Gaussian:

$$(1.11) \quad \mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n).$$

For concreteness, our results are concerned with TWF based on the Poisson log-likelihood function

$$(1.12) \quad \ell_i(\mathbf{z}) := \ell(\mathbf{z}; y_i) = y_i \log(|\mathbf{a}_i^\top \mathbf{z}|^2) - |\mathbf{a}_i^\top \mathbf{z}|^2,$$

where we shall use $\ell_i(\mathbf{z})$ as a shorthand for $\ell(\mathbf{z}; y_i)$ from now on. We begin with the performance guarantees of TWF in the absence of noise.

THEOREM 1.1 (Exact Recovery). *Consider the noiseless case (1.1) with an arbitrary signal $\mathbf{x} \in \mathbb{R}^n$. Suppose that the step size μ_t is either taken to be a positive constant $\mu_t \equiv \mu$ or chosen via a backtracking line search. Then there exist some universal constants $0 < \rho, \nu < 1$ and $\mu_0, c_0, c_1, c_2 > 0$ such that with probability exceeding $1 - c_1 \exp(-c_2 m)$, the truncated Wirtinger flow estimates (Algorithm 1 with parameters specified in Table 2.1) obey*

$$(1.13) \quad \text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \nu(1 - \rho)^t \|\mathbf{x}\| \quad \forall t \in \mathbb{N}$$

provided that

$$m \geq c_0 n \quad \text{and} \quad 0 < \mu \leq \mu_0.$$

As explained below, we can often take $\mu_0 \approx 0.3$.

Remark 1.2. As will be made precise in Section 5 (and in particular Proposition 5.1), one can take

$$\mu_0 = \frac{0.994 - \zeta_1 - \zeta_2 - \sqrt{2/(9\pi)} \alpha_h^{-1}}{2(1.02 + 0.665/\alpha_h)}$$

for some small quantities ζ_1, ζ_2 and some predetermined threshold α_h that is usually taken to be $\alpha_h \geq 5$. Under appropriate conditions, one can treat μ_0 as $\mu_0 \approx 0.3$.

Theorem 1.1 justifies at least two appealing features of TWF: (i) *minimal sample complexity* and (ii) *linear-time computational cost*. Specifically, TWF allows exact recovery from $O(n)$ quadratic equations, which is optimal since one needs at least n measurements to have a well-posed problem. Also, because of the geometric convergence rate, TWF achieves ϵ -accuracy (i.e., $\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \epsilon \|\mathbf{x}\|$) within at most

$O(\log(1/\epsilon))$ iterations. The total computational cost is therefore $O(mn \log(1/\epsilon))$, which is linear in the problem size. These outperform the performance guarantees of WF [13], which runs in $O(mn^2 \log(1/\epsilon))$ time and requires $O(n \log n)$ sample complexity.

We emphasize that enhanced performance vis-à-vis WF is not the result of a sharper analysis, but rather the result of key algorithmic changes. In both the initialization and iterative refinement stages, TWF proceeds in a more prudent manner by means of proper regularization, which effectively trims away those components that are too influential on either the initial guess or search directions, thus reducing the volatility of each movement. With a tighter initialization and better-controlled search directions in place, we take the step size in a far more liberal fashion—which is some constant bounded away from 0—compared to a step size that is $O(1/n)$ as explained in [13]. In fact, what enables the movement to be more aggressive is exactly the cautious choice of \mathcal{T}_t , which precludes adverse effects from high-leverage samples.

To be broadly applicable, the proposed algorithm must guarantee reasonably faithful estimates in the presence of noise. Suppose that

$$(1.14) \quad y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2 + \eta_i, \quad 1 \leq i \leq m,$$

where η_i represents an error term. We claim that TWF is stable against additive noise, as demonstrated in the theorem below.

THEOREM 1.3 (Stability). *Consider the noisy case (1.14). Suppose that the step size μ_t is either taken to be a positive constant $\mu_t \equiv \mu$ or chosen via a backtracking line search. If*

$$(1.15) \quad m \geq c_0 n, \quad \mu \leq \mu_0, \quad \text{and} \quad \|\boldsymbol{\eta}\|_\infty \leq c_1 \|\mathbf{x}\|^2,$$

then with probability at least $1 - c_2 \exp(-c_3 m)$, the truncated Wirtinger flow estimates (Algorithm 1 with parameters specified in Table 2.1) satisfy

$$(1.16) \quad \text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \lesssim \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{x}\|} + (1 - \rho)^t \|\mathbf{x}\| \quad \forall t \in \mathbb{N}$$

simultaneously for all $\mathbf{x} \in \mathbb{R}^n$. Here, $0 < \rho < 1$ and $\mu_0, c_0, c_1, c_2, c_3 > 0$ are some universal constants.

Under the Poisson noise model (1.4), one has

$$(1.17) \quad \text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \lesssim 1 + (1 - \rho)^t \|\mathbf{x}\| \quad \forall t \in \mathbb{N}$$

with probability approaching 1, provided that $\|\mathbf{x}\| \geq \log^{1.5} m$.

Remark 1.4. In the main text, we will prove Theorem 1.3 only for the case where \mathbf{x} is fixed and independent of the design vectors $\{\mathbf{a}_i\}$. Interested readers are referred to the supplemental materials [15] for the proof of the universal theory (i.e., the case simultaneously accommodating all $\mathbf{x} \in \mathbb{R}^n$). Note that when there is no noise ($\boldsymbol{\eta} = \mathbf{0}$), this stronger result guarantees the universality of the noiseless recovery.

Remark 1.5. [45] establishes stability estimates using the WF approach under Gaussian noise. There, the sample and computational complexities are still on the order of $n \log n$ and mn^2 , respectively, whereas the computational complexity in Theorem 1.3 is linear, i.e., on the order of mn .

Theorem 1.3 essentially reveals that the estimation error of TWF rapidly shrinks to $O((\|\boldsymbol{\eta}\|/\sqrt{m})/\|\mathbf{x}\|)$ within logarithmic iterations. Put another way, since the SNR for the model (1.14) is captured by

$$(1.18) \quad \text{SNR} := \frac{\sum_{i=1}^m |\langle \mathbf{a}_i, \mathbf{x} \rangle|^4}{\|\boldsymbol{\eta}\|^2} \approx \frac{3m\|\mathbf{x}\|^4}{\|\boldsymbol{\eta}\|^2},$$

we immediately arrive at an alternative form of the performance guarantee:

$$(1.19) \quad \text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \lesssim \frac{1}{\sqrt{\text{SNR}}} \|\mathbf{x}\| + (1 - \rho)^t \|\mathbf{x}\| \quad \forall t \in \mathbb{N},$$

revealing the stability of TWF as a function of SNR. We emphasize that this estimate holds for any error term $\boldsymbol{\eta}$ —i.e., any noise structure, even deterministic. This being said, specializing this estimate to the Poisson noise model (1.4) with $\|\mathbf{x}\| \gtrsim \log^{1.5} m$ gives an estimation error that will eventually approach a numerical constant, independent of n and m .

Encouragingly, this is already the best statistical guarantee any algorithm can achieve. We formalize this claim by deriving a fundamental lower bound on the minimax estimation error.

THEOREM 1.6 (Lower Bound on the Minimax Risk). *Suppose that $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $m = \kappa n$ for some fixed κ independent of n , and n is sufficiently large. For any $K \geq \log^{1.5} m$, define⁵*

$$\Upsilon(K) := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \in (1 \pm 0.1)K\}.$$

With probability approaching 1, the minimax risk under the Poisson model (1.4) obeys

$$(1.20) \quad \inf_{\hat{\mathbf{x}}} \sup_{\mathbf{x} \in \Upsilon(K)} \mathbb{E}[\text{dist}(\hat{\mathbf{x}}, \mathbf{x}) \mid \{\mathbf{a}_i\}_{1 \leq i \leq m}] \geq \frac{\varepsilon_1}{\sqrt{\kappa}},$$

where the infimum is over all estimators $\hat{\mathbf{x}}$. Here, $\varepsilon_1 > 0$ is a numerical constant independent of n and m .

When the number m of measurements is proportional to n and the energy of the planted solution exceeds $\log^3 m$, Theorem 1.6 asserts that there exists absolutely no estimator that can achieve an estimation error that vanishes as n increases. This lower limit matches the estimation error of TWF, which corroborates the optimality of TWF under noisy data.

Recall that in many optical imaging applications, the output data we collect are the intensities of the diffractive waves scattered by the sample or specimen under

⁵ Here, 0.1 can be replaced by any positive constant within $(0, \frac{1}{2})$.

study. The Poisson noise model employs the input \mathbf{x} and output \mathbf{y} to describe the numbers of photons diffracted by the specimen and detected by the optical sensor, respectively. Each specimen needs to be sufficiently illuminated in order for the receiver to sense the diffracted light. In such settings, the low-intensity regime $\|\mathbf{x}\| \leq \log^{1.5} m$ is of little practical interest as it corresponds to an illumination with just very few photons. We forego the details.

It is worth noting that apart from WF, various other nonconvex procedures have been proposed as well for phase retrieval, including the error reduction schemes dating back to Gerchberg-Saxton and Fienup [23, 24], iterated projections [22], alternating minimization [36], generalized approximate message passing [41], the Kaczmarz method [53], and greedy methods that exploit additional sparsity constraint [42], to name just a few. While these paradigms enjoy favorable empirical behavior, most of them fall short of theoretical support except for a version of alternating minimization (called AltMinPhase) [36] that requires fresh samples for each iteration. In comparison, AltMinPhase attains ϵ -accuracy when the sample complexity exceeds the order of $n \log^3 n + n \log^2 n \log(1/\epsilon)$, which is at least a factor of $\log^3 n$ from optimal and is empirically largely outperformed by the variant that reuses all samples.

In contrast, our algorithm uses the same set of samples all the time and is therefore practically appealing. Furthermore, none of these algorithms come with provable stability guarantees, which are particularly important in most realistic scenarios. Numerically, each iteration of Fienup's algorithm (or alternating minimization) involves solving a least squares problem, and the algorithm converges in tens or hundreds of iterations. This is computationally more expensive than TWF, whose computational complexity is merely about 4 times that of solving a least squares problem. Interested readers are referred to [13] for a comparison of several nonconvex schemes and to [11] for a discussion of other alternative approaches (e.g., [1, 5]) and performance lower bounds (e.g., [6, 21]).

2 Algorithm: Truncated Wirtinger Flow

This section describes the two stages of TWF in detail, presented in a reverse order. For each stage, we start with some algorithmic issues encountered by WF, which is then used to motivate and explain the basic principles of TWF. Here and throughout, we let $\mathcal{A} : \mathbb{R}^{n \times n} \mapsto \mathbb{R}^m$ be the linear map

$$\mathbf{M} \in \mathbb{R}^{n \times n} \quad \mapsto \quad \mathcal{A}(\mathbf{M}) := \{\mathbf{a}_i^\top \mathbf{M} \mathbf{a}_i\}_{1 \leq i \leq m}$$

and \mathbf{A} the design matrix

$$\mathbf{A} := [\mathbf{a}_1, \dots, \mathbf{a}_m]^\top.$$




FIGURE 2.1. The locus of

$$-\frac{1}{2}\nabla\ell_i(z) = \frac{|a_i^\top z|^2 - |a_i^\top x|^2}{a_i^\top z} a_i$$

when a_i ranges over all unit vectors, where $x = (2.7, 8)$ and $z = (3, 6)$. For each direction a_i , $-\frac{1}{2}\nabla\ell_i(z)$ is aligned with a_i , and its length represents the weight assigned to this component. In particular, the red arrows depict a few directions that behave like outliers, whereas the blue arrows depict several directions whose resulting gradients take typical sizes.

2.1 Regularized Gradient Stage

For independent samples, the gradient of the real-valued Poisson log-likelihood obeys

$$(2.1) \quad \sum_{i=1}^m \nabla\ell_i(z) = \sum_{i=1}^m 2 \underbrace{\frac{y_i - |a_i^\top z|^2}{a_i^\top z}}_{:=v_i} a_i,$$

where v_i represents the weight assigned to each a_i . This forms the descent direction of WF updates.

Unfortunately, WF moving along the preceding direction might not come close to the truth unless z is already very close to x . To see this, it is helpful to consider any fixed vector $z \in \mathbb{R}^n$ independent of the design vectors. The typical size of $\min_{1 \leq i \leq m} |a_i^\top z|$ is about on the order of $\frac{1}{m}\|z\|$. This introduces some unreasonably large weights v_i , which can be as large as $m\|x\|^2/\|z\|$. Consequently, the iterative updates based on (1.2) often overshoot, and this arises starting from the very initial stage.⁶

Figure 2.1 illustrates this phenomenon by showing the locus of $-\nabla\ell_i(z)$ when a_i has unit norm and ranges over all possible directions. Examination of the figure seems to suggest that most of the gradient components $\nabla\ell_i(z)$ are more or less pointing towards the truth x and forming reasonable search directions. But there

⁶For complex-valued data where $a_i \sim \mathcal{N}(\mathbf{0}, I) + j\mathcal{N}(\mathbf{0}, I)$, WF converges empirically, as $\min_i |a_i^* z|$ is much larger than the real-valued case.




FIGURE 2.2. Relative MSE vs. SNR in dB. The curves are shown for WF and TWF, both employing the Poisson log-likelihood. Here, $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $n = 1000$, $m = 8n$, and each point is averaged over 100 Monte Carlo trials.

exist a few outlier components that are excessively large, which lead to unstable search directions. Notably, an underlying premise for a nonconvex procedure to succeed is to ensure all iterates reside within a *basin of attraction*, that is, a neighborhood surrounding \mathbf{x} within which \mathbf{x} is the unique stationary point of the objective. When a gradient is not well controlled, the iterative procedure might overshoot and end up leaving this basin of attraction. This intuition is corroborated by numerical experiments under *real-valued* data. As illustrated in Figure 2.2, the solutions returned by the WF (designed for a real-valued Poisson log-likelihood and $m = 8n$) are very far from the ground truth.

Hence, to remedy the aforementioned stability issue, it would be natural to separate the small fraction of abnormal gradient components by regularizing the weights v_i , possibly via data-dependent trimming rules. This gives rise to the update rule of TWF:

$$(2.2) \quad \mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{\mu_t}{m} \nabla \ell_{\text{tr}}(\mathbf{z}^{(t)}) \quad \forall t \in \mathbb{N},$$

where $\nabla \ell_{\text{tr}}(\cdot)$ denotes the regularized gradient given by⁷

$$(2.3) \quad \nabla \ell_{\text{tr}}(\mathbf{z}) := \sum_{i=1}^m 2 \frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i(\mathbf{z}) \cap \mathcal{E}_2^i(\mathbf{z})}$$

⁷In the complex-valued case, the trimming rule is enforced upon the Wirtinger derivative, which reads

$$\nabla \ell_{\text{tr}}(\mathbf{z}) := \sum_{i=1}^m 2 \frac{y_i - |\mathbf{z}^* \mathbf{a}_i|^2}{\mathbf{z}^* \mathbf{a}_i} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i(\mathbf{z}) \cap \mathcal{E}_2^i(\mathbf{z})}.$$

for some trimming criteria specified by $\mathcal{E}_1^i(\cdot)$ and $\mathcal{E}_2^i(\cdot)$. In our algorithm, we take $\mathcal{E}_1^i(\mathbf{z})$ and $\mathcal{E}_2^i(\mathbf{z})$ to be two collections of events given by

$$(2.4) \quad \mathcal{E}_1^i(\mathbf{z}) := \left\{ \alpha_z^{\text{lb}} \leq \frac{|\mathbf{a}_i^\top \mathbf{z}|}{\|\mathbf{z}\|} \leq \alpha_z^{\text{ub}} \right\},$$

$$(2.5) \quad \mathcal{E}_2^i(\mathbf{z}) := \left\{ |y_i - |\mathbf{a}_i^\top \mathbf{z}|| \leq \frac{\alpha_h}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{z}\mathbf{z}^\top)\|_1 \frac{|\mathbf{a}_i^\top \mathbf{z}|}{\|\mathbf{z}\|} \right\},$$

where α_z^{lb} , α_z^{ub} , α_z are predetermined thresholds. To keep notation light, we shall use \mathcal{E}_1^i and \mathcal{E}_2^i rather than $\mathcal{E}_1^i(\mathbf{z})$ and $\mathcal{E}_2^i(\mathbf{z})$ whenever it is clear from context.

We emphasize that the above trimming procedure simply throws away those components whose weights v_i 's fall outside some confidence range, so as to remove the influence of outlier components. To achieve this, we regularize both the numerator and denominator of v_i by enforcing separate trimming rules. Recognize that for any fixed \mathbf{z} , the denominator obeys

$$\mathbb{E}[|\mathbf{a}_i^\top \mathbf{z}|] = \sqrt{2/\pi} \|\mathbf{z}\|,$$

leading up to the rule (2.4). Regarding the numerator, by the law of large numbers one would expect

$$\mathbb{E}[|y_i - |\mathbf{a}_i^\top \mathbf{z}||^2] \approx \frac{1}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{z}\mathbf{z}^\top)\|_1,$$

and hence it is natural to regularize the numerator by ensuring

$$|y_i - |\mathbf{a}_i^\top \mathbf{z}||^2 \lesssim \frac{1}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{z}\mathbf{z}^\top)\|_1.$$

As a remark, we include an extra term $|\mathbf{a}_i^\top \mathbf{z}|/\|\mathbf{z}\|$ in (2.5) to sharpen the theory, but all our results continue to hold (up to some modification of constants) if we drop this term in (2.5). Detailed procedures are summarized in Algorithm 1.⁸

The proposed paradigm could be counterintuitive at first glance, since one might expect the larger terms to be better aligned with the desired search direction. The issue, however, is that the large terms are extremely volatile and could have too high of a leverage on the descent directions. In contrast, TWF discards these high-leverage data, slightly increasing the bias but remarkably reducing the variance of the descent direction. We expect such gradient regularization and variance reduction schemes to be beneficial for solving a broad family of nonconvex problems.

⁸Careful readers might note that we include some extra factor $\sqrt{n}/\|\mathbf{a}_i\|$ (which is approximately 1 in the Gaussian model) in Algorithm 1. This occurs since we present Algorithm 1 in a more general fashion that applies beyond the model $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, but all results and proofs continue to hold in the presence of this extra term.

Algorithm 1 Truncated Wirtinger Flow.

Input: Measurements $\{y_i \mid 1 \leq i \leq m\}$ and sampling vectors $\{\mathbf{a}_i \mid 1 \leq i \leq m\}$; trimming thresholds $\alpha_z^{\text{lb}}, \alpha_z^{\text{ub}}, \alpha_h$, and α_y (see default values in Table 2.1).

Initialize $\mathbf{z}^{(0)}$ to be $\sqrt{\frac{mn}{\sum_{i=1}^m \|\mathbf{a}_i\|^2}} \lambda_0 \tilde{\mathbf{z}}$, where $\lambda_0 = \sqrt{\frac{1}{m} \sum_{i=1}^m y_i}$ and $\tilde{\mathbf{z}}$ is the leading eigenvector of

$$(2.6) \quad \mathbf{Y} = \frac{1}{m} \sum_{i=1}^m y_i \mathbf{a}_i \mathbf{a}_i^* \mathbb{1}_{\{|y_i| \leq \alpha_y^2 \lambda_0^2\}}.$$

Loop: for $t = 0 : T$ do

$$(2.7) \quad \mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{2\mu_t}{m} \sum_{i=1}^m \frac{y_i - |\mathbf{a}_i^* \mathbf{z}^{(t)}|^2}{\mathbf{z}^{(t)*} \mathbf{a}_i} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i},$$

where

$$(2.8) \quad \mathcal{E}_1^i := \left\{ \alpha_z^{\text{lb}} \leq \frac{\sqrt{n}}{\|\mathbf{a}_i\|} \frac{|\mathbf{a}_i^* \mathbf{z}^{(t)}|}{\|\mathbf{z}^{(t)}\|} \leq \alpha_z^{\text{ub}} \right\},$$

$$\mathcal{E}_2^i := \left\{ |y_i - |\mathbf{a}_i^* \mathbf{z}^{(t)}|^2| \leq \alpha_h K_t \frac{\sqrt{n}}{\|\mathbf{a}_i\|} \frac{|\mathbf{a}_i^* \mathbf{z}^{(t)}|}{\|\mathbf{z}^{(t)}\|} \right\},$$

$$\text{and } K_t := \frac{1}{m} \sum_{l=1}^m |y_l - |\mathbf{a}_l^* \mathbf{z}^{(t)}|^2|.$$

Output \mathbf{z}_T .

2.2 Truncated Spectral Initialization

In order for the gradient stage to converge rapidly, we need to seed it with a suitable initialization. One natural alternative is the spectral method adopted in [13, 36], which amounts to computing the leading eigenvector of

$$\tilde{\mathbf{Y}} := \frac{1}{m} \sum_{i=1}^m y_i \mathbf{a}_i \mathbf{a}_i^\top.$$

This arises from the observation that when $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\|\mathbf{x}\| = 1$,

$$\mathbb{E}[\tilde{\mathbf{Y}}] = \mathbf{I} + 2\mathbf{x}\mathbf{x}^\top,$$

whose leading eigenvector is exactly \mathbf{x} with an eigenvalue of 3.

Unfortunately, this spectral technique converges to a good initial point only when $m \gtrsim n \log n$, due to the fact that $(\mathbf{a}_i^\top \mathbf{x})^2 \mathbf{a}_i \mathbf{a}_i^\top$ is heavy-tailed, a random quantity that does not have a moment-generating function. To be more precise, consider the noiseless case $y_i = |\mathbf{a}_i^\top \mathbf{x}|^2$ and recall that $\max_i y_i \approx 2 \log m$. Letting

$k = \arg \max_i y_i$, one can calculate

$$\left(\frac{\mathbf{a}_k}{\|\mathbf{a}_k\|} \right)^\top \tilde{\mathbf{Y}} \frac{\mathbf{a}_k}{\|\mathbf{a}_k\|} \geq \left(\frac{\mathbf{a}_k}{\|\mathbf{a}_k\|} \right)^\top \left(\frac{1}{m} \mathbf{a}_k \mathbf{a}_k^\top \right) (\mathbf{a}_k^\top \mathbf{x})^2 \left(\frac{\mathbf{a}_k}{\|\mathbf{a}_k\|} \right) \approx \frac{2n \log m}{m},$$

which is much larger than $\mathbf{x}^\top \tilde{\mathbf{Y}} \mathbf{x} = 3$ unless m/n is very large. This tells us that in the regime where $m \asymp n$, there exists some unit vector $\mathbf{a}_k/\|\mathbf{a}_k\|$ that is closer to the leading eigenvector of $\tilde{\mathbf{Y}}$ than \mathbf{x} . This phenomenon happens because the summands of $\tilde{\mathbf{Y}}$ have huge tails so that even one large term could end up dominating the empirical sum, thus preventing the spectral method from returning a meaningful initial guess.

To address this issue, we propose a more robust version of spectral method, which discards those observations y_i that are several times larger than the mean during spectral initialization. Specifically, the initial estimate is obtained by computing the leading eigenvector $\tilde{\mathbf{z}}$ of the truncated sum

$$(2.9) \quad \mathbf{Y} := \frac{1}{m} \sum_{i=1}^m y_i \mathbf{a}_i \mathbf{a}_i^\top \mathbb{1}_{\{|y_i| \leq \alpha_y^2 (\frac{1}{m} \sum_{l=1}^m y_l)\}}$$

for some predetermined threshold α_y , and then rescaling $\tilde{\mathbf{z}}$ so as to have roughly the same norm as \mathbf{x} (which is estimated to be $\frac{1}{m} \sum_{l=1}^m y_l$); see Algorithm 1 for the detailed procedure.

