

Complete Dictionary Recovery over the Sphere II: Recovery by Riemannian Trust-Region Method

Ju Sun, *Student Member, IEEE*, Qing Qu, *Student Member, IEEE*, and John Wright, *Member, IEEE*

Abstract

We consider the problem of recovering a complete (i.e., square and invertible) matrix \mathbf{A}_0 , from $\mathbf{Y} \in \mathbb{R}^{n \times p}$ with $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$, provided \mathbf{X}_0 is sufficiently sparse. This recovery problem is central to theoretical understanding of dictionary learning, which seeks a sparse representation for a collection of input signals and finds numerous applications in modern signal processing and machine learning. We give the first efficient algorithm that provably recovers \mathbf{A}_0 when \mathbf{X}_0 has $O(n)$ nonzeros per column, under suitable probability model for \mathbf{X}_0 .

Our algorithmic pipeline centers around solving a certain nonconvex optimization problem with a spherical constraint, and hence is naturally phrased in the language of manifold optimization. In a companion paper [3], we have showed that with high probability our nonconvex formulation has no “spurious” local minimizers and around any saddle point the objective function has a negative directional curvature. In this paper, we take advantage of the particular geometric structure, and describe a Riemannian trust region algorithm that provably converges to a local minimizer with from arbitrary initializations. Such minimizers give excellent approximations to rows of \mathbf{X}_0 . The rows are then recovered by linear programming rounding and deflation.

Index Terms

Dictionary learning, Nonconvex optimization, Spherical constraint, Escaping saddle points, Trust-region method, Manifold optimization, Function landscape, Second-order geometry, Inverse problems, Structured signals, Nonlinear approximation

I. INTRODUCTION

Recently, there is a surge of research studying nonconvex formulations and provable algorithms for a number of central problems in signal processing and machine learning, including, e.g., low-rank matrix completion/recovery [4]–[25], phase retrieval [26]–[36], tensor recovery [37]–[41], mixed regression [42], [43], structured element pursuit [41], [44], blind deconvolution [45]–[49], noisy phase synchronization and community detection [50]–[52], deep learning [53], [54], numerical linear algebra and optimization [55], [56]. The research efforts are fruitful in producing more practical and scalable algorithms and even significantly better performance guarantees than known convex methods.

In a companion paper [3], we set out to understand the surprising effectiveness of nonconvex heuristics on the dictionary learning (DL) problem. In particular, we have focused on the complete dictionary recovery (DR) setting: given $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$, with $\mathbf{A}_0 \in \mathbb{R}^{n \times n}$ complete (i.e., square and invertible), and $\mathbf{X}_0 \in \mathbb{R}^{n \times p}$ obeying an i.i.d. Bernoulli-Gaussian (BG) model with rate θ (i.e., $[X_0]_{ij} = \Omega_{ij} Z_{ij}$ with $\Omega_{ij} \sim \text{Ber}(\theta)$ and $Z_{ij} \sim \mathcal{N}(0, 1)$), recover \mathbf{A}_0 and \mathbf{X}_0 . In this setting, $\text{row}(\mathbf{Y}) = \text{row}(\mathbf{X}_0)$, where $\text{row}(\cdot)$ denotes the row space. To first recover rows of \mathbf{X}_0 , we have tried to find the sparsest vectors in $\text{row}(\mathbf{Y})$, and proposed solving the nonconvex formulation

$$\text{minimize} \quad f(\mathbf{q}; \widehat{\mathbf{Y}}) \doteq \frac{1}{p} \sum_{k=1}^p h_{\mu}(\mathbf{q}^* \widehat{\mathbf{y}}_k) \quad \text{subject to} \quad \mathbf{q} \in \mathbb{S}^{n-1}, \quad (\text{I.1})$$

where $\widehat{\mathbf{Y}}$ is a proxy of \mathbf{Y} (i.e., after appropriate processing), $\widehat{\mathbf{y}}_k$ is the k -th column of $\widehat{\mathbf{Y}}$, and $h_{\mu}(z) \doteq \mu \log \cosh(z/\mu)$ is a (convex) smooth approximation to the absolute-value function. The spherical constraint renders the problem nonconvex.

JS, QQ, and JW are all with Electrical Engineering, Columbia University, New York, NY 10027, USA. Email: {js4038, qq2105, jw2966}@columbia.edu. An extended abstract of the current work has been published in [1]. Proofs of some secondary results are contained in the combined technical report [2].

Manuscript received xxx; revised xxx.

Despite the apparent nonconvexity, our prior analysis in [3] has showed that all local minimizers of (I.1) are qualitatively equally good, because each of them produces a close approximation to certain row of \mathbf{X}_0 (Corollary II.4 in [3]). So the central issue is how to escape from saddle points. Fortunately, our previous results (Theorem II.3 in [3]) imply that all saddle points under consideration are *ridable*, i.e., the associated Hessians have both strictly positive and strictly negative values (see the recapitulation in Section II-B). Particularly, eigenvectors of the negative eigenvalues are direction of negative curvature, which intuitively serve as directions of local descent.

Second-order methods can naturally exploit the curvature information to escape from ridable saddle points. To gain some intuition, consider an unconstrained optimization problem

$$\text{minimize}_{\mathbf{x} \in \mathbb{R}^n} \phi(\mathbf{x}).$$

The second-order Taylor expansion of ϕ at a saddle point \mathbf{x}_0 is

$$\widehat{\phi}(\boldsymbol{\delta}; \mathbf{x}_0) = \phi(\mathbf{x}_0) + \frac{1}{2} \boldsymbol{\delta}^* \nabla^2 \phi(\mathbf{x}_0) \boldsymbol{\delta}.$$

When $\boldsymbol{\delta}$ is chosen to align with an eigenvector of a negative eigenvalue $\lambda_{\text{neg}}[\nabla^2 \phi(\mathbf{x}_0)] < 0$, it holds that

$$\widehat{\phi}(\boldsymbol{\delta}; \mathbf{x}_0) - \phi(\mathbf{x}_0) \leq -|\lambda_{\text{neg}}[\nabla^2 \phi(\mathbf{x}_0)]| \|\boldsymbol{\delta}\|^2.$$

Thus, minimizing $\widehat{\phi}(\boldsymbol{\delta}; \mathbf{x}_0)$ returns a direction $\boldsymbol{\delta}_*$ that tends to decrease the objective ϕ , provided local approximation of $\widehat{\phi}$ to ϕ is reasonably accurate. Based on this intuition, we derive a (second-order) Riemannian trust-region algorithm that exploits the second-order information to escape from saddle points and provably returns a local minimizer to (I.1), from arbitrary initializations. We provide rigorous guarantees for recovering a local minimizer in Section II.

Obtaining a local minimizer only helps approximate one row of \mathbf{X}_0 . To recover the row, we derive a simple linear programming rounding procedure that provably works. To recover all rows of \mathbf{X}_0 , one repeats the above process based on a carefully designed deflation process. The whole algorithmic pipeline and the related recovery guarantees are provided in Section III. Particularly, we show that when p is reasonably large, with high probability (w.h.p.), our pipeline efficiently recovers \mathbf{A}_0 and \mathbf{X}_0 , even when each column of \mathbf{X}_0 contains $O(n)$ nonzeros.

A. Prior Arts and Connections

In Section II.E of the companion paper [3], we provide detailed comparisons of our results with prior theoretical results on DR; we conclude that this is the first algorithmic framework that guarantees efficient recovery of complete dictionaries when the coefficients have up to constant fraction of nonzeros. We also draw methodological connections to work on understanding nonconvex heuristics, and other nonconvex problems with similar geometric structures. Here we focus on drawing detailed connections to the optimization literature.

Trust-region method (TRM) has a rich history dating back to 40's; see the monograph [57] for accounts of the history and developments. The main motivation for early developments was to address limitations of the classic Newton's method (see, e.g., Section 3 of [58]). The limitations include the technical subtleties to establish local and global convergence results. Moreover, when the Hessian is singular or indefinite, the movement direction is either not well-defined, or does not improve the objective function. [59]–[64] initialized the line of work that addresses the limitations. Particularly, [58], [65] proposed using local second-order Taylor approximation as model function in the trust-region framework for unconstrained optimization. They showed that under mild conditions, the trust-region iterate sequence has a limit point that is critical and has positive semidefinite Hessian; see also Section 6.5-6.6 of [57]. Upon inspecting the relevant proofs, it seems not hard to strengthen the results to sequence convergence to local minimizers, under a ridable saddle condition as ours, for unconstrained optimization.

Research activities to port theories and algorithms of optimization in Euclidean space to Riemannian manifolds are best summarized by three monographs: [66]–[68]. [69] developed Newton and conjugate-gradient methods for the Stiefel manifolds, of which the sphere is a special case; [68] presents a complete set of first- and second-order Riemannian algorithms and convergence analyses; see also the excellent associated optimization software toolbox [70]. Among these, trust-region method was first ported to the Riemannian setting in [71], with emphasis on efficient implementation which only approximately solves the trust-region subproblem according to the Cauchy point scheme. The Cauchy point definition adopted there was the usual form based on the gradient, not strong enough to ensure the algorithm escape from ridable saddle points even if the true Hessian is in use in local approximation. In comparison,

in this work we assume that the trust-region subproblem is *exactly* solved, such that ridable saddles (the only possible saddles for our problem) are properly skipped. By this, we obtain the strong guarantee that the iterate sequence converges to a local minimizer, in contrast to the weak global convergence (gradient sequence converging to zero) or local convergence (sequence converging to a local minimizer within a small radius) established in [71]. To the best of our knowledge, our convergence result is first of its kind for a specific problem on sphere. After our initial submission, [72] has recently established worst-case iteration complexity of Riemannian TRM to converge to second-order critical points (i.e., critical points with positive semidefinite Hessians), echoing the results in the Euclidean case [73]. Their results are under mild Lipschitz-type assumptions and allow inexact subproblem solvers, and hence are very practical and general. However, on our particular problem, their result is considerably pessimistic, compared to our convergence result obtained from a specialized analysis.

Solving the trust-region subproblem exactly is expensive. Practically, often a reasonable approximate solution with controlled quality is adequate to guarantee convergence. In this regard, the truncated conjugate gradient (tCG) solver with a good initial search direction is commonly employed in practice (see, e.g., Section 7.5 in [57]). To ensure ridable saddle points be properly escaped from, the eigenpoint idea (see, e.g., Section 6.6 of [57]) is particularly relevant; see also Algorithm 3 and Lemma 10 in [72].

The benign function landscape we characterized in the first paper allows any reasonable iterative method that is capable of escaping from ridable saddles to find a local minimizer, with possibly different performance guarantees. The trust-region method we focus on here, and the curvilinear search method [61] are second-order methods that guarantee global optimization from arbitrary initializations. Typical first-order methods such as the vanilla gradient descent can only guarantee convergence to a critical point. Nonetheless, for our particular function, noisy/stochastic gradient method guarantees to find a local minimizer from an arbitrary initialization with high probability [74].

B. Notations, and Reproducible Research

We use bold capital and small letters such as \mathbf{X} and \mathbf{x} to denote matrices and vectors, respectively. Small letters are reserved for scalars. Several specific mathematical objects we will frequently work with: O_k for the orthogonal group of order k , \mathbb{S}^{n-1} for the unit sphere in \mathbb{R}^n , \mathbb{B}^n for the unit ball in \mathbb{R}^n , and $[m] \doteq \{1, \dots, m\}$ for positive integers m . We use $(\cdot)^*$ for matrix transposition, causing no confusion as we will work entirely on the real field. We use superscript to index rows of a matrix, such as \mathbf{x}^i for the i -th row of the matrix \mathbf{X} , and subscript to index columns, such as \mathbf{x}_j . All vectors are defaulted to column vectors. So the i -th row of \mathbf{X} as a row vector will be written as $(\mathbf{x}^i)^*$. For norms, $\|\cdot\|$ is the usual ℓ^2 norm for a vector and the operator norm (i.e., $\ell^2 \rightarrow \ell^2$) for a matrix; all other norms will be indexed by subscript, for example the Frobenius norm $\|\cdot\|_F$ for matrices and the element-wise max-norm $\|\cdot\|_\infty$. We use $\mathbf{x} \sim \mathcal{L}$ to mean that the random variable \mathbf{x} is distributed according to the law \mathcal{L} . Let \mathcal{N} denote the Gaussian law. Then $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ means that \mathbf{x} is a standard Gaussian vector. Similarly, we use $\mathbf{x} \sim_{i.i.d.} \mathcal{L}$ to mean elements of \mathbf{x} are independently and identically distributed according to the law \mathcal{L} . So the fact $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is equivalent to that $\mathbf{x} \sim_{i.i.d.} \mathcal{N}(0, 1)$. One particular distribution of interest for this paper is the Bernoulli-Gaussian with rate θ : $Z \sim B \cdot G$, with $G \sim \mathcal{N}(0, 1)$ and $B \sim \text{Ber}(\theta)$. We also write this compactly as $Z \sim \text{BG}(\theta)$. We frequently use indexed C and c for numerical constants when stating and proving technical results. The scopes of such constants are local unless otherwise noted. We use standard notations for most other cases, with exceptions clarified locally.

The codes to reproduce all the figures and experimental results are available online:

https://github.com/sunju/dl_focm .

II. FINDING ONE LOCAL MINIMIZER VIA THE RIEMANNIAN TRUST-REGION METHOD

We are interested to seek a local minimizer of (I.1). The presence of saddle points have motivated us to develop a second-order Riemannian trust-region algorithm over the sphere; the existence of descent directions at nonoptimal points drives the trust-region iteration sequence towards one of the minimizers asymptotically. We will prove that under our modeling assumptions, this algorithm with an arbitrary initialization efficiently produces an accurate approximation¹ to one of the minimizers. Throughout the exposition, basic knowledge of Riemannian geometry is assumed. We will try to keep the technical requirement minimal possible; the reader can consult the excellent monograph [68] for relevant background and details.

¹By “accurate” we mean one can achieve an arbitrary numerical accuracy $\varepsilon > 0$ with a reasonable amount of time. Here the running time of the algorithm is on the order of $\log \log(1/\varepsilon)$ in the target accuracy ε , and polynomial in other problem parameters.

A. Some Basic Facts about the Sphere and f

For any point $\mathbf{q} \in \mathbb{S}^{n-1}$, the tangent space $T_{\mathbf{q}}\mathbb{S}^{n-1}$ and the orthoprojector $\mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}}$ onto $T_{\mathbf{q}}\mathbb{S}^{n-1}$ are given by

$$\begin{aligned} T_{\mathbf{q}}\mathbb{S}^{n-1} &= \{\boldsymbol{\delta} \in \mathbb{R}^n : \mathbf{q}^* \boldsymbol{\delta} = 0\}, \\ \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} &= \mathbf{I} - \mathbf{q}\mathbf{q}^* = \mathbf{U}\mathbf{U}^*, \end{aligned}$$

where $\mathbf{U} \in \mathbb{R}^{n \times (n-1)}$ is an arbitrary orthonormal basis for $T_{\mathbf{q}}\mathbb{S}^{n-1}$ (note that the orthoprojector is independent of the basis \mathbf{U} we choose). Consider any $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$. The map

$$\gamma(t) : t \mapsto \mathbf{q} \cos(t \|\boldsymbol{\delta}\|) + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \sin(t \|\boldsymbol{\delta}\|)$$

defines a smooth curve on the sphere that satisfies $\gamma(0) = \mathbf{q}$ and $\dot{\gamma}(0) = \boldsymbol{\delta}$. Geometrically, $\gamma(t)$ is a segment of the great circle that passes \mathbf{q} and has $\boldsymbol{\delta}$ as its tangent vector at \mathbf{q} . The exponential map for $\boldsymbol{\delta}$ is defined as

$$\exp_{\mathbf{q}}(\boldsymbol{\delta}) \doteq \gamma(1) = \mathbf{q} \cos \|\boldsymbol{\delta}\| + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \sin \|\boldsymbol{\delta}\|.$$

It is a canonical way of pulling $\boldsymbol{\delta}$ to the sphere.




Fig. 1: Illustrations of the tangent space $T_{\mathbf{q}}\mathbb{S}^{n-1}$ and exponential map $\exp_{\mathbf{q}}(\boldsymbol{\delta})$ defined on the sphere \mathbb{S}^{n-1} .

In this paper we are interested in the restriction of f to the unit sphere \mathbb{S}^{n-1} . For the sake of performing optimization, we need local approximations of f . Instead of directly approximating the function in \mathbb{R}^n , we form quadratic approximations of f in the tangent spaces of \mathbb{S}^{n-1} . We consider the smooth function $f \circ \exp_{\mathbf{q}}(\boldsymbol{\delta}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{R}$, where \circ is the usual function composition operator. An applications of vector space Taylor's theorem gives

$$f \circ \exp_{\mathbf{q}}(\boldsymbol{\delta}) \approx f(\mathbf{q}; \hat{\mathbf{Y}}) + \langle \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \boldsymbol{\delta} \rangle + \frac{1}{2} \boldsymbol{\delta}^* \left(\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I} \right) \boldsymbol{\delta}$$

when $\|\boldsymbol{\delta}\|$ is small. Thus, we form a quadratic approximation $\hat{f}(\boldsymbol{\delta}; \mathbf{q}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{R}$ as

$$\hat{f}(\boldsymbol{\delta}; \mathbf{q}, \hat{\mathbf{Y}}) \doteq f(\mathbf{q}; \hat{\mathbf{Y}}) + \left\langle \underbrace{\nabla f(\mathbf{q}; \hat{\mathbf{Y}})}_{\text{gradient}}, \boldsymbol{\delta} \right\rangle + \frac{1}{2} \boldsymbol{\delta}^* \left(\underbrace{\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I}}_{\text{Hessian}} \right) \boldsymbol{\delta}. \quad (\text{II.1})$$

Here $\nabla f(\mathbf{q})$ and $\nabla^2 f(\mathbf{q})$ denote the usual (Euclidean) gradient and Hessian of f w.r.t. \mathbf{q} in \mathbb{R}^n . For our specific f defined in (I.1), it is easy to check that

$$\nabla f(\mathbf{q}; \hat{\mathbf{Y}}) = \frac{1}{p} \sum_{k=1}^p \tanh\left(\frac{\mathbf{q}^* \hat{\mathbf{y}}_k}{\mu}\right) \hat{\mathbf{y}}_k, \quad (\text{II.2})$$

$$\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) = \frac{1}{p} \sum_{k=1}^p \frac{1}{\mu} \left[1 - \tanh^2\left(\frac{\mathbf{q}^* \hat{\mathbf{y}}_k}{\mu}\right) \right] \hat{\mathbf{y}}_k \hat{\mathbf{y}}_k^*. \quad (\text{II.3})$$

The quadratic approximation also naturally gives rise to the Riemannian gradient and Riemannian Hessian defined on $T_{\mathbf{q}}\mathbb{S}^{n-1}$ as

$$\text{grad } f(\mathbf{q}; \hat{\mathbf{Y}}) = \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \quad (\text{II.4})$$

$$\text{Hess } f(\mathbf{q}; \hat{\mathbf{Y}}) = \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \left(\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I} \right) \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}}. \quad (\text{II.5})$$

Thus, the above quadratic approximation can be rewritten compactly as

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{q}, \widehat{\mathbf{Y}}) = f(\mathbf{q}; \widehat{\mathbf{Y}}) + \langle \boldsymbol{\delta}, \text{grad } f(\mathbf{q}; \widehat{\mathbf{Y}}) \rangle + \frac{1}{2} \boldsymbol{\delta}^* \text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}}) \boldsymbol{\delta}, \quad \forall \boldsymbol{\delta} \in T_{\mathbf{q}} \mathbb{S}^{n-1}.$$

The first order necessary condition for *unconstrained* minimization of function \widehat{f} over $T_{\mathbf{q}} \mathbb{S}^{n-1}$ is

$$\text{grad } f(\mathbf{q}; \widehat{\mathbf{Y}}) + \text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}}) \boldsymbol{\delta}_* = \mathbf{0}. \quad (\text{II.6})$$

If $\text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}})$ is positive semidefinite and has “full rank” $n - 1$ (hence “nondegenerate”²), the unique solution $\boldsymbol{\delta}_*$ is

$$\boldsymbol{\delta}_* = -\mathbf{U} \left(\mathbf{U}^* \left[\text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}}) \right] \mathbf{U} \right)^{-1} \mathbf{U}^* \text{grad } f(\mathbf{q}),$$

which is also invariant to the choice of basis \mathbf{U} . Given a tangent vector $\boldsymbol{\delta} \in T_{\mathbf{q}} \mathbb{S}^{n-1}$, let $\gamma(t) \doteq \exp_{\mathbf{q}}(t\boldsymbol{\delta})$ denote a geodesic curve on \mathbb{S}^{n-1} . Following the notation of [68], let

$$\mathcal{P}_{\gamma}^{\tau \leftarrow 0} : T_{\mathbf{q}} \mathbb{S}^{n-1} \rightarrow T_{\gamma(\tau)} \mathbb{S}^{n-1}$$

denotes the parallel translation operator, which translates the tangent vector $\boldsymbol{\delta}$ at $\mathbf{q} = \gamma(0)$ to a tangent vector at $\gamma(\tau)$, in a “parallel” manner. In the sequel, we identify $\mathcal{P}_{\gamma}^{\tau \leftarrow 0}$ with the following $n \times n$ matrix, whose restriction to $T_{\mathbf{q}} \mathbb{S}^{n-1}$ is the parallel translation operator (the detailed derivation can be found in Chapter 8.1 of [68]):

$$\begin{aligned} \mathcal{P}_{\gamma}^{\tau \leftarrow 0} &= \left(\mathbf{I} - \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} \right) - \mathbf{q} \sin(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \cos(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} \\ &= \mathbf{I} + (\cos(\tau \|\boldsymbol{\delta}\|) - 1) \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \sin(\tau \|\boldsymbol{\delta}\|) \frac{\mathbf{q} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|}. \end{aligned} \quad (\text{II.7})$$

Similarly, following the notation of [68], we denote the inverse of this matrix by $\mathcal{P}_{\gamma}^{0 \leftarrow \tau}$, where its restriction to $T_{\gamma(\tau)} \mathbb{S}^{n-1}$ is the inverse of the parallel translation operator $\mathcal{P}_{\gamma}^{\tau \leftarrow 0}$.

B. The Geometric Results from [3]

We reproduce the main geometric theorems from [3] here for the sake of completeness. To characterize the function landscape of $f(\mathbf{q}; \mathbf{X}_0)$ over \mathbb{S}^{n-1} , we mostly work with the function

$$g(\mathbf{w}; \mathbf{X}_0) \doteq f(\mathbf{q}(\mathbf{w}); \mathbf{X}_0) = \frac{1}{p} \sum_{k=1}^p h_{\mu}(\mathbf{q}(\mathbf{w})^* (\mathbf{x}_0)_k), \quad (\text{II.8})$$

induced by the reparametrization

$$\mathbf{q}(\mathbf{w}) = \left(\mathbf{w}, \sqrt{1 - \|\mathbf{w}\|^2} \right), \quad \mathbf{w} \in \mathbb{B}^{n-1}. \quad (\text{II.9})$$

Geometrically, this corresponds to projection of the function f above the equatorial section e_n^{\perp} onto e_n^{\perp} (see Fig. 2 (right) for illustration). In particular, we focus our attention to the smaller set of the ball:

$$\Gamma = \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\} \subsetneq \mathbb{B}^{n-1}, \quad (\text{II.10})$$

because $\mathbf{q}(\Gamma)$ contains all points $\mathbf{q} \in \mathbb{S}^{n-1}$ with $n \in \arg \max_{i \in \pm[n]} \mathbf{q}^* \mathbf{e}_i$. We can similarly characterize other parts of f on \mathbb{S}^{n-1} using projection onto other equatorial sections.

Theorem II.1 (High-dimensional landscape - orthogonal dictionary). *Suppose $\mathbf{A}_0 = \mathbf{I}$ and hence $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0 = \mathbf{X}_0$. There exist positive constants c_* and C , such that for any $\theta \in (0, 1/2)$ and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, whenever*

$$p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta}, \quad (\text{II.11})$$

²Note that the $n \times n$ matrix $\text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}})$ has rank at most $n - 1$, as the nonzero \mathbf{q} obviously is in its null space. When $\text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}})$ has rank $n - 1$, it has no null direction in the tangent space. Thus, in this case it acts on the tangent space like a full-rank matrix.

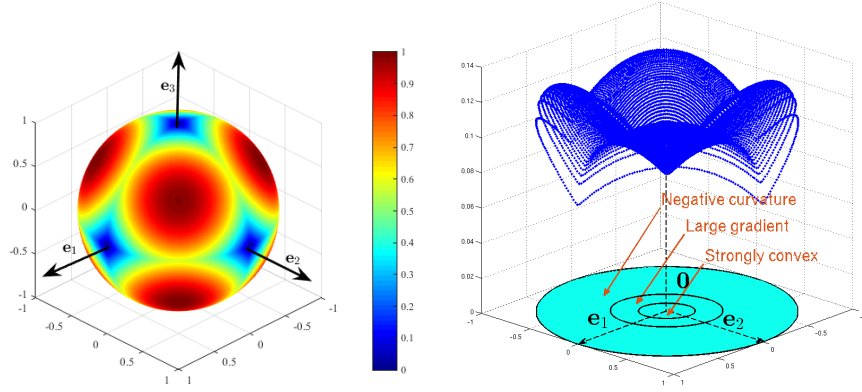


Fig. 2: **Why is dictionary learning over \mathbb{S}^{n-1} tractable?** Assume the target dictionary $\mathbf{A}_0 = \mathbf{I}$. **Left:** Large sample objective function $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q})]$. The only local minimizers are the standard basis vectors \mathbf{e}_i 's and their negatives. **Right:** A visualization of the function as a height above the equatorial section \mathbf{e}_3^\perp , i.e., $\text{span}\{\mathbf{e}_1, \mathbf{e}_2\} \cap \mathbb{B}^3$. The derived function is obtained by assigning values of points on the upper hemisphere to their corresponding projections on the equatorial section \mathbf{e}_3^\perp . The minimizers for the derived function are $\mathbf{0}, \pm\mathbf{e}_1, \pm\mathbf{e}_2$. Around $\mathbf{0}$ in \mathbf{e}_3^\perp , the function exhibits a small region of strong convexity, a region of large gradient, and finally a region in which the direction away from $\mathbf{0}$ is a direction of negative curvature.

the following hold simultaneously with probability at least $1 - c_b p^{-6}$:

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0) \succeq \frac{c_* \theta}{\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (\text{II.12})$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \geq c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (\text{II.13})$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (\text{II.14})$$

and the function $g(\mathbf{w}; \mathbf{X}_0)$ has exactly one local minimizer \mathbf{w}_* over the open set $\Gamma \doteq \{\mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}}\}$, which satisfies

$$\|\mathbf{w}_* - \mathbf{0}\| \leq \min \left\{ \frac{c_c \mu}{\theta} \sqrt{\frac{n \log p}{p}}, \frac{\mu}{16} \right\}. \quad (\text{II.15})$$

Here c_a through c_c are all positive constants.

Recall that the reason we just need to characterize the geometry for the case $\mathbf{A}_0 = \mathbf{I}$ is that for other orthogonal \mathbf{A}_0 , the function landscape is simply a rotated version of that of $\mathbf{A}_0 = \mathbf{I}$.

