

Complete Dictionary Recovery over the Sphere

Ju Sun, Qing Qu, and John Wright

{js4038, qq2105, jw2966}@columbia.edu

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

April 26, 2015

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 the 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 . In contrast, prior results based on efficient algorithms provide recovery guarantees when \mathbf{X}_0 has only $O(\sqrt{n})$ nonzeros per column.

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. To show this apparently hard problem is tractable, we first provide a geometric characterization of the high-dimensional objective landscape, which shows that with high probability there are no “spurious” local minima. This particular geometric structure allows us to design a Riemannian trust region algorithm over the sphere that provably converges to one local minimizer with an arbitrary initialization, despite the presence of saddle points. The geometric approach we develop here may also shed light on other problems arising from nonconvex recovery of structured signals.

Keywords. Dictionary learning, Nonconvex optimization, Spherical constraint, Trust region method, Manifold optimization, Function landscape, Second-order geometry, Inverse problem, Structured signal, Nonlinear approximation

Mathematics Subject Classification. 68P30, 58C05, 94A12, 94A08, 68T05, 90C26, 90C48, 90C55

Acknowledgement. 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, and funding from the Moore and Sloan Foundations. The authors thank Cun Mu and Henry Kuo for discussions related to this project.

Contents

1	Introduction	3
1.1	Theoretical and Algorithmic Challenges	4
1.2	An Intriguing Numerical Experiment with Real Images	4
1.3	Dictionary Recovery and Our Results	6
1.4	Main Ingredients and Innovations	7
1.4.1	A Nonconvex Formulation	7
1.4.2	A Glimpse into High-dimensional Function Landscape	8
1.4.3	A Second-order Algorithm on Manifold: Riemannian Trust Region Method	10
1.5	Prior Arts and Connections	12
1.6	Notations, Organization, and Reproducible Research	15
2	High-dimensional Function Landscapes	16
2.1	Main Geometric Theorems	16
2.2	Useful Technical Lemmas and Proof Ideas for Orthogonal Dictionaries	18
2.3	Extending to Complete Dictionaries	20
3	Finding One Local Minimizer via the Riemannian Trust-Region Method	21
3.1	The Riemannian Trust-Region Algorithm over the Sphere	22
3.2	Main Convergence Results	24
3.3	Useful Technical Results and Proof Ideas for Orthogonal Dictionaries	26
3.3.1	Basic Facts about the Sphere	26
3.3.2	Key Steps towards the Proof	27
3.4	Extending to Convergence for Complete Dictionaries	37
4	Complete Algorithm Pipeline and Main Results	40
4.1	Recovering Orthogonal Dictionaries	41
4.2	Recovering Complete Dictionaries	42
5	Simulations	45
6	Discussion	46
7	Proofs of Main Technical Results for High Dimensional Geometry	47
7.1	Proofs for Section 2.2	47
7.1.1	Proof of Proposition 2.5	47
7.1.2	Proof of Proposition 2.6	52
7.1.3	Proof of Proposition 2.7	54
7.1.4	Proof of Pointwise Concentration Results	58
7.1.5	Proof of Lipschitz Results	61
7.2	Proofs of Theorem 2.1	66
7.3	Proofs for Section 2.3 and Theorem 2.3	69
8	Proof of Convergence for the Trust-Region Algorithm	72
9	Proofs of Technical Results for Section 4	83

Appendices	88
A Technical Tools and Basic Facts Used in Proofs	88
B Auxillary Results for Proofs	94

1 Introduction

Given p signal samples from \mathbb{R}^n , i.e., $\mathbf{Y} \doteq [\mathbf{y}_1, \dots, \mathbf{y}_p]$, is it possible to construct a dictionary $\mathbf{A} \doteq [\mathbf{a}_1, \dots, \mathbf{a}_m]$ with m much smaller than p , such that $\mathbf{Y} \approx \mathbf{A}\mathbf{X}$ and the coefficient matrix \mathbf{X} has as few nonzeros as possible? In other words, this model *dictionary learning* (DL) problem seeks a concise representation for a collection of input signals. Concise signal representations play a central role in compression, and also prove useful for many other important tasks, such as signal acquisition, denoising, and classification.

Traditionally, concise signal representations have relied heavily on explicit analytic bases constructed in nonlinear approximation and harmonic analysis. This constructive approach has proved highly successfully; the numerous theoretical advances in these fields (see, e.g., [DeV98, Tem03, DeV09, Can02, MP10a] for summary of relevant results) provide ever more powerful representations, ranging from the classic Fourier to modern multidimensional, multidirectional, multiresolution bases, including wavelets, curvelets, ridgelets, and so on. However, two challenges confront practitioners in adapting these results to new domains: which function class best describes signals at hand, and consequently which representation is most appropriate. These challenges are coupled, as function classes with known “good” analytic bases are rare.¹

Around 1996, neuroscientists Olshausen and Field discovered that sparse coding, the principle of encoding a signal with few atoms from a learned dictionary, reproduces important properties of the receptive fields of the simple cells that perform early visual processing [OF96, OF97]. The discovery has spurred a flurry of algorithmic developments and successful applications for DL in the past two decades, spanning classical image processing, visual recognition, compressive signal acquisition, and also recent deep architectures for signal classification (see, e.g., [Ela10, MBP14] for review this development).

The learning approach is particularly relevant to modern signal processing and machine learning, which deal with data of huge volume and great variety (e.g., images, audios, graphs, texts, genome sequences, time series, etc). The proliferation of problems and data seems to preclude analytically deriving optimal representations for each new class of data in a timely manner. On the other hand, as datasets grow, learning dictionaries directly from data looks increasingly attractive and promising. When armed with sufficiently many data samples of one signal class, by solving the model DL problem, one would expect to obtain a dictionary that allows sparse representation for the whole class. This hope has been borne out in a number of successful examples [Ela10, MBP14] and theories [MP10b, VMB11, MG13, GJB⁺13].

¹As Donoho et al [DVDD98] put it, “...in effect, uncovering the optimal codebook structure of naturally occurring data involves more challenging empirical questions than any that have ever been solved in empirical work in the mathematical sciences.”

1.1 Theoretical and Algorithmic Challenges

In contrast to the above empirical successes, the theoretical study of dictionary learning is still developing. For applications in which dictionary learning is to be applied in a “hands-free” manner, it is desirable to have efficient algorithms which are guaranteed to perform correctly, when the input data admit a sparse model. There have been several important recent results in this direction, which we will review in Section 1.5, after our sketching main results. Nevertheless, obtaining algorithms that provably succeed under broad and realistic conditions remains an important research challenge.

To understand where the difficulties arise, we can consider a model formulation, in which we attempt to obtain the dictionary \mathbf{A} and coefficients \mathbf{X} which best trade-off sparsity and fidelity to the observed data:

$$\text{minimize}_{\mathbf{A} \in \mathbb{R}^{n \times m}, \mathbf{X} \in \mathbb{R}^{m \times p}} \lambda \|\mathbf{X}\|_1 + \frac{1}{2} \|\mathbf{A}\mathbf{X} - \mathbf{Y}\|_F^2, \text{ subject to } \mathbf{A} \in \mathcal{A}. \quad (1.1)$$

Here, $\|\mathbf{X}\|_1 \doteq \sum_{i,j} |X_{ij}|$ promotes sparsity of the coefficients, $\lambda \geq 0$ trades off the level of coefficient sparsity and quality of approximation, and \mathcal{A} imposes desired structures on the dictionary.

This formulation is nonconvex: the admissible set \mathcal{A} is typically nonconvex (e.g., orthogonal group, matrices with normalized columns)², while the most daunting nonconvexity comes from the bilinear mapping: $(\mathbf{A}, \mathbf{X}) \mapsto \mathbf{A}\mathbf{X}$. Because (\mathbf{A}, \mathbf{X}) and $(\mathbf{A}\mathbf{\Pi}\mathbf{\Sigma}, \mathbf{\Sigma}^{-1}\mathbf{\Pi}^*\mathbf{X})$ result in the same objective value for the conceptual formulation (1.1), where $\mathbf{\Pi}$ is any permutation matrix, and $\mathbf{\Sigma}$ any diagonal matrix with diagonal entries in $\{\pm 1\}$, and $(\cdot)^*$ denotes matrix transpose. Thus, we should expect the problem to have combinatorially many global minima. Because there are multiple isolated global minima, the problem does not appear to be amenable to convex relaxation (see similar discussions in, e.g., [GS10] and [GW11]).³ This contrasts sharply with problems in sparse recovery and compressed sensing, in which simple convex relaxations are often provably effective [DT09, OH10, CLMW11, DGM13, MT14, MHWG13, CRPW12, CSV13, ALMT14, Can14]. Is there any hope to obtain global solutions to the DL problem?

1.2 An Intriguing Numerical Experiment with Real Images

We provide empirical evidence in support of a positive answer to the above question. Specifically, we learn orthogonal bases (orthobases) for real images patches. Orthobases are of interest because typical hand-designed dictionaries such as discrete cosine (DCT) and wavelet bases are orthogonal, and orthobases seem competitive in performance for applications such as image denoising, as compared to overcomplete dictionaries [BCJ13]⁴.

²For example, in nonlinear approximation and harmonic analysis, orthonormal basis or (tight-)frames are preferred; to fix the scale ambiguity discussed in the text, a common practice is to require that \mathbf{A} to be column-normalized. There is no obvious reason to believe that convexifying these constraint sets would leave the optima unchanged. For example, the convex hull of the orthogonal group O_n is the operator norm ball $\{\mathbf{X} \in \mathbb{R}^{n \times n} : \|\mathbf{X}\| \leq 1\}$. If there are no effective symmetry breaking constraints, any convex objective function tends to have minimizers inside the ball, which obviously will not be orthogonal matrices. Other ideas such as lifting may not play together with the objective function, nor yield tight relaxations (see, e.g., [BKS13, BR14]).

³Semidefinite programming (SDP) lifting may be one useful general strategy to convexify bilinear inverse problems, see, e.g., [ARR14, CM14]. However, for problems with general nonlinear constraints, it is unclear whether the lifting always yield tight relaxation, consider, e.g., [BKS13, BR14] again.

⁴See Section 1.3 for more detailed discussions of this point. [LGBB05] also gave motivations and algorithms for learning (union of) orthobases as dictionaries.




Figure 1: Alternating direction method for (1.2) on uncompressed real images seems to always produce the same solution! Top: Each image is 512×512 in resolution and encoded in the uncompressed pgm format (uncompressed images to prevent possible bias towards standard bases used for compression, such as DCT or wavelet bases). Each image is evenly divided into 8×8 non-overlapping image patches (4096 in total), and these patches are all vectorized and then stacked as columns of the data matrix \mathbf{Y} . **Bottom:** Given each \mathbf{Y} , we solve (1.2) 100 times with independent and randomized (uniform over the orthogonal group) initialization \mathbf{A}_0 . The plots show the values of $\|\mathbf{A}_\infty^* \mathbf{Y}\|_1$ across the independent repetitions. They are virtually the same and the relative differences are less than 10^{-3} !

We divide a given greyscale image into 8×8 non-overlapping patches, which are converted into 64-dimensional vectors and stacked column-wise into a data matrix \mathbf{Y} . Specializing (1.1) to this setting, we obtain the optimization problem:

$$\text{minimize}_{\mathbf{A} \in \mathbb{R}^{n \times n}, \mathbf{X} \in \mathbb{R}^{n \times p}} \lambda \|\mathbf{X}\|_1 + \frac{1}{2} \|\mathbf{A}\mathbf{X} - \mathbf{Y}\|_F^2, \text{ subject to } \mathbf{A} \in O_n. \quad (1.2)$$

To derive a concrete algorithm for (1.2), one can deploy the alternating direction method (ADM)⁵, i.e., alternately minimizing the objective function with respect to (w.r.t.) one variable while fixing the other. The iteration sequence actually takes very simple form: for $k = 1, 2, 3, \dots$,

$$\mathbf{X}_k = \mathcal{S}_\lambda [\mathbf{A}_{k-1}^* \mathbf{Y}], \quad \mathbf{A}_k = \mathbf{U}\mathbf{V}^* \text{ for } \mathbf{U}\mathbf{D}\mathbf{V}^* = \text{SVD}(\mathbf{Y}\mathbf{X}_k^*)$$

where $\mathcal{S}_\lambda [\cdot]$ denotes the well-known soft-thresholding operator acting elementwise on matrices, i.e., $\mathcal{S}_\lambda [x] \doteq \text{sign}(x) \max(|x| - \lambda, 0)$ for any scalar x .

Figure 1 shows what we obtained using the simple ADM algorithm, with *independent and randomized initializations*:

⁵This method is also called alternating minimization or (block) coordinate descent method. see, e.g., [BT89, Tse01] for classic results and [ABRS10, BST14] for several interesting recent developments.

The algorithm seems to always produce the same solution, regardless of the initialization.

This observation implies the heuristic ADM algorithm may *always converge to one global minimizer!*⁶ Equally surprising is that the phenomenon has been observed on real images⁷. One may imagine only random data typically have “favorable” structures; in fact, almost all existing theories for DL pertain only to random data [SWW12, AAJ⁺13, AGM13, AAN13, ABGM14, AGMM15].

1.3 Dictionary Recovery and Our Results

In this paper, we take a step towards explaining the surprising effectiveness of nonconvex optimization heuristics for DL. We focus on the *dictionary recovery* (DR) setting: given a data matrix \mathbf{Y} generated as $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$, where $\mathbf{A}_0 \in \mathcal{A} \subseteq \mathbb{R}^{n \times m}$ and $\mathbf{X}_0 \in \mathbb{R}^{m \times p}$ is “reasonably sparse”, try to recover \mathbf{A}_0 and \mathbf{X}_0 . Here recovery means to return any pair $(\mathbf{A}_0 \mathbf{\Pi} \mathbf{\Sigma}, \mathbf{\Sigma}^{-1} \mathbf{\Pi}^* \mathbf{X}_0)$, where $\mathbf{\Pi}$ is a permutation matrix and $\mathbf{\Sigma}$ is a nonsingular diagonal matrix, i.e., recovering up to sign, scale, and permutation.

To define a reasonably simple and structured problem, we make the following assumptions:

- The target dictionary \mathbf{A}_0 is complete, i.e., square and invertible ($m = n$). In particular, this class includes orthogonal dictionaries. Admittedly overcomplete dictionaries tend to be more powerful for modeling and to allow sparser representations. Nevertheless, most classic hand-designed dictionaries in common use are orthogonal. Orthobases are competitive in performance for certain tasks such as image denoising [BCJ13], and admit faster algorithms for learning and encoding.⁸
- The coefficient matrix \mathbf{X}_0 follows the Bernoulli-Gaussian (BG) model with rate θ : $[\mathbf{X}_0]_{ij} = \Omega_{ij} V_{ij}$, with $\Omega_{ij} \sim \text{Ber}(\theta)$ and $V_{ij} \sim \mathcal{N}(0, 1)$, where all the different random variables are mutually independent. We write compactly $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$.

We prove the following result:

Theorem 1.1 (Informal statement of our results) *For any $\theta \in (0, 1/3)$, given $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$ with \mathbf{A}_0 a complete dictionary and $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$, there is a polynomial time algorithm that recovers \mathbf{A}_0 and \mathbf{X}_0 with high probability (at least $1 - O(p^{-6})$) whenever $p \geq p_*(n, 1/\theta, \kappa(\mathbf{A}_0), 1/\mu)$ for a fixed polynomial $p_*(\cdot)$, where $\kappa(\mathbf{A}_0)$ is the condition number of \mathbf{A}_0 and μ is a parameter that can be set as $cn^{-5/4}$ for a fixed positive numerical constant c .*

Obviously, even if \mathbf{X}_0 is known, one needs $p \geq n$ to make the identification problem well posed. Under our particular probabilistic model, a simple coupon collection argument implies that one needs $p \geq \Omega(\frac{1}{\theta} \log n)$ to ensure all atoms in \mathbf{A}_0 are observed with high probability (w.h.p.). To

⁶Technically, the converge to global solutions is surprising because even convergence of ADM to critical points is atypical, see, e.g., [ABRS10, BST14] and references therein. Section 6 includes more detailed discussions on this point.

⁷Actually the same phenomenon is also observed for simulated data when the coefficient matrix obeys the Bernoulli-Gaussian model, which is defined later. The result on real images supports that previously claimed empirical successes over two decades may be non-incidental.

⁸Empirically, there is no systematic evidence supporting that overcomplete dictionaries are strictly necessary for good performance in all published applications (though [OF97] argues for the necessity from neuroscience perspective). Some of the ideas and tools developed here for complete dictionaries may also apply to certain classes of structured overcomplete dictionaries, such as tight frames. See Section 6 for relevant discussion.

ensure that an efficient algorithm exists may demand more. Our result implies when p is polynomial in n , $1/\theta$ and $\kappa(\mathbf{A}_0)$, recovery with efficient algorithm is possible.

The parameter θ controls the sparsity level of \mathbf{X}_0 . Intuitively, the recovery problem is easy for small θ and becomes harder for large θ .⁹ It is perhaps surprising that an efficient algorithm can succeed up to constant θ , i.e., linear sparsity in \mathbf{X}_0 . Compared to the case when \mathbf{A}_0 is known, there is only at most a constant gap in the sparsity level one can deal with.

For DL, our result gives the first efficient algorithm that provably recovers complete \mathbf{A}_0 and \mathbf{X}_0 when \mathbf{X}_0 has $O(n)$ nonzeros per column under appropriate probability model. Section 1.5 provides detailed comparison of our result with other recent recovery results for complete and overcomplete dictionaries.

1.4 Main Ingredients and Innovations

In this section we describe three main ingredients that we use to obtain the stated result.

1.4.1 A Nonconvex Formulation

Since $\mathbf{Y} = \mathbf{A}_0\mathbf{X}_0$ and \mathbf{A}_0 is complete, $\text{row}(\mathbf{Y}) = \text{row}(\mathbf{X}_0)$ ($\text{row}(\cdot)$ denotes the row space of a matrix) and hence rows of \mathbf{X}_0 are sparse vectors in the known (linear) subspace $\text{row}(\mathbf{Y})$. We can use this fact to first recover the rows of \mathbf{X}_0 , and subsequently recover \mathbf{A}_0 by solving a system of linear equations. In fact, for $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$, rows of \mathbf{X}_0 are the n sparsest vectors (directions) in $\text{row}(\mathbf{Y})$ w.h.p. whenever $p \geq \Omega(n \log n)$ [SWW12]. Thus one might try to recover rows of \mathbf{X}_0 by solving

$$\text{minimize } \|\mathbf{q}^*\mathbf{Y}\|_0 \quad \text{subject to } \mathbf{q} \neq \mathbf{0}. \quad (1.3)$$

The objective is discontinuous, and the domain is an open set. In particular, the homogeneous constraint is nonconventional and tricky to deal with. Since the recovery is up to scale, one can remove the homogeneity by fixing the scale of \mathbf{q} . Known relaxations [SWW12, DH14] fix the scale by setting $\|\mathbf{q}^*\mathbf{Y}\|_\infty = 1$, where $\|\cdot\|_\infty$ is the elementwise ℓ^∞ norm. The optimization problem reduces to a sequence of convex programs, which recover $(\mathbf{A}_0, \mathbf{X}_0)$ for very sparse \mathbf{X}_0 , but provably break down when columns of \mathbf{X}_0 has more than $O(\sqrt{n})$ nonzeros, or $\theta \geq \Omega(1/\sqrt{n})$. Inspired by our previous image experiment, we work with a *nonconvex* alternative¹⁰:

$$\text{minimize } f(\mathbf{q}; \hat{\mathbf{Y}}) \doteq \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{q}^*\hat{\mathbf{y}}_k), \quad \text{subject to } \|\mathbf{q}\| = 1, \quad (1.4)$$

where $\hat{\mathbf{Y}} \in \mathbb{R}^{n \times p}$ is a proxy for \mathbf{Y} (i.e., after appropriate processing), k indexes columns of $\hat{\mathbf{Y}}$, and $\|\cdot\|$ is the usual ℓ^2 norm for vectors. Here $h_\mu(\cdot)$ is chosen to be a convex smooth approximation to $|\cdot|$, namely,

$$h_\mu(z) = \mu \log \left(\frac{\exp(z/\mu) + \exp(-z/\mu)}{2} \right) = \mu \log \cosh(z/\mu), \quad (1.5)$$

⁹Indeed, when θ is small enough such that columns of \mathbf{X}_0 are predominately 1-sparse, one directly observes scaled versions of the atoms (i.e., columns of \mathbf{X}_0); when \mathbf{X}_0 is fully dense corresponding to $\theta = 1$, recovery is never possible as one can easily find another complete \mathbf{A}'_0 and fully dense \mathbf{X}'_0 such that $\mathbf{Y} = \mathbf{A}'_0\mathbf{X}'_0$ with \mathbf{A}'_0 not equivalent to \mathbf{A}_0 .

¹⁰A similar formulation has been proposed in [ZP01] in the context of blind source separation; see also [QSW14].




Figure 2: Why is dictionary learning over \mathbb{S}^{n-1} tractable? Assume the target dictionary \mathbf{A}_0 is orthogonal. **Left:** Large sample objective function $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q})]$. The only local minima are the columns of \mathbf{A}_0 and their negatives. **Center:** the same function, visualized as a height above the plane \mathbf{a}_1^\perp (\mathbf{a}_1 is the first column of \mathbf{A}_0). **Right:** Around the optimum, the function exhibits a small region of positive curvature, a region of large gradient, and finally a region in which the direction away from \mathbf{a}_1 is a direction of negative curvature.

which is infinitely differentiable and μ controls the smoothing level.¹¹ The spherical constraint is nonconvex. Hence, a-priori, it is unclear whether (1.4) admits efficient algorithms that attain global optima. Surprisingly, simple descent algorithms for (1.4) exhibit very striking behavior: on many practical numerical examples¹², they appear to produce global solutions. Our next section will uncover interesting geometrical structures underlying the phenomenon.

1.4.2 A Glimpse into High-dimensional Function Landscape

For the moment, suppose \mathbf{A}_0 is orthogonal, and take $\hat{\mathbf{Y}} = \mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$ in (1.4). Figure 2 (left) plots $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q}; \mathbf{Y})]$ over $\mathbf{q} \in \mathbb{S}^2$ ($n = 3$). Remarkably, $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q}; \mathbf{Y})]$ has no spurious local minima. In fact, every local minimizer $\hat{\mathbf{q}}$ produces a row of \mathbf{X}_0 : $\hat{\mathbf{q}}^* \mathbf{Y} = \alpha e_i^* \mathbf{X}_0$ for some $\alpha \neq 0$.

To better illustrate the point, we take the particular case $\mathbf{A}_0 = \mathbf{I}$ and project the upper hemisphere above the equatorial plane e_3^\perp onto e_3^\perp . The projection is bijective and we equivalently define a reparameterization $g : e_3^\perp \mapsto \mathbb{R}$ of f . Figure 2 (center) plots the graph of g . Obviously the only local minimizers are $\mathbf{0}, \pm e_1, \pm e_2$, and they are also global minimizers. Moreover, the apparent nonconvex landscape has interesting structures around $\mathbf{0}$: when moving away from $\mathbf{0}$, one sees successively a strongly convex region, a nonzero gradient region, and a region where at each point one can always find a direction of negative curvature, as shown schematically in Figure 2 (right). This geometry implies that at any nonoptimal point, there is always at least one direction of descent. Thus, any algorithm that can take advantage of the descent directions will likely converge to one global minimizer, irrespective of initialization.

Two challenges stand out when implementing this idea. For geometry, one has to show similar structure exists for general complete \mathbf{A}_0 , in high dimensions ($n \geq 3$), when the number of observations p is finite (vs. the expectation in the experiment). For algorithms, we need to be able to

¹¹In fact, there is nothing special about this choice and we believe that any valid smooth (twice continuously differentiable) approximation to $|\cdot|$ would work and yield qualitatively similar results. We also have some preliminary results showing the latter geometric picture remains the same for certain nonsmooth functions, such as a modified version of the Huber function, though the analysis involves handling a different set of technical subtleties. The algorithm also needs additional modifications.

¹²... not restricted to the model we assume here for \mathbf{A}_0 and \mathbf{X}_0 .

take advantage of this structure without knowing \mathbf{A}_0 ahead of time. In Section 1.4.3, we describe a Riemannian trust region method which addresses the latter challenge.

Geometry for orthogonal \mathbf{A}_0 . In this case, we take $\widehat{\mathbf{Y}} = \mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$. Since $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0) = f(\mathbf{A}_0^* \mathbf{q}; \mathbf{X}_0)$, the landscape of $f(\mathbf{q}; \mathbf{A}_0 \mathbf{X}_0)$ is simply a rotated version of that of $f(\mathbf{q}; \mathbf{X}_0)$, i.e., when $\mathbf{A}_0 = \mathbf{I}$. Hence we will focus on the case when $\mathbf{A}_0 = \mathbf{I}$. Among the $2n$ symmetric sections of \mathbb{S}^{n-1} centered around the signed basis vectors $\pm e_1, \dots, \pm e_n$, we work with the symmetric section around e_n as an example. The result will carry over to all sections with the same argument; together this provides a complete characterization of the function $f(\mathbf{q}; \mathbf{X}_0)$ over \mathbb{S}^{n-1} .

We again invoke the projection trick described above, this time onto the equatorial plane e_n^\perp . This can be formally captured by the reparameterization mapping:

$$\mathbf{q}(\mathbf{w}) = \left(\mathbf{w}, \sqrt{1 - \|\mathbf{w}\|^2} \right), \quad \mathbf{w} \in \mathbb{B}^{n-1}, \quad (1.6)$$

where \mathbf{w} is the new variable in $e_n^\perp \cap \mathbb{B}^{n-1}$ and \mathbb{B}^{n-1} is the unit ball in \mathbb{R}^{n-1} . We first study the composition $g(\mathbf{w}; \mathbf{X}_0) \doteq f(\mathbf{q}(\mathbf{w}); \mathbf{X}_0)$ over the set

$$\Gamma \doteq \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\}. \quad (1.7)$$

It can be verified the section we chose to work with is contained in this set¹³.

Our analysis characterizes the properties of $g(\mathbf{w}; \mathbf{X}_0)$ by studying three quantities

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0), \quad \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|}, \quad \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2}$$

respectively over three consecutive regions moving away from the origin, corresponding to the three regions in Figure 2 (right). In particular, through typical expectation-concentration style argument, we show that there exists a positive constant c such that

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0) \succeq \frac{1}{\mu} c \theta \mathbf{I}, \quad \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \geq c \theta, \quad \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -c \theta \quad (1.8)$$

over the respective regions w.h.p., confirming our low-dimensional observations described above. In particular, the favorable structure we observed for $n = 3$ persists in high dimensions, w.h.p., even when p is large *yet finite*, for the case \mathbf{A}_0 is orthogonal. Moreover, the local minimizer of $g(\mathbf{w}; \mathbf{X}_0)$ over Γ is very close to $\mathbf{0}$, within a distance of $O(\mu)$.

Geometry for complete \mathbf{A}_0 . For general complete dictionaries \mathbf{A}_0 , we hope that the function f retains the nice geometric structure discussed above. We can ensure this by “preconditioning” \mathbf{Y} such that the output looks as if being generated from a certain orthogonal matrix, possibly plus

¹³Indeed, if $\langle \mathbf{q}, e_n \rangle \geq |\langle \mathbf{q}, e_i \rangle|$ for any $i \neq n$, $1 - \|\mathbf{w}\|^2 = q_n^2 \geq 1/n$, implying $\|\mathbf{w}\|^2 \leq \frac{n-1}{n} < \frac{4n-1}{4n}$. The reason we have defined an open set instead of a closed (compact) one is to avoid potential trivial local minimizers located on the boundary.

a small perturbation. We can then argue that the perturbation does not significantly affect the properties of the graph of the objective function. Write

$$\bar{\mathbf{Y}} = \left(\frac{1}{p\theta} \mathbf{Y} \mathbf{Y}^* \right)^{-1/2} \mathbf{Y}. \quad (1.9)$$

Note that for $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$, $\mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^*] / (p\theta) = \mathbf{I}$. Thus, one expects $\frac{1}{p\theta} \mathbf{Y} \mathbf{Y}^* = \frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*$ to behave roughly like $\mathbf{A}_0 \mathbf{A}_0^*$ and hence $\bar{\mathbf{Y}}$ to behave like

$$(\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \mathbf{A}_0 \mathbf{X}_0 = \mathbf{U} \mathbf{V}^* \mathbf{X}_0 \quad (1.10)$$

where we write the SVD of \mathbf{A}_0 as $\mathbf{A}_0 = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$. It is easy to see $\mathbf{U} \mathbf{V}^*$ is an orthogonal matrix. Hence the preconditioning scheme we have introduced is technically sound.

Our analysis shows that $\bar{\mathbf{Y}}$ can be written as

$$\bar{\mathbf{Y}} = \mathbf{U} \mathbf{V}^* \mathbf{X}_0 + \mathbf{\Xi} \mathbf{X}_0, \quad (1.11)$$

where $\mathbf{\Xi}$ is a matrix with small magnitude. Simple perturbation argument shows that the constant c in (1.8) is at most shrunk to $c/2$ for all \mathbf{w} when p is sufficiently large. Thus, the qualitative aspects of the geometry have not been changed by the perturbation.