Notably, the aforementioned drawback of the spectral method is not merely a theoretical concern but rather a substantial practical issue. We have seen this in Figure 1.2 (main quad example) showing the enormous advantage of truncated spectral initialization. This is also further illustrated in Figure 2.3, which compares the empirical efficiency of both methods with $\alpha_y = 3$ set to be the truncation threshold. For both Gaussian designs and CDP models, the empirical loss incurred by the original spectral method increases as n grows, which is in stark contrast to the truncated spectral method that achieves almost identical accuracy over the same range of n .

2.3 Choice of Algorithmic Parameters

One implementation detail to specify is the step size μ_t at each iteration t . There are two alternatives that work well in both theory and practice:

- (1) *Fixed step size.* Take $\mu_t \equiv \mu$ ($\forall t \in \mathbb{N}$) for some constant $\mu > 0$. As long as μ is not too large, our main results state that this strategy always works—although the convergence rate depends on μ . Under appropriate conditions, our theorems hold for any constant $0 < \mu < 0.28$.
- (2) *Backtracking line search with truncated objective.* This strategy performs a line search along the descent direction

$$\mathbf{p}_t := \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}_t)$$




FIGURE 2.3. The empirical relative error for both the spectral and the truncated spectral methods. The results are averaged over 50 Monte Carlo runs, and are shown for: (a) 1-D Gaussian measurement where $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $m = 6n$; (b) 2-D CDP model (1.9) where the diagonal entries of $\mathbf{D}^{(l)}$ are uniformly drawn from $\{1, -1, j, -j\}$, $n = n_1 \times n_2$ with $n_1 = 300$ and n_2 ranging from 64 to 1280, and $m = 12n$.

and determines an appropriate step size that guarantees a sufficient improvement. In contrast to the conventional search strategy that determines the sufficient progress with respect to the true objective function, we propose to evaluate instead a regularized version of the objective function. Specifically, put

$$(2.10) \quad \hat{\ell}(\mathbf{z}) := \sum_{i \in \hat{\mathcal{T}}(\mathbf{z})} \{y_i \log(|\mathbf{a}_i^\top \mathbf{z}|^2) - |\mathbf{a}_i^\top \mathbf{z}|^2\},$$

where

$$\hat{\mathcal{T}}(\mathbf{z}) := \{i \mid |\mathbf{a}_i^\top \mathbf{z}| \geq \alpha_z^{\text{lb}} \|\mathbf{z}\| \text{ and } |\mathbf{a}_i^\top \mathbf{p}| \leq \alpha_p \|\mathbf{p}\|\}.$$

Then the backtracking line search proceeds as follows:

- (a) Start with $\tau = 1$;
- (b) Repeat $\tau \leftarrow \beta\tau$ until

$$(2.11) \quad \frac{1}{m} \hat{\ell}(\mathbf{z}^{(t)} + \tau \mathbf{p}^{(t)}) \geq \frac{1}{m} \hat{\ell}(\mathbf{z}^{(t)}) + \frac{1}{2} \tau \|\mathbf{p}^{(t)}\|^2,$$

where $\beta \in (0, 1)$ is some predetermined constant;

- (c) Set $\mu_t = \tau$.

By definition (2.10), evaluating $\hat{\ell}(\mathbf{z}^{(t)} + \tau \mathbf{p}^{(t)})$ mainly consists in calculating the matrix-vector product $\mathbf{A}(\mathbf{z}^{(t)} + \tau \mathbf{p}^{(t)})$. In total, we are going to evaluate $\hat{\ell}(\mathbf{z}^{(t)} + \tau \mathbf{p}^{(t)})$ for $O(\log(1/\beta))$ different τ 's, and hence the total cost amounts to computing $\mathbf{A}\mathbf{z}^{(t)}$, $\mathbf{A}\mathbf{p}^{(t)}$ as well as $O(m \log(1/\beta))$ additional flops. Note that the matrix-vector products $\mathbf{A}\mathbf{z}^{(t)}$ and $\mathbf{A}\mathbf{p}^{(t)}$ need to

TABLE 2.1. Range of algorithmic parameters

(a) **When a fixed step size $\mu_t \equiv \mu$ is employed:** $(\alpha_z^{\text{lb}}, \alpha_z^{\text{ub}}, \alpha_h, \alpha_y)$ obeys

$$(2.13) \quad \begin{cases} \zeta_1 := \max\{\mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| \leq \sqrt{1.01}\alpha_z^{\text{lb}} \text{ or } |\xi| \geq \sqrt{0.99}\alpha_z^{\text{ub}}\}}], \\ \mathbb{P}(|\xi| \leq \sqrt{1.01}\alpha_z^{\text{lb}} \text{ or } |\xi| \geq \sqrt{0.99}\alpha_z^{\text{ub}})\}, \\ \zeta_2 := \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| > 0.473\alpha_h\}}], \\ 2(\zeta_1 + \zeta_2) + \sqrt{8/(9\pi)}\alpha_h^{-1} < 1.99, \\ \alpha_y \geq 3, \end{cases}$$

where $\xi \sim \mathcal{N}(0, 1)$. By default, $\alpha_z^{\text{lb}} = 0.3$, $\alpha_z^{\text{ub}} = \alpha_h = 5$, and $\alpha_y = 3$.

(b) **When μ_t is chosen by a backtracking line search:** $(\alpha_z^{\text{lb}}, \alpha_z^{\text{ub}}, \alpha_h, \alpha_y, \alpha_p)$ obeys

$$(2.14) \quad 0 < \alpha_z^{\text{lb}} \leq 0.1, \quad \alpha_z^{\text{ub}} \geq 5, \quad \alpha_h \geq 6, \quad \alpha_y \geq 3, \quad \text{and} \quad \alpha_p \geq 5.$$

By default, $\alpha_z^{\text{lb}} = 0.1$, $\alpha_z^{\text{ub}} = 5$, $\alpha_h = 6$, $\alpha_y = 3$, and $\alpha_p = 5$.

be computed even when one adopts a predetermined step size. Hence, the extra cost incurred by a backtracking line search, which is $O(m \log(1/\beta))$ flops, is negligible compared to that of computing the gradient even once.

Another set of important algorithmic parameters to determine is the trimming thresholds α_h , α_z^{lb} , α_z^{ub} , α_y , and α_p (for a backtracking line search only). The present paper isolates the set of $(\alpha_h, \alpha_z^{\text{lb}}, \alpha_z^{\text{ub}}, \alpha_y)$ obeying (2.13) as given in Table 2.1 when a fixed step size is employed. More concretely, this range subsumes as special cases all parameters obeying the following constraints:

$$(2.12) \quad 0 < \alpha_z^{\text{lb}} \leq 0.5, \quad \alpha_z^{\text{ub}} \geq 5, \quad \alpha_h \geq 5, \quad \text{and} \quad \alpha_y \geq 3.$$

When a backtracking line search is adopted, an extra parameter α_p is needed, which we take to be $\alpha_p \geq 5$. In all theory presented herein, we assume that the parameters fall within the range singled out in Table 2.1.

3 Why Does TWF Work?

Before proceeding, it is best to develop an intuitive understanding of the TWF iterations. We start with a notation representing the (unrecoverable) global phase [13] for real-valued data

$$(3.1) \quad \phi(\mathbf{z}) := \begin{cases} 0 & \text{if } \|\mathbf{z} - \mathbf{x}\| \leq \|\mathbf{z} + \mathbf{x}\|, \\ \pi & \text{otherwise.} \end{cases}$$

It is self-evident that

$$(-z) + \frac{\mu}{m} \nabla_{\text{tr}} \ell(-z) = -\left\{z + \frac{\mu}{m} \nabla_{\text{tr}} \ell(z)\right\},$$

and hence (cf. Definition (1.8))

$$\text{dist}\left((-z) + \frac{\mu}{m} \nabla_{\text{tr}} \ell(-z), x\right) = \text{dist}\left(z + \frac{\mu}{m} \nabla_{\text{tr}} \ell(z), x\right)$$

despite the global phase uncertainty. To simplify presentation, we shall drop the phase term by letting z be $e^{-j\phi(z)}z$ and setting $\mathbf{h} = z - x$ whenever it is clear from context.

The first object to consider is the descent direction. To this end, we find it convenient to work with a fixed z independent of the design vectors \mathbf{a}_i , which is of course heuristic but helpful in developing some intuition. Rewrite

$$\begin{aligned} \nabla \ell_i(z) &= 2 \frac{(\mathbf{a}_i^\top x)^2 - (\mathbf{a}_i^\top z)^2}{\mathbf{a}_i^\top z} \mathbf{a}_i \stackrel{(i)}{=} -2 \frac{(\mathbf{a}_i^\top \mathbf{h})(2\mathbf{a}_i^\top z - \mathbf{a}_i^\top \mathbf{h})}{\mathbf{a}_i^\top z} \mathbf{a}_i \\ (3.2) \quad &= -4(\mathbf{a}_i^\top \mathbf{h}) \mathbf{a}_i + \underbrace{2 \frac{(\mathbf{a}_i^\top \mathbf{h})^2}{\mathbf{a}_i^\top z}}_{:= \mathbf{r}_i} \mathbf{a}_i, \end{aligned}$$

where (i) follows from the identity $a^2 - b^2 = (a+b)(a-b)$. The first component of (3.2), which on average gives $-4\mathbf{h}$, makes a good search direction when averaged over all the observations $i = 1, \dots, m$. The issue is that the other term \mathbf{r}_i —which is in general nonintegrable—could be devastating. The reason is that $\mathbf{a}_i^\top z$ could be arbitrarily small, thus resulting in an unbounded \mathbf{r}_i . As a consequence, a nonnegligible portion of the \mathbf{r}_i 's may exert a very strong influence on the descent direction in an undesired manner.

Such an issue can be prevented if one can detect and separate those gradient components bearing abnormal \mathbf{r}_i 's. Since we cannot observe the individual components of the decomposition (3.2), we cannot reject indices with large values of \mathbf{r}_i directly. Instead, we examine each gradient component as a whole and discard it if its size is not absolutely controlled. Fortunately, such a strategy is sufficient to ensure that most of the contribution from the regularized gradient comes from the first component of (3.2), namely, $-4(\mathbf{a}_i^\top \mathbf{h}) \mathbf{a}_i$. As will be made precise in Proposition 5.5 and Lemma 5.9, the regularized gradient obeys

$$(3.3) \quad -\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(z), \mathbf{h} \right\rangle \geq (4 - \epsilon) \|\mathbf{h}\|^2 - O\left(\frac{\|\mathbf{h}\|^3}{\|\mathbf{z}\|}\right)$$

$$(3.4) \quad \text{and} \quad \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(z) \right\| \lesssim \|\mathbf{h}\|.$$

Here, one has $(4 - \epsilon)\|\mathbf{h}\|^2$ in (3.3) instead of $4\|\mathbf{h}\|^2$ to account for the bias introduced by adaptive trimming, where ϵ is small as long as we only throw away a small fraction of data. Looking at (3.3) and (3.4), we see that the search direction is sufficiently aligned with the deviation $-\mathbf{h} = x - z$ of the current iterate; i.e.,




FIGURE 3.1. A function $-\ell(z)$ satisfying RC: $-\ell(z) = z^2$ for any $z \in [-6, 6]$, and $-\ell(z) = z^2 + 1.5|z|(\cos(|z| - 6) - 1)$ otherwise.

they form a reasonably good angle that is bounded away from 90° . Consequently, z is expected to be dragged towards x provided that the step size is appropriately chosen.

The observations (3.3) and (3.4) are reminiscent of a (local) regularity condition given in [13], which is a fundamental criterion that dictates rapid convergence of iterative procedures (including WF and other gradient descent schemes). When specialized to TWF, we say that $-\frac{1}{m}\nabla\ell_{\text{tr}}(\cdot)$ satisfies the *regularity condition*, denoted by $\text{RC}(\mu, \lambda, \epsilon)$, if

$$(3.5) \quad \left\langle \mathbf{h}, -\frac{1}{m}\nabla\ell_{\text{tr}}(z) \right\rangle \geq \frac{\mu}{2} \left\| \frac{1}{m}\nabla\ell_{\text{tr}}(z) \right\|^2 + \frac{\lambda}{2} \|\mathbf{h}\|^2$$

holds for all z obeying $\|z - \mathbf{x}\| \leq \epsilon\|\mathbf{x}\|$, where $0 < \epsilon < 1$ is some constant. Such an ϵ -ball around \mathbf{x} forms a basin of attraction. Formally, under $\text{RC}(\mu, \lambda, \epsilon)$, a little algebra gives

$$\begin{aligned} \text{dist}^2\left(z + \frac{\mu}{m}\nabla\ell_{\text{tr}}(z), \mathbf{x}\right) &\leq \left\| z + \frac{\mu}{m}\nabla\ell_{\text{tr}}(z) - \mathbf{x} \right\|^2 \\ &= \|\mathbf{h}\|^2 + \left\| \frac{\mu}{m}\nabla\ell_{\text{tr}}(z) \right\|^2 + 2\mu \left\langle \mathbf{h}, \frac{1}{m}\nabla\ell_{\text{tr}}(z) \right\rangle \\ &\leq \|\mathbf{h}\|^2 + \left\| \frac{\mu}{m}\nabla\ell_{\text{tr}}(z) \right\|^2 - \mu^2 \left\| \frac{1}{m}\nabla\ell_{\text{tr}}(z) \right\|^2 - \mu\lambda\|\mathbf{h}\|^2 \\ (3.6) \quad &= (1 - \mu\lambda) \text{dist}^2(z, \mathbf{x}) \end{aligned}$$

for any z with $\|z - \mathbf{x}\| \leq \epsilon$. In words, the TWF update rule is locally contractive around the planted solution, provided that $\text{RC}(\mu, \lambda, \epsilon)$ holds for some nonzero μ and λ . Apparently, conditions (3.3) and (3.4) already imply the validity of RC for some constants $\mu, \lambda \asymp 1$ when $\|\mathbf{h}\|/\|z\|$ is reasonably small, which in turn allows us to take a constant step size μ and enables a constant contraction rate $1 - \mu\lambda$.

Finally, caution must be exercised when connecting RC with strong convexity, since the former does not necessarily guarantee the latter within the basin of attraction. As an illustration, Figure 3.1 plots the graph of a nonconvex function obeying




FIGURE 4.1. Empirical success rate under real-valued Gaussian sampling $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$.

RC. The distinction stems from the fact that RC is stated only for those pairs \mathbf{z} and $\mathbf{h} = \mathbf{z} - \mathbf{x}$ with \mathbf{x} being a fixed component, rather than simultaneously accommodating all possible \mathbf{z} and $\mathbf{h} = \mathbf{z} - \tilde{\mathbf{z}}$ with $\tilde{\mathbf{z}}$ being an arbitrary vector. In contrast, RC says that the only stationary point of the truncated objective in a neighborhood of \mathbf{x} is \mathbf{x} , which often suffices for a gradient-descent-type scheme to succeed.

4 Numerical Experiments

In this section, we report additional numerical results to verify the practical applicability of TWF. In all numerical experiments conducted in the current paper, we set

$$(4.1) \quad \alpha_z^{\text{lb}} = 0.3, \quad \alpha_z^{\text{ub}} = 5, \quad \alpha_h = 5, \quad \text{and} \quad \alpha_y = 3.$$

This is a concrete combination of parameters satisfying our condition (2.13). Unless otherwise noted, we employ 50 power iterations for initialization, adopt a fixed step size $\mu_t \equiv 0.2$ when updating TWF iterates, and set the maximum number of iterations to be $T = 1000$ for the iterative refinement stage.

The first series of experiments concerns exact recovery from noise-free data. Set $n = 1000$ and generate a real-valued signal \mathbf{x} at random. Then for m varying between $2n$ and $6n$, generate m design vectors \mathbf{a}_i independently drawn from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. An experiment is claimed to succeed if the returned estimate $\hat{\mathbf{x}}$ satisfies $\text{dist}(\hat{\mathbf{x}}, \mathbf{x}) / \|\mathbf{x}\| \leq 10^{-5}$. Figure 4.1 illustrates the empirical success rate of TWF (over 100 Monte Carlo trials for each m) revealing that exact recovery is practically guaranteed from fewer than 1000 iterations when the number of quadratic constraints is about 5 times the ambient dimension.

To see how special the real-valued Gaussian designs are to our theoretical finding, we perform experiments on two other types of measurement models. In the first, TWF is applied to complex-valued data by generating $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \frac{1}{2}\mathbf{I}) + j\mathcal{N}(\mathbf{0}, \frac{1}{2}\mathbf{I})$. The other is the model of coded diffraction patterns described in (1.9).

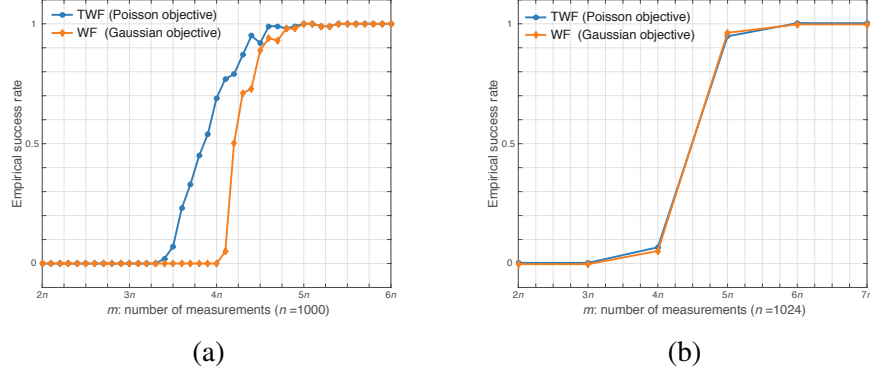


FIGURE 4.2. Empirical success rate for exact recovery using TWF. The results are shown for (a) complex-valued Gaussian sampling $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \frac{1}{2}\mathbf{I}_n) + j\mathcal{N}(\mathbf{0}, \frac{1}{2}\mathbf{I}_n)$, and (b) CDP with masks uniformly drawn from $\{1, -1, j, -j\}$.




FIGURE 4.3. Relative MSE vs. SNR when the y_i 's follow the Poisson model.

Figure 4.2 depicts the average success rate for both types of measurements over 100 Monte Carlo trials, indicating that $m > 4.5n$ and $m \geq 6n$ are often sufficient under complex-valued Gaussian and CDP models, respectively.

For the sake of comparison, we also report the empirical performance of WF in all the above settings, where the step size is set to be the default choice of [13], that is, $\mu_t = \min\{1 - e^{-t/330}, 0.2\}$. As can be seen, the empirical success rates of TWF outperform WF when $T = 1000$ under Gaussian models, suggesting that TWF either converges faster or exhibits better phase transition behavior.

Another series of experiments has been carried out to demonstrate the stability of TWF when the number m of quadratic equations varies. We consider the case where $n = 1000$ and vary the SNR (cf. (1.10)) from 15 dB to 55 dB. The design vectors are real-valued independent Gaussian $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, while the measurements y_i are generated according to the Poisson noise model (1.4). Figure 4.3 shows the relative mean square error—in the dB scale—as a function of SNR,

when averaged over 100 independent runs. For all choices of m , the numerical experiments demonstrate that the relative MSE scales inversely proportional to SNR, which matches our stability guarantees in Theorem 1.3 (since we observe that on the dB scale, the slope is about -1 as predicted by the theory (1.19)).

5 Exact Recovery from Noiseless Data

This section proves the theoretical guarantees of TWF in the absence of noise (i.e., Theorem 1.1). We separate the noiseless case mainly out of pedagogical reasons, as most of the steps carry over to the noisy case with slight modification.

The analysis for the truncated spectral method relies on the celebrated Davis-Kahan $\sin \Theta$ theorem [19], which we defer to Appendix C. In short, for any fixed $\delta > 0$ and $\mathbf{x} \in \mathbb{R}^n$, the initial point $\mathbf{z}^{(0)}$ returned by the truncated spectral method obeys

$$\text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq \delta \|\mathbf{x}\|$$

with high probability, provided that m/n exceeds some numerical constant. With this in place, it suffices to demonstrate that the TWF update rule is locally contractive, as stated in the following proposition.

PROPOSITION 5.1 (Local Error Contraction). *Consider the noiseless case (1.1). Under condition (2.13), there exist some universal constants $0 < \rho_0 < 1$ and $c_0, c_1, c_2 > 0$ such that with probability exceeding $1 - c_1 \exp(-c_2 m)$,*

$$(5.1) \quad \text{dist}^2\left(\mathbf{z} + \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{x}\right) \leq (1 - \rho_0) \text{dist}^2(\mathbf{z}, \mathbf{x})$$

holds simultaneously for all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ obeying

$$(5.2) \quad \frac{\text{dist}(\mathbf{z}, \mathbf{x})}{\|\mathbf{z}\|} \leq \min \left\{ \frac{1}{11}, \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{5.7(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}} \right\},$$

provided that $m \geq c_0 n$ and that μ is some constant obeying

$$0 < \mu \leq \mu_0 := \frac{0.994 - \zeta_1 - \zeta_2 - \sqrt{2/(9\pi)} \alpha_h^{-1}}{2(1.02 + 0.665/\alpha_h)}.$$

Proposition 5.1 reveals the monotonicity of the estimation error: once entering a neighborhood around \mathbf{x} of a reasonably small size, the iterative updates will remain within this neighborhood all the time and be attracted towards \mathbf{x} at a geometric rate.

As shown in Section 3, under the hypothesis $\text{RC}(\mu, \lambda, \epsilon)$ one can conclude

$$(5.3) \quad \text{dist}^2\left(\mathbf{z} + \frac{\mu}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{x}\right) \leq (1 - \mu\lambda) \text{dist}^2(\mathbf{z}, \mathbf{x}) \quad \forall (\mathbf{z}, \mathbf{x}) \text{ with } \text{dist}(\mathbf{z}, \mathbf{x}) \leq \epsilon.$$

Thus, everything now boils down to showing that $\text{RC}(\mu, \lambda, \epsilon)$ for some constants $\mu, \lambda, \epsilon > 0$. This occupies the rest of this section.




FIGURE 5.1. $\frac{f(t)}{\sqrt{1+t^2}}$ as a function of t .

5.1 Preliminary Facts about $\{\mathcal{E}_1^i\}$ and $\{\mathcal{E}_2^i\}$

Before proceeding, we gather a few properties of the events \mathcal{E}_1^i and \mathcal{E}_2^i , which will prove crucial in establishing $\text{RC}(\mu, \lambda, \epsilon)$. To begin with, recall that the truncation level given in \mathcal{E}_2^i depends on $\frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1$. Instead of working with this random variable directly, we use deterministic quantities that are more amenable to analysis. Specifically, we claim that $\frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1$ offers a uniform and orderwise tight estimate on $\|\mathbf{h}\|\|\mathbf{z}\|$, which can be seen from the following two facts.

LEMMA 5.2. Fix $\zeta \in (0, 1)$. If $m > c_0 n \zeta^{-2} \log \frac{1}{\zeta}$, then with probability at least $1 - C \exp(-c_1 \zeta^2 m)$,

$$(5.4) \quad 0.9(1 - \zeta) \|\mathbf{M}\|_F \leq \frac{1}{m} \|\mathcal{A}(\mathbf{M})\|_1 \leq (1 + \zeta) \sqrt{2} \|\mathbf{M}\|_F$$

holds for all symmetric rank 2 matrices $\mathbf{M} \in \mathbb{R}^{n \times n}$. Here, $c_0, c_1, C > 0$ are some universal constants.