Theorem II.2 (High-dimensional landscape - complete dictionary). *Suppose \mathbf{A}_0 is complete with its condition number $\kappa(\mathbf{A}_0)$. There exist positive constants c_* (particularly, the same constant as in Theorem II.1) and C , such that for any $\theta \in (0, 1/2)$ and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, when*

$$p \geq \frac{C}{c_*^2 \theta^2} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right) \quad (\text{II.16})$$

and $\bar{\mathbf{Y}} \doteq \sqrt{p\theta} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y}$, $\mathbf{U}\Sigma\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$, the following hold simultaneously with probability at least $1 - c_b p^{-6}$:

$$\nabla^2 g(\mathbf{w}; \mathbf{V}\mathbf{U}^* \bar{\mathbf{Y}}) \succeq \frac{c_* \theta}{2\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (\text{II.17})$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{V}\mathbf{U}^* \bar{\mathbf{Y}})}{\|\mathbf{w}\|} \geq \frac{1}{2} c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (\text{II.18})$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{V} \mathbf{U}^* \bar{\mathbf{Y}}) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -\frac{1}{2} c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (\text{II.19})$$

and the function $g(\mathbf{w}; \mathbf{V} \mathbf{U}^* \bar{\mathbf{Y}})$ has exactly one local minimizer \mathbf{w}_* over the open set $\Gamma \doteq \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\}$, which satisfies

$$\|\mathbf{w}_* - \mathbf{0}\| \leq \mu/7. \quad (\text{II.20})$$

Here c_a, c_b are both positive constants.

From the above theorems, it is clear that for any saddle point in the \mathbf{w} space, the Hessian of g has at least one negative eigenvalue. Now the problem is whether all saddle points of f on \mathbb{S}^{n-1} are “ridable”, because as alluded to in previous discussion, we need to perform actual optimization in the \mathbf{q} space. Instead of presenting a rigorous technical statement and detailed proof, we include here just an informal argument; our actual proof for algorithmic convergence runs back and forth in \mathbf{w} and \mathbf{q} space and such lack will not affect our arguments there.

It is very easy to verify the following fact (see proof of Lemma II.13 on page 33):

$$\langle \text{grad } f(\mathbf{q}), \mathbf{q} - \mathbf{e}_n/q_n \rangle = \langle \mathbf{w}, \nabla g(\mathbf{w}) \rangle.$$

Thus, $\langle \text{grad } f(\mathbf{q}), \mathbf{q} - \mathbf{e}_n/q_n \rangle \neq 0$ if and only if $\langle \mathbf{w}, \nabla g(\mathbf{w}) \rangle \neq 0$, implying that $\text{grad } f(\mathbf{q})$ will never be zero in the spherical region corresponding to $\left\{ \mathbf{w} : \mu/(4\sqrt{2}) \leq \|\mathbf{w}\| \leq 1/(20\sqrt{5}) \right\}$. Moreover, it is shown in Lemma II.11 below that the Riemannian Hessian is positive definite for the spherical region corresponding to $\left\{ \mathbf{w} : \|\mathbf{w}\| \leq \mu/(4\sqrt{2}) \right\}$, so there is no saddle point in this region either. Over $\mathbf{q}(\Gamma) \cap \mathbb{S}^{n-1}$, potential saddle points lie only in the region corresponding to $\left\{ \mathbf{w} : 1/(20\sqrt{5}) \leq \|\mathbf{w}\| \leq \sqrt{(4n-1)/(4n)} \right\}$. Theorem II.1 and Theorem II.2 imply that around each point in this region, a cross section of the function $g(\mathbf{w})$ is *strictly* concave locally. Intuitively, by the $\mathbf{q}(\mathbf{w})$ mapping the same happens in the \mathbf{q} space, i.e., the Riemannian Hessian has a strictly negative eigenvalue.

C. The Riemannian Trust-Region Algorithm over the Sphere

For a function f in the Euclidean space, the typical TRM starts from some initialization $\mathbf{q}^{(0)} \in \mathbb{R}^n$, and produces a sequence of iterates $\mathbf{q}^{(1)}, \mathbf{q}^{(2)}, \dots$, by repeatedly minimizing a quadratic approximation \hat{f} to the objective function $f(\mathbf{q})$, over a ball centered around the current iterate.

For our f defined over \mathbb{S}^{n-1} , given the previous iterate $\mathbf{q}^{(r-1)}$, the TRM produces the next movement by generating a solution $\hat{\delta}$ to

$$\text{minimize}_{\delta \in T_{\mathbf{q}^{(r-1)}} \mathbb{S}^{n-1}, \|\delta\| \leq \Delta} \hat{f}(\delta; \mathbf{q}^{(r-1)}), \quad (\text{II.21})$$

where $\hat{f}(\delta; \mathbf{q}^{(r-1)})$ is the local quadratic approximation defined in (II.1). The solution $\hat{\delta}$ is then pulled back to \mathbb{S}^{n-1} from $T_{\mathbf{q}^{(r-1)}} \mathbb{S}^{n-1}$. If we choose the exponential map to pull back the movement $\hat{\delta}$,³ the next iterate then reads

$$\mathbf{q}^{(r)} = \mathbf{q}^{(r-1)} \cos \|\hat{\delta}\| + \frac{\hat{\delta}}{\|\hat{\delta}\|} \sin \|\hat{\delta}\|. \quad (\text{II.22})$$

To solve the subproblem (II.21) numerically, we can take any matrix $\mathbf{U} \in \mathbb{R}^{n \times (n-1)}$ whose columns form an orthonormal basis for $T_{\mathbf{q}^{(r-1)}} \mathbb{S}^{n-1}$, and produce a solution $\hat{\xi}$ to

$$\text{minimize}_{\|\xi\| \leq \Delta} \hat{f}(\mathbf{U}\xi; \mathbf{q}^{(r-1)}). \quad (\text{II.23})$$

Solution to (II.21) can then be recovered as $\hat{\delta} = \mathbf{U}\hat{\xi}$.

The problem (II.23) is an instance of the classic *trust region subproblem*, i.e., minimizing a quadratic function subject to a single quadratic constraint. Albeit potentially nonconvex, this notable subproblem can be solved in polynomial time by several numerical methods [57], [65], [75]–[78]. Approximate solution of the subproblem suffices to guarantee convergence in theory, and lessens the storage and computational burden in practice. We will deploy the approximate version in simulations. For simplicity, however, our subsequent analysis assumes the subproblem is solved *exactly*. We next briefly describe how one can deploy the semidefinite programming (SDP) approach [75]–[78]

³The exponential map is only one of the many possibilities; also for general manifolds other retraction schemes may be more practical. See exposition on retraction in Chapter 4 of [68].

to solve the subproblem exactly. This choice is due to the well-known effectiveness and robustness of the SDP approach on this problem. We introduce

$$\tilde{\boldsymbol{\xi}} = [\boldsymbol{\xi}^*, 1]^*, \quad \boldsymbol{\Theta} = \tilde{\boldsymbol{\xi}}\tilde{\boldsymbol{\xi}}^*, \quad \mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^* & 0 \end{bmatrix}, \quad (\text{II.24})$$

where $\mathbf{A} = \mathbf{U}^* \text{Hess } f(\mathbf{q}^{(r-1)}; \widehat{\mathbf{Y}}) \mathbf{U}$ and $\mathbf{b} = \mathbf{U}^* \text{grad } \nabla f(\mathbf{q}^{(r-1)}; \widehat{\mathbf{Y}})$. The resulting SDP to solve is

$$\text{minimize } \langle \mathbf{M}, \boldsymbol{\Theta} \rangle, \quad \text{subject to } \text{tr}(\boldsymbol{\Theta}) \leq \Delta^2 + 1, \quad \langle \mathbf{E}_{n+1}, \boldsymbol{\Theta} \rangle = 1, \quad \boldsymbol{\Theta} \succeq \mathbf{0}, \quad (\text{II.25})$$

where $\mathbf{E}_{n+1} = \mathbf{e}_{n+1}\mathbf{e}_{n+1}^*$. Once the problem (II.25) is solved to its optimum $\boldsymbol{\Theta}_*$, one can provably recover the minimizer $\boldsymbol{\xi}_*$ of (II.23) by computing the SVD of $\boldsymbol{\Theta}_* = \widetilde{\mathbf{U}}\boldsymbol{\Sigma}\widetilde{\mathbf{V}}^*$, and extract as a subvector the first $n-1$ coordinates of the principal eigenvector $\widetilde{\mathbf{u}}_1$ (see Appendix B of [79]).

D. Main Convergence Results

Using general convergence results on Riemannian TRM (see, e.g., Chapter 7 of [68]), it is not difficult to prove that the gradient sequence $\text{grad } f(\mathbf{q}^{(r)}; \widehat{\mathbf{Y}})$ produced by TRM converges to zero (i.e., global convergence), or the sequence converges (at quadratic rate) to a local minimizer if the initialization is already close a local minimizer (i.e., local convergence). In this section, we show that under our probabilistic assumptions, these results can be substantially strengthened. In particular, the algorithm is guaranteed to produce an accurate approximation to a local minimizer of the objective function, in a number of iterations that is polynomial in the problem size, from arbitrary initializations. The arguments in the companion paper [3] showed that w.h.p. every local minimizer of f produces a close approximation to a row of \mathbf{X}_0 . Taken together, this implies that the algorithm efficiently produces a close approximation to one row of \mathbf{X}_0 .

Thorough the analysis, we assume the trust-region subproblem is exactly solved and the step size parameter Δ is fixed. Our next two theorems summarize the convergence results for orthogonal and complete dictionaries, respectively.

Theorem II.3 (TRM convergence - orthogonal dictionary). *Suppose the dictionary \mathbf{A}_0 is orthogonal. There exists a positive constant C , such that for all $\theta \in (0, 1/2)$ and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, whenever*

$$p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta},$$

with probability at least $1 - c_b p^{-6}$, the Riemannian trust-region algorithm with input data matrix $\widehat{\mathbf{Y}} = \mathbf{Y}$, any initialization $\mathbf{q}^{(0)}$ on the sphere, and a step size satisfying

$$\Delta \leq \frac{c_c c_*^3 \theta^3 \mu^2}{n^{7/2} \log^{7/2}(np)} \quad (\text{II.26})$$

returns a solution $\widehat{\mathbf{q}} \in \mathbb{S}^{n-1}$ which is ε near to one of the local minimizers \mathbf{q}_* (i.e., $\|\widehat{\mathbf{q}} - \mathbf{q}_*\| \leq \varepsilon$) in at most

$$\max \left\{ \frac{c_d n^6 \log^3(np)}{c_*^3 \theta^3 \mu^4}, \frac{c_e n}{c_*^2 \theta^2 \Delta^2} \right\} f(\mathbf{q}^{(0)}) + \log \log \frac{c_f c_* \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \quad (\text{II.27})$$

iterations. Here c_* is as defined in Theorem II.1, and c_a through c_f are all positive constants.

Theorem II.4 (TRM convergence - complete dictionary). *Suppose the dictionary \mathbf{A}_0 is complete with condition number $\kappa(\mathbf{A}_0)$. There exists a positive constant C , such that for all $\theta \in (0, 1/2)$, and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, whenever*

$$p \geq \frac{C}{c_*^2 \theta^2} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right),$$

with probability at least $1 - c_b p^{-6}$, the Riemannian trust-region algorithm with input data matrix $\overline{\mathbf{Y}} \doteq \sqrt{p\theta} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y}$ where $\mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$, any initialization $\mathbf{q}^{(0)}$ on the sphere and a step size satisfying

$$\Delta \leq \frac{c_c c_*^3 \theta^3 \mu^2}{n^{7/2} \log^{7/2}(np)} \quad (\text{II.28})$$

returns a solution $\hat{\mathbf{q}} \in \mathbb{S}^{n-1}$ which is ε near to one of the local minimizers \mathbf{q}_\star (i.e., $\|\hat{\mathbf{q}} - \mathbf{q}_\star\| \leq \varepsilon$) in at most

$$\max \left\{ \frac{c_d n^6 \log^3(np)}{c_\star^3 \theta^3 \mu^4}, \frac{c_e n}{c_\star^2 \theta^2 \Delta^2} \right\} f(\mathbf{q}^{(0)}) + \log \log \frac{c_f c_\star \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \quad (\text{II.29})$$

iterations. Here c_\star is as in Theorem II.1, and c_a through c_f are all positive constants.

Our convergence result shows that for any target accuracy $\varepsilon > 0$ the algorithm terminates within polynomially many steps. Specifically, the first summand in (II.27) or (II.29) is the number of steps the sequence takes to enter the strongly convex region and be “reasonably” close to a local minimizer. All subsequent trust-region subproblems are then unconstrained (proved below) – the constraint is inactive at optimal point, and hence the steps behave like Newton steps. The second summand reflects the typical quadratic local convergence of the Newton steps.

Our estimate of the number of steps is pessimistic: the running time is a relatively high-degree polynomial in p and n . We will discuss practical implementation details that help speed up in Section IV. Our goal in stating the above results is not to provide a tight analysis, but to prove that the Riemannian TRM algorithm finds a local minimizer in polynomial time. For nonconvex problems, this is not entirely trivial – results of [80] show that in general it is NP-hard to find a local minimizer of a nonconvex function.

E. Sketch of Proof for Orthogonal Dictionaries

The reason that our algorithm is successful derives from the geometry formalized in Theorem II.1. Basically, the sphere \mathbb{S}^{n-1} can be divided into three regions. Near each local minimizer, the function is strongly convex, and the algorithm behaves like a standard (Euclidean) TRM algorithm applied to a strongly convex function – in particular, it exhibits a quadratic asymptotic rate of convergence. Away from local minimizers, the function always exhibits either a strong gradient, or a direction of negative curvature (i.e., the Hessian has a strictly negative eigenvalue). The Riemannian TRM algorithm is capable of exploiting these quantities to reduce the objective value by at least a fixed amount in each iteration. The total number of iterations spent away from the vicinity of the local minimizers can be bounded by comparing this amount to the initial objective value. Our proofs follow exactly this line and make the various quantities precise.

Note that for any orthogonal \mathbf{A}_0 , $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0) = f(\mathbf{A}_0^* \mathbf{q}; \mathbf{X}_0)$. In words, this is the established fact that the function landscape of $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0)$ is a rotated version of that of $f(\mathbf{q}; \mathbf{X}_0)$. Thus, any local minimizer \mathbf{q}_\star of $f(\mathbf{q}; \mathbf{X}_0)$ is rotated to $\mathbf{A}_0 \mathbf{q}_\star$, a local minimizer of $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0)$. Also if our algorithm generates iteration sequence $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{X}_0)$ upon initialization \mathbf{q}_0 , it will generate the iteration sequence $\mathbf{A}_0 \mathbf{q}_0, \mathbf{A}_0 \mathbf{q}_1, \mathbf{A}_0 \mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0)$. So w.l.o.g. it is adequate that we prove the convergence results for the case $\mathbf{A}_0 = \mathbf{I}$. So in this section (Section II-E), we write $f(\mathbf{q})$ to mean $f(\mathbf{q}; \mathbf{X}_0)$.

We partition the sphere into three regions, for which we label as R_{I} , R_{II} , R_{III} , corresponding to the strongly convex, nonzero gradient, and negative curvature regions, respectively (see Theorem II.1). That is, R_{I} consists of a union of $2n$ spherical caps of radius $\mu/(4\sqrt{2})$, each centered around a signed standard basis vector $\pm e_i$. R_{II} consist of the set difference of a union of $2n$ spherical caps of radius $1/(20\sqrt{5})$, centered around the standard basis vectors $\pm e_i$, and R_{I} . Finally, R_{III} covers the rest of the sphere. We say a trust-region step takes an R_{I} step if the current iterate is in R_{I} ; similarly for R_{II} and R_{III} steps. Since we use the geometric structures derived in Theorem II.1 and Corollary II.2 in [3], the conditions

$$\theta \in (0, 1/2), \quad \mu < c \min \left\{ \theta n^{-1}, n^{-5/4} \right\}, \quad p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta} \quad (\text{II.30})$$

are always in force.

At step r of the algorithm, suppose $\delta^{(r)}$ is the minimizer of the trust-region subproblem (II.21). We call the step “constrained” if $\|\delta^{(r)}\| = \Delta$ (the minimizer lies on the boundary and hence the constraint is active), and call it “unconstrained” if $\|\delta^{(r)}\| < \Delta$ (the minimizer lies in the relative interior and hence the constraint is not in force). Thus, in the unconstrained case the optimality condition is (II.6).

The next lemma provides some estimates about ∇f and $\nabla^2 f$ that are useful in various contexts.

Lemma II.5. *We have the following estimates about ∇f and $\nabla^2 f$:*

$$\begin{aligned} \sup_{\mathbf{q} \in \mathbb{S}^{n-1}} \|\nabla f(\mathbf{q})\| &\doteq M_{\nabla} \leq \sqrt{n} \|\mathbf{X}_0\|_{\infty}, \\ \sup_{\mathbf{q} \in \mathbb{S}^{n-1}} \|\nabla^2 f(\mathbf{q})\| &\doteq M_{\nabla^2} \leq \frac{n}{\mu} \|\mathbf{X}_0\|_{\infty}^2, \\ \sup_{\mathbf{q}, \mathbf{q}' \in \mathbb{S}^{n-1}, \mathbf{q} \neq \mathbf{q}'} \frac{\|\nabla f(\mathbf{q}) - \nabla f(\mathbf{q}')\|}{\|\mathbf{q} - \mathbf{q}'\|} &\doteq L_{\nabla} \leq \frac{n}{\mu} \|\mathbf{X}_0\|_{\infty}^2, \\ \sup_{\mathbf{q}, \mathbf{q}' \in \mathbb{S}^{n-1}, \mathbf{q} \neq \mathbf{q}'} \frac{\|\nabla^2 f(\mathbf{q}) - \nabla^2 f(\mathbf{q}')\|}{\|\mathbf{q} - \mathbf{q}'\|} &\doteq L_{\nabla^2} \leq \frac{2}{\mu^2} n^{3/2} \|\mathbf{X}_0\|_{\infty}^3. \end{aligned}$$

Proof. See Page 25 under Section VI. □

Our next lemma says if the trust-region step size Δ is small enough, one Riemannian trust-region step reduces the objective value by a certain amount when there is any descent direction.

Lemma II.6. *Suppose that the trust region size $\Delta \leq 1$, and there exists a tangent vector $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ with $\|\boldsymbol{\delta}\| \leq \Delta$, such that*

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) \leq f(\mathbf{q}) - s$$

for some positive scalar $s \in \mathbb{R}$. Then the trust region subproblem produces a point $\boldsymbol{\delta}_{\star}$ with

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta}_{\star})) \leq f(\mathbf{q}) - s + \frac{1}{3}\eta_f\Delta^3,$$

where $\eta_f \doteq M_{\nabla} + 2M_{\nabla^2} + L_{\nabla} + L_{\nabla^2}$ and M_{∇} , M_{∇^2} , L_{∇} , L_{∇^2} are the quantities defined in Lemma II.5.

Proof. See Page 26 under Section VI. □

To show decrease in objective value for R_{II} and R_{III} , now it is enough to exhibit a descent direction for each point in these regions. The next two lemmas help us almost accomplish the goal. For convenience again we choose to state the results for the ‘‘canonical’’ section that is in the vicinity of \mathbf{e}_n and the projection map $\mathbf{q}(\mathbf{w}) = [\mathbf{w}; (1 - \|\mathbf{w}\|^2)^{1/2}]$, with the idea that similar statements hold for other symmetric sections.

Lemma II.7. *Suppose that the trust region size $\Delta \leq 1$, $\mathbf{w}^* \nabla g(\mathbf{w}) / \|\mathbf{w}\| \geq \beta_g$ for some scalar β_g , and that $\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} / \|\mathbf{w}\|^2$ is L_g -Lipschitz on an open ball $\mathcal{B}\left(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}}\right)$ centered at \mathbf{w} . Then there exists a tangent vector $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ with $\|\boldsymbol{\delta}\| \leq \Delta$, such that*

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) \leq f(\mathbf{q}) - \min \left\{ \frac{\beta_g^2}{2L_g}, \frac{3\beta_g\Delta}{4\pi\sqrt{n}} \right\}.$$

Proof. See Page 27 under Section VI. □

Lemma II.8. *Suppose that the trust-region size $\Delta \leq 1$, $\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} / \|\mathbf{w}\|^2 \leq -\beta_{\wedge}$, for some β_{\wedge} , and that $\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} / \|\mathbf{w}\|^2$ is L_{\wedge} Lipschitz on the open ball $\mathcal{B}\left(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}}\right)$ centered at \mathbf{w} . Then there exists a tangent vector $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ with $\|\boldsymbol{\delta}\| \leq \Delta$, such that*

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) \leq f(\mathbf{q}) - \min \left\{ \frac{2\beta_{\wedge}^3}{3L_{\wedge}^2}, \frac{3\Delta^2\beta_{\wedge}}{8\pi^2 n} \right\}.$$

Proof. See Page 28 under Section VI. □

One can take $\beta_g = \beta_{\wedge} = c_{\star}\theta$ as shown in Theorem II.1, and take the Lipschitz results in Proposition B.4 and Proposition B.3 (note that $\|\mathbf{X}_0\|_{\infty} \leq 4 \log^{1/2}(np)$ w.h.p. by Lemma B.6), repeat the argument for other $2n - 1$ symmetric regions, and conclude that w.h.p. the objective value decreases by at least a constant amount. The next proposition summarizes the results.

Proposition II.9. *Assume (II.30). In regions R_{II} and R_{III} , each trust-region step reduces the objective value by at least*

$$d_{\text{II}} = \frac{1}{2} \min \left(\frac{c_*^2 c_a \theta^2 \mu}{n^2 \log(np)}, \frac{3\Delta c_* \theta}{4\pi\sqrt{n}} \right), \quad \text{and} \quad d_{\text{III}} = \frac{1}{2} \min \left(\frac{c_*^3 c_b \theta^3 \mu^4}{n^6 \log^3(np)}, \frac{3\Delta^2 c_* \theta}{8\pi^2 n} \right) \quad (\text{II.31})$$

respectively, provided that

$$\Delta < \frac{c_c c_* \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \quad (\text{II.32})$$

where c_a to c_c are positive constants, and c_* is as defined in Theorem II.1.

Proof. We only consider the symmetric section in the vicinity of e_n and the claims carry on to others by symmetry. If the current iterate $\mathbf{q}^{(r)}$ is in the region R_{II} , by Theorem II.1, w.h.p., we have $\mathbf{w}^* g(\mathbf{w}) / \|\mathbf{w}\| \geq c_* \theta$ for the constant c_* . By Proposition B.4 and Lemma B.6, w.h.p., $\mathbf{w}^* g(\mathbf{w}) / \|\mathbf{w}\|$ is $C_2 n^2 \log(np) / \mu$ -Lipschitz. Therefore, By Lemma II.6 and Lemma II.7, a trust-region step decreases the objective value by at least

$$d_{\text{II}} \doteq \min \left(\frac{c_*^2 \theta^2 \mu}{2C_2 n^2 \log(np)}, \frac{3c_* \theta \Delta}{4\pi\sqrt{n}} \right) - \frac{c_0 n^{3/2} \log^{3/2}(np)}{3\mu^2} \Delta^3.$$

Similarly, if $\mathbf{q}^{(r)}$ is in the region R_{III} , by Proposition B.3, Theorem II.1 and Lemma B.6, w.h.p., $\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} / \|\mathbf{w}\|^2$ is $C_3 n^3 \log^{3/2}(np) / \mu^2$ -Lipschitz and upper bounded by $-c_* \theta$. By Lemma II.6 and Lemma II.8, a trust-region step decreases the objective value by at least

$$d_{\text{III}} \doteq \min \left(\frac{2c_*^3 \theta^3 \mu^4}{3C_3^2 n^6 \log^3(np)}, \frac{3\Delta^2 c_* \theta}{8\pi^2 n} \right) - \frac{c_0 n^{3/2} \log^{3/2}(np)}{3\mu^2} \Delta^3.$$

It can be easily verified that when Δ obeys (II.31), (II.32) holds. \square

The analysis for R_{I} is slightly trickier. In this region, near each local minimizer, the objective function is strongly convex. So we still expect each trust-region step decreases the objective value. On the other hand, it is very unlikely that we can provide a universal lower bound for the amount of decrease - as the iteration sequence approaches a local minimizer, the movement is expected to be diminishing. Nevertheless, close to the minimizer the trust-region algorithm takes “unconstrained” steps. For constrained R_{I} steps, we will again show reduction in objective value by at least a fixed amount; for unconstrained step, we will show the distance between the iterate and the nearest local minimizer drops down rapidly.

The next lemma concerns the function value reduction for constrained R_{I} steps.

Lemma II.10. *Suppose the trust-region size $\Delta \leq 1$, and that at a given iterate r , $\text{Hess } f(\mathbf{q}^{(r)}) \succeq m_H \mathcal{P}_{T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}}$, and $\|\text{Hess } f(\mathbf{q}^{(r)})\| \leq M_H$. Further assume the optimal solution $\delta_* \in T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}$ to the trust-region subproblem (II.21) satisfies $\|\delta_*\| = \Delta$, i.e., the norm constraint is active. Then there exists a tangent vector $\delta \in T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}$ with $\|\delta\| \leq \Delta$, such that*

$$f(\exp_{\mathbf{q}^{(r)}}(\delta)) \leq f(\mathbf{q}^{(r)}) - \frac{m_H^2 \Delta^2}{M_H} + \frac{1}{6} \eta_f \Delta^3,$$

where η_f is defined the same as Lemma II.6.

Proof. See Page 28 under Section VI. \square

The next lemma provides an estimate of m_H . Again we will only state the result for the “canonical” section with the “canonical” $\mathbf{q}(\mathbf{w})$ mapping.

Lemma II.11. *There exists a positive constant C , such that for all $\theta \in (0, 1/2)$ and $\mu < \theta/10$, whenever $p \geq Cn^3 \log \frac{n}{\theta\mu} / (\mu\theta^2)$, it holds with probability at least $1 - cp^{-7}$ that for all \mathbf{q} with $\|\mathbf{w}(\mathbf{q})\| \leq \mu/(4\sqrt{2})$,*

$$\text{Hess } f(\mathbf{q}) \succeq c_* \frac{\theta}{\mu} \mathcal{P}_{T_{\mathbf{q}} \mathbb{S}^{n-1}}.$$

Here c_* is as in Theorem II.1 and Theorem II.2, and $c > 0$ is another constant.

Proof. See Page 29 under Section VI. \square

We know that $\|\mathbf{X}_0\|_\infty \leq 4 \log^{1/2}(np)$ w.h.p., and hence by the definition of Riemannian Hessian and Lemma II.5,

$$M_H \doteq \|\text{Hess } f(\mathbf{q})\| \leq \|\nabla^2 f(\mathbf{q})\| + \|\nabla f(\mathbf{q})\| \leq M_{\nabla^2} + M_{\nabla} \leq \frac{2n}{\mu} \|\mathbf{X}_0\|_\infty^2 \leq \frac{16n}{\mu} \log(np).$$

Combining this estimate and Lemma II.11 and Lemma II.6, we obtain a concrete lower bound for the reduction of objective value for each constrained R_I step.

Proposition II.12. *Assume (II.30). Each constrained R_I trust-region step (i.e., $\|\delta\| = \Delta$) reduces the objective value by at least*

$$d_I = \frac{cc_*^2\theta^2}{\mu n \log(np)} \Delta^2, \quad (\text{II.33})$$

provided

$$\Delta \leq \frac{c'c_*^2\theta^2\mu}{n^{5/2} \log^{5/2}(np)}. \quad (\text{II.34})$$

Here c_* is as in Theorem II.1 and Theorem II.2, and c, c' are positive constants.