1.4.3 A Second-order Algorithm on Manifold: Riemannian Trust Region Method

We do not know \mathbf{A}_0 ahead of time, so our algorithm needs to take advantage of the structure described above without knowledge of \mathbf{A}_0 . Intuitively, this seems possible as the descent direction in the \mathbf{w} space appears to also be a local descent direction for f over the sphere. Another issue is that although the optimization problem has no spurious local minima, it does have many saddle points (Figure 2). We can use second-order information to guarantee to escape saddle points. We derive an algorithm based on the Riemannian trust region method (TRM) [ABG07, AMS09] over the sphere for this purpose.

For a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and an unconstrained optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}),$$

typical (second-order) TRM proceeds by successively forming second-order approximations to f at the current iterate,

$$\hat{f}(\boldsymbol{\delta}; \mathbf{x}^{(k-1)}) \doteq f(\mathbf{x}^{(k-1)}) + \nabla^* f(\mathbf{x}^{(k-1)}) \boldsymbol{\delta} + \frac{1}{2} \boldsymbol{\delta}^* \mathbf{Q}(\mathbf{x}^{(k-1)}) \boldsymbol{\delta}, \quad (1.12)$$

where $\mathbf{Q}(\mathbf{x}^{(k-1)})$ is a proxy for the Hessian matrix $\nabla^2 f(\mathbf{x}^{(k-1)})$, which encodes the second-order geometry. The next movement direction is determined by seeking a minimum of $\hat{f}(\boldsymbol{\delta}; \mathbf{x}^{(k-1)})$ over a small region, normally a norm ball $\|\boldsymbol{\delta}\|_p \leq \Delta$, called the trust region, inducing the well studied trust-region subproblem:

$$\boldsymbol{\delta}^{(k)} \doteq \arg \min_{\boldsymbol{\delta} \in \mathbb{R}^n, \|\boldsymbol{\delta}\|_p \leq \Delta} \hat{f}(\boldsymbol{\delta}; \mathbf{x}^{(k-1)}), \quad (1.13)$$

where Δ is called the trust-region radius that controls how far the movement can be made. A ratio

$$\rho_k \doteq \frac{f(\mathbf{x}^{(k-1)}) - f(\mathbf{x}^{(k-1)} + \boldsymbol{\delta}^{(k)})}{\widehat{f}(\mathbf{0}) - \widehat{f}(\boldsymbol{\delta}^{(k-1)})} \quad (1.14)$$

is defined to measure the progress and typically the radius Δ is updated dynamically according to ρ_k to adapt to the local function behavior. Detailed introductions to the classical TRM can be found in the texts [CGT00a, NW06].




Figure 3: 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} .

To generalize the idea to smooth manifolds, one natural choice is to form the approximation over the tangent spaces [ABG07, AMS09]. Specific to our spherical manifold, for which the tangent space at an iterate $\mathbf{q}^{(k)} \in \mathbb{S}^{n-1}$ is $T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1} \doteq \{\mathbf{v} : \mathbf{v}^* \mathbf{q}^{(k)} = 0\}$ (see Figure 3), we work with a “quadratic” approximation $\widehat{f} : T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1} \mapsto \mathbb{R}$ defined as

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{q}^{(k)}) \doteq f(\mathbf{q}^{(k)}) + \langle \nabla f(\mathbf{q}^{(k)}), \boldsymbol{\delta} \rangle + \frac{1}{2} \boldsymbol{\delta}^* \left(\nabla^2 f(\mathbf{q}^{(k)}) - \langle \nabla f(\mathbf{q}^{(k)}), \mathbf{q}^{(k)} \rangle \mathbf{I} \right) \boldsymbol{\delta}. \quad (1.15)$$

To interpret this approximation, let $\mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \doteq \left(\mathbf{I} - \mathbf{q}^{(k)} (\mathbf{q}^{(k)})^* \right)$ be the orthoprojector onto $T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}$ and write (3.2) into an equivalent form:

$$\begin{aligned} \widehat{f}(\boldsymbol{\delta}; \mathbf{q}^{(k)}) &\doteq f(\mathbf{q}^{(k)}) + \langle \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \nabla f(\mathbf{q}^{(k)}), \boldsymbol{\delta} \rangle \\ &\quad + \frac{1}{2} \boldsymbol{\delta}^* \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \left(\nabla^2 f(\mathbf{q}^{(k)}) - \langle \nabla f(\mathbf{q}^{(k)}), \mathbf{q}^{(k)} \rangle \mathbf{I} \right) \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \boldsymbol{\delta}. \end{aligned}$$

The two terms

$$\begin{aligned} \text{grad} f \left(\mathbf{q}^{(k)} \right) &\doteq \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \nabla f(\mathbf{q}^{(k)}), \\ \text{Hess} f \left(\mathbf{q}^{(k)} \right) &\doteq \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \left(\nabla^2 f(\mathbf{q}^{(k)}) - \langle \nabla f(\mathbf{q}^{(k)}), \mathbf{q}^{(k)} \rangle \mathbf{I} \right) \mathcal{P}_{T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}} \end{aligned}$$

are the Riemannian gradient and Riemannian Hessian of f w.r.t. \mathbb{S}^{n-1} , respectively [ABG07, AMS09]; the above approximation is reminiscent of the usual quadratic approximation described in (1.12).

Then the Riemannian trust-region subproblem is

$$\min_{\boldsymbol{\delta} \in T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}, \|\boldsymbol{\delta}\| \leq \Delta} \widehat{f}(\boldsymbol{\delta}; \mathbf{q}^{(k)}), \quad (1.16)$$

where we take the simple ℓ^2 norm ball for the trust region. This can be transformed into a classical trust region subproblem: indeed, taking any orthonormal basis $\mathbf{U}_{\mathbf{q}^{(k)}}$ for $T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}$, the above problem is equivalent to

$$\min_{\|\boldsymbol{\xi}\|\leq\Delta} \widehat{f}\left(\mathbf{U}_{\mathbf{q}^{(k)}}\boldsymbol{\xi}, \mathbf{q}^{(k)}\right), \quad (1.17)$$

where the objective is quadratic in $\boldsymbol{\xi}$. This is the classical trust region problem (with ℓ^2 norm ball constraint) that admits very efficient numerical algorithms [MS83, HK14]. Once we obtain the minimizer $\boldsymbol{\xi}_*$, we set $\boldsymbol{\delta}_* = \mathbf{U}\boldsymbol{\xi}_*$, which solves (1.16).

One additional issue as compared to the Euclidean setting is that now $\boldsymbol{\delta}_*$ is one vector in the tangent space and additive update leads to a point outside the sphere. We resort to the natural exponential map to pull the tangent vector to a point on the sphere:

$$\mathbf{q}^{(k+1)} \doteq \exp_{\mathbf{q}^{(k)}}(\boldsymbol{\delta}_*) = \mathbf{q}^{(k)} \cos \|\boldsymbol{\delta}_*\| + \frac{\boldsymbol{\delta}_*}{\|\boldsymbol{\delta}_*\|} \sin \|\boldsymbol{\delta}_*\|. \quad (1.18)$$

As seen from Figure 3, the movement to the next iterate is “along the direction”¹⁴ of $\boldsymbol{\delta}_*$ while staying over the sphere.

Using the above geometric characterizations, we prove that w.h.p., the algorithm converges to a local minimizer when the parameter Δ is sufficiently small¹⁵. In particular, we show that (1) the trust region step induces at least a fixed amount of decrease to the objective value in the negative curvature and nonzero gradient region; (2) the trust region iterate sequence will eventually move to and stay in the strongly convex region, and converge to the local minimizer contained in the region with an asymptotic quadratic rate. In short, the geometric structure implies that from *any initialization*, the iterate sequence converges to a close approximation to the target solution in a polynomial number of steps.

1.5 Prior Arts and Connections

It is far too ambitious to include here a comprehensive review of the exciting developments of DL algorithms and applications after the pioneer work [OF96]. We refer the reader to Chapter 12 - 15 of the book [Ela10] and the survey paper [MBP14] for summaries of relevant developments in image analysis and visual recognition. In the following, we focus on reviewing recent developments on the theoretical side of dictionary learning, and draw connections to problems and techniques that are relevant to the current work.

Theoretical Dictionary Learning. The theoretical study of DL in the recovery setting started only very recently. [AEB06] was the first to provide an algorithmic procedure to correctly extract the generating dictionary. The algorithm requires exponentially many samples and has exponential running time; see also [HS11]. Subsequent work [GS10, GW11, Sch14a, Sch14b, Sch15] studied when the target dictionary is a local optimum of natural recovery criteria. These meticulous analyses show that polynomially many samples are sufficient to ensure local correctness under natural assumptions. However, these results do not imply that one can design efficient algorithms to obtain the desired local optimum and hence the dictionary.

¹⁴Technically, moving along the geodesic whose velocity at time zero is $\boldsymbol{\delta}_*$.

¹⁵For simplicity of analysis, we have assumed Δ is fixed throughout the analysis. In practice, dynamic updates to Δ lead to faster convergence.

[SWW12] initiated the on-going research effort to provide efficient algorithms that globally solve DR. They showed that one can recover a complete dictionary \mathbf{A}_0 from $\mathbf{Y} = \mathbf{A}_0\mathbf{X}_0$ by solving a certain sequence of linear programs, when \mathbf{X}_0 is a sparse random matrix with $O(\sqrt{n})$ nonzeros per column. [AAJ⁺13, AAN13] and [AGM13, AGMM15] give efficient algorithms that provably recover overcomplete ($m \geq n$) and incoherent dictionaries, based on a combination of {clustering or spectral initialization} and local refinement. These algorithms again succeed when \mathbf{X}_0 has $\tilde{O}(\sqrt{n})$ ¹⁶ nonzeros per column. Recent work, including [ABGM14] and [BKS14], guarantees recovery with (almost) $O(n)$ nonzeros per column, but runs in super-polynomial (quasipolynomial) time. By comparison, we give the first *polynomial-time* algorithm that provably recovers complete dictionary \mathbf{A}_0 when \mathbf{X}_0 has $O(n)$ nonzeros per column.

Aside from efficient recovery, other theoretical work on DL includes results on identifiability [AEB06, HS11], generalization bounds [MP10b, VMB11, MG13, GJB⁺13], and noise stability [GJB14].

Finding Sparse Vectors in a Linear Subspace. We have followed [SWW12] and cast the core problem as finding the sparsest vectors in a given linear subspace, which is also of independent interest. Under a planted sparse model¹⁷, [DH14] shows solving a sequence of linear programs similar to [SWW12] can recover sparse vectors with sparsity up to $O(p/\sqrt{n})$, sublinear in the vector dimension. [QSW14] improved the recovery limit to $O(p)$ by solving a nonconvex spherical constrained problem similar to (1.4)¹⁸ via an ADM algorithm. The idea of seeking rows of \mathbf{X}_0 sequentially by solving the above core problem sees precursors in [ZP01] for blind source separation, and [GN10] for matrix sparsification. [ZP01] also proposed a nonconvex optimization similar to (1.4) here and that employed in [QSW14].

Nonconvex Optimization Problems. For other nonconvex optimization problems of recovery of structured signals¹⁹, including low-rank matrix completion/recovery [KMO10, JNS13, Har14, HW14, NNS⁺14, JN14], phase retrieval [NJS13, CLS14], tensor recovery [JO14, AGJ14], mixed regression [YCS13, LWB13], structured element pursuit [QSW14], and recovery of simultaneously structured signals [LWB13], the initialization plus local refinement strategy adopted in theoretical DL [AAJ⁺13, AAN13, AGM13, AGMM15, ABGM14] is also crucial: nearness to the target solution enables exploiting the local geometry of the target to analyze the local refinement.²⁰ By comparison, we provide a complete characterization of the global geometry, which admits efficient algorithms without any special initialization. The idea of separating the geometric analysis and algorithmic design may also prove valuable for other nonconvex problems discussed above.

Optimization over Riemannian Manifolds. Our trust-region algorithm on the sphere builds on the extensive research efforts to generalize Euclidean numerical algorithms to (Riemannian) manifold settings. We refer the reader to the monographs [Udr94, HMG94, AMS09] for survey

¹⁶The \tilde{O} suppresses some logarithm factors.

¹⁷... where one sparse vector embedded in an otherwise random subspace.

¹⁸The only difference is that they chose to work with the Huber function as a proxy of the $\|\cdot\|_1$ function.

¹⁹This is a body of recent work studying nonconvex recovery up to statistical precision, including, e.g., [LW11, LW13, WLL14, BWY14, WGNL14, LW14, Loh15, SLLC15].

²⁰The powerful framework [ABRS10, BST14] to establish local convergence of ADM algorithms to critical points applies to DL/DR also, see, e.g., [BJS14, BQJ14, BJS14]. However, these results do not guarantee to produce global optima.

of developments in this field. In particular, [EAS98] developed Newton and conjugate-gradient methods for the Stiefel manifolds, of which the spherical manifold is a special case. [ABG07] generalized the trust-region methods to Riemannian manifolds. We cannot, however, adopt the existing convergence results that concern either global convergence (convergence to critical points) or local convergence (convergence to a local minimum within a radius). The particular geometric structure forces us to piece together different arguments to obtain the global result.

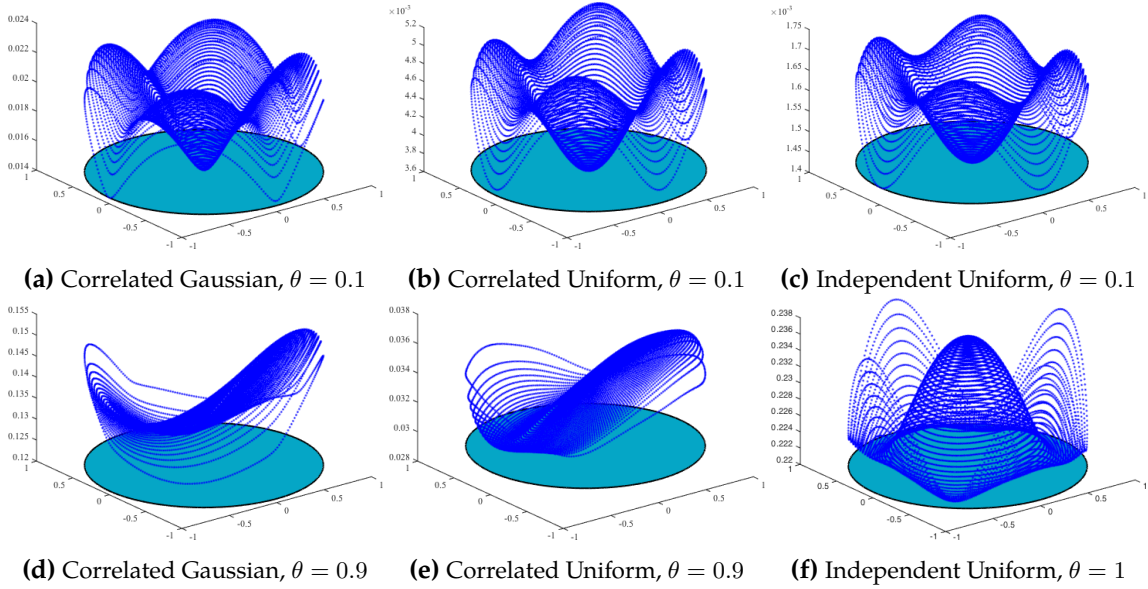


Figure 4: Asymptotic function landscapes when rows of X_0 are not independent. W.l.o.g., we again assume $A_0 = I$. In (a) and (d), $X_0 = \Omega \odot V$, with $\Omega \sim_{i.i.d.} \text{Ber}(\theta)$ and columns of X_0 i.i.d. Gaussian vectors obeying $v_i \sim \mathcal{N}(0, \Sigma^2)$ for symmetric Σ with 1's on the diagonal and i.i.d. off-diagonal entries distributed as $\mathcal{N}(0, \sqrt{2}/20)$. Similarly, in (b) and (e), $X_0 = \Omega \odot W$, with $\Omega \sim_{i.i.d.} \text{Ber}(\theta)$ and columns of X_0 i.i.d. vectors generated as $w_i = \Sigma u^i$ with $u_i \sim_{i.i.d.} \text{Uniform}[-0.5, 0.5]$. For comparison, in (c) and (f), $X_0 = \Omega \odot W$ with $\Omega \sim_{i.i.d.} \text{Ber}(\theta)$ and $W \sim_{i.i.d.} \text{Uniform}[-0.5, 0.5]$. Here \odot denote the elementwise product, and the objective function is still based on the log cosh function as in (1.4).

Independent Component Analysis (ICA) and Other Matrix Factorization Problems. DL can also be considered in the general framework of matrix factorization problems, which encompass the classic principal component analysis (PCA), ICA, and clustering, and more recent problems such as nonnegative matrix factorization (NMF), multi-layer neural nets (deep learning architectures). Most of these problems are NP-hard. Identifying tractable cases of practical interest and providing provable efficient algorithms are subject of on-going research endeavors; see, e.g., recent progresses on NMF [AGKM12], and learning deep neural nets [ABGM13, SA14, NP13, LSSS14].

ICA factors a data matrix Y as $Y = AX$ such that A is square and rows of X are as independent as possible [HO00, HO01]. In theoretical study of the recovery problem, it is often assumed that rows of X_0 are (weakly) independent (see, e.g., [Com94, FJK96, AGMS12]). Our i.i.d. probability model on X_0 implies rows of X_0 are independent, aligning our problem perfectly with the ICA problem. More interestingly, the log cosh objective we analyze here was proposed as a general-purpose *contrast function* in ICA that has not been thoroughly analyzed [Hyv99], and algorithm and

analysis with another popular contrast function, the fourth-order cumulants, indeed overlap with ours considerably [FJK96, AGMS12]²¹. While this interesting connection potentially helps port our analysis to ICA, it is a fundamental question to ask what is playing the vital role for DR, sparsity or independence.

Figure 4 helps shed some light in this direction, where we again plot the asymptotic objective landscape with the natural reparameterization as in Section 1.4.2. From the left and central panels, it is evident even without independence, \mathbf{X}_0 with sparse columns induces the familiar geometric structures we saw in Figure 2; such structures are broken when the sparsity level becomes large. We believe all our later analyses can be generalized to the correlated cases we experimented with. On the other hand, from the right panel²², it seems with independence, the function landscape undergoes a transition as sparsity level grows - target solution goes from minimizers of the objective to the maximizers of the objective. Without adequate knowledge of the true sparsity, it is unclear whether one would like to minimize or maximize the objective.²³ This suggests sparsity, instead of independence, makes our current algorithm for DR work.

1.6 Notations, Organization, 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, n, k . 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 to 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 rest of the paper is organized as follows. In Section 2 we present major technical results for a complete characterization of the geometry sketched in Section 1.4.2. Similarly in Section 3 we present necessary technical machinery and results for convergence proof of the Riemannian

²¹Nevertheless, the objective functions are apparently different. Moreover, we have provided a complete geometric characterization of the objective, in contrast to [FJK96, AGMS12]. We believe the geometric characterization could not only provide insight to the algorithm, but also help improve the algorithm in terms of stability and also finding all components.

²²We have not showed the results on the BG model here, as it seems the structure persists even when θ approaches 1. We suspect the “phase transition” of the landscape occurs at different points for different distributions and Gaussian is the outlying case where the transition occurs at 1.

²³For solving the ICA problem, this suggests the log cosh contrast function, that works well empirically [Hyv99], may not work for all distributions (rotation-invariant Gaussian excluded of course).

trust-region algorithm over the sphere, corresponding to Section 1.4.3. In Section 4, we discuss the whole algorithmic pipeline for recovering complete dictionaries given \mathbf{Y} , and present the main theorems. After presenting a simple simulation to corroborate our theory in Section 5, we wrap up the main content in Section 6 by discussing possible improvement and future directions after this work. All major proofs of geometrical and algorithmic results are deferred to Section 7 and Section 8, respectively. Section 9 augments the main results. The appendices cover some recurring technical tools and auxiliary results for the proofs.

The codes to reproduce all the figures and experimental results can be found online:

https://github.com/sunju/dl_focm

2 High-dimensional Function Landscapes

To characterize the function landscape of $f(\mathbf{q}; \mathbf{X}_0)$ over \mathbb{S}^{n-1} , we mostly work with the function

$$g(\mathbf{w}) \doteq f(\mathbf{q}(\mathbf{w}); \mathbf{X}_0) = \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{q}(\mathbf{w})^* (\mathbf{x}_0)_k), \quad (2.1)$$

induced by the reparametrization

$$\mathbf{q}(\mathbf{w}) = \left(\mathbf{w}, \sqrt{1 - \|\mathbf{w}\|^2} \right), \quad \mathbf{w} \in \mathbb{B}^{n-1}. \quad (2.2)$$

In particular, we focus our attention to the smaller set

$$\Gamma = \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\}, \quad (2.3)$$

because $\mathbf{q}(\Gamma)$ contains all points $\mathbf{q} \in \mathbb{S}^{n-1}$ with $n \in \arg \max_{i \in \pm[n]} \mathbf{q}^* \mathbf{e}_i$ and we can characterize other parts of f on \mathbb{S}^{n-1} using projection onto other equatorial planes. Note that over Γ , $q_n = (1 - \|\mathbf{w}\|^2)^{1/2} \geq \frac{1}{2\sqrt{n}}$.

2.1 Main Geometric Theorems

Theorem 2.1 (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_\star and C , such that for any $\theta \in (0, 1/2)$ and $\mu < \min \{c_a \theta n^{-1}, c_b n^{-5/4}\}$, whenever $p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta}$, the following hold simultaneously with high probability:*

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0) \succeq \frac{c_\star \theta}{\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (2.4)$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \geq c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (2.5)$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (2.6)$$

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

$$\|\mathbf{w}_\star - \mathbf{0}\| \leq \min \left\{ \frac{c_c \mu}{\theta} \sqrt{\frac{n \log p}{p}}, \frac{\mu}{16} \right\}. \quad (2.7)$$

In particular, with this choice of p , the probability the claim fails to hold is at most $4np^{-10} + \theta(np)^{-7} + \exp(-0.3\theta np) + c_d \exp(-c_e p \mu^2 \theta^2 / n^2)$. Here c_a to c_e are all positive numerical constants.

Here $\mathbf{q}(\mathbf{0}) = \mathbf{e}_n$, which exactly recovers the last row of \mathbf{X}_0 , \mathbf{x}_0^n . Though the unique local minimizer \mathbf{w}_\star may not be $\mathbf{0}$, it is very near to $\mathbf{0}$. Hence the resulting $\mathbf{q}(\mathbf{w}_\star)$ produces a close approximation to \mathbf{x}_0^n . Note that $\mathbf{q}(\Gamma)$ (strictly) contains all points $\mathbf{q} \in \mathbb{S}^{n-1}$ such that $n = \arg \max_{i \in \pm[n]} \mathbf{q}^* \mathbf{e}_i$. We can characterize the graph of the function $f(\mathbf{q}; \mathbf{X}_0)$ in the vicinity of other signed basis vector $\pm \mathbf{e}_i$ simply by changing the plane \mathbf{e}_n^\perp to \mathbf{e}_i^\perp . Doing this $2n$ times (and multiplying the failure probability in Theorem 2.1 by $2n$), we obtain a characterization of $f(\mathbf{q}; \mathbf{X}_0)$ over the entirety of \mathbb{S}^{n-1} .²⁴ The result is captured by the next corollary.

Corollary 2.2 *Suppose $\mathbf{A}_0 = \mathbf{I}$ and hence $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0 = \mathbf{X}_0$. There exist positive constant C , such that for any $\theta \in (0, 1/2)$ and $\mu < \min \{c_a \theta n^{-1}, c_b n^{-5/4}\}$, whenever $p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta}$, with probability at least $1 - 8n^2 p^{-10} - \theta(np)^{-7} - \exp(-0.3\theta np) - c_c \exp(-c_d p \mu^2 \theta^2 / n^2)$, the function $f(\mathbf{q}; \mathbf{X}_0)$ has exactly $2n$ local minimizers over the sphere \mathbb{S}^{n-1} . In particular, there is a bijective map between these minimizers and signed basis vectors $\{\pm \mathbf{e}_i\}_i$, such that the corresponding local minimizer \mathbf{q}_\star and $\mathbf{b} \in \{\pm \mathbf{e}_i\}_i$ satisfy*

$$\|\mathbf{q}_\star - \mathbf{b}\| \leq \sqrt{2} \min \left\{ \frac{c_c \mu}{\theta} \sqrt{\frac{n \log p}{p}}, \frac{\mu}{16} \right\}. \quad (2.8)$$

Here c_a to c_d are numerical constants (possibly different from that in the above theorem).

Proof By Theorem 2.1, over $\mathbf{q}(\Gamma)$, $\mathbf{q}(\mathbf{w}_\star)$ is the unique local minimizer. Suppose not. Then there exist $\mathbf{q}' \in \mathbf{q}(\Gamma)$ with $\mathbf{q}' \neq \mathbf{q}(\mathbf{w}_\star)$ and $\varepsilon > 0$, such that $f(\mathbf{q}'; \mathbf{X}_0) \leq f(\mathbf{q}; \mathbf{X}_0)$ for all $\mathbf{q} \in \mathbf{q}(\Gamma)$ satisfying $\|\mathbf{q}' - \mathbf{q}\| < \varepsilon$. Since the mapping $\mathbf{w} \mapsto \mathbf{q}(\mathbf{w})$ is $2\sqrt{n}$ -Lipschitz (Lemma 7.7), $g(\mathbf{w}(\mathbf{q}'); \mathbf{X}_0) \leq g(\mathbf{w}(\mathbf{q}); \mathbf{X}_0)$ for all $\mathbf{w} \in \Gamma$ satisfying $\|\mathbf{w}(\mathbf{q}') - \mathbf{w}(\mathbf{q})\| < \varepsilon / (2\sqrt{n})$, implying $\mathbf{w}(\mathbf{q}')$ is a local minimizer different from \mathbf{w}_\star , a contradiction. Let $\|\mathbf{w}_\star - \mathbf{0}\| = \eta$. Straightforward calculation shows

$$\|\mathbf{q}(\mathbf{w}_\star) - \mathbf{e}_n\|^2 = \left(1 - \sqrt{1 - \eta^2}\right)^2 + \eta^2 = 2 - 2\sqrt{1 - \eta^2} \leq 2\eta^2.$$

Repeating the argument $2n$ times in the vicinity of other signed basis vectors $\pm \mathbf{e}_i$ gives $2n$ local minimizers of f . Indeed, the $2n$ symmetric sections cover the sphere with certain overlaps, and a simple calculation shows that no such local minimizer lies in the overlapped regions (due to nearness to a signed basis vector). There is no extra local minimizer, as such local minimizer is contained in at least one of the $2n$ symmetric sections, resulting two different local minimizers in one section, contradicting the uniqueness result we obtained above. ■

Though the $2n$ isolated local minimizers may have different objective values, they are equally good in the sense any of them produces a close approximation to a certain row of \mathbf{X}_0 . As discussed

²⁴In fact, it is possible to pull the very detailed geometry captured in (2.4) through (2.6) back to the sphere (i.e., the \mathbf{q} space) also; analysis of the Riemannian trust-region algorithm later does part of these. We will stick to this simple global version here.

in Section 1.4.2, for cases \mathbf{A}_0 is an orthobasis other than \mathbf{I} , the landscape of $f(\mathbf{q}; \mathbf{Y})$ is simply a rotated version of the one we characterized above.

Theorem 2.3 (High-dimensional landscape - complete dictionary) *Suppose \mathbf{A}_0 is complete with its condition number $\kappa(\mathbf{A}_0)$. There exist positive constants c_* and C , such that for any $\theta \in (0, 1/2)$ and $\mu < \min\{c_a\theta n^{-1}, c_b n^{-5/4}\}$, when $p \geq \frac{C}{c_*^2\theta} \max\left\{\frac{n^4}{\mu^4}, \frac{n^5}{\mu^2}\right\} \kappa^8(\mathbf{A}_0) \log^4\left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta}\right)$ 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 high probability:*

$$\nabla^2 g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}}) \succeq \frac{c_*\theta}{2\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (2.9)$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}})}{\|\mathbf{w}\|} \geq \frac{1}{2} c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (2.10)$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}}) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -\frac{1}{2} c_* \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (2.11)$$

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

$$\|\mathbf{w}_* - \mathbf{0}\| \leq \frac{\mu}{7}. \quad (2.12)$$

In particular, with this choice of p , the probability the claim fails to hold is at most $4np^{-10} + \theta(np)^{-7} + \exp(-0.3\theta np) + p^{-8} + c_d \exp(-c_e p \mu^2 \theta^2 / n^2)$. Here c_a to c_e are all positive numerical constants.