PROOF. Since [14, lemma 3.1] already establishes the upper bound, it suffices to prove the lower tail bound. Consider all symmetric rank 2 matrices \mathbf{M} with eigenvalues 1 and $-t$ for some $-1 \leq t \leq 1$. When $t \in [0, 1]$, it has been shown in the proof of [14, lemma 3.2] that with high probability

$$(5.5) \quad \frac{1}{m} \|\mathcal{A}(\mathbf{M})\|_1 \geq (1 - \zeta) f(t)$$

for all such rank 2 matrices \mathbf{M} , where

$$f(t) := \frac{2}{\pi} \left\{ 2\sqrt{t} + (1 - t) \left(\frac{\pi}{2} - 2 \arctan(\sqrt{t}) \right) \right\}.$$

The lower bound in this case can then be justified by noting that $f(t)/\sqrt{1+t^2} \geq 0.9$ for all $t \in [0, 1]$, as illustrated in Figure 5.1. The case where $t \in [-1, 0]$ is an immediate consequence from [14, lemma 3.1]. \square

LEMMA 5.3. Consider any $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ obeying $\|\mathbf{z} - \mathbf{x}\| \leq \delta \|\mathbf{z}\|$ for some $\delta < \frac{1}{2}$. Then one has

$$(5.6) \quad \sqrt{2 - 4\delta} \|\mathbf{z} - \mathbf{x}\| \|\mathbf{z}\| \leq \|\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top\|_F \leq (2 + \delta) \|\mathbf{z} - \mathbf{x}\| \|\mathbf{z}\|.$$

PROOF. Take $\mathbf{h} = \mathbf{z} - \mathbf{x}$ and write

$$\begin{aligned} \|\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top\|_F^2 &= \|-\mathbf{h}\mathbf{z}^\top - \mathbf{z}\mathbf{h}^\top + \mathbf{h}\mathbf{h}^\top\|_F^2 \\ &= \|\mathbf{h}\mathbf{z}^\top + \mathbf{z}\mathbf{h}^\top\|_F^2 + \|\mathbf{h}\mathbf{h}^\top\|_F^2 - 2\langle \mathbf{h}\mathbf{z}^\top + \mathbf{z}\mathbf{h}^\top, \mathbf{h}\mathbf{h}^\top \rangle \\ &= 2\|\mathbf{z}\|^2 \|\mathbf{h}\|^2 + 2|\mathbf{h}^\top \mathbf{z}|^2 + \|\mathbf{h}\mathbf{h}^\top\|_F^2 - 2\|\mathbf{h}\|^2 (\mathbf{h}^\top \mathbf{z} + \mathbf{z}^\top \mathbf{h}). \end{aligned}$$

When $\|\mathbf{h}\| < \frac{1}{2} \|\mathbf{z}\|$, the Cauchy-Schwarz inequality gives

$$(5.7) \quad \begin{aligned} 2\|\mathbf{z}\|^2 \|\mathbf{h}\|^2 - 4\|\mathbf{z}\| \|\mathbf{h}\|^3 &\leq \|\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top\|_F^2 \\ &\leq 4\|\mathbf{z}\|^2 \|\mathbf{h}\|^2 + 4\|\mathbf{h}\|^3 \|\mathbf{z}\| + \|\mathbf{h}\mathbf{h}^\top\|_F^2 \end{aligned}$$

$$(5.8) \quad \begin{aligned} \Rightarrow \sqrt{(2\|\mathbf{z}\| - 4\|\mathbf{h}\|)\|\mathbf{z}\|} \cdot \|\mathbf{h}\| &\leq \|\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top\|_F \\ &\leq (2\|\mathbf{z}\| + \|\mathbf{h}\|) \cdot \|\mathbf{h}\| \end{aligned}$$

as claimed. \square

With probability $1 - \exp(-\Omega(m))$, the above two facts taken together demonstrate that

$$(5.9) \quad 1.15 \|\mathbf{z} - \mathbf{x}\| \|\mathbf{z}\| \leq \frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1 \leq 3 \|\mathbf{z} - \mathbf{x}\| \|\mathbf{z}\|$$

holds simultaneously for all \mathbf{z} and \mathbf{x} satisfying $\|\mathbf{h}\| \leq \frac{1}{11} \|\mathbf{z}\|$. Conditional on (5.9), the inclusion

$$(5.10) \quad \mathcal{E}_3^i \subseteq \mathcal{E}_2^i \subseteq \mathcal{E}_4^i$$

holds with respect to the following events:

$$(5.11) \quad \mathcal{E}_3^i := \{ \left| |\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2 \right| \leq 1.15 \alpha_h \|\mathbf{h}\| \cdot |\mathbf{a}_i^\top \mathbf{z}| \},$$

$$(5.12) \quad \mathcal{E}_4^i := \{ \left| |\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2 \right| \leq 3 \alpha_h \|\mathbf{h}\| \cdot |\mathbf{a}_i^\top \mathbf{z}| \}.$$

The point of introducing these new events is that the \mathcal{E}_3^i 's (respectively, \mathcal{E}_4^i 's) are statistically independent for any fixed \mathbf{x} and \mathbf{z} and are therefore easier to work with.

Note that each \mathcal{E}_3^i (respectively, \mathcal{E}_4^i) is specified by a quadratic inequality. A closer inspection reveals that in order to satisfy these quadratic inequalities, the quantity $\mathbf{a}_i^\top \mathbf{h}$ must fall within two intervals centered around 0 and $2\mathbf{a}_i^\top \mathbf{z}$, respectively. One can thus facilitate analysis by decoupling each quadratic inequality of interest into two simple linear inequalities, as stated in the following lemma.

LEMMA 5.4. For any $\gamma > 0$, define

$$(5.13) \quad \mathcal{D}_\gamma^i := \{ \left| |\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2 \right| \leq \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}| \},$$

$$(5.14) \quad \mathcal{D}_\gamma^{i,1} := \left\{ \frac{|\mathbf{a}_i^\top \mathbf{h}|}{\|\mathbf{h}\|} \leq \gamma \right\},$$

$$(5.15) \quad \text{and } \mathcal{D}_\gamma^{i,2} := \left\{ \left| \frac{\mathbf{a}_i^\top \mathbf{h}}{\|\mathbf{h}\|} - \frac{2\mathbf{a}_i^\top \mathbf{z}}{\|\mathbf{h}\|} \right| \leq \gamma \right\}.$$

Thus, $\mathcal{D}_\gamma^{i,1}$ and $\mathcal{D}_\gamma^{i,2}$ represent the two intervals on $\mathbf{a}_i^\top \mathbf{h}$ centered around 0 and $2\mathbf{a}_i^\top \mathbf{z}$. If $\frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \frac{\alpha_z^{\text{lb}}}{\gamma}$, then the following inclusion holds:

$$(5.16) \quad (\mathcal{D}_{\gamma/(1+\sqrt{2})}^{i,1} \cap \mathcal{E}_1^i) \cup (\mathcal{D}_{\gamma/(1+\sqrt{2})}^{i,2} \cap \mathcal{E}_1^i) \subseteq \mathcal{D}_\gamma^i \cap \mathcal{E}_1^i \\ \subseteq (\mathcal{D}_\gamma^{i,1} \cap \mathcal{E}_1^i) \cup (\mathcal{D}_\gamma^{i,2} \cap \mathcal{E}_1^i).$$

5.2 Proof of the Regularity Condition

By definition, one step towards proving the regularity condition (3.5) is to control the norm of the regularized gradient. In fact, a crude argument already reveals that $\|\frac{1}{m}\nabla \ell_{\text{tr}}(\mathbf{z})\| \lesssim \|\mathbf{h}\|$. To see this, introduce $\mathbf{v} = [v_i]_{1 \leq i \leq m}$ with

$$v_i := 2 \frac{|\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}.$$

It comes from the trimming rule \mathcal{E}_1^i as well as the inclusion property (5.10) that

$$|\mathbf{a}_i^\top \mathbf{z}| \gtrsim \|\mathbf{z}\| \quad \text{and} \quad |y_i - |\mathbf{a}_i^\top \mathbf{x}|^2| \lesssim \frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1 \asymp \|\mathbf{h}\| \|\mathbf{z}\|,$$

implying $|v_i| \lesssim \|\mathbf{h}\|$ and hence $\|\mathbf{v}\| \lesssim \sqrt{m}\|\mathbf{h}\|$. The Marchenko-Pastur law gives $\|\mathbf{A}\| \lesssim \sqrt{m}$, whence

$$(5.17) \quad \frac{1}{m} \|\nabla \ell_{\text{tr}}(\mathbf{z})\| = \frac{1}{m} \|\mathbf{A}^\top \mathbf{v}\| \leq \frac{1}{m} \|\mathbf{A}\| \cdot \|\mathbf{v}\| \lesssim \|\mathbf{h}\|.$$

A more refined estimate will be provided in Lemma 5.9.

The above argument essentially tells us that to establish RC, it suffices to verify a uniform lower bound of the form

$$(5.18) \quad -\left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle \gtrsim \|\mathbf{h}\|^2,$$

as formally derived in the following proposition.

PROPOSITION 5.5. *Consider the noise-free measurements $y_i = |\mathbf{a}_i^\top \mathbf{x}|^2$ and any fixed constant $\epsilon > 0$. Under the condition (2.13), if $m > c_1 n$, then with probability exceeding $1 - C \exp(-c_0 m)$,*

$$(5.19) \quad -\left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle \geq 2\{1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/(9\pi)}\alpha_h^{-1} - \epsilon\} \|\mathbf{h}\|^2$$

holds uniformly over all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ obeying

$$(5.20) \quad \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \min \left\{ \frac{1}{11}, \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{5.7(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}} \right\}.$$

Here, $c_0, c_1, C > 0$ are some universal constants, and ζ_1 and ζ_2 are defined in (2.13).

The basic starting point is noting that $(\mathbf{a}_i^\top \mathbf{z}) - (\mathbf{a}_i^\top \mathbf{x})^2 = (\mathbf{a}_i^\top \mathbf{h})(2\mathbf{a}_i^\top \mathbf{z} - \mathbf{a}_i^\top \mathbf{h})$ and hence

$$(5.21) \quad \begin{aligned} -\frac{1}{2m} \nabla \ell_{\text{tr}}(\mathbf{z}) &= \frac{1}{m} \sum_{i=1}^m \frac{(\mathbf{a}_i^\top \mathbf{z})^2 - (\mathbf{a}_i^\top \mathbf{x})^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} \\ &= \frac{1}{m} \sum_{i=1}^m 2(\mathbf{a}_i^\top \mathbf{h}) \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} - \frac{1}{m} \sum_{i=1}^m \frac{(\mathbf{a}_i^\top \mathbf{h})^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}. \end{aligned}$$

One would expect the contribution of the second term of (5.21) (which is a second-order quantity) to be small as $\|\mathbf{h}\|/\|\mathbf{z}\|$ decreases.

To facilitate analysis, we rewrite (5.21) in terms of the more convenient events $\mathcal{D}_{\gamma}^{i,1}$ and $\mathcal{D}_{\gamma}^{i,2}$. Specifically, the inclusion property (5.10) together with Lemma 5.4 reveals that

$$(5.22) \quad \mathcal{D}_{\gamma_3}^{i,1} \cap \mathcal{E}_1^i \subseteq \mathcal{E}_3^i \cap \mathcal{E}_1^i \subseteq \mathcal{E}_2^i \cap \mathcal{E}_1^i \subseteq \mathcal{E}_4^i \cap \mathcal{E}_1^i \subseteq (\mathcal{D}_{\gamma_4}^{i,1} \cup \mathcal{D}_{\gamma_4}^{i,2}) \cap \mathcal{E}_1^i,$$

where the parameters γ_3, γ_4 are given by

$$(5.23) \quad \gamma_3 := 0.476\alpha_h \quad \text{and} \quad \gamma_4 := 3\alpha_h.$$

This taken collectively with the identity (5.21) leads to a lower estimate

$$(5.24) \quad \begin{aligned} -\left\langle \frac{1}{2m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle &\geq \frac{2}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_3}^{i,1}} - \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{h}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_4}^{i,1}} \\ &\quad - \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{h}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_4}^{i,2}}, \end{aligned}$$

leaving us with three quantities in the right-hand side to deal with. We pause here to explain and compare the influences of these three terms.

To begin with, as long as the trimming step does not discard too many data, the first term should be close to $\frac{2}{m} \sum_i |\mathbf{a}_i^\top \mathbf{h}|^2$, which approximately gives $2\|\mathbf{h}\|^2$ from the law of large numbers. This term turns out to be dominant in the right-hand side of (5.24) as long as $\|\mathbf{h}\|/\|\mathbf{z}\|$ is reasonably small. To see this, please recognize that the second term in the right-hand side is $O(\|\mathbf{h}\|^3/\|\mathbf{z}\|)$, simply because both $\mathbf{a}_i^\top \mathbf{h}$ and $\mathbf{a}_i^\top \mathbf{z}$ are absolutely controlled on $\mathcal{D}_{\gamma_4}^{i,1} \cap \mathcal{E}_1^i$. However, $\mathcal{D}_{\gamma_4}^{i,2}$ does not share such a desired feature. By the very definition of $\mathcal{D}_{\gamma_4}^{i,2}$, each nonzero summand of the last term of (5.24) must obey $|\mathbf{a}_i^\top \mathbf{h}| \approx 2|\mathbf{a}_i^\top \mathbf{z}|$ and therefore $(|\mathbf{a}_i^\top \mathbf{h}|^3/|\mathbf{a}_i^\top \mathbf{z}|) \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_4}^{i,2}}$ is

roughly of the order of $\|\mathbf{z}\|^2$; this could be much larger than our target level $\|\mathbf{h}\|^2$. Fortunately, $\mathcal{D}_{\gamma_4}^{i,2}$ is a rare event, thus precluding a noticeable influence upon the descent direction. All of this is made rigorous in Lemma 5.6 (first term), Lemma 5.7 (second term), and Lemma 5.8 (third term) together with subsequent analysis.

LEMMA 5.6. *Fix $\gamma > 0$, and let \mathcal{E}_1^i and $\mathcal{D}_{\gamma}^{i,1}$ be defined in (2.4) and (5.14), respectively. Set*

$$(5.25) \quad \zeta_1 := 1 - \min \left\{ \mathbb{E} \left[\xi^2 \mathbb{1}_{\{\sqrt{1.01}\alpha_z^{\text{lb}} \leq |\xi| \leq \sqrt{0.99}\alpha_z^{\text{ub}}\}} \right], \right. \\ \left. \mathbb{E} \left[\mathbb{1}_{\{\sqrt{1.01}\alpha_z^{\text{lb}} \leq |\xi| \leq \sqrt{0.99}\alpha_z^{\text{ub}}\}} \right] \right\}$$

$$(5.26) \quad \text{and} \quad \zeta_2 := \mathbb{E} \left[\xi^2 \mathbb{1}_{\{|\xi| > \sqrt{0.99}\gamma\}} \right],$$

where $\xi \sim \mathcal{N}(0, 1)$. For any $\epsilon > 0$, if $m > c_1 n \epsilon^{-2} \log \epsilon^{-1}$, then with probability at least $1 - C \exp(-c_0 \epsilon^2 m)$,

$$(5.27) \quad \frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^{\top} \mathbf{h}|^2 \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}} \geq (1 - \zeta_1 - \zeta_2 - \epsilon) \|\mathbf{h}\|^2$$

holds for all nonzero vectors $\mathbf{h}, \mathbf{z} \in \mathbb{R}^n$. Here, $c_0, c_1, C > 0$ are some universal constants.

We now move on to the second term in the right-hand side of (5.24). For any fixed $\gamma > 0$, the definition of \mathcal{E}_1^i gives rise to an upper estimate

$$(5.28) \quad \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^{\top} \mathbf{h}|^3}{|\mathbf{a}_i^{\top} \mathbf{z}|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}} \leq \frac{1}{\alpha_z^{\text{lb}} \|\mathbf{z}\|} \frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^{\top} \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_{\gamma}^{i,1}} \\ \leq \frac{(1 + \epsilon) \sqrt{8/\pi} \|\mathbf{h}\|^3}{\alpha_z^{\text{lb}} \|\mathbf{z}\|},$$

where $\sqrt{8/\pi} \|\mathbf{h}\|^3$ is exactly the untruncated moment $\mathbb{E}[|\mathbf{a}_i^{\top} \mathbf{h}|^3]$. The second inequality is a consequence of the lemma below, which arises by observing that the summands $|\mathbf{a}_i^{\top} \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_{\gamma}^{i,1}}$ are independent sub-Gaussian random variables.

LEMMA 5.7. *For any constant $\gamma > 0$, if $m/n \geq c_0 \cdot \epsilon^{-2} \log \epsilon^{-1}$, then*

$$(5.29) \quad \frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^{\top} \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_{\gamma}^{i,1}} \leq (1 + \epsilon) \sqrt{8/\pi} \|\mathbf{h}\|^3 \quad \forall \mathbf{h} \in \mathbb{R}^n$$

with probability at least $1 - C \exp(-c_1 \epsilon^2 m)$ for some universal constants $c_0, c_1, C > 0$.

It remains to control the last term of (5.24). As mentioned above, the influence of this term is small since the set of \mathbf{a}_i 's satisfying $\mathcal{D}_{\gamma}^{i,2}$ accounts for a small fraction of measurements. Put formally, the number of equations satisfying $|\mathbf{a}_i^{\top} \mathbf{h}| \geq \gamma \|\mathbf{h}\|$ decays rapidly for large γ (at least at a quadratic rate), as stated below.

LEMMA 5.8. *For any $0 < \epsilon < 1$, there exist some universal constants $c_0, c_1, C > 0$ such that*

$$(5.30) \quad \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq \gamma \|\mathbf{h}\|\}} \leq \frac{1}{0.49\gamma} \exp(-0.485\gamma^2) + \frac{\epsilon}{\gamma^2}$$

$$\forall \mathbf{h} \in \mathbb{R}^n \setminus \{\mathbf{0}\} \text{ and } \gamma \geq 2$$

with probability at least $1 - C \exp(-c_0 \epsilon^2 m)$. This holds with the proviso $m/n \geq c_1 \cdot \epsilon^{-2} \log \epsilon^{-1}$.

To connect this lemma with the last term of (5.24), we recognize that when $\gamma \leq \alpha_z^{\text{lb}} \|\mathbf{z}\| / \|\mathbf{h}\|$, one has

$$(5.31) \quad \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_\gamma^{i,2}} \leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq \alpha_z^{\text{lb}} \|\mathbf{z}\|\}}.$$

The constraint

$$\left| \frac{\mathbf{a}_i^\top \mathbf{h}}{\|\mathbf{h}\|} - \frac{2\mathbf{a}_i^\top \mathbf{z}}{\|\mathbf{h}\|} \right| \leq \gamma$$

of $\mathcal{D}_\gamma^{i,2}$ necessarily requires

$$(5.32) \quad \frac{|\mathbf{a}_i^\top \mathbf{h}|}{\|\mathbf{h}\|} \geq \frac{2|\mathbf{a}_i^\top \mathbf{z}|}{\|\mathbf{h}\|} - \gamma \geq \frac{2\alpha_z^{\text{lb}} \|\mathbf{z}\|}{\|\mathbf{h}\|} - \gamma \geq \frac{\alpha_z^{\text{lb}} \|\mathbf{z}\|}{\|\mathbf{h}\|},$$

where the last inequality comes from our assumption on γ . With Lemma 5.8 in place, (5.31) immediately gives

$$(5.33) \quad \begin{aligned} \sum_{i=1}^m \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_\gamma^{i,2}} &\leq \frac{\|\mathbf{h}\|}{0.49\alpha_z^{\text{lb}} \|\mathbf{z}\|} \exp\left(-0.485\left(\frac{\alpha_z^{\text{lb}} \|\mathbf{z}\|}{\|\mathbf{h}\|}\right)^2\right) + \frac{\epsilon \|\mathbf{h}\|^2}{(\alpha_z^{\text{lb}})^2 \|\mathbf{z}\|^2} \\ &\leq \frac{1}{9800} \left(\frac{\|\mathbf{h}\|}{\alpha_z^{\text{lb}} \|\mathbf{z}\|}\right)^4 + \frac{\epsilon}{(\alpha_z^{\text{lb}})^2} \left(\frac{\|\mathbf{h}\|}{\|\mathbf{z}\|}\right)^2 \end{aligned}$$

as long as $\|\mathbf{h}\|/\|\mathbf{z}\| \leq \alpha_z^{\text{lb}}/6$, where the last inequality uses the majorization $\frac{1}{20000x^4} \geq \frac{1}{x} \exp(-0.485x^2)$ holding for any $x \geq 6$.

In addition, on $\mathcal{E}_1^i \cap \mathcal{D}_\gamma^{i,2}$, the amplitude of each summand can be bounded in such a way that

$$(5.34) \quad \frac{|\mathbf{a}_i^\top \mathbf{h}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \leq \frac{|2\mathbf{a}_i^\top \mathbf{z}| + \gamma \|\mathbf{h}\|}{|\mathbf{a}_i^\top \mathbf{z}|} (2\alpha_z^{\text{ub}} \|\mathbf{z}\| + \gamma \|\mathbf{h}\|)^2$$

$$(5.35) \quad \leq \left(2 + \frac{\gamma}{\alpha_z^{\text{lb}}} \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|}\right) \left(2\alpha_z^{\text{ub}} + \gamma \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|}\right)^2 \|\mathbf{z}\|^2,$$

where both inequalities are immediate consequences from the definitions of $\mathcal{D}_\gamma^{i,2}$ and \mathcal{E}_1^i (see (5.15) and (2.4)). Taking this together with the cardinality bound (5.33)

and picking ϵ appropriately, we get

$$(5.36) \quad \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{h}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_y^{i,2}} \leq \underbrace{\left\{ \frac{\left(2 + \frac{\gamma}{\alpha_z^{\text{lb}}} \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|}\right) \left(2\alpha_z^{\text{ub}} + \gamma \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|}\right)^2}{9800(\alpha_z^{\text{lb}})^4} \frac{\|\mathbf{h}\|^2}{\|\mathbf{z}\|^2} + \epsilon \right\}}_{\vartheta_1} \|\mathbf{h}\|^2.$$

Furthermore, under the condition that

$$\gamma \leq \alpha_z^{\text{lb}} \frac{\|\mathbf{z}\|}{\|\mathbf{h}\|} \quad \text{and} \quad \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \frac{\sqrt{98}(\alpha_z^{\text{lb}})^2}{\sqrt{3}(2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}})},$$

one can simplify (5.36) by observing that $\vartheta_1 \leq \frac{1}{100}$, which results in

$$(5.37) \quad \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{h}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_y^{i,2}} \leq \left(\frac{1}{100} + \epsilon \right) \|\mathbf{h}\|^2.$$

Putting all preceding results in this subsection together reveals that with probability exceeding $1 - \exp(-\Omega(m))$,

$$(5.38) \quad \begin{aligned} -\left\langle \mathbf{h}, \frac{1}{2m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle &\geq \left\{ 1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/\pi} \frac{\|\mathbf{h}\|}{\alpha_z^{\text{lb}} \|\mathbf{z}\|} - 3\epsilon \right\} \|\mathbf{h}\|^2 \\ &\geq \left\{ 1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/\pi} (3\alpha_h)^{-1} - 3\epsilon \right\} \|\mathbf{h}\|^2 \end{aligned}$$

holds simultaneously over all \mathbf{x} and \mathbf{z} satisfying

$$(5.39) \quad \frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \min \left\{ \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98/3}(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}}, \frac{1}{11} \right\}$$

as claimed in Proposition 5.5.