Proof. We only consider the symmetric section in the vicinity of e_n and the claims carry on to others by symmetry. We have that w.h.p.

$$\|\text{Hess } f(\mathbf{q})\| \leq \frac{16n}{\mu} \log(np), \quad \text{and} \quad \text{Hess } f(\mathbf{q}) \succeq c_* \frac{\theta}{\mu} \mathcal{P}_{T_q \mathbb{S}^{n-1}}.$$

Combining these estimates with Lemma II.6 and Lemma II.10, one trust-region step will find next iterate $\mathbf{q}^{(r+1)}$ that decreases the objective value by at least

$$d_I \doteq \frac{c_*^2\theta^2/\mu^2}{2n \log(np)/\mu} \Delta^2 - \frac{c_0 n^{3/2} \log^{3/2}(np)}{\mu^2} \Delta^3.$$

Finally, by the condition on Δ in (II.34) and the assumed conditions (II.30), we obtain

$$d_I \geq \frac{c_*^2\theta^2}{2\mu n \log(np)} \Delta^2 - \frac{c_0 n^{3/2} \log^{3/2}(np)}{\mu^2} \Delta^3 \geq \frac{c_*^2\theta^2}{4\mu n \log(np)} \Delta^2,$$

as desired. \square

By the proof strategy for R_I we sketched before Lemma II.10, we expect the iteration sequence ultimately always takes unconstrained steps when it moves very close to a local minimizer. We will show that the following is true: when Δ is small enough, once the iteration sequence starts to take unconstrained R_I step, it will take consecutive unconstrained R_I steps afterwards. It takes two steps to show this: (1) upon an unconstrained R_I step, the next iterate will stay in R_I . It is obvious we can make $\Delta \in O(1)$ to ensure the next iterate stays in $R_I \cup R_{II}$. To strengthen the result, we use the gradient information. From Theorem II.1, we expect the magnitudes of the gradients in R_{II} to be lower bounded; on the other hand, in R_I where points are near local minimizers, continuity argument implies that the magnitudes of gradients should be upper bounded. We will show that when Δ is small enough, there is a gap between these two bounds, implying the next iterate stays in R_I ; (2) when Δ is small enough, the step is in fact unconstrained. Again we will only state the result for the ‘‘canonical’’ section with the ‘‘canonical’’ $\mathbf{q}(\mathbf{w})$ mapping. The next lemma exhibits an absolute lower bound for magnitudes of gradients in R_{II} .

Lemma II.13. *For all \mathbf{q} satisfying $\mu/(4\sqrt{2}) \leq \|\mathbf{w}(\mathbf{q})\| \leq 1/(20\sqrt{5})$, it holds that*

$$\|\text{grad } f(\mathbf{q})\| \geq \frac{9}{10} \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|}.$$

Proof. See Page 33 under Section VI. \square

Assuming (II.30), Theorem II.1 gives that w.h.p. $\mathbf{w}^* \nabla g(\mathbf{w}) / \|\mathbf{w}\| \geq c_* \theta$. Thus, w.h.p., $\|\text{grad } f(\mathbf{q})\| \geq 9c_* \theta / 10$ for all $\mathbf{q} \in R_{II}$. The next lemma compares the magnitudes of gradients before and after taking one unconstrained R_I

step. This is crucial to providing upper bound for magnitude of gradient for the next iterate, and also to establishing the ultimate (quadratic) sequence convergence.

Lemma II.14. *Suppose the trust-region size $\Delta \leq 1$, and at a given iterate r , $\text{Hess } f(\mathbf{q}^{(r)}) \succeq m_H \mathcal{P}_{T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}}$, and that the unique minimizer $\delta_* \in T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}$ to the trust region subproblem (II.21) satisfies $\|\delta_*\| < \Delta$ (i.e., the constraint is inactive). Then, for $\mathbf{q}^{(r+1)} = \exp_{\mathbf{q}^{(r)}}(\delta_*)$, we have*

$$\|\text{grad } f(\mathbf{q}^{(r+1)})\| \leq \frac{L_H}{2m_H^2} \|\text{grad } f(\mathbf{q}^{(r)})\|^2,$$

where $L_H \doteq 5n^{3/2}/(2\mu^2) \|\mathbf{X}_0\|_\infty^3 + 9n/\mu \|\mathbf{X}_0\|_\infty^2 + 9\sqrt{n} \|\mathbf{X}_0\|_\infty$.

Proof. See Page 34 under Section VI. □

We can now bound the Riemannian gradient of the next iterate as

$$\begin{aligned} \|\text{grad } f(\mathbf{q}^{(r+1)})\| &\leq \frac{L_H}{2m_H^2} \|\text{grad } f(\mathbf{q}^{(r)})\|^2 \\ &\leq \frac{L_H}{2m_H^2} \|[U^* \text{Hess } f(\mathbf{q}^{(r)}) U][U^* \text{Hess } f(\mathbf{q}^{(r)}) U]^{-1} \text{grad } f(\mathbf{q}^{(r)})\|^2 \\ &\leq \frac{L_H}{2m_H^2} \|\text{Hess } f(\mathbf{q}^{(r)})\|^2 \Delta^2 = \frac{L_H M_H^2}{2m_H^2} \Delta^2. \end{aligned}$$

Obviously, one can make the upper bound small by tuning down Δ . Combining the above lower bound for $\|\text{grad } f(\mathbf{q})\|$ for $\mathbf{q} \in R_{\text{II}}$, one can conclude that when Δ is small, the next iterate $\mathbf{q}^{(r+1)}$ stays in R_{I} . Another application of the optimality condition (II.6) gives conditions on Δ that guarantees the next trust-region step is also unconstrained. Detailed argument can be found in proof of the following proposition.

Proposition II.15. *Assume (II.30). W.h.p, once the trust-region algorithm takes an unconstrained R_{I} step (i.e., $\|\delta\| < \Delta$), it always takes unconstrained R_{I} steps, provided that*

$$\Delta \leq \frac{cc_*^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)}, \quad (\text{II.35})$$

Here c_* is as in Theorem II.1 and Theorem II.2, and $c > 0$ is another constant.

Proof. We only consider the symmetric section in the vicinity of e_n and the claims carry on to others by symmetry. Suppose that step k is an unconstrained R_{I} step. Then

$$\begin{aligned} \|\mathbf{w}(\mathbf{q}^{(r+1)}) - \mathbf{w}(\mathbf{q}^{(r)})\| &\leq \|\mathbf{q}^{(r+1)} - \mathbf{q}^{(r)}\| = \|\exp_{\mathbf{q}^{(r)}}(\delta) - \mathbf{q}^{(r)}\| \\ &= \sqrt{2 - 2 \cos \|\delta\|} = 2 \sin(\|\delta\|/2) \leq \|\delta\| < \Delta. \end{aligned}$$

Thus, if $\Delta \leq \frac{1}{20\sqrt{5}} - \frac{\mu}{4\sqrt{2}}$, $\mathbf{q}^{(r+1)}$ will be in $R_{\text{I}} \cup R_{\text{II}}$. Next, we show that if Δ is sufficiently small, $\mathbf{q}^{(r+1)}$ will be indeed in R_{I} . By Lemma II.14,

$$\begin{aligned} \|\text{grad } f(\mathbf{q}^{(r+1)})\| &\leq \frac{L_H}{2m_H^2} \|\text{grad } f(\mathbf{q}^{(r)})\|^2 \\ &\leq \frac{L_H M_H^2}{2m_H^2} \|[U^* \text{Hess } f(\mathbf{q}^{(r)}) U]^{-1} U^* \text{grad } f(\mathbf{q}^{(r)})\|^2 \leq \frac{L_H M_H^2}{2m_H^2} \Delta^2, \end{aligned} \quad (\text{II.36})$$

where we have used the fact that

$$\|\delta^{(r)}\| = \|[U^* \text{Hess } f(\mathbf{q}^{(r)}) U]^{-1} U^* \text{grad } f(\mathbf{q}^{(r)})\| < \Delta,$$

as the step is unconstrained. On the other hand, by Theorem II.1 and Lemma II.13, w.h.p.

$$\|\text{grad } f(\mathbf{q})\| \geq \beta_{\text{grad}} \doteq \frac{9}{10} c_* \theta, \quad \forall \mathbf{q} \in R_{\text{II}}. \quad (\text{II.37})$$

Hence, provided

$$\Delta < \frac{m_H}{M_H} \sqrt{\frac{2\beta_{\text{grad}}}{L_H}}, \quad (\text{II.38})$$

we have $\mathbf{q}^{(r+1)} \in R_{\mathcal{I}}$.

We next show that when Δ is small enough, the next step is also unconstrained. Straight forward calculations give

$$\left\| \mathbf{U} \left[\mathbf{U}^* \text{Hess } f \left(\mathbf{q}^{(r+1)} \right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f \left(\mathbf{q}^{(r+1)} \right) \right\| \leq \frac{L_H M_H^2}{2m_H^3} \Delta^2.$$

Hence, provided that

$$\Delta < \frac{2m_H^3}{L_H M_H^2}, \quad (\text{II.39})$$

we will have

$$\left\| \mathbf{U} \left[\mathbf{U}^* \text{Hess } f \left(\mathbf{q}^{(r+1)} \right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f \left(\mathbf{q}^{(r+1)} \right) \right\| < \Delta;$$

in words, the minimizer to the trust-region subproblem for the next step lies in the relative interior of the trust region - the constraint is inactive. By Lemma II.14 and Lemma B.6, we have

$$L_H = C_1 n^{3/2} \log^{3/2}(np) / \mu^2, \quad (\text{II.40})$$

w.h.p.. Combining this and our previous estimates of m_H , M_H , we conclude whenever

$$\Delta \leq \min \left\{ \frac{1}{20\sqrt{5}} - \frac{\mu}{4\sqrt{2}}, \frac{c_1 c_\star^{3/2} \theta^{3/2}}{n^{7/4} \log^{7/4}(np)}, \frac{c_2 \mu c_\star^3 \theta^3}{n^{7/2} \log^{7/2}(np)} \right\},$$

w.h.p., our next trust-region step is also an unconstrained $R_{\mathcal{I}}$ step. Simplifying the above bound completes the proof. \square

Finally, we want to show that ultimate unconstrained $R_{\mathcal{I}}$ iterates actually converges to one nearby local minimizer rapidly. Lemma II.14 has established the gradient is diminishing. The next lemma shows the magnitude of gradient serves as a good proxy for distance to the local minimizer.

Lemma II.16. *Let $\mathbf{q}_\star \in \mathbb{S}^{n-1}$ such that $\text{grad } f(\mathbf{q}_\star) = \mathbf{0}$, and $\boldsymbol{\delta} \in T_{\mathbf{q}_\star} \mathbb{S}^{n-1}$. Consider a geodesic $\gamma(t) = \exp_{\mathbf{q}_\star}(t\boldsymbol{\delta})$, and suppose that on $[0, \tau]$, $\text{Hess } f(\gamma(t)) \succeq m_H \mathcal{P}_{T_{\gamma(t)} \mathbb{S}^{n-1}}$. Then*

$$\|\text{grad } f(\gamma(\tau))\| \geq m_H \tau \|\boldsymbol{\delta}\|.$$

Proof. See Page 34 under Section VI. \square

To see this relates the magnitude of gradient to the distance away from the nearby local minimizer, w.l.o.g., one can assume $\tau = 1$ and consider the point $\mathbf{q} = \exp_{\mathbf{q}_\star}(\boldsymbol{\delta})$. Then

$$\|\mathbf{q}_\star - \mathbf{q}\| = \|\exp_{\mathbf{q}_\star}(\boldsymbol{\delta}) - \mathbf{q}_\star\| = \sqrt{2 - 2 \cos \|\boldsymbol{\delta}\|} = 2 \sin(\|\boldsymbol{\delta}\|/2) \leq \|\boldsymbol{\delta}\| \leq \|\text{grad } f(\mathbf{q})\| / m_H,$$

where at the last inequality above we have used Lemma II.16. Hence, combining this observation with Lemma II.14, we can derive the asymptotic sequence convergence rate as follows.

Proposition II.17. *Assume (II.30) and the conditions in Lemma II.15. Let $\mathbf{q}^{(r_0)} \in R_{\mathcal{I}}$ and the r_0 -th step the first unconstrained $R_{\mathcal{I}}$ step and \mathbf{q}_\star be the unique local minimizer of f over one connected component of $R_{\mathcal{I}}$ that contains $\mathbf{q}^{(r_0)}$. Then w.h.p., for any positive integer $r' \geq 1$,*

$$\left\| \mathbf{q}^{(r_0+r')} - \mathbf{q}_\star \right\| \leq \frac{cc_\star \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2r'}, \quad (\text{II.41})$$

provided that

$$\Delta \leq \frac{c' c_\star^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \quad (\text{II.42})$$

Here c_* is as in Theorem II.1 and Theorem II.2, and c, c' are both positive constants.

Proof. By the geometric characterization in Theorem II.1 and corollary II.2 in [3], f has $2n$ separated local minimizers, each located in R_{I} and within distance $\sqrt{2}\mu/16$ of one of the $2n$ signed basis vectors $\{\pm e_i\}_{i \in [n]}$. Moreover, it is obvious when $\mu \leq 1$, R_{I} consists of $2n$ disjoint connected components. We only consider the symmetric component in the vicinity of e_n and the claims carry on to others by symmetry.

Suppose that r_0 is the index of the first unconstrained iterate in region R_{I} , i.e., $\mathbf{q}^{(r_0)} \in R_{\text{I}}$. By Lemma II.14, for any integer $r' \geq 1$, we have

$$\left\| \text{grad } f \left(\mathbf{q}^{(r_0+r')} \right) \right\| \leq \frac{2m_H^2}{L_H} \left(\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(r_0)} \right) \right\| \right)^{2^{r'}}. \quad (\text{II.43})$$

where L_H is as defined in Lemma II.14, m_H as the strong convexity parameter for R_{I} defined above.

Now suppose \mathbf{q}_* is the unique local minimizer of f , lies in the same R_{I} component that $\mathbf{q}^{(r_0)}$ is located. Let $\gamma_{r'}(t) = \exp_{\mathbf{q}_*}(t\delta)$ to be the unique geodesic that connects \mathbf{q}_* and $\mathbf{q}^{(r_0+r')}$ with $\gamma_{r'}(0) = \mathbf{q}_*$ and $\gamma_{r'}(1) = \mathbf{q}^{(r_0+r')}$. We have

$$\begin{aligned} \left\| \mathbf{q}^{(r_0+r')} - \mathbf{q}_* \right\| &\leq \left\| \exp_{\mathbf{q}_*}(\delta) - \mathbf{q}_* \right\| = \sqrt{2 - 2 \cos \|\delta\|} = 2 \sin(\|\delta\|/2) \\ &\leq \|\delta\| \leq \frac{1}{m_H} \left\| \text{grad } f \left(\mathbf{q}^{(r_0+r')} \right) \right\| \leq \frac{2m_H}{L_H} \left(\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(r_0)} \right) \right\| \right)^{2^{r'}}, \end{aligned}$$

where at the second line we have repeatedly applied Lemma II.16.

By the optimality condition (II.6) and the fact that $\|\delta^{(r_0)}\| < \Delta$, we have

$$\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(r_0)} \right) \right\| \leq \frac{L_H}{2m_H^2} M_H \left\| \left[\mathbf{U}^* \text{Hess } f \left(\mathbf{q}^{(r_0)} \right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f \left(\mathbf{q}^{(r_0)} \right) \right\| \leq \frac{L_H M_H}{2m_H^2} \Delta.$$

Thus, provided

$$\Delta < \frac{m_H^2}{L_H M_H}, \quad (\text{II.44})$$

we can combine the above results and obtain

$$\left\| \mathbf{q}^{(r_0+r')} - \mathbf{q}_* \right\| \leq \frac{2m_H}{L_H} 2^{-2^{r'}}.$$

Based on the previous estimates for m_H, M_H and L_H , we obtain that w.h.p.,

$$\left\| \mathbf{q}^{(r_0+r')} - \mathbf{q}_* \right\| \leq \frac{c_1 c_* \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2^{r'}}.$$

Moreover, by (II.44), w.h.p., it is sufficient to have the trust region size

$$\Delta \leq \frac{c_2 c_*^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}.$$

Thus, we complete the proof. \square

Now we are ready to piece together the above technical proposition to prove Theorem II.3.

Proof. (of Theorem II.3) Assuming (II.30) and in addition that

$$\Delta < \frac{c_1 c_*^3 \theta^3 \mu^2}{n^{7/2} \log^{7/2}(np)},$$

it can be verified that the conditions of all the above propositions are satisfied.

By the preceding four propositions, a step will either be R_{III} , R_{II} , or constrained R_{I} step that decreases the objective value by at least a certain fixed amount (we call this *Type A*), or be an unconstrained R_{I} step (*Type B*), such that all future steps are unconstrained R_{I} and the sequence converges to a local minimizer quadratically. Hence, regardless the initialization, the whole iteration sequence consists of consecutive *Type A* steps, followed by

consecutive Type B steps. Depending on the initialization, either the Type A phase or the Type B phase can be absent. In any case, from $\mathbf{q}^{(0)}$ it takes at most (note $f(\mathbf{q}) \geq 0$ always holds)

$$\frac{f(\mathbf{q}^{(0)})}{\min\{d_I, d_{II}, d_{III}\}} \quad (\text{II.45})$$

steps for the iterate sequence to start take consecutive unconstrained R_I steps, or to already terminate. In case the iterate sequence continues to take consecutive unconstrained R_I steps, Proposition II.17 implies that it takes at most

$$\log \log \left(\frac{c_2 c_* \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right) \quad (\text{II.46})$$

steps to obtain an ε -near solution to the \mathbf{q}_* that is contained in the connected subset of R_I that the sequence entered.

Thus, the number of iterations to obtain an ε -near solution to \mathbf{q}_* can be grossly bounded by

$$\begin{aligned} \#\text{Iter} &\leq \frac{f(\mathbf{q}^{(0)})}{\min\{d_I, d_{II}, d_{III}\}} + \log \log \left(\frac{c_2 c_* \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right) \\ &\leq \left[\min \left\{ \frac{c_3 c_*^3 \theta^3 \mu^4}{n^6 \log^3(np)}, \frac{c_4 c_*^2 \theta^2}{n} \Delta^2 \right\} \right]^{-1} f(\mathbf{q}^{(0)}) + \log \log \left(\frac{c_2 c_* \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right). \end{aligned}$$

Finally, the claimed failure probability comes from a simple union bound with careful bookkeeping. \square

F. Extending to Convergence for Complete Dictionaries

Recall that in this case we consider the preconditioned input

$$\bar{\mathbf{Y}} \doteq \sqrt{p\theta}(\mathbf{Y}\mathbf{Y}^*)^{-1/2}\mathbf{Y}. \quad (\text{II.47})$$

Note that for any complete \mathbf{A}_0 with condition number $\kappa(\mathbf{A}_0)$, from Lemma B.7 we know when p is large enough, w.h.p. one can write the preconditioned $\bar{\mathbf{Y}}$ as

$$\bar{\mathbf{Y}} = \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \mathbf{\Xi}\mathbf{X}_0$$

for a certain $\mathbf{\Xi}$ with small magnitude, and $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$. Particularly, when p is chosen by Theorem II.2, the perturbation is bounded as

$$\|\tilde{\mathbf{\Xi}}\| \leq c c_* \theta \left(\max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1} \quad (\text{II.48})$$

for a certain constant c which can be made arbitrarily small by making the constant C in p large. Since $\mathbf{U}\mathbf{V}^*$ is orthogonal,

$$f(\mathbf{q}; \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \mathbf{\Xi}\mathbf{X}_0) = f(\mathbf{V}\mathbf{U}^*\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0).$$

In words, the function landscape of $f(\mathbf{q}; \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \mathbf{\Xi}\mathbf{X}_0)$ is a rotated version of that of $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0)$. Thus, any local minimizer \mathbf{q}_* of $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0)$ is rotated to $\mathbf{U}\mathbf{V}^*\mathbf{q}_*$, one minimizer of $f(\mathbf{q}; \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \mathbf{\Xi}\mathbf{X}_0)$. Also if our algorithm generates iteration sequence $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0)$ upon initialization \mathbf{q}_0 , it will generate the iteration sequence $\mathbf{U}\mathbf{V}^*\mathbf{q}_0, \mathbf{U}\mathbf{V}^*\mathbf{q}_1, \mathbf{U}\mathbf{V}^*\mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \mathbf{\Xi}\mathbf{X}_0)$. So w.l.o.g. it is adequate that we prove the convergence results for the case $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0)$, corresponding to $\mathbf{A}_0 = \mathbf{I}$ with perturbation $\tilde{\mathbf{\Xi}} \doteq \mathbf{V}\mathbf{U}^*\mathbf{\Xi}$. So in this section (Section II-F), we write $f(\mathbf{q}; \tilde{\mathbf{X}}_0)$ to mean $f(\mathbf{q}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0)$.

Theorem II.2 has shown that when

$$\theta \in \left(0, \frac{1}{2}\right), \mu \leq c \min \left\{ \frac{\theta}{n}, \frac{1}{n^{5/4}} \right\}, p \geq \frac{C}{c_*^2 \theta^2} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta} \right), \quad (\text{II.49})$$

the geometric structure of the landscape is qualitatively unchanged from the orthogonal case, and the parameter c_* constant can be replaced with $c_*/2$. Particularly, for this choice of p , Lemma B.7 implies

$$\|\tilde{\mathbf{\Xi}}\| = \|\mathbf{V}\mathbf{U}^*\mathbf{\Xi}\| \leq c c_* \theta \left(\max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1} \quad (\text{II.50})$$

for a constant c that can be made arbitrarily small by setting the constant C in p sufficiently large. The whole proof is quite similar to that of orthogonal case in the last section. We will only sketch the major changes below. To distinguish with the corresponding quantities in the last section, we use $\tilde{\cdot}$ to denote the corresponding perturbed quantities here.

- Lemma II.5: Note that

$$\|\mathbf{X}_0 + \tilde{\Xi}\mathbf{X}_0\|_\infty \leq \|\mathbf{X}_0\|_\infty + \|\tilde{\Xi}\mathbf{X}_0\|_\infty \leq \|\mathbf{X}_0\|_\infty + \sqrt{n}\|\tilde{\Xi}\|\|\mathbf{X}_0\|_\infty \leq 3\|\mathbf{X}_0\|_\infty/2,$$

where by (II.50) we have used $\|\tilde{\Xi}\| \leq 1/(2\sqrt{n})$ to simplify the above result. So we obtain

$$\tilde{M}_\nabla \leq \frac{3}{2}M_\nabla, \quad \tilde{M}_{\nabla^2} \leq \frac{9}{4}M_{\nabla^2}, \quad \tilde{L}_\nabla \leq \frac{9}{4}L_\nabla, \quad \tilde{L}_{\nabla^2} \leq \frac{27}{8}L_{\nabla^2}.$$

- Lemma II.6: Now we have

$$\tilde{\eta}_f \doteq \tilde{M}_\nabla + 2\tilde{M}_{\nabla^2} + \tilde{L}_\nabla + \tilde{L}_{\nabla^2} \leq 4\eta_f.$$

- Lemma II.7 and Lemma II.8 are generic and nothing changes.
- Proposition II.9: We have now $\mathbf{w}^*g(\mathbf{w}; \tilde{\mathbf{X}}_0)/\|\mathbf{w}\| \geq c_*\theta/2$ by Theorem II.2, w.h.p. $\mathbf{w}^*\nabla g(\mathbf{w}; \tilde{\mathbf{X}}_0)/\|\mathbf{w}\|$ is $C_1n^2 \log(np)/\mu$ -Lipschitz by Proposition B.4, and $\|\mathbf{X}_0 + \tilde{\Xi}\mathbf{X}_0\|_\infty \leq 3\|\mathbf{X}_0\|_\infty/2$ as shown above. Similarly, $\mathbf{w}^*g(\mathbf{w}; \tilde{\mathbf{X}}_0)/\|\mathbf{w}\| \leq -c_*\theta/2$ by Theorem II.2, and $\mathbf{w}^*\nabla^2 g(\mathbf{w}; \tilde{\mathbf{X}}_0)\mathbf{w}/\|\mathbf{w}\|^2$ is $C_2n^3 \log^{3/2}(np)/\mu^2$ -Lipschitz. Moreover, $\tilde{\eta}_f \leq 4\eta_f$ as shown above. Since there are only multiplicative constant changes to the various quantities, we conclude

$$\tilde{d}_{\text{II}} = c_1d_{\text{II}}, \quad \tilde{d}_{\text{III}} = c_1d_{\text{III}} \quad (\text{II.51})$$

provided

$$\Delta < \frac{c_2c_*\theta\mu^2}{n^{5/2} \log^{3/2}(np)}. \quad (\text{II.52})$$

- Lemma II.10: η_f is changed to $\tilde{\eta}_f$ with $\tilde{\eta}_f \leq 4\eta_f$ as shown above.
- Lemma II.11: By (II.3), we have

$$\begin{aligned} \left\| \nabla^2 f(\mathbf{q}; \mathbf{X}_0) - \nabla^2 f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| &\leq \frac{1}{p} \sum_{k=1}^p \left\{ L_{\tilde{h}} \|\tilde{\Xi}\| \|(\mathbf{x}_0)_k\|^2 + \frac{1}{\mu} \left\| (\mathbf{x}_0)_k (\mathbf{x}_0)_k^* - (\tilde{\mathbf{x}}_0)_k (\tilde{\mathbf{x}}_0)_k^* \right\| \right\} \\ &\leq \|\tilde{\Xi}\| \left(L_{\tilde{h}} + 2/\mu + \|\tilde{\Xi}\|/\mu \right) \sum_{k=1}^p \|(\mathbf{x}_0)_k\|^2 \leq \|\tilde{\Xi}\| (L_{\tilde{h}} + 3/\mu) n \|\mathbf{X}_0\|_\infty^2, \end{aligned}$$

where $L_{\tilde{h}}$ is the Lipschitz constant for the function $\tilde{h}_\mu(\cdot)$ and we have used the fact that $\|\tilde{\Xi}\| \leq 1$. Similarly, by II.2,

$$\left\| \nabla f(\mathbf{q}; \mathbf{X}_0) - \nabla f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| \leq \frac{1}{p} \sum_{k=1}^p \left\{ L_{\tilde{h}_\mu} \|\tilde{\Xi}\| \|(\mathbf{x}_0)_k\| + \|\tilde{\Xi}\| \|(\mathbf{x}_0)_k\| \right\} \leq (L_{\tilde{h}_\mu} + 1) \|\tilde{\Xi}\| \sqrt{n} \|\mathbf{X}_0\|_\infty,$$

where $L_{\tilde{h}}$ is the Lipschitz constant for the function $\tilde{h}_\mu(\cdot)$. Since $L_{\tilde{h}} \leq 2/\mu^2$ and $L_{\tilde{h}_\mu} \leq 1/\mu$, and $\|\mathbf{X}_0\|_\infty \leq 4\sqrt{\log(np)}$ w.h.p. (Lemma B.6). By (II.50), w.h.p. we have

$$\left\| \nabla f(\mathbf{q}; \mathbf{X}_0) - \nabla f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| \leq \frac{1}{2}c_*\theta, \quad \text{and} \quad \left\| \nabla^2 f(\mathbf{q}; \mathbf{X}_0) - \nabla^2 f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| \leq \frac{1}{2}c_*\theta,$$

provided the constant C in (II.49) for p is large enough. Thus, by (II.5) and the above estimates we have

$$\begin{aligned} \left\| \text{Hess} f(\mathbf{q}; \mathbf{X}_0) - \text{Hess} f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| &\leq \left\| \nabla f(\mathbf{q}; \mathbf{X}_0) - \nabla f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| + \left\| \nabla^2 f(\mathbf{q}; \mathbf{X}_0) - \nabla^2 f(\mathbf{q}; \tilde{\mathbf{X}}_0) \right\| \\ &\leq c_*\theta \leq \frac{1}{2}c_*\frac{\theta}{\mu}, \end{aligned}$$

provided $\mu \leq 1/2$. So we conclude

$$\text{Hess } f(\mathbf{q}; \widetilde{\mathbf{X}}_0) \succeq \frac{1}{2} c_* \frac{\theta}{\mu} \mathcal{P}_{T_q} \mathbb{S}^{n-1} \implies \widetilde{m}_H \geq \frac{1}{2} c_* \frac{\theta}{\mu}. \quad (\text{II.53})$$

- Proposition II.12: From the estimate of M_H above Proposition II.12 and the last point, we have

$$\left\| \text{Hess } f(\mathbf{q}; \widetilde{\mathbf{X}}_0) \right\| \leq \frac{36}{\mu} \log(np), \quad \text{and} \quad \text{Hess } f(\mathbf{q}; \widetilde{\mathbf{X}}_0) \succeq \frac{1}{2} c_* \frac{\theta}{\mu} \mathcal{P}_{T_q} \mathbb{S}^{n-1}.$$

Also since $\widetilde{\eta}_f \leq 4\eta_f$ in Lemma II.6 and Lemma II.10, there are only multiplicative constant change to the various quantities. We conclude that

$$\widetilde{d}_I = c_3 d_I \quad (\text{II.54})$$

provided that

$$\Delta \leq \frac{c_4 c_*^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \quad (\text{II.55})$$

- Lemma II.13 is generic and nothing changes.
- Lemma II.14: $\widetilde{L}_H \leq 27L_H/8$.
- Proposition II.15: All the quantities involved in determining Δ , m_H , M_H , and L_H , β_{grad} are modified by at most constant multiplicative factors and changed to their respective tilde version, so we conclude that the TRM algorithm always takes unconstrained R_I step after taking one, provided that

$$\Delta \leq \frac{c_5 c_*^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)}. \quad (\text{II.56})$$

- Lemma II.16: is generic and nothing changes.
- Proposition II.17: Again m_H , M_H , L_H are changed to \widetilde{m}_H , \widetilde{M}_H , and \widetilde{L}_H , respectively, differing by at most constant multiplicative factors. So we conclude for any integer $k' \geq 1$,

$$\left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_* \right\| \leq \frac{c_6 c_* \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2k'}, \quad (\text{II.57})$$

provided

$$\Delta \leq \frac{c_7 c_*^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \quad (\text{II.58})$$

The final proof to Theorem II.2 is almost identical to that of Theorem II.1, except that d_I , d_{II} , and d_{III} are changed to \widetilde{d}_I , \widetilde{d}_{II} , and \widetilde{d}_{III} as defined above, respectively. The final iteration complexity to each an ε -near solution is hence

$$\#\text{Iter} \leq \left[\min \left\{ \frac{c_8 c_*^3 \theta^3 \mu^4}{n^6 \log^3(np)}, \frac{c_9 c_*^2 \theta^2}{n} \Delta^2 \right\} \right]^{-1} \left(f(\mathbf{q}^{(0)}) - f(\mathbf{q}_*) \right) + \log \log \left(\frac{c_{10} c_* \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right).$$

Hence overall the qualitative behavior of the algorithm is not changed, as compared to that for the orthogonal case.