Corollary 2.4 *Suppose \mathbf{A}_0 is complete with its condition number $\kappa(\mathbf{A}_0)$. There exist positive constants c_* and C , such that for any $\theta \in (0, 1/2)$ and $\mu < \min\{c_a\theta n^{-1}, c_b n^{-5/4}\}$, when $p \geq \frac{C}{c_*^2\theta} \max\left\{\frac{n^4}{\mu^4}, \frac{n^5}{\mu^2}\right\} \kappa^8(\mathbf{A}_0) \log^4\left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta}\right)$ and $\bar{\mathbf{Y}} \doteq \sqrt{p\theta}(\mathbf{Y}\mathbf{Y}^*)^{-1/2}\mathbf{Y}$, $\mathbf{U}\Sigma\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$, with probability at least $1 - 8n^2 p^{-10} - \theta(np)^{-7} - \exp(-0.3\theta np) - p^{-8} - c_d \exp(-c_e p \mu^2 \theta^2 / n^2)$, the function $f(\mathbf{q}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}})$ has exactly $2n$ local minimizers over the sphere \mathbb{S}^{n-1} . In particular, there is a bijective map between these minimizers and signed basis vectors $\{\pm \mathbf{e}_i\}_i$, such that the corresponding local minimizer \mathbf{q}_* and $\mathbf{b} \in \{\pm \mathbf{e}_i\}_i$ satisfy*

$$\|\mathbf{q}_* - \mathbf{b}\| \leq \frac{\sqrt{2}\mu}{7}. \quad (2.13)$$

Here c_a to c_d are numerical constants (possibly different from that in the above theorem).

We will omit the proof as it is almost identical to that of corollary 2.2.

2.2 Useful Technical Lemmas and Proof Ideas for Orthogonal Dictionaries

The proof of Theorem 2.1 is conceptually straightforward: one shows that $\mathbb{E}_{\mathbf{X}_0} [g(\mathbf{w}; \mathbf{X}_0)]$ has the claimed properties, and then proves that each of the quantities of interest concentrates uniformly about its expectation. The detailed calculations are nontrivial.

The next three propositions show that in the expected function landscape, we see successively strongly convex region, nonzero gradient region, and directional negative curvature region when moving away from zero, as depicted in Figure 2 and sketched in Section 1.4.2. Note that in this case

$$\mathbb{E}_{\mathbf{X}_0} [g(\mathbf{q}; \mathbf{X}_0)] = \mathbb{E}_{\mathbf{x} \sim i.i.d. \text{BG}(\theta)} [h_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x})].$$

Proposition 2.5 *There exists a positive constant c , such that for every $\theta \in (0, \frac{1}{2})$ and any $R_h \in (0, \sqrt{\frac{4n-1}{4n}})$, if $\mu \leq c \min \{\theta R_h^2 n^{-1}, R_h n^{-5/4}\}$, it holds for every \mathbf{w} satisfying $R_h \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}$ that*

$$\frac{\mathbf{w}^* \nabla_{\mathbf{w}}^2 \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \mathbf{w}}{\|\mathbf{w}\|^2} \leq -\frac{\theta}{2\sqrt{2\pi}}.$$

Proof See Section 7.1.1 on Page 47. ■

Proposition 2.6 *For every $\theta \in (0, \frac{1}{2})$ and every $\mu \leq 9/50$, it holds for every \mathbf{w} satisfying $r_g \leq \|\mathbf{w}\| \leq R_g$, where $r_g = \frac{\mu}{6\sqrt{2}}$ and $R_g = \frac{1-\theta}{10\sqrt{5}}$, that*

$$\frac{\mathbf{w}^* \nabla_{\mathbf{w}} \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})]}{\|\mathbf{w}\|} \geq \frac{\theta}{20\sqrt{2\pi}}.$$

Proof See Section 7.1.2 on Page 52. ■

Proposition 2.7 *For every $\theta \in (0, \frac{1}{2})$, and every $\mu \leq \frac{1}{20\sqrt{n}}$, it holds for every \mathbf{w} satisfying $\|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$ that*

$$\mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \succeq \frac{\theta}{25\sqrt{2\pi}\mu} \mathbf{I}.$$

Proof See Section 7.1.3 on Page 54. ■

To prove that the above hold qualitatively for finite p , i.e., the function $g(\mathbf{w}; \mathbf{X}_0)$, we will need first prove that for a fixed \mathbf{w} each of the quantity of interest concentrate about their expectation w.h.p., and the function is nice enough (Lipschitz) such that we can extend the results to all \mathbf{w} via a discretization argument. The next three propositions provide the desired pointwise concentration results.

Proposition 2.8 *Suppose $0 < \mu \leq \frac{1}{\sqrt{n}}$. For every $\mathbf{w} \in \Gamma$, it holds that for any $t > 0$,*

$$\mathbb{P} \left[\left| \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \right| \geq t \right] \leq 4 \exp \left(-\frac{p\mu^2 t^2}{512n^2 + 32n\mu t} \right).$$

Proof See Page 58 under Section 7.1.4. ■

Proposition 2.9 *For every $\mathbf{w} \in \Gamma$, it holds that for any $t > 0$,*

$$\mathbb{P} \left[\left| \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \right] \right| \geq t \right] \leq 2 \exp \left(-\frac{pt^2}{8n + 4\sqrt{nt}} \right).$$

Proof See Page 59 under Section 7.1.4. ■

Proposition 2.10 Suppose $0 < \mu \leq \frac{1}{\sqrt{n}}$. For every $\mathbf{w} \in \Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \leq 1/4\}$, it holds that for any $t > 0$,

$$\mathbb{P} [\|\nabla^2 g(\mathbf{w}; \mathbf{X}_0) - \mathbb{E} [\nabla^2 g(\mathbf{w}; \mathbf{X}_0)]\| \geq t] \leq 4n \exp\left(-\frac{p\mu^2 t^2}{512n^2 + 32\mu nt}\right).$$

Proof See Page 60 under Section 7.1.4. ■

The next three propositions provide the desired Lipschitz results.

Proposition 2.11 (Hessian Lipschitz) Fix any $r_\wedge \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_\wedge\}$, $\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2}$ is L_\wedge -Lipschitz with

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

Proof See Page 64 under Section 7.1.5. ■

Proposition 2.12 (Gradient Lipschitz) Fix any $r_g \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_g\}$, $\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|}$ is L_g -Lipschitz with

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

Proof See Page 65 under Section 7.1.5. ■

Proposition 2.13 (Lipschitz for Hessian around zero) Fix any $r_\vee \in (0, \frac{1}{2})$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \leq r_\vee\}$, $\nabla^2 g(\mathbf{w}; \mathbf{X}_0)$ is L_\vee -Lipschitz with

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

Proof See Page 65 under Section 7.1.5. ■

Integrating the above pieces, Section 7.2 provides a complete proof of Theorem 2.1.

2.3 Extending to Complete Dictionaries

As hinted in Section 1.4.2, instead of proving things from scratch, we build on the results we have obtained for orthogonal dictionaries. In particular, we will work with the preconditioned data matrix

$$\bar{\mathbf{Y}} \doteq \left(\frac{1}{p\theta} \mathbf{Y} \mathbf{Y}^*\right)^{-1/2} \mathbf{Y} \tag{2.14}$$

and show that the function landscape $f(\mathbf{q}; \bar{\mathbf{Y}})$ looks qualitatively like that of orthogonal dictionaries (up to a global rotation), provided that p is large enough.

The next lemma shows $\bar{\mathbf{Y}}$ can be treated as being generated from an orthobasis with the same BG coefficients, plus small noise.

Lemma 2.14 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 (2.14) as

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

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

Proof See Page 69 under Section 7.3. ■

Notice that \mathbf{UV}^* above is orthogonal, and that landscape of $f(\mathbf{q}; \bar{\mathbf{Y}})$ is simply a rotated version of that of $f(\mathbf{q}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}})$, or using the notation in the above lemma, that of $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^*\mathbf{\Xi}\mathbf{X}_0) = f(\mathbf{q}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0)$ assuming $\tilde{\mathbf{\Xi}} \doteq \mathbf{V}\mathbf{U}^*\mathbf{\Xi}$. So similar to the orthogonal case, it is enough to consider this “canonical” case, and its “canonical” reparametrization:

$$g(\mathbf{w}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0) = \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{q}^*(\mathbf{w})(\mathbf{x}_0)_k + \mathbf{q}^*(\mathbf{w})\tilde{\mathbf{\Xi}}(\mathbf{x}_0)_k).$$

The following lemma provides quantitative comparison between the gradient and Hessian of $g(\mathbf{w}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0)$ and that of $g(\mathbf{w}; \mathbf{X}_0)$.

Lemma 2.15 There exist positive constants C_a and C_b , such that for all $\mathbf{w} \in \Gamma$,

$$\begin{aligned} \left\| \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0) - \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}_0) \right\| &\leq C_a \frac{n}{\mu} \log(np) \|\tilde{\mathbf{\Xi}}\|, \\ \left\| \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}}\mathbf{X}_0) - \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}_0) \right\| &\leq C_b \max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \|\tilde{\mathbf{\Xi}}\| \end{aligned}$$

with probability at least $1 - \theta(np)^{-7} - \exp(-0.3\theta np)$.

Proof See Page 70 under Section 7.3. ■

Combining the above two lemmas, it is easy to see when p is large enough, $\|\tilde{\mathbf{\Xi}}\| = \|\mathbf{\Xi}\|$ is then small enough (Lemma 2.14), and hence the changes to the gradient and Hessian caused by the perturbation are small. This gives the results presented in Theorem 2.3; see Section 7.3 for the detailed proof. In particular, for the p chosen in Theorem 2.3, it holds that

$$\|\tilde{\mathbf{\Xi}}\| \leq cc_*\theta \left(\max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1} \quad (2.15)$$

for a certain constant c which can be made arbitrarily small by making the constant C in p large.

3 Finding One Local Minimizer via the Riemannian Trust-Region Method

The above geometric results show every local minimizer of $f(\mathbf{q}; \hat{\mathbf{Y}})$ over \mathbb{S}^{n-1} approximately recovers one row of \mathbf{X}_0 . So the crucial problem left now is how to efficiently obtain one of the local minimizers.

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 efficiently produces an accurate approximation²⁵ 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 [AMS09] for relevant background and details.

3.1 The Riemannian Trust-Region Algorithm over the Sphere

We are interested to seek one local minimizer of the problem

$$\text{minimize } f(\mathbf{q}; \widehat{\mathbf{Y}}) \doteq \frac{1}{p} \sum_{k=1}^p h_{\mu}(\mathbf{q}^* \widehat{\mathbf{y}}_k) \quad \text{subject to } \mathbf{q} \in \mathbb{S}^{n-1}. \quad (3.1)$$

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 \widehat{f} to the objective function $f(\mathbf{q})$, over a ball centered about the current iterate.

Here, we are interested in the restriction of f to the unit sphere \mathbb{S}^{n-1} . Instead of directly approximating the function in \mathbb{R}^n , we form quadratic approximations of f in the tangent space of \mathbb{S}^{n-1} . Recall that the tangent space of a sphere at a point $\mathbf{q} \in \mathbb{S}^{n-1}$ is $T_{\mathbf{q}}\mathbb{S}^{n-1} = \{\boldsymbol{\delta} \in \mathbb{R}^n \mid \mathbf{q}^* \boldsymbol{\delta} = 0\}$, i.e., the set of vectors that are orthogonal to \mathbf{q} . Consider $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ with $\|\boldsymbol{\delta}\| = 1$. The map $\gamma(t) : t \mapsto \mathbf{q} \cos t + \boldsymbol{\delta} \sin t$ defines a smooth curve on the sphere that satisfies $\gamma(0) = \mathbf{q}$ and $\dot{\gamma}(0) = \boldsymbol{\delta}$. The function $f \circ \gamma(t)$ obviously is smooth and we expect Taylor expansion around 0 a good approximation of the function, at least in the vicinity of 0. Taylor's theorem gives

$$f \circ \gamma(t) = f(\mathbf{q}) + t \langle \nabla f(\mathbf{q}), \boldsymbol{\delta} \rangle + \frac{t^2}{2} (\boldsymbol{\delta}^* \nabla^2 f(\mathbf{q}) \boldsymbol{\delta} - \langle \nabla f(\mathbf{q}), \mathbf{q} \rangle) + O(t^3).$$

We therefore form the ‘‘quadratic’’ approximation $\widehat{f}(\boldsymbol{\delta}; \mathbf{q}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{R}$ as

$$\widehat{f}(\boldsymbol{\delta}; \mathbf{q}, \widehat{\mathbf{Y}}) \doteq f(\mathbf{q}) + \langle \nabla f(\mathbf{q}; \widehat{\mathbf{Y}}), \boldsymbol{\delta} \rangle + \frac{1}{2} \boldsymbol{\delta}^* \left(\nabla^2 f(\mathbf{q}; \widehat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \widehat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I} \right) \boldsymbol{\delta}. \quad (3.2)$$

Given the previous iterate $\mathbf{q}^{(k-1)}$, the TRM produces the next iterate by generating a solution $\widehat{\boldsymbol{\delta}}$ to

$$\min_{\boldsymbol{\delta} \in T_{\mathbf{q}^{(k-1)}}\mathbb{S}^{n-1}, \|\boldsymbol{\delta}\| \leq \Delta} \widehat{f}(\boldsymbol{\delta}; \mathbf{q}^{(k-1)}), \quad (3.3)$$

and then ‘‘pull’’ the solution $\widehat{\boldsymbol{\delta}}$ from $T_{\mathbf{q}}\mathbb{S}^{n-1}$ back to \mathbb{S}^{n-1} . Moreover, for any vector $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$, the exponential map $\exp_{\mathbf{q}}(\boldsymbol{\delta}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{S}^{n-1}$ is

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

²⁵By ‘‘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.

If we choose the exponential map to pull back the movement $\widehat{\delta}^{26}$, the next iterate then reads

$$\mathbf{q}^{(k)} = \mathbf{q}^{(k-1)} \cos \|\widehat{\delta}\| + \frac{\widehat{\delta}}{\|\widehat{\delta}\|} \sin \|\widehat{\delta}\|. \quad (3.4)$$

We have motivated (3.2) and hence the algorithm in an intuitive way from the Taylor approximation to the function f over \mathbb{S}^{n-1} . To understand its properties, it is useful to interpret it as a *Riemannian trust-region method* over the manifold \mathbb{S}^{n-1} . The class of algorithm is discussed in detail in the monograph [AMS09]. In particular, the quadratic approximation (3.2) can be obtained by noting that the function $f \circ \exp_{\mathbf{q}}(\delta; \widehat{\mathbf{Y}}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{R}$ obeys

$$f \circ \exp_{\mathbf{q}}(\delta; \widehat{\mathbf{Y}}) = f(\mathbf{q}; \widehat{\mathbf{Y}}) + \langle \delta, \text{grad } f(\mathbf{q}; \widehat{\mathbf{Y}}) \rangle + \frac{1}{2} \delta^* \text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}}) \delta + O(\|\delta\|^3),$$

where $\text{grad } f(\mathbf{q}; \widehat{\mathbf{Y}})$ and $\text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}})$ are the Riemannian gradient and Riemannian Hessian [AMS09] respectively, defined as

$$\begin{aligned} \text{grad } f(\mathbf{q}; \widehat{\mathbf{Y}}) &\doteq \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \nabla f(\mathbf{q}; \widehat{\mathbf{Y}}), \\ \text{Hess } f(\mathbf{q}; \widehat{\mathbf{Y}}) &\doteq \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \left(\nabla^2 f(\mathbf{q}; \widehat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \widehat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I} \right) \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}}, \end{aligned}$$

with $\mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \doteq \mathbf{I} - \mathbf{q}\mathbf{q}^*$ the orthoprojector onto the tangent space $T_{\mathbf{q}}\mathbb{S}^{n-1}$. We will use these standard notions in analysis of the algorithm.

To solve the subproblem (3.3) numerically, we can take any matrix $\mathbf{U} \in \mathbb{R}^{n \times (n-1)}$ whose columns form an orthonormal basis for $T_{\mathbf{q}^{(k-1)}}\mathbb{S}^{n-1}$, and produce a solution $\widehat{\xi}$ to

$$\min_{\|\xi\| \leq \Delta} \widehat{f}(\mathbf{U}\xi; \mathbf{q}^{(k-1)}), \quad (3.5)$$

where by (3.2),

$$\begin{aligned} \widehat{f}(\mathbf{U}\xi; \mathbf{q}^{(k-1)}) &= f(\mathbf{q}) + \langle \mathbf{U}^* \nabla f(\mathbf{q}^{(k-1)}), \xi \rangle + \\ &\quad \frac{1}{2} \xi^* \left(\mathbf{U}^* \nabla^2 f(\mathbf{q}^{(k-1)}; \widehat{\mathbf{Y}}) \mathbf{U} - \langle \nabla f(\mathbf{q}^{(k-1)}; \widehat{\mathbf{Y}}), \mathbf{q}^{(k-1)} \rangle \mathbf{I}_{n-1} \right) \xi. \end{aligned}$$

Solution to (3.3) can then be recovered as $\widehat{\delta} = \mathbf{U}\widehat{\xi}$. The problem (3.5) is an instance of the classic *trust region subproblem*, i.e., minimizing a quadratic function subject to a single quadratic constraint, which can be solved in polynomial time, either by root finding methods [MS83, CGT00b] or by semidefinite programming (SDP) [RW97, YZ03, FW04, HK14]. As the root finding methods numerically suffer from the so-called “hard case” [MS83], we deploy the SDP approach here. We introduce

$$\tilde{\xi} = [\xi^*, 1]^*, \quad \Theta = \tilde{\xi} \tilde{\xi}^*, \quad \mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^* & 0 \end{bmatrix}, \quad (3.6)$$

where $\mathbf{A} = \mathbf{U}^* (\nabla^2 f(\mathbf{q}^{(k-1)}; \widehat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}^{(k-1)}; \widehat{\mathbf{Y}}), \mathbf{q}^{(k-1)} \rangle \mathbf{I}) \mathbf{U}$ and $\mathbf{b} = \mathbf{U}^* \nabla f(\mathbf{q}^{(k-1)}; \widehat{\mathbf{Y}})$. The resulting SDP to solve is

$$\text{minimize } \Theta \langle \mathbf{M}, \Theta \rangle, \quad \text{subject to } \text{tr}(\Theta) \leq \Delta^2 + 1, \quad \langle \mathbf{E}_{n+1}, \Theta \rangle = 1, \quad \Theta \succeq \mathbf{0}, \quad (3.7)$$

²⁶The 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 [AMS09].

where $\mathbf{E}_{n+1} = e_{n+1}e_{n+1}^*$. Once the problem (3.7) is solved to its optimal Θ_* , one can provably recover the optimal solution ξ_* of (3.5) by computing the SVD of $\Theta_* = \tilde{U}\Sigma\tilde{V}^*$, and extract as a subvector by the first $n - 1$ coordinates of the principal eigenvector \tilde{u}_1 (see Appendix B of [BV04]).

The choice of trust region size Δ is important both for the convergence theory and practical effectiveness of TRMs. Following standard recommendations (see, e.g., Chapter 4 of [NW06]), we use a backtracking approach which modifies Δ from iteration to iteration based on the accuracy of the approximation \hat{f} . The whole algorithmic procedure is described as pseudocode as Algorithm 1. In our numerical implementation, we randomly initialize $\mathbf{q}^{(0)}$ and set $\Delta^{(0)} = 0.1, \eta_{vs} =$

Algorithm 1 Riemannian TRM Algorithm for Finding One Local Minimizer

Input: Data matrix $\mathbf{Y} \in \mathbb{R}^{n \times p}$, smoothing parameter μ and parameters $\eta_{vs}, \eta_s, \gamma_i, \gamma_d, \Delta_{\max}, \Delta_{\min}$

Output: $\hat{\mathbf{q}} \in \mathbb{S}^{n-1}$

- 1: Initialize $\mathbf{q}^{(0)} \in \mathbb{S}^{n-1}, \Delta^{(0)}$ and $k = 1$,
- 2: **while** not converged **do**
- 3: Set $\mathbf{U} \in \mathbb{R}^{n \times (n-1)}$ to be an orthonormal basis for $T_{\mathbf{q}^{(k-1)}}\mathbb{S}^{n-1}$
- 4: Solve the trust region subproblem

$$\hat{\xi} = \arg \min_{\|\xi\| \leq \Delta^{(k-1)}} \hat{f}(\mathbf{U}\xi; \mathbf{q}^{(k-1)}, \hat{\mathbf{Y}})$$

- 5: Set

$$\hat{\delta} \leftarrow \mathbf{U}\hat{\xi}, \quad \hat{\mathbf{q}} \leftarrow \mathbf{q}^{(k-1)} \cos \|\hat{\delta}\| + \frac{\hat{\delta}}{\|\hat{\delta}\|} \sin \|\hat{\delta}\|.$$

- 6: Set

$$\rho_k \leftarrow \frac{f(\mathbf{q}^{(k-1)}; \hat{\mathbf{Y}}) - f(\hat{\mathbf{q}}; \hat{\mathbf{Y}})}{f(\mathbf{q}^{(k-1)}; \hat{\mathbf{Y}}) - \hat{f}(\hat{\delta}; \mathbf{q}^{(k-1)}, \hat{\mathbf{Y}})}$$

- 7: **if** $\rho_k \geq \eta_{vs}$ **then**
 - 8: Set $\mathbf{q}^{(k)} \leftarrow \hat{\mathbf{q}}$ and $\Delta^{(k)} \leftarrow \min(\gamma_i \Delta^{(k-1)}, \Delta_{\max})$. ▷ very successful
 - 9: **else if** $\rho_k \geq \eta_s$ **then**
 - 10: Set $\mathbf{q}^{(k)} \leftarrow \hat{\mathbf{q}}$ and $\Delta^{(k)} \leftarrow \Delta^{(k-1)}$. ▷ successful
 - 11: **else**
 - 12: Set $\mathbf{q}^{(k)} \leftarrow \mathbf{q}^{(k-1)}$ and $\Delta^{(k)} \leftarrow \max(\gamma_d \Delta^{(k-1)}, \Delta_{\min})$. ▷ unsuccessful
 - 13: **end if**
 - 14: Set $k = k + 1$.
 - 15: **end while**
-

0.9, $\eta_s = 0.1, \gamma_d = 1/2, \gamma_i = 2, \Delta_{\max} = 1$ and $\Delta_{\min} = 10^{-16}$, and the algorithm is stopped when $(f(\hat{\mathbf{q}}) - f(\mathbf{q}^{(k-1)})) / \|\hat{\delta}\| \leq 10^{-6}$.

3.2 Main Convergence Results

By using general results on the Riemannian TRM (see, e.g., Chapter 7 of [AMS09]), it is not difficult to prove that the iterates $\mathbf{q}^{(k)}$ produced by Algorithm 1 converge to a critical point of the objective $f(\mathbf{q})$ over \mathbb{S}^{n-1} . In this section, we show that under our probabilistic assumptions, this claim can be 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. The arguments described in Section 2 show that with high probability every local minimizer of f produces a close approximation of one row of \mathbf{X}_0 . Taken together, this implies that the algorithm efficiently produces a close approximation to one row of \mathbf{X}_0 .

Our next two theorems summarize the convergence results for orthogonal and complete dictionaries, respectively.

Theorem 3.1 (TRM convergence - orthogonal dictionary) *Suppose the dictionary \mathbf{A}_0 is orthogonal. Then there exists a positive constant C , such that for all $\theta \in (0, 1/2)$, and $\mu < \min \{c_a \theta n^{-1}, c_b n^{-5/4}\}$, whenever $\exp(n) \geq p \geq C n^3 \log \frac{n}{\mu \theta} / (\mu^2 \theta^2)$, with probability at least $1 - 8n^2 p^{-10} - \theta (np)^{-7} - \exp(-0.3\theta np) - p^{-10} - c_c \exp(-c_d p \mu^2 \theta^2 / n^2)$, the Riemannian trust-region algorithm with input data matrix $\hat{\mathbf{Y}} = \mathbf{Y}$, any initialization $\mathbf{q}^{(0)}$ on the sphere, and a step size satisfying*

$$\Delta \leq \min \left\{ \frac{c_e c_\star \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \frac{c_f c_\#^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)} \right\}. \quad (3.8)$$

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

$$\max \left\{ \frac{c_g n^6 \log^3(np)}{c_\star^3 \theta^3 \mu^4}, \frac{c_h n}{c_\#^2 \theta^2 \Delta^2} \right\} \left(f(\mathbf{q}^{(0)}) - f(\mathbf{q}_\star) \right) + \log \log \frac{c_i c_\star \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \quad (3.9)$$

iterations. Here $c_\star, c_\#$ as defined in Theorem 2.1 and Lemma 3.9 respectively (c_\star and $c_\#$ can be set to the same constant value), and c_a, c_b are the same numerical constants as defined in Theorem 2.1, c_c to c_i are other positive numerical constants.

Theorem 3.2 (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 $\mu < \min \{c_a \theta n^{-1}, c_b n^{-5/4}\}$, whenever $\exp(n) \geq p \geq \frac{C}{c_\#^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right)$, with probability at least $1 - 8n^2 p^{-10} - \theta (np)^{-7} - \exp(-0.3\theta np) - 2p^{-8} - c_c \exp(-c_d p \mu^2 \theta^2 / n^2)$, the Riemannian trust-region algorithm with input data matrix $\bar{\mathbf{Y}} \doteq \sqrt{p\theta} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y}$ where $\mathbf{U}\Sigma\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$, any initialization $\mathbf{q}^{(0)}$ on the sphere and a step size satisfying*

$$\Delta \leq \min \left\{ \frac{c_e c_\star \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \frac{c_f c_\#^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)} \right\}. \quad (3.10)$$

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

$$\max \left\{ \frac{c_g n^6 \log^3(np)}{c_\star^3 \theta^3 \mu^4}, \frac{c_h n}{c_\#^2 \theta^2 \Delta^2} \right\} \left(f(\mathbf{q}^{(0)}) - f(\mathbf{q}_\star) \right) + \log \log \frac{c_i c_\star \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \quad (3.11)$$

iterations. Here $c_\star, c_\#$ as defined in Theorem 2.1 and Lemma 3.9 respectively (c_\star and $c_\#$ can be set to the same constant value), and c_a, c_b are the same numerical constants as defined in Theorem 2.1, c_c to c_i are other positive numerical constants.

Our convergence result shows that for any target accuracy $\varepsilon > 0$ the algorithm terminates within polynomially many steps. Our estimate of the number of steps is pessimistic: our analysis has assumed a fixed step size Δ and the running time is relatively large degree polynomial in p and n , while on typical numerical examples (e.g., $\mu = 10^{-2}$, $n \sim 100$, and $\varepsilon = O(\mu)$), the algorithm with adaptive step size as described in Algorithm 1 produces an accurate solution in relatively few (20-50) iterations. Nevertheless, 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 [MK87] show that in general it is NP-hard to find a local minimum of a nonconvex function.

3.3 Useful Technical Results and Proof Ideas for Orthogonal Dictionaries

The reason that our algorithm is successful derives from the geometry depicted in Figure 2 and formalized in Theorem 2.1. 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 (an eigenvalue of the Hessian which is bounded below zero). The Riemannian TRM algorithm is capable of exploiting these quantities to reduce the objective value by at least a constant in each iteration. The total number of iterations spent away from the vicinity of the local minimizers can be bounded by comparing this constant to the initial objective value. Our proofs follow exactly this line and make the various quantities precise.

3.3.1 Basic Facts about the Sphere

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

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

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). Moreover, for any $\boldsymbol{\delta} \in T_{\mathbf{q}}\mathbb{S}^{n-1}$, the exponential map $\exp_{\mathbf{q}}(\boldsymbol{\delta}) : T_{\mathbf{q}}\mathbb{S}^{n-1} \mapsto \mathbb{S}^{n-1}$ is given by

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

Let $\nabla f(\mathbf{q})$ 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 (3.1), it is easy to check that

$$\nabla f(\mathbf{q}; \hat{\mathbf{Y}}) = \frac{1}{p} \sum_{k=1}^p \tanh\left(\frac{\mathbf{q}^* \hat{\mathbf{y}}_k}{\mu}\right) \hat{\mathbf{y}}_k, \quad (3.12)$$

$$\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) = \frac{1}{p} \sum_{k=1}^p \frac{1}{\mu} \left[1 - \tanh^2\left(\frac{\mathbf{q}^* \hat{\mathbf{y}}_k}{\mu}\right) \right] \hat{\mathbf{y}}_k \hat{\mathbf{y}}_k^*. \quad (3.13)$$

Since \mathbb{S}^{n-1} is an embedded submanifold of \mathbb{R}^n , the Riemannian gradient and Riemannian Hessian defined on $T_{\mathbf{q}}\mathbb{S}^{n-1}$ are given by

$$\text{grad } f(\mathbf{q}; \hat{\mathbf{Y}}) = \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \quad (3.14)$$

$$\text{Hess } f(\mathbf{q}; \hat{\mathbf{Y}}) = \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}} \left(\nabla^2 f(\mathbf{q}; \hat{\mathbf{Y}}) - \langle \nabla f(\mathbf{q}; \hat{\mathbf{Y}}), \mathbf{q} \rangle \mathbf{I} \right) \mathcal{P}_{T_{\mathbf{q}}\mathbb{S}^{n-1}}; \quad (3.15)$$

so the second-order Taylor approximation for the function f is

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

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

$$\text{grad } f(\mathbf{q}; \hat{\mathbf{Y}}) + \text{Hess } f(\mathbf{q}; \hat{\mathbf{Y}}) \boldsymbol{\delta}_* = \mathbf{0}; \quad (3.16)$$

if $\text{Hess } f(\mathbf{q})$ is positive semidefinite and has full rank $n - 1$ (hence “nondegenerate”²⁷), the unique solution $\boldsymbol{\delta}_*$ is

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

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 [AMS09], 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 [AMS09]):

$$\begin{aligned} \mathcal{P}_{\gamma}^{\tau \leftarrow 0} &= \left(\mathbf{I} - \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} \right) - \mathbf{q} \sin(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \cos(\tau \|\boldsymbol{\delta}\|) \frac{\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|} \\ &= \mathbf{I} + (\cos(\tau \|\boldsymbol{\delta}\|) - 1) \frac{\boldsymbol{\delta} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \sin(\tau \|\boldsymbol{\delta}\|) \frac{\mathbf{q} \boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|}. \end{aligned} \quad (3.17)$$

Similarly, following the notation of [AMS09], 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 $\mathcal{P}_{\gamma}^{\tau \leftarrow 0}$.