To conclude this section, we provide a tighter estimate about the norm of the regularized gradient.

LEMMA 5.9. *Fix $\delta > 0$ and assume that $y_i = (\mathbf{a}_i^\top \mathbf{x})^2$. Suppose that $m \geq c_0 n$ for some large constant $c_0 > 0$. There exist some universal constants $c, C > 0$ such that with probability at least $1 - C \exp(-cm)$,*

$$(5.40) \quad \frac{1}{m} \|\nabla \ell_{\text{tr}}(\mathbf{z})\| \leq (1 + \delta) \cdot 4\sqrt{1.02 + 0.665/\alpha_h} \|\mathbf{h}\|$$

holds simultaneously for all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ satisfying

$$\frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \min \left\{ \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98/3}(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}}, \frac{1}{11} \right\}.$$

Lemma 5.9 complements the preceding arguments by allowing us to identify a concrete plausible range for the step size. Specifically, putting Lemma 5.9 and Proposition 5.5 together suggests that

$$(5.41) \quad -\left\langle \mathbf{h}, \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\rangle \geq \frac{2\{1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/(9\pi)}\alpha_h^{-1} - \epsilon\}}{(1 + \delta)^2 \cdot 16(1.02 + 0.665/\alpha_h)} \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\|^2.$$

Taking ϵ and δ to be sufficiently small we arrive at a feasible range (cf. Definition (3.5))

$$(5.42) \quad \mu \leq \frac{0.994 - \zeta_1 - \zeta_2 - \sqrt{2/(9\pi)}\alpha_h^{-1}}{2(1.02 + 0.665/\alpha_h)} := \mu_0.$$

This establishes Proposition 5.1 and in turn Theorem 1.1 when μ_t is taken to be a fixed constant.

To justify the contraction under backtracking line search, it suffices to prove that the resulting step size falls within this range (5.42), which we defer to Appendix D.

6 Stability

This section goes in the direction of establishing stability guarantees of TWF. We concentrate on the iterative gradient stage, and defer the analysis for the initialization stage to Appendix C.

Before continuing, we collect two bounds that we shall use several times. The first is the observation that

$$(6.1) \quad \begin{aligned} \frac{1}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{z}\mathbf{z}^\top)\|_1 &\leq \frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1 + \frac{1}{m} \|\boldsymbol{\eta}\|_1 \\ &\lesssim \|\mathbf{h}\| \|\mathbf{z}\| + \frac{1}{m} \|\boldsymbol{\eta}\|_1 \lesssim \|\mathbf{h}\| \|\mathbf{z}\| + \frac{1}{\sqrt{m}} \|\boldsymbol{\eta}\|, \end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. Setting

$$v_i := 2 \frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}$$

as usual, this inequality together with the trimming rules \mathcal{E}_1^i and \mathcal{E}_2^i gives

$$(6.2) \quad \begin{aligned} |v_i| &\lesssim \|\mathbf{h}\| + \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{z}\|} \\ \implies \left\| \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}) \right\| &= \frac{1}{m} \|\mathbf{A}^\top \mathbf{v}\| \leq \left\| \frac{1}{\sqrt{m}} \mathbf{A} \right\| \left\| \frac{1}{\sqrt{m}} \mathbf{v} \right\| \\ &\stackrel{(i)}{\lesssim} \frac{1}{\sqrt{m}} \|\mathbf{v}\| \lesssim \|\mathbf{h}\| + \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{z}\|}, \end{aligned}$$

where (i) arises from [51, cor. 5.35].

As discussed in Section 3, the estimation error is contractive if $-\frac{1}{m}\nabla\ell_{\text{tr}}(\mathbf{z})$ satisfies the regularity condition. With (6.2) in place, RC reduces to

$$(6.3) \quad -\frac{1}{m}\langle \nabla\ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \rangle \gtrsim \|\mathbf{h}\|^2.$$

Unfortunately, (6.3) does not hold for all \mathbf{z} within the neighborhood of \mathbf{x} due to the existence of noise. Instead we establish the following:

- The condition (6.3) holds for all \mathbf{h} obeying

$$(6.4) \quad c_3 \frac{\|\boldsymbol{\eta}\|/\sqrt{m}}{\|\mathbf{z}\|} \leq \|\mathbf{h}\| \leq c_4 \|\mathbf{x}\|$$

for some constants $c_3, c_4 > 0$ (we shall call it *Regime 1*); this will be proved later. In this regime, the reasoning in Section 3 gives

$$(6.5) \quad \text{dist}\left(\mathbf{z} + \frac{\mu}{m}\nabla\ell_{\text{tr}}(\mathbf{z}), \mathbf{x}\right) \leq (1 - \rho) \text{dist}(\mathbf{z}, \mathbf{x})$$

for some appropriate constants $\mu, \rho > 0$, and hence error contraction occurs as in the noiseless setting.

- However, once the iterate enters *Regime 2* where

$$(6.6) \quad \|\mathbf{h}\| \leq \frac{c_3\|\boldsymbol{\eta}\|}{\sqrt{m}\|\mathbf{z}\|},$$

the estimation error might no longer be contractive. Fortunately, in this regime each move by $\frac{\mu}{m}\nabla\ell_{\text{tr}}(\mathbf{z})$ is of size at most $O(\|\boldsymbol{\eta}\|/\sqrt{m}\|\mathbf{z}\|)$; compare (6.2). As a result, at each iteration the estimation error cannot increase by more than a numerical constant times $\|\boldsymbol{\eta}\|/\sqrt{m}\|\mathbf{z}\|$ before possibly jumping out (of this regime). Therefore,

$$(6.7) \quad \text{dist}\left(\mathbf{z} + \frac{\mu}{m}\nabla\ell_{\text{tr}}(\mathbf{z}), \mathbf{x}\right) \leq c_5 \frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\mathbf{x}\|}$$

for some constant $c_5 > 0$. Moreover, as long as $\|\boldsymbol{\eta}\|_\infty/\|\mathbf{x}\|^2$ is sufficiently small, one can guarantee that

$$c_5 \frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\mathbf{x}\|} \leq c_5 \frac{\|\boldsymbol{\eta}\|_\infty}{\|\mathbf{x}\|} \leq c_4 \|\mathbf{x}\|.$$

In other words, if the iterate jumps out of Regime 2, it will still fall within Regime 1.

To summarize, suppose the initial guess $\mathbf{z}^{(0)}$ obeys $\text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq c_4 \|\mathbf{x}\|$. Then the estimation error will shrink at a geometric rate $1 - \rho$ before it enters Regime 2. Afterwards, $\mathbf{z}^{(t)}$ will either stay within Regime 2 or jump back and forth between Regimes 1 and 2. Because of the bounds (6.7) and (6.5), the estimation errors will never exceed the order of $\|\boldsymbol{\eta}\|/\sqrt{m}\|\mathbf{x}\|$ from then on. Putting these together establishes (1.16), namely, the first part of the theorem.

Below we justify the condition (6.3) for Regime 1, for which we start by gathering additional properties of the trimming rules. By Cauchy-Schwarz, $\frac{1}{m}\|\boldsymbol{\eta}\|_1 \leq$

$\frac{1}{\sqrt{m}}\|\boldsymbol{\eta}\| \leq \frac{1}{c_3}\|\mathbf{h}\|\|\mathbf{z}\|$. When c_3 is sufficiently large, applying Lemmas 5.2 and 5.3 gives

$$(6.8) \quad \begin{aligned} \frac{1}{m} \sum_{l=1}^m |y_l - |\mathbf{a}_l^\top \mathbf{z}|^2| &\leq \frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1 + \frac{1}{m} \|\boldsymbol{\eta}\|_1 \leq 2.98 \|\mathbf{h}\|\|\mathbf{z}\|, \\ \frac{1}{m} \sum_{l=1}^m |y_l - |\mathbf{a}_l^\top \mathbf{z}|^2| &\geq \frac{1}{m} \|\mathcal{A}(\mathbf{x}\mathbf{x}^\top - \mathbf{z}\mathbf{z}^\top)\|_1 - \frac{1}{m} \|\boldsymbol{\eta}\|_1 \geq 1.151 \|\mathbf{h}\|\|\mathbf{z}\|. \end{aligned}$$

From now on, we shall denote

$$\tilde{\mathcal{E}}_2^i := \left\{ \left| |\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2 \right| \leq \frac{\alpha_h}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{z}\mathbf{z}^\top)\|_1 \frac{|\mathbf{a}_i^\top \mathbf{z}|}{\|\mathbf{z}\|} \right\}$$

to differentiate from \mathcal{E}_2^i . For any small constant $\epsilon > 0$, we introduce the index set $\mathcal{G} := \{i : |\eta_i| \leq C_\epsilon \|\boldsymbol{\eta}\|/\sqrt{m}\}$ that satisfies $|\mathcal{G}| = (1 - \epsilon)m$. Note that C_ϵ must be bounded as n scales, since

$$(6.9) \quad \begin{aligned} \|\boldsymbol{\eta}\|^2 &\geq \sum_{i \notin \mathcal{G}} \eta_i^2 \geq (m - |\mathcal{G}|) \cdot C_\epsilon^2 \|\boldsymbol{\eta}\|^2 / m \\ &\geq \epsilon C_\epsilon^2 \|\boldsymbol{\eta}\|^2 \end{aligned} \quad \Rightarrow \quad C_\epsilon \leq 1/\sqrt{\epsilon}.$$

We are now ready to analyze the regularized gradient, which we separate into several components as follows:

$$(6.10) \quad \begin{aligned} \nabla_{\text{tr}} \ell(\mathbf{z}) &= 2 \underbrace{\sum_{i \in \mathcal{G}} \frac{|\mathbf{a}_i^\top \mathbf{x}|^2 b^2 - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} + \sum_{i \notin \mathcal{G}} \frac{|\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \tilde{\mathcal{E}}_2^i}}_{:= \nabla_{\text{tr}}^{\text{clean}} \ell(\mathbf{z})} \\ &\quad + 2 \underbrace{\sum_{i \in \mathcal{G}} \frac{\eta_i}{\mathbf{a}_i^\top \mathbf{z}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}}_{:= \nabla_{\text{tr}}^{\text{noise}} \ell(\mathbf{z})} \\ &\quad + 2 \underbrace{\sum_{i \notin \mathcal{G}} \left(\frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} - \frac{|\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \tilde{\mathcal{E}}_2^i} \right) \mathbf{a}_i}_{:= \nabla_{\text{tr}}^{\text{extra}} \ell(\mathbf{z})}. \end{aligned}$$

- For each index $i \in \mathcal{G}$, the inclusion property (5.10) (i.e., $\mathcal{E}_3^i \subseteq \mathcal{E}_2^i \subseteq \mathcal{E}_4^i$) holds. To see this, observe that

$$|y_i - |\mathbf{a}_i^\top \mathbf{z}|^2| \in [||\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2| \pm |\eta_i|].$$

Since $|\eta_i| \leq C_\epsilon \|\boldsymbol{\eta}\|/\sqrt{m} \ll \|\mathbf{h}\|\|\mathbf{z}\|$ when c_3 is sufficiently large, one can derive the inclusion (5.10) immediately from (6.8). As a result, all the proof arguments for Proposition 5.5 carry over to $\nabla_{\text{tr}}^{\text{clean}} \ell(\mathbf{z})$, suggesting that

$$(6.11) \quad -\left\langle \mathbf{h}, \frac{1}{m} \nabla_{\text{tr}}^{\text{clean}} \ell(\mathbf{z}) \right\rangle \geq 2\{1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/(9\pi)} \alpha_h^{-1} - \epsilon\} \|\mathbf{h}\|^2.$$

- Next, letting $w_i = \frac{2\eta_i}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} \mathbb{1}_{\{i \in \mathcal{G}\}}$, we see that for any constant $\delta > 0$, the noise component obeys

$$(6.12) \quad \begin{aligned} \left\| \frac{1}{m} \nabla_{\text{tr}}^{\text{noise}} \ell(\mathbf{z}) \right\| &= \left\| \frac{1}{m} \mathbf{A}^\top \mathbf{w} \right\| \leq \left\| \frac{1}{\sqrt{m}} \mathbf{A} \right\| \left\| \frac{1}{\sqrt{m}} \mathbf{w} \right\| \\ &\stackrel{\text{(ii)}}{\leq} \frac{1+\delta}{\sqrt{m}} \|\mathbf{w}\| \leq (1+\delta) \frac{2\|\boldsymbol{\eta}\|/\sqrt{m}}{\alpha_z^{\text{lb}} \|\mathbf{z}\|} \end{aligned}$$

provided that m/n is sufficiently large. Here, (ii) arises from [51, cor. 5.35], and the last inequality is a consequence of the upper estimate

$$(6.13) \quad \|\mathbf{w}\|^2 \leq 4 \sum_{i=1}^m \frac{|\eta_i|^2}{(\mathbf{a}_i^\top \mathbf{z})^2} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} \leq 4 \sum_{i=1}^m \frac{|\eta_i|^2}{(\alpha_z^{\text{lb}} \|\mathbf{z}\|)^2} = \frac{4\|\boldsymbol{\eta}\|^2}{(\alpha_z^{\text{lb}} \|\mathbf{z}\|)^2}.$$

In turn, this immediately gives

$$(6.14) \quad \left| \left\langle \mathbf{h}, \frac{1}{m} \nabla_{\text{tr}}^{\text{noise}} \ell(\mathbf{z}) \right\rangle \right| \leq \|\mathbf{h}\| \left\| \frac{1}{m} \nabla_{\text{tr}}^{\text{noise}} \ell(\mathbf{z}) \right\| \leq \frac{2(1+\delta)}{\alpha_z^{\text{lb}}} \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{z}\|} \|\mathbf{h}\|.$$

- We now turn to the last term $\nabla_{\text{tr}}^{\text{extra}} \ell(\mathbf{z})$. According to the definition of \mathcal{E}_2^i and $\tilde{\mathcal{E}}_2^i$ as well as the property (6.8), the weight

$$q_i := 2 \left(\frac{y_i - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} - \frac{|\mathbf{a}_i^\top \mathbf{x}|^2 - |\mathbf{a}_i^\top \mathbf{z}|^2}{\mathbf{a}_i^\top \mathbf{z}} \mathbb{1}_{\mathcal{E}_1^i \cap \tilde{\mathcal{E}}_2^i} \right) \mathbb{1}_{\{i \notin \mathcal{G}\}}$$

is bounded in magnitude by $6\|\mathbf{h}\|$. This gives

$$\|\mathbf{q}\| \leq \sqrt{m - |\mathcal{G}|} \cdot 6\|\mathbf{h}\| \leq 6\sqrt{\epsilon m} \|\mathbf{h}\|,$$

and hence

$$(6.15) \quad \begin{aligned} \left| \left\langle \frac{1}{m} \nabla_{\text{tr}}^{\text{extra}} \ell(\mathbf{z}), \mathbf{h} \right\rangle \right| &\leq \|\mathbf{h}\| \cdot \left\| \frac{1}{m} \nabla_{\text{tr}}^{\text{extra}} \ell(\mathbf{z}) \right\| = \frac{1}{m} \|\mathbf{h}\| \cdot \|\mathbf{A}^\top \mathbf{q}\| \\ &\leq 6(1+\delta) \sqrt{\epsilon} \|\mathbf{h}\|^2. \end{aligned}$$

Taking the above bounds together yields

$$\begin{aligned} -\frac{1}{m} \langle \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \rangle &\geq 2 \left\{ 1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{\frac{8}{9\pi}} \frac{1}{\alpha_h} - 6(1+\delta) \sqrt{\epsilon} - \epsilon \right\} \|\mathbf{h}\|^2 \\ &\quad - \frac{2(1+\delta)}{\alpha_z^{\text{lb}}} \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{z}\|} \|\mathbf{h}\|. \end{aligned}$$

Since $\|\mathbf{h}\| \geq c_3 \frac{\|\boldsymbol{\eta}\|}{\sqrt{m} \|\mathbf{z}\|}$ for some large constant $c_3 > 0$, setting ϵ to be small one obtains

$$(6.16) \quad -\frac{1}{m} \langle \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \rangle \geq 2 \{ 1.95 - 2(\zeta_1 + \zeta_2) - \sqrt{8/(9\pi)} \alpha_h^{-1} \} \|\mathbf{h}\|^2$$

for all \mathbf{h} obeying

$$\frac{c_3 \|\boldsymbol{\eta}\| / \sqrt{m}}{\|\mathbf{z}\|} \leq \|\mathbf{h}\| \leq \min \left\{ \frac{1}{11}, \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98/3}(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}} \right\} \|\mathbf{z}\|,$$

which finishes the proof of Theorem 1.3 for general $\boldsymbol{\eta}$.

Up until now, we have established the theorem for general $\boldsymbol{\eta}$, and it remains to specialize it to the Poisson model. Standard concentration results, which we omit, give

$$(6.17) \quad \frac{1}{m} \|\boldsymbol{\eta}\|^2 \approx \frac{1}{m} \sum_{i=1}^m \mathbb{E}[\eta_i^2] = \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{x})^2 \approx \|\mathbf{x}\|^2$$

with high probability. Substitution into (1.16) completes the proof.

7 Minimax Lower Bound

The goal of this section is to establish the minimax lower bound given in Theorem 1.6. For notational simplicity, we denote by $\mathbb{P}(\mathbf{y} \mid \mathbf{w})$ the likelihood of $y_i \stackrel{\text{ind.}}{\sim} \text{Poisson}(|\mathbf{a}_i^\top \mathbf{w}|^2)$, $1 \leq i \leq m$ conditional on $\{\mathbf{a}_i\}$. For any two probability measures P and Q , we denote by $\text{KL}(P \parallel Q)$ the Kullback-Leibler (KL) divergence between them:

$$(7.1) \quad \text{KL}(P \parallel Q) := \int \log \left(\frac{dP}{dQ} \right) dP.$$

The basic idea is to adopt the general reduction scheme discussed in [50, sec. 2.2], which amounts to finding a finite collection of hypotheses that are minimally separated. Below we gather one result useful for constructing and analyzing such hypotheses.

LEMMA 7.1. *Suppose that $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, n is sufficiently large, and $m = \kappa n$ for some sufficiently large constant $\kappa > 0$. Consider any $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$. On an event \mathcal{B} of probability approaching 1, there exists a collection \mathcal{M} of $M = \exp(n/30)$ distinct vectors obeying the following properties:*

- (i) $\mathbf{x} \in \mathcal{M}$;
- (ii) for all $\mathbf{w}^{(l)}, \mathbf{w}^{(j)} \in \mathcal{M}$,

$$(7.2) \quad \frac{1}{\sqrt{8}} - (2n)^{-1/2} \leq \|\mathbf{w}^{(l)} - \mathbf{w}^{(j)}\| \leq \frac{3}{2} + n^{-1/2};$$

- (iii) for all $\mathbf{w} \in \mathcal{M}$,

$$(7.3) \quad \frac{|\mathbf{a}_i^\top (\mathbf{w} - \mathbf{x})|^2}{|\mathbf{a}_i^\top \mathbf{x}|^2} \leq \frac{\|\mathbf{w} - \mathbf{x}\|^2}{\|\mathbf{x}\|^2} \{2 + 17 \log^3 m\}, \quad 1 \leq i \leq m.$$

In words, Lemma 7.1 constructs a set \mathcal{M} of exponentially many vectors/hypotheses scattered around \mathbf{x} and yet well separated. From (ii) we see that each pair of hypotheses in \mathcal{M} is separated by a distance roughly on the order of 1, and all hypotheses reside within a spherical ball centered at \mathbf{x} of radius $\frac{3}{2} + o(1)$. When $\|\mathbf{x}\| \geq \log^{1.5} m$, every hypothesis $\mathbf{w} \in \mathcal{M}$ satisfies $\|\mathbf{w}\| \approx \|\mathbf{x}\| \gg 1$. In addition, (iii) says that the quantities $|\mathbf{a}_i^\top(\mathbf{w} - \mathbf{x})|/|\mathbf{a}_i^\top \mathbf{x}|$ are all very well controlled (modulo some logarithmic factor). In particular, when $\|\mathbf{x}\| \geq \log^{1.5} m$, one must have

$$(7.4) \quad \frac{|\mathbf{a}_i^\top(\mathbf{w} - \mathbf{x})|^2}{|\mathbf{a}_i^\top \mathbf{x}|^2} \lesssim \frac{\|\mathbf{w} - \mathbf{x}\|^2}{\|\mathbf{x}\|^2} \log^3 m \lesssim \frac{1}{\log^3 m} \log^3 m \lesssim 1.$$

In the Poisson model, such a quantity turns out to be crucial in controlling the information divergence between two hypotheses, as demonstrated in the following lemma.

LEMMA 7.2. *Fix a family of design vectors $\{\mathbf{a}_i\}$. Then for any \mathbf{w} and $\mathbf{r} \in \mathbb{R}^n$,*

$$(7.5) \quad \text{KL}(\mathbb{P}(\mathbf{y} | \mathbf{w} + \mathbf{r}) \| \mathbb{P}(\mathbf{y} | \mathbf{w})) \leq \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{r}|^2 \left(8 + \frac{2|\mathbf{a}_i^\top \mathbf{r}|^2}{|\mathbf{a}_i^\top \mathbf{w}|^2} \right).$$

Lemma 7.2 and (7.4) taken collectively suggest that on the event $\mathcal{B} \cap \mathcal{C}$ (\mathcal{B} is in Lemma 7.1 and $\mathcal{C} := \{\|\mathbf{A}\| \leq \sqrt{2m}\}$), the conditional KL divergence (we condition on the \mathbf{a}_i 's obeys

$$(7.6) \quad \begin{aligned} \text{KL}(\mathbb{P}(\mathbf{y} | \mathbf{w}) \| \mathbb{P}(\mathbf{y} | \mathbf{x})) &\leq c_3 \sum_{i=1}^m |\mathbf{a}_i^\top(\mathbf{w} - \mathbf{x})|^2 \\ &\leq 2c_3 m \|\mathbf{w} - \mathbf{x}\|^2 \quad \forall \mathbf{w} \in \mathcal{M}; \end{aligned}$$

here, the inequality holds for some constant $c_3 > 0$ provided that $\|\mathbf{x}\| \geq \log^{1.5} m$, and the last inequality is a result of \mathcal{C} (which occurs with high probability). We now use hypotheses as in Lemma 7.1 but rescaled in such a way that

$$(7.7) \quad \|\mathbf{w} - \mathbf{x}\| \asymp \delta \quad \text{and} \quad \|\mathbf{w} - \tilde{\mathbf{w}}\| \asymp \delta \quad \forall \mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{M} \text{ with } \mathbf{w} \neq \tilde{\mathbf{w}}$$

for some $0 < \delta < 1$. This is achieved via the substitution $\mathbf{w} \leftarrow \mathbf{x} + \delta(\mathbf{w} - \mathbf{x})$; with a slight abuse of notation, \mathcal{M} denotes the new set.