III. COMPLETE ALGORITHM PIPELINE AND MAIN RESULTS

For orthogonal dictionaries, from Theorem II.1 (and Corollary II.2 in [3]), we know that all the minimizers $\widehat{\mathbf{q}}_*$ are $O(\mu)$ away from their respective nearest ‘‘target’’ \mathbf{q}_* , with $\mathbf{q}_*^* \widehat{\mathbf{Y}} = \alpha e_i^* \mathbf{X}_0$ for a certain $\alpha \neq 0$ and $i \in [n]$; in Theorem II.3, we have shown that w.h.p. the Riemannian TRM algorithm produces a solution $\widehat{\mathbf{q}} \in \mathbb{S}^{n-1}$ that is ε away to one of the minimizers, say $\widehat{\mathbf{q}}_*$. Thus, the $\widehat{\mathbf{q}}$ returned by the TRM algorithm is $O(\varepsilon + \mu)$ away from \mathbf{q}_* . For exact recovery, we use a simple linear programming rounding procedure, which guarantees to produce the target \mathbf{q}_* . We then use deflation to sequentially recover other rows of \mathbf{X}_0 . Overall, w.h.p. both the dictionary \mathbf{A}_0 and sparse coefficient \mathbf{X}_0 are exactly recovered up to sign permutation, when $\theta \in \Omega(1)$, for orthogonal dictionaries. We summarize relevant technical lemmas and main results in Section III-A. The same procedure can be used to recover complete dictionaries, though the analysis is slightly more complicated; we present the results in Section III-B. Our overall algorithmic pipeline for recovering orthogonal dictionaries is sketched as follows.

- 1) **Estimating one row of \mathbf{X}_0 by the Riemannian TRM algorithm.** By Theorem II.1 (resp. Theorem II.2) and Theorem II.3 (resp. Theorem II.4), starting from any $\mathbf{q} \in \mathbb{S}^{n-1}$, when the relevant parameters are set appropriately (say as μ_* and Δ_*), w.h.p., our Riemannian TRM algorithm finds a local minimizer $\hat{\mathbf{q}}$, with \mathbf{q}_* the nearest target that exactly recovers a row of \mathbf{X}_0 and $\|\hat{\mathbf{q}} - \mathbf{q}_*\| \in O(\mu)$ (by setting the target accuracy of the TRM as, say, $\varepsilon = \mu$).
- 2) **Recovering one row of \mathbf{X}_0 by rounding.** To obtain the target solution \mathbf{q}_* and hence recover (up to scale) one row of \mathbf{X}_0 , we solve the following linear program:

$$\text{minimize}_{\mathbf{q}} \left\| \mathbf{q}^* \hat{\mathbf{Y}} \right\|_1, \quad \text{subject to} \quad \langle \mathbf{r}, \mathbf{q} \rangle = 1, \quad (\text{III.1})$$

with $\mathbf{r} = \hat{\mathbf{q}}$. We show in Lemma III.2 (resp. Lemma III.4) that when $\langle \hat{\mathbf{q}}, \mathbf{q}_* \rangle$ is sufficiently large, implied by μ being sufficiently small, w.h.p. the minimizer of (III.1) is exactly \mathbf{q}_* , and hence one row of \mathbf{X}_0 is recovered by $\mathbf{q}_*^* \hat{\mathbf{Y}}$.

- 3) **Recovering all rows of \mathbf{X}_0 by deflation.** Once ℓ rows of \mathbf{X}_0 ($1 \leq \ell \leq n-2$) have been recovered, say, by unit vectors $\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell$, one takes an orthonormal basis \mathbf{U} for $[\text{span}(\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell)]^\perp$, and minimizes the new function $h(\mathbf{z}) \doteq f(\mathbf{U}\mathbf{z}; \hat{\mathbf{Y}})$ on the sphere $\mathbb{S}^{n-\ell-1}$ with the Riemannian TRM algorithm (though conservative, one can again set parameters as μ_* , Δ_* , as in Step 1) to produce a $\hat{\mathbf{z}}$. Another row of \mathbf{X}_0 is then recovered via the LP rounding (III.1) with input $\mathbf{r} = \mathbf{U}\hat{\mathbf{z}}$ (to produce $\mathbf{q}_*^{\ell+1}$). Finally, by repeating the procedure until depletion, one can recover all the rows of \mathbf{X}_0 .
- 4) **Reconstructing the dictionary \mathbf{A}_0 .** By solving the linear system $\mathbf{Y} = \mathbf{A}\mathbf{X}_0$, one can obtain the dictionary $\mathbf{A}_0 = \mathbf{Y}\mathbf{X}_0^* (\mathbf{X}_0\mathbf{X}_0^*)^{-1}$.

A. Recovering Orthogonal Dictionaries

Theorem III.1 (Main theorem - recovering orthogonal dictionaries). *Assume the dictionary \mathbf{A}_0 is orthogonal and we take $\hat{\mathbf{Y}} = \mathbf{Y}$. Suppose $\theta \in (0, 1/3)$, $\mu_* < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, and $p \geq Cn^3 \log \frac{n}{\mu_* \theta} / (\mu_*^2 \theta^2)$. The above algorithmic pipeline with parameter setting*

$$\Delta_* = \frac{c_b c_*^3 \theta^3 \mu_*^2}{n^{7/2} \log^{7/2}(np)}, \quad (\text{III.2})$$

recovers the dictionary \mathbf{A}_0 and \mathbf{X}_0 in polynomial time, with failure probability bounded by $c_c p^{-6}$. Here c_ is as defined in Theorem II.1, and c_a through c_c , and C are all positive constants.*

Towards a proof of the above theorem, it remains to be shown the correctness of the rounding and deflation procedures.

a) *Proof of LP rounding.*: The following lemma shows w.h.p. the rounding will return the desired \mathbf{q}_* , provided the estimated $\hat{\mathbf{q}}$ is already near to it.

Lemma III.2 (LP rounding - orthogonal dictionary). *For any $\theta \in (0, 1/3)$, whenever $p \geq Cn^2 \log(n/\theta)/\theta$, with probability at least $1 - cp^{-6}$, the rounding procedure (III.1) returns \mathbf{q}_* for any input vector \mathbf{r} that satisfies*

$$\langle \mathbf{r}, \mathbf{q}_* \rangle \geq 249/250.$$

Here C, c are both positive constants.

Proof. See Page 36 under Section VII. □

Since $\langle \hat{\mathbf{q}}, \mathbf{q}_* \rangle = 1 - \|\hat{\mathbf{q}} - \mathbf{q}_*\|^2/2$, and $\|\hat{\mathbf{q}} - \mathbf{q}_*\| \in O(\mu)$, it is sufficient when μ is smaller than some small constant.

b) *Proof sketch of deflation.*: We show the deflation works by induction. To understand the deflation procedure, it is important to keep in mind that the ‘‘target’’ solutions $\{\mathbf{q}_*^i\}_{i=1}^n$ are orthogonal to each other. W.l.o.g., suppose we have found the first ℓ unit vectors $\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell$ which recover the first ℓ rows of \mathbf{X}_0 . Correspondingly, we partition the target dictionary \mathbf{A}_0 and \mathbf{X}_0 as

$$\mathbf{A}_0 = [\mathbf{V}, \mathbf{V}^\perp], \quad \mathbf{X}_0 = \begin{bmatrix} \mathbf{X}_0^{[\ell]} \\ \mathbf{X}_0^{[n-\ell]} \end{bmatrix}, \quad (\text{III.3})$$

where $\mathbf{V} \in \mathbb{R}^{n \times \ell}$, and $\mathbf{X}_0^{[\ell]} \in \mathbb{R}^{\ell \times n}$ denotes the submatrix with the first ℓ rows of \mathbf{X}_0 . Let us define a function: $f_{n-\ell}^\downarrow : \mathbb{R}^{n-\ell} \mapsto \mathbb{R}$ by

$$f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{W}) \doteq \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{z}^* \mathbf{w}_k), \quad (\text{III.4})$$

for any matrix $\mathbf{W} \in \mathbb{R}^{(n-\ell) \times p}$. Then by (I.1), our objective function is equivalent to

$$h(\mathbf{z}) = f(\mathbf{U}\mathbf{z}; \mathbf{A}_0\mathbf{X}_0) = f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*\mathbf{A}_0\mathbf{X}_0) = f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[\ell]} + \mathbf{U}^*\mathbf{V}^\perp\mathbf{X}_0^{[n-\ell]}).$$

Since the columns of the orthogonal matrix $\mathbf{U} \in \mathbb{R}^{n \times (n-\ell)}$ forms the orthogonal complement of $\text{span}(\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell)$, it is obvious that $\mathbf{U}^*\mathbf{V} = \mathbf{0}$. Therefore, we obtain

$$h(\mathbf{z}) = f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*\mathbf{V}^\perp\mathbf{X}_0^{[n-\ell]}).$$

Since $\mathbf{U}^*\mathbf{V}^\perp$ is orthogonal and $\mathbf{X}_0^{[n-\ell]} \sim_{i.i.d.} \text{BG}(\theta)$, this is another instance of orthogonal dictionary learning problem with reduced dimension. If we keep the parameter settings μ_* and Δ_* as Theorem III.1, the conditions of Theorem II.1 and Theorem II.3 for all cases with reduced dimensions are still valid. So w.h.p., the TRM algorithm returns a $\widehat{\mathbf{z}}$ such that $\|\widehat{\mathbf{z}} - \mathbf{z}_*\| \in O(\mu_*)$ where \mathbf{z}_* is a ‘‘target’’ solution that recovers a row of \mathbf{X}_0 :

$$\mathbf{z}_*^* \mathbf{U}^* \mathbf{V}^\perp \mathbf{X}_0^{[n-\ell]} = \mathbf{z}_*^* \mathbf{U}^* \mathbf{A}_0 \mathbf{X}_0 = \alpha \mathbf{e}_i^* \mathbf{X}_0, \quad \text{for some } i \notin [\ell].$$

So pulling everything back in the original space, the effective target is $\mathbf{q}_*^{\ell+1} \doteq \mathbf{U}\mathbf{z}_*$, and $\mathbf{U}\widehat{\mathbf{z}}$ is our estimation obtained from the TRM algorithm. Moreover,

$$\|\mathbf{U}\widehat{\mathbf{z}} - \mathbf{U}\mathbf{z}_*\| = \|\widehat{\mathbf{z}} - \mathbf{z}_*\| \in O(\mu_*).$$

Thus, by Lemma III.2, one successfully recovers $\mathbf{U}\mathbf{z}_*$ from $\mathbf{U}\widehat{\mathbf{z}}$ w.h.p. when μ_* is smaller than a constant. The overall failure probability can be obtained via a simple union bound and simplification of the exponential tails with inverse polynomials in p .

B. Recovering Complete Dictionaries

By working with the preconditioned data samples $\widehat{\mathbf{Y}} = \overline{\mathbf{Y}} \doteq \sqrt{\theta p} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y}$,⁴ we can use the same procedure as described above to recover complete dictionaries.

Theorem III.3 (Main theorem - recovering complete dictionaries). *Assume the dictionary \mathbf{A}_0 is complete with a condition number $\kappa(\mathbf{A}_0)$ and we take $\widehat{\mathbf{Y}} = \overline{\mathbf{Y}}$. Suppose $\theta \in (0, 1/3)$, $\mu_* < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, and $p \geq \frac{C}{c_*^2 \theta^2} \max\left\{\frac{n^4}{\mu^4}, \frac{n^5}{\mu^2}\right\} \kappa^8(\mathbf{A}_0) \log^4\left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta}\right)$. The algorithmic pipeline with parameter setting*

$$\Delta_* = \frac{c_d c_*^3 \theta^3 \mu_*^2}{n^{7/2} \log^{7/2}(np)} \quad (\text{III.5})$$

recovers the dictionary \mathbf{A}_0 and \mathbf{X}_0 in polynomial time, with failure probability bounded by $c_b p^{-6}$. Here c_ is as defined in Theorem II.1, and c_a, c_b are both positive constants.*

Similar to the orthogonal case, we need to show the correctness of the rounding and deflation procedures so that the theorem above holds.

⁴In practice, the parameter θ might not be known beforehand. However, because it only scales the problem, it does not affect the overall qualitative aspect of results.

a) *Proof of LP rounding:* The result of the LP rounding is only slightly different from that of the orthogonal case in Lemma III.2, so is the proof.

Lemma III.4 (LP rounding - complete dictionary). *For any $\theta \in (0, 1/3)$, whenever*

$$p \geq \frac{C}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right),$$

with probability at least $1 - cp^{-6}$, the rounding procedure (III.1) returns \mathbf{q}_* for any input vector \mathbf{r} that satisfies

$$\langle \mathbf{r}, \mathbf{q}_* \rangle \geq 249/250.$$

Here C, c are both positive constants.

Proof. See Page 37 under Section VII. □

b) *Proof sketch of deflation.:* We use a similar induction argument to show the deflation works. Compared to the orthogonal case, the tricky part here is that the target vectors $\{\mathbf{q}_*^i\}_{i=1}^n$ are not necessarily orthogonal to each other, but they are almost so. W.l.o.g., let us again assume that $\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell$ recover the first ℓ rows of \mathbf{X}_0 , and similarly partition the matrix \mathbf{X}_0 as in (III.3).

By Lemma B.7 and (II.48), we can write $\bar{\mathbf{Y}} = (\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0$ for some orthogonal matrix \mathbf{Q} and small perturbation $\mathbf{\Xi}$ with $\|\mathbf{\Xi}\| \leq \delta < 1/10$ for some large p as usual. Similar to the orthogonal case, we have

$$h(\mathbf{z}) = f(\mathbf{U}\mathbf{z}; (\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0) = f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0),$$

where $f_{n-\ell}^\downarrow$ is defined the same as in (III.4). Next, we show that the matrix $\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0$ can be decomposed as $\mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}$, where $\mathbf{V} \in \mathbb{R}^{(n-\ell) \times n}$ is orthogonal and $\mathbf{\Delta}$ is a small perturbation matrix. More specifically, we show that

Lemma III.5. *Suppose the matrices $\mathbf{U} \in \mathbb{R}^{n \times (n-\ell)}$, $\mathbf{Q} \in \mathbb{R}^{n \times n}$ are orthogonal as defined above, $\mathbf{\Xi}$ is a perturbation matrix with $\|\mathbf{\Xi}\| \leq 1/20$, then*

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}, \quad (\text{III.6})$$

where $\mathbf{V} \in \mathbb{R}^{n \times (n-\ell)}$ is an orthogonal matrix spanning the same subspace as that of \mathbf{U} , and the norms of $\mathbf{\Delta}$ is bounded by

$$\|\mathbf{\Delta}\|_{\ell^1 \rightarrow \ell^2} \leq 16\sqrt{n}\|\mathbf{\Xi}\|\|\mathbf{X}_0\|_\infty, \quad \|\mathbf{\Delta}\| \leq 16\|\mathbf{\Xi}\|\|\mathbf{X}_0\|, \quad (\text{III.7})$$

where $\|\mathbf{W}\|_{\ell^1 \rightarrow \ell^2} = \sup_{\|\mathbf{z}\|_1=1} \|\mathbf{W}\mathbf{z}\| = \max_k \|\mathbf{w}_k\|$ denotes the max column ℓ^2 -norm of a matrix \mathbf{W} .

Proof. See Page 38 under Section VII. □

Since $\mathbf{U}\mathbf{V}$ is orthogonal and $\mathbf{X}_0^{[n-\ell]} \sim_{i.i.d.} \text{BG}(\theta)$, we come into another instance of perturbed dictionary learning problem with a reduced dimension

$$h(\mathbf{z}) = f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}).$$

Since our perturbation analysis in proving Theorem II.2 and Theorem II.4 solely relies on the fact that $\|\mathbf{\Delta}\|_{\ell^1 \rightarrow \ell^2} \leq C\|\mathbf{\Xi}\|\sqrt{n}\|\mathbf{X}_0\|_\infty$, it is enough to make p large enough so that the theorems are still applicable for the reduced version $f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta})$. Thus, by invoking Theorem II.2 and Theorem II.4, the TRM algorithm provably returns one $\hat{\mathbf{z}}$ such that $\hat{\mathbf{z}}$ is near to a perturbed optimal $\hat{\mathbf{z}}_*$ with

$$\hat{\mathbf{z}}_*^* \mathbf{U}^* \mathbf{V} \mathbf{X}_0^{[n-\ell]} = \mathbf{z}_*^* \mathbf{U}^* \mathbf{V} \mathbf{X}_0^{[n-\ell]} + \mathbf{z}_*^* \mathbf{\Delta} = \alpha \mathbf{e}_i^* \mathbf{X}_0, \quad \text{for some } i \notin [\ell], \quad (\text{III.8})$$

where \mathbf{z}_* with $\|\mathbf{z}_*\| = 1$ is the exact solution. More specifically, Corollary II.4 in [3] implies that

$$\|\hat{\mathbf{z}} - \hat{\mathbf{z}}_*\| \leq \sqrt{2}\mu_*/7.$$

Next, we show that \widehat{z}_* is also very near to the exact solution z_* . Indeed, the identity (III.8) suggests

$$\begin{aligned} (\widehat{z}_* - z_*)^* U^* V X_0^{[n-\ell]} &= z_*^* \Delta \\ \implies \widehat{z}_* - z_* &= \left[(X_0^{[n-\ell]})^* V^* U \right]^\dagger \Delta^* z_* = U^* V \left[(X_0^{[n-\ell]})^* \right]^\dagger \Delta^* z_* \end{aligned} \quad (\text{III.9})$$

where $W^\dagger = (W^* W)^{-1} W^*$ denotes the pseudo inverse of a matrix W with full column rank. Hence, by (III.9) we can bound the distance between \widehat{z}_* and z_* by

$$\|\widehat{z}_* - z_*\| \leq \left\| \left[(X_0^{[n-\ell]})^* \right]^\dagger \right\| \|\Delta\| \leq \sigma_{\min}^{-1}(X_0^{[n-\ell]}) \|\Delta\|$$

By Lemma B.1, when $p \geq \Omega(n^2 \log n)$, w.h.p.,

$$\theta p/2 \leq \sigma_{\min}(X_0^{[n-\ell]}(X_0^{[n-\ell]})^*) \leq \left\| X_0^{[n-\ell]}(X_0^{[n-\ell]})^* \right\| \leq \|X_0 X_0^*\| \leq 3\theta p/2.$$

Hence, combined with Lemma III.5, we obtain

$$\sigma_{\min}^{-1}(X_0^{[n-\ell]}) \leq \sqrt{\frac{2}{\theta p}}, \quad \|\Delta\| \leq 28\sqrt{\theta p} \|\Xi\| / \sqrt{2},$$

which implies that $\|\widehat{z}_* - z_*\| \leq 28 \|\Xi\|$. Thus, combining the results above, we obtain

$$\|\widehat{z} - z_*\| \leq \|\widehat{z} - \widehat{z}_*\| + \|\widehat{z}_* - z_*\| \leq \sqrt{2}\mu_*/7 + 28 \|\Xi\|.$$

Lemma B.7, and in particular (II.48), for our choice of p as in Theorem II.2, $\|\Xi\| \leq c\mu_*^2 n^{-3/2}$, where c can be made smaller by making the constant in p larger. For μ_* sufficiently small, we conclude that

$$\|U\widehat{z} - Uz_*\| = \|\widehat{z} - z_*\| \leq 2\mu_*/7.$$

In words, the TRM algorithm returns a \widehat{z} such that $U\widehat{z}$ is very near to one of the unit vectors $\{q_*^i\}_{i=1}^n$, such that $(q_*^i)^* \bar{Y} = \alpha e_i^* X_0$ for some $\alpha \neq 0$. For μ_* smaller than a fixed constant, one will have

$$\langle U\widehat{z}, q_*^i \rangle \geq 249/250,$$

and hence by Lemma III.4, the LP rounding exactly returns the optimal solution q_*^i upon the input $U\widehat{z}$.

The proof sketch above explains why the recursive TRM plus rounding works. The overall failure probability can be obtained via a simple union bound and simplifications of the exponential tails with inverse polynomials in p .

IV. SIMULATIONS

A. Practical TRM Implementation

Fixing a small step size and solving the trust-region subproblem exactly eases the analysis, but also renders the TRM algorithm impractical. In practice, the trust-region subproblem is never exactly solved, and the trust-region step size is adjusted to the local geometry, say by backtracking. It is possible to modify our algorithmic analysis to account for inexact subproblem solvers and adaptive step size; for sake of brevity, we do not pursue it here. Recent theoretical results on the practical version include [72], [73].

Here we describe a practical implementation based on the *Manopt* toolbox [70]⁵. *Manopt* is a user-friendly Matlab toolbox that implements several sophisticated solvers for tackling optimization problems over Riemannian manifolds. The most developed solver is based on the TRM. This solver uses the truncated conjugate gradient (tCG; see, e.g., Section 7.5.4 of [57]) method to (approximately) solve the trust-region subproblem (vs. the exact solver in our analysis). It also dynamically adjusts the step size using backtracking. However, the original implementation (*Manopt* 2.0) is not adequate for our purposes. Their tCG solver uses the gradient as the initial search direction, which does not ensure that the TRM solver can escape from saddle points [68], [71]. We modify the tCG solver, such that when the current gradient is small and there is a negative curvature direction (i.e., the current point is near a saddle point or a local maximizer of $f(q)$), the tCG solver explicitly uses the negative curvature direction⁶

⁵Available online: <http://www.manopt.org>.

⁶...adjusted in sign to ensure positive correlation with the gradient – if it does not vanish.




Fig. 3: Phase transition for recovering a single sparse vector. **Left:** We fix $p = 5n^2 \log n$ and vary the dimension n and sparsity level k ; **Right:** We fix the sparsity level as $\lceil 0.2 \cdot n \rceil$ and vary the dimension n and number of samples p . For each configuration, the experiment is independently repeated for five times. White indicates success, and black indicates failure.

as the initial search direction. This modification ensures the TRM solver always escape from saddle points/local maximizers with negative directional curvature. Hence, the modified TRM algorithm based on Manopt is expected to have the same qualitative behavior as the idealized version we analyzed above, with better scalability. We will perform our numerical simulations using the modified TRM algorithm whenever necessary. Algorithm 3 together with Lemmas 9 and 10 and the surrounding discussion in the very recent work [72] provides a detailed description of this practical version.

B. Simulated Data

To corroborate our theory, we experiment with dictionary recovery on simulated data.⁷ For simplicity, we focus on recovering orthogonal dictionaries and we declare success once a single row of the coefficient matrix is recovered.

Since the problem is invariant to rotations, w.l.o.g. we set the dictionary as $A_0 = I \in \mathbb{R}^{n \times n}$. For any fixed sparsity k , each column of the coefficient matrix $X_0 \in \mathbb{R}^{n \times p}$ has exactly k nonzero entries, chosen uniformly random from $\binom{[n]}{k}$. These nonzero entries are i.i.d. standard normals. This is slightly different from the Bernoulli-Gaussian model we assumed for analysis. For n reasonably large, these two models have similar behaviors. For our sparsity surrogate, we fix the smoothing parameter as $\mu = 10^{-2}$. Because the target points are the signed basis vector $\pm e_i$'s (to recover rows of X_0), for a solution \hat{q} returned by the TRM algorithm, we define the reconstruction error (RE) to be

$$\text{RE} = \min_{i \in [n]} (\|\hat{q} - e_i\|, \|\hat{q} + e_i\|). \quad (\text{IV.1})$$

One trial is determined to be a success once $\text{RE} \leq \mu$, with the idea that this indicates \hat{q} is already very near the target and the target can likely be recovered via the LP rounding we described (which we do not implement here).

We consider two settings: (1) fix $p = 5n^2 \log n$ and vary the dimension n and sparsity k ; (2) fix the sparsity level as $\lceil 0.2 \cdot n \rceil$ and vary the dimension n and number of samples p . For each pair of (k, n) for (1), and each pair of (p, n) for (2), we repeat the simulations independently for $T = 5$ times. Fig. 3 shows the phase transition for the two settings. It seems that our TRM algorithm can work well into the linear region whenever $p \in O(n^2 \log n)$ (Fig. 3-Left), but p should have order greater than $\Omega(n)$ (Fig. 3-Right). The sample complexity from our theory is significantly suboptimal compared to this.