3.3.2 Key Steps towards the Proof

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 above 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}_* of $f(\mathbf{q}; \mathbf{X}_0)$ is rotated to $\mathbf{A}_0 \mathbf{q}_*$, one 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

²⁷Note that the $n \times n$ matrix $\text{Hess } f(\mathbf{q}; \hat{\mathbf{Y}})$ has rank at most $n - 1$, as the nonzero \mathbf{q} obviously is in its null space. When $\text{Hess } f(\mathbf{q}; \hat{\mathbf{Y}})$ 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.

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 3.3), 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 2.1). That is, R_I consists of a union of $2n$ spherical caps of radius $\frac{\mu}{4\sqrt{2}}$, 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 $\frac{1}{20\sqrt{5}}$, 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 2.1 and Corollary 2.2, the conditions

$$\theta \in (0, 1/2), \quad \mu < \min \left\{ c_a \theta n^{-1}, c_b n^{-5/4} \right\}, \quad p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta} \quad (3.18)$$

are always in force.

At each step k of the algorithm, suppose $\delta^{(k)}$ is the minimizer of the trust-region subproblem (3.3). We call the step “constrained” if $\|\delta^{(k)}\| = \Delta$ (the minimizer lies on the boundary and hence the constraint is active), and call it “unconstrained” if $\|\delta^{(k)}\| < \Delta$ (the minimizer lies in the relative interior and hence the constraint is not in force). Thus, in the unconstrained case the optimality condition is (3.16).

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

Lemma 3.3 *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 72 under Section 8. ■

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

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

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

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

where $\eta_f \doteq M_{\nabla} + 2M_{\nabla^2} + L_{\nabla} + L_{\nabla^2}$ and $M_{\nabla}, M_{\nabla^2}, L_{\nabla}, L_{\nabla^2}$ are the quantities defined in Lemma 3.3.

Proof See Page 72 under Section 8. ■

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 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 3.5 *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 $\mathbf{w}^* \nabla g(\mathbf{w}) / \|\mathbf{w}\|$ is L_g -Lipschitz on an open ball $\mathcal{B}(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}})$ 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 74 under Section 8. ■

Lemma 3.6 *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 $\mathcal{B}(\mathbf{w}, \frac{3\Delta}{2\pi\sqrt{n}})$ 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*

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

Proof See Page 74 under Section 8. ■

One can take $\beta_g = \beta_\wedge = c_\star\theta$ as shown in Theorem 2.1, and take the Lipschitz results in Section 2.2 (note that $\|\mathbf{X}_0\|_\infty \leq 4 \log^{1/2}(np)$ w.h.p. by Lemma 7.11), 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 3.7 *Assume (3.18). In regions R_{II} and R_{III} , each trust-region step reduces the objective value by at least*

$$d_{\text{II}} = \frac{1}{2} \min \left(\frac{c_\star^2 c_a \theta^2 \mu}{n^2 \log(np)}, \frac{3\Delta c_\star \theta}{4\pi\sqrt{n}} \right), \quad \text{and} \quad d_{\text{III}} = \frac{1}{2} \min \left(\frac{c_\star^3 c_b \theta^3 \mu^4}{n^6 \log^3(np)}, \frac{3\Delta^2 c_\star \theta}{8\pi^2 n} \right) \quad (3.19)$$

respectively, provided that

$$\Delta < \frac{c_c c_\star \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \quad (3.20)$$

where c_a to c_c are positive numerical constants, and c_\star is as defined in Theorem 2.1.

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 $\mathbf{q}^{(k)}$ is in the region R_{II} , by Theorem 2.1, w.h.p., we have $\mathbf{w}^* g(\mathbf{w}) / \|\mathbf{w}\| \geq c_\star \theta$ for the constant c_\star . By Proposition 2.12 and Lemma 7.11, w.h.p., $\mathbf{w}^* g(\mathbf{w}) / \|\mathbf{w}\|$

is $C_2 n^2 \log(np) / \mu$ -Lipschitz. Therefore, By Lemma 3.4 and Lemma 3.5, 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 $\mathbf{q}^{(k)}$ is in the region R_{III} , by Proposition 2.11, Theorem 2.1 and Lemma 7.11, 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 3.4 and Lemma 3.6, 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 (3.19), (3.20) holds. \blacksquare

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 one 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 3.8 *Suppose the trust-region size $\Delta \leq 1$, and that at a given iterate k , $\text{Hess } f(\mathbf{q}^{(k)}) \succeq m_H \mathcal{P}_{T_{\mathbf{q}^{(k)}} \mathbb{S}^{n-1}}$, and $\|\text{Hess } f(\mathbf{q}^{(k)})\| \leq M_H$. Further assume the optimal solution $\delta_* \in T_{\mathbf{q}^{(k)}} \mathbb{S}^{n-1}$ to the trust-region subproblem (3.3) satisfies $\|\delta_*\| = \Delta$, i.e., the norm constraint is active. Then there exists a tangent vector $\delta \in T_{\mathbf{q}^{(k)}} \mathbb{S}^{n-1}$ with $\|\delta\| \leq \Delta$, such that*

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

where η_f is defined the same as Lemma 3.4.

Proof See Page 75 under Section 8. \blacksquare

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 3.9 *There exist positive constants C and c_{\sharp} , such that for all $\theta \in (0, 1/2)$ and $\mu < \min\{\theta, 1/10\}$, whenever $p \geq C n^3 \log \frac{n}{\theta \mu} / (\mu \theta^2)$, it holds with probability at least $1 - \theta (np)^{-7} - \exp(-0.3\theta np) - p^{-10}$ that for all \mathbf{q} with $\|\mathbf{w}(\mathbf{q})\| \leq \frac{\mu}{4\sqrt{2}}$,*

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

Proof See Page 78 under Section 8. ■

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 3.3,

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

Combining this estimate and Lemma 3.9, and Lemma 3.4, we obtain a concrete lower bound for the reduction of objective value for each constrained R_I step.

Proposition 3.10 *Assume (3.18). Each constrained R_I trust-region step (i.e., $\|\delta\| = \Delta$) reduces the objective value by at least*

$$d_I = \frac{cc_*^2\theta^2}{\mu n \log(np)} \Delta^2, \quad (3.21)$$

provided

$$\Delta \leq \frac{c'c_\#^2\theta^2\mu}{n^{5/2} \log^{5/2}(np)}. \quad (3.22)$$

The constant $c_\#$ is as defined in Lemma 3.9 and c, c' are a positive numerical 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.

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

where $c_\#$ is as defined in Lemma 3.9. Combining these estimates with Lemma 3.4 and Lemma 3.8, one trust-region step will find next iterate $\mathbf{q}^{(k+1)}$ 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 (3.22) and the assumed conditions (3.18), 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. ■

By the proof strategy for R_I we sketched before Lemma 3.8, we expect the iteration sequence ultimately always takes unconstrained steps when it moves very near 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 2.1, 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 3.11 *For all \mathbf{q} satisfying $\frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}(\mathbf{q})\| \leq \frac{1}{20\sqrt{5}}$, it holds that*

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

Proof See Page 82 under Section 8. ■

Assuming (3.18), Theorem 2.1 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 3.12 *Suppose the trust-region size $\Delta \leq 1$, and at a given iterate k , $\text{Hess } f(\mathbf{q}^{(k)}) \succeq m_H \mathcal{P}_{T_{\mathbf{q}^{(k)}} \mathbb{S}^{n-1}}$, and that the unique minimizer $\delta_* \in T_{\mathbf{q}^{(k)}} \mathbb{S}^{n-1}$ to the trust region subproblem (3.3) satisfies $\|\delta_*\| < \Delta$ (i.e., the constraint is inactive). Then, for $\mathbf{q}^{(k+1)} = \exp_{\mathbf{q}^{(k)}}(\delta_*)$, we have*

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

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

Proof See Page 82 under Section 8. ■

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

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

Obviously, one can make the upper bound small by tuning down Δ . Combining the above lower bound for $\|\text{grad } f(\mathbf{q})\|$ for $\mathbf{q} \in R_{II}$, one can conclude that when Δ is small, the next iterate $\mathbf{q}^{(k+1)}$ stays in R_I . Another application of the optimality condition (3.16) 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 3.13 *Assume (3.18). 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*

$$\Delta \leq \frac{cc_\sharp^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)}, \quad (3.23)$$

Here c is a positive numerical constant, and c_\sharp is as defined in Lemma 3.9.

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

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

Thus, if $\Delta \leq \frac{1}{20\sqrt{5}} - \frac{\mu}{4\sqrt{2}}$, $\mathbf{q}^{(k+1)}$ will be in $R_I \cup R_{II}$. Next, we show that if Δ is sufficiently small, $\mathbf{q}^{(k+1)}$ will be indeed in R_I . By Lemma 3.12,

$$\begin{aligned}\|\text{grad } f(\mathbf{q}^{(k+1)})\| &\leq \frac{L_H}{2m_H^2} \|\text{grad } f(\mathbf{q}^{(k)})\|^2 \\ &\leq \frac{L_H M_H^2}{2m_H^2} \left\| \left[\mathbf{U}^* \text{Hess } f(\mathbf{q}^{(k)}) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f(\mathbf{q}^{(k)}) \right\|^2 \leq \frac{L_H M_H^2}{2m_H^2} \Delta^2,\end{aligned}\quad (3.24)$$

where we have used the fact that

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

as the step is unconstrained. On the other hand, by Theorem 2.1 and Lemma 3.11, w.h.p.

$$\|\text{grad } f(\mathbf{q})\| \geq \beta_{\text{grad}} \doteq \frac{9}{10} c_* \theta, \quad \forall \mathbf{q} \in R_{II}.\quad (3.25)$$

Hence, provided

$$\Delta < \frac{m_H}{M_H} \sqrt{\frac{2\beta_{\text{grad}}}{L_H}},\quad (3.26)$$

we have $\mathbf{q}^{(k+1)} \in R_I$.

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

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

Hence, provided that

$$\Delta < \frac{2m_H^3}{L_H M_H^2},\quad (3.27)$$

we will have

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

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 3.12 and Lemma 7.11, we have

$$L_H = C_1 n^{3/2} \log^{3/2}(np) / \mu^2,\quad (3.28)$$

w.h.p. for some numerical constant C_1 . Combining this and our previous estimates of m_H, M_H , we conclude whenever

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

for some positive numerical constants c_1 and c_2 , w.h.p. our next trust-region step is also an unconstrained R_{Γ} step. Noting that c_{\star} and $c_{\#}$ can be made the same by our definition, we make the claimed simplification on Δ . This completes the proof. \blacksquare

Finally, we want to show that ultimate unconstrained R_{Γ} iterates actually converges to one nearby local minimizer rapidly. Lemma 3.12 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 3.14 *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 83 under Section 8. \blacksquare

To see this relates the magnitude of gradient to the distance away from the critical point, 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 3.14. Hence, combining this observation with Lemma 3.12, we can derive the asymptotic sequence convergence result as follows.

Proposition 3.15 *Assume (3.18) and the conditions in Lemma 3.13. Let $\mathbf{q}^{(k_0)} \in R_{\Gamma}$ and the k_0 -th step the first unconstrained R_{Γ} step and \mathbf{q}_{\star} be the unique local minimizer of f over one connected component of R_{Γ} that contains $\mathbf{q}^{(k_0)}$. Then w.h.p., for any positive integer $k' \geq 1$,*

$$\left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_{\star} \right\| \leq \frac{c c_{\#} \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2k'}, \quad (3.29)$$

provided that

$$\Delta \leq \frac{c' c_{\#}^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \quad (3.30)$$

Here $c_{\#}$ is as defined in Lemma 3.9 that can be made equal to $c_{s\star}$ as defined in Theorem 2.1, and c, c' are positive numerical constants.

Proof By the geometric characterization in Theorem 2.1 and corollary 2.2, f has $2n$ separated local minimizers, each located in R_{Γ} 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_{Γ} 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 k_0 is the index of the first unconstrained iterate in region R_I , i.e., $\mathbf{q}^{(k_0)} \in R_I$. By Lemma 3.12, for any integer $k' \geq 1$, we have

$$\left\| \text{grad } f \left(\mathbf{q}^{(k_0+k')} \right) \right\| \leq \frac{2m_H^2}{L_H} \left(\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(k_0)} \right) \right\| \right)^{2^{k'}}. \quad (3.31)$$

where L_H is as defined in Lemma 3.12, m_H as the strong convexity parameter for R_I defined above.

Now suppose \mathbf{q}_\star is the unique local minimizer of f , lies in the same R_I component that $\mathbf{q}^{(k_0)}$ is located. Let $\gamma_{k'}(t) = \exp_{\mathbf{q}_\star}(t\boldsymbol{\delta})$ to be the unique geodesic that connects \mathbf{q}_\star and $\mathbf{q}^{(k_0+k')}$ with $\gamma_{k'}(0) = \mathbf{q}_\star$ and $\gamma_{k'}(1) = \mathbf{q}^{(k_0+k')}$. We have

$$\begin{aligned} \left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_\star \right\| &\leq \left\| \exp_{\mathbf{q}_\star}(\boldsymbol{\delta}) - \mathbf{q}_\star \right\| = \sqrt{2 - 2 \cos \|\boldsymbol{\delta}\|} = 2 \sin(\|\boldsymbol{\delta}\|/2) \\ &\leq \|\boldsymbol{\delta}\| \leq \frac{1}{m_H} \left\| \text{grad } f \left(\mathbf{q}^{(k_0+k')} \right) \right\| \leq \frac{2m_H}{L_H} \left(\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(k_0)} \right) \right\| \right)^{2^{k'}}, \end{aligned}$$

where at the second line we have repeatedly applied Lemma 3.14.

By the optimality condition (3.16) and the fact that $\|\boldsymbol{\delta}^{(k_0)}\| < \Delta$, we have

$$\frac{L_H}{2m_H^2} \left\| \text{grad } f \left(\mathbf{q}^{(k_0)} \right) \right\| \leq \frac{L_H}{2m_H^2} M_H \left\| \left[\mathbf{U}^* \text{Hess } f \left(\mathbf{q}^{(k_0)} \right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f \left(\mathbf{q}^{(k_0)} \right) \right\| \leq \frac{L_H M_H}{2m_H^2} \Delta.$$

Thus, provided

$$\Delta < \frac{m_H^2}{L_H M_H}, \quad (3.32)$$

we can combine the above results and obtain

$$\left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_\star \right\| \leq \frac{2m_H}{L_H} 2^{-2^{k'}}.$$

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

$$\left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_\star \right\| \leq \frac{c_1 c_\# \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2^{k'}}.$$

Moreover, by (3.32), 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. ■

Now we are ready to piece together the above technical proposition to prove Theorem 3.1.

Proof [of Theorem 3.1] Assuming (3.18) and in addition that

$$\Delta < \min \left\{ \frac{c_1 c_\# \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \frac{c_2 c_\#^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)} \right\}$$

for small enough numerical constants c_1 and c_2 and $c_*, c_{\#}$ as defined in Theorem 2.1 and Lemma 3.9 respectively (c_* and $c_{\#}$ can be set to the same constant value), it can be verified that the conditions of all the above propositions are satisfied. Since each of the local minimizers is contained in the relative interior of one connected component of R_I (comparing distance of local minimizers to their respective signed basis vector, as stated in Corollary 2.2, with size of each connected R_I component yields this), we can define a threshold value

$$\zeta \doteq \min \left\{ \min_{\mathbf{q} \in \overline{R_{II}} \cup \overline{R_{III}}} f(\mathbf{q}), \max_{\mathbf{q} \in \overline{R_I}} f(\mathbf{q}) \right\}$$

where overline $\bar{\cdot}$ here denotes set closure. Obviously ζ is well-defined as the function f is continuous, and both sets $\overline{R_{II}} \cup \overline{R_{III}}$ and $\overline{R_I}$ are compact. Also for any of the local minimizers, say \mathbf{q}_* , it holds that $\zeta > f(\mathbf{q}_*)$.

By the four propositions above, 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 one 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, *in a finite number of steps, the function value must drop below ζ and all future iterates stay in R_I* . Indeed, if the function value never drops below ζ , by continuity the whole sequence must be of entirely Type A - whereby either the finite-length sequence converges to one local minimizer, or every iterate of the infinite sequence steadily decreases the objective value by at least a *fixed* amount - in either case, the objective value should ever drop below ζ in finitely many steps; hence contradiction arises. Once the function value drops below ζ , type A future steps decreases the objective value further down below ζ - by definition of ζ , these iterates stay within R_I , and type B future steps, aka unconstrained R_I steps obviously keep all subsequent iterates in R_I .

There are three possibilities after the objective value drop below ζ and all future iterates stay in R_I . Assume \mathbf{q}_* is the unique local minimizer in the same connected component of R_I as the current iterate: (1) the sequence always take constrained R_I steps and hits \mathbf{q}_* exactly in finitely many steps; (2) the sequence takes constrained R_I steps until reaching certain point $\mathbf{q}' \in R_I$ such that $f(\mathbf{q}') < f(\mathbf{q}_*) + d_I$, where d_I is as defined in Proposition 3.10. Since each unconstrained R_I step must decrease the objective value by at least d_I , the next and all future steps must be unconstrained R_I steps and the sequence converges to \mathbf{q}_* ; (3) the sequence starts to take unconstrained R_I steps at a certain point $\mathbf{q}'' \in R_I$ such that $f(\mathbf{q}'') \geq f(\mathbf{q}_*) + d_I$. In any case, the sequence converges to the local minimizer \mathbf{q}_* . By Proposition 3.7, Proposition 3.10, and Proposition 3.15, the number of iterations to obtain an ε -near solution to \mathbf{q}_* can be grossly bounded by

$$\begin{aligned} \#\text{Iter} &\leq \frac{f(\mathbf{q}^{(0)}) - f(\mathbf{q}_*)}{\min\{d_I, d_{II}, d_{III}\}} + \log \log \left(\frac{c_5 c_{\#} \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right) \\ &\leq \left[\min \left\{ \frac{c_3 c_*^3 \theta^3 \mu^4}{n^6 \log^3(np)}, \frac{c_4 c_{\#}^2 \theta^2}{n} \Delta^2 \right\} \right]^{-1} \left(f(\mathbf{q}^{(0)}) - f(\mathbf{q}_*) \right) + \log \log \left(\frac{c_5 c_{\#} \theta \mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right), \end{aligned}$$

where we have assumed $p \leq \exp(n)$ when comparing the various bounds. Finally, the claimed failure probability comes from a simple union bound with careful bookkeeping. ■

3.4 Extending to Convergence for Complete Dictionaries

Note that for any complete \mathbf{A}_0 with condition number $\kappa(\mathbf{A}_0)$, from Lemma 2.14 we know when p is large enough, w.h.p. one can write the preconditioned $\bar{\mathbf{Y}}$ as

$$\bar{\mathbf{Y}} = \mathbf{UV}^* \mathbf{X}_0 + \boldsymbol{\Xi} \mathbf{X}_0$$

for a certain $\boldsymbol{\Xi}$ with small magnitude, and $\mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^* = \text{SVD}(\mathbf{A}_0)$. Since \mathbf{UV}^* is orthogonal,

$$f(\mathbf{q}; \mathbf{UV}^* \mathbf{X}_0 + \boldsymbol{\Xi} \mathbf{X}_0) = f(\mathbf{VU}^* \mathbf{q}; \mathbf{X}_0 + \mathbf{VU}^* \boldsymbol{\Xi} \mathbf{X}_0).$$

In words, the function landscape of $f(\mathbf{q}; \mathbf{UV}^* \mathbf{X}_0 + \boldsymbol{\Xi} \mathbf{X}_0)$ is a rotated version of that of $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{VU}^* \boldsymbol{\Xi} \mathbf{X}_0)$. Thus, any local minimizer \mathbf{q}_* of $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{VU}^* \boldsymbol{\Xi} \mathbf{X}_0)$ is rotated to $\mathbf{UV}^* \mathbf{q}_*$, one minimizer of $f(\mathbf{q}; \mathbf{UV}^* \mathbf{X}_0 + \boldsymbol{\Xi} \mathbf{X}_0)$. Also if our algorithm generates iteration sequence $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{VU}^* \boldsymbol{\Xi} \mathbf{X}_0)$ upon initialization \mathbf{q}_0 , it will generate the iteration sequence $\mathbf{UV}^* \mathbf{q}_0, \mathbf{UV}^* \mathbf{q}_1, \mathbf{UV}^* \mathbf{q}_2, \dots$ for $f(\mathbf{q}; \mathbf{UV}^* \mathbf{X}_0 + \boldsymbol{\Xi} \mathbf{X}_0)$. So w.l.o.g. it is adequate that we prove the convergence results for the case $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{VU}^* \boldsymbol{\Xi} \mathbf{X}_0)$, corresponding to $\mathbf{A}_0 = \mathbf{I}$ with perturbation $\tilde{\boldsymbol{\Xi}} \doteq \mathbf{VU}^* \boldsymbol{\Xi}$. So in this section (Section 3.4), we write $f(\mathbf{q}; \tilde{\mathbf{X}}_0)$ to mean $f(\mathbf{q}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0)$.

Theorem 2.3 has shown that when

$$\theta \in \left(0, \frac{1}{2}\right), \mu \leq \min \left\{ \frac{c_a \theta}{n}, \frac{c_b}{n^{5/4}} \right\}, p \geq \frac{C}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right), \quad (3.33)$$

the geometric structure of the landscape is qualitatively unchanged and the c_* constant can be replaced with $c_*/2$. Particularly, for this choice of p , Lemma 2.14 implies

$$\|\tilde{\boldsymbol{\Xi}}\| = \|\mathbf{VU}^* \boldsymbol{\Xi}\| \leq \|\tilde{\boldsymbol{\Xi}}\| \leq cc_* \theta \left(\max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1} \quad (3.34)$$

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 3.3: Note that

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

where by (3.34) we have used $\|\tilde{\boldsymbol{\Xi}}\| \leq 1/(2\sqrt{n})$ to simplify the above result. So we obtain

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

- Lemma 3.4: 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 3.5 and Lemma 3.6 are generic and nothing changes.

- Proposition 3.7: We have now $\mathbf{w}^* \mathbf{g}(\mathbf{w}) / \|\mathbf{w}\| \geq c_* \theta / 2$ by Theorem 2.3 and w.h.p. $\mathbf{w}^* \nabla g(\mathbf{w}) / \|\mathbf{w}\|$ is $C_1 n^2 \log(np) / \mu$ -Lipschitz by Proposition 2.12 and the fact $\|\mathbf{X}_0 + \tilde{\mathbf{\Xi}} \mathbf{X}_0\|_\infty \leq 3 \|\mathbf{X}_0\|_\infty / 2$ shown above. Similarly, $\mathbf{w}^* \mathbf{g}(\mathbf{w}) / \|\mathbf{w}\| \leq -c_* \theta / 2$ by Theorem 2.3 and $\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} / \|\mathbf{w}\|^2$ is $C_2 n^3 \log^{3/2}(np) / \mu^2$ -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

$$\widetilde{d}_{\text{II}} = c_1 d_{\text{II}}, \quad \widetilde{d}_{\text{III}} = c_1 d_{\text{III}} \quad (3.35)$$

provided

$$\Delta < \frac{c_2 c_* \theta \mu^2}{n^{5/2} \log^{3/2}(np)}. \quad (3.36)$$

- Lemma 3.8: η_f is changed to $\tilde{\eta}_f$ with $\tilde{\eta}_f \leq 4\eta_f$ as shown above.
- Lemma 3.9: By (3.13), we have

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

where $L_{\tilde{h}}$ is the Lipschitz constant for the function $\tilde{h}_\mu(\cdot)$ and we have used the fact that $\|\tilde{\mathbf{\Xi}}\| \leq 1$. Similarly, by 3.12,

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

where L_{h_μ} is the Lipschitz constant for the function $h_\mu(\cdot)$. Since $L_{\tilde{h}} \leq 2/\mu^2$ and $L_{h_\mu} \leq 1/\mu$, and $\|\mathbf{X}_0\|_\infty \leq 4\sqrt{\log(np)}$ w.h.p. (Lemma 7.11). By (3.34), w.h.p. we have

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

provided the constant C in (3.33) for p is large enough. Thus, by (3.15) and the above estimates we have

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

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

$$\text{Hess } f(\mathbf{q}; \widetilde{\mathbf{X}}_0) \succeq \frac{1}{2} c_\# \frac{\theta}{\mu} \mathcal{P}_{T_q} \mathbb{S}^{n-1} \implies \widetilde{m}_H \geq \frac{1}{2} c_\# \frac{\theta}{\mu}. \quad (3.37)$$

- Proposition 3.10: From the estimate of M_H above Proposition 3.10 and the last point, we have

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

Also since $\widetilde{\eta}_f \leq 4\eta_f$ in Lemma 3.4 and Lemma 3.8, there are only multiplicative constant change to the various quantities. We conclude that

$$\widetilde{d}_{\text{I}} = c_3 d_{\text{I}} \tag{3.38}$$

provided that

$$\Delta \leq \frac{c_4 c_{\#}^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \tag{3.39}$$

- Lemma 3.11 is generic and nothing changes.
- Lemma 3.12: $\widetilde{L}_H \leq 27L_H/8$.
- Proposition 3.13: 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 RTM algorithm always takes unconstrained R_{I} step after taking one, provided that

$$\Delta \leq \frac{c_5 c_{\#}^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)}. \tag{3.40}$$

- Lemma 3.14: is generic and nothing changes.
- Proposition 3.15: 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 $k' \geq 1$,

$$\left\| \mathbf{q}^{(k_0+k')} - \mathbf{q}_{\star} \right\| \leq \frac{c_6 c_{\#} \theta \mu}{n^{3/2} \log^{3/2}(np)} 2^{-2^{k'}}, \tag{3.41}$$

provided

$$\Delta \leq \frac{c_7 c_{\#}^2 \theta^2 \mu}{n^{5/2} \log^{5/2}(np)}. \tag{3.42}$$

The final proof to Theorem 2.3 is almost identical to that of Theorem 2.1, except for

$$\Delta \leq \min \left\{ \frac{c_8 c_{\star} \theta \mu^2}{n^{5/2} \log^{3/2}(np)}, \frac{c_9 c_{\#}^3 \theta^3 \mu}{n^{7/2} \log^{7/2}(np)} \right\}, \tag{3.43}$$

$$\widetilde{\zeta} \doteq \min \left\{ \min_{\mathbf{q} \in R_{\text{II}} \cup R_{\text{III}}} f(\mathbf{q}; \widetilde{\mathbf{X}}_0), \max_{\mathbf{q} \in R_{\text{I}}} f(\mathbf{q}; \widetilde{\mathbf{X}}_0) \right\}, \tag{3.44}$$

and hence all ζ is now changed to $\tilde{\zeta}$, and also d_{I} , d_{II} , and d_{III} are changed to \tilde{d}_{I} , \tilde{d}_{II} , and \tilde{d}_{III} as defined above, respectively. The final iteration complexity to each an ε -near solution is hence

$$\#\text{Iter} \leq \left[\min \left\{ \frac{c_{10}c_*^3\theta^3\mu^4}{n^6 \log^3(np)}, \frac{c_{11}c_*^2\theta^2}{n} \Delta^2 \right\} \right]^{-1} \left(f(\mathbf{q}^{(0)}) - f(\mathbf{q}_*) \right) + \log \log \left(\frac{c_{12}c_*\theta\mu}{\varepsilon n^{3/2} \log^{3/2}(np)} \right).$$

Hence overall the qualitative behavior of the algorithm is not changed, as compared to that for the orthogonal case. Above c_1 through c_{12} are all numerical constants.