The hardness of a minimax estimation problem is known to be dictated by information divergence inequalities such as (7.6). Indeed, suppose that

$$(7.8) \quad \frac{1}{M-1} \sum_{\mathbf{w} \in \mathcal{M} \setminus \{\mathbf{x}\}} \text{KL}(\mathbb{P}(\mathbf{y} | \mathbf{w}) \| \mathbb{P}(\mathbf{y} | \mathbf{x})) \leq \frac{1}{10} \log(M-1)$$

holds; then the Fano-type minimax lower bound [50, theorem 2.7] asserts that

$$(7.9) \quad \inf_{\hat{\mathbf{x}}} \sup_{\mathbf{x} \in \mathcal{M}} \mathbb{E}[\|\hat{\mathbf{x}} - \mathbf{x}\| | \{\mathbf{a}_i\}] \gtrsim \min_{\substack{\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{M}, \\ \mathbf{w} \neq \tilde{\mathbf{w}}}} \|\mathbf{w} - \tilde{\mathbf{w}}\|.$$

Since $M = \exp(n/30)$, (7.8) would follow from

$$(7.10) \quad 2c_3 \|\mathbf{w} - \mathbf{x}\|^2 \leq \frac{n}{300m}, \quad \mathbf{w} \in \mathcal{M}.$$

Hence we just need to select δ to be a small multiple of $\sqrt{n/m}$. This in turn gives

$$(7.11) \quad \inf_{\hat{\mathbf{x}}} \sup_{\mathbf{x} \in \mathcal{M}} \mathbb{E}[\|\hat{\mathbf{x}} - \mathbf{x}\| \mid \{\mathbf{a}_i\}] \gtrsim \min_{\substack{\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{M}, \\ \mathbf{w} \neq \tilde{\mathbf{w}}}} \|\mathbf{w} - \tilde{\mathbf{w}}\| \gtrsim \sqrt{n/m}.$$

Finally, it remains to connect $\|\hat{\mathbf{x}} - \mathbf{x}\|$ with $\text{dist}(\hat{\mathbf{x}}, \mathbf{x})$. Since all the $\mathbf{w} \in \mathcal{M}$ are clustered around \mathbf{x} and are at a mutual distance about δ that is much smaller than $\|\mathbf{x}\|$, we can see that for any reasonable estimator, $\text{dist}(\hat{\mathbf{x}}, \mathbf{x}) = \|\hat{\mathbf{x}} - \mathbf{x}\|$. This finishes the proof.

8 Discussion

To keep our treatment concise, this paper does not strive to explore all possible generalizations of the theory. There are nevertheless a few extensions worth pointing out.

- **MORE GENERAL OBJECTIVE FUNCTIONS.** For concreteness, we restrict our analysis to the Poisson log-likelihood function, but the analysis framework we laid out easily carries over to a broad class of (nonconvex) objective functions. For instance, all results continue to hold if we replace the Poisson log-likelihood by the Gaussian log-likelihood, that is, the polynomial function $-\sum_{i=1}^m (y_i - |\mathbf{a}_i^\top \mathbf{z}|^2)^2$ studied in [13].

A general guideline is to first check whether the expected regularity condition

$$\mathbb{E}\left[-\left\langle \frac{1}{m} \nabla \ell_{\text{tr}}(\mathbf{z}), \mathbf{h} \right\rangle\right] \gtrsim \|\mathbf{h}\|^2$$

holds for any fixed \mathbf{z} within a neighborhood around \mathbf{x} . If so, then often times RC holds uniformly within this neighborhood due to sharp concentration of measure ensured by the regularization procedure.

- **SUB-GAUSSIAN MEASUREMENTS.** The theory extends to the situation where the \mathbf{a}_i 's are i.i.d. sub-Gaussian random vectors, although the truncation threshold might need to be tweaked based on the sub-Gaussian norm of \mathbf{a}_i . A more challenging scenario, however, is the case where the \mathbf{a}_i 's are generated according to the CDP model, since there is much less randomness to exploit in the mathematical analysis. We leave this to future research.

Having demonstrated the power of TWF in recovering a rank 1 matrix $\mathbf{x}\mathbf{x}^*$ from quadratic equations, we remark on the potential of TWF for recovering low-rank matrices from rank 1 measurements.

Imagine that we wish to estimate a rank- r matrix $\mathbf{X} \succeq \mathbf{0}$ and that all we know about \mathbf{X} is

$$y_i = \mathbf{a}_i^\top \mathbf{X} \mathbf{a}_i, \quad 1 \leq i \leq m.$$

It is known that this problem can be efficiently solved by using more computationally intensive semidefinite programs [9, 16]. With the hope of developing a linear-time algorithm, one might consider a modified TWF scheme, which would maintain a rank- r matrix variable and operate as follows: perform truncated spectral initialization and then successively update the current guess via a regularized gradient descent rule applied to a presumed log-likelihood function.

Moving away from i.i.d. sub-Gaussian measurements, there is a proliferation of problems that involve completion of a low-rank matrix X from partial entries, where the rank is known a priori. It is self-evident that such entrywise observations can also be cast as rank 1 measurements of X . Therefore, the preceding modified TWF may add to recent literature in applying nonconvex schemes for low-rank matrix completion [27, 30, 31, 46], robust principal component analysis (PCA) [37], or even a broader family of latent-variable models (e.g., dictionary learning [47, 48], sparse coding [2], and mixture problems [4, 55]).

A concrete application of this flavor is a simple form of the fundamental alignment/matching problem [7, 17, 28]. Imagine a collection of n instances, each representing an image of the same physical object but with different shifts $r_i \in \{0, \dots, M-1\}$. The goal is to align all these instances from observations on the relative shift between pairs of them. Denoting by X_i the cyclic shift by an amount r_i of \mathbf{I}_M , one sees that the collection matrix $X := [X_i^\top X_j]_{1 \leq i, j \leq k}$ is a rank- M matrix, and the relative shift observations can be treated as rank 1 measurements of X . Running TWF over this problem instance might result in a statistically and computationally efficient solution. This would be of great practical interest.

Appendix A Proofs for Section 5

A.1 Proof of Lemma 5.4

First, we make the observation that $(\mathbf{a}_i^\top \mathbf{z})^2 - (\mathbf{a}_i^\top \mathbf{x})^2 = (2\mathbf{a}_i^\top \mathbf{z} - \mathbf{a}_i^\top \mathbf{h})\mathbf{a}_i^\top \mathbf{h}$ is a quadratic function in $\mathbf{a}_i^\top \mathbf{h}$. If we assume $\gamma \leq \alpha_z^{\text{lb}} \|\mathbf{z}\| / \|\mathbf{h}\|$, then on the event \mathcal{E}_1^i one has

$$(A.1) \quad (\mathbf{a}_i^\top \mathbf{z})^2 \geq \alpha_z^{\text{lb}} \|\mathbf{z}\| \cdot |\mathbf{a}_i^\top \mathbf{z}| \geq \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|.$$

Solving the quadratic inequality that specifies \mathcal{D}_γ^i gives

$$\begin{aligned} \mathbf{a}_i^\top \mathbf{h} &\in \left[\mathbf{a}_i^\top \mathbf{z} - \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 + \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}, \mathbf{a}_i^\top \mathbf{z} - \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 - \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|} \right], \\ \text{or } \mathbf{a}_i^\top \mathbf{h} &\in \left[\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 - \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}, \mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 + \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|} \right], \end{aligned}$$

which we will simplify in what follows.

Suppose for the moment that $\mathbf{a}_i^\top \mathbf{z} \geq 0$; then the preceding two intervals are respectively equivalent to

$$\begin{aligned} \mathbf{a}_i^\top \mathbf{h} &\in \underbrace{\left[\frac{-\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 + \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}}, \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 - \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}} \right]}_{:=I_1}, \\ \mathbf{a}_i^\top \mathbf{h} - 2\mathbf{a}_i^\top \mathbf{z} &\in \underbrace{\left[\frac{-\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 - \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}}, \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 + \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}} \right]}_{:=I_2}. \end{aligned}$$

Assuming (A.1) and making use of the observations

$$\begin{aligned} \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 - \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}} &\leq \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z}} = \gamma \|\mathbf{h}\| \\ \text{and } \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{\mathbf{a}_i^\top \mathbf{z} + \sqrt{(\mathbf{a}_i^\top \mathbf{z})^2 + \gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}} &\geq \frac{\gamma \|\mathbf{h}\| |\mathbf{a}_i^\top \mathbf{z}|}{(1 + \sqrt{2}) |\mathbf{a}_i^\top \mathbf{z}|} = \frac{\gamma}{1 + \sqrt{2}} \|\mathbf{h}\|, \end{aligned}$$

we obtain the inner and outer bounds

$$[\pm(1 + \sqrt{2})^{-1} \gamma \|\mathbf{h}\|] \subseteq I_1, I_2 \subseteq [\pm \gamma \|\mathbf{h}\|].$$

Setting $\gamma_1 := \frac{\gamma}{1 + \sqrt{2}}$ gives

$$(\mathcal{D}_{\gamma_1}^{i,1} \cap \mathcal{E}_{i,1}) \cup (\mathcal{D}_{\gamma_1}^{i,2} \cap \mathcal{E}_{i,1}) \subseteq \mathcal{D}_\gamma \cap \mathcal{E}_{i,1} \subseteq (\mathcal{D}_\gamma^{i,1} \cap \mathcal{E}_{i,1}) \cup (\mathcal{D}_\gamma^{i,2} \cap \mathcal{E}_{i,1}).$$

Proceeding with the same argument, we can derive exactly the same inner and outer bounds in the regime where $\mathbf{a}_i^\top \mathbf{z} < 0$, concluding the proof.

A.2 Proof of Lemma 5.6

By homogeneity, it suffices to establish the claim for the case where both \mathbf{h} and \mathbf{z} are *unit vectors*.

Suppose for the moment that \mathbf{h} and \mathbf{z} are *statistically independent* from $\{\mathbf{a}_i\}$. We introduce two auxiliary Lipschitz functions approximating indicator functions:

$$(A.2) \quad \chi_z(\tau) := \begin{cases} 1 & \text{if } |\tau| \in [\sqrt{1.01}\alpha_z^{\text{lb}}, \sqrt{0.99}\alpha_z^{\text{ub}}], \\ -100(\alpha_z^{\text{ub}})^{-2}\tau^2 + 100 & \text{if } |\tau| \in [\sqrt{0.99}\alpha_z^{\text{ub}}, \alpha_z^{\text{ub}}], \\ 100(\alpha_z^{\text{lb}})^{-2}\tau^2 - 100 & \text{if } |\tau| \in [\alpha_z^{\text{lb}}, \sqrt{1.01}\alpha_z^{\text{lb}}], \\ 0 & \text{otherwise.} \end{cases}$$

$$(A.3) \quad \chi_h(\tau) := \begin{cases} 1 & \text{if } |\tau| \in [0, \sqrt{0.99}\gamma], \\ -\frac{100}{\gamma^2}\tau^2 + 100 & \text{if } |\tau| \in [\sqrt{0.99}\gamma, \gamma], \\ 0 & \text{otherwise.} \end{cases}$$

Since \mathbf{h} and \mathbf{z} are assumed to be unit vectors, these two functions obey

$$(A.4) \quad 0 \leq \chi_z(\mathbf{a}_i^\top \mathbf{z}) \leq \mathbb{1}_{\mathcal{E}_1^i} \quad \text{and} \quad 0 \leq \chi_h(\mathbf{a}_i^\top \mathbf{h}) \leq \mathbb{1}_{\mathcal{D}_Y^{i,1}},$$

and thus

$$(A.5) \quad \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_Y^{i,1}} \geq \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h}).$$

We proceed to lower bound $\frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h})$.

Firstly, to compute the mean of $(\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h})$, we introduce an auxiliary orthonormal matrix

$$(A.6) \quad \mathbf{U}_z = \begin{bmatrix} \mathbf{z}^\top / \|\mathbf{z}\| \\ \vdots \end{bmatrix}$$

whose first row is along the direction of \mathbf{z} , and set

$$(A.7) \quad \tilde{\mathbf{h}} := \mathbf{U}_z \mathbf{h} \quad \text{and} \quad \tilde{\mathbf{a}}_i := \mathbf{U}_z \mathbf{a}_i.$$

Also, denote by $\tilde{a}_{i,1}$ (respectively, \tilde{h}_1) the first entry of $\tilde{\mathbf{a}}_i$ (respectively, $\tilde{\mathbf{h}}$), and $\tilde{\mathbf{a}}_{i,\setminus 1}$ (respectively, $\tilde{\mathbf{h}}_{\setminus 1}$) the remaining entries of $\tilde{\mathbf{a}}_i$ (respectively, $\tilde{\mathbf{h}}$), and let $\xi \sim \mathcal{N}(0, 1)$. We have

$$(A.8) \quad \begin{aligned} & \mathbb{E}[(\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h})] \\ & \geq \mathbb{E}[(\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z})] - \mathbb{E}[(\mathbf{a}_i^\top \mathbf{h})^2 (1 - \chi_h(\mathbf{a}_i^\top \mathbf{h}))] \\ & \geq \mathbb{E}[(\tilde{a}_{i,1} \tilde{h}_1)^2 \chi_z(\mathbf{a}_i^\top \mathbf{z})] + \mathbb{E}[(\tilde{\mathbf{a}}_{i,\setminus 1}^\top \tilde{\mathbf{h}}_{\setminus 1})^2] \mathbb{E}[\chi_z(\mathbf{a}_i^\top \mathbf{z})] \\ & \quad - \|\mathbf{h}\|^2 \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| > \sqrt{0.99}\gamma\}}] \\ & \geq |\tilde{h}_1|^2 (1 - \zeta_1) + \|\tilde{\mathbf{h}}_{\setminus 1}\|^2 (1 - \zeta_1) - \zeta_2 \|\mathbf{h}\|^2 \\ & \geq (1 - \zeta_1 - \zeta_2) \|\mathbf{h}\|^2, \end{aligned}$$

where identity (A.8) arises from (5.25) and (5.26). Since $(\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h})$ is bounded in magnitude by $\gamma^2 \|\mathbf{h}\|^2$, it is a sub-Gaussian random variable with sub-Gaussian norm $O(\gamma^2 \|\mathbf{h}\|^2)$. Apply the Hoeffding-type inequality [51, prop. 5.10] to deduce that for any $\epsilon > 0$,

$$(A.9) \quad \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h}) \geq \mathbb{E}[(\mathbf{a}_i^\top \mathbf{h})^2 \chi_z(\mathbf{a}_i^\top \mathbf{z}) \chi_h(\mathbf{a}_i^\top \mathbf{h})] - \epsilon \|\mathbf{h}\|^2$$

$$(A.10) \quad \geq (1 - \zeta_1 - \zeta_2 - \epsilon) \|\mathbf{h}\|^2$$

with probability at least $1 - \exp(-\Omega(\epsilon^2 m))$.

The next step is to obtain uniform control over all *unit vectors*, for which we adopt a basic version of an ϵ -net argument. Specifically, we construct an ϵ -net \mathcal{N}_ϵ with cardinality $|\mathcal{N}_\epsilon| \leq (1 + 2/\epsilon)^{2n}$ (cf. [51]) such that for any (\mathbf{h}, \mathbf{z}) with $\|\mathbf{h}\| = \|\mathbf{z}\| = 1$, there exists a pair $\mathbf{h}_0, \mathbf{z}_0 \in \mathcal{N}_\epsilon$ satisfying $\|\mathbf{h} - \mathbf{h}_0\| \leq \epsilon$ and

$\|z - z_0\| \leq \epsilon$. Now that we have discretized the unit spheres using a finite set, taking the union bound gives

$$(A.11) \quad \frac{1}{m} \sum_{i=1}^m (a_i^\top h_0)^2 \chi_z(a_i^\top z_0) \chi_h(a_i^\top h_0) \geq (1 - \zeta_1 - \zeta_2 - \epsilon) \|h_0\|^2 \quad \forall h_0, z_0 \in \mathcal{N}_\epsilon$$

with probability at least $1 - (1 + 2/\epsilon)^{2n} \exp(-\Omega(\epsilon^2 m))$.

Define $f_1(\cdot)$ and $f_2(\cdot)$ such that $f_1(\tau) := \tau \chi_h(\sqrt{\tau})$ and $f_2(\tau) := \chi_z(\sqrt{\tau})$, which are both bounded functions with Lipschitz constant $O(1)$. This guarantees that for each *unit* vector pair h and z ,

$$\begin{aligned} & |(a_i^\top h)^2 \chi_z(a_i^\top z) \chi_h(a_i^\top h) - (a_i^\top h_0)^2 \chi_z(a_i^\top z_0) \chi_h(a_i^\top h_0)| \\ & \leq |\chi_h(a_i^\top z)| \cdot |(a_i^\top h)^2 \chi_h(a_i^\top h) - (a_i^\top h_0)^2 \chi_h(a_i^\top h_0)| \\ & \quad + |(a_i^\top h_0)^2 \chi_h(a_i^\top h_0)| \cdot |\chi_h(a_i^\top z) - \chi_h(a_i^\top z_0)| \\ & \leq |\chi_h(a_i^\top z)| \cdot |f_1(|a_i^\top h|^2) - f_1(|a_i^\top h_0|^2)| \\ & \quad + |(a_i^\top h_0)^2 \chi_h(a_i^\top h_0)| \cdot |f_2(|a_i^\top z|^2) - f_2(|a_i^\top z_0|^2)| \\ & \lesssim |(a_i^\top h)^2 - (a_i^\top h_0)^2| + (a_i^\top z)^2 - (a_i^\top z_0)^2|. \end{aligned}$$

Consequently, there exists some universal constant $c_3 > 0$ such that

$$\begin{aligned} & \left| \frac{1}{m} \sum_{i=1}^m (a_i^\top h)^2 \chi_z(a_i^\top z) \chi_h(a_i^\top h) - \frac{1}{m} \sum_{i=1}^m (a_i^\top h_0)^2 \chi_z(a_i^\top z_0) \chi_h(a_i^\top h_0) \right| \\ & \lesssim \frac{1}{m} \|\mathcal{A}(hh^\top - h_0 h_0^\top)\|_1 + \frac{1}{m} \|\mathcal{A}(zz^\top - z_0 z_0^\top)\|_1 \\ & \stackrel{(i)}{\leq} c_3 \{\|hh^\top - h_0 h_0^\top\|_F + \|zz^\top - z_0 z_0^\top\|_F\} \\ & \stackrel{(ii)}{\leq} 2.5c_3 \{\|h - h_0\| \cdot \|h\| + \|z - z_0\| \cdot \|z\|\} \leq 5c_3 \epsilon, \end{aligned}$$

where (i) results from Lemma 5.2, and (ii) arises from Lemma 5.3 whenever $\epsilon < \frac{1}{2}$.

With the assertion (A.11) in place, we see that with high probability,

$$\frac{1}{m} \sum_{i=1}^m (a_i^\top h)^2 \chi_z(a_i^\top z) \chi_h(a_i^\top h) \geq (1 - \zeta_1 - \zeta_2 - (5c_3 + 1)\epsilon) \|h\|^2$$

for all unit vectors h and z . Since ϵ can be arbitrary, putting this and (A.5) together completes the proof.

A.3 Proof of Lemma 5.7

The proof makes use of standard concentration of measure and covering arguments, and it suffices to restrict our attention to *unit vectors* h . We find it convenient

to work with an auxiliary function

$$\chi_2(\tau) = \begin{cases} |\tau|^{3/2} & \text{if } |\tau| \leq \gamma^2, \\ -\gamma(|\tau| - \gamma^2) + \gamma^3 & \text{if } \gamma^2 < |\tau| \leq 2\gamma^2, \\ 0 & \text{otherwise.} \end{cases}$$

Apparently, $\chi_2(\tau)$ is a Lipschitz function of τ with Lipschitz norm $O(\gamma)$. Recalling the definition of $\mathcal{D}_\gamma^{i,1}$, we see that each summand is bounded above by

$$|\mathbf{a}_i^\top \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_\gamma^{i,1}} \leq \chi_2(|\mathbf{a}_i^\top \mathbf{h}|^2).$$

For each fixed \mathbf{h} and $\epsilon > 0$, applying the Bernstein inequality [51, prop. 5.16] gives

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_\gamma^{i,1}} &\leq \frac{1}{m} \sum_{i=1}^m \chi_2(|\mathbf{a}_i^\top \mathbf{h}|^2) \leq \mathbb{E}[\chi_2(|\mathbf{a}_i^\top \mathbf{h}|^2)] + \epsilon \\ &\leq \mathbb{E}[|\mathbf{a}_i^\top \mathbf{h}|^3] + \epsilon = \sqrt{8/\pi} + \epsilon \end{aligned}$$

with probability exceeding $1 - \exp(-\Omega(\epsilon^2 m))$.

From [51, lemma 5.2], there exists an ϵ -net \mathcal{N}_ϵ of the unit sphere with cardinality $|\mathcal{N}_\epsilon| \leq (1 + \frac{2}{\epsilon})^n$. For each \mathbf{h} , suppose that $\|\mathbf{h}_0 - \mathbf{h}\| \leq \epsilon$ for some $\mathbf{h}_0 \in \mathcal{N}_\epsilon$. The Lipschitz property of χ_2 implies

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \{\chi_2(|\mathbf{a}_i^\top \mathbf{h}|^2) - \chi_2(|\mathbf{a}_i^\top \mathbf{h}_0|^2)\} &\lesssim \frac{1}{m} \sum_{i=1}^m \left| |\mathbf{a}_i^\top \mathbf{h}|^2 - |\mathbf{a}_i^\top \mathbf{h}_0|^2 \right| \\ &\stackrel{(i)}{\asymp} \|\mathbf{h} - \mathbf{h}_0\| \|\mathbf{h}\| \asymp \epsilon, \end{aligned}$$

where (i) arises by combining Lemmas 5.2 and 5.3. This demonstrates that with high probability,

$$\frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{h}|^3 \mathbb{1}_{\mathcal{D}_\gamma^{i,1}} \leq \frac{1}{m} \sum_{i=1}^m \chi_2(|\mathbf{a}_i^\top \mathbf{h}|^2) \leq \sqrt{8/\pi} + O(\epsilon)$$

for all unit vectors \mathbf{h} , as claimed.

A.4 Proof of Lemma 5.8

Without loss of generality, the proof focuses on the case where $\|\mathbf{h}\| = 1$. Fix an arbitrary small constant $\delta > 0$. One can eliminate the difficulty of handling the discontinuous indicator functions by working with the following auxiliary function:

$$(A.12) \quad \chi_3(\tau, \gamma) := \begin{cases} 1 & \text{if } \sqrt{\tau} \geq \psi_{\text{lb}}(\gamma), \\ \frac{100\tau}{\psi_{\text{lb}}^2(\gamma)} - 99 & \text{if } \sqrt{\tau} \in [\sqrt{0.99}\psi_{\text{lb}}(\gamma), \psi_{\text{lb}}(\gamma)], \\ 0 & \text{otherwise.} \end{cases}$$

Here, $\psi_{\text{lb}}(\cdot)$ is a piecewise constant function defined as

$$\psi_{\text{lb}}(\gamma) := (1 + \delta) \left\lfloor \frac{\log \gamma}{\log(1 + \delta)} \right\rfloor,$$

which clearly satisfy $\frac{\gamma}{1+\delta} \leq \psi_{\text{lb}}(\gamma) \leq \gamma$. Such a function is useful for our purpose since for any $0 < \delta \leq 0.005$,

$$(A.13) \quad \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq \gamma\}} \leq \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma) \leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq \sqrt{0.99} \psi_{\text{lb}}(\gamma)\}} \leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}}.$$

For any fixed unit vector \mathbf{h} , the above argument leads to an upper tail estimate: for any $0 < t \leq 1$,

$$\begin{aligned} \mathbb{P}\{\chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma) \geq t\} &\leq \mathbb{P}\{\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}} \geq t\} = \mathbb{P}\{\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}} = 1\} \\ &= 2 \Delta_{0.99\gamma}^\infty \phi(x) dx \leq \frac{2}{0.99\gamma} \phi(0.99\gamma), \end{aligned}$$

where $\phi(x)$ is the density of a standard normal, and (A.14) follows from the tail bound $\Delta_x^\infty \phi(x) dx \leq \frac{1}{x} \phi(x)$ for all $x > 0$. This implies that when $\gamma \geq 2$, both $\chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma)$ and $\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}}$ are subexponential with subexponential norm $O(\gamma^{-2})$ (cf. [51, def. 5.13]). We apply the Bernstein-type inequality for the sum of subexponential random variables [51, cor. 5.17], which indicates that for any fixed \mathbf{h} and γ as well as any sufficiently small $\epsilon \in (0, 1)$,

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma) &\leq \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}} \leq \mathbb{E}[\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq 0.99\gamma\}}] + \epsilon \frac{1}{\gamma^2} \\ &\leq \frac{2}{0.99\gamma} \exp(-0.49\gamma^2) + \epsilon \frac{1}{\gamma^2} \end{aligned}$$

holds with probability exceeding $1 - \exp(-\Omega(\epsilon^2 m))$.