C. Image Data Again

Our algorithmic framework has been derived based on the BG model on the coefficients. Real data may not admit sparse representations w.r.t. complete dictionaries, or even so, the coefficients may not obey the BG model. In this

⁷The code is available online: https://github.com/sunju/dl_focm

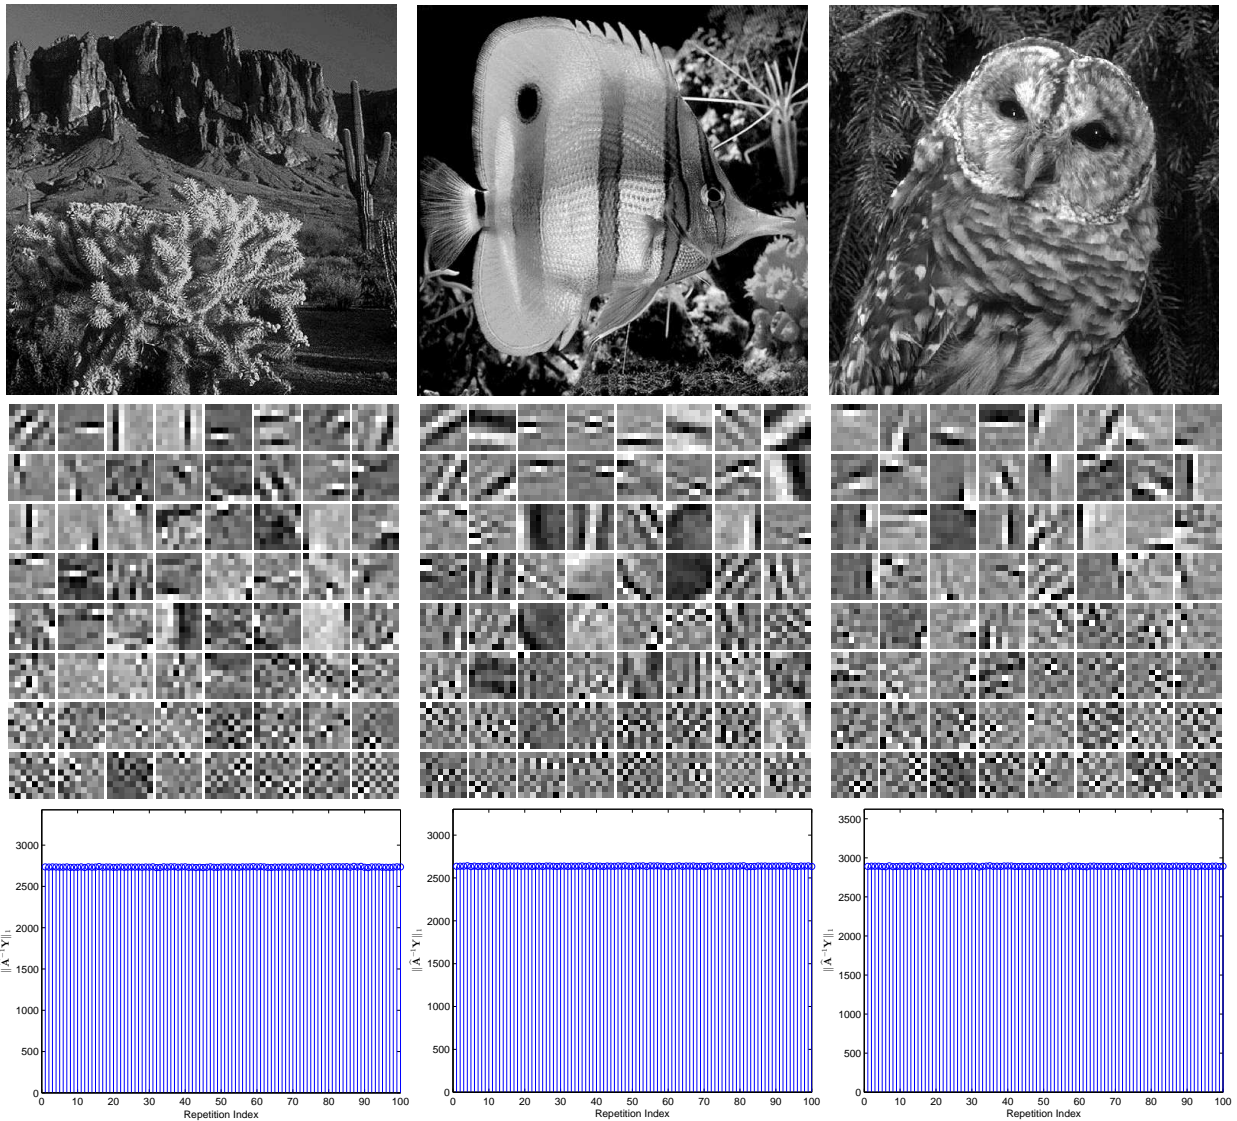


Fig. 4: Results of learning complete dictionaries from image patches, using the algorithmic pipeline in Section III. **Top:** Images we used for the experiment. These are the three images we used in our motivational experiment (Section I.B) in the companion paper [3]. The way we formed the data matrix \mathbf{Y} is exactly the same as in that experiment. **Middle:** The 64 dictionary elements we learned. **Bottom:** Let $\hat{\mathbf{A}}$ be the final dictionary matrix at convergence. This row shows the value $\|\hat{\mathbf{A}}^{-1}\mathbf{Y}\|_1$ across one hundred independent runs. The values are almost the same, with a relative difference less than 10^{-3} .

experiment, we explore how our algorithm performs in learning complete dictionaries for image patches, emulating our motivational experiment in the companion paper [3] (Section I.B). Thanks to research on image compression, we know patches of natural images tend to admit sparse representation, even w.r.t. simple orthogonal bases, such as Fourier basis or wavelets.

We take the three images that we used in the motivational experiment. For each image, we divide it into 8×8 non-overlapping patches, vectorize the patches, and then stack the vectorized patches into a data matrix \mathbf{Y} . \mathbf{Y} is preconditioned as

$$\bar{\mathbf{Y}} = (\mathbf{Y}\mathbf{Y}^\top)^{-1/2} \mathbf{Y},$$

and the resulting $\bar{\mathbf{Y}}$ is fed to the dictionary learning pipeline described in Section III. The smoothing parameter μ is fixed to 10^{-2} . Fig. 4 contains the learned dictionaries: the dictionaries generally contain localized, directional

features that resemble subset of wavelets and generalizations. These are very reasonable representing elements for natural images. Thus, the BG coefficient model may be a sensible, simple model for natural images.

Another piece of strong evidence in support of the above claim is as follows. For each image, we repeat the learning pipeline for one hundred times, with independent initializations across the runs. Let $\hat{\mathbf{A}}$ be the final learned dictionary for each run, we plot the value of $\|\hat{\mathbf{A}}^{-1}\mathbf{Y}\|_1$ across the one hundred independent runs. Strikingly, the values are virtually the same, with a relative difference of 10^{-3} ! This is predicted by our theory, under the BG model. If the model is unreasonable for natural images, the preconditioning, benign function landscape, LP rounding, and the deflation process that hinge on this model would have completely fallen down.

For this image experiment, $n = 64$ and $p = 4096$. A single run of the learning pipeline, including solving 64 instances of the optimization over the sphere (with varying dimensions) and solving 64 instances of the LP rounding (using CVX), lasts about 20 minutes on a mid-range modern laptop. So with careful implementation we discussed above, the learning pipeline is actually not far from practical.

V. DISCUSSION

For recovery of complete dictionaries, the LP program approach in [81] that works with $\theta \leq O(1/\sqrt{n})$ only demands $p \geq \Omega(n^2 \log^2 n)$. The sample complexity has recently been improved to $p \geq \Omega(n \log n)$ [82], matching the information-theoretic lower bound $\Omega(n \log n)$ (i.e., when $\theta \sim 1/n$; see also [83]), albeit still in the $\theta \in O(1/\sqrt{n})$ regime. The sample complexity of our method working in the $\theta \sim \Theta(1)$ regime as stated in Theorem III.3 is obviously much higher. As already discussed in [3], working with other ℓ^1 proxies, directly in the \mathbf{q} space would likely save the sample complexity both for geometric characterization and algorithm analysis. Another possibility is to analyze the complete case directly, instead of treating it as transformed version of perturbed orthogonal case.

Our experiments seem to suggest the necessary complexity level lies between $\Omega(n)$ and $O(n^2 \log n)$ even for the orthogonal case. While it is interesting to determine the true complexity requirement for the TRM, there could be other efficient algorithms that demand less. For example, simulations in [44] seem to suggest $O(n \log n)$ samples suffice to guarantee efficient recovery. The simulations run an alternating direction algorithm fed with problem-specific initializations, nevertheless.

Our analysis is based on exact trust-region subproblem solver and fixed step size. The convergence result for the practical version from [72], based on approximate solver and adaptive step size, is general, but pessimistic. It seems not difficult to adapt their analysis according to our objective geometry, and obtain a tight, practical convergence result.

Our motivating experiment on real images in introduction of our companion paper [3] remains mysterious. If we were to believe that real image data are “nice” and our objective there does not have spurious local minima either, it is surprising ADM would escape all other critical points – this is not predicted by classic or modern theories. One reasonable place to start is to look at how gradient descent algorithms with generic initializations can escape from ridable saddle points (at least with high probability). The recent work [74] has showed that randomly perturbing each iterate can help gradient algorithm to achieve this with high probability.

VI. PROOF OF CONVERGENCE FOR THE TRUST-REGION ALGORITHM

A. Proof of Lemma II.5

Proof. Using the fact $\tanh(\cdot)$ and $1 - \tanh^2(\cdot)$ are bounded by one in magnitude, by (II.2) and (II.3) we have

$$\begin{aligned} \|\nabla f(\mathbf{q})\| &\leq \frac{1}{p} \sum_{k=1}^p \|(\mathbf{x}_0)_k\| \leq \sqrt{n} \|\mathbf{X}_0\|_\infty, \\ \|\nabla^2 f(\mathbf{q})\| &\leq \frac{1}{p} \sum_{k=1}^p \frac{1}{\mu} \|(\mathbf{x}_0)_k\|^2 \leq \frac{n}{\mu} \|\mathbf{X}_0\|_\infty^2, \end{aligned}$$

for any $\mathbf{q} \in \mathbb{S}^{n-1}$. Moreover,

$$\begin{aligned} \sup_{\mathbf{q}, \mathbf{q}' \in \mathbb{S}^{n-1}, \mathbf{q} \neq \mathbf{q}'} \frac{\|\nabla f(\mathbf{q}) - \nabla f(\mathbf{q}')\|}{\|\mathbf{q} - \mathbf{q}'\|} &\leq \frac{1}{p} \sum_{k=1}^p \|(\mathbf{x}_0)_k\| \sup_{\mathbf{q}, \mathbf{q}' \in \mathbb{S}^{n-1}, \mathbf{q} \neq \mathbf{q}'} \frac{\left| \tanh\left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu}\right) - \tanh\left(\frac{\mathbf{q}'^*(\mathbf{x}_0)_k}{\mu}\right) \right|}{\|\mathbf{q} - \mathbf{q}'\|} \\ &\leq \frac{1}{p} \sum_{k=1}^p \|(\mathbf{x}_0)_k\| \frac{\|(\mathbf{x}_0)_k\|}{\mu} \leq \frac{n}{\mu} \|\mathbf{X}_0\|_\infty^2, \end{aligned}$$

where at the last line we have used the fact the mapping $\mathbf{q} \mapsto \mathbf{q}^*(\mathbf{x}_0)_k/\mu$ is $\|(\mathbf{x}_0)_k\|/\mu$ Lipschitz, and $x \mapsto \tanh(x)$ is 1-Lipschitz, and the composition rule of Lipschitz functions (i.e., Lemma V.5 of [3]). Similar argument yields the final bound. \square

B. Proof of Lemma II.6

Proof. Suppose we can establish

$$\left| f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) - \widehat{f}(\mathbf{q}, \boldsymbol{\delta}) \right| \leq \frac{1}{6} \eta_f \|\boldsymbol{\delta}\|^3.$$

Applying this twice we obtain

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta}_*)) \leq \widehat{f}(\mathbf{q}, \boldsymbol{\delta}_*) + \frac{1}{6} \eta_f \Delta^3 \leq \widehat{f}(\mathbf{q}, \boldsymbol{\delta}) + \frac{1}{6} \eta_f \Delta^3 \leq f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) + \frac{1}{3} \eta_f \Delta^3 \leq f(\mathbf{q}) - s + \frac{1}{3} \eta_f \Delta^3,$$

as claimed. Next we establish the first result. Let $\boldsymbol{\delta}_0 = \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|}$, and $t = \|\boldsymbol{\delta}\|$. Consider the composite function

$$\zeta(t) \doteq f(\exp_{\mathbf{q}}(t\boldsymbol{\delta}_0)) = f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)),$$

and also

$$\begin{aligned} \dot{\zeta}(t) &= \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) \rangle \\ \ddot{\zeta}(t) &= \langle \nabla^2 f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)) (-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t)), -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) \rangle \\ &\quad + \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), -\mathbf{q} \cos(t) - \boldsymbol{\delta}_0 \sin(t) \rangle. \end{aligned}$$

In particular, this gives that

$$\begin{aligned} \zeta(0) &= f(\mathbf{q}) \\ \dot{\zeta}(0) &= \langle \boldsymbol{\delta}_0, \nabla f(\mathbf{q}) \rangle \\ \ddot{\zeta}(0) &= \boldsymbol{\delta}_0^* (\nabla^2 f(\mathbf{q}) - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \mathbf{I}) \boldsymbol{\delta}_0. \end{aligned}$$

We next develop a bound on $|\ddot{\zeta}(t) - \ddot{\zeta}(0)|$. Using the triangle inequality, we can casually bound this difference as

$$\begin{aligned} & \left| \ddot{\zeta}(t) - \ddot{\zeta}(0) \right| \\ & \leq \left| \langle \nabla^2 f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)) (-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t)), -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) \rangle - \boldsymbol{\delta}_0^* \nabla^2 f(\mathbf{q}) \boldsymbol{\delta}_0 \right| \\ & \quad + \left| \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), -\mathbf{q} \cos(t) - \boldsymbol{\delta}_0 \sin(t) \rangle + \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \right| \\ & \leq \left| \langle [\nabla^2 f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)) - \nabla^2 f(\mathbf{q})] (-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t)), -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) \rangle \right| \\ & \quad + \left| \langle \nabla^2 f(\mathbf{q}) (-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) - \boldsymbol{\delta}_0), -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) \rangle \right| \\ & \quad + \left| \langle \nabla^2 f(\mathbf{q}) \boldsymbol{\delta}_0, -\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) - \boldsymbol{\delta}_0 \rangle \right| \\ & \quad + \left| \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), -\mathbf{q} \cos(t) - \boldsymbol{\delta}_0 \sin(t) \rangle + \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), \mathbf{q} \rangle \right| \\ & \quad + \left| \langle \nabla f(\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t)), \mathbf{q} \rangle - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \right| \\ & \leq L_{\nabla^2} \|\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t) - \mathbf{q}\| \\ & \quad + M_{\nabla^2} \|-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) - \boldsymbol{\delta}_0\| \\ & \quad + M_{\nabla^2} \|-\mathbf{q} \sin(t) + \boldsymbol{\delta}_0 \cos(t) - \boldsymbol{\delta}_0\| \\ & \quad + M_{\nabla} \|-\mathbf{q} \cos(t) - \boldsymbol{\delta}_0 \sin(t) + \mathbf{q}\| \end{aligned}$$

$$\begin{aligned}
 & + L_{\nabla} \|\mathbf{q} \cos(t) + \boldsymbol{\delta}_0 \sin(t) - \mathbf{q}\| \\
 & = (L_{\nabla^2} + 2M_{\nabla^2} + M_{\nabla} + L_{\nabla}) \sqrt{(1 - \cos(t))^2 + \sin^2(t)} \\
 & = \eta_f \sqrt{2 - 2 \cos t} \leq \eta_f \sqrt{4 \sin^2(t/2)} \leq \eta_f t,
 \end{aligned}$$

where in the final line we have used the fact $1 - \cos x = 2 \sin^2(x/2)$ and that $\sin x \leq x$ for $x \in [0, 1]$, and M_{∇} , M_{∇^2} , L_{∇} and L_{∇^2} are the quantities defined in Lemma II.5. By the integral form of Taylor's theorem in Lemma A.7 and the result above, we have

$$\begin{aligned}
 \left| f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) - \widehat{f}(\mathbf{q}, \boldsymbol{\delta}) \right| & = \left| \zeta(t) - \left(\zeta(0) + t\dot{\zeta}(0) + \frac{t^2}{2}\ddot{\zeta}(0) \right) \right| \\
 & = \left| t^2 \int_0^1 (1-s) \ddot{\zeta}(st) ds - \frac{t^2}{2} \ddot{\zeta}(0) \right| \\
 & = t^2 \left| \int_0^1 (1-s) \left[\ddot{\zeta}(st) - \ddot{\zeta}(0) \right] ds \right| \\
 & \leq t^2 \int_0^1 (1-s) st \eta_f ds = \frac{\eta_f t^3}{6},
 \end{aligned}$$

with $t = \|\boldsymbol{\delta}\|$ we obtain the desired result. \square

C. Proof of Lemma II.7

Proof. By the integral form of Taylor's theorem in Lemma A.7, for any $t \in \left[0, \frac{3\Delta}{2\pi\sqrt{n}}\right]$, we have

$$\begin{aligned}
 & g\left(\mathbf{w} - t \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \\
 & = g(\mathbf{w}) - t \int_0^1 \left\langle \nabla g\left(\mathbf{w} - st \frac{\mathbf{w}}{\|\mathbf{w}\|}\right), \frac{\mathbf{w}}{\|\mathbf{w}\|} \right\rangle ds \\
 & = g(\mathbf{w}) - t \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} + t \int_0^1 \left\langle \nabla g(\mathbf{w}) - \nabla g\left(\mathbf{w} - st \frac{\mathbf{w}}{\|\mathbf{w}\|}\right), \frac{\mathbf{w}}{\|\mathbf{w}\|} \right\rangle ds \\
 & = g(\mathbf{w}) - t \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} + t \int_0^1 \left(\left\langle \nabla g(\mathbf{w}), \frac{\mathbf{w}}{\|\mathbf{w}\|} \right\rangle - \left\langle \nabla g\left(\mathbf{w} - st \frac{\mathbf{w}}{\|\mathbf{w}\|}\right), \frac{\mathbf{w} - st\mathbf{w}/\|\mathbf{w}\|}{\|\mathbf{w} - st\mathbf{w}/\|\mathbf{w}\|} \right\rangle \right) ds \\
 & \leq g(\mathbf{w}) - t \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} + \frac{L_g}{2} t^2 \leq g(\mathbf{w}) - t\beta_g + \frac{L_g}{2} t^2.
 \end{aligned}$$

Minimizing this function over $t \in \left[0, \frac{3\Delta}{2\pi\sqrt{n}}\right]$, we obtain that there exists a $\mathbf{w}' \in \mathcal{B}\left(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}}\right)$ such that

$$g(\mathbf{w}') \leq g(\mathbf{w}) - \min \left\{ \frac{\beta_g^2}{2L_g}, \frac{3\beta_g\Delta}{4\pi\sqrt{n}} \right\}.$$

Given such a $\mathbf{w}' \in \mathcal{B}\left(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}}\right)$, there must exist some $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ such that $\mathbf{q}(\mathbf{w}') = \exp_{\mathbf{q}}(\boldsymbol{\delta})$. It remains to show that $\|\boldsymbol{\delta}\| \leq \Delta$. It is easy to verify that $\|\mathbf{q}(\mathbf{w}') - \mathbf{q}(\mathbf{w})\| \leq 2\sqrt{n}\|\mathbf{w}' - \mathbf{w}\| \leq 3\Delta/\pi$. Hence,

$$\|\exp_{\mathbf{q}}(\boldsymbol{\delta}) - \mathbf{q}\|^2 = \left\| \mathbf{q}(1 - \cos\|\boldsymbol{\delta}\|) + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \sin\|\boldsymbol{\delta}\| \right\|^2 = 2 - 2\cos\|\boldsymbol{\delta}\| = 4\sin^2\frac{\|\boldsymbol{\delta}\|}{2} \leq \frac{9\Delta^2}{\pi^2},$$

which means that $\sin(\|\boldsymbol{\delta}\|/2) \leq 3\Delta/(2\pi)$. Because $\sin x \geq \frac{3}{\pi}x$ over $x \in [0, \pi/6]$, it implies that $\|\boldsymbol{\delta}\| \leq \Delta$. Since $g(\mathbf{w}) = f(\mathbf{q}(\mathbf{w}))$, by summarizing all the results, we conclude that there exists a $\boldsymbol{\delta}$ with $\|\boldsymbol{\delta}\| \leq \Delta$, such that

$$f(\exp_{\mathbf{q}}(\boldsymbol{\delta})) \leq f(\mathbf{q}) - \min \left\{ \frac{\beta_g^2}{2L_g}, \frac{3\beta_g\Delta}{4\pi\sqrt{n}} \right\},$$

as claimed. \square

D. Proof of Lemma II.8

Proof. Let $\sigma = \text{sign}(\mathbf{w}^* \nabla g(\mathbf{w}))$. For any $t \in \left[0, \frac{\Delta}{2\sqrt{n}}\right]$, by integral form of Taylor's theorem in Lemma A.7, we have

$$\begin{aligned}
& g\left(\mathbf{w} - t\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \\
&= g(\mathbf{w}) - t\sigma \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} + t^2 \int_0^1 (1-s) \frac{\mathbf{w}^* \nabla^2 g\left(\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \mathbf{w}}{\|\mathbf{w}\|^2} ds \\
&\leq g(\mathbf{w}) + \frac{t^2}{2} \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} + t^2 \int_0^1 \left[(1-s) \frac{\mathbf{w}^* \nabla^2 g\left(\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \mathbf{w}}{\|\mathbf{w}\|^2} - (1-s) \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \right] ds \\
&= g(\mathbf{w}) + \frac{t^2}{2} \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \\
&\quad + t^2 \int_0^1 (1-s) \left[\frac{\left(\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right)^* \nabla^2 g\left(\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \left(\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right)}{\left\|\mathbf{w} - st\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right\|^2} - \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \right] ds \\
&\leq g(\mathbf{w}) - \frac{t^2}{2} \beta_\wedge + t^2 \int_0^1 (1-s) s L_\wedge t ds \leq g(\mathbf{w}) - \frac{t^2}{2} \beta_\wedge + \frac{t^3}{6} L_\wedge.
\end{aligned}$$

Minimizing this function over $t \in \left[0, \frac{3\Delta}{2\pi\sqrt{n}}\right]$, we obtain

$$t_\star = \min \left\{ \frac{2\beta_\wedge}{L_\wedge}, \frac{3\Delta}{2\pi\sqrt{n}} \right\},$$

and there exists a $\mathbf{w}' = \mathbf{w} - t_\star \sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}$ such that

$$g\left(\mathbf{w} - t_\star \sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right) \leq g(\mathbf{w}) - \min \left\{ \frac{2\beta_\wedge^3}{3L_\wedge^2}, \frac{3\Delta^2 \beta_\wedge}{8\pi^2 n} \right\}.$$

By arguments identical to those used in Lemma II.7, there exists a tangent vector $\boldsymbol{\delta} \in T_{\mathbf{q}} \mathbb{S}^{n-1}$ such that $\mathbf{q}(\mathbf{w}') = \exp_{\mathbf{q}}(\boldsymbol{\delta})$ and $\|\boldsymbol{\delta}\| \leq \Delta$. This completes the proof. \square

E. Proof of Lemma II.10

Proof. For any $t \in \left[0, \frac{\Delta}{\|\text{grad } f(\mathbf{q}^{(r)})\|}\right]$, it holds that $\|t \text{ grad } f(\mathbf{q}^{(r)})\| \leq \Delta$, and the quadratic approximation

$$\begin{aligned}
\widehat{f}\left(\mathbf{q}^{(r)}, -t \text{ grad } f\left(\mathbf{q}^{(r)}\right)\right) &\leq f\left(\mathbf{q}^{(r)}\right) - t \left\| \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\|^2 + \frac{M_H}{2} t^2 \left\| \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\|^2 \\
&= f\left(\mathbf{q}^{(r)}\right) - t \left(1 - \frac{1}{2} M_H t\right) \left\| \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\|^2.
\end{aligned}$$

Taking $t_0 = \min \left\{ \frac{\Delta}{\|\text{grad } f(\mathbf{q}^{(r)})\|}, \frac{1}{M_H} \right\}$, we obtain

$$\widehat{f}\left(\mathbf{q}^{(r)}, -t_0 \text{ grad } f\left(\mathbf{q}^{(r)}\right)\right) \leq f\left(\mathbf{q}^{(r)}\right) - \frac{1}{2} \min \left\{ \frac{\Delta}{\|\text{grad } f(\mathbf{q}^{(r)})\|}, \frac{1}{M_H} \right\} \left\| \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\|^2. \quad (\text{VI.1})$$

Now let \mathbf{U} be an arbitrary orthonormal basis for $T_{\mathbf{q}^{(r)}} \mathbb{S}^{n-1}$. Since the norm constraint is active, by the optimality condition in (II.6), we have

$$\begin{aligned}
\Delta &\leq \left\| \left[\mathbf{U}^* \text{Hess } f\left(\mathbf{q}^{(r)}\right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\| \\
&\leq \left\| \left[\mathbf{U}^* \text{Hess } f\left(\mathbf{q}^{(r)}\right) \mathbf{U} \right]^{-1} \right\| \left\| \mathbf{U}^* \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\| \leq \frac{\left\| \text{grad } f\left(\mathbf{q}^{(r)}\right) \right\|}{m_H},
\end{aligned}$$

which means that $\|\text{grad } f(\mathbf{q}^{(r)})\| \geq m_H \Delta$. Substituting this into (VI.1), we obtain

$$\widehat{f}\left(\mathbf{q}^{(r)}, -t_0 \text{grad } f\left(\mathbf{q}^{(r)}\right)\right) \leq f\left(\mathbf{q}^{(r)}\right) - \frac{1}{2} \min\left\{m_H \Delta^2, \frac{m_H^2}{M_H} \Delta^2\right\} \leq f\left(\mathbf{q}^{(r)}\right) - \frac{m_H^2 \Delta^2}{2M_H}.$$

By the key comparison result established in proof of Lemma II.6, we have

$$\begin{aligned} f\left(\exp_{\mathbf{q}^{(r)}}\left(-t_0 \text{grad } f\left(\mathbf{q}^{(r)}\right)\right)\right) &\leq \widehat{f}\left(\mathbf{q}^{(r)}, -t_0 \text{grad } f\left(\mathbf{q}^{(r)}\right)\right) + \frac{1}{6} \eta_f \Delta^3 \\ &\leq f\left(\mathbf{q}^{(r)}\right) - \frac{m_H^2 \Delta^2}{M_H} + \frac{1}{6} \eta_f \Delta^3. \end{aligned}$$

This completes the proof. \square

F. Proof of Lemma II.11

It takes certain delicate work to prove Lemma II.11. Basically to use discretization argument, the degree of continuity of the Hessian is needed. The tricky part is that for continuity, we need to compare the Hessian operators at different points, while these Hessian operators are only defined on the respective tangent planes. This is the place where parallel translation comes into play. The next two lemmas compute spectral bounds for the forward and inverse parallel translation operators.