4 Complete Algorithm Pipeline and Main Results

For orthogonal dictionaries, from Theorem 2.1 and its corollary, we know that all the minimizers $\hat{\mathbf{q}}_*$ are $O(\mu)$ away from their respective nearest “target” \mathbf{q}_* , with $\mathbf{q}_*^* \hat{\mathbf{Y}} = \alpha \mathbf{e}_i^* \mathbf{X}_0$ for certain $\alpha \neq 0$ and $i \in [n]$; in Theorem 3.1, we have shown that w.h.p. the Riemannian TRM algorithm produces a solution $\hat{\mathbf{q}} \in \mathbb{S}^{n-1}$ that is ε away to one of the minimizers, say $\hat{\mathbf{q}}_*$. Thus, the $\hat{\mathbf{q}}$ 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 exactly produce the optimizer \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 4.1. The same procedure can be used to recover complete dictionaries, though the analysis is slightly more complicated; we present the results in Section 4.2. 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 2.1 (resp. Theorem 2.3) and Theorem 3.1 (resp. Theorem 3.2), starting from any, 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 one 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:

$$\text{minimize}_{\mathbf{q}} \left\| \mathbf{q}^* \hat{\mathbf{Y}} \right\|_1, \quad \text{subject to} \quad \langle \mathbf{r}, \mathbf{q} \rangle = 1, \quad (4.1)$$

with $\mathbf{r} = \hat{\mathbf{q}}$. We show in Lemma 4.2 (resp. Lemma 4.4) that when $\langle \hat{\mathbf{q}}, \mathbf{q}_* \rangle$ is sufficiently large, implied by μ being sufficiently small, w.h.p. the minimizer of (4.1) 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 (4.1) 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 $\mathbf{A}_0 = \mathbf{Y}\mathbf{X}_0^* (\mathbf{X}_0\mathbf{X}_0^*)^{-1}$.

4.1 Recovering Orthogonal Dictionaries

Theorem 4.1 (Main theorem - recovering orthogonal dictionaries) *Assume the dictionary \mathbf{A}_0 is orthogonal and we take $\hat{\mathbf{Y}} = \mathbf{Y}$. Suppose $\theta \in (0, 1/3)$, $\mu_\star < \min \{c_a\theta n^{-1}, c_b n^{-5/4}\}$, and $p \geq Cn^3 \log \frac{n}{\mu_\star\theta} / (\mu_\star^2\theta^2)$. The above algorithmic pipeline with parameter setting*

$$\Delta_\star \leq \min \left\{ \frac{c_c c_\star \theta \mu_\star^2}{n^{5/2} \log^{5/2}(np)}, \frac{c_d c_\star^3 \theta^3 \mu_\star}{n^{7/2} \log^{7/2}(np)} \right\}, \quad (4.2)$$

recovers the dictionary \mathbf{A}_0 and \mathbf{X}_0 in polynomial time, with failure probability bounded by $c_e p^{-6}$. Here c_\star is as defined in Theorem 2.1, and c_a through c_e , and C are all positive numerical constants.

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

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

Lemma 4.2 (LP rounding - orthogonal dictionary) *There exists a positive constant C , such that for all $\theta \in (0, 1/3)$, and $p \geq Cn^2 \log(n/\theta)/\theta$, with probability at least $1 - 2p^{-10} - \theta(n-1)^{-7}p^{-7} - \exp(-0.3\theta(n-1)p)$, the rounding procedure (4.1) returns \mathbf{q}_\star for any input vector \mathbf{r} that satisfies*

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

Proof See Page 85 under Section 9. ■

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

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 $\{\mathbf{q}_\star^i\}_{i=1}^n$ are orthogonal to each other. W.l.o.g., suppose we have found the first ℓ unit vectors $\mathbf{q}_\star^1, \dots, \mathbf{q}_\star^\ell$ which recover the first ℓ rows of \mathbf{X}_0 . Correspondingly, we partition the target dictionary \mathbf{A}_0 and \mathbf{X}_0 as

$$\mathbf{A}_0 = [\mathbf{V}, \mathbf{V}^\perp], \quad \mathbf{X}_0 = \begin{bmatrix} \mathbf{X}_0^{[\ell]} \\ \mathbf{X}_0^{[n-\ell]} \end{bmatrix}, \quad (4.3)$$

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

$$f_{n-\ell}^\downarrow(\mathbf{z}; \mathbf{W}) \doteq \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{z}^* \mathbf{w}_k), \quad (4.4)$$

for any matrix $\mathbf{W} \in \mathbb{R}^{(n-\ell) \times p}$. Then by (1.4), 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 4.1, the conditions of Theorem 2.1 and Theorem 3.1 for all cases with reduced dimensions are still valid. So w.h.p., the TRM algorithm returns a $\hat{\mathbf{z}}$ such that $\|\hat{\mathbf{z}} - \mathbf{z}_*\| \in O(\mu_*)$ 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 $\mathbf{U}\hat{\mathbf{z}}$ is our estimation obtained from the TRM algorithm. Moreover,

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

Thus, by Lemma 4.2, one successfully recovers $\mathbf{U}\mathbf{z}_*$ from $\mathbf{U}\hat{\mathbf{z}}$ 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 .

4.2 Recovering Complete Dictionaries

By working with the preconditioned data samples $\hat{\mathbf{Y}} = \bar{\mathbf{Y}} \doteq \sqrt{\theta p} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y}$,²⁸ we can use a similar procedure described above to recover complete dictionaries.

Theorem 4.3 (Main theorem - recovering complete dictionaries) *Assume the dictionary \mathbf{A}_0 is complete with condition number $\kappa(\mathbf{A}_0)$ and we take $\hat{\mathbf{Y}} = \bar{\mathbf{Y}}$. Suppose $\theta \in (0, 1/3)$, $\mu_* < \min\{c_a\theta n^{-1}, c_b n^{-5/4}\}$, and $p \geq \frac{C}{c_*^2\theta} \max\left\{\frac{n^4}{\mu^4}, \frac{n^5}{\mu^2}\right\} \kappa^8(\mathbf{A}_0) \log^4\left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta}\right)$. The algorithmic pipeline with parameter setting*

$$\Delta_* \leq \min \left\{ \frac{c_c c_* \theta \mu_*^2}{n^{5/2} \log^{5/2}(np)}, \frac{c_d c_*^3 \theta^3 \mu_*}{n^{7/2} \log^{7/2}(np)} \right\}, \quad (4.5)$$

recovers the dictionary \mathbf{A}_0 and \mathbf{X}_0 in polynomial time, with failure probability bounded by $c_e p^{-6}$. Here c_ is as defined in Theorem 2.1, and c_a through c_f , and C are all positive numerical constants.*

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

²⁸In 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.

Proof of LP rounding The result of the LP rounding is only slightly different from that of the orthogonal case in Lemma 4.2, so is the proof.

Lemma 4.4 (LP rounding - complete dictionary) *There exists a positive constant C , such that for all $\theta \in (0, 1/3)$, and $p \geq \frac{C}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta} \right)$, with probability at least $1 - 3p^{-8} - \theta(n-1)^{-7}p^{-7} - \exp(-0.3\theta(n-1)p)$, the rounding procedure (4.1) returns \mathbf{q}_* for any input vector \mathbf{r} that satisfies*

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

Proof See Page 86 under Section 9. ■

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 (4.3).

By Lemma 2.14 and (2.15), 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 (4.4). 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 4.5 *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*

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}, \quad (4.6)$$

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

$$\|\mathbf{\Delta}\|_{\ell^1 \rightarrow \ell^2} \leq 16\sqrt{n}\|\mathbf{\Xi}\|\|\mathbf{X}_0\|_\infty, \quad \|\mathbf{\Delta}\| \leq 16\|\mathbf{\Xi}\|\|\mathbf{X}_0\|, \quad (4.7)$$

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 87 under Section 9. ■

Since \mathbf{UV} 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 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 2.3 and Theorem 3.2 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 2.3

and Theorem 3.2, the TRM algorithm provably returns one \hat{z} such that \hat{z} is near to a perturbed optimal \hat{z}_* with

$$\hat{z}_*^* U^* V X_0^{[n-\ell]} = z_*^* U^* V X_0^{[n-\ell]} + z_*^* \Delta = \alpha e_i^* X_0, \quad \text{for some } i \notin [\ell], \quad (4.8)$$

where z_* with $\|z_*\| = 1$ is the exact solution. More specifically, Corollary 2.4 implies

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

Next, we show that \hat{z} is also very near to the exact solution z_* . Indeed, the identity (4.8) suggests

$$\begin{aligned} (\hat{z}_* - z_*)^* U^* V X_0^{[n-\ell]} &= z_*^* \Delta \\ \implies \hat{z}_* - z_* &= \left[(X_0^{[n-\ell]})^* V^* U \right]^\dagger \Delta^* z_* = U^* V \left[(X_0^{[n-\ell]})^* \right]^\dagger \Delta^* z_* \end{aligned} \quad (4.9)$$

where $W^\dagger = (W^* W)^{-1} W^*$ denotes the pseudo inverse of a matrix W with full column rank. Hence, by (4.9) we can bound the distance between \hat{z}_* and z_* by

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

By Lemma B.3, 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 4.5, 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 $\|\hat{z}_* - z_*\| \leq 28 \|\Xi\|$. Thus, combining the results above, we obtain

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

Lemma 2.14, and in particular (2.15), for our choice of p as in Theorem 2.3, $\|\Xi\| \leq c\mu_*^2 n^{-3/2}$, where c can be made smaller by making the constant in p larger. For μ_* sufficiently small, we conclude that

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

In words, the TRM algorithm returns a \hat{z} such that $U\hat{z}$ 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

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

and hence by Lemma 4.4, the LP rounding exactly returns the optimal solution q_*^i upon the input $U\hat{z}$.

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 .




Figure 5: Phase transition for recovering a single sparse vector under the dictionary learning model with the sample complexity $p = 5n^3$

5 Simulations

To corroborate our theory, we experiment with dictionary recovery on simulated data. 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 $\mathbf{A}_0 = \mathbf{I} \in \mathbb{R}^{n \times n}$. We fix $p = 5n^3$, and each column of the coefficient matrix $\mathbf{X}_0 \in \mathbb{R}^{n \times p}$ 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 produce similar behavior. For the sparsity surrogate defined in (1.5), we fix the parameter $\mu = 10^{-2}$. We implement Algorithm 1 with adaptive step size instead of the fixed step size in our analysis.

To see how the allowable sparsity level varies with the dimension, which our theory primarily is about, we vary the dictionary dimension n and the sparsity k both between 1 and 120; for every pair of (k, n) we repeat the simulations independently for $T = 5$ times. Because the optimal solutions are signed coordinate vectors $\{e_i\}_{i=1}^n$, for a solution $\hat{\mathbf{q}}$ returned by the TRM algorithm, we define the reconstruction error (RE) to be

$$\text{RE} = \min_{1 \leq i \leq n} (\|\hat{\mathbf{q}} - e_i\|, \|\hat{\mathbf{q}} + e_i\|). \quad (5.1)$$

The trial is determined to be a success once $\text{RE} \leq \mu$, with the idea that this indicates $\hat{\mathbf{q}}$ 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). Figure 5 shows the phase transition in the (n, k) plane for the orthogonal case. It is obvious that our TRM algorithm can work well into the linear region whenever $p \in O(n^3)$. Our analysis is tight up to logarithm factors, and also the polynomial dependency on $1/\mu$, which under the theory is polynomial in n .

6 Discussion

For recovery of complete dictionaries, the LP program approach in [SWW12] that works with $\theta \leq O(1/\sqrt{n})$ only demands $p \geq \Omega(n^2 \log n^2)$, which is recently improved to $p \geq \Omega(n \log^4 n)$ [LV15], almost matching the lower bound $\Omega(n \log n)$ (i.e., when $\theta \sim 1/n$). The sample complexity stated in Theorem 4.3 is obviously much higher. It is interesting to see whether such growth in complexity is intrinsic to working in the linear regime. Though our experiments seemed to suggest the necessity of $p \sim O(n^3)$ even for the orthogonal case, there could be other efficient algorithms that demand much less. Tweaking these three points will likely improve the complexity: (1) The ℓ^1 proxy. The derivative and Hessians of the log cosh function we adopted entail the tanh function, which is not amenable to effective approximation and affects the sample complexity; (2) Geometric characterization and algorithm analysis. It seems working directly on the sphere (i.e., in the \mathbf{q} space) could simplify and possibly improve certain parts of the analysis; (3) treating the complete case directly, rather than using (pessimistic) bounds to treat it as a perturbation of the orthogonal case. Particularly, general linear transforms may change the space significantly, such that preconditioning and comparing to the orthogonal transforms may not be the most efficient way to proceed.

It is possible to extend the current analysis to other dictionary settings. Our geometric structures and algorithms allow plug-and-play noise analysis. Nevertheless, we believe a more stable way of dealing with noise is to directly extract the whole dictionary, i.e., to consider geometry and optimization (and perturbation) over the orthogonal group. This will require additional nontrivial technical work, but likely feasible thanks to the relatively complete knowledge of the orthogonal group [EAS98, AMS09]. A substantial leap forward would be to extend the methodology to recovery of *structured* overcomplete dictionaries, such as tight frames. Though there is no natural elimination of one variable, one can consider the marginalization of the objective function wrt the coefficients and work with hidden functions.²⁹ For the coefficient model, as we alluded to in Section 1.5, our analysis and results likely can be carried through to coefficients with statistical dependence and physical constraints.

The connection to ICA we discussed in Section 1.5 suggests our geometric characterization and algorithms can be modified for the ICA problem. This likely will provide new theoretical insights and computational schemes to ICA. In the surge of theoretical understanding of nonconvex heuristics [KMO10, JNS13, Har14, HW14, NNS⁺14, JN14, NJS13, CLS14, JO14, AGJ14, YCS13, LWB13, QSW14, LWB13, AAJ⁺13, AAN13, AGM13, AGMM15, ABGM14], the initialization plus local refinement strategy mostly differs from practice, whereby random initializations seem to work well, and the analytic techniques developed are mostly fragmented and highly specialized. The analytic and algorithmic we developed here hold promise to provide a coherent account of these problems. It is interesting to see to what extent we can streamline and generalize the framework.

Our motivating experiment on real images in Section 1.2 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 local maxima and saddle points (at least with high probability). The recent work [GHJY15] has showed that randomly perturbing each iterate can help gradient

²⁹This recent work [AGMM15] on overcomplete DR has used a similar idea. The marginalization taken there is near to the global optimum of one variable, where the function is well-behaved. Studying the global properties of the marginalization may introduce additional challenges.

algorithm to escape saddle points with high probability. It would be interesting to know whether similar results can be obtained for gradient descent algorithms with random initialization. The continuous counterpart seems well understood; see, e.g., [HMG94] for discussions of Morse-Bott theorem and gradient flow convergence.

7 Proofs of Main Technical Results for High Dimensional Geometry

In this section, we provide complete proofs for technical results stated in Section 2. Before that, let us introduce some notations and common results that will be used later throughout this section. Since we deal with BG random variables and random vectors, it is often convenient to write such vector explicitly as $\mathbf{x} = [\Omega_1 v_1, \dots, \Omega_n v_n] = \boldsymbol{\Omega} \odot \mathbf{v}$, where $\Omega_1, \dots, \Omega_n$ are i.i.d. Bernoulli random variables and v_1, \dots, v_n are i.i.d. standard normal. For a particular realization of such random vector, we will denote the support as $\mathcal{I} \subset [n]$. Due to the particular coordinate map in use, we will often refer to subset $\mathcal{J} \doteq \mathcal{I} \setminus \{n\}$ and the random vectors $\bar{\mathbf{x}} \doteq [\Omega_1 v_1, \dots, \Omega_{n-1} v_{n-1}]$ and $\bar{\mathbf{v}} \doteq [v_1, \dots, v_{n-1}]$ in \mathbb{R}^{n-1} . By Lemma A.1, it is not hard to see that

$$\nabla_{\mathbf{w}} h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) = \tanh\left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu}\right) \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w}\right), \quad (7.1)$$

$$\begin{aligned} \nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) &= \frac{1}{\mu} \left[1 - \tanh^2\left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu}\right)\right] \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w}\right) \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w}\right)^* \\ &\quad - x_n \tanh\left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu}\right) \left(\frac{1}{q_n(\mathbf{w})} \mathbf{I} + \frac{1}{q_n^3(\mathbf{w})} \mathbf{w} \mathbf{w}^*\right). \end{aligned} \quad (7.2)$$

7.1 Proofs for Section 2.2

7.1.1 Proof of Proposition 2.5

The proof involves some delicate analysis, particularly polynomial approximation of the function $f(t) = \frac{1}{(1+t)^2}$ over $t \in [0, 1]$. This is naturally induced by the $1 - \tanh^2(\cdot)$ function. The next lemma characterizes one polynomial approximation of $f(t)$.

Lemma 7.1 *Consider $f(t) = \frac{1}{(1+t)^2}$ for $t \in [0, 1]$. For every $T > 1$, there is a sequence b_0, b_1, \dots , with $\|\mathbf{b}\|_{\ell^1} = T < \infty$, such that the polynomial $p(t) = \sum_{k=0}^{\infty} b_k t^k$ satisfies*

$$\|f - p\|_{L^1[0,1]} \leq \frac{1}{2\sqrt{T}}, \quad \|f - p\|_{L^\infty[0,1]} \leq \frac{1}{\sqrt{T}},$$

In particular, one can choose $b_k = (-1)^k (k+1) \beta^k$ with $\beta = 1 - 1/\sqrt{T} < 1$ such that

$$p(t) = \frac{1}{(1 + \beta t)^2} = \sum_{k=0}^{\infty} (-1)^k (k+1) \beta^k t^k.$$

Moreover, such sequence satisfies $0 < \sum_{k=0}^{\infty} \frac{b_k}{(1+k)^3} < \sum_{k=0}^{\infty} \frac{|b_k|}{(1+k)^3} < 2$.

Lemma 7.2 *Let $X \sim \mathcal{N}(0, \sigma_X^2)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$. We have*

$$\mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) X^2 \mathbb{1}_{X+Y>0} \right] \leq \frac{1}{\sqrt{2\pi}} \frac{\mu \sigma_X^2 \sigma_Y^2}{(\sigma_X^2 + \sigma_Y^2)^{3/2}} + \frac{\mu^3 \sigma_X^2 \sigma_Y^2}{(\sigma_X^2 + \sigma_Y^2)^{3/2}} + \frac{3}{4\sqrt{2\pi}} \frac{\sigma_X^2 \mu^3}{(\sigma_X^2 + \sigma_Y^2)^{5/2}} (3\mu^2 + 4\sigma_X^2).$$

Proof For $x + y > 0$, let $z = \exp\left(-2\frac{x+y}{\mu}\right) \in [0, 1]$, then $1 - \tanh^2\left(\frac{x+y}{\mu}\right) = \frac{4z}{(1+z)^2}$. Fix any $T > 1$ to be determined later, by Lemma 7.1, we choose the polynomial $p_\beta(z) = \frac{1}{(1+\beta z)^2}$ with $\beta = 1 - 1/\sqrt{T}$ to upper bound $f(z) = \frac{1}{(1+z)^2}$. So we have

$$\begin{aligned} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) X^2 \mathbb{1}_{X+Y>0} \right] &= 4\mathbb{E} [Zf(Z) X^2 \mathbb{1}_{X+Y>0}] \\ &\leq 4\mathbb{E} [Zp_\beta(Z) X^2 \mathbb{1}_{X+Y>0}] \\ &= 4 \sum_{k=0}^{\infty} \left\{ b_k \mathbb{E} [Z^{k+1} X^2 \mathbb{1}_{X+Y>0}] \right\}, \end{aligned}$$

where $b_k = (-1)^k (k+1)\beta^k$, and the exchange of infinite summation and expectation above is justified in view that

$$\sum_{k=0}^{\infty} |b_k| \mathbb{E} [Z^{k+1} X^2 \mathbb{1}_{X+Y>0}] \leq \sum_{k=0}^{\infty} |b_k| \mathbb{E} [X^2 \mathbb{1}_{X+Y>0}] \leq \sigma_X^2 \sum_{k=0}^{\infty} |b_k| < \infty$$

and the dominated convergence theorem (see, e.g., theorem 2.24 and 2.25 of [Fol99]). By Lemma B.1, we have

$$\begin{aligned} &\sum_{k=0}^{\infty} \left\{ b_k \mathbb{E} [Z^{k+1} X^2 \mathbb{1}_{X+Y>0}] \right\} \\ &= \sum_{k=0}^{\infty} (-\beta)^k (k+1) \left[\left(\sigma_X^2 + \frac{4(k+1)^2}{\mu^2} \sigma_X^4 \right) \exp \left(\frac{2(k+1)^2}{\mu^2} (\sigma_X^2 + \sigma_Y^2) \right) \Phi^c \left(\frac{2(k+1)}{\mu} \sqrt{\sigma_X^2 + \sigma_Y^2} \right) \right. \\ &\quad \left. - \frac{2(k+1)}{\mu} \frac{\sigma_X^4}{\sqrt{2\pi} \sqrt{\sigma_X^2 + \sigma_Y^2}} \right] \\ &\leq \frac{1}{\sqrt{2\pi}} \sum_{k=0}^{\infty} (-\beta)^k (k+1) \left[\frac{\sigma_X^2 \mu}{2(k+1) \sqrt{\sigma_X^2 + \sigma_Y^2}} - \frac{\sigma_X^2 \mu^3}{8(k+1)^3 (\sigma_X^2 + \sigma_Y^2)^{3/2}} - \frac{\mu \sigma_X^4}{2(k+1) (\sigma_X^2 + \sigma_Y^2)^{3/2}} \right] \\ &\quad + \frac{3}{\sqrt{2\pi}} \sum_{k=0}^{\infty} \beta^k (k+1) \left(\sigma_X^2 + \frac{4(k+1)^2}{\mu^2} \sigma_X^4 \right) \frac{\mu^5}{32(k+1)^5 (\sigma_X^2 + \sigma_Y^2)^{5/2}}, \end{aligned}$$

where we have applied Type I upper and lower bounds for $\Phi^c(\cdot)$ to even k and odd k respectively and rearrange the terms to obtain the last line. Using the following estimates (see Lemma 7.1)

$$\sum_{k=0}^{\infty} (-\beta)^k = \frac{1}{1+\beta}, \quad \sum_{k=0}^{\infty} \frac{b_k}{(k+1)^3} \geq 0, \quad \sum_{k=0}^{\infty} \frac{|b_k|}{(k+1)^5} \leq \sum_{k=0}^{\infty} \frac{|b_k|}{(k+1)^3} \leq 2,$$

we obtain

$$\sum_{k=0}^{\infty} \left\{ b_k \mathbb{E} \left[Z^{k+1} X^2 \mathbb{1}_{X+Y>0} \right] \right\} \leq \frac{1}{2\sqrt{2\pi}} \frac{\mu \sigma_X^2 \sigma_Y^2}{(\sigma_X^2 + \sigma_Y^2)^{3/2}} \frac{1}{1+\beta} + \frac{3}{16\sqrt{2\pi}} \frac{\sigma_X^2 \mu^3}{(\sigma_X^2 + \sigma_Y^2)^{5/2}} (3\mu^2 + 4\sigma_X^2).$$

Noticing $\frac{1}{1+\beta} < \frac{1}{2} + \frac{1}{2\sqrt{T}}$ and choosing $T = \mu^{-4}$, we obtain the desired result. \blacksquare

Lemma 7.3 *Let $X \sim \mathcal{N}(0, \sigma_X^2)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$. We have*

$$\mathbb{E} \left[\tanh \left(\frac{X+Y}{\mu} \right) X \right] \geq \frac{2\sigma_X^2}{\sqrt{2\pi} \sqrt{\sigma_X^2 + \sigma_Y^2}} - \frac{4\mu^2 \sigma_X^2}{\sqrt{2\pi} \sqrt{\sigma_X^2 + \sigma_Y^2}} - \frac{2\sigma_X^2 \mu^2}{\sqrt{2\pi} (\sigma_X^2 + \sigma_Y^2)^{3/2}} - \frac{3\sigma_X^2 \mu^4}{2\sqrt{2\pi} (\sigma_X^2 + \sigma_Y^2)^{5/2}}.$$

Proof By Lemma B.1, we know

$$\mathbb{E} \left[\tanh \left(\frac{X+Y}{\mu} \right) X \right] = \frac{\sigma_X^2}{\mu} \mathbb{E} \left[1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right]$$

Similar to the proof of the above lemma, for $x+y > 0$, let $z = \exp\left(-2\frac{x+y}{\mu}\right)$ and $f(z) = \frac{1}{(1+z)^2}$. Fixing any $T > 1$, we will use $4zp_\beta(z) = \frac{4z}{(1+\beta z)^2}$ to approximate the $1 - \tanh^2\left(\frac{x+y}{\mu}\right) = 4zf(z)$ function from above, where again $\beta = 1 - 1/\sqrt{T}$. So we obtain

$$\begin{aligned} \mathbb{E} \left[1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right] &= 8\mathbb{E} [f(Z) Z \mathbb{1}_{X+Y>0}] \\ &= 8\mathbb{E} [p_\beta(Z) Z \mathbb{1}_{X+Y>0}] - 8\mathbb{E} [(p_\beta(Z) - f(Z)) Z \mathbb{1}_{X+Y>0}]. \end{aligned}$$

Now for the first term, we have

$$\mathbb{E} [p_\beta(Z) Z \mathbb{1}_{X+Y>0}] = \sum_{k=0}^{\infty} b_k \mathbb{E} \left[Z^{k+1} \mathbb{1}_{X+Y>0} \right],$$

justified as $\sum_{k=0}^{\infty} |b_k| \mathbb{E} [Z^{k+1} \mathbb{1}_{X+Y>0}] \leq \sum_{k=0}^{\infty} |b_k| < \infty$ making the dominated convergence theorem (see, e.g., theorem 2.24 and 2.25 of [Fol99]) applicable. To proceed, from Lemma B.1, we obtain

$$\begin{aligned} &\sum_{k=0}^{\infty} b_k \mathbb{E} \left[Z^{k+1} \mathbb{1}_{X+Y>0} \right] \\ &= \sum_{k=0}^{\infty} (-\beta)^k (k+1) \exp \left(\frac{2}{\mu^2} (k+1)^2 (\sigma_X^2 + \sigma_Y^2) \right) \Phi^c \left(\frac{2}{\mu} (k+1) \sqrt{\sigma_X^2 + \sigma_Y^2} \right) \end{aligned}$$

$$\begin{aligned} &\geq \frac{1}{\sqrt{2\pi}} \sum_{k=0}^{\infty} (-\beta)^k (k+1) \left(\frac{\mu}{2(k+1)\sqrt{\sigma_X^2 + \sigma_Y^2}} - \frac{\mu^3}{8(k+1)^3(\sigma_X^2 + \sigma_Y^2)^{3/2}} \right) \\ &\quad - \frac{3}{\sqrt{2\pi}} \sum_{k=0}^{\infty} \beta^k (k+1) \frac{\mu^5}{32(k+1)^5(\sigma_X^2 + \sigma_Y^2)^{5/2}}, \end{aligned}$$

where we have applied Type I upper and lower bounds for $\Phi^c(\cdot)$ to odd k and even k respectively and rearrange the terms to obtain the last line. Using the following estimates (see Lemma 7.1)

$$\sum_{k=0}^{\infty} (-\beta)^k = \frac{1}{1+\beta}, \quad 0 \leq \sum_{k=0}^{\infty} \frac{b_k}{(k+1)^3} \leq \sum_{k=0}^{\infty} \frac{|b_k|}{(k+1)^5} \leq \sum_{k=0}^{\infty} \frac{|b_k|}{(k+1)^3} \leq 2,$$

we obtain

$$\begin{aligned} \sum_{k=0}^{\infty} b_k \mathbb{E} \left[Z^{k+1} \mathbb{1}_{X+Y>0} \right] &\geq \\ &\frac{\mu}{2\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}} \frac{1}{1+\beta} - \frac{\mu^3}{4\sqrt{2\pi}(\sigma_X^2 + \sigma_Y^2)^{3/2}} - \frac{3\mu^5}{16\sqrt{2\pi}(\sigma_X^2 + \sigma_Y^2)^{5/2}}. \end{aligned}$$

To proceed, by Lemma B.1 and Lemma 7.1, we have

$$\mathbb{E}[(p_\beta(Z) - f(Z)) Z \mathbb{1}_{X+Y>0}] \leq \|p - f\|_{L^\infty[0,1]} \mathbb{E}[Z \mathbb{1}_{X+Y>0}] \leq \frac{\mu}{2\sqrt{2\pi T} \sqrt{\sigma_X^2 + \sigma_Y^2}},$$

where we have also used Type I upper bound for $\Phi^c(\cdot)$. Combining the above estimates, we get

$$\begin{aligned} \mathbb{E} \left[\tanh \left(\frac{X+Y}{\mu} \right) X \right] &\geq \\ &\frac{4\sigma_X^2}{\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}} \left(\frac{1}{1+\beta} - \frac{1}{\sqrt{T}} \right) - \frac{2\sigma_X^2\mu^2}{\sqrt{2\pi}(\sigma_X^2 + \sigma_Y^2)^{3/2}} - \frac{3\sigma_X^2\mu^4}{2\sqrt{2\pi}(\sigma_X^2 + \sigma_Y^2)^{5/2}}. \end{aligned}$$