We now proceed to obtain uniform control over all \mathbf{h} and $2 \leq \gamma \leq 2^n$. To begin with, we consider all $2 \leq \gamma \leq m$ and construct an ϵ -net \mathcal{N}_ϵ over the unit sphere such that: (i) $|\mathcal{N}_\epsilon| \leq (1 + \frac{2}{\epsilon})^n$; (ii) for any \mathbf{h} with $\|\mathbf{h}\| = 1$, there exists a unit vector $\mathbf{h}_0 \in \mathcal{N}_\epsilon$ obeying $\|\mathbf{h} - \mathbf{h}_0\| \leq \epsilon$. Taking the union bound gives the following: with probability at least $1 - \frac{\log m}{\log(1+\delta)} (1 + \frac{2}{\epsilon})^n \exp(-\Omega(\epsilon^2 m))$,

$$\frac{1}{m} \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}_0|^2, \gamma_0) \leq (0.495\gamma_0)^{-1} \exp(-0.49\gamma_0^2) + \epsilon\gamma_0^{-2}$$

holds simultaneously for all $\mathbf{h}_0 \in \mathcal{N}_\epsilon$ and $\gamma_0 \in \{(1 + \delta)^k \mid 1 \leq k \leq \frac{\log m}{\log(1+\delta)}\}$.

Note that $\chi_3(\tau, \gamma_0)$ is a Lipschitz function in τ with the Lipschitz constant bounded above by $100/\psi_{\text{lb}}^2(\gamma_0)$. With this in mind, for any (\mathbf{h}, γ) with $\|\mathbf{h}\| = 1$ and $\gamma_0 := (1 + \delta)^k \leq \gamma < (1 + \delta)^{k+1}$, one has

$$\begin{aligned} |\chi_3(|\mathbf{a}_i^\top \mathbf{h}_0|^2, \gamma_0) - \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma)| &= |\chi_3(|\mathbf{a}_i^\top \mathbf{h}_0|^2, \gamma_0) - \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma_0)| \\ &\leq \frac{100}{\psi_{\text{lb}}^2(\gamma_0)} \left| |\mathbf{a}_i^\top \mathbf{h}|^2 - |\mathbf{a}_i^\top \mathbf{h}_0|^2 \right|. \end{aligned}$$

It then follows from Lemmas 5.2 and 5.3 that

$$\begin{aligned} & \frac{1}{m} \left| \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}_0|^2, \gamma_0) - \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma) \right| \\ & \leq \frac{100}{\psi_{\text{lb}}^2(\gamma_0)} \frac{1}{m} \|\mathbf{A}(\mathbf{h}\mathbf{h}^\top - \mathbf{h}_0\mathbf{h}_0^\top)\|_1 \\ & \leq \frac{250(1+\delta)^2}{\gamma^2} \|\mathbf{h} - \mathbf{h}_0\| \|\mathbf{h}\| \leq \frac{250(1+\delta)^2 \epsilon}{\gamma^2}. \end{aligned}$$

Putting the above results together gives that for all $2 \leq \gamma \leq (1+\delta)^{\frac{\log m}{\log(1+\delta)}} = m$,

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}|^2, \gamma) & \leq \frac{1}{m} \sum_{i=1}^m \chi_3(|\mathbf{a}_i^\top \mathbf{h}_0|^2, \gamma_0) + \frac{250(1+\delta)^2}{\gamma^2} \epsilon \\ & \leq \frac{1}{0.495\gamma_0} \exp(-0.49\gamma_0^2) + 251(1+\delta)^2 \frac{\epsilon}{\gamma^2} \\ & \leq \frac{1}{0.49\gamma} \exp(-0.485\gamma^2) + 251(1+\delta)^2 \frac{\epsilon}{\gamma^2} \end{aligned}$$

with probability exceeding $1 - \frac{\log m}{\log(1+\delta)} (1 + \frac{2}{\epsilon})^n \exp(-c\epsilon^2 m)$. This establishes (5.30) for all $2 \leq \gamma \leq m$.

It remains to deal with the case where $\gamma > m$. To this end, we rely on the following observation:

$$\frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq m\}} \leq \frac{1}{m} \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{h}|^2}{m^2} \stackrel{(i)}{\leq} \frac{1+\delta}{m^2} \|\mathbf{h}\|^2 \ll \frac{1}{m} \quad \forall \mathbf{h} \text{ with } \|\mathbf{h}\| = 1,$$

where (i) comes from [14, lemma 3.1]. This basically tells us that with high probability, none of the indicator variables can be equal to 1. Consequently,

$$\frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{h}| \geq m\}} = 0,$$

which proves the claim.

A.5 Proof of Lemma 5.9

Fix $\delta > 0$. Recalling the notation

$$v_i := 2 \left\{ 2\mathbf{a}_i^\top \mathbf{h} - \frac{|\mathbf{a}_i^\top \mathbf{h}|^2}{\mathbf{a}_i^\top \mathbf{z}} \right\} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i},$$

we see from the expansion (5.21) that

$$(A.14) \quad \left\| \frac{1}{m} \nabla_{\text{tr}} \ell(\mathbf{z}) \right\| = \left\| \frac{1}{m} \mathbf{A}^\top \mathbf{v} \right\| \leq \frac{1}{m} \|\mathbf{A}\| \cdot \|\mathbf{v}\| \leq (1+\delta) \frac{\|\mathbf{v}\|}{\sqrt{m}}$$

as soon as $m \geq c_1 n$ for some sufficiently large $c_1 > 0$. Here, the norm estimate $\|A\| \leq \sqrt{m}(1 + \delta)$ arises from standard random matrix results [51, cor. 5.35].

Everything then comes down to controlling $\|v\|$. To this end, making use of the inclusion (5.22) yields

$$\begin{aligned} \frac{1}{4m} \|v\|^2 &= \frac{1}{m} \sum_{i=1}^m \left(2a_i^\top h - \frac{|a_i^\top h|^2}{a_i^\top z} \right)^2 \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i} \\ &\leq \frac{1}{m} \sum_{i=1}^m \left(2|a_i^\top h| + \frac{|a_i^\top h|^2}{|a_i^\top z|} \right)^2 \mathbb{1}_{\mathcal{E}_1^i \cap (\mathcal{D}_{\gamma_4}^{i,1} \cup \mathcal{D}_{\gamma_4}^{i,2})} \\ &\leq \frac{1}{m} \sum_{i=1}^m \left\{ 4(a_i^\top h)^2 + \left(\frac{4|a_i^\top h|^3}{|a_i^\top z|} + \frac{|a_i^\top h|^4}{|a_i^\top z|^2} \right) \mathbb{1}_{\mathcal{E}_1^i \cap (\mathcal{D}_{\gamma_4}^{i,1} \cup \mathcal{D}_{\gamma_4}^{i,2})} \right\} \\ &= \frac{1}{m} \sum_{i=1}^m \left\{ 4(a_i^\top h)^2 + \left(4 + \frac{|a_i^\top h|}{|a_i^\top z|} \right) \frac{|a_i^\top h|^3}{|a_i^\top z|} (\mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_4}^{i,1}} + \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma_4}^{i,2}}) \right\}. \end{aligned}$$

The first term is controlled by [14, lemma 3.1] in such a way that with probability $1 - \exp(-\Omega(m))$,

$$\frac{1}{m} \sum_{i=1}^m 4(a_i^\top h)^2 \leq 4(1 + \delta) \|h\|^2.$$

Turning to the remaining terms, we see from the definitions of $\mathcal{D}_{\gamma}^{i,1}$ and $\mathcal{D}_{\gamma}^{i,2}$ that

$$\frac{|a_i^\top h|}{|a_i^\top z|} \leq \begin{cases} \frac{\gamma \|h\|}{\alpha_z^{\text{lb}} \|z\|} & \text{on } \mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}, \\ 2 + \frac{\gamma \|h\|}{\alpha_z^{\text{lb}} \|z\|} & \text{on } \mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,2}, \end{cases} \leq \begin{cases} 1 & \text{on } \mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}, \\ 3 & \text{on } \mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,2}, \end{cases}$$

as long as $\gamma \leq \frac{\alpha_z^{\text{lb}} \|z\|}{\|h\|}$. Consequently, one can bound

$$\begin{aligned} &\frac{1}{m} \sum_{i=1}^m \left(4 + \frac{|a_i^\top h|}{|a_i^\top z|} \right) \frac{|a_i^\top h|^3}{|a_i^\top z|} (\mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}} + \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,2}}) \\ &\leq \frac{5}{m} \sum_{i=1}^m \frac{|a_i^\top h|^3}{|a_i^\top z|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,1}} + \frac{7}{m} \sum_{i=1}^m \frac{|a_i^\top h|^3}{|a_i^\top z|} \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{D}_{\gamma}^{i,2}} \\ &\leq \frac{5(1 + \delta) \sqrt{8/\pi} \|h\|^3}{\alpha_z^{\text{lb}} \|z\|} + \frac{7}{100} (1 + \delta) \|h\|^2, \end{aligned}$$

where the last inequality follows from (5.28) and (5.37).

Recall that $\gamma_4 = 3\alpha_h$. Taken together all these bounds lead to the upper bound

$$\begin{aligned} \frac{1}{4m} \|v\|^2 &\leq (1 + \delta) \left\{ 4 + \frac{5\sqrt{8/\pi} \|h\|}{\alpha_z^{\text{lb}} \|z\|} + \frac{7}{100} \right\} \|h\|^2 \\ &\leq (1 + \delta) \left\{ 4 + \frac{5\sqrt{8/\pi}}{3\alpha_h} + \frac{7}{100} \right\} \|h\|^2 \end{aligned}$$

whenever

$$\frac{\|\mathbf{h}\|}{\|\mathbf{z}\|} \leq \min \left\{ \frac{\alpha_z^{\text{lb}}}{3\alpha_h}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98/3}(\alpha_z^{\text{lb}})^2}{2\alpha_z^{\text{ub}} + \alpha_z^{\text{lb}}}, \frac{1}{11} \right\}.$$

Substituting this into (A.14) completes the proof.

Appendix B Proofs for Section 7

B.1 Proof of Lemma 7.1

First, we collect a few results on the magnitudes of $\mathbf{a}_i^\top \mathbf{x}$ ($1 \leq i \leq m$) that will be useful in constructing the hypotheses. Observe that for any given \mathbf{x} and any sufficiently large m ,

$$\begin{aligned} \mathbb{P} \left\{ \min_{1 \leq i \leq m} |\mathbf{a}_i^\top \mathbf{x}| \geq \frac{1}{m \log m} \|\mathbf{x}\| \right\} &= \left(\mathbb{P} \left\{ |\mathbf{a}_i^\top \mathbf{x}| \geq \frac{1}{m \log m} \|\mathbf{x}\| \right\} \right)^m \\ &\geq \left(1 - \frac{2}{\sqrt{2\pi}} \frac{1}{m \log m} \right)^m \geq 1 - o(1). \end{aligned}$$

Besides, since

$$\mathbb{E}[\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \frac{\|\mathbf{x}\|}{5 \log m}\}}] \leq \frac{1}{\sqrt{2\pi}} \frac{2}{5 \log m} \leq \frac{1}{5 \log m},$$

applying Hoeffding's inequality yields

$$\begin{aligned} &\mathbb{P} \left\{ \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \frac{\|\mathbf{x}\|}{5 \log m}\}} > \frac{m}{4 \log m} \right\} \\ &= \mathbb{P} \left\{ \frac{1}{m} \sum_{i=1}^m (\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \frac{\|\mathbf{x}\|}{5 \log m}\}} - \mathbb{E}[\mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \frac{\|\mathbf{x}\|}{5 \log m}\}}]) > \frac{1}{20 \log m} \right\} \\ &\leq \exp \left(-\Omega \left(\frac{m}{\log^2 m} \right) \right). \end{aligned}$$

To summarize, with probability $1 - o(1)$, one has

$$(B.1) \quad \min_{1 \leq i \leq m} |\mathbf{a}_i^\top \mathbf{x}| \geq \frac{1}{m \log m} \|\mathbf{x}\|,$$

$$(B.2) \quad \sum_{i=1}^m \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \frac{\|\mathbf{x}\|}{\log m}\}} \leq \frac{m}{4 \log m} := k.$$

In what follows, we will first produce a set \mathcal{M}_1 of exponentially many vectors surrounding \mathbf{x} in such a way that every pair is separated by about the same distance, and then verify that a nontrivial fraction of \mathcal{M}_1 obeys (7.3). Without loss of generality, we assume that \mathbf{x} takes the form $\mathbf{x} = [b, 0, \dots, 0]^\top$ for some $b > 0$.

The construction of \mathcal{M}_1 follows a standard random packing argument. Let $\mathbf{w} = [w_1, \dots, w_n]^\top$ be a random vector with

$$w_i = x_i + \frac{1}{\sqrt{2n}} z_i, \quad 1 \leq i \leq n,$$

where $z_i \stackrel{\text{ind.}}{\sim} \mathcal{N}(0, 1)$. The collection \mathcal{M}_1 is then obtained by generating $M_1 = \exp(\frac{n}{20})$ independent copies $\mathbf{w}^{(l)}$ ($1 \leq l < M_1$) of \mathbf{w} . For any $\mathbf{w}^{(l)}, \mathbf{w}^{(j)} \in \mathcal{M}_1$, the concentration inequality [51, cor. 5.35] gives

$$\mathbb{P}\{0.5\sqrt{n} - 1 \leq \sqrt{n}\|\mathbf{w}^{(l)} - \mathbf{w}^{(j)}\| \leq 1.5\sqrt{n} + 1\} \geq 1 - 2\exp(-n/8),$$

$$\mathbb{P}\{0.5\sqrt{n} - 1 \leq \sqrt{2n}\|\mathbf{w}^{(l)} - \mathbf{x}\| \leq 1.5\sqrt{n} + 1\} \geq 1 - 2\exp(-n/8).$$

Taking the union bound over all $\binom{M_1}{2}$ pairs we obtain

$$(B.3) \quad 0.5 - n^{-1/2} \leq \|\mathbf{w}^{(l)} - \mathbf{w}^{(j)}\| \leq 1.5 + n^{-1/2} \quad \forall l \neq j,$$

$$(B.4) \quad 1/\sqrt{8} - (2n)^{-1/2} \leq \|\mathbf{w}^{(l)} - \mathbf{x}\| \leq \sqrt{9/8} + (2n)^{-1/2}, \quad 1 \leq l \leq M_1,$$

with probability exceeding $1 - 2M_1^2 \exp(-\frac{n}{8}) \geq 1 - 2\exp(-\frac{n}{40})$.

The next step is to show that many vectors in \mathcal{M}_1 satisfy (7.3). For any given \mathbf{w} with $\mathbf{r} := \mathbf{w} - \mathbf{x}$, by letting $\mathbf{a}_{i,\perp} := [a_{i,2}, \dots, a_{i,n}]^\top$, $r_\parallel := r_1$, and $\mathbf{r}_\perp := [r_2, \dots, r_n]^\top$, we derive

$$(B.5) \quad \frac{|\mathbf{a}_i^\top \mathbf{r}|^2}{|\mathbf{a}_i^\top \mathbf{x}|^2} \leq \frac{2|a_{i,1}r_\parallel|^2 + 2|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp|^2}{|a_{i,1}|^2 \|\mathbf{x}\|^2} \leq \frac{2|r_\parallel|^2}{\|\mathbf{x}\|^2} + \frac{2|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp|^2}{|a_{i,1}|^2 \|\mathbf{x}\|^2} \\ \leq \frac{2\|\mathbf{r}\|^2}{\|\mathbf{x}\|^2} + \frac{2|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp|^2}{|a_{i,1}|^2 \|\mathbf{x}\|^2}.$$

It then boils down to developing an upper bound on $|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp|^2 / |a_{i,1}|^2$. This ratio is convenient to work with since the numerator and denominator are stochastically independent. To simplify presentation, we reorder $\{\mathbf{a}_i\}$ in a way that

$$(m \log m)^{-1} \|\mathbf{x}\| \leq |\mathbf{a}_1^\top \mathbf{x}| \leq |\mathbf{a}_2^\top \mathbf{x}| \leq \dots \leq |\mathbf{a}_m^\top \mathbf{x}|;$$

this will not affect our subsequent analysis concerning $\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp$ since it is independent of $\mathbf{a}_i^\top \mathbf{x}$.

To proceed, we let $\mathbf{r}_\perp^{(l)}$ consist of all but the first entry of $\mathbf{w}^{(l)} - \mathbf{x}$, and introduce the indicator variables

$$(B.6) \quad \xi_i^l := \begin{cases} \mathbb{1}_{\{|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp^{(l)}| \leq \frac{1}{m} \sqrt{\frac{n-1}{2n}}\}}, & 1 \leq i \leq k, \\ \mathbb{1}_{\{|\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp^{(l)}| \leq \sqrt{\frac{2(n-1) \log n}{n}}\}}, & i > k, \end{cases}$$

where $k = \frac{m}{4 \log m}$ as before. In words, we divide $\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp^{(l)}$, $1 \leq i \leq m$, into two groups, with the first group enforcing far more stringent control than the second group. These indicator variables are useful since any $\mathbf{w}^{(l)}$ obeying $\prod_{i=1}^m \xi_i^l = 1$

will satisfy (7.3) when n is sufficiently large. To see this, note that for the first group of indices, $\xi_i^l = 1$ requires

$$(B.7) \quad |a_{i,\perp}^\top r^{(l)}| \leq \frac{1}{m} \sqrt{\frac{n-1}{2n}} \leq \frac{2}{m} \frac{\sqrt{n-1}}{\sqrt{n}-2} \|r^{(l)}\| \leq \frac{3}{m} \|r^{(l)}\|, \quad 1 \leq i \leq k,$$

where the second inequality follows from (B.4). This taken collectively with (B.1) and (B.5) yields

$$\frac{|a_i^\top r^{(l)}|^2}{|a_i^\top x|^2} \leq \frac{2\|r^{(l)}\|^2}{\|x\|^2} + \frac{\frac{9}{m^2}\|r^{(l)}\|^2}{\frac{1}{m^2 \log^2 m}\|x\|^2} \leq \frac{(2 + 9 \log^2 m)\|r^{(l)}\|^2}{\|x\|^2}, \quad 1 \leq i \leq k.$$

Regarding the second group of indices, $\xi_i^l = 1$ gives

$$(B.8) \quad |a_{i,\perp}^\top r^{(l)}| \leq \sqrt{\frac{2(n-1) \log n}{n}} \leq \sqrt{17 \log n} \|r^{(l)}\|, \quad i = k+1, \dots, m,$$

where the last inequality again follows from (B.4). Plugging (B.8) and (B.2) into (B.5) gives

$$\frac{|a_i^\top r^{(l)}|^2}{|a_i^\top x|^2} \leq \frac{2\|r^{(l)}\|^2}{\|x\|^2} + \frac{17\|r^{(l)}\|^2 \log n}{\|x\|^2 / \log^2 m} \leq \frac{(2 + 17 \log^3 m)\|r^{(l)}\|^2}{\|x\|^2}, \quad i \geq k+1.$$

Consequently, (7.3) is satisfied for all $1 \leq i \leq m$. It then suffices to guarantee the existence of exponentially many vectors obeying $\prod_{i=1}^m \xi_i^l = 1$.

Note that the first group of indicator variables is quite stringent, namely, for each i only a fraction $O(1/m)$ of the equations could satisfy $\xi_i^l = 1$. Fortunately, M_1 is exponentially large, and hence even M_1/m^k is exponentially large. Put formally, we claim that the first group satisfies

$$(B.9) \quad \sum_{l=1}^{M_1} \prod_{i=1}^k \xi_i^l \geq \frac{1}{2} \frac{M_1}{(2\pi)^{k/2} (1 + 4\sqrt{k/n})^{k/2}} \left(\frac{1}{\sqrt{2\pi m}} \right)^k := \tilde{M}_1$$

with probability exceeding $1 - \exp(-\Omega(k)) - \exp(-\tilde{M}_1/4)$. With this claim in place (which will be proved later), one has

$$\begin{aligned} \sum_{l=1}^{M_1} \prod_{i=1}^k \xi_i^l &\geq \frac{1}{2} M_1 \frac{1}{(e^2 m)^k} = \frac{1}{2} \exp\left(\left(\frac{1}{20} - \frac{k(2 + \log m)}{n}\right)n\right) \\ &\geq \frac{1}{2} \exp\left(\frac{1}{25}n\right) \end{aligned}$$

when n and m/n are sufficiently large. In light of this, we will let \mathcal{M}_2 be a collection comprising all $w^{(l)}$ obeying $\prod_{i=1}^k \xi_i^l = 1$, which has size $M_2 \geq \frac{1}{2} \exp(\frac{1}{25}n)$ based on the preceding argument. For notational simplicity, it will be assumed that the vectors in \mathcal{M}_2 are exactly $w^{(j)}$ ($1 \leq j \leq M_2$).

We now move on to the second group by examining how many vectors $w^{(j)}$ in \mathcal{M}_2 further satisfy $\prod_{i=k+1}^m \xi_i^j = 1$. Notably, the above construction of \mathcal{M}_2 relies

only on $\{\mathbf{a}_i\}_{1 \leq i \leq k}$ and is independent of the remaining vectors $\{\mathbf{a}_i\}_{i > k}$. In what follows, the argument proceeds conditionally on \mathcal{M}_2 and $\{\mathbf{a}_i\}_{1 \leq i \leq k}$. Applying the union bound gives

$$\begin{aligned} & \mathbb{E} \left[\sum_{j=1}^{M_2} \left(1 - \prod_{i=k+1}^m \xi_i^j \right) \right] \\ &= \sum_{j=1}^{M_2} \mathbb{P} \left\{ \exists i (k < i \leq m) : |\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp^{(l)}| > \sqrt{\frac{2(n-1) \log n}{n}} \right\} \\ &\leq \sum_{j=1}^{M_2} \sum_{i=k+1}^m \mathbb{P} \left\{ |\mathbf{a}_{i,\perp}^\top \mathbf{r}_\perp^{(l)}| > \sqrt{\frac{2(n-1) \log n}{n}} \right\} \leq M_2 m \frac{1}{n^2}. \end{aligned}$$

This combined with Markov's inequality gives

$$\sum_{j=1}^{M_2} \left(1 - \prod_{i=k+1}^m \xi_i^j \right) \leq \frac{m \log m}{n^2} \cdot M_2$$

with probability $1 - o(1)$.