Lemma VI.1. For $\tau \in [0, 1]$ and $\|\boldsymbol{\delta}\| \leq 1/2$, we have

$$\|\mathcal{P}_\gamma^{\tau \leftarrow 0} - \mathbf{I}\| \leq \frac{5}{4} \tau \|\boldsymbol{\delta}\|, \quad (\text{VI.2})$$

$$\|\mathcal{P}_\gamma^{0 \leftarrow \tau} - \mathbf{I}\| \leq \frac{3}{2} \tau \|\boldsymbol{\delta}\|. \quad (\text{VI.3})$$

Proof. By (II.7), we have

$$\begin{aligned} \|\mathcal{P}_\gamma^{\tau \leftarrow 0} - \mathbf{I}\| &= \left\| (\cos(\tau \|\boldsymbol{\delta}\|) - 1) \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \sin(\tau \|\boldsymbol{\delta}\|) \frac{\mathbf{q} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} \right\| \\ &\leq 1 - \cos(\tau \|\boldsymbol{\delta}\|) + \sin(\tau \|\boldsymbol{\delta}\|) \\ &\leq 2 \sin^2\left(\frac{\tau \|\boldsymbol{\delta}\|}{2}\right) + \sin(\tau \|\boldsymbol{\delta}\|) \leq \frac{1}{4} \tau \|\boldsymbol{\delta}\| + \tau \|\boldsymbol{\delta}\| \leq \frac{5}{4} \tau \|\boldsymbol{\delta}\|, \end{aligned}$$

where we have used the fact $\sin(t) \leq t$ and $1 - \cos x = 2 \sin^2(x/2)$. Moreover, $\mathcal{P}_\gamma^{0 \leftarrow \tau}$ is in the form of $(\mathbf{I} + \mathbf{u} \mathbf{v}^*)^{-1}$ for some vectors \mathbf{u} and \mathbf{v} . By the Sherman-Morrison matrix inverse formula, i.e., $(\mathbf{I} + \mathbf{u} \mathbf{v}^*)^{-1} = \mathbf{I} - \mathbf{u} \mathbf{v}^* / (1 + \mathbf{v}^* \mathbf{u})$ (justified as $\left\| (\cos(\tau \|\boldsymbol{\delta}\|) - 1) \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \mathbf{q} \sin(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} \right\| \leq 5\tau \|\boldsymbol{\delta}\| / 4 \leq 5/8 < 1$ as shown above), we have

$$\begin{aligned} &\|\mathcal{P}_\gamma^{0 \leftarrow \tau} - \mathbf{I}\| \\ &= \left\| (\cos(\tau \|\boldsymbol{\delta}\|) - 1) \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \mathbf{q} \sin(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} \right\| \frac{1}{1 + (\cos(\tau \|\boldsymbol{\delta}\|) - 1)} \quad (\text{as } \mathbf{q}^* \boldsymbol{\delta} = 0) \\ &\leq \frac{5}{4} \tau \|\boldsymbol{\delta}\| \frac{1}{\cos(\tau \|\boldsymbol{\delta}\|)} \leq \frac{5}{4} \tau \|\boldsymbol{\delta}\| \frac{1}{\cos(1/2)} \leq \frac{3}{2} \tau \|\boldsymbol{\delta}\|, \end{aligned}$$

completing the proof. \square

The next lemma establishes the ‘‘local-Lipschitz’’ property of the Riemannian Hessian.

Lemma VI.2. Let $\gamma(t) = \exp_{\mathbf{q}}(t\boldsymbol{\delta})$ denotes a geodesic curve on \mathbb{S}^{n-1} . Whenever $\|\boldsymbol{\delta}\| \leq 1/2$ and $\tau \in [0, 1]$,

$$\left\| \mathcal{P}_\gamma^{0 \leftarrow \tau} \text{Hess } f(\gamma(\tau)) \mathcal{P}_\gamma^{\tau \leftarrow 0} - \text{Hess } f(\mathbf{q}) \right\| \leq L_H \cdot \tau \|\boldsymbol{\delta}\|, \quad (\text{VI.4})$$

where $L_H = 5n^{3/2} \|\mathbf{X}_0\|_\infty^3 / (2\mu^2) + \frac{9}{\mu} n \|\mathbf{X}_0\|_\infty^2 + 9\sqrt{n} \|\mathbf{X}_0\|_\infty$.

Proof. First of all, by (II.5) and using the fact that the operator norm of a projection operator is unitary bounded, we have

$$\begin{aligned}
 & \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\| \\
 & \leq \left\| \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} [\nabla^2 f(\gamma(\tau)) - \nabla^2 f(\mathbf{q}) - (\langle \nabla f(\gamma(\tau)), \gamma(\tau) \rangle - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle) \mathbf{I}] \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} \right\| \\
 & \quad + \left\| \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} (\nabla^2 f(\mathbf{q}) - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \mathbf{I}) \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} \right. \\
 & \quad \left. - \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} (\nabla^2 f(\mathbf{q}) - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \mathbf{I}) \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \right\| \\
 & \leq \|\nabla^2 f(\gamma(\tau)) - \nabla^2 f(\mathbf{q})\| + |\langle \nabla f(\gamma(\tau)), \gamma(\tau) \rangle - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle| \\
 & \quad + \left\| \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} - \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \right\| \left\| \mathcal{P}_{T_{\gamma(\tau)}\mathbb{S}^{n-1}} + \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \right\| \|\nabla^2 f(\mathbf{q}) - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle \mathbf{I}\|.
 \end{aligned}$$

By the estimates in Lemma II.5, we obtain

$$\begin{aligned}
 & \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\| \\
 & \leq \frac{2}{\mu^2} n^{3/2} \|\mathbf{X}_0\|_\infty^3 \|\gamma(\tau) - \mathbf{q}\| + \frac{n}{\mu} \|\mathbf{X}_0\|_\infty^2 \|\gamma(\tau) - \mathbf{q}\| + \sqrt{n} \|\mathbf{X}_0\|_\infty \|\gamma(\tau) - \mathbf{q}\| \\
 & \quad + 2 \|\gamma(\tau) \gamma^*(\tau) - \mathbf{q} \mathbf{q}^*\| \left(\frac{n}{\mu} \|\mathbf{X}_0\|_\infty^2 + \sqrt{n} \|\mathbf{X}_0\|_\infty \right) \\
 & \leq \left(\frac{5}{2\mu^2} n^{3/2} \|\mathbf{X}_0\|_\infty^3 + \frac{25n}{4\mu} \|\mathbf{X}_0\|_\infty^2 + \frac{25}{4} \sqrt{n} \|\mathbf{X}_0\|_\infty \right) \tau \|\boldsymbol{\delta}\|, \tag{VI.5}
 \end{aligned}$$

where at the last line we have used the following estimates:

$$\begin{aligned}
 \|\gamma(\tau) - \mathbf{q}\| & = \left\| \mathbf{q} (\cos(\tau \|\boldsymbol{\delta}\|) - 1) + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \sin(\tau \|\boldsymbol{\delta}\|) \right\| \leq \frac{5}{4} \tau \|\boldsymbol{\delta}\|, \text{ (Proof of Lemma VI.1)} \\
 \|\gamma(\tau) \gamma^*(\tau) - \mathbf{q} \mathbf{q}^*\| & \leq \left\| \left(\frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \mathbf{q} \mathbf{q}^* \right) \sin^2(\tau \|\boldsymbol{\delta}\|) \right\| + 2 \sin(\tau \|\boldsymbol{\delta}\|) \cos(\tau \|\boldsymbol{\delta}\|) \\
 & \leq \sin^2(\tau \|\boldsymbol{\delta}\|) + \sin(2\tau \|\boldsymbol{\delta}\|) \leq \frac{5}{2} \tau \|\boldsymbol{\delta}\|.
 \end{aligned}$$

Therefore, by Lemma VI.1, we obtain

$$\begin{aligned}
 & \left\| \mathcal{P}_\gamma^{0 \leftarrow \tau} \text{Hess } f(\gamma(\tau)) \mathcal{P}_\gamma^{\tau \leftarrow 0} - \text{Hess } f(\mathbf{q}) \right\| \\
 & \leq \left\| \mathcal{P}_\gamma^{0 \leftarrow \tau} \text{Hess } f(\gamma(\tau)) \mathcal{P}_\gamma^{\tau \leftarrow 0} - \text{Hess } f(\gamma(\tau)) \mathcal{P}_\gamma^{\tau \leftarrow 0} \right\| + \left\| \text{Hess } f(\gamma(\tau)) \mathcal{P}_\gamma^{\tau \leftarrow 0} - \text{Hess } f(\gamma(\tau)) \right\| \\
 & \quad + \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\| \\
 & \leq \left\| \mathcal{P}_\gamma^{0 \leftarrow \tau} - \mathbf{I} \right\| \|\text{Hess } f(\gamma(\tau))\| + \left\| \mathcal{P}_\gamma^{\tau \leftarrow 0} - \mathbf{I} \right\| \|\text{Hess } f(\gamma(\tau))\| + \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\| \\
 & \leq \frac{11}{4} \tau \|\boldsymbol{\delta}\| \|\nabla^2 f(\gamma(\tau)) - \langle \nabla f(\gamma(\tau)), \gamma(\tau) \rangle \mathbf{I}\| + \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\|.
 \end{aligned}$$

By Lemma II.5 and substituting the estimate in (VI.5), we obtain the claimed result. \square

Proof. (of Lemma II.11) For any given \mathbf{q} with $\|\mathbf{w}(\mathbf{q})\| \leq \mu/(4\sqrt{2})$, assume \mathbf{U} is an orthonormal basis for its tangent space $T_{\mathbf{q}}\mathbb{S}^{n-1}$. Again we first work with the ‘‘canonical’’ section in the vicinity of \mathbf{e}_n with the ‘‘canonical’’ reparametrization $\mathbf{q}(\mathbf{w}) = [\mathbf{w}; \sqrt{1 - \|\mathbf{w}\|^2}]$.

- 1) **Expectation of the operator.** By definition of the Riemannian Hessian in (II.5), expressions of $\nabla^2 f$ and ∇f in (II.2) and (II.3), and exchange of differential and expectation operators, we obtain

$$\begin{aligned}
 \mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} & = \mathbb{E}[\mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}] \\
 & = \mathbb{E}[\mathbf{U}^* \nabla^2 f(\mathbf{q}) \mathbf{U} - \langle \mathbf{q}, \nabla f(\mathbf{q}) \rangle \mathbf{I}_{n-1}] \\
 & = \mathbf{U}^* \mathbb{E} \left[\frac{1}{\mu} \left\{ 1 - \tanh^2 \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \right\} \mathbf{x} \mathbf{x}^* \right] \mathbf{U} - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \mathbf{q}^* \mathbf{x} \right] \mathbf{I}_{n-1}.
 \end{aligned}$$

Let $\bar{\mathbf{x}} \in \mathbb{R}^{n-1}$ be the first $n-1$ coordinates of \mathbf{x} and $\mathbf{w} \in \mathbb{R}^{n-1}$ the similar subvector of \mathbf{q} (as used in [3]). We have

$$\mathbf{U}^* \mathbb{E} \left[\frac{1}{\mu} \left\{ 1 - \tanh^2 \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \right\} \mathbf{x} \mathbf{x}^* \right] \mathbf{U} \succeq \frac{1-\theta}{\mu} \mathbf{U}^* \mathbb{E} \left[\left\{ 1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) \right\} \begin{bmatrix} \bar{\mathbf{x}} \bar{\mathbf{x}}^* & \mathbf{0} \\ \mathbf{0}^* & 0 \end{bmatrix} \right] \mathbf{U}.$$

Now consider any vector $\mathbf{z} \in T_{\mathbf{q}} \mathbb{S}^{n-1}$ such that $\mathbf{z} = \mathbf{U} \mathbf{v}$ for some $\mathbf{v} \in \mathbb{R}^{n-1}$ and $\|\mathbf{z}\| = 1$. Then

$$\mathbf{z}^* \mathbb{E} \left[\left\{ 1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) \right\} \begin{bmatrix} \bar{\mathbf{x}} \bar{\mathbf{x}}^* & \mathbf{0} \\ \mathbf{0}^* & 0 \end{bmatrix} \right] \mathbf{z} \geq \frac{\theta}{\sqrt{2\pi}} \left(2 - \frac{3}{4}\sqrt{2} \right) \|\bar{\mathbf{z}}\|^2$$

by proof of Proposition II.7 in [3], where $\bar{\mathbf{z}} \in \mathbb{R}^{n-1}$ as above is the first $n-1$ coordinates of \mathbf{z} . Now we know that $\langle \mathbf{q}, \mathbf{z} \rangle = 0$, or

$$\mathbf{w}^* \bar{\mathbf{z}} + q_n z_n = 0 \implies \frac{\|\bar{\mathbf{z}}\|}{|z_n|} = \frac{q_n}{\|\mathbf{w}\|} = \frac{\sqrt{1 - \|\mathbf{w}\|^2}}{\|\mathbf{w}\|} \geq 50,$$

where we have used $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$ and $\mu \leq 1/10$ to obtain the last lower bound. Combining the above with the fact that $\|\mathbf{z}\| = 1$, we obtain

$$\begin{aligned} \mathbf{U}^* \mathbb{E} \left[\frac{1}{\mu} \left\{ 1 - \tanh^2 \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \right\} \mathbf{x} \mathbf{x}^* \right] \mathbf{U} &\succeq \frac{99}{100} \frac{1-\theta}{\mu} \frac{\theta}{\sqrt{2\pi}} (2 - 3\sqrt{2}/4) \mathbf{I}_{n-1} \\ &\succeq \frac{99}{200\sqrt{2\pi}} (2 - 3\sqrt{2}/4) \frac{\theta}{\mu} \mathbf{I}_{n-1}, \end{aligned} \quad (\text{VI.6})$$

where we have simplified the expression using $\theta \leq 1/2$. To bound the second term,

$$\begin{aligned} &\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \mathbf{q}^* \mathbf{x} \right] \\ &= \mathbb{E}_{\mathcal{I}} \left[\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [\tanh(Z/\mu) Z] \right] \\ &= \frac{1}{\mu} \mathbb{E}_{\mathcal{I}} \left[\|\mathbf{q}_{\mathcal{I}}\|^2 \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [1 - \tanh^2(Z/\mu)] \right] \quad (\text{by Lemma B.1 in [3]}) \\ &\leq \frac{1}{\mu} \mathbb{E}_{\mathcal{I}} \left[\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [1 - \tanh^2(Z/\mu)] \right]. \end{aligned}$$

Now we have the following estimate:

$$\begin{aligned} &\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [1 - \tanh^2(Z/\mu)] \\ &= 2 \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [(1 - \tanh^2(Z/\mu)) \mathbb{1}_{Z>0}] \\ &\leq 8 \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [\exp(-2Z/\mu) \mathbb{1}_{Z>0}] \\ &= 8 \exp \left(\frac{2 \|\mathbf{w}_{\mathcal{J}}\|^2 + 2q_n^2}{\mu^2} \right) \Phi^c \left(\frac{2\sqrt{\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2}}{\mu} \right) \quad (\text{by Lemma B.1 in [3]}) \\ &\leq \frac{4}{\sqrt{2\pi}} \frac{\mu}{\sqrt{\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2}}, \end{aligned}$$

where at the last inequality we have applied Gaussian tail upper bound of Type II in Lemma A.2. Since $\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2 \geq q_n^2 = 1 - \|\mathbf{w}\|^2 \geq 1 - \mu^2/32 \geq 31/32$ for $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$ and $\mu \leq 1$, we obtain

$$\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [1 - \tanh^2(Z/\mu)] \leq \frac{4}{\sqrt{2\pi}} \frac{\mu}{\sqrt{31/32}} \leq \frac{4}{\sqrt{2\pi}} \mu. \quad (\text{VI.7})$$

Collecting the above estimates, we obtain

$$\mathbf{U}^* \text{Hess} \mathbb{E} [f(\mathbf{q})] \mathbf{U} \succeq \frac{99}{200\sqrt{2\pi}} (2 - 3\sqrt{2}/4) \frac{\theta}{\mu} \mathbf{I}_{n-1} - \frac{1}{\mu} \frac{4}{\sqrt{2\pi}} \mu \mathbf{I}_{n-1} \succeq \frac{1}{4\sqrt{2\pi}} \frac{\theta}{\mu} \mathbf{I}_{n-1}, \quad (\text{VI.8})$$

where we have used the fact $\mu \leq \theta/10$ to obtain the final lower bound.

2) **Concentration.** Next we perform concentration analysis. For any \mathbf{q} , we can write

$$\mathbf{U}^* \nabla^2 f(\mathbf{q}) \mathbf{U} = \frac{1}{p} \sum_{k=1}^p \mathbf{W}_k, \quad \text{with } \mathbf{W}_k \doteq \frac{1}{\mu} \left[1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu} \right) \right] \mathbf{U}^* (\mathbf{x}_0)_k (\mathbf{x}_0)_k^* \mathbf{U}.$$

For any integer $m \geq 2$, we have

$$\begin{aligned} \mathbf{0} &\preceq \mathbb{E} [\mathbf{W}_k^m] \preceq \frac{1}{\mu^m} \mathbb{E} [(\mathbf{U}^* (\mathbf{x}_0)_k (\mathbf{x}_0)_k^* \mathbf{U})^m] \\ &\preceq \frac{1}{\mu^m} \mathbb{E} [\|(\mathbf{x}_0)_k (\mathbf{x}_0)_k^*\|^m] \mathbf{I} = \frac{1}{\mu^m} \mathbb{E} [\|(\mathbf{x}_0)_k\|^{2m}] \mathbf{I} \preceq \frac{1}{\mu^m} \mathbb{E}_{Z \sim \chi^2(n)} [Z^m] \mathbf{I}, \end{aligned}$$

where we have used Lemma B.8 to obtain the last inequality. By Lemma A.4, we obtain

$$\mathbf{0} \preceq \mathbb{E} [\mathbf{W}_k^m] \preceq \frac{1}{\mu^m} \frac{m!}{2} (2n)^m \mathbf{I} \preceq \frac{m!}{2} \left(\frac{2n}{\mu} \right)^m \mathbf{I}.$$

Taking $R_{\mathbf{W}} = 2n/\mu$, and $\sigma_{\mathbf{W}}^2 = 4n^2/\mu^2 \geq \mathbb{E} [\mathbf{W}_k^2]$, by Lemma A.6, we obtain

$$\mathbb{P} \left[\left\| \frac{1}{p} \sum_{k=1}^p \mathbf{W}_k - \frac{1}{p} \sum_{k=1}^p \mathbb{E} [\mathbf{W}_k] \right\| > t/2 \right] \leq 2n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right) \quad (\text{VI.9})$$

for any $t > 0$. Similarly, we write

$$\langle \nabla f(\mathbf{q}), \mathbf{q} \rangle = \frac{1}{p} \sum_{k=1}^p Z_k, \quad \text{with } Z_k \doteq \tanh \left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu} \right) \mathbf{q}^*(\mathbf{x}_0)_k.$$

For any integer $m \geq 2$, we have

$$\mathbb{E} [|Z_k|^m] \leq \mathbb{E} [|\mathbf{q}^*(\mathbf{x}_0)_k|^m] \leq \mathbb{E}_{Z \sim \mathcal{N}(0,1)} [|Z|^m] \leq \frac{m!}{2},$$

where at the first inequality we used the fact $|\tanh(\cdot)| \leq 1$, at the second we invoked Lemma B.8, and at the third we invoked Lemma A.3. Taking $R_Z = \sigma_Z^2 = 1$, by Lemma A.5, we obtain

$$\mathbb{P} \left[\left| \frac{1}{p} \sum_{k=1}^p Z_k - \frac{1}{p} \sum_{k=1}^p \mathbb{E} [Z_k] \right| > t/2 \right] \leq 2 \exp(-pt^2/16) \quad (\text{VI.10})$$

for any $t > 0$. Gathering (VI.9) and (VI.10), we obtain that for any $t > 0$,

$$\begin{aligned} &\mathbb{P} [\|\mathbf{U}^* \text{Hess} \mathbb{E} [f(\mathbf{q})] \mathbf{U} - \mathbf{U}^* \text{Hess} f(\mathbf{q}) \mathbf{U}\| > t] \\ &\leq \mathbb{P} [\|\mathbf{U}^* \nabla^2 f(\mathbf{q}) \mathbf{U} - \nabla^2 \mathbb{E} [f(\mathbf{q})]\| > t/2] + \mathbb{P} [|\langle \nabla f(\mathbf{q}), \mathbf{q} \rangle - \langle \nabla \mathbb{E} [f(\mathbf{q})], \mathbf{q} \rangle| > t/2] \\ &\leq 2n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right) + 2 \exp \left(-\frac{pt^2}{16} \right) \leq 4n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right). \end{aligned} \quad (\text{VI.11})$$

3) **Uniformizing the bound.** Now we are ready to pull above results together for a discretization argument. For any $\varepsilon \in (0, \mu/(4\sqrt{2}))$, there is an ε -net N_ε of size at most $(3\mu/(4\sqrt{2}\varepsilon))^n$ that covers the region $\{\mathbf{q} : \|\mathbf{w}(\mathbf{q})\| \leq \mu/(4\sqrt{2})\}$. By Lemma VI.2, the function $\text{Hess} f(\mathbf{q})$ is locally Lipschitz within each normal ball of radius

$$\|\mathbf{q} - \exp_{\mathbf{q}}(1/2)\| = \sqrt{2 - 2\cos(1/2)} \geq 1/\sqrt{5}$$

with Lipschitz constant L_H (as defined in Lemma VI.2). Note that $\varepsilon < \mu/(4\sqrt{2}) < 1/(4\sqrt{2}) < 1/\sqrt{5}$ for $\mu < 1$, so any choice of $\varepsilon \in (0, \mu/(4\sqrt{2}))$ makes the Lipschitz constant L_H valid within each ε -ball centered around one element of the ε -net. Let

$$\mathcal{E}_\infty \doteq \left\{ 1 \leq \|\mathbf{X}_0\|_\infty \leq 4\sqrt{\log(np)} \right\}.$$

From Lemma B.6, $\mathbb{P}[\mathcal{E}_\infty^c] \leq \theta (np)^{-7} + \exp(-0.3\theta np)$. By Lemma VI.2, with at least the same probability,

$$L_H \leq C_1 \frac{n^{3/2}}{\mu^2} \log^{3/2}(np).$$

Set $\varepsilon = \frac{\theta}{12\sqrt{2\pi}\mu L_H} < \mu/(4\sqrt{2})$, so

$$\#N_\varepsilon \leq \exp\left(n \log \frac{C_2 n^{3/2} \log^{3/2}(np)}{\theta}\right).$$

Let \mathcal{E}_H denote the event that

$$\mathcal{E}_H \doteq \left\{ \max_{\mathbf{q} \in N_\varepsilon} \|\mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} - \mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}\| \leq \frac{\theta}{12\sqrt{2\pi}\mu} \right\}.$$

On $\mathcal{E}_\infty \cap \mathcal{E}_H$,

$$\sup_{\mathbf{q}: \|\mathbf{w}(\mathbf{q})\| \leq \mu/(4\sqrt{2})} \|\mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} - \mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}\| \leq \frac{\theta}{6\sqrt{2\pi}\mu}.$$

So on $\mathcal{E}_\infty \cap \mathcal{E}_H$, we have

$$\mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U} \succeq c_\# \frac{\theta}{\mu} \tag{VI.12}$$

for any $c_\# \leq 1/(12\sqrt{2\pi})$. We take $c_\# = c_\star$ for simplicity. Setting $t = \frac{\theta}{12\sqrt{2\pi}\mu}$ in (VI.11), we obtain that for any fixed \mathbf{q} in this region,

$$\mathbb{P}[\|\mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} - \mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}\| > t] \leq 4n \exp\left(-\frac{p\theta^2}{c_3 n^2 + c_4 n\theta/\mu}\right).$$

Taking a union bound, we obtain that

$$\mathbb{P}[\mathcal{E}_H^c] \leq 4n \exp\left(-\frac{p\theta^2}{c_3 n^2 + c_4 n\theta/\mu} + C_5 n \log n + C_6 n \log \log p\right).$$

It is enough to make $p \geq C_7 n^3 \log(n/(\mu\theta))/(\mu\theta^2)$ to make the failure probability small, completing the proof. \square

G. Proof of Lemma II.13

Proof. For a given \mathbf{q} , consider the vector $\mathbf{r} \doteq \mathbf{q} - \mathbf{e}_n/q_n$. It is easy to verify that $\langle \mathbf{q}, \mathbf{r} \rangle = 0$, and hence $\mathbf{r} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$. Now, by (II.2) and (II.4), we have

$$\begin{aligned} \langle \text{grad } f(\mathbf{q}), \mathbf{r} \rangle &= \langle (\mathbf{I} - \mathbf{q}\mathbf{q}^*) \nabla f(\mathbf{q}), \mathbf{q} - \mathbf{e}_n/q_n \rangle \\ &= \langle (\mathbf{I} - \mathbf{q}\mathbf{q}^*) \nabla f(\mathbf{q}), -\mathbf{e}_n/q_n \rangle \\ &= \frac{1}{p} \sum_{k=1}^p \left\langle (\mathbf{I} - \mathbf{q}\mathbf{q}^*) \tanh\left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu}\right) (\mathbf{x}_0)_k, -\mathbf{e}_n/q_n \right\rangle \\ &= \frac{1}{p} \sum_{k=1}^p \tanh\left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu}\right) \left(-\frac{x_k(n)}{q_n} + \mathbf{q}^*(\mathbf{x}_0)_k\right) \\ &= \frac{1}{p} \sum_{k=1}^p \tanh\left(\frac{\mathbf{q}^*(\mathbf{x}_0)_k}{\mu}\right) \left(\mathbf{w}^*(\mathbf{q}) \bar{x}_k - \frac{x_k(n)}{q_n} \|\mathbf{w}(\mathbf{q})\|^2\right) \\ &= \mathbf{w}^*(\mathbf{q}) \nabla g(\mathbf{w}), \end{aligned}$$

where an explicit expression for $g(\mathbf{w})$ can be found at the start of Section IV in [3]. Thus,

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} = \frac{\langle \text{grad } f(\mathbf{q}), \mathbf{r} \rangle}{\|\mathbf{w}\|} \leq \|\text{grad } f(\mathbf{q})\| \frac{\|\mathbf{r}\|}{\|\mathbf{w}\|},$$

where

$$\frac{\|\mathbf{r}\|^2}{\|\mathbf{w}\|^2} = \frac{\|\mathbf{w}\|^2 + \left(q_n - \frac{1}{q_n}\right)^2}{\|\mathbf{w}\|^2} = \frac{\|\mathbf{w}\|^2 + \|\mathbf{w}\|^4/q_n^2}{\|\mathbf{w}\|^2} = \frac{1}{q_n^2} = \frac{1}{1 - \|\mathbf{w}\|^2} \leq \frac{1}{1 - \frac{1}{2000}} = \frac{2000}{1999},$$

where we have invoked our assumption that $\|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}}$. Therefore we obtain

$$\|\text{grad } f(\mathbf{q})\| \geq \frac{\|\mathbf{w}\|}{\|\mathbf{r}\|} \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \geq \sqrt{\frac{1999}{2000}} \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \geq \frac{9}{10} \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|},$$

completing the proof. \square

H. Proof of Lemma II.14

Proof of Lemma II.14 combines the local Lipschitz property of Hess $f(\mathbf{q})$ in Lemma VI.2, and the Taylor's theorem (manifold version, Lemma 7.4.7 of [68]).