Noticing $\frac{1}{1+\beta} > \frac{1}{2}$ and taking $T = \mu^{-4}$, we obtain the claimed result. \blacksquare

Proof [of Proposition 2.5] For any $i \in [n-1]$, we have

$$\int_0^1 \int_{\mathbf{x}} \left| \frac{\partial}{\partial w_i} h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) \right| \mu(d\mathbf{x}) dw_i \leq \int_0^1 \int_{\mathbf{x}} \left(|x_i| + |x_n| \frac{1}{q_n(\mathbf{w})} \right) \mu(d\mathbf{x}) dw_i < \infty.$$

Hence by Lemma A.4 we obtain $\frac{\partial}{\partial w_i} \mathbb{E}[h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] = \mathbb{E} \left[\frac{\partial}{\partial w_i} h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) \right]$. Moreover for any $j \in [n-1]$,

$$\begin{aligned} \int_0^1 \int_{\mathbf{x}} \left| \frac{\partial^2}{\partial w_j \partial w_i} h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) \right| \mu(d\mathbf{x}) dw_j &\leq \\ \int_0^1 \int_{\mathbf{x}} \left[\frac{1}{\mu} \left(|x_i| + \frac{|x_n|}{q_n(\mathbf{w})} \right) \left(|x_j| + \frac{|x_n|}{q_n(\mathbf{w})} \right) + |x_n| \left(\frac{1}{q_n(\mathbf{w})} + \frac{1}{q_n^3(\mathbf{w})} \right) \right] \mu(d\mathbf{x}) dw_i &< \infty. \end{aligned}$$

Invoking Lemma A.4 again we obtain

$$\frac{\partial^2}{\partial w_j \partial w_i} \mathbb{E} [h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] = \frac{\partial}{\partial w_j} \mathbb{E} \left[\frac{\partial}{\partial w_i} h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) \right] = \mathbb{E} \left[\frac{\partial^2}{\partial w_j \partial w_i} h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x}) \right].$$

The above holds for any pair of $i, j \in [n-1]$, so it follows that

$$\nabla_{\mathbf{w}}^2 \mathbb{E} [h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] = \mathbb{E} [\nabla_{\mathbf{w}}^2 h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})].$$

Hence it is easy to see that

$$\begin{aligned} & \mathbf{w}^* \nabla_{\mathbf{w}}^2 \mathbb{E} [h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \mathbf{w} \\ &= \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \left(\mathbf{w}^* \bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \|\mathbf{w}\|^2 \right)^2 \right] - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \frac{x_n}{q_n^3(\mathbf{w})} \|\mathbf{w}\|^2 \right]. \end{aligned}$$

Now the first term is

$$\begin{aligned} & \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \left(\mathbf{w}^* \bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \|\mathbf{w}\|^2 \right)^2 \right] \\ &= \frac{2(1-\theta)}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) \right) (\mathbf{w}^* \bar{\mathbf{x}})^2 \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} > 0} \right] \\ & \quad - \frac{4\theta}{\mu} \frac{\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}} + q_n(\mathbf{w}) x_n}{\mu} \right) \right) (\mathbf{w}^* \bar{\mathbf{x}}) (q_n(\mathbf{w}) x_n) \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} + q_n(\mathbf{w}) x_n > 0} \right] \\ & \quad + \frac{2\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) \right) (\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}})^2 \mathbb{1}_{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n > 0} \right] \\ & \quad + \frac{2\theta}{\mu} \frac{\|\mathbf{w}\|^4}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) \right) (q_n(\mathbf{w}) v_n)^2 \mathbb{1}_{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n > 0} \right] \\ &\leq \frac{8(1-\theta)}{\mu} \mathbb{E} \left[\exp \left(-2 \frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}})^2 \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} > 0} \right] \\ & \quad + \frac{8\theta}{\mu} \frac{\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \mathbb{E} \left[\exp \left(-\frac{2}{\mu} (\mathbf{w}^* \bar{\mathbf{x}} + q_n(\mathbf{w}) x_n) \right) (\mathbf{w}^* \bar{\mathbf{x}} + q_n(\mathbf{w}) x_n)^2 \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} + q_n(\mathbf{w}) x_n > 0} \right] \\ & \quad + \frac{2\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{X,Y} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) Y^2 \mathbb{1}_{X+Y > 0} \right] \\ & \quad + \frac{2\theta}{\mu} \frac{\|\mathbf{w}\|^4}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{X,Y} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) X^2 \mathbb{1}_{X+Y > 0} \right], \end{aligned}$$

where conditioned on each support set \mathcal{J} , we let $X \doteq q_n(\mathbf{w}) v_n \sim \mathcal{N}(0, q_n^2(\mathbf{w}))$ and $Y \doteq \mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)$. Noticing the fact $t \mapsto \exp(-2t/\mu) t^2$ for $t > 0$ is maximized at $t = \mu$ with maximum value $\exp(-2) \mu^2$, and in view of the estimate in Lemma 7.2, we obtain

$$\begin{aligned} & \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \left(\mathbf{w}^* \bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \|\mathbf{w}\|^2 \right)^2 \right] \\ &\leq 8 \exp(-2) \left(1 - \theta + \frac{\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \theta \right) \mu \end{aligned}$$

$$\begin{aligned}
& + \frac{2\theta}{\mu} \mathbb{E}_{\mathcal{J}} \left[\frac{1}{\sqrt{2\pi}} \frac{\mu \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{\mu^3 \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{3}{4\sqrt{2\pi}} \frac{\|\mathbf{w}_{\mathcal{J}}\|^2 \mu^3}{\|\mathbf{q}_{\mathcal{I}}\|^5} (3\mu^2 + 4\|\mathbf{w}_{\mathcal{J}}\|^2) \right] \\
& + \frac{2\theta}{\mu} \frac{\|\mathbf{w}\|^4}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{1}{\sqrt{2\pi}} \frac{\mu \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{\mu^3 \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{3}{4\sqrt{2\pi}} \frac{q_n^2(\mathbf{w}) \mu^3}{\|\mathbf{q}_{\mathcal{I}}\|^5} (3\mu^2 + 4q_n^2(\mathbf{w})) \right] \\
& \leq \frac{2\theta}{\sqrt{2\pi} q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + \frac{11}{20} \mu \left(2 + \frac{1}{q_n^2(\mathbf{w})} \right) + 2\theta \mu^2 \left(1 + \frac{3}{\sqrt{2\pi} q_n(\mathbf{w})} + \frac{1}{q_n^3(\mathbf{w})} + \frac{3}{\sqrt{2\pi} q_n^5(\mathbf{w})} \right),
\end{aligned}$$

where we have used $\mu < q_n(\mathbf{w}) \leq \|\mathbf{q}_{\mathcal{I}}\|$ and $\|\mathbf{w}_{\mathcal{J}}\| \leq \|\mathbf{q}_{\mathcal{I}}\|$ and $\|\mathbf{w}\| \leq 1$ and $\theta \in (0, 1/2)$ to simplify the intermediate quantities to obtain the last line. Similarly for the second term, we obtain

$$\begin{aligned}
& \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \frac{x_n}{q_n^3(\mathbf{w})} \|\mathbf{w}\|^2 \right] \\
& = \frac{\|\mathbf{w}\|^2 \theta}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \mathbb{E}_v \left[\tanh \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) x_n q_n(\mathbf{w}) \right] \\
& \geq \frac{\|\mathbf{w}\|^2 \theta}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{2q_n^2(\mathbf{w})}{\sqrt{2\pi} \|\mathbf{q}_{\mathcal{I}}\|} - \frac{4\mu^2 q_n^2(\mathbf{w})}{\sqrt{2\pi} \|\mathbf{q}_{\mathcal{I}}\|} - \frac{2q_n^2(\mathbf{w}) \mu^2}{\sqrt{2\pi} \|\mathbf{q}_{\mathcal{I}}\|^3} - \frac{3q_n^2(\mathbf{w}) \mu^4}{2\sqrt{2\pi} \|\mathbf{q}_{\mathcal{I}}\|^5} \right] \\
& \geq \sqrt{\frac{2}{\pi}} \frac{\theta}{q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|} \right] - \frac{4\theta \mu^2}{\sqrt{2\pi}} \left(\frac{1}{q_n^3(\mathbf{w})} + \frac{1}{q_n^5(\mathbf{w})} \right).
\end{aligned}$$

Collecting the above estimates, we obtain

$$\begin{aligned}
& \mathbf{w}^* \nabla_{\mathbf{w}}^2 \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \mathbf{w} \\
& \leq \sqrt{\frac{2}{\pi}} \frac{\theta}{q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} - \frac{\|\mathbf{w}\|^2 (\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2(\mathbf{w}))}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] \\
& \quad + \frac{11}{20} \mu \left(2 + \frac{1}{q_n^2(\mathbf{w})} \right) + 2\theta \mu^2 \left(1 + \frac{3}{\sqrt{2\pi} q_n(\mathbf{w})} + \frac{2}{q_n^3(\mathbf{w})} + \frac{5}{\sqrt{2\pi} q_n^5(\mathbf{w})} \right) \\
& \leq -\sqrt{\frac{2}{\pi}} \theta \mathbb{E} \left[\frac{\|\mathbf{w}_{\mathcal{J}^c}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + \frac{11}{10} \mu + \frac{11}{20} \frac{\mu}{q_n(\mathbf{w})} + 2\theta \mu^2 \left(1 + \frac{6}{q_n^5(\mathbf{w})} \right) \\
& \leq -\sqrt{\frac{2}{\pi}} \theta (1 - \theta) \|\mathbf{w}\|^2 \mathbb{E} \left[\frac{1}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + \frac{11}{10} \mu + \frac{11}{20} \frac{\mu}{q_n^2(\mathbf{w})} + 2\theta \mu^2 \left(1 + \frac{6}{q_n^5(\mathbf{w})} \right), \tag{7.3}
\end{aligned}$$

where to obtain the last line we have invoked the association inequality in Lemma A.3, as both $\|\mathbf{w}_{\mathcal{J}^c}\|^2$ and $1/\|\mathbf{q}_{\mathcal{I}}\|^3$ both coordinatewise nonincreasing w.r.t. the index set. Substituting the upper bound for μ into (7.3) and noting $R_h \leq \|\mathbf{w}\|$ and also noting the fact $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ (implied by the assumption $\|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}$), we obtain the claimed result. \blacksquare

7.1.2 Proof of Proposition 2.6

Proof By similar consideration as proof of the above proposition, the following is justified:

$$\nabla_{\mathbf{w}} \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] = \mathbb{E} [\nabla_{\mathbf{w}} h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})].$$

Now consider

$$\begin{aligned} \mathbf{w}^* \nabla \mathbb{E} [h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] &= \nabla \mathbb{E} [\mathbf{w}^* h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \\ &= \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}}) \right] - \frac{\|\mathbf{w}\|^2}{q_n} \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) x_n \right]. \end{aligned} \quad (7.4)$$

For (7.4), we next provide a lower bound for the first expectation and an upper bound for the second expectation. For the first, we have

$$\begin{aligned} &\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}}) \right] \\ &= \theta \mathbb{E}_{\mathcal{J}} \left[\mathbb{E}_v \left[\tanh \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) (\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}}) \right] \right] + (1 - \theta) \mathbb{E}_{\mathcal{J}} \left[\mathbb{E}_v \left[\tanh \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}}}{\mu} \right) (\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}}) \right] \right] \\ &= \theta \mathbb{E}_{\mathcal{J}} \left[\mathbb{E}_{X,Y} \left[\tanh \left(\frac{X+Y}{\mu} \right) Y \right] \right] + (1 - \theta) \mathbb{E}_{\mathcal{J}} \left[\mathbb{E}_Y \left[\tanh \left(\frac{Y}{\mu} \right) Y \right] \right], \end{aligned}$$

where $X \doteq q_n(\mathbf{w}) v_n \sim \mathcal{N}(0, q_n^2(\mathbf{w}))$ and $Y \doteq \mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)$. Now by Lemma A.3 we obtain

$$\mathbb{E} \left[\tanh \left(\frac{X+Y}{\mu} \right) Y \right] \geq \mathbb{E} \left[\tanh \left(\frac{X+Y}{\mu} \right) \right] \mathbb{E}[Y] = 0,$$

as $\tanh \left(\frac{X+Y}{\mu} \right)$ and X are both coordinatewise nondecreasing function of X and Y . Using the $\tanh(z) \geq (1 - \exp(-2z))/2$ lower bound for $z > 0$ and integral results in Lemma B.1, we obtain

$$\begin{aligned} \mathbb{E} \left[\tanh \left(\frac{Y}{\mu} \right) Y \right] &= 2 \mathbb{E} \left[\tanh \left(\frac{Y}{\mu} \right) Y \mathbb{1}_{Y>0} \right] \\ &\geq \mathbb{E} \left[\left(1 - \exp \left(-\frac{2Y}{\mu} \right) \right) Y \mathbb{1}_{Y>0} \right] \\ &= \frac{2\sigma_Y^2}{\mu} \exp \left(\frac{2\sigma_Y^2}{\mu^2} \right) \Phi^c \left(\frac{2\sigma_Y}{\mu} \right) \\ &\geq \frac{2\sigma_Y^2}{\mu\sqrt{2\pi}} \left(\sqrt{1 + \frac{\sigma_Y^2}{\mu^2}} - \frac{\sigma_Y}{\mu} \right) \\ &\geq \frac{2\sigma_Y^2}{\mu\sqrt{2\pi}} \left(\sqrt{1 + \frac{\|\mathbf{w}\|^2}{\mu^2}} - \frac{\|\mathbf{w}\|}{\mu} \right), \end{aligned}$$

where at the second last inequality we have used Type III lower bound for Gaussian upper tail $\Phi^c(\cdot)$ (Lemma A.5), and at the last we have used the fact that $t \mapsto \sqrt{1+t^2} - t$ is a monotonic decreasing function over $t > 0$ and that $\sigma_Y = \|\mathbf{w}_{\mathcal{J}}\| \leq \|\mathbf{w}\|$. Collecting the above estimates, we have

$$\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}}) \right] \geq (1 - \theta) \mathbb{E}_{\mathcal{J}} \left[\frac{2\|\mathbf{w}_{\mathcal{J}}\|^2}{\mu\sqrt{2\pi}} \left(\sqrt{1 + \frac{\|\mathbf{w}\|_2^2}{\mu^2}} - \frac{\|\mathbf{w}\|}{\mu} \right) \right]$$

$$\begin{aligned}
&\geq (1 - \theta) \mathbb{E}_{\mathcal{J}} \left[\frac{2 \|\mathbf{w}_{\mathcal{J}}\|^2}{\mu \sqrt{2\pi}} \frac{\mu}{10 \|\mathbf{w}\|} \right] \\
&\geq \frac{\theta (1 - \theta) \|\mathbf{w}\|}{5\sqrt{2\pi}}, \tag{7.5}
\end{aligned}$$

where at the second line we have used the assumption that $\|\mathbf{w}\| \geq \frac{\mu}{6\sqrt{2}}$ and also the fact that $\sqrt{1+x^2} \geq x + \frac{1}{10x}$ for $x \geq \frac{1}{6\sqrt{2}}$.

For the second expectation of (7.4), we have

$$\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) x_n \right] \leq \theta \mathbb{E} \left[\left| \tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right| |v_n| \right] \leq \theta \sqrt{\frac{2}{\pi}}, \tag{7.6}$$

as $\tanh(\cdot)$ is bounded by one in magnitude. Plugging the results of (7.5) and (7.6) into (7.4) and noticing that $q_n(\mathbf{w})^2 + \|\mathbf{w}\|^2 = 1$ we obtain

$$\mathbf{w}^* \nabla \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \geq \frac{\theta \|\mathbf{w}\|}{\sqrt{2\pi}} \left[\frac{1 - \theta}{5} - \frac{2 \|\mathbf{w}\|}{\sqrt{1 - \|\mathbf{w}\|^2}} \right] \geq \frac{\theta (1 - \theta) \|\mathbf{w}\|}{10\sqrt{2\pi}},$$

where we have invoked the assumption that $\|\mathbf{w}\| \leq \frac{1}{10\sqrt{5}} (1 - \theta)$ to provide the upper bound $\frac{2\|\mathbf{w}\|}{\sqrt{1-\|\mathbf{w}\|^2}} \leq \frac{1}{10} (1 - \theta)$. We then choose the particular ranges as stated for μ and θ to ensure $r_g < R_g$, completing the proof. \blacksquare

7.1.3 Proof of Proposition 2.7

Proof By consideration similar to proof of Proposition 2.5, we can exchange the hessian and expectation, i.e.,

$$\nabla_{\mathbf{w}}^2 \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] = \mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})].$$

We are interested in the expected Hessian matrix

$$\begin{aligned}
\nabla_{\mathbf{w}}^2 \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] &= \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w} \right) \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w} \right)^* \right] \\
&\quad - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \left(\frac{x_n}{q_n(\mathbf{w})} \mathbf{I} + \frac{x_n}{q_n^3(\mathbf{w})} \mathbf{w} \mathbf{w}^* \right) \right]
\end{aligned}$$

in the region that $0 \leq \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$.

When $\mathbf{w} = \mathbf{0}$, by Lemma B.1, we have

$$\begin{aligned}
&\mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] |_{\mathbf{w}=\mathbf{0}} \\
&= \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{x_n}{\mu} \right) \right) \bar{\mathbf{x}} \bar{\mathbf{x}}^* \right] - \mathbb{E} \left[\tanh \left(\frac{x_n}{\mu} \right) x_n \right] \mathbf{I} \\
&= \frac{\theta(1-\theta)}{\mu} \mathbf{I} + \frac{\theta^2}{\mu} \mathbb{E}_{v_n} \left[1 - \tanh^2 \left(\frac{v_n}{\mu} \right) \right] \mathbf{I} - \frac{\theta}{\mu} \mathbb{E}_{v_n} \left[1 - \tanh^2 \left(\frac{v_n}{\mu} \right) \right] \mathbf{I}
\end{aligned}$$

$$= \frac{\theta(1-\theta)}{\mu} \mathbb{E}_{v_n} \left[\tanh^2 \left(\frac{q_n(\mathbf{w}) v_n}{\mu} \right) \right] \mathbf{I}.$$

Simple calculation based on Lemma B.1 shows

$$\mathbb{E}_{v_n} \left[\tanh^2 \left(\frac{v_n}{\mu} \right) \right] \geq 2 \left(1 - 4 \exp \left(-\frac{2}{\mu^2} \right) \Phi^c \left(\frac{2}{\mu} \right) \right) \geq 2 \left(1 - \frac{2}{\sqrt{2\pi}\mu} \right).$$

Invoking the assumptions $\mu \leq \frac{1}{20\sqrt{n}} \leq 1/20$ and $\theta < 1/2$, we obtain

$$\mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] |_{\mathbf{w}=0} \succeq \frac{\theta(1-\theta)}{\mu} \left(2 - \frac{4}{\sqrt{2\pi}\mu} \right) \mathbf{I} \succeq \frac{\theta}{\mu} \left(1 - \frac{1}{10\sqrt{2\pi}} \right) \mathbf{I}.$$

When $0 < \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$, we aim to derive a semidefinite lower bound for

$$\begin{aligned} & \mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \\ &= \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \bar{\mathbf{x}} \bar{\mathbf{x}}^* \right] - \frac{1}{q_n^2(\mathbf{w})} \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n(\mathbf{w}) x_n \right] \mathbf{I} \\ & \quad - \frac{1}{\mu q_n^2(\mathbf{w})} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n (\mathbf{w} \bar{\mathbf{x}}^* + \bar{\mathbf{x}} \mathbf{w}^*) \right] \\ & \quad + \frac{1}{q_n^4(\mathbf{w})} \left\{ \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) (q_n(\mathbf{w}) x_n)^2 \right] - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n(\mathbf{w}) x_n \right] \right\} \mathbf{w} \mathbf{w}^*. \end{aligned} \tag{7.7}$$

We will first provide bounds for the last two lines and then tackle the first which is slightly more tricky. For the second line, we have

$$\begin{aligned} & \frac{1}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n (\mathbf{w} \bar{\mathbf{x}}^* + \bar{\mathbf{x}} \mathbf{w}^*) \right] \right\| \\ & \leq \frac{2}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n \bar{\mathbf{x}} \right] \mathbf{w}^* \right\| \\ & \leq \frac{2}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n \bar{\mathbf{x}} \right] \right\| \|\mathbf{w}\| \\ & \leq \frac{2}{\mu q_n(\mathbf{w})} \theta^2 \mathbb{E} [|v_n|] \mathbb{E} [\|\bar{\mathbf{v}}\|] \|\mathbf{w}\| \\ & \leq \frac{4\theta^2}{\pi \mu q_n(\mathbf{w})} \sqrt{n} \|\mathbf{w}\| \leq \frac{\theta}{\mu} \frac{4\theta \sqrt{n} \|\mathbf{w}\|}{\pi \sqrt{1 - \|\mathbf{w}\|^2}} \leq \frac{\theta}{\mu} \frac{1}{40\pi}, \end{aligned}$$

where from the third to the fourth line we have used $\left\| 1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right\| \leq 1$, Jensen's inequality for the $\|\cdot\|$ function, and independence of x_n and $\bar{\mathbf{x}}$, and to obtain the last bound we have invoked the $\|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$, $\mu \leq \frac{1}{20\sqrt{n}}$, and $\theta < \frac{1}{2}$ assumptions. For the third line in (7.7), by Lemma A.1 and Lemma B.1,

$$\left| \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) (q_n(\mathbf{w}) x_n)^2 \right] - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n x_n \right] \right|$$

$$\begin{aligned}
&= \left[\frac{\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) (q_n(\mathbf{w}) v_n)^2 \right) \right] \right. \\
&\quad \left. - \theta \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\tanh \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) q_n(\mathbf{w}) v_n \right] \right] \\
&= \frac{\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) \right) ((q_n(\mathbf{w}) v_n)^2 + q_n^2(\mathbf{w})) \right] \\
&\leq \frac{8\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\exp \left(-\frac{2}{\mu} (\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n) \right) ((q_n(\mathbf{w}) v_n)^2 + q_n^2(\mathbf{w})) \mathbb{1}_{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n > 0} \right] \\
&\leq \frac{8\theta}{\sqrt{2\pi}} \mathbb{E}_{\mathcal{J}} \left[\frac{q_n^2(\mathbf{w})}{\sqrt{q_n^2(\mathbf{w}) + \|\mathbf{w}_{\mathcal{J}}\|^2}} \right] \leq \frac{8\theta q_n(\mathbf{w})}{\sqrt{2\pi}}.
\end{aligned}$$

Thus, we have

$$\begin{aligned}
&\frac{1}{q_n^4(\mathbf{w})} \left\{ \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) \right) (q_n x_n)^2 \right] - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^* \mathbf{x}}{\mu} \right) q_n x_n \right] \right\} \mathbf{w} \mathbf{w}^* \\
&\succeq -\frac{8\theta}{q_n^3(\mathbf{w}) \sqrt{2\pi}} \|\mathbf{w}\|^2 \mathbf{I} \succeq -\frac{\theta}{\mu} \left(\frac{64n^{3/2} \mu \|\mathbf{w}\|^2}{q_n^3(\mathbf{w}) \sqrt{2\pi}} \right) \mathbf{I} \succeq -\frac{\theta}{\mu} \frac{1}{4000\sqrt{2\pi}} \mathbf{I},
\end{aligned}$$

where we have again used $\|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$, $\mu \leq \frac{1}{20\sqrt{n}}$, and $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ assumptions to simplify the final bound.

To derive a lower bound for the first line of (7.7), we lower bound the first term and upper bound the second. The latter is easy: using Lemma A.1 and Lemma B.1,

$$\begin{aligned}
&\frac{1}{q_n^2(\mathbf{w})} \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n(\mathbf{w}) x_n \right] \\
&= \frac{\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[1 - \tanh^2 \left[\frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right] \right] \\
&\leq \frac{8\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{\mathbf{v}} \left[\exp \left(-2 \frac{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n}{\mu} \right) \mathbb{1}_{\mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} + q_n(\mathbf{w}) v_n > 0} \right] \\
&\leq \frac{4\theta}{\sqrt{2\pi} q_n(\mathbf{w})} \leq \frac{\theta}{\mu} \frac{8\sqrt{n}\mu}{\sqrt{2\pi}} \leq \frac{\theta}{\mu} \frac{2}{5\sqrt{2\pi}},
\end{aligned}$$

where we have again used assumptions that $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ and $\mu \leq \frac{1}{20\sqrt{n}}$ to simplify the last bound. To lower bound the first term, first note that

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

We set out to lower bound the expectation as

$$\mathbb{E}_{\bar{\mathbf{x}}} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) \right) \bar{\mathbf{x}} \bar{\mathbf{x}}^* \right] \succeq \theta \beta \mathbf{I}$$

for some scalar $\beta > 0$. Suppose \mathbf{w} has $k \in [n-1]$ nonzeros, w.l.o.g., further assume the first k elements of \mathbf{w} are these nonzeros. It is easy to see the expectation above has a block diagonal

structure $\text{diag}(\boldsymbol{\Sigma}; \theta \mathbf{I}_{n-1-k})$. So in order to derive the $\theta\beta\mathbf{I}$ lower bound as desired, it is sufficient to show $\boldsymbol{\Sigma} \succeq \theta\beta\mathbf{I}$ for some $0 < \beta < 1$, i.e., letting $\tilde{\mathbf{w}} \in \mathbb{R}^k$ be the subvector of nonzero elements,

$$\mathbb{E}_{\tilde{\mathbf{x}} \sim i.i.d. \text{BG}(\theta)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}^* \tilde{\mathbf{x}}}{\mu} \right) \right) \tilde{\mathbf{x}} \tilde{\mathbf{x}}^* \right] \succeq \theta\beta\mathbf{I},$$

which is equivalent to that for all $\mathbf{z} \in \mathbb{R}^k$ such that $\|\mathbf{z}\| = 1$,

$$\mathbb{E}_{\tilde{\mathbf{x}} \sim i.i.d. \text{BG}(\theta)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}^* \tilde{\mathbf{x}}}{\mu} \right) \right) (\tilde{\mathbf{x}}^* \mathbf{z})^2 \right] \geq \theta\beta.$$

It is then sufficient to show that for any nontrivial support set $\mathcal{S} \subset [k]$ and any vector $\mathbf{z} \in \mathbb{R}^k$ such that $\text{supp}(\mathbf{z}) = \mathcal{S}$ with $\|\mathbf{z}\| = 1$,

$$\mathbb{E}_{\tilde{\mathbf{v}} \sim i.i.d. \mathcal{N}(0,1)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \mathbf{z})^2 \right] \geq \beta.$$

To see the implication, suppose the latter claimed holds, then for any \mathbf{z} with unit norm,

$$\begin{aligned} & \mathbb{E}_{\tilde{\mathbf{x}} \sim i.i.d. \text{BG}(\theta)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}^* \tilde{\mathbf{x}}}{\mu} \right) \right) (\tilde{\mathbf{x}}^* \mathbf{z})^2 \right] \\ &= \sum_{s=1}^k \theta^s (1-\theta)^{k-s} \sum_{\mathcal{S} \in \binom{[k]}{s}} \mathbb{E}_{\tilde{\mathbf{v}} \sim i.i.d. \mathcal{N}(0,1)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \mathbf{z}_{\mathcal{S}})^2 \right] \\ &\geq \sum_{s=1}^k \theta^s (1-\theta)^{k-s} \sum_{\mathcal{S} \in \binom{[k]}{s}} \beta \|\mathbf{z}_{\mathcal{S}}\|^2 = \beta \mathbb{E}_{\mathcal{S}} [\|\mathbf{z}_{\mathcal{S}}\|^2] = \theta\beta. \end{aligned}$$

Now for any fixed support set $\mathcal{S} \subset [k]$, $\mathbf{z} = \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}} \mathbf{z} + (\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z}$. So we have

$$\begin{aligned} & \mathbb{E}_{\tilde{\mathbf{v}} \sim i.i.d. \mathcal{N}(0,1)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \mathbf{z})^2 \right] \\ &= \mathbb{E}_{\tilde{\mathbf{v}}} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}} \mathbf{z})^2 \right] + \mathbb{E}_{\tilde{\mathbf{v}}} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* (\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z})^2 \right] \\ &= \frac{(\tilde{\mathbf{w}}_{\mathcal{S}}^* \mathbf{z})^2}{\|\mathbf{w}_{\mathcal{S}}\|^4} \mathbb{E}_{\tilde{\mathbf{v}}} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \tilde{\mathbf{w}}_{\mathcal{S}})^2 \right] + \mathbb{E}_{\tilde{\mathbf{v}}} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) \right] \mathbb{E}_{\tilde{\mathbf{v}}} \left[(\tilde{\mathbf{v}}^* (\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z})^2 \right] \\ &\geq 2 \frac{(\tilde{\mathbf{w}}_{\mathcal{S}}^* \mathbf{z})^2}{\|\mathbf{w}_{\mathcal{S}}\|^4} \mathbb{E}_{\tilde{\mathbf{v}}} \left[\exp \left(-\frac{2\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) (\tilde{\mathbf{v}}^* \tilde{\mathbf{w}}_{\mathcal{S}})^2 \mathbb{1}_{\tilde{\mathbf{v}}^* \tilde{\mathbf{w}}_{\mathcal{S}} > 0} \right] + 2 \mathbb{E}_{\tilde{\mathbf{v}}} \left[\exp \left(-\frac{2\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \mathbb{1}_{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}} > 0} \right] \|(\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z}\|^2. \end{aligned}$$

Using expectation result from Lemma B.1, and applying Type III lower bound for Gaussian tails, we obtain

$$\begin{aligned} & \mathbb{E}_{\tilde{\mathbf{v}} \sim i.i.d. \mathcal{N}(0,1)} \left[\left(1 - \tanh^2 \left(\frac{\tilde{\mathbf{w}}_{\mathcal{S}}^* \tilde{\mathbf{v}}}{\mu} \right) \right) (\tilde{\mathbf{v}}^* \mathbf{z})^2 \right] \\ &\geq \frac{1}{\sqrt{2\pi}} \left(\sqrt{4 + \frac{4\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^2}{\mu^2}} - \frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|}{\mu} \right) - \frac{4(\tilde{\mathbf{w}}_{\mathcal{S}}^* \mathbf{z})^2}{\mu\sqrt{2\pi}\|\tilde{\mathbf{w}}_{\mathcal{S}}\|} \end{aligned}$$

$$\geq \frac{1}{\sqrt{2\pi}} \left(2 - \frac{3}{4}\sqrt{2}\right),$$

where we have used Cauchy-Schwarz to obtain $(\tilde{\mathbf{v}}^* \mathbf{z})^2 \leq \|\tilde{\mathbf{v}}^*\|^2$ and invoked the assumption $\|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}$ to simplify the last bound. So $\beta = \frac{1}{\sqrt{2\pi}} \left(2 - \frac{3}{4}\sqrt{2}\right) < 1$.