Putting the above inequalities together suggests that with probability $1 - o(1)$, there exist at least

$$\left(1 - \frac{m \log m}{n^2} \right) M_2 \geq \frac{1}{2} \left(1 - \frac{m \log m}{n^2} \right) \exp\left(\frac{1}{25}n\right) \geq \exp\left(\frac{n}{30}\right)$$

vectors in \mathcal{M}_2 satisfying $\prod_{l=k+1}^m \xi_i^l = 1$. We then choose \mathcal{M} to be the set consisting of all these vectors, which forms a valid collection satisfying the properties of Lemma 7.1.

Finally, the only remaining step is to establish the claim (B.9). To start with, consider an $n \times k$ matrix $\mathbf{B} := [\mathbf{b}_1, \dots, \mathbf{b}_k]$ of i.i.d. standard normal entries, and let $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \frac{1}{n} \mathbf{I}_n)$. Conditional on the $\{\mathbf{b}_i\}$,

$$\mathbf{b}_\mathbf{u} = \begin{bmatrix} b_{1,\mathbf{u}} \\ \vdots \\ b_{k,\mathbf{u}} \end{bmatrix} := \begin{bmatrix} \mathbf{b}_1^\top \mathbf{u} \\ \vdots \\ \mathbf{b}_k^\top \mathbf{u} \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \frac{1}{n} \mathbf{B}^\top \mathbf{B}\right).$$

For sufficiently large m , one has $k = \frac{m}{4 \log m} \leq \frac{1}{4}n$. Using [51, cor. 5.35] we get

$$(B.10) \quad \left\| \frac{1}{n} \mathbf{B}^\top \mathbf{B} - \mathbf{I} \right\| \leq 4\sqrt{k/n}$$

with probability $1 - \exp(-\Omega(k))$. Thus, for any constant $0 < \epsilon < \frac{1}{2}$, conditional on $\{\mathbf{b}_i\}$ and (B.10), we obtain

$$\begin{aligned} & \mathbb{P} \left\{ \bigcap_{i=1}^k \left\{ |\mathbf{b}_i^\top \mathbf{u}| \leq \frac{1}{m} \right\} \right\} \\ & \geq (2\pi)^{-\frac{k}{2}} \det^{-\frac{1}{2}} \left(\frac{1}{n} \mathbf{B}^\top \mathbf{B} \right) \Delta_{\mathbf{b}_u \in \Upsilon} \exp \left(-\frac{1}{2} \mathbf{b}_u^\top \left(\frac{1}{n} \mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{b}_u \right) d\mathbf{b}_u \\ (B.11) \quad & \geq (2\pi)^{-\frac{k}{2}} (1 + 4\sqrt{k/n})^{-\frac{k}{2}} \Delta_{\mathbf{b}_u \in \Upsilon} \exp \left(-\frac{1}{2} (1 + 4\sqrt{k/n})^{-1} \sum_{i=1}^k b_{i,u}^2 \right) d\mathbf{b}_u \end{aligned}$$

$$(B.12) \quad \geq (2\pi)^{-\frac{k}{2}} (1 + 4\sqrt{k/n})^{-\frac{k}{2}} (\sqrt{2\pi m})^{-k},$$

where $\Upsilon := \{\tilde{\mathbf{b}} \mid |\tilde{b}_i| \leq m^{-1}, 1 \leq i \leq k\}$ and (B.11) is a direct consequence from (B.10).

When it comes to our quantity of interest, the above lower bound (B.12) indicates that on an event (defined via $\{\mathbf{a}_i\}$) of probability approaching 1, we have

$$(B.13) \quad \mathbb{E} \left[\sum_{l=1}^{M_1} \prod_{i=1}^k \xi_i^l \right] \geq M_1 (2\pi)^{-\frac{k}{2}} (1 + 4\sqrt{k/n})^{-\frac{k}{2}} (\sqrt{2\pi m})^{-k}.$$

Since conditional on $\{\mathbf{a}_i\}$, $\prod_{i=1}^k \xi_i^l$ are independent across l , applying the Chernoff-type bound [35, theorem 4.5] gives

$$\sum_{l=1}^{M_1} \prod_{i=1}^k \xi_i^l \geq \frac{M_1}{2} (2\pi)^{-\frac{k}{2}} (1 + 4\sqrt{k/n})^{-\frac{k}{2}} (\sqrt{2\pi m})^{-k}$$

with probability exceeding

$$1 - \exp \left(-\frac{1}{8} \frac{M_1}{(2\pi)^{k/2} (1 + 4\sqrt{k/n})^{k/2}} \left(\frac{1}{\sqrt{2\pi m}} \right)^k \right).$$

This concludes the proof.

B.2 Proof of Lemma 7.2

Before proceeding, we introduce the χ^2 -divergence between two probability measures P and Q as

$$(B.14) \quad \chi^2(P \| Q) := \Delta \left(\frac{dP}{dQ} \right)^2 dQ - 1.$$

It is well known (e.g., [50, lemma 2.7]) that

$$(B.15) \quad \text{KL}(P \| Q) \leq \log(1 + \chi^2(P \| Q)),$$

and hence it suffices to develop an upper bound on the χ^2 divergence.

Under independence, for any $\mathbf{w}_0, \mathbf{w}_1 \in \mathbb{R}^n$, the decoupling identity of the χ^2 divergence [50, p. 96] gives

$$\begin{aligned} \chi^2(\mathbb{P}(\mathbf{y} \mid \mathbf{w}_1) \parallel \mathbb{P}(\mathbf{y} \mid \mathbf{w}_0)) &= \prod_{i=1}^m (1 + \chi^2(\mathbb{P}(y_i \mid \mathbf{w}_1) \parallel \mathbb{P}(y_i \mid \mathbf{w}_0))) - 1 \\ (B.16) \quad &= \exp\left(\sum_{i=1}^m \frac{(|\mathbf{a}_i^\top \mathbf{w}_1|^2 - |\mathbf{a}_i^\top \mathbf{w}_0|^2)^2}{|\mathbf{a}_i^\top \mathbf{w}_0|^2}\right) - 1. \end{aligned}$$

The preceding identity (B.16) arises from the following computation: by definition of $\chi^2(\cdot \parallel \cdot)$,

$$\begin{aligned} &\chi^2(\text{Poisson}(\lambda_1) \parallel \text{Poisson}(\lambda_0)) \\ &= \left\{ \sum_{k=0}^{\infty} \frac{(\lambda_1^k \exp(-\lambda_1))^2}{\lambda_0^k \exp(-\lambda_0) k!} \right\} - 1 \\ &= \exp\left(\lambda_0 - 2\lambda_1 + \frac{\lambda_1^2}{\lambda_0}\right) \left\{ \sum_{k=0}^{\infty} \frac{(\lambda_1^2/\lambda_0)^k}{k!} \exp\left(-\frac{\lambda_1^2}{\lambda_0}\right) \right\} - 1 \\ &= \exp\left(\frac{(\lambda_1 - \lambda_0)^2}{\lambda_0}\right) - 1. \end{aligned}$$

Set $\mathbf{r} := \mathbf{w}_1 - \mathbf{w}_0$. To summarize,

$$\begin{aligned} \text{KL}(\mathbb{P}(\mathbf{y} \mid \mathbf{w}_1) \parallel \mathbb{P}(\mathbf{y} \mid \mathbf{w}_0)) &\leq \sum_{i=1}^m \frac{(|\mathbf{a}_i^\top \mathbf{w}_1|^2 - |\mathbf{a}_i^\top \mathbf{w}_0|^2)^2}{|\mathbf{a}_i^\top \mathbf{w}_0|^2} \\ &\leq \sum_{i=1}^m \frac{|\mathbf{a}_i^\top \mathbf{r}|^2 (2|\mathbf{a}_i^\top \mathbf{w}_0| + |\mathbf{a}_i^\top \mathbf{r}|)^2}{|\mathbf{a}_i^\top \mathbf{w}_0|^2} \\ (B.17) \quad &= \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{r}|^2 \left(\frac{8|\mathbf{a}_i^\top \mathbf{w}_0|^2 + 2|\mathbf{a}_i^\top \mathbf{r}|^2}{|\mathbf{a}_i^\top \mathbf{w}_0|^2} \right). \end{aligned}$$

Appendix C Initialization via Truncated Spectral Method

This section demonstrates that the truncated spectral method works when $m \asymp n$, as stated in the proposition below.

PROPOSITION C.1. Fix $\delta > 0$ and $\mathbf{x} \in \mathbb{R}^n$. Consider the model where $y_i = |\mathbf{a}_i^\top \mathbf{x}|^2 + \eta_i$ and $\mathbf{a}_i \stackrel{\text{ind.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I})$. Suppose that

$$(C.1) \quad |\eta_i| \leq \varepsilon \max\{\|\mathbf{x}\|^2, |\mathbf{a}_i^\top \mathbf{x}|^2\}, \quad 1 \leq i \leq m,$$

for some sufficiently small constant $\varepsilon > 0$. The solution $\mathbf{z}^{(0)}$ returned by the truncated spectral method obeys

$$(C.2) \quad \text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq \delta \|\mathbf{x}\|,$$

with probability exceeding $1 - \exp(-\Omega(m))$ provided that $m > c_0 n$ for some constant $c_0 > 0$.

PROOF. By homogeneity, it suffices to consider the case where $\|\mathbf{x}\| = 1$. Recall from [14, lemma 3.1] that $\frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{x}|^2 \in [1 \pm \varepsilon] \|\mathbf{x}\|^2$ holds with probability $1 - \exp(-\Omega(m))$. Under the hypothesis (C.1), one has

$$\frac{1}{m} \|\boldsymbol{\eta}\|_1 \leq \frac{1}{m} \sum_{i=1}^m \varepsilon (\|\mathbf{x}\|^2 + |\mathbf{a}_i^\top \mathbf{x}|^2) \leq \varepsilon \|\mathbf{x}\|^2 + \varepsilon(1 + \varepsilon) \|\mathbf{x}\|^2 \leq 3\varepsilon \|\mathbf{x}\|^2,$$

which yields

$$(C.3) \quad \lambda_0^2 := \frac{1}{m} \sum_{l=1}^m y_l = \frac{1}{m} \sum_{l=1}^m |\mathbf{a}_l^\top \mathbf{x}|^2 + \frac{1}{m} \sum_{l=1}^m \eta_l \in [1 \pm 4\varepsilon] \|\mathbf{x}\|^2$$

with probability $1 - \exp(-\Omega(m))$.

Consequently, when $|\eta_i| \leq \varepsilon \|\mathbf{x}\|^2$, one has

$$(C.4) \quad \begin{aligned} \mathbb{1}_{\{(\mathbf{a}_i^\top \mathbf{x})^2 + \eta_i \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} &\leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l) + |\eta_i|\}} \\ &\leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq (1+4\varepsilon)\alpha_y^2 + \varepsilon\}}, \end{aligned}$$

$$(C.5) \quad \begin{aligned} \mathbb{1}_{\{(\mathbf{a}_i^\top \mathbf{x})^2 + \eta_i \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} &\geq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l) - |\eta_i|\}} \\ &\geq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq (1-4\varepsilon)\alpha_y^2 - \varepsilon\}}. \end{aligned}$$

In addition, in the case where $|\eta_i| \leq \varepsilon |\mathbf{a}_i^\top \mathbf{x}|^2$,

$$(C.6) \quad \begin{aligned} \mathbb{1}_{\{(\mathbf{a}_i^\top \mathbf{x})^2 + \eta_i \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} &\leq \mathbb{1}_{\{(1-\varepsilon)|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} \\ &\leq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \frac{1+4\varepsilon}{1-\varepsilon} \alpha_y^2\}}, \end{aligned}$$

$$(C.7) \quad \begin{aligned} \mathbb{1}_{\{(\mathbf{a}_i^\top \mathbf{x})^2 + \eta_i \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} &\geq \mathbb{1}_{\{(1+\varepsilon)|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \alpha_y^2 (\frac{1}{m} \sum_l y_l)\}} \\ &\geq \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}|^2 \leq \frac{1-4\varepsilon}{1+\varepsilon} \alpha_y^2\}}. \end{aligned}$$

Taken collectively, these inequalities imply that

$$(C.8) \quad \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbf{a}_i \mathbf{a}_i^\top (\mathbf{a}_i^\top \mathbf{x})^2 \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \gamma_2\}}}_{:= Y_2} \leq Y \leq \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbf{a}_i \mathbf{a}_i^\top (\mathbf{a}_i^\top \mathbf{x})^2 \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq \gamma_1\}}}_{:= Y_1},$$

where

$$\begin{aligned} \gamma_1 &:= \max \left\{ \sqrt{(1+4\varepsilon)\alpha_y^2 + \varepsilon}, \sqrt{\frac{1+4\varepsilon}{1-\varepsilon}} \alpha_y \right\}, \\ \gamma_2 &:= \min \left\{ \sqrt{(1-4\varepsilon)\alpha_y^2 - \varepsilon}, \sqrt{\frac{1-4\varepsilon}{1+\varepsilon}} \alpha_y \right\}. \end{aligned}$$

Letting $\xi \sim \mathcal{N}(0, 1)$, one can compute

$$(C.9) \quad \mathbb{E}[Y_1] = \beta_1 \mathbf{x} \mathbf{x}^\top + \beta_2 \mathbf{I} \quad \text{and} \quad \mathbb{E}[Y_2] = \beta_3 \mathbf{x} \mathbf{x}^\top + \beta_4 \mathbf{I},$$

where

$$\begin{aligned} \beta_1 &:= \mathbb{E}[\xi^4 \mathbb{1}_{\{|\xi| \leq \gamma_1\}}] - \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| \leq \gamma_1\}}], & \beta_2 &:= \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| \leq \gamma_1\}}], \\ \beta_3 &:= \mathbb{E}[\xi^4 \mathbb{1}_{\{|\xi| \leq \gamma_2\}}] - \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| \leq \gamma_2\}}], & \beta_4 &:= \mathbb{E}[\xi^2 \mathbb{1}_{\{|\xi| \leq \gamma_2\}}]. \end{aligned}$$

Recognizing that $\mathbf{a}_i \mathbf{a}_i^\top (\mathbf{a}_i^\top \mathbf{x})^2 \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq c\}}$ can be rewritten as $\mathbf{b}_i \mathbf{b}_i^\top$ for some sub-Gaussian vector $\mathbf{b}_i := \mathbf{a}_i (\mathbf{a}_i^\top \mathbf{x}) \mathbb{1}_{\{|\mathbf{a}_i^\top \mathbf{x}| \leq c\}}$, we apply standard results on random matrices with nonisotropic sub-Gaussian rows [51, eq. (5.26)] to deduce

$$(C.10) \quad \|\mathbf{Y}_1 - \mathbb{E}[\mathbf{Y}_1]\| \leq \delta, \quad \|\mathbf{Y}_2 - \mathbb{E}[\mathbf{Y}_2]\| \leq \delta,$$

with probability $1 - \exp(-\Omega(m))$, provided that m/n exceeds some large constant. In addition, when ε is sufficiently small, one also has $\|\mathbb{E}[\mathbf{Y}_1] - \mathbb{E}[\mathbf{Y}_2]\| \leq \delta$. These taken collectively with (C.8) give

$$(C.11) \quad \|\mathbf{Y} - \beta_1 \mathbf{x} \mathbf{x}^\top - \beta_2 \mathbf{I}\| \leq 3\delta.$$

Fix some small $\tilde{\delta} > 0$. With (C.11) in place, applying the Davis-Kahan sin Θ theorem [19] and taking δ, ε to be sufficiently small, we obtain

$$\text{dist}(\tilde{\mathbf{z}}, \mathbf{x}) \leq \tilde{\delta},$$

where $\tilde{\mathbf{z}}$ is the leading eigenvector of \mathbf{Y} . Since $\mathbf{z}^{(0)} := \lambda_0 \tilde{\mathbf{z}}$, one derives

$$(C.12) \quad \begin{aligned} \text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) &\leq \text{dist}(\lambda_0 \tilde{\mathbf{z}}, \tilde{\mathbf{z}}) + \text{dist}(\tilde{\mathbf{z}}, \mathbf{x}) \\ &\leq |\lambda_0 - 1| + \tilde{\delta} \leq \max\{\sqrt{1 + 2\varepsilon} - 1, 1 - \sqrt{1 - 2\varepsilon}\} + \tilde{\delta} \end{aligned}$$

as long as m/n is sufficiently large, where the last inequality follows from (C.3). Picking $\tilde{\delta}$ and ε to be sufficiently small, we establish the claim. \square

We now justify that the Poisson model (1.4) satisfies the condition (C.1). Suppose that $\mu_i = (\mathbf{a}_i^\top \mathbf{x})^2$ and hence $y_i \sim \text{Poisson}(\mu_i)$. It follows from the Chernoff bound that for all $t \geq 0$,

$$\begin{aligned} \mathbb{P}(y_i - \mu_i \geq \tau) &\leq \frac{\mathbb{E}[e^{t y_i}]}{\exp(t(\mu_i + \tau))} = \frac{\exp(\mu_i(e^t - 1))}{\exp(t(\mu_i + \tau))} \\ &= \exp(\mu_i(e^t - t - 1) - t\tau). \end{aligned}$$

Taking $\tau = 2\tilde{\varepsilon}\mu_i$ and $t = \tilde{\varepsilon}$ for any $0 \leq \tilde{\varepsilon} \leq 1$ gives

$$\begin{aligned} \mathbb{P}(y_i - \mu_i \geq 2\tilde{\varepsilon}\mu_i) &\leq \exp(\mu_i(e^{\tilde{\varepsilon}} - \tilde{\varepsilon} - 1 - 2\tilde{\varepsilon}\tilde{\varepsilon})) \\ &\stackrel{(i)}{\leq} \exp(\mu_i(t^2 - 2\tilde{\varepsilon}t)) = \exp(-\mu_i\tilde{\varepsilon}^2), \end{aligned}$$

where (i) follows since $e^t \leq 1 + t + t^2$ ($0 \leq t \leq 1$). In addition, for any $\tilde{\varepsilon} > 1$, taking $t = 1$ we get

$$\mathbb{P}(y_i - \mu_i \geq 2\tilde{\varepsilon}\mu_i) \leq \exp(\mu_i(e^1 - 1 - 2\tilde{\varepsilon})) \leq \exp(-0.5\mu_i\tilde{\varepsilon}).$$

Suppose that $\|\mathbf{x}\| \gtrsim \log m$. When $\mu_i \gtrsim \|\mathbf{x}\|^2$, setting $\tilde{\varepsilon} = 0.5\varepsilon < 1$ yields

$$\mathbb{P}(y_i - \mu_i \geq \varepsilon \mu_i) \leq \exp(-\mu_i \varepsilon^2 / 4) \leq \exp(-\Omega(\varepsilon^2 \|\mathbf{x}\|^2)) = \exp(-\Omega(\varepsilon^2 \log^2 m)).$$

When $\mu_i \lesssim \|\mathbf{x}\|^2$, letting $\kappa_i = \mu_i / \|\mathbf{x}\|^2$ and setting $\tilde{\varepsilon} = \varepsilon / 2\kappa_i \gtrsim \varepsilon$, we obtain

$$\begin{aligned} \mathbb{P}(y_i - \mu_i \geq \varepsilon \|\mathbf{x}\|^2) &= \mathbb{P}(y_i - \mu_i \geq 2\tilde{\varepsilon} \mu_i) \\ &\leq \exp(-\min\{\tilde{\varepsilon}, 0.5\} \cdot \tilde{\varepsilon} \mu_i) \\ &= \exp(-0.5 \min\{\tilde{\varepsilon}, 0.5\} \varepsilon \|\mathbf{x}\|^2) = \exp(-\Omega(\varepsilon^2 \log^2 m)). \end{aligned}$$

In view of the union bound,

$$(C.13) \quad \mathbb{P}(\exists i : \eta_i \geq \varepsilon \max\{\|\mathbf{x}\|^2, |\mathbf{a}_i^\top \mathbf{x}|^2\}) \leq m \exp(-\Omega(\varepsilon^2 \log^2 m)) \rightarrow 0.$$

Applying the same argument on $-y_i$ we get $\eta_i \geq -\varepsilon \max\{\|\mathbf{x}\|^2, |\mathbf{a}_i^\top \mathbf{x}|^2\}$ for all i , which together with (C.13) establishes the condition (C.1) with high probability. In conclusion, the claim (C.2) applies to the Poisson model.

Appendix D Local Error Contraction with Backtracking Line Search

In this section, we verify the effectiveness of a backtracking line search strategy by showing local error contraction. To keep it concise, we only sketch the proof for the noiseless case, but the proof extends to the noisy case without much difficulty. Also, we do not strive to obtain an optimized constant. For concreteness, we prove the following proposition.

PROPOSITION D.1. *The claim in Proposition 5.1 continues to hold if $\alpha_h \geq 6$, $\alpha_z^{\text{ub}} \geq 5$, $\alpha_z^{\text{lb}} \leq 0.1$, $\alpha_p \geq 5$, and*

$$(D.1) \quad \|\mathbf{h}\| / \|\mathbf{z}\| \leq \epsilon_{\text{tr}}$$

for some constant $\epsilon_{\text{tr}} > 0$ independent of n and m .

Note that if $\alpha_h \geq 6$, $\alpha_z^{\text{ub}} \geq 5$, and $\alpha_z^{\text{lb}} \leq 0.1$, then the boundary step size μ_0 given in Proposition 5.1 satisfies

$$\frac{0.994 - \zeta_1 - \zeta_2 - \sqrt{2/(9\pi)} \alpha_h^{-1}}{2(1.02 + 0.665 \alpha_h^{-1})} \geq 0.384.$$

Thus, it suffices to show that the step size obtained by a backtracking line search lies within $(0, 0.384)$. For notational convenience, we will set

$$\mathbf{p} := m^{-1} \nabla \ell_{\text{tr}}(\mathbf{z}) \quad \text{and} \quad \mathcal{E}_3^i := \{|\mathbf{a}_i^\top \mathbf{z}| \geq \alpha_z^{\text{lb}} \|\mathbf{z}\| \text{ and } |\mathbf{a}_i^\top \mathbf{p}| \leq \alpha_p \|\mathbf{p}\|\}$$

throughout the rest of the proof. We also impose the assumption

$$(D.2) \quad \|\mathbf{p}\| / \|\mathbf{z}\| \leq \epsilon$$

for some sufficiently small constant $\epsilon > 0$, so that $|\mathbf{a}_i^\top \mathbf{p}|/|\mathbf{a}_i^\top \mathbf{z}|$ is small for all nontruncated terms. It is self-evident from (5.38) that in the regime under study, one has

$$(D.3) \quad \|\mathbf{p}\| \geq 2\{1.99 - 2(\zeta_1 + \zeta_2) - \sqrt{8/\pi}(3\alpha_h)^{-1} - o(1)\}\|\mathbf{h}\| \geq 3.64\|\mathbf{h}\|.$$

To begin, consider three scalars h , b , and δ . Setting $b_\delta := [(b + \delta)^2 - b^2]/b^2$, we get

$$\begin{aligned} (D.4) \quad & (b + h)^2 \log \frac{(b + \delta)^2}{b^2} - (b + \delta)^2 + b^2 \\ &= (b + h)^2 \log(1 + b_\delta) - b^2 b_\delta \\ &\stackrel{(i)}{\leq} (b + h)^2 \{b_\delta - 0.4875b_\delta^2\} - b^2 b_\delta \\ &= ((b + h)^2 - b^2)b_\delta - 0.4875(b + h)^2 b_\delta^2 \\ &= h\delta(2 + h/b)(2 + \delta/b) - 0.4875(1 + h/b)^2 |\delta(2 + \delta/b)|^2 \\ &= 4h\delta + \frac{2h^2\delta}{b} + \frac{2h\delta^2}{b} + \frac{h^2\delta^2}{b^2} - 0.4875\delta^2 \left(1 + \frac{h}{b}\right)^2 \left(2 + \frac{\delta}{b}\right)^2, \end{aligned}$$

where (i) follows from the inequality $\log(1 + x) \leq x - 0.4875x^2$ for sufficiently small x .