Proof. (of Lemma II.14) Let $\gamma(t)$ be the unique geodesic that satisfies $\gamma(0) = \mathbf{q}^{(r)}$, $\gamma(1) = \mathbf{q}^{(r+1)}$, and its directional derivative $\dot{\gamma}(0) = \boldsymbol{\delta}_*$. Since the parallel translation defined by the Riemannian connection is an isometry, then $\|\text{grad } f(\mathbf{q}^{(r+1)})\| = \|\mathcal{P}_\gamma^{0 \leftarrow 1} \text{grad } f(\mathbf{q}^{(r+1)})\|$. Moreover, since $\|\boldsymbol{\delta}_*\| \leq \Delta$, the unconstrained optimality condition in (II.6) implies that $\text{grad } f(\mathbf{q}^{(r)}) + \text{Hess } f(\mathbf{q}^{(r)}) \boldsymbol{\delta}_* = \mathbf{0}_{\mathbf{q}^{(r)}}$. Thus, by using Taylor's theorem in [68], we have

$$\begin{aligned} \|\text{grad } f(\mathbf{q}^{(r+1)})\| &= \left\| \mathcal{P}_\gamma^{0 \leftarrow 1} \text{grad } f(\mathbf{q}^{(r+1)}) - \text{grad } f(\mathbf{q}^{(r)}) - \text{Hess } f(\mathbf{q}^{(r)}) \boldsymbol{\delta}_* \right\| \\ &= \left\| \int_0^1 \left[\mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) [\dot{\gamma}(t)] - \text{Hess } f(\mathbf{q}^{(r)}) \boldsymbol{\delta}_* \right] dt \right\| \quad (\text{Taylor's theorem}) \\ &= \left\| \int_0^1 \left(\mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) \mathcal{P}_\gamma^{t \leftarrow 0} \boldsymbol{\delta}_* - \text{Hess } f(\mathbf{q}^{(r)}) \boldsymbol{\delta}_* \right) dt \right\| \\ &\leq \|\boldsymbol{\delta}_*\| \int_0^1 \left\| \mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) \mathcal{P}_\gamma^{t \leftarrow 0} - \text{Hess } f(\mathbf{q}^{(r)}) \right\| dt. \end{aligned}$$

From the Lipschitz bound in Lemma VI.2 and the optimality condition in (II.6), we obtain

$$\|\text{grad } f(\mathbf{q}^{(r+1)})\| \leq \frac{1}{2} \|\boldsymbol{\delta}_*\|^2 L_H = \frac{L_H}{2m_H^2} \|\text{grad } f(\mathbf{q}^{(r)})\|^2.$$

This completes the proof. \square

I. Proof of Lemma II.16

Proof. By invoking Taylor's theorem in [68], we have

$$\mathcal{P}_\gamma^{0 \leftarrow \tau} \text{grad } f(\gamma(\tau)) = \int_0^\tau \mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) [\dot{\gamma}(t)] dt.$$

Hence, we have

$$\begin{aligned} \langle \mathcal{P}_\gamma^{0 \leftarrow \tau} \text{grad } f(\gamma(\tau)), \boldsymbol{\delta} \rangle &= \int_0^\tau \langle \mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) [\dot{\gamma}(t)], \boldsymbol{\delta} \rangle dt \\ &= \int_0^\tau \langle \mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) [\dot{\gamma}(t)], \mathcal{P}_\gamma^{0 \leftarrow t} \dot{\gamma}(t) \rangle dt \\ &= \int_0^\tau \langle \text{Hess } f(\gamma(t)) [\dot{\gamma}(t)], \dot{\gamma}(t) \rangle dt \\ &\geq m_H \int_0^\tau \|\dot{\gamma}(t)\|^2 dt \geq m_H \tau \|\boldsymbol{\delta}\|^2, \end{aligned}$$

where we have used the fact that the parallel transport $\mathcal{P}_\gamma^{0 \leftarrow t}$ defined by the Riemannian connection is an isometry. On the other hand, we have

$$\langle \mathcal{P}_\gamma^{0 \leftarrow \tau} \text{grad } f(\gamma(\tau)), \boldsymbol{\delta} \rangle \leq \|\mathcal{P}_\gamma^{0 \leftarrow \tau} \text{grad } f(\gamma(\tau))\| \|\boldsymbol{\delta}\| = \|\text{grad } f(\gamma(\tau))\| \|\boldsymbol{\delta}\|,$$

where again used the isometry property of the operator $\mathcal{P}_\gamma^{0 \leftarrow \tau}$. Combining the two bounds above, we obtain

$$\|\text{grad } f(\gamma(\tau))\| \|\boldsymbol{\delta}\| \geq m_{H\tau} \|\boldsymbol{\delta}\|^2,$$

which implies the claimed result. \square

VII. PROOFS OF TECHNICAL RESULTS FOR SECTION III

We need one technical lemma to prove Lemma III.2 and the relevant lemma for complete dictionaries.

Lemma VII.1. *For all integer $n_1 \in \mathbb{N}$, $\theta \in (0, 1/3)$, and $n_2 \in \mathbb{N}$ with $n_2 \geq C n_1 \log(n_1/\theta)/\theta^2$, any random matrix $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2} \sim_{i.i.d.} \text{BG}(\theta)$ obeys the following: for any fixed index set $\mathcal{I} \subset [n_2]$ with $|\mathcal{I}| \leq \frac{9}{8}\theta n_2$, it holds that*

$$\|\mathbf{v}^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|\mathbf{v}^* \mathbf{M}_{\mathcal{I}}\|_1 \geq \frac{n_2}{6} \sqrt{\frac{2}{\pi}} \theta \|\mathbf{v}\| \quad \text{for all } \mathbf{v} \in \mathbb{R}^{n_1},$$

with probability at least $1 - cp^{-6}$. Here C, c are both positive constants.

Proof. By homogeneity, it is sufficient to consider all $\mathbf{v} \in \mathbb{S}^{n_1}$. For any $i \in [n_2]$, let $\mathbf{m}_i \in \mathbb{R}^{n_1}$ be a column of \mathbf{M} . For a fixed \mathbf{v} such that $\|\mathbf{v}\| = 1$, we have

$$T(\mathbf{v}) \doteq \|\mathbf{v}^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|\mathbf{v}^* \mathbf{M}_{\mathcal{I}}\|_1 = \sum_{i \in \mathcal{I}^c} |\mathbf{v}^* \mathbf{m}_i| - \sum_{i \in \mathcal{I}} |\mathbf{v}^* \mathbf{m}_i|,$$

namely as a sum of independent random variables. Since $|\mathcal{I}| \leq 9n_2\theta/8$, we have

$$\mathbb{E}[T(\mathbf{v})] \geq \left(n_2 - \frac{9}{8}\theta n_2 - \frac{9}{8}\theta n_2\right) \mathbb{E}[|\mathbf{v}^* \mathbf{m}_1|] = \left(1 - \frac{9}{4}\theta\right) n_2 \mathbb{E}[|\mathbf{v}^* \mathbf{m}_1|] \geq \frac{1}{4} n_2 \mathbb{E}[|\mathbf{v}^* \mathbf{m}_1|],$$

where the expectation $\mathbb{E}[|\mathbf{v}^* \mathbf{m}_1|]$ can be lower bounded as

$$\begin{aligned} \mathbb{E}[|\mathbf{v}^* \mathbf{m}_1|] &= \sum_{k=0}^{n_1} \theta^k (1-\theta)^{n_1-k} \sum_{\mathcal{J} \in \binom{[n_1]}{k}} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [|\mathbf{v}_{\mathcal{J}}^* \mathbf{g}|] \\ &= \sum_{k=0}^{n_1} \theta^k (1-\theta)^{n_1-k} \sum_{\mathcal{J} \in \binom{[n_1]}{k}} \sqrt{\frac{2}{\pi}} \|\mathbf{v}_{\mathcal{J}}\| \geq \sqrt{\frac{2}{\pi}} \|\mathbb{E}_{\mathcal{J}}[\mathbf{v}_{\mathcal{J}}]\| = \sqrt{\frac{2}{\pi}} \theta. \end{aligned}$$

Moreover, by Lemma B.8 and Lemma A.3, for any $i \in [n_2]$ and any integer $m \geq 2$,

$$\mathbb{E}[|\mathbf{v}^* \mathbf{m}_i|^m] \leq \mathbb{E}_{\mathbf{Z} \sim \mathcal{N}(0,1)} [|\mathbf{Z}|^m] \leq (m-1)!! \leq \frac{m!}{2}.$$

So invoking the moment-control Bernstein's inequality in Lemma A.5, we obtain

$$\mathbb{P}\left[T(\mathbf{v}) < \frac{n_2}{4} \sqrt{\frac{2}{\pi}} \theta - t\right] \leq \mathbb{P}[T(\mathbf{v}) < \mathbb{E}[T(\mathbf{v})] - t] \leq \exp\left(-\frac{t^2}{2n_2 + 2t}\right).$$

Taking $t = \frac{n_2}{20} \sqrt{\frac{2}{\pi}} \theta$ and simplifying, we obtain that

$$\mathbb{P}\left[T(\mathbf{v}) < \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta\right] \leq \exp(-c_1 \theta^2 n_2). \quad (\text{VII.1})$$

Fix $\varepsilon = \sqrt{\frac{2}{\pi}} \frac{\theta}{120} [n_1 \log(n_1 n_2)]^{-1/2} < 1$. The unit sphere \mathbb{S}^{n_1} has an ε -net N_ε of cardinality at most $(3/\varepsilon)^{n_1}$. Consider the event

$$\mathcal{E}_{bg} \doteq \left\{ T(\mathbf{v}) \geq \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta \quad \forall \mathbf{v} \in N_\varepsilon \right\}.$$

A simple union bound implies

$$\mathbb{P}[\mathcal{E}_{bg}^c] \leq \exp\left(-c_1\theta^2n_2 + n_1 \log\left(\frac{3}{\varepsilon}\right)\right) \leq \exp\left(-c_1\theta^2n_2 + c_2n_1 \log\frac{n_1 \log n_2}{\theta}\right). \quad (\text{VII.2})$$

Conditioned on \mathcal{E}_{bg} , we have that any $\mathbf{z} \in \mathbb{S}^{n_1-1}$ can be written as $\mathbf{z} = \mathbf{v} + \mathbf{e}$ for some $\mathbf{v} \in N_\varepsilon$ and $\|\mathbf{e}\| \leq \varepsilon$. Moreover,

$$\begin{aligned} T(\mathbf{z}) &= \|(\mathbf{v} + \mathbf{e})^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|(\mathbf{v} + \mathbf{e})^* \mathbf{M}_{\mathcal{I}}\|_1 \geq T(\mathbf{v}) - \|\mathbf{e}^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|\mathbf{e}^* \mathbf{M}_{\mathcal{I}}\|_1 \\ &= \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \|\mathbf{e}^* \mathbf{M}\|_1 = \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \sum_{k=1}^{n_2} |\mathbf{e}^* \mathbf{m}_k| \\ &\geq \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \varepsilon \sum_{k=1}^{n_2} \|\mathbf{m}_k\|. \end{aligned}$$

By Lemma B.6, with probability at least $1 - \theta(n_1n_2)^{-7} - \exp(-0.3\theta n_1n_2)$, $\|\mathbf{M}\|_\infty \leq 4\sqrt{\log(n_1n_2)}$. Thus,

$$T(\mathbf{z}) \geq \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \sqrt{\frac{2}{\pi}} \frac{\theta}{120} \frac{n_2 \sqrt{n_1} 4 \sqrt{\log(n_1n_2)}}{\sqrt{n_1} \sqrt{\log(n_1n_2)}} = \frac{n_2}{6} \sqrt{\frac{2}{\pi}} \theta. \quad (\text{VII.3})$$

Thus, by (VII.2), it is enough to take $n_2 > Cn_1 \log(n_1/\theta)/\theta^2$ for sufficiently large $C > 0$ to make the overall failure probability small enough so that the lower bound (VII.3) holds. \square

A. Proof of Lemma III.2

Proof. The proof is similar to that of [44]. First, let us assume the dictionary $\mathbf{A}_0 = \mathbf{I}$. W.l.o.g., suppose that the Riemannian TRM algorithm returns a solution $\hat{\mathbf{q}}$, to which \mathbf{e}_n is the nearest signed basis vector. Thus, the rounding LP (III.1) takes the form:

$$\text{minimize}_{\mathbf{q}} \|\mathbf{q}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad \langle \mathbf{r}, \mathbf{q} \rangle = 1. \quad (\text{VII.4})$$

where the vector $\mathbf{r} = \hat{\mathbf{q}}$. Next, We will show whenever $\hat{\mathbf{q}}$ is close enough to \mathbf{e}_n , w.h.p., the above linear program returns \mathbf{e}_n . Let $\mathbf{X}_0 = [\overline{\mathbf{X}}; \mathbf{x}_n^*]$, where $\overline{\mathbf{X}} \in \mathbb{R}^{(n-1) \times p}$ and \mathbf{x}_n^* is the last row of \mathbf{X}_0 . Set $\mathbf{q} = [\overline{\mathbf{q}}, q_n]$, where $\overline{\mathbf{q}}$ denotes the first $n-1$ coordinates of \mathbf{q} and q_n is the last coordinate; similarly for \mathbf{r} . Let us consider a relaxation of the problem (VII.4),

$$\text{minimize}_{\mathbf{q}} \|\mathbf{q}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad q_n r_n + \langle \overline{\mathbf{q}}, \overline{\mathbf{r}} \rangle \geq 1, \quad (\text{VII.5})$$

It is obvious that the feasible set of (VII.5) contains that of (VII.4). So if \mathbf{e}_n is the unique optimal solution (UOS) of (VII.5), it is the UOS of (VII.4). Suppose $\mathcal{I} = \text{supp}(\mathbf{x}_n)$ and define an event $\mathcal{E}_0 = \{|\mathcal{I}| \leq \frac{9}{8}\theta p\}$. By Hoeffding's inequality, we know that $\mathbb{P}[\mathcal{E}_0^c] \leq \exp(-\theta^2 p/2)$. Now conditioned on \mathcal{E}_0 and consider a fixed support \mathcal{I} . (VII.5) can be further relaxed as

$$\text{minimize}_{\mathbf{q}} \|\mathbf{x}_n\|_1 |q_n| - \|\overline{\mathbf{q}}^* \overline{\mathbf{X}}_{\mathcal{I}}\|_1 + \|\overline{\mathbf{q}}^* \overline{\mathbf{X}}_{\mathcal{I}^c}\|_1, \quad \text{subject to} \quad q_n r_n + \|\overline{\mathbf{q}}\| \|\overline{\mathbf{r}}\| \geq 1. \quad (\text{VII.6})$$

The objective value of (VII.6) lower bounds that of (VII.5), and are equal when $\mathbf{q} = \mathbf{e}_n$. So if $\mathbf{q} = \mathbf{e}_n$ is UOS of (VII.6), it is UOS of (VII.4). By Lemma VII.1, we know that

$$\|\overline{\mathbf{q}}^* \overline{\mathbf{X}}_{\mathcal{I}^c}\|_1 - \|\overline{\mathbf{q}}^* \overline{\mathbf{X}}_{\mathcal{I}}\|_1 \geq \frac{p}{6} \sqrt{\frac{2}{\pi}} \theta \|\overline{\mathbf{q}}\|$$

holds w.h.p. when $p \geq C_1(n-1) \log((n-1)/\theta)/\theta^2$. Let $\zeta = \frac{p}{6} \sqrt{\frac{2}{\pi}} \theta$, thus we can further lower bound the objective value in (VII.6) by

$$\text{minimize}_{\mathbf{q}} \|\mathbf{x}_n\|_1 |q_n| + \zeta \|\overline{\mathbf{q}}\|, \quad \text{subject to} \quad q_n r_n + \|\overline{\mathbf{q}}\| \|\overline{\mathbf{r}}\| \geq 1. \quad (\text{VII.7})$$

By similar arguments, if e_n is the UOS of (VII.7), it is also the UOS of (VII.4). For the optimal solution of (VII.7), notice that it is necessary to have $\text{sign}(q_n) = \text{sign}(r_n)$ and $q_n r_n + \|\tilde{\mathbf{q}}\| \|\tilde{\mathbf{r}}\| = 1$. Therefore, the problem (VII.7) is equivalent to

$$\text{minimize}_{q_n} \|\mathbf{x}_n\|_1 |q_n| + \zeta \frac{1 - |r_n| |q_n|}{\|\tilde{\mathbf{r}}\|}, \quad \text{subject to} \quad |q_n| \leq \frac{1}{|r_n|}. \quad (\text{VII.8})$$

Notice that the problem (VII.8) is a linear program in $|q_n|$ with a compact feasible set, which indicates that the optimal solution only occurs at the boundary points $|q_n| = 0$ and $|q_n| = 1/|r_n|$. Therefore, $\mathbf{q} = e_n$ is the UOS of (VII.8) if and only if

$$\frac{1}{|r_n|} \|\mathbf{x}_n\|_1 < \frac{\zeta}{\|\tilde{\mathbf{r}}\|}. \quad (\text{VII.9})$$

Conditioned on \mathcal{E}_0 , by using the Gaussian concentration bound, we have

$$\mathbb{P} \left[\|\mathbf{x}_n\|_1 \geq \frac{9}{8} \sqrt{\frac{2}{\pi}} \theta p + t \right] \leq \mathbb{P} [\|\mathbf{x}_n\|_1 \geq \mathbb{E}[\|\mathbf{x}_n\|_1] + t] \leq \exp\left(-\frac{t^2}{2p}\right),$$

which means that

$$\mathbb{P} \left[\|\mathbf{x}_n\|_1 \geq \frac{5}{4} \sqrt{\frac{2}{\pi}} \theta p \right] \leq \exp\left(-\frac{\theta^2 p}{64\pi}\right). \quad (\text{VII.10})$$

Therefore, by (VII.9) and (VII.10), for $\mathbf{q} = e_n$ to be the UOS of (VII.4) w.h.p., it is sufficient to have

$$\frac{5}{4|r_n|} \sqrt{\frac{2}{\pi}} \theta p < \frac{\theta p}{6\sqrt{1-|r_n|^2}} \sqrt{\frac{2}{\pi}}, \quad (\text{VII.11})$$

which is implied by

$$|r_n| > \frac{249}{250}.$$

The failure probability can be estimated via a simple union bound. Since the above argument holds uniformly for any fixed support set \mathcal{I} , we obtain the desired result.

When our dictionary \mathbf{A}_0 is an arbitrary orthogonal matrix, it only rotates the row subspace of \mathbf{X}_0 . Thus, w.l.o.g., suppose the TRM algorithm returns a solution $\hat{\mathbf{q}}$, to which $\mathbf{A}_0 \mathbf{q}_*$ is the nearest ‘‘target’’ with \mathbf{q}_* a signed basis vector. By a change of variable $\tilde{\mathbf{q}} = \mathbf{A}_0^* \mathbf{q}$, the problem (VII.4) is of the form

$$\text{minimize}_{\tilde{\mathbf{q}}} \|\tilde{\mathbf{q}}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad \langle \mathbf{A}_0^* \mathbf{r}, \tilde{\mathbf{q}} \rangle = 1,$$

obviously our target solution for $\tilde{\mathbf{q}}$ is again the standard basis \mathbf{q}_* . By a similar argument above, we only need $\langle \mathbf{A}_0^* \mathbf{r}, e_n \rangle > 249/250$ to exactly recover the target, which is equivalent to $\langle \mathbf{r}, \hat{\mathbf{q}}_* \rangle > 249/250$. This implies that our rounding (III.1) is invariant to change of basis, completing the proof. \square

B. Proof of Lemma III.4

Proof. Define $\tilde{\mathbf{q}} \doteq (\mathbf{UV}^* + \mathbf{\Xi})^* \mathbf{q}$. By Lemma B.7, and in particular (II.48), when

$$p \geq \frac{C_1}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right),$$

$\|\mathbf{\Xi}\| \leq 1/2$ so that $\mathbf{UV}^* + \mathbf{\Xi}$ is invertible. Then the LP rounding can be written as

$$\text{minimize}_{\tilde{\mathbf{q}}} \|\tilde{\mathbf{q}}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad \langle (\mathbf{UV}^* + \mathbf{\Xi})^{-1} \mathbf{r}, \tilde{\mathbf{q}} \rangle = 1.$$

By Lemma III.2, to obtain $\tilde{\mathbf{q}} = e_n$ from this LP, it is enough to have

$$\langle (\mathbf{UV}^* + \mathbf{\Xi})^{-1} \mathbf{r}, e_n \rangle \geq 249/250,$$

and $p \geq C_2 n^2 \log(n/\theta)/\theta$ for some large enough C_2 . This implies that to obtain \mathbf{q}_* for the original LP, such that $(\mathbf{UV}^* + \mathbf{\Xi})^* \mathbf{q}_* = e_n$, it is enough that

$$\langle (\mathbf{UV}^* + \mathbf{\Xi})^{-1} \mathbf{r}, (\mathbf{UV}^* + \mathbf{\Xi})^* \mathbf{q}_* \rangle = \langle \mathbf{r}, \mathbf{q}_* \rangle \geq 249/250,$$

completing the proof. \square

C. Proof of Lemma III.5

Proof. Note that $[\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell] = (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}[\mathbf{e}_1, \dots, \mathbf{e}_\ell]$, we have

$$\begin{aligned} \mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 &= \mathbf{U}^*(\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}(\mathbf{Q} + \mathbf{\Xi})^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 \\ &= \mathbf{U}^* \left[\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell \mid \widehat{\mathbf{V}} \right] (\mathbf{I} + \mathbf{\Delta}_1)\mathbf{X}_0, \end{aligned}$$

where $\widehat{\mathbf{V}} \doteq (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}[\mathbf{e}_{\ell+1}, \dots, \mathbf{e}_n]$, and the matrix $\mathbf{\Delta}_1 = \mathbf{Q}^*\mathbf{\Xi} + \mathbf{\Xi}^*\mathbf{Q} + \mathbf{\Xi}^*\mathbf{\Xi}$ so that $\|\mathbf{\Delta}_1\| \leq 3\|\mathbf{\Xi}\|$. Since $\mathbf{U}^* \left[\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell \mid \widehat{\mathbf{V}} \right] = \left[\mathbf{0} \mid \mathbf{U}^*\widehat{\mathbf{V}} \right]$, we have

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \left[\mathbf{0} \mid \mathbf{U}^*\widehat{\mathbf{V}} \right] \mathbf{X}_0 + \left[\mathbf{0} \mid \mathbf{U}^*\widehat{\mathbf{V}} \right] \mathbf{\Delta}_1 \mathbf{X}_0 = \mathbf{U}^*\widehat{\mathbf{V}} \mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}_2 \mathbf{X}_0, \quad (\text{VII.12})$$

where $\mathbf{\Delta}_2 = \left[\mathbf{0} \mid \mathbf{U}^*\widehat{\mathbf{V}} \right] \mathbf{\Delta}_1$. Let $\delta = \|\mathbf{\Xi}\|$, so that

$$\|\mathbf{\Delta}_2\| \leq \frac{\|\mathbf{\Delta}_1\|}{\sigma_{\min}(\mathbf{Q} + \mathbf{\Xi})} \leq \frac{3\|\mathbf{\Xi}\|}{\sigma_{\min}(\mathbf{Q} + \mathbf{\Xi})} \leq \frac{3\delta}{1-\delta}. \quad (\text{VII.13})$$

Since the matrix $\widehat{\mathbf{V}}$ is near orthogonal, it can be decomposed as $\widehat{\mathbf{V}} = \mathbf{V} + \mathbf{\Delta}_3$, where \mathbf{V} is orthogonal, and $\mathbf{\Delta}_3$ is a small perturbation. Obviously, $\mathbf{V} = \mathbf{U}\mathbf{R}$ for some orthogonal matrix \mathbf{R} , so that \mathbf{V} spans the same subspace as that of \mathbf{U} . Next, we control the spectral norm of $\mathbf{\Delta}_3$:

$$\|\mathbf{\Delta}_3\| = \min_{\mathbf{R} \in \mathcal{O}_\ell} \left\| \mathbf{U}\mathbf{R} - \widehat{\mathbf{V}} \right\| \leq \min_{\mathbf{R} \in \mathcal{O}_\ell} \left\| \mathbf{U}\mathbf{R} - \mathbf{Q}_{[n-\ell]} \right\| + \left\| \mathbf{Q}_{[n-\ell]} - \widehat{\mathbf{V}} \right\|, \quad (\text{VII.14})$$

where $\mathbf{Q}_{[n-\ell]}$ collects the last $n - \ell$ columns of \mathbf{Q} , i.e., $\mathbf{Q} = [\mathbf{Q}_{[\ell]}, \mathbf{Q}_{[n-\ell]}]$.

To bound the second term on the right, we have

$$\left\| \mathbf{Q}_{[n-\ell]} - \widehat{\mathbf{V}} \right\| \leq \left\| \mathbf{Q}^{-1} - (\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| \leq \frac{\|\mathbf{Q}^{-1}\| \|\mathbf{Q}^{-1}\mathbf{\Xi}\|}{1 - \|\mathbf{Q}^{-1}\mathbf{\Xi}\|} \leq \frac{\delta}{1-\delta},$$

where we have used perturbation bound for matrix inverse (see, e.g., Theorem 2.5 of Chapter III in [84]).