Putting together the above estimates for the case $\mathbf{w} \neq \mathbf{0}$, we obtain

$$\mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \succeq \frac{\theta}{\mu\sqrt{2\pi}} \left(1 - \frac{3}{8}\sqrt{2} - \frac{\sqrt{2\pi}}{40\pi} - \frac{1}{4000} - \frac{2}{5}\right) \mathbf{I} \succeq \frac{1}{25\sqrt{2\pi}} \frac{\theta}{\mu} \mathbf{I}.$$

Hence for all \mathbf{w} , we can take the $\frac{1}{25\sqrt{2\pi}} \frac{\theta}{\mu}$ as the lower bound, completing the proof. \blacksquare

7.1.4 Proof of Pointwise Concentration Results

To avoid clutter of notations, in this subsection we write \mathbf{X} to mean \mathbf{X}_0 ; similarly \mathbf{x}_k for $(\mathbf{x}_0)_k$, the k -th column of \mathbf{X}_0 . The function $g(\mathbf{w})$ means $g(\mathbf{w}; \mathbf{X}_0)$. We first establish a useful comparison lemma between random i.i.d. Bernoulli random vectors random i.i.d. normal random vectors.

Lemma 7.4 *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 $m \geq 1$.

Now, we are ready to prove Proposition 2.8 to Proposition 2.10 as follows.

Proof [of Proposition 2.8] Let

$$Y_k = \frac{1}{\|\mathbf{w}\|^2} \mathbf{w}^* \nabla^2 h_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) \mathbf{w},$$

then $\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} = \frac{1}{p} \sum_{k=1}^p Y_k$. For each Y_k ($k \in [p]$), from (7.2), we know that

$$Y_k = \frac{1}{\mu} \left(1 - \tanh^2 \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu}\right)\right) \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{x_k(n) \|\mathbf{w}\|}{q_n(\mathbf{w})}\right)^2 - \tanh \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu}\right) \frac{x_k(n)}{q_n^3(\mathbf{w})}.$$

Writing $Y_k = W_k + V_k$, where

$$\begin{aligned} W_k &= \frac{1}{\mu} \left(1 - \tanh^2 \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu}\right)\right) \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{x_k(n) \|\mathbf{w}\|}{q_n(\mathbf{w})}\right)^2, \\ V_k &= - \tanh \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu}\right) \frac{x_k(n)}{q_n^3(\mathbf{w})}. \end{aligned}$$

Then by similar argument as in proof to Proposition 2.9, we have for all integers $m \geq 2$ that

$$\mathbb{E} [|W_k|^m] \leq \frac{1}{\mu^m} \mathbb{E} \left[\left| \frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{x_k(n) \|\mathbf{w}\|}{q_n(\mathbf{w})} \right|^{2m} \right] \leq \frac{1}{\mu^m} \mathbb{E}_{Z \sim \mathcal{N}(0, 1/q_n^2(\mathbf{w}))} [|Z|^{2m}]$$

$$\begin{aligned} &\leq \frac{1}{\mu^m} (2m-1)!! (4n)^m \leq \frac{m!}{2} \left(\frac{4n}{\mu}\right)^m, \\ \mathbb{E}[|V_k|^m] &\leq \frac{1}{q_n^{3m}(\mathbf{w})} \mathbb{E}[|v_k(n)|^m] \leq (2\sqrt{n})^{3m} (m-1)!! \leq \frac{m!}{2} (8n\sqrt{n})^m, \end{aligned}$$

where we have again used the assumption that $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to simplify the result. Taking $\sigma_W^2 = 16n^2/\mu^2 \geq \mathbb{E}[W_k^2]$, $R_W = 4n/\mu$ and $\sigma_V^2 = 64n^3 \geq \mathbb{E}[V_k^2]$, $R_V = 8n\sqrt{n}$, and considering $S_W = \frac{1}{p} \sum_{k=1}^p W_k$ and $S_V = \frac{1}{p} \sum_{k=1}^p V_k$, then by Lemma A.9, we obtain

$$\begin{aligned} \mathbb{P}\left[|S_W - \mathbb{E}[S_W]| \geq \frac{t}{2}\right] &\leq 2 \exp\left(-\frac{p\mu^2 t^2}{128n^2 + 16n\mu t}\right), \\ \mathbb{P}\left[|S_V - \mathbb{E}[S_V]| \geq \frac{t}{2}\right] &\leq 2 \exp\left(-\frac{pt^2}{512n^3 + 32n\sqrt{nt}}\right). \end{aligned}$$

Combining the above results, we obtain

$$\begin{aligned} \mathbb{P}\left[\left|\frac{1}{p} \sum_{k=1}^p X_k - \mathbb{E}[X_k]\right| \geq t\right] &= \mathbb{P}[|S_W - \mathbb{E}[S_W] + S_V - \mathbb{E}[S_V]| \geq t] \\ &\leq \mathbb{P}\left[|S_W - \mathbb{E}[S_W]| \geq \frac{t}{2}\right] + \mathbb{P}\left[|S_V - \mathbb{E}[S_V]| \geq \frac{t}{2}\right] \\ &\leq 2 \exp\left(-\frac{p\mu^2 t^2}{128n^2 + 16n\mu t}\right) + 2 \exp\left(-\frac{pt^2}{512n^3 + 32n\sqrt{nt}}\right) \\ &\leq 4 \exp\left(-\frac{p\mu^2 t^2}{512n^2 + 32n\mu t}\right), \end{aligned}$$

provided that $\mu \leq \frac{1}{\sqrt{n}}$, as desired. ■

Proof [of Proposition 2.9] Let

$$X_k = \frac{\mathbf{w}^*}{\|\mathbf{w}\|_2} \nabla h_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k),$$

then $\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|_2} = \frac{1}{p} \sum_{k=1}^p X_k$. For each $X_k, k \in [p]$, from (7.1), we know that

$$|X_k| = \left| \tanh\left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu}\right) \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{\|\mathbf{w}\|_2 x_k(n)}{q_n(\mathbf{w})}\right) \right| \leq \left| \frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{\|\mathbf{w}\|_2 x_k(n)}{q_n(\mathbf{w})} \right|,$$

as the magnitude of $\tanh(\cdot)$ is bounded by one. Because $\frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{\|\mathbf{w}\|_2 x_k(n)}{q_n(\mathbf{w})} = \left(\frac{\mathbf{w}}{\|\mathbf{w}\|}, -\frac{\|\mathbf{w}\|}{q_n(\mathbf{w})}\right)^* \mathbf{x}_k$ and $\mathbf{x}_k \sim_{i.i.d.} \text{BG}(\theta)$, invoking Lemma 7.4, we obtain for every integer $m \geq 2$ that

$$\mathbb{E}[|X_k|^m] \leq \mathbb{E}_{Z \sim \mathcal{N}(0, 1/q_n^2(\mathbf{w}))}[|Z|^m] \leq \frac{1}{q_n(\mathbf{w})^m} (m-1)!! \leq \frac{m!}{2} (4n) (2\sqrt{n})^{m-2},$$

where the Gaussian moment can be looked up in Lemma A.6 and we used the fact that $(m-1)!! \leq m!/2$ and the assumption that $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to get the result. Thus, by taking $\sigma^2 = 4n \geq \mathbb{E}[X_k^2]$ and $R = 2\sqrt{n}$, and we obtain the claimed result by invoking Lemma A.9. ■

Proof [of Proposition 2.10] Let $\mathbf{Z}_k = \nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k)$, then $\nabla_{\mathbf{w}}^2 g(\mathbf{w}) = \frac{1}{p} \sum_{k=1}^p \mathbf{Z}_k$. From (7.2), we know that

$$\mathbf{Z}_k = \mathbf{W}_k + \mathbf{V}_k$$

where

$$\begin{aligned} \mathbf{W}_k &= \frac{1}{\mu} \left(1 - \tanh^2 \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu} \right) \right) \left(\bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right) \left(\bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right)^* \\ \mathbf{V}_k &= -\tanh \left(\frac{\mathbf{q}(\mathbf{w})^* \mathbf{x}_k}{\mu} \right) \left(\frac{x_k(n)}{q_n(\mathbf{w})} \mathbf{I} + \frac{x_k(n) \mathbf{w} \mathbf{w}^*}{q_n^3(\mathbf{w})} \right). \end{aligned}$$

For \mathbf{W}_k , we have

$$\begin{aligned} \mathbf{0} \preceq \mathbb{E}[\mathbf{W}_k^m] &\preceq \frac{1}{\mu^m} \mathbb{E} \left[\left\| \bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right\|^{2m-2} \left(\bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right) \left(\bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right)^* \right] \\ &\preceq \frac{1}{\mu^m} \mathbb{E} \left[\left\| \bar{\mathbf{x}}_k - \frac{x_k(n) \mathbf{w}}{q_n(\mathbf{w})} \right\|^{2m} \right] \mathbf{I} \\ &\preceq \frac{2^m}{\mu^m} \mathbb{E} \left[\left(\|\bar{\mathbf{x}}_k\|^2 + \frac{x_k^2(n) \|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \right)^m \right] \mathbf{I} \\ &\preceq \frac{2^m}{\mu^m} \mathbb{E} \left[\|\mathbf{x}_k\|^{2m} \right] \mathbf{I} \preceq \frac{2^m}{\mu^m} \mathbb{E}_{Z \sim \chi^2(n)} [Z^m] \mathbf{I}, \end{aligned}$$

where we have used the fact that $\|\mathbf{w}\|^2 / q_n^2(\mathbf{w}) = \|\mathbf{w}\|^2 / (1 - \|\mathbf{w}\|^2) \leq 1$ for $\|\mathbf{w}\|_2 \leq \frac{1}{4}$ and Lemma 7.4 to obtain the last line. By Lemma A.7, we obtain

$$\mathbf{0} \preceq \mathbb{E}[\mathbf{W}_k^m] \preceq \left(\frac{2}{\mu} \right)^m \frac{m!}{2} (2n)^m \mathbf{I} = \frac{m!}{2} \left(\frac{4n}{\mu} \right)^m \mathbf{I}.$$

Taking $R_W = \frac{4n}{\mu}$ and $\sigma_W^2 = \frac{16n^2}{\mu^2} \geq \mathbb{E}[\mathbf{W}_k^2]$, and letting $\mathbf{S}_W \doteq \frac{1}{p} \sum_{k=1}^p \mathbf{W}_k$, by Lemma A.10, we obtain

$$\mathbb{P} \left[\|\mathbf{S}_W - \mathbb{E}[\mathbf{S}_W]\| \geq \frac{t}{2} \right] \leq 2n \exp \left(-\frac{p\mu^2 t^2}{128n^2 + 16\mu n t} \right).$$

Similarly, for \mathbf{V}_k , we have

$$\begin{aligned} \mathbb{E}[\mathbf{V}_k^m] &\preceq \left(\frac{1}{q_n(\mathbf{w})} + \frac{\|\mathbf{w}\|^2}{q_n^3(\mathbf{w})} \right)^m \mathbb{E}[|x_k(n)|^m] \mathbf{I} \\ &\preceq (8n\sqrt{n})^m (m-1)!! \mathbf{I} \\ &\preceq \frac{m!}{2} (8n\sqrt{n})^m \mathbf{I}, \end{aligned}$$

where we have used the fact $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to simplify the result. Similar argument also shows $-\mathbb{E}[\mathbf{V}_k^m] \preceq m! (8n\sqrt{n})^m \mathbf{I} / 2$. Taking $R_V = 8n\sqrt{n}$ and $\sigma_V^2 = 64n^3$, and letting $\mathbf{S}_V \doteq \frac{1}{p} \sum_{k=1}^p \mathbf{V}_k$, again by Lemma A.10, we obtain

$$\mathbb{P} \left[\|\mathbf{S}_V - \mathbb{E}[\mathbf{S}_V]\| \geq \frac{t}{2} \right] \leq 2n \exp \left(-\frac{pt^2}{512n^3 + 32n\sqrt{nt}} \right).$$

Combining the above results, we obtain

$$\begin{aligned}
\mathbb{P} \left[\left\| \frac{1}{p} \sum_{k=1}^p \mathbf{Z}_k - \mathbb{E}[\mathbf{Z}_k] \right\| \geq t \right] &= \mathbb{P} [\|\mathbf{S}_W - \mathbb{E}[\mathbf{S}_W] + \mathbf{S}_V - \mathbb{E}[\mathbf{S}_V]\| \geq t] \\
&\leq \mathbb{P} \left[\|\mathbf{S}_W - \mathbb{E}[\mathbf{S}_W]\| \geq \frac{t}{2} \right] + \mathbb{P} \left[\|\mathbf{S}_V - \mathbb{E}[\mathbf{S}_V]\| \geq \frac{t}{2} \right] \\
&\leq 2n \exp \left(-\frac{p\mu^2 t^2}{128n^2 + 16\mu nt} \right) + 2n \exp \left(-\frac{pt^2}{512n^3 + 32n\sqrt{nt}} \right) \\
&\leq 4n \exp \left(-\frac{p\mu^2 t^2}{512n^2 + 32\mu nt} \right),
\end{aligned}$$

where we have simplified the final result based on the fact that $\mu \leq \frac{1}{\sqrt{n}}$. ■

7.1.5 Proof of Lipschitz Results

To avoid clutter of notations, in this subsection we write \mathbf{X} to mean \mathbf{X}_0 ; similarly x_k for $(x_0)_k$, the k -th column of \mathbf{X}_0 . The function $g(\mathbf{w})$ means $g(\mathbf{w}; \mathbf{X}_0)$. We need the following lemmas to prove the Lipschitz results.

Lemma 7.5 *Suppose that $\varphi_1 : U \rightarrow V$ is an L -Lipschitz map from a normed space U to a normed space V , and that $\varphi_2 : V \rightarrow W$ is an L' -Lipschitz map from V to a normed space W . Then the composition $\varphi_2 \circ \varphi_1 : U \rightarrow W$ is LL' -Lipschitz.*

Lemma 7.6 *Fix any $\mathcal{D} \subseteq \mathbb{R}^{n-1}$. Let $g_1, g_2 : \mathcal{D} \rightarrow \mathbb{R}$, and assume that g_1 is L_1 -Lipschitz, and g_2 is L_2 -Lipschitz, and that g_1 and g_2 are bounded over \mathcal{D} , i.e., $|g_1(\mathbf{x})| \leq M_1$ and $|g_2(\mathbf{x})| \leq M_2$ for all $x \in \mathcal{D}$ with some constants $M_1 > 0$ and $M_2 > 0$. Then the function $h(\mathbf{x}) = g_1(\mathbf{x})g_2(\mathbf{x})$ is L -Lipschitz, with*

$$L = M_1 L_2 + M_2 L_1.$$

Lemma 7.7 *For every $\mathbf{w}, \mathbf{w}' \in \Gamma$, and every fixed \mathbf{x} , we have*

$$\begin{aligned}
\left| \dot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}) - \dot{h}_\mu(\mathbf{q}(\mathbf{w}')^* \mathbf{x}) \right| &\leq \frac{2\sqrt{n}}{\mu} \|\mathbf{x}\| \|\mathbf{w} - \mathbf{w}'\|, \\
\left| \ddot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}) - \ddot{h}_\mu(\mathbf{q}(\mathbf{w}')^* \mathbf{x}) \right| &\leq \frac{4\sqrt{n}}{\mu^2} \|\mathbf{x}\| \|\mathbf{w} - \mathbf{w}'\|.
\end{aligned}$$

Proof We have

$$\begin{aligned}
|q_n(\mathbf{w}) - q_n(\mathbf{w}')| &= \left| \sqrt{1 - \|\mathbf{w}\|^2} - \sqrt{1 - \|\mathbf{w}'\|^2} \right| = \frac{\|\mathbf{w} + \mathbf{w}'\| \|\mathbf{w} - \mathbf{w}'\|}{\sqrt{1 - \|\mathbf{w}\|^2} + \sqrt{1 - \|\mathbf{w}'\|^2}} \\
&\leq \frac{\max(\|\mathbf{w}\|, \|\mathbf{w}'\|)}{\min(q_n(\mathbf{w}), q_n(\mathbf{w}'))} \|\mathbf{w} - \mathbf{w}'\|.
\end{aligned}$$

Hence it holds that

$$\|\mathbf{q}(\mathbf{w}) - \mathbf{q}(\mathbf{w}')\|^2 = \|\mathbf{w} - \mathbf{w}'\|^2 + |q_n(\mathbf{w}) - q_n(\mathbf{w}')|^2 \leq \left(1 + \frac{\max(\|\mathbf{w}\|^2, \|\mathbf{w}'\|^2)}{\min(q_n^2(\mathbf{w}), q_n^2(\mathbf{w}'))} \right) \|\mathbf{w} - \mathbf{w}'\|^2$$

$$= \frac{1}{\min(q_n^2(\mathbf{w}), q_n^2(\mathbf{w}'))} \|\mathbf{w} - \mathbf{w}'\|^2 \leq 4n \|\mathbf{w} - \mathbf{w}'\|^2,$$

where we have used the fact $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to get the final result. Hence the mapping $\mathbf{w} \mapsto \mathbf{q}(\mathbf{w})$ is $2\sqrt{n}$ -Lipschitz over Γ . Moreover it is easy to see $\mathbf{q} \mapsto \mathbf{q}^* \mathbf{x}$ is $\|\mathbf{x}\|_2$ -Lipschitz. By Lemma A.1 and the composition rule in Lemma 7.5, we obtain the desired claims. ■

Lemma 7.8 *For any fixed \mathbf{x} , consider the function*

$$t_{\mathbf{x}}(\mathbf{w}) \doteq \frac{\mathbf{w}^* \bar{\mathbf{x}}}{\|\mathbf{w}\|} - \frac{x_n}{q_n(\mathbf{w})} \|\mathbf{w}\|$$

defined over $\mathbf{w} \in \Gamma$. Then, for all \mathbf{w}, \mathbf{w}' in Γ such that $\|\mathbf{w}\| \geq r$ and $\|\mathbf{w}'\| \geq r$ for some constant $r \in (0, 1)$, it holds that

$$\begin{aligned} |t_{\mathbf{x}}(\mathbf{w}) - t_{\mathbf{x}}(\mathbf{w}')| &\leq 2 \left(\frac{\|\mathbf{x}\|}{r} + 4n^{3/2} \|\mathbf{x}\|_{\infty} \right) \|\mathbf{w} - \mathbf{w}'\|, \\ |t_{\mathbf{x}}(\mathbf{w})| &\leq 2\sqrt{n} \|\mathbf{x}\|, \\ |t_{\mathbf{x}}^2(\mathbf{w}) - t_{\mathbf{x}}^2(\mathbf{w}')| &\leq 8\sqrt{n} \|\mathbf{x}\| \left(\frac{\|\mathbf{x}\|}{r} + 4n^{3/2} \|\mathbf{x}\|_{\infty} \right) \|\mathbf{w} - \mathbf{w}'\|, \\ |t_{\mathbf{x}}^2(\mathbf{w})| &\leq 4n \|\mathbf{x}\|^2. \end{aligned}$$

Proof First of all, we have

$$|t_{\mathbf{x}}(\mathbf{w})| = \left[\frac{\mathbf{w}^*}{\|\mathbf{w}\|}, -\frac{\|\mathbf{w}\|}{q_n(\mathbf{w})} \right] \mathbf{x} \leq \|\mathbf{x}\| \left(1 + \frac{\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \right)^{1/2} = \frac{\|\mathbf{x}\|}{|q_n(\mathbf{w})|} \leq 2\sqrt{n} \|\mathbf{x}\|,$$

where we have used the assumption that $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to simplify the final result. The claim about $|t_{\mathbf{x}}^2(\mathbf{w})|$ follows immediately. Now

$$|t_{\mathbf{x}}(\mathbf{w}) - t_{\mathbf{x}}(\mathbf{w}')| \leq \left| \left(\frac{\mathbf{w}}{\|\mathbf{w}\|} - \frac{\mathbf{w}'}{\|\mathbf{w}'\|} \right)^* \bar{\mathbf{x}} \right| + |x_n| \left| \frac{\|\mathbf{w}\|}{q_n(\mathbf{w})} - \frac{\|\mathbf{w}'\|}{q_n(\mathbf{w}')} \right|.$$

Moreover we have

$$\begin{aligned} \left| \left(\frac{\mathbf{w}}{\|\mathbf{w}\|} - \frac{\mathbf{w}'}{\|\mathbf{w}'\|} \right)^* \bar{\mathbf{x}} \right| &\leq \|\bar{\mathbf{x}}\| \left\| \frac{\mathbf{w}}{\|\mathbf{w}\|} - \frac{\mathbf{w}'}{\|\mathbf{w}'\|} \right\| \leq \|\mathbf{x}\| \frac{\|\mathbf{w} - \mathbf{w}'\| \|\mathbf{w}'\| + \|\mathbf{w}'\| \|\|\mathbf{w}\| - \|\mathbf{w}'\|\|}{\|\mathbf{w}\| \|\mathbf{w}'\|} \\ &\leq \frac{2\|\mathbf{x}\|}{r} \|\mathbf{w} - \mathbf{w}'\|, \end{aligned}$$

where we have used the assumption that $\|\mathbf{w}\| \geq r$ to simplify the result. Noticing that $t \mapsto t/\sqrt{1-t^2}$ is continuous over $[a, b]$ and differentiable over (a, b) for any $0 < a < b < 1$, by mean value theorem,

$$\left| \frac{\|\mathbf{w}\|}{q_n(\mathbf{w})} - \frac{\|\mathbf{w}'\|}{q_n(\mathbf{w}')} \right| \leq \sup_{\mathbf{w} \in \Gamma} \frac{1}{(1 - \|\mathbf{w}\|^2)^{3/2}} \|\mathbf{w} - \mathbf{w}'\| \leq 8n^{3/2} \|\mathbf{w} - \mathbf{w}'\|,$$

where we have again used the assumption that $q_n(\mathbf{w}) \geq \frac{1}{2\sqrt{n}}$ to simplify the last result. Collecting the above estimates, we obtain

$$|t_x(\mathbf{w}) - t_x(\mathbf{w}')| \leq \left(2\frac{\|\mathbf{x}\|}{r} + 8n^{3/2}\|\mathbf{x}\|_\infty\right)\|\mathbf{w} - \mathbf{w}'\|,$$

as desired. For the last one, we have

$$\begin{aligned} |t_x^2(\mathbf{w}) - t_x^2(\mathbf{w}')| &= |t_x(\mathbf{w}) - t_x(\mathbf{w}')||t_x(\mathbf{w}) + t_x(\mathbf{w}')| \\ &\leq 2 \sup_{s \in \Gamma} |t_x(s)| |t_x(\mathbf{w}) - t_x(\mathbf{w}')|, \end{aligned}$$

leading to the claimed result once we substitute estimates of the involved quantities. \blacksquare

Lemma 7.9 *For any fixed \mathbf{x} , consider the function*

$$\Phi_x(\mathbf{w}) = \frac{x_n}{q_n(\mathbf{w})}\mathbf{I} + \frac{x_n}{q_n^3(\mathbf{w})}\mathbf{w}\mathbf{w}^*$$

defined over $\mathbf{w} \in \Gamma$. Then, for all $\mathbf{w}, \mathbf{w}' \in \Gamma$ such that $\|\mathbf{w}\| < r$ and $\|\mathbf{w}'\| < r$ with some constant $r \in (0, \frac{1}{2})$, it holds that

$$\begin{aligned} \|\Phi_x(\mathbf{w})\| &\leq 2\|\mathbf{x}\|_\infty, \\ \|\Phi_x(\mathbf{w}) - \Phi_x(\mathbf{w}')\| &\leq 4\|\mathbf{x}\|_\infty\|\mathbf{w} - \mathbf{w}'\|. \end{aligned}$$

Proof Simple calculation shows

$$\|\Phi_x(\mathbf{w})\| \leq \|\mathbf{x}\|_\infty \left(\frac{1}{q_n(\mathbf{w})} + \frac{\|\mathbf{w}\|^2}{q_n^3(\mathbf{w})} \right) = \frac{\|\mathbf{x}\|_\infty}{q_n^3(\mathbf{w})} \leq \frac{\|\mathbf{x}\|_\infty}{(1-r^2)^{3/2}} \leq 2\|\mathbf{x}\|_\infty.$$

For the second one, we have

$$\begin{aligned} \|\Phi_x(\mathbf{w}) - \Phi_x(\mathbf{w}')\| &\leq \|\mathbf{x}\|_\infty \left\| \frac{1}{q_n(\mathbf{w})}\mathbf{I} + \frac{1}{q_n^3(\mathbf{w})}\mathbf{w}\mathbf{w}^* - \frac{1}{q_n(\mathbf{w}')}\mathbf{I} - \frac{1}{q_n^3(\mathbf{w}')}\mathbf{w}'(\mathbf{w}')^* \right\| \\ &\leq \|\mathbf{x}\|_\infty \left(\left| \frac{1}{q_n(\mathbf{w})} - \frac{1}{q_n(\mathbf{w}')} \right| + \left| \frac{\|\mathbf{w}\|^2}{q_n^3(\mathbf{w})} - \frac{\|\mathbf{w}'\|^2}{q_n^3(\mathbf{w}')} \right| \right). \end{aligned}$$

Now

$$\left| \frac{1}{q_n(\mathbf{w})} - \frac{1}{q_n(\mathbf{w}')} \right| = \frac{|q_n(\mathbf{w}) - q_n(\mathbf{w}')|}{q_n(\mathbf{w})q_n(\mathbf{w}')} \leq \frac{\max(\|\mathbf{w}\|, \|\mathbf{w}'\|)}{\min(q_n^3(\mathbf{w}), q_n^3(\mathbf{w}'))} \|\mathbf{w} - \mathbf{w}'\| \leq \frac{4}{3\sqrt{3}} \|\mathbf{w} - \mathbf{w}'\|,$$

where we have applied the estimate for $|q_n(\mathbf{w}) - q_n(\mathbf{w}')|$ as established in Lemma 7.7 and also used $\|\mathbf{w}\| \leq 1/2$ and $\|\mathbf{w}'\| \leq 1/2$ to simplify the above result. Further noticing $t \mapsto t^2/(1-t^2)^{3/2}$ is differentiable over $t \in (0, 1)$, we apply the mean value theorem and obtain

$$\left| \frac{\|\mathbf{w}\|^2}{q_n^3(\mathbf{w})} - \frac{\|\mathbf{w}'\|^2}{q_n^3(\mathbf{w}')} \right| \leq \sup_{s \in \Gamma, \|s\| \leq r < \frac{1}{2}} \frac{\|s\|^3 + 2\|s\|}{(1-\|s\|^2)^{5/2}} \|\mathbf{w} - \mathbf{w}'\| \leq \frac{4}{\sqrt{3}} \|\mathbf{w} - \mathbf{w}'\|.$$

Combining the above estimates gives the claimed result. \blacksquare

Lemma 7.10 For any fixed \mathbf{x} , consider the function

$$\zeta_{\mathbf{x}}(\mathbf{w}) = \bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})} \mathbf{w}$$

defined over $\mathbf{w} \in \Gamma$. Then, for all $\mathbf{w}, \mathbf{w}' \in \Gamma$ such that $\|\mathbf{w}\| \leq r$ and $\|\mathbf{w}'\| \leq r$ for some constant $r \in (0, \frac{1}{2})$, it holds that

$$\begin{aligned} \|\zeta_{\mathbf{x}}(\mathbf{w})\zeta_{\mathbf{x}}(\mathbf{w})^*\| &\leq 2n \|\mathbf{x}\|_{\infty}^2, \\ \|\zeta_{\mathbf{x}}(\mathbf{w})\zeta_{\mathbf{x}}(\mathbf{w})^* - \zeta_{\mathbf{x}}(\mathbf{w}')\zeta_{\mathbf{x}}(\mathbf{w}')^*\| &\leq 8\sqrt{2}\sqrt{n} \|\mathbf{x}\|_{\infty}^2 \|\mathbf{w} - \mathbf{w}'\|. \end{aligned}$$

Proof We have $\|\mathbf{w}\|^2 / q_n^2(\mathbf{w}) \leq 1/3$ when $\|\mathbf{w}\| \leq r < 1/2$, hence it holds that

$$\|\zeta_{\mathbf{x}}(\mathbf{w})\zeta_{\mathbf{x}}(\mathbf{w})^*\| \leq \|\zeta_{\mathbf{x}}(\mathbf{w})\|^2 \leq 2\|\bar{\mathbf{x}}\|^2 + 2x_n^2 \frac{\|\mathbf{w}\|}{q_n^2(\mathbf{w})} \leq 2n \|\mathbf{x}\|_{\infty}^2.$$

For the second, we first estimate

$$\begin{aligned} \|\zeta(\mathbf{w}) - \zeta(\mathbf{w}')\| &= \left\| x_n \left(\frac{\mathbf{w}}{q_n(\mathbf{w})} - \frac{\mathbf{w}'}{q_n(\mathbf{w}')} \right) \right\| \leq \|\mathbf{x}\|_{\infty} \left\| \frac{\mathbf{w}}{q_n(\mathbf{w})} - \frac{\mathbf{w}'}{q_n(\mathbf{w}')} \right\| \\ &\leq \|\mathbf{x}\|_{\infty} \left(\frac{1}{q_n(\mathbf{w})} \|\mathbf{w} - \mathbf{w}'\| + \|\mathbf{w}'\| \left| \frac{1}{q_n(\mathbf{w})} - \frac{1}{q_n(\mathbf{w}')} \right| \right) \\ &\leq \|\mathbf{x}\|_{\infty} \left(\frac{1}{q_n(\mathbf{w})} + \frac{\|\mathbf{w}'\|}{\min\{q_n^3(\mathbf{w}), q_n^3(\mathbf{w}')\}} \right) \|\mathbf{w} - \mathbf{w}'\| \\ &\leq \|\mathbf{x}\|_{\infty} \left(\frac{2}{\sqrt{3}} + \frac{4}{3\sqrt{3}} \right) \|\mathbf{w} - \mathbf{w}'\| \leq 4\|\mathbf{x}\|_{\infty} \|\mathbf{w} - \mathbf{w}'\|. \end{aligned}$$

Thus, we have

$$\begin{aligned} \|\zeta_{\mathbf{x}}(\mathbf{w})\zeta_{\mathbf{x}}(\mathbf{w})^* - \zeta_{\mathbf{x}}(\mathbf{w}')\zeta_{\mathbf{x}}(\mathbf{w}')^*\| &\leq \|\zeta(\mathbf{w})\| \|\zeta(\mathbf{w}) - \zeta(\mathbf{w}')\| + \|\zeta(\mathbf{w}) - \zeta(\mathbf{w}')\| \|\zeta(\mathbf{w}')\| \\ &\leq 8\sqrt{2}\sqrt{n} \|\mathbf{x}\|_{\infty}^2 \|\mathbf{w} - \mathbf{w}'\|, \end{aligned}$$

as desired. ■

Now, we are ready to prove all the Lipschitz propositions.