To further simplify the bound, observe that

$$\begin{aligned} \delta^2 \left(1 + \frac{h}{b}\right)^2 \left(2 + \frac{\delta}{b}\right)^2 &\geq 4\delta^2 \left(1 + \frac{h}{b}\right)^2 + \delta^2 \left(1 + \frac{h}{b}\right)^2 \frac{4\delta}{b} \\ \text{and } \frac{2h\delta^2}{b} + \frac{h^2\delta^2}{b^2} &= \left(\left(1 + \frac{h}{b}\right)^2 - 1\right)\delta^2. \end{aligned}$$

Plugging these two identities into (D.4) yields

$$\begin{aligned} (D.4) &\leq 4h\delta + \frac{2h^2\delta}{b} - \left(0.95\left(1 + \frac{h}{b}\right)^2 + 1\right)\delta^2 - 0.4875\delta^2 \left(1 + \frac{h}{b}\right)^2 \frac{4\delta}{b} \\ &\leq 4h\delta - 1.95\delta^2 + \frac{2h^2|\delta|}{|b|} + \frac{1.9|h|}{|b|}\delta^2 + \frac{1.95|\delta^3|}{|b|} \left(1 + \frac{h}{b}\right)^2. \end{aligned}$$

Replacing, respectively, b , δ , and h with $\mathbf{a}_i^\top \mathbf{z}$, $\tau \mathbf{a}_i^\top \mathbf{p}$, and $-\mathbf{a}_i^\top \mathbf{h}$, one sees that the log-likelihood $\ell_i(\mathbf{z}) = y_i \log(|\mathbf{a}_i^\top \mathbf{z}|^2) - |\mathbf{a}_i^\top \mathbf{z}|^2$ obeys

$$\begin{aligned} \ell_i(\mathbf{z} + \tau \mathbf{p}) - \ell_i(\mathbf{z}) &= y_i \log \frac{|\mathbf{a}_i^\top (\mathbf{z} + \tau \mathbf{p})|^2}{|\mathbf{a}_i^\top \mathbf{z}|^2} - |\mathbf{a}_i^\top (\mathbf{z} + \tau \mathbf{p})|^2 + |\mathbf{a}_i^\top \mathbf{z}|^2 \\ &\leq \underbrace{-4\tau(\mathbf{a}_i^\top \mathbf{h})(\mathbf{a}_i^\top \mathbf{p})}_{:=I_{1,i}} - \underbrace{1.95\tau^2(\mathbf{a}_i^\top \mathbf{p})^2}_{:=I_{2,i}} + \underbrace{\frac{2\tau(\mathbf{a}_i^\top \mathbf{h})^2|\mathbf{a}_i^\top \mathbf{p}|}{|\mathbf{a}_i^\top \mathbf{z}|}}_{:=I_{3,i}} \\ &\quad + \underbrace{\frac{1.9\tau^2|\mathbf{a}_i^\top \mathbf{h}|}{|\mathbf{a}_i^\top \mathbf{z}|}(\mathbf{a}_i^\top \mathbf{p})^2}_{:=I_{4,i}} + \underbrace{\frac{1.95\tau^3|\mathbf{a}_i^\top \mathbf{p}|^3}{|\mathbf{a}_i^\top \mathbf{z}|}\left(1 - \frac{\mathbf{a}_i^\top \mathbf{h}}{\mathbf{a}_i^\top \mathbf{z}}\right)^2}_{:=I_{5,i}}. \end{aligned}$$

The next step is then to bound each of these terms separately. Most of the following bounds are straightforward consequences from [14, lemma 3.1] combined with the truncation rule. For the first term, applying the AM-GM inequality we get

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m I_{1,i} \mathbb{1}_{\mathcal{E}_3^i} &\leq \frac{4\tau}{3.64m} \sum_{i=1}^m \left\{ \frac{3.64^2}{2} (\mathbf{a}_i^\top \mathbf{h})^2 + \frac{1}{2} (\mathbf{a}_i^\top \mathbf{p})^2 \right\} \\ &\leq \frac{4\tau(1+\delta)}{3.64} \left\{ \frac{3.64^2}{2} \|\mathbf{h}\|^2 + \frac{1}{2} \|\mathbf{p}\|^2 \right\}. \end{aligned}$$

Secondly, it follows from Lemma 5.6 that

$$\frac{1}{m} \sum_{i=1}^m I_{2,i} \mathbb{1}_{\mathcal{E}_3^i} = -1.95\tau^2 \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{p})^2 \mathbb{1}_{\mathcal{E}_3^i} \leq -1.95(1 - \tilde{\zeta}_1 - \tilde{\zeta}_2)\tau^2 \|\mathbf{p}\|^2,$$

where

$$\begin{aligned} \tilde{\zeta}_1 &:= \max \left\{ \mathbb{E} \left[\xi^2 \mathbb{1}_{\{|\xi| \leq \sqrt{1.01}\alpha_z^{\text{lb}}\}} \right], \mathbb{E} \left[\mathbb{1}_{\{|\xi| \leq \sqrt{1.01}\alpha_z^{\text{lb}}\}} \right] \right\}, \\ \tilde{\zeta}_2 &:= \mathbb{E} \left[\xi^2 \mathbb{1}_{\{|\xi| > \sqrt{0.99}\alpha_h\}} \right]. \end{aligned}$$

The third term is controlled by

$$\frac{1}{m} \sum_{i=1}^m I_{3,i} \mathbb{1}_{\mathcal{E}_3^i} \leq 2\tau \frac{\alpha_p \|\mathbf{p}\|}{\alpha_z^{\text{lb}} \|\mathbf{z}\|} \left\{ \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{h})^2 \right\} \lesssim \tau \epsilon \|\mathbf{h}\|^2.$$

Fourthly, from the AM-GM inequality we get

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m I_{4,i} \mathbb{1}_{\mathcal{E}_3^i} &\leq \frac{1.9\tau^2 \alpha_p \|\mathbf{p}\|}{\alpha_z^{\text{lb}} \|\mathbf{z}\|} \frac{1}{m} \sum_{i=1}^m |\mathbf{a}_i^\top \mathbf{h}| |\mathbf{a}_i^\top \mathbf{p}| \\ &\lesssim \epsilon \tau^2 \frac{1}{m} \sum_{i=1}^m \left\{ 2|\mathbf{a}_i^\top \mathbf{h}|^2 + \frac{1}{8} |\mathbf{a}_i^\top \mathbf{p}|^2 \right\} \lesssim \epsilon \tau^2 \|\mathbf{p}\|^2. \end{aligned}$$

Finally, the last term is bounded by

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m I_{5,i} \mathbb{1}_{\mathcal{E}_3^i} &\leq \frac{1}{m} \sum_{i=1}^m \frac{1.95\tau^3 |\mathbf{a}_i^\top \mathbf{p}|^3}{|\mathbf{a}_i^\top \mathbf{z}|} \left(\frac{\mathbf{a}_i^\top \mathbf{x}}{\mathbf{a}_i^\top \mathbf{z}} \right)^2 \\ &\leq \frac{1.95\tau^3 \alpha_p^3 \|\mathbf{p}\|^3}{(\alpha_z^{\text{lb}})^3 \|\mathbf{z}\|^3} \frac{1}{m} \sum_{i=1}^m (\mathbf{a}_i^\top \mathbf{x})^2 \lesssim \tau^3 \epsilon \frac{\|\mathbf{x}\|^2}{\|\mathbf{z}\|^2} \|\mathbf{p}\|^2. \end{aligned}$$

Under the hypothesis (D.3), we can further derive

$$\frac{1}{m} \sum_{i=1}^m I_{1,i} \mathbb{1}_{\mathcal{E}_3^i} \leq \tau(1.1 + \delta) \|\mathbf{p}\|^2.$$

Putting all the above bounds together yields that the truncated objective function is majorized by

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \{\ell_i(\mathbf{z} + \tau \mathbf{p}) - \ell_i(\mathbf{z})\} \mathbb{1}_{\mathcal{E}_3^i} &\leq \frac{1}{m} \sum_{i=1}^m (I_{1,i} + I_{2,i} + I_{3,i} + I_{4,i} + I_{5,i}) \mathbb{1}_{\mathcal{E}_3^i} \\ &\leq \tau(1.1 + \delta) \|\mathbf{p}\|^2 - 1.95(1 - \tilde{\xi}_1 - \tilde{\xi}_2) \tau^2 \|\mathbf{p}\|^2 + \tau \tilde{\epsilon} \|\mathbf{p}\|^2 \\ &= \{\tau(1.1 + \delta) - 1.95(1 - \tilde{\xi}_1 - \tilde{\xi}_2) \tau^2 + \tau \tilde{\epsilon}\} \|\mathbf{p}\|^2 \end{aligned} \tag{D.5}$$

for some constant $\tilde{\epsilon} > 0$ that is linear in ϵ .

Note that the backtracking line search seeks a point satisfying

$$\frac{1}{m} \sum_{i=1}^m \{\ell_i(\mathbf{z} + \tau \mathbf{p}) - \ell_i(\mathbf{z})\} \mathbb{1}_{\mathcal{E}_3^i} \geq \frac{1}{2} \tau \|\mathbf{p}\|^2.$$

Given the above majorization (D.5), this search criterion is satisfied only if

$$\tau/2 \leq \tau(1.1 + \delta) - 1.95(1 - \tilde{\xi}_1 - \tilde{\xi}_2) \tau^2 + \tau \tilde{\epsilon}$$

or, equivalently,

$$\tau \leq \frac{0.6 + \delta + \tilde{\epsilon}}{1.95(1 - \tilde{\xi}_1 - \tilde{\xi}_2)} := \tau_{\text{ub}}.$$

Taking δ and $\tilde{\epsilon}$ to be sufficiently small, we see that $\tau \leq \tau_{\text{ub}} \leq 0.384$, provided that $\alpha_z^{\text{lb}} \leq 0.1$, $\alpha_z^{\text{ub}} \geq 5$, $\alpha_h \geq 6$, and $\alpha_p \geq 5$.

Using very similar arguments, one can also show that

$$\frac{1}{m} \sum_{i=1}^m \{\ell_i(\mathbf{z} + \tau \mathbf{p}) - \ell_i(\mathbf{z})\} \mathbb{1}_{\mathcal{E}_3^i}$$

is minorized by a similar quadratic function, which combined with the stopping criterion $\frac{1}{m} \sum_{i=1}^m \{\ell_i(\mathbf{z} + \tau \mathbf{p}) - \ell_i(\mathbf{z})\} \mathbb{1}_{\mathcal{E}_3} \geq \frac{1}{2} \tau \|\mathbf{p}\|^2$ suggests that τ is bounded away from 0. We omit this part for conciseness.

Acknowledgment. E.C. is partially supported by the National Science Foundation under grant CCF-0963835 and by the Math + X Award from the Simons Foundation. Y.C. is supported by the same National Science Foundation grant. We thank Carlos Sing-Long and Weijie Su for their constructive comments about an early version of the manuscript. E.C. is grateful to Xiaodong Li and Mahdi Soltanolkotabi for many discussions about Wirtinger flows. We thank the anonymous reviewers for helpful comments.

Bibliography

- [1] Alexeev, B.; Bandeira, A. S.; Fickus, M.; Mixon, D. G. Phase retrieval with polarization. *SIAM Journal on Imaging Sciences* **7** (2014), no. 1, 35–66. doi:10.1137/12089939X
- [2] Arora, S.; Ge, R.; Ma, T.; Moitra, A. Simple, efficient, and neural algorithms for sparse coding. Preprint, 2015. arXiv:1503.00778 [cs.LG]
- [3] Bahmani, S.; Romberg, J. Efficient compressive phase retrieval with constrained sensing vectors. Preprint, 2015. arXiv:1507.08254 [cs.IT]
- [4] Balakrishnan, S.; Wainwright, M. J.; Yu, B. Statistical guarantees for the EM algorithm: From population to sample-based analysis. Preprint, 2014. arXiv:1408.2156 [math.ST]
- [5] Balan, R.; Bodmann, B. G.; Casazza, P. G.; Edidin, D. Painless reconstruction from magnitudes of frame coefficients. *J. Fourier Anal. Appl.* **15** (2009), no. 4, 488–501. doi:10.1007/s00041-009-9065-1
- [6] Bandeira, A. S.; Cahill, J.; Mixon, D. G.; Nelson, A. A. Saving phase: injectivity and stability for phase retrieval. *Appl. Comput. Harmon. Anal.* **37** (2014), no. 1, 106–125. doi:10.1016/j.acha.2013.10.002
- [7] Bandeira, A. S.; Charikar, M.; Singer, A.; Zhu, A. Multireference alignment using semidefinite programming. *ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science*, 459–470. ACM, New York, 2014.
- [8] Ben-Tal, A.; Nemirovski, A. *Lectures on modern convex optimization*. Analysis, algorithms, and engineering applications. MPS/SIAM Series on Optimization. Society for Industrial and Applied Mathematics (SIAM), Philadelphia; Mathematical Programming Society (MPS), Philadelphia, 2001. doi:10.1137/1.9780898718829
- [9] Cai, T. T.; Zhang, A. ROP: matrix recovery via rank-one projections. *Ann. Statist.* **43** (2015), no. 1, 102–138. doi:10.1214/14-AOS1267
- [10] Candès, E. J.; Li, X. Solving quadratic equations via PhaseLift when there are about as many equations as unknowns. *Found. Comput. Math.* **14** (2014), no. 5, 1017–1026. doi:10.1007/s10208-013-9162-z
- [11] Candès, E. J.; Li, X.; Soltanolkotabi, M. Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.* **39** (2015), no. 2, 277–299. doi:10.1016/j.acha.2014.09.004
- [12] Candès, E. J.; Li, X.; Soltanolkotabi, M. Phase retrieval via Wirtinger flow, 2014 Available at: <http://www-bcf.usc.edu/~soltanol/PRWF.html>.
- [13] Candès, E. J.; Li, X.; Soltanolkotabi, M. Phase retrieval via Wirtinger flow: theory and algorithms. *IEEE Trans. Inform. Theory* **61** (2015), no. 4, 1985–2007. doi:10.1109/TIT.2015.2399924

- [14] Candès, E. J.; Strohmer, T.; Voroninski, V. PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.* **66** (2013), no. 8, 1241–1274. doi:10.1002/cpa.21432
- [15] Chen, Y.; Candès, E. J. Supplemental materials for: “Solving random quadratic systems of equations is nearly as easy as solving linear systems” (2015). Available at: <http://statweb.stanford.edu/~candes/publications.html>
- [16] Chen, Y.; Chi, Y.; Goldsmith, A. J. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Trans. Inform. Theory* **61** (2015), no. 7, 4034–4059. doi:10.1109/TIT.2015.2429594
- [17] Chen, Y.; Guibas, L. J.; Huang, Q. Near-optimal joint optimal matching via convex relaxation. *International Conference on Machine Learning (ICML)* (2014), 100–108.
- [18] Chen, Y.; Yi, X.; Caramanis, C. A convex formulation for mixed regression with two components: Minimax optimal rates. Preprint, 2013. arXiv:1312.7006 [stat.ML]
- [19] Davis, C.; Kahan, W. M. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* **7** (1970), 1–46. doi:10.1137/0707001
- [20] Demanet, L.; Hand, P. Stable optimizationless recovery from phaseless linear measurements. *J. Fourier Anal. Appl.* **20** (2014), no. 1, 199–221.
- [21] Eldar, Y. C.; Mendelson, S. Phase retrieval: stability and recovery guarantees. *Appl. Comput. Harmon. Anal.* **36** (2014), no. 3, 473–494. doi:10.1016/j.acha.2013.08.003
- [22] Elser, V. Phase retrieval by iterated projections. *J. Opt. Soc. Am. A* **20** (2003), no. 1, 40–55. doi:10.1364/JOSAA.20.000040
- [23] Fienup, J. R. Phase retrieval algorithms: a comparison. *Applied Optics* **21** (1982), 2758–2769. doi:10.1364/AO.21.002758
- [24] Gerchberg, R. W. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* **35** (1972), 237–246.
- [25] Gross, D.; Krahmer, F.; Kueng, R. Improved recovery guarantees for phase retrieval from coded diffraction patterns. Preprint, 2014. arXiv:1402.6286 [cs.IT]
- [26] Gross, D.; Krahmer, F.; Kueng, R. A partial derandomization of PhaseLift using spherical designs. *J. Fourier Anal. Appl.* **21** (2015), no. 2, 229–266. doi:10.1007/s00041-014-9361-2
- [27] Hardt, M.; Wootters, M. Fast matrix completion without the condition number. *Conference on Learning Theory* (2014), 638–678.
- [28] Huang, Q.; Guibas, L. Consistent shape maps via semidefinite programming. *Computer Graphics Forum* **32** (2013), no. 5, 177–186. doi:10.1111/cgf.12184
- [29] Jaganathan, K.; Oymak, S.; Hassibi, B. Recovery of sparse 1-D signals from the magnitudes of their Fourier transform. *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 1473–1477. IEEE, Cambridge, 2012.
- [30] Jain, P.; Netrapalli, P.; Sanghavi, S. Low-rank matrix completion using alternating minimization (extended abstract). *STOC’13—Proceedings of the 2013 ACM Symposium on Theory of Computing*, 665–674. ACM, New York, 2013. doi:10.1145/2488608.2488693
- [31] Keshavan, R. H.; Montanari, A.; Oh, S. Matrix completion from a few entries. *IEEE Trans. Inform. Theory* **56** (2010), no. 6, 2980–2998. doi:10.1109/TIT.2010.2046205
- [32] Kueng, R.; Rauhut, H.; Terstiege, U. Low rank matrix recovery from rank one measurements. Preprint, 2014. arXiv:1410.6913 [cs.IT]
- [33] Li, X.; Voroninski, V. Sparse signal recovery from quadratic measurements via convex programming. *SIAM J. Math. Anal.* **45** (2013), no. 5, 3019–3033. doi:10.1137/120893707
- [34] Marchesin, S.; Tu, Y.; Wu, H. Alternating projection, ptychographic imaging and phase synchronization. Preprint, 2014. arXiv:1402.0550 [math.OC]
- [35] Mitzenmacher, M.; Upfal, E. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, Cambridge, 2005. doi:10.1017/CBO9780511813603

- [36] Netrapalli, P.; Jain, P.; Sanghavi, S. Phase retrieval using alternating minimization. *IEEE Trans. Signal Process.* **63** (2015), no. 18, 4814–4826. doi:10.1109/TSP.2015.2448516
- [37] Netrapalli, P.; Niranjan, U.; Sanghavi, S.; Anandkumar, A.; Jain, P.: Non-convex robust PCA. *Advances in Neural Information Processing Systems* (2014), 1107–1115.
- [38] Nocedal, J.; Wright, S. J. *Numerical optimization*. Second edition. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2006.
- [39] Ohlsson, H.; Yang, A.; Dong, R.; Sastry, S. S. Compressive phase retrieval from squared output measurements via semidefinite programming. Preprint, 2011. arXiv:1111.6323 [math.ST]
- [40] Repetti, A.; Chouzenoux, E.; Pesquet, J.-C. A nonconvex regularized approach for phase retrieval. *International Conference on Image Processing* (2014), 1753–1757. doi:10.1109/ICIP.2014.7025351
- [41] Schniter, P.; Rangan, S. Compressive phase retrieval via generalized approximate message passing. *IEEE Trans. Signal Process.* **63** (2015), no. 4, 1043–1055. doi:10.1109/TSP.2014.2386294
- [42] Shechtman, Y.; Beck, A.; Eldar, Y. C. GESPAR: efficient phase retrieval of sparse signals. *IEEE Trans. Signal Process.* **62** (2014), no. 4, 928–938. doi:10.1109/TSP.2013.2297687
- [43] Shechtman, Y.; Eldar, Y. C.; Cohen, O.; Chapman, H.; Miao, J.; Segev, M. Phase retrieval with application to optical imaging. *IEEE Signal Processing Magazine* **32** (2015), no. 3, 87–109. doi:10.1109/MSP.2014.2352673
- [44] Shechtman, Y.; Eldar, Y. C.; Szameit, A.; Segev, M. Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing. *Optics Express* **19** (2011), no. 16, 14807–14822. doi:10.1364/OE.19.014807
- [45] Soltanolkotabi, M. Algorithms and theory for clustering and nonconvex quadratic programming. Ph.D. thesis, Stanford University, 2014.
- [46] Sun, R.; Luo, Z. Guaranteed matrix completion via non-convex factorization. Preprint, 2014. arXiv:1411.8003 [cs.LG]
- [47] Sun, J.; Qu, Q.; Wright, J. Complete dictionary recovery using nonconvex optimization. *International Conference on Machine Learning* (2015), 2351–2360.
- [48] Sun, J.; Qu, Q.; Wright, J. When are nonconvex problems not scary? Preprint, 2015. arXiv:1510.06096 [math.OC]
- [49] Trefethen, L. N.; Bau, D., III. *Numerical linear algebra*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1997. doi:10.1137/1.9780898719574
- [50] Tsybakov, A. B. *Introduction to nonparametric estimation*. Springer Series in Statistics. Springer, New York, 2009. doi:10.1007/b13794
- [51] Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *Compressed sensing*, 210–268. Cambridge University Press, Cambridge, 2012.
- [52] Waldspurger, I.; d’Aspremont, A.; Mallat, S. Phase recovery, MaxCut and complex semidefinite programming. *Math. Program.* **149** (2015), no. 1-2, Ser. A, 47–81. doi:10.1007/s10107-013-0738-9
- [53] Wei, K. Phase retrieval via Kaczmarz methods. Preprint, 2015. arXiv:1502.01822 [math.NA]
- [54] White, C. D.; Ward, R.; Sanghavi, S. The local convexity of solving quadratic equations. Preprint, 2015. arXiv:1506.07868 [math.NA]
- [55] Yi, X.; Caramanis, C.; Sanghavi, S. Alternating minimization for mixed linear regression. Preprint, 2013. arXiv:1310.3745 [stat.ML]
- [56] Zheng, Q.; Lafferty, J. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. *Advances in Neural Information Processing Systems* **28**, 109–117. Curran Associates, Inc., 2015. Available at: <http://papers.nips.cc/paper/5830-a-convergent-gradient-descent-algorithm-for-rank-minimization-and-semidefinite-programming-from-random-linear-measurements.pdf>

YUXIN CHEN
Packard Electrical Engineering
Building, Room 338
350 Serra Mall
Stanford University
Stanford, CA 94305
USA
E-mail: yxchen@stanford.edu

EMMANUEL CANDÈS
Department of Statistics
Stanford University
390 Serra Mall, Sequoia Hall
Stanford CA 94305
USA
E-mail: candes@stanford.edu

Received May 2015.