To bound the first term, from Lemma B.2, it is enough to upper bound the largest principal angle θ_1 between the subspaces $\text{span}([\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell])$, and that spanned by $\mathbf{Q}[\mathbf{e}_1, \dots, \mathbf{e}_\ell]$. Write $\mathbf{I}_{[\ell]} \doteq [\mathbf{e}_1, \dots, \mathbf{e}_\ell]$ for short, we bound $\sin \theta_1$ as

$$\begin{aligned} \sin \theta_1 &\leq \left\| \mathbf{Q}\mathbf{I}_{[\ell]}\mathbf{I}_{[\ell]}^*\mathbf{Q}^* - (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}\mathbf{I}_{[\ell]} \left(\mathbf{I}_{[\ell]}^*(\mathbf{Q} + \mathbf{\Xi})^{-1}(\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}\mathbf{I}_{[\ell]} \right)^{-1} \mathbf{I}_{[\ell]}^*(\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| \\ &= \left\| \mathbf{Q}\mathbf{I}_{[\ell]}\mathbf{I}_{[\ell]}^*\mathbf{Q}^* - (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}\mathbf{I}_{[\ell]} \left(\mathbf{I}_{[\ell]}^*(\mathbf{I} + \mathbf{\Delta}_1)^{-1}\mathbf{I}_{[\ell]} \right)^{-1} \mathbf{I}_{[\ell]}^*(\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| \\ &\leq \left\| \mathbf{Q}\mathbf{I}_{[\ell]}\mathbf{I}_{[\ell]}^*\mathbf{Q}^* - (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}\mathbf{I}_{[\ell]}\mathbf{I}_{[\ell]}^*(\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| \\ &\quad + \left\| (\mathbf{Q}^* + \mathbf{\Xi}^*)^{-1}\mathbf{I}_{[\ell]} \left[\mathbf{I} - \left(\mathbf{I}_{[\ell]}^*(\mathbf{I} + \mathbf{\Delta}_1)^{-1}\mathbf{I}_{[\ell]} \right)^{-1} \right] \mathbf{I}_{[\ell]}^*(\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| \\ &\leq \left(1 + \frac{1}{\sigma_{\min}(\mathbf{Q} + \mathbf{\Xi})} \right) \left\| \mathbf{Q}^{-1} - (\mathbf{Q} + \mathbf{\Xi})^{-1} \right\| + \frac{1}{\sigma_{\min}^2(\mathbf{Q} + \mathbf{\Xi})} \left\| \mathbf{I} - \left(\mathbf{I}_{[\ell]}^*(\mathbf{I} + \mathbf{\Delta}_1)^{-1}\mathbf{I}_{[\ell]} \right)^{-1} \right\| \\ &\leq \left(1 + \frac{1}{1-\delta} \right) \frac{\delta}{1-\delta} + \frac{1}{(1-\delta)^2} \frac{\left\| \mathbf{I}_{[\ell]}^*(\mathbf{I} + \mathbf{\Delta}_1)^{-1}\mathbf{I}_{[\ell]} - \mathbf{I} \right\|}{1 - \left\| \mathbf{I}_{[\ell]}^*(\mathbf{I} + \mathbf{\Delta}_1)^{-1}\mathbf{I}_{[\ell]} - \mathbf{I} \right\|} \\ &\leq \left(1 + \frac{1}{1-\delta} \right) \frac{\delta}{1-\delta} + \frac{1}{(1-\delta)^2} \frac{\|\mathbf{\Delta}_1\|}{1 - 2\|\mathbf{\Delta}_1\|}, \end{aligned}$$

where to obtain the first line we used that for any full column rank matrix \mathbf{M} , $\mathbf{M}(\mathbf{M}^*\mathbf{M})^{-1}\mathbf{M}^*$ is the orthogonal projector onto the its column span, and to obtain the fifth and six lines we invoked the matrix inverse perturbation bound again. Using $\delta < 1/20$ and $\|\mathbf{\Delta}_1\| \leq 3\delta < 1/2$, we have

$$\sin \theta_1 \leq \frac{(2-\delta)\delta}{(1-\delta)^2} + \frac{3\delta}{(1-\delta)^2(1-6\delta)} = \frac{5\delta - 13\delta^2 + 6\delta^3}{(1-\delta)^2(1-6\delta)} \leq 8\delta.$$

For $\delta < 1/20$, the upper bound is nontrivial. By Lemma B.2,

$$\min_{\mathbf{R} \in O_\ell} \|\mathbf{UR} - \mathbf{Q}_{[n-\ell]}\| \leq \sqrt{2 - 2 \cos \theta_1} \leq \sqrt{2 - 2 \cos^2 \theta_1} = \sqrt{2} \sin \theta_1 \leq 8\sqrt{2}\delta.$$

Put the estimates above, there exists an orthogonal matrix $\mathbf{R} \in O_\ell$ such that $\mathbf{V} = \mathbf{UR}$ and $\widehat{\mathbf{V}} = \mathbf{V} + \mathbf{\Delta}_3$ with

$$\|\mathbf{\Delta}_3\| \leq \delta/(1 - \delta) + 8\sqrt{2}\delta \leq 25/(2\delta). \quad (\text{VII.15})$$

Therefore, by (VII.12), we obtain

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}, \quad \text{with} \quad \mathbf{\Delta} \doteq \mathbf{U}^*\mathbf{\Delta}_3\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}_2\mathbf{X}_0. \quad (\text{VII.16})$$

By using the results in (VII.13) and (VII.15), we get the desired result. \square

APPENDIX A

TECHNICAL TOOLS AND BASIC FACTS USED IN PROOFS

In this section, we summarize some basic calculations that are useful throughout, and also record major technical tools we use in proofs.

Lemma A.1 (Derivates and Lipschitz Properties of $h_\mu(z)$). *For the sparsity surrogate*

$$h_\mu(z) = \mu \log \cosh(z/\mu),$$

the first two derivatives are

$$\dot{h}_\mu(z) = \tanh(z/\mu), \quad \ddot{h}_\mu(z) = [1 - \tanh^2(z/\mu)]/\mu.$$

Also, for any $z > 0$, we have

$$\begin{aligned} (1 - \exp(-2z/\mu))/2 &\leq \tanh(z/\mu) \leq 1 - \exp(-2z/\mu), \\ \exp(-2z/\mu) &\leq 1 - \tanh^2(z/\mu) \leq 4 \exp(-2z/\mu). \end{aligned}$$

Moreover, for any $z, z' \in \mathbb{R}$, we have

$$|\dot{h}_\mu(z) - \dot{h}_\mu(z')| \leq |z - z'|/\mu, \quad |\ddot{h}_\mu(z) - \ddot{h}_\mu(z')| \leq 2|z - z'|/\mu^2.$$

Lemma A.2 (Gaussian Tail Estimates). *Let $X \sim \mathcal{N}(0, 1)$ and $\Phi(x)$ be CDF of X . For any $x \geq 0$, we have the following estimates for $\Phi^c(x) \doteq 1 - \Phi(x)$:*

$$\begin{aligned} \left(\frac{1}{x} - \frac{1}{x^3}\right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}} &\leq \Phi^c(x) \leq \left(\frac{1}{x} - \frac{1}{x^3} + \frac{3}{x^5}\right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}}, \quad (\text{Type I}) \\ \frac{x}{x^2 + 1} \frac{\exp(-x^2/2)}{\sqrt{2\pi}} &\leq \Phi^c(x) \leq \frac{1}{x} \frac{\exp(-x^2/2)}{\sqrt{2\pi}}, \quad (\text{Type II}) \\ \frac{\sqrt{x^2 + 4} - x}{2} \frac{\exp(-x^2/2)}{\sqrt{2\pi}} &\leq \Phi^c(x) \leq \left(\sqrt{2 + x^2} - x\right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}} \quad (\text{Type III}). \end{aligned}$$

Proof. See proof of Lemma A.5 in the technical report [2]. \square

Lemma A.3 (Moments of the Gaussian RV). *If $X \sim \mathcal{N}(0, \sigma^2)$, then it holds for all integer $p \geq 1$ that*

$$\mathbb{E}[|X|^p] \leq \sigma^p (p-1)!!.$$

Lemma A.4 (Moments of the χ^2 RV). *If $X \sim \chi^2(n)$, then it holds for all integer $p \geq 1$ that*

$$\mathbb{E}[X^p] = 2^p \frac{\Gamma(p + n/2)}{\Gamma(n/2)} = \prod_{k=1}^p (n + 2k - 2) \leq p! (2n)^p / 2.$$

Lemma A.5 (Moment-Control Bernstein's Inequality for Scalar RVs, Theorem 2.10 of [85]). *Let X_1, \dots, X_p be i.i.d. real-valued random variables. Suppose that there exist some positive numbers R and σ^2 such that*

$$\mathbb{E}[|X_k|^m] \leq m! \sigma^2 R^{m-2} / 2, \quad \text{for all integers } m \geq 2.$$

Let $S \doteq \frac{1}{p} \sum_{k=1}^p X_k$, then for all $t > 0$, it holds that

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

Lemma A.6 (Moment-Control Bernstein's Inequality for Matrix RVs, Theorem 6.2 of [86]). Let $\mathbf{X}_1, \dots, \mathbf{X}_p \in \mathbb{R}^{d \times d}$ be i.i.d. random, symmetric matrices. Suppose there exist some positive number R and σ^2 such that

$$\mathbb{E}[\mathbf{X}_k^m] \preceq m! \sigma^2 R^{m-2} / 2 \cdot \mathbf{I} \text{ and } -\mathbb{E}[\mathbf{X}_k^m] \preceq m! \sigma^2 R^{m-2} / 2 \cdot \mathbf{I}, \text{ for all integers } m \geq 2.$$

Let $S \doteq \frac{1}{p} \sum_{k=1}^p \mathbf{X}_k$, then for all $t > 0$, it holds that

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

Proof. See proof of Lemma A.10 in the technical report [2]. □

Lemma A.7 (Integral Form of Taylor's Theorem). Let $f(\mathbf{x}) : \mathbb{R}^n \mapsto \mathbb{R}$ be a twice continuously differentiable function, then for any direction $\mathbf{y} \in \mathbb{R}^n$, we have

$$\begin{aligned} f(\mathbf{x} + t\mathbf{y}) &= f(\mathbf{x}) + t \int_0^1 \langle \nabla f(\mathbf{x} + s t \mathbf{y}), \mathbf{y} \rangle ds, \\ f(\mathbf{x} + t\mathbf{y}) &= f(\mathbf{x}) + t \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle + t^2 \int_0^1 (1-s) \langle \nabla^2 f(\mathbf{x} + s t \mathbf{y}) \mathbf{y}, \mathbf{y} \rangle ds. \end{aligned}$$

APPENDIX B

AUXILLARY RESULTS FOR PROOFS

Lemma B.1. There exists a positive constant C such that for any $\theta \in (0, 1/2)$ and $n_2 > C n_1^2 \log n_1$, the random matrix $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with $\mathbf{X} \sim_{i.i.d.} \text{BG}(\theta)$ obeys

$$\left\| \frac{1}{n_2 \theta} \mathbf{X} \mathbf{X}^* - \mathbf{I} \right\| \leq 10 \sqrt{\frac{\theta n_1 \log n_2}{n_2}} \quad (\text{B.1})$$

with probability at least $1 - n_2^{-8}$.

Proof. See proof of Lemma B.3 in [3]. □

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

- i) $\min_{\mathbf{Q} \in O_k} \|\mathbf{U} - \mathbf{V}\mathbf{Q}\| \leq \sqrt{2 - 2 \cos \theta_1}$;
- ii) $\sin \theta_1 = \|\mathbf{U}\mathbf{U}^* - \mathbf{V}\mathbf{V}^*\|$;
- iii) Let \mathcal{U}^\perp and \mathcal{V}^\perp be the orthogonal complement of \mathcal{U} and \mathcal{V} , respectively. Then $\theta_1(\mathcal{U}, \mathcal{V}) = \theta_1(\mathcal{U}^\perp, \mathcal{V}^\perp)$.

Proof. See proof of Lemma B.4 in [2]. □

Below are restatements of several technical results from [3] that are important for proofs in this paper.

Proposition B.3 (Hessian Lipschitz). Fix any $r_\wedge \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_\wedge\}$, $\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w} / \|\mathbf{w}\|^2$ is L_\wedge -Lipschitz with

$$L_\wedge \leq \frac{16n^3}{\mu^2} \|\mathbf{X}_0\|_\infty^3 + \frac{8n^{3/2}}{\mu r_\wedge} \|\mathbf{X}_0\|_\infty^2 + \frac{48n^{5/2}}{\mu} \|\mathbf{X}_0\|_\infty^2 + 96n^{5/2} \|\mathbf{X}_0\|_\infty.$$

Proposition B.4 (Gradient Lipschitz). Fix any $r_g \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_g\}$, $\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0) / \|\mathbf{w}\|$ is L_g -Lipschitz with

$$L_g \leq \frac{2\sqrt{n} \|\mathbf{X}_0\|_\infty}{r_g} + 8n^{3/2} \|\mathbf{X}_0\|_\infty + \frac{4n^2}{\mu} \|\mathbf{X}_0\|_\infty^2.$$

Proposition B.5 (Lipschitz for Hessian around zero). *Fix any $r_\cup \in (0, 1/2)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \leq r_\cup\}$, $\nabla^2 g(\mathbf{w}; \mathbf{X}_0)$ is L_\cup -Lipschitz with*

$$L_\cup \leq \frac{4n^2}{\mu^2} \|\mathbf{X}_0\|_\infty^3 + \frac{4n}{\mu} \|\mathbf{X}_0\|_\infty^2 + \frac{8\sqrt{2}\sqrt{n}}{\mu} \|\mathbf{X}_0\|_\infty^2 + 8 \|\mathbf{X}_0\|_\infty.$$

Lemma B.6. *For any $\theta \in (0, 1)$, consider the random matrix $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with $\mathbf{X} \sim_{i.i.d.} \text{BG}(\theta)$. Define the event $\mathcal{E}_\infty \doteq \left\{1 \leq \|\mathbf{X}\|_\infty \leq 4\sqrt{\log(np)}\right\}$. It holds that*

$$\mathbb{P}[\mathcal{E}_\infty^c] \leq \theta (np)^{-7} + \exp(-0.3\theta np).$$

Lemma B.7. *For any $\theta \in (0, 1/2)$, suppose \mathbf{A}_0 is complete with condition number $\kappa(\mathbf{A}_0)$ and $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$. Provided $p \geq C\kappa^4(\mathbf{A}_0)\theta n^2 \log(n\theta\kappa(\mathbf{A}_0))$, one can write $\bar{\mathbf{Y}}$ as defined in (II.47) as*

$$\bar{\mathbf{Y}} = \mathbf{U}\mathbf{V}^*\mathbf{X}_0 + \Xi\mathbf{X}_0,$$

for a certain Ξ obeying $\|\Xi\| \leq 20\kappa^4(\mathbf{A})\sqrt{\frac{\theta n \log p}{p}}$, with probability at least $1 - p^{-8}$. Here $\mathbf{U}\Sigma\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$, and C is a positive numerical constant.

Lemma B.8. *Suppose $\mathbf{z}, \mathbf{z}' \in \mathbb{R}^n$ are independent and obey $\mathbf{z} \sim_{i.i.d.} \text{BG}(\theta)$ and $\mathbf{z}' \sim_{i.i.d.} \mathcal{N}(0, 1)$. Then, for any fixed vector $\mathbf{v} \in \mathbb{R}^n$, it holds that*

$$\begin{aligned} \mathbb{E}[\|\mathbf{v}^*\mathbf{z}\|^m] &\leq \mathbb{E}[\|\mathbf{v}^*\mathbf{z}'\|^m] = \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{v}\|^2)}[|Z|^m], \\ \mathbb{E}[\|\mathbf{z}\|^m] &\leq \mathbb{E}[\|\mathbf{z}'\|^m], \end{aligned}$$

for all integers $m \geq 1$.

ACKNOWLEDGMENT

We thank Dr. Boaz Barak for pointing out an inaccurate comment made on overcomplete dictionary learning using SOS. We thank Cun Mu and Henry Kuo of Columbia University for discussions related to this project. We also thank the anonymous reviewers for their careful reading of the paper, and for comments which have helped us to substantially improve the presentation. JS thanks the Wei Family Private Foundation for their generous support. This work was partially supported by grants ONR N00014-13-1-0492, NSF 1343282, NSF CCF 1527809, NSF IIS 1546411, and funding from the Moore and Sloan Foundations.

REFERENCES

- [1] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery using nonconvex optimization," in *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015, pp. 2351–2360.
- [2] —, "Complete dictionary recovery over the sphere," *arXiv preprint arXiv:1504.06785*, 2015.
- [3] —, "Complete dictionary recovery over the sphere I: Overview and the geometric picture," *arXiv preprint arXiv:1511.03607*, 2015.
- [4] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *Information Theory, IEEE Transactions on*, vol. 56, no. 6, pp. 2980–2998, 2010.
- [5] P. Jain, P. Netrapalli, and S. Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*. ACM, 2013, pp. 665–674.
- [6] M. Hardt, "Understanding alternating minimization for matrix completion," in *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 2014, pp. 651–660.
- [7] M. Hardt and M. Wootters, "Fast matrix completion without the condition number," in *Proceedings of The 27th Conference on Learning Theory*, 2014, pp. 638–678.
- [8] P. Netrapalli, U. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain, "Non-convex robust pca," in *Advances in Neural Information Processing Systems*, 2014, pp. 1107–1115.
- [9] P. Jain and P. Netrapalli, "Fast exact matrix completion with finite samples," *arXiv preprint arXiv:1411.1087*, 2014.
- [10] R. Sun and Z.-Q. Luo, "Guaranteed matrix completion via non-convex factorization," *arXiv preprint arXiv:1411.8003*, 2014.
- [11] Q. Zheng and J. Lafferty, "A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements," *arXiv preprint arXiv:1506.06081*, 2015.
- [12] S. Tu, R. Boczar, M. Soltanolkotabi, and B. Recht, "Low-rank solutions of linear matrix equations via procrustes flow," *arXiv preprint arXiv:1507.03566*, 2015.
- [13] Y. Chen and M. J. Wainwright, "Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees," *arXiv preprint arXiv:1509.03025*, 2015.

- [14] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, “Guarantees of riemannian optimization for low rank matrix recovery,” *arXiv preprint arXiv:1511.01562*, 2015.
- [15] Y. Li, Y. Liang, and A. Risteski, “Recovery guarantee of weighted low-rank approximation via alternating minimization,” *arXiv preprint arXiv:1602.02262*, 2016.
- [16] D. Gamarnik and S. Misra, “A note on alternating minimization algorithm for the matrix completion problem,” *arXiv preprint arXiv:1602.02164*, 2016.
- [17] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, “Guarantees of riemannian optimization for low rank matrix completion,” *arXiv preprint arXiv:1603.06610*, 2016.
- [18] Q. Zheng and J. Lafferty, “Convergence analysis for rectangular matrix completion using burer-monteiro factorization and gradient descent,” *arXiv preprint arXiv:1605.07051*, 2016.
- [19] X. Yi, D. Park, Y. Chen, and C. Caramanis, “Fast algorithms for robust pca via gradient descent,” *arXiv preprint arXiv:1605.07784*, 2016.
- [20] C. Jin, S. M. Kakade, and P. Netrapalli, “Provable efficient online matrix completion via non-convex stochastic gradient descent,” *arXiv preprint arXiv:1605.08370*, 2016.
- [21] D. Park, A. Kyrillidis, C. Caramanis, and S. Sanghavi, “Finding low-rank solutions to matrix problems, efficiently and provably,” *arXiv preprint arXiv:1606.03168*, 2016.
- [22] D. Park, A. Kyrillidis, S. Bhojanapalli, C. Caramanis, and S. Sanghavi, “Provable non-convex projected gradient descent for a class of constrained matrix optimization problems,” *arXiv preprint arXiv:1606.01316*, 2016.
- [23] Y. Cherapanamjeri, K. Gupta, and P. Jain, “Nearly-optimal robust matrix completion,” *arXiv preprint arXiv:1606.07315*, 2016.
- [24] R. Ge, J. D. Lee, and T. Ma, “Matrix completion has no spurious local minimum,” *arXiv preprint arXiv:1605.07272*, 2016.
- [25] S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Global optimality of local search for low rank matrix recovery,” *arXiv preprint arXiv:1605.07221*, 2016.
- [26] P. Netrapalli, P. Jain, and S. Sanghavi, “Phase retrieval using alternating minimization,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2796–2804.
- [27] E. Candès, X. Li, and M. Soltanolkotabi, “Phase retrieval via wirtinger flow: Theory and algorithms,” *Information Theory, IEEE Transactions on*, vol. 61, no. 4, pp. 1985–2007, April 2015.
- [28] Y. Chen and E. J. Candes, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” *arXiv preprint arXiv:1505.05114*, 2015.
- [29] C. D. White, R. Ward, and S. Sanghavi, “The local convexity of solving quadratic equations,” *arXiv preprint arXiv:1506.07868*, 2015.
- [30] J. Sun, Q. Qu, and J. Wright, “A geometric analysis of phase retrieval,” *arXiv preprint arXiv:1602.06664*, 2016.
- [31] H. Zhang, Y. Chi, and Y. Liang, “Provable non-convex phase retrieval with outliers: Median truncated wirtinger flow,” *arXiv preprint arXiv:1603.03805*, 2016.
- [32] H. Zhang and Y. Liang, “Reshaped wirtinger flow for solving quadratic systems of equations,” *arXiv preprint arXiv:1605.07719*, 2016.
- [33] G. Wang, G. B. Giannakis, and Y. C. Eldar, “Solving systems of random quadratic equations via truncated amplitude flow,” *arXiv preprint arXiv:1605.08285*, 2016.
- [34] R. Kolte and A. Özgür, “Phase retrieval via incremental truncated wirtinger flow,” *arXiv preprint arXiv:1606.03196*, 2016.
- [35] B. Gao and Z. Xu, “Gauss-newton method for phase retrieval,” *arXiv preprint arXiv:1606.08135*, 2016.
- [36] T. Bendory and Y. C. Eldar, “Non-convex phase retrieval from stft measurements,” *arXiv preprint arXiv:1607.08218*, 2016.
- [37] P. Jain and S. Oh, “Provable tensor factorization with missing data,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1431–1439.
- [38] A. Anandkumar, R. Ge, and M. Janzamin, “Guaranteed non-orthogonal tensor decomposition via alternating rank-1 updates,” *arXiv preprint arXiv:1402.5180*, 2014.
- [39] —, “Analyzing tensor power method dynamics: Applications to learning overcomplete latent variable models,” *arXiv preprint arXiv:1411.1488*, 2014.
- [40] A. Anandkumar, P. Jain, Y. Shi, and U. Niranjan, “Tensor vs matrix methods: Robust tensor decomposition under block sparse perturbations,” *arXiv preprint arXiv:1510.04747*, 2015.
- [41] S. B. Hopkins, T. Schramm, J. Shi, and D. Steurer, “Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors,” *arXiv preprint arXiv:1512.02337*, 2015.
- [42] X. Yi, C. Caramanis, and S. Sanghavi, “Alternating minimization for mixed linear regression,” *arXiv preprint arXiv:1310.3745*, 2013.
- [43] H. Sedghi and A. Anandkumar, “Provable tensor methods for learning mixtures of classifiers,” *arXiv preprint arXiv:1412.3046*, 2014.
- [44] Q. Qu, J. Sun, and J. Wright, “Finding a sparse vector in a subspace: Linear sparsity using alternating directions,” in *Advances in Neural Information Processing Systems*, 2014, pp. 3401–3409.
- [45] K. Lee, Y. Wu, and Y. Bresler, “Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization,” *arXiv preprint arXiv:1312.0525*, 2013.
- [46] K. Lee, Y. Li, M. Junge, and Y. Bresler, “Blind recovery of sparse signals from subsampled convolution,” *arXiv preprint arXiv:1511.06149*, 2015.
- [47] K. Lee and M. Junge, “Rip-like properties in subsampled blind deconvolution,” *arXiv preprint arXiv:1511.06146*, 2015.
- [48] V. Cambareri and L. Jacques, “A non-convex blind calibration method for randomised sensing strategies,” *arXiv preprint arXiv:1605.02615*, 2016.
- [49] X. Li, S. Ling, T. Strohmer, and K. Wei, “Rapid, robust, and reliable blind deconvolution via nonconvex optimization,” *arXiv preprint arXiv:1606.04933*, 2016.
- [50] N. Boumal, V. Voroninski, and A. S. Bandeira, “The non-convex burer-monteiro approach works on smooth semidefinite programs,” *arXiv preprint arXiv:1606.04970*, 2016.
- [51] N. Boumal, “Nonconvex phase synchronization,” *arXiv preprint arXiv:1601.06114*, 2016.
- [52] A. S. Bandeira, N. Boumal, and V. Voroninski, “On the low-rank approach for semidefinite programs arising in synchronization and community detection,” *arXiv preprint arXiv:1602.04426*, 2016.

- [53] K. Kawaguchi, "Deep learning without poor local minima," *arXiv preprint arXiv:1605.07110*, 2016.
- [54] D. Soudry and Y. Carmon, "No bad local minima: Data independent training error guarantees for multilayer neural networks," *arXiv preprint arXiv:1605.08361*, 2016.
- [55] P. Jain, C. Jin, S. M. Kakade, and P. Netrapalli, "Computing matrix squareroot via non convex local search," *arXiv preprint arXiv:1507.05854*, 2015.
- [56] S. Bhojanapalli, A. Kyriillidis, and S. Sanghavi, "Dropping convexity for faster semi-definite optimization," *arXiv preprint arXiv:1509.03917*, 2015.
- [57] A. R. Conn, N. I. Gould, and P. L. Toint, *Trust region methods*. Siam, 2000, vol. 1.
- [58] D. C. Sorensen, "Newton's method with a model trust region modification," *SIAM Journal on Numerical Analysis*, vol. 19, no. 2, pp. 409–426, 1982.
- [59] R. Fletcher, "Modified marquardt subroutine for nonlinear least squares," Atomic Energy Research Establishment, Harwell (England), Tech. Rep., 1971.
- [60] W. Gander, "On the linear least squares problem with a quadratic constraint," 1978.
- [61] D. Goldfarb, "Curvilinear path steplength algorithms for minimization which use directions of negative curvature," *Mathematical programming*, vol. 18, no. 1, pp. 31–40, 1980.
- [62] G. P. McCormick, "A modification of armijo's step-size rule for negative curvature," *Mathematical Programming*, vol. 13, no. 1, pp. 111–115, 1977.
- [63] H. Mukai and E. Polak, "A second-order method for unconstrained optimization," *Journal of Optimization Theory and Applications*, vol. 26, no. 4, pp. 501–513, 1978.
- [64] J. J. Moré and D. C. Sorensen, "On the use of directions of negative curvature in a modified newton method," *Mathematical Programming*, vol. 16, no. 1, pp. 1–20, 1979.
- [65] —, "Computing a trust region step," *SIAM Journal on Scientific and Statistical Computing*, vol. 4, no. 3, pp. 553–572, 1983.
- [66] U. Helmke, J. B. Moore, and W. Germany, *Optimization and dynamical systems*. Springer, 1994.
- [67] C. Udriste, *Convex functions and optimization methods on Riemannian manifolds*. Springer Science & Business Media, 1994, vol. 297.
- [68] P.-A. Absil, R. Mahoney, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [69] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998.
- [70] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a matlab toolbox for optimization on manifolds," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1455–1459, 2014.
- [71] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on Riemannian manifolds," *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 303–330, 2007.
- [72] N. Boumal, P.-A. Absil, and C. Cartis, "Global rates of convergence for nonconvex optimization on manifolds," *arXiv preprint arXiv:1605.08101*, 2016.
- [73] C. Cartis, N. I. Gould, and P. L. Toint, "Complexity bounds for second-order optimality in unconstrained optimization," *Journal of Complexity*, vol. 28, no. 1, pp. 93–108, 2012.
- [74] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—online stochastic gradient for tensor decomposition," in *Proceedings of The 28th Conference on Learning Theory*, 2015, pp. 797–842.
- [75] F. Rendl and H. Wolkowicz, "A semidefinite framework for trust region subproblems with applications to large scale minimization," *Mathematical Programming*, vol. 77, no. 1, pp. 273–299, 1997.
- [76] Y. Ye and S. Zhang, "New results on quadratic minimization," *SIAM Journal on Optimization*, vol. 14, no. 1, pp. 245–267, 2003.
- [77] C. Fortin and H. Wolkowicz, "The trust region subproblem and semidefinite programming," *Optimization methods and software*, vol. 19, no. 1, pp. 41–67, 2004.
- [78] E. Hazan and T. Koren, "A linear-time algorithm for trust region problems," *arXiv preprint arXiv:1401.6757*, 2014.
- [79] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.
- [80] K. G. Murty and S. N. Kabadi, "Some np-complete problems in quadratic and nonlinear programming," *Mathematical programming*, vol. 39, no. 2, pp. 117–129, 1987.
- [81] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proceedings of the 25th Annual Conference on Learning Theory*, 2012.
- [82] R. Adamczak, "A note on the sample complexity of the er-spud algorithm by spielman, wang and wright for exact recovery of sparsely used dictionaries," *arXiv preprint arXiv:1601.02049*, 2016.
- [83] K. Luh and V. Vu, "Dictionary learning with few samples and matrix concentration," *arXiv preprint arXiv:1503.08854*, 2015.
- [84] G. W. Stewart and J.-g. Sun, *Matrix Perturbation Theory*. Academic press, 1990.
- [85] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Springer, 2013.
- [86] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Foundations of Computational Mathematics*, vol. 12, no. 4, pp. 389–434, 2012.