Proof [of Proposition 2.11] Let

$$F_k(\mathbf{w}) = \ddot{h}_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) t_{\mathbf{x}_k}^2(\mathbf{w}) + \dot{h}_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) \frac{x_k(n)}{q_n^3(\mathbf{w})}.$$

Then, $\frac{1}{\|\mathbf{w}\|^2} \mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w} = \frac{1}{p} \sum_{k=1}^p F_k(\mathbf{w})$. Noticing that $\ddot{h}_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k)$ is bounded by $1/\mu$ and $\dot{h}_{\mu}(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k)$ is bounded by 1, both in magnitude. Applying Lemma 7.6, Lemma 7.7 and Lemma 7.8, we can see $F_k(\mathbf{w})$ is $L_{\hat{\cdot}}^k$ -Lipschitz with

$$\begin{aligned} L_{\hat{\cdot}}^k &= 4n \|\mathbf{x}_k\|^2 \frac{4\sqrt{n}}{\mu^2} \|\mathbf{x}_k\| + \frac{1}{\mu} 8\sqrt{n} \|\mathbf{x}_k\| \left(\frac{\|\mathbf{x}_k\|}{r_{\hat{\cdot}}} + 4n^{3/2} \|\mathbf{x}_k\|_{\infty} \right) \\ &\quad + (2\sqrt{n})^3 \|\mathbf{x}_k\|_{\infty} \frac{2\sqrt{n}}{\mu} \|\mathbf{x}_k\| + \sup_{r_{\hat{\cdot}} < a < \sqrt{\frac{2n-1}{2n}}} \frac{3}{(1-a^2)^{5/2}} \|\mathbf{x}_k\|_{\infty} \end{aligned}$$

$$= \frac{16n^{3/2}}{\mu^2} \|\mathbf{x}_k\|^3 + \frac{8\sqrt{n}}{\mu r_\wedge} \|\mathbf{x}_k\|^2 + \frac{48n^2}{\mu} \|\mathbf{x}_k\| \|\mathbf{x}_k\|_\infty + 96n^{5/2} \|\mathbf{x}_k\|_\infty.$$

Thus, $\frac{1}{\|\mathbf{w}\|_2} \mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}$ is L_\wedge -Lipschitz with

$$L_\wedge \leq \frac{1}{p} \sum_{k=1}^p L_\wedge^k \leq \frac{16n^3}{\mu^2} \|\mathbf{X}\|_\infty^3 + \frac{8n^{3/2}}{\mu r_\wedge} \|\mathbf{X}\|_\infty^2 + \frac{48n^{5/2}}{\mu} \|\mathbf{X}\|_\infty^2 + 96n^{5/2} \|\mathbf{X}\|_\infty,$$

as desired. ■

Proof [of Proposition 2.12] We have

$$\left\| \frac{\mathbf{w}^*}{\|\mathbf{w}\|} \nabla g(\mathbf{w}) - \frac{\mathbf{w}'^*}{\|\mathbf{w}'\|} \nabla g(\mathbf{w}') \right\| \leq \frac{1}{p} \sum_{k=1}^p \left\| \dot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) t_{\mathbf{x}_k}(\mathbf{w}) - \dot{h}_\mu(\mathbf{q}(\mathbf{w}')^* \mathbf{x}_k) t_{\mathbf{x}_k}(\mathbf{w}') \right\|$$

where $\dot{h}_\mu(t) = \tanh(t/\mu)$ is bounded by one in magnitude, and $t_{\mathbf{x}_k}(\mathbf{w})$ and $t_{\mathbf{x}_k'}(\mathbf{w})$ is defined as in Lemma 7.8. By Lemma 7.6, Lemma 7.7 and Lemma 7.8, we know that $\dot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) t_{\mathbf{x}_k}(\mathbf{w})$ is L_k -Lipschitz with constant

$$L_k = \frac{2\|\mathbf{x}_k\|}{r_g} + 8n^{3/2} \|\mathbf{x}_k\|_\infty + \frac{4n}{\mu} \|\mathbf{x}_k\|^2.$$

Therefore, we have

$$\begin{aligned} \left\| \frac{\mathbf{w}^*}{\|\mathbf{w}\|} \nabla g(\mathbf{w}) - \frac{\mathbf{w}'^*}{\|\mathbf{w}'\|} \nabla g(\mathbf{w}') \right\| &\leq \frac{1}{p} \sum_{k=1}^p \left(\frac{2\|\mathbf{x}_k\|}{r_g} + 8n^{3/2} \|\mathbf{x}_k\|_\infty + \frac{4n}{\mu} \|\mathbf{x}_k\|^2 \right) \|\mathbf{w} - \mathbf{w}'\| \\ &\leq \left(\frac{2\sqrt{n}}{r_g} \|\mathbf{X}\|_\infty + 8n^{3/2} \|\mathbf{X}\|_\infty + \frac{4n^2}{\mu} \|\mathbf{X}\|_\infty^2 \right) \|\mathbf{w} - \mathbf{w}'\|, \end{aligned}$$

as desired. ■

Proof [of Proposition 2.13] Let

$$\mathbf{F}_k(\mathbf{w}) = \ddot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) \zeta_k(\mathbf{w}) \zeta_k(\mathbf{w})^* - \dot{h}_\mu(\mathbf{q}(\mathbf{w})^* \mathbf{x}_k) \Phi_k(\mathbf{w})$$

with $\zeta_k(\mathbf{w}) = \bar{\mathbf{x}}_k - \frac{\mathbf{x}_k(n)}{q_n(\mathbf{w})} \mathbf{w}$ and $\Phi_k(\mathbf{w}) = \frac{\mathbf{x}_k(n)}{q_n(\mathbf{w})} \mathbf{I} + \frac{\mathbf{x}_{n,k}}{q_n(\mathbf{w})} \mathbf{w} \mathbf{w}^*$. Then, $\nabla^2 g(\mathbf{w}) = \frac{1}{p} \sum_{k=1}^p \mathbf{F}_k(\mathbf{w})$.

Using Lemma 7.6, Lemma 7.7, Lemma 7.9 and Lemma 7.10, and the facts that $\dot{h}_\mu(t)$ is bounded by $1/\mu$ and that $\ddot{h}_\mu(t)$ is bounded by 1 in magnitude, we can see $\mathbf{F}_k(\mathbf{w})$ is L_\cup^k -Lipschitz continuous with

$$\begin{aligned} L_\cup^k &= \frac{1}{\mu} \times 8\sqrt{2}\sqrt{n} \|\mathbf{x}_k\|_\infty^2 + \frac{2\sqrt{n}}{\mu^2} \|\mathbf{x}_k\| \times 2n \|\mathbf{x}_k\|_\infty^2 + 4 \|\mathbf{x}_k\|_\infty + \frac{2\sqrt{n}}{\mu} \|\mathbf{x}_k\| \times 2 \|\mathbf{x}_k\|_\infty \\ &\leq \frac{4n^{3/2}}{\mu^2} \|\mathbf{x}_k\| \|\mathbf{x}_k\|_\infty^2 + \frac{4\sqrt{n}}{\mu} \|\mathbf{x}_k\| \|\mathbf{x}_k\|_\infty + \frac{8\sqrt{2}\sqrt{n}}{\mu} \|\mathbf{x}_k\|_\infty^2 + 4 \|\mathbf{x}_k\|_\infty. \end{aligned}$$

Thus, we have

$$L_\cup \leq \frac{1}{p} \sum_{k=1}^p L_\cup^k \leq \frac{4n^2}{\mu^2} \|\mathbf{X}\|_\infty^3 + \frac{4n}{\mu} \|\mathbf{X}\|_\infty^2 + \frac{8\sqrt{2}\sqrt{n}}{\mu} \|\mathbf{X}\|_\infty^2 + 8 \|\mathbf{X}\|_\infty,$$

as desired. ■

7.2 Proofs of Theorem 2.1

To avoid clutter of notations, in this subsection we write \mathbf{X} to mean \mathbf{X}_0 ; similarly \mathbf{x}_k for $(\mathbf{x}_0)_k$, the k -th column of \mathbf{X}_0 . The function $g(\mathbf{w})$ means $g(\mathbf{w}; \mathbf{X}_0)$. Before proving Theorem 2.1, we record one useful lemma.

Lemma 7.11 *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*

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

For convenience, we define three regions for the range of \mathbf{w} :

$$\begin{aligned} R_1 &\doteq \left\{ \mathbf{w} \mid \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}} \right\}, & R_2 &\doteq \left\{ \mathbf{w} \mid \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \right\}, \\ R_3 &\doteq \left\{ \mathbf{w} \mid \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}} \right\}. \end{aligned}$$

Proof [of Theorem 2.1] We will focus on deriving the qualitative result and hence be sloppy about constants. All indexed capital C or small c are numerical constants.

Strong convexity in region R_1 . Proposition 2.7 shows that for any $\mathbf{w} \in R_1$, $\mathbb{E}[\nabla^2 g(\mathbf{w})] \succeq \frac{c_1\theta}{\mu} \mathbf{I}$. For any $\varepsilon \in (0, \mu/(4\sqrt{2}))$, R_1 has an ε -net N_1 of size at most $(3\mu/(4\sqrt{2}\varepsilon))^n$. On \mathcal{E}_∞ , $\nabla^2 g$ is

$$L_1 \doteq \frac{C_2 n^2}{\mu^2} \log^{3/2}(np)$$

Lipschitz by Proposition 2.13. Set $\varepsilon = \frac{c_1\theta}{3\mu L_1}$, so

$$\#N_1 \leq \exp\left(2n \log\left(\frac{C_3 n \log(np)}{\theta}\right)\right).$$

Let \mathcal{E}_1 denote the event

$$\mathcal{E}_1 = \left\{ \max_{\mathbf{w} \in N_1} \|\nabla^2 g(\mathbf{w}) - \mathbb{E}[\nabla^2 g(\mathbf{w})]\| \leq \frac{c_1\theta}{3\mu} \right\}.$$

On $\mathcal{E}_1 \cap \mathcal{E}_\infty$,

$$\sup_{\|\mathbf{w}\| \leq \mu/(4\sqrt{2})} \|\nabla^2 g(\mathbf{w}) - \mathbb{E}[\nabla^2 g(\mathbf{w})]\| \leq \frac{2c_1\theta}{3\mu},$$

and so on $\mathcal{E}_1 \cap \mathcal{E}_\infty$, (2.4) holds for any constant $c_* \leq c_1/3$. Setting $t = c_1\theta/3\mu$ in Proposition 2.10, we obtain that for any fixed \mathbf{w} ,

$$\mathbb{P}\left[\|\nabla^2 g(\mathbf{w}) - \mathbb{E}[\nabla^2 g(\mathbf{w})]\| \geq \frac{c_1\theta}{3\mu}\right] \leq 4n \exp\left(-\frac{c_4 p \theta^2}{n^2}\right).$$

Taking a union bound, we obtain that

$$\mathbb{P}[\mathcal{E}_1^c] \leq 4n \exp\left(-\frac{c_4 p \theta^2}{n^2} + C_5 n \log(n) + C_5 n \log \log(p)\right).$$

Large gradient in region R_2 . Similarly, for the gradient quantity, for $\mathbf{w} \in R_2$, Proposition 2.6 shows that

$$\mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \right] \geq c_6 \theta.$$

Moreover, on \mathcal{E}_∞ , $\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|}$ is

$$L_2 \doteq \frac{C_7 n^2}{\mu} \log(np)$$

Lipschitz by Proposition 2.12. For any $\varepsilon < \frac{1}{20\sqrt{5}}$, the set R_2 has an ε -net N_2 of size at most $\left(\frac{3}{20\varepsilon\sqrt{5}}\right)^n$.

Set $\varepsilon = \frac{c_6 \theta}{3L_2}$, so

$$\#N_2 \leq \exp \left(n \log \left(\frac{C_8 n^2 \log(np)}{\theta \mu} \right) \right).$$

Let \mathcal{E}_2 denote the event

$$\mathcal{E}_2 = \left\{ \max_{\mathbf{w} \in N_2} \left| \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \right] \right| \leq \frac{c_6 \theta}{3} \right\}.$$

On $\mathcal{E}_2 \cap \mathcal{E}_\infty$,

$$\sup_{\mathbf{w} \in R_2} \left| \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \right] \right| \leq \frac{2c_6 \theta}{3}, \quad (7.8)$$

and so on $\mathcal{E}_2 \cap \mathcal{E}_\infty$, (2.5) holds for any constant $c_* \leq c_6/3$. Setting $t = c_6 \theta/3$ in Proposition 2.9, we obtain that for any fixed $\mathbf{w} \in R_2$,

$$\mathbb{P} \left[\left| \frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w})}{\|\mathbf{w}\|} \right] \right| \geq t \right] \leq 2 \exp \left(-\frac{c_9 p \theta^2}{n} \right),$$

and so

$$\mathbb{P}[\mathcal{E}_2^c] \leq 2 \exp \left(-\frac{c_9 p \theta^2}{n} + n \log \left(\frac{C_8 n^2 \log(np)}{\theta \mu} \right) \right). \quad (7.9)$$

Existence of negative curvature direction in R_3 . Finally, for any $\mathbf{w} \in R_3$, Proposition 2.5 shows that

$$\mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \leq -c_9 \theta.$$

On \mathcal{E}_∞ , $\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2}$ is

$$L_3 = \frac{C_{10} n^3}{\mu^2} \log^{3/2}(np)$$

Lipschitz by Proposition 2.11. As above, for any $\varepsilon \leq \sqrt{\frac{4n-1}{4n}}$, R_3 has an ε -net N_3 of size at most $(3/\varepsilon)^n$. Set $\varepsilon = c_9 \theta/3L_3$. Then

$$\#N_3 \leq \exp \left(n \log \left(\frac{C_{11} n^3 \log^{3/2}(np)}{\theta \mu^2} \right) \right).$$

Let \mathcal{E}_3 denote the event

$$\mathcal{E}_3 = \left\{ \max_{\mathbf{w} \in N_3} \left| \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \right| \leq \frac{c_9 \theta}{3} \right\}$$

On $\mathcal{E}_3 \cap \mathcal{E}_\infty$,

$$\sup_{\mathbf{w} \in R_3} \left| \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \right| \leq \frac{2c_9 \theta}{3},$$

and (2.6) holds with any constant $c_\star < c_9/3$. Setting $t = c_9 \theta/3$ in Proposition 2.8 and taking a union bound, we obtain

$$\mathbb{P}[\mathcal{E}_3^c] \leq 4 \exp \left(-\frac{c_{12} p \mu^2 \theta^2}{n^2} + n \log \left(\frac{C_{11} n^3 \log^{3/2}(np)}{\theta \mu^2} \right) \right).$$

The unique local minimizer located near $\mathbf{0}$. Let \mathcal{E}_g be the event that the bounds (2.4)-(2.6) hold. On \mathcal{E}_g , the function g is $\frac{c_\star \theta}{\mu}$ -strongly convex over $R_1 = \{\mathbf{w} \mid \|\mathbf{w}\| \leq \mu / (4\sqrt{2})\}$. This implies that f has at most one local minimum on R_1 . It also implies that for any $\mathbf{w} \in R_1$,

$$g(\mathbf{w}) \geq g(\mathbf{0}) + \langle \nabla g(\mathbf{0}), \mathbf{w} \rangle + \frac{c\theta}{2\mu} \|\mathbf{w}\|^2 \geq g(\mathbf{0}) - \|\mathbf{w}\| \|\nabla g(\mathbf{0})\| + \frac{c_\star \theta}{2\mu} \|\mathbf{w}\|^2.$$

So, if $g(\mathbf{w}) \leq g(\mathbf{0})$, we necessarily have

$$\|\mathbf{w}\| \leq \frac{2\mu}{c_\star \theta} \|\nabla g(\mathbf{0})\|.$$

Suppose that

$$\|\nabla g(\mathbf{0})\| \leq \frac{c_\star \theta}{32}. \tag{7.10}$$

Then $g(\mathbf{w}) \leq g(\mathbf{0})$ implies that $\|\mathbf{w}\| \leq \mu/16$. By Weierstrass's theorem, $g(\mathbf{w})$ has at least one minimizer \mathbf{w}_\star over the compact set $S = \{\mathbf{w} \mid \|\mathbf{w}\| \leq \mu/10\}$. By the above reasoning, $\|\mathbf{w}_\star\| \leq \mu/16$, and hence \mathbf{w}_\star does not lie on the boundary of S . This implies that \mathbf{w}_\star is a local minimizer of g . Moreover, as above,

$$\|\mathbf{w}_\star\| \leq \frac{2\mu}{c_\star \theta} \|\nabla g(\mathbf{0})\|.$$

We now use the vector Bernstein inequality to show that with our choice of p , (7.10) is satisfied with high probability. Notice that

$$\nabla g(\mathbf{0}) = \frac{1}{p} \sum_{i=1}^p \dot{h}_\mu(x_i(n)) \bar{\mathbf{x}}_i,$$

and \dot{h}_μ is bounded by one in magnitude, so for any integer $m \geq 2$,

$$\mathbb{E} \left[\left\| \dot{h}_\mu(x_i(n)) \bar{\mathbf{x}}_i \right\|^m \right] \leq \mathbb{E} [\|\mathbf{x}_i\|^m] \leq \mathbb{E}_{Z \sim \chi(n)} [Z^m] \leq m! n^{m/2},$$

where we have applied the moment estimate for the $\chi(n)$ distribution shown in Lemma A.8. Applying the vector Bernstein inequality in Corollary A.11 with $R = \sqrt{n}$ and $\sigma^2 = 2n$, we obtain

$$\mathbb{P}[\|\nabla g(\mathbf{0})\| \geq t] \leq 2(n+1) \exp\left(-\frac{pt^2}{4n + 2\sqrt{nt}}\right)$$

for all $t > 0$. Using this inequality, it is not difficult to show that there exist constants $C_{13}, C_{14} > 0$ such that when $p \geq C_{13}n \log n$, with probability at least $1 - 4np^{-10}$,

$$\|\nabla g(\mathbf{0})\| \leq C_3 \sqrt{\frac{n \log p}{p}}. \quad (7.11)$$

When $\frac{p}{\log p} \geq \frac{C_{14}n}{\theta^2}$, for appropriately large C_{14} , (7.11) implies (7.10). Summing up failure probabilities completes the proof. \blacksquare

7.3 Proofs for Section 2.3 and Theorem 2.3

Proof [of Lemma 2.14] By the generative model,

$$\bar{\mathbf{Y}} = \left(\frac{1}{p\theta} \mathbf{Y} \mathbf{Y}^*\right)^{-1/2} \mathbf{Y} = \left(\frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*\right)^{-1/2} \mathbf{A}_0 \mathbf{X}_0.$$

Since $\mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^* / (p\theta)] = \mathbf{I}$, we will compare $\left(\frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*\right)^{-1/2} \mathbf{A}_0$ with $(\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \mathbf{A}_0 = \mathbf{U} \mathbf{V}^*$. By Lemma B.2, we have

$$\begin{aligned} & \left\| \left(\frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*\right)^{-1/2} \mathbf{A}_0 - (\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \mathbf{A}_0 \right\| \\ & \leq \|\mathbf{A}_0\| \left\| \left(\frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*\right)^{-1/2} - (\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \right\| \\ & \leq \|\mathbf{A}_0\| \frac{2 \|\mathbf{A}_0\|^3}{\sigma_{\min}^4(\mathbf{A}_0)} \left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| = 2\kappa^4(\mathbf{A}_0) \left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| \end{aligned}$$

provided

$$\|\mathbf{A}_0\|^2 \left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| \leq \frac{\sigma_{\min}^2(\mathbf{A}_0)}{2} \iff \left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| \leq \frac{1}{2\kappa^2(\mathbf{A}_0)}.$$

On the other hand, by Lemma B.3, when $p \geq C_1 n^2 \log n$ for some large constant C_1 , $\left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| \leq 10 \sqrt{\frac{\theta n \log p}{p}}$ with probability at least $1 - p^{-8}$. Thus, when $p \geq C_2 \kappa^4(\mathbf{A}_0) \theta n^2 \log(n\theta\kappa(\mathbf{A}_0))$,

$$\left\| \left(\frac{1}{p\theta} \mathbf{A}_0 \mathbf{X}_0 \mathbf{X}_0^* \mathbf{A}_0^*\right)^{-1/2} \mathbf{A}_0 - (\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \mathbf{A}_0 \right\| \leq 20\kappa^4(\mathbf{A}_0) \sqrt{\frac{\theta n \log p}{p}},$$

as desired. \blacksquare

Proof [of Lemma 2.15] To avoid clutter in notation, we write \mathbf{X} to mean \mathbf{X}_0 , and \mathbf{x}_k to mean $(\mathbf{x}_0)_k$ in this proof. We also let $\tilde{\mathbf{Y}} \doteq \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}}\mathbf{X}_0$. Note the Jacobian matrix for the mapping $\mathbf{q}(\mathbf{w})$ is $\nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w}) = \left[\mathbf{I}, -\mathbf{w}/\sqrt{1-\|\mathbf{w}\|^2} \right]$. Hence for any vector $\mathbf{z} \in \mathbb{R}^n$ and all $\mathbf{w} \in \Gamma$,

$$\|\nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})\mathbf{z}\| \leq \sqrt{n-1}\|\mathbf{z}\|_{\infty} + \frac{\|\mathbf{w}\|}{\sqrt{1-\|\mathbf{w}\|^2}}\|\mathbf{z}\|_{\infty} \leq 3\sqrt{n}\|\mathbf{z}\|_{\infty}.$$

Now we have

$$\begin{aligned} & \left\| \nabla_{\mathbf{w}}g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}}g(\mathbf{w}; \mathbf{X}) \right\| \\ &= \left\| \frac{1}{p} \sum_{k=1}^p \dot{h}_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k + \mathbf{q}^*(\mathbf{w})\tilde{\boldsymbol{\Xi}}\mathbf{x}_k) \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})(\mathbf{x}_k + \tilde{\boldsymbol{\Xi}}\mathbf{x}_k) - \frac{1}{p} \sum_{k=1}^p \dot{h}_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k) \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})\mathbf{x}_k \right\| \\ &\leq \left\| \frac{1}{p} \sum_{k=1}^p \dot{h}_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k + \mathbf{q}^*(\mathbf{w})\tilde{\boldsymbol{\Xi}}\mathbf{x}_k) \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})(\mathbf{x}_k + \tilde{\boldsymbol{\Xi}}\mathbf{x}_k - \mathbf{x}_k) \right\| \\ &\quad + \left\| \frac{1}{p} \sum_{k=1}^p \left[\dot{h}_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k + \mathbf{q}^*(\mathbf{w})\tilde{\boldsymbol{\Xi}}\mathbf{x}_k) - \dot{h}_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k) \right] \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})\mathbf{x}_k \right\| \\ &\leq \left\| \tilde{\boldsymbol{\Xi}} \right\| \left(\max_t \dot{h}_{\mu}(t) 3n \|\mathbf{X}\|_{\infty} + L_{\dot{h}_{\mu}} 3n \|\mathbf{X}\|_{\infty}^2 \right), \end{aligned}$$

where $L_{\dot{h}_{\mu}}$ denotes the Lipschitz constant for $\dot{h}_{\mu}(\cdot)$. Similarly, suppose $\left\| \tilde{\boldsymbol{\Xi}} \right\| \leq \frac{1}{2n}$, and also notice that

$$\left\| \frac{\mathbf{I}}{q_n(\mathbf{w})} + \frac{\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right\| \leq \frac{1}{q_n(\mathbf{w})} + \frac{\|\mathbf{w}\|^2}{q_n^3(\mathbf{w})} = \frac{1}{q_n^3(\mathbf{w})} \leq 2\sqrt{2}n^{3/2},$$

we obtain that

$$\begin{aligned} & \left\| \nabla_{\mathbf{w}}^2g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}}^2g(\mathbf{w}; \mathbf{X}) \right\| \\ &\leq \left\| \frac{1}{p} \sum_{k=1}^p \left[\ddot{h}(\mathbf{q}^*(\mathbf{w})\tilde{\mathbf{y}}_k) \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})\tilde{\mathbf{y}}_k\tilde{\mathbf{y}}_k^* (\nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w}))^* - \ddot{h}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k) \nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w})\mathbf{x}_k\mathbf{x}_k^* (\nabla_{\mathbf{w}}\mathbf{q}(\mathbf{w}))^* \right] \right\| \\ &\quad + \left\| \frac{1}{p} \sum_{k=1}^p \left[\dot{h}(\mathbf{q}^*(\mathbf{w})\tilde{\mathbf{y}}_k) \left(\frac{\mathbf{I}}{q_n(\mathbf{w})} + \frac{\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right) \tilde{\mathbf{y}}_k(n) - \dot{h}(\mathbf{q}^*(\mathbf{w})\mathbf{x}_k) \left(\frac{\mathbf{I}}{q_n(\mathbf{w})} + \frac{\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right) \mathbf{x}_k(n) \right] \right\| \\ &\leq \frac{45}{2}L_{\ddot{h}_{\mu}}n^{3/2}\|\mathbf{X}\|_{\infty}^3\|\tilde{\boldsymbol{\Xi}}\| + \max_t \ddot{h}_{\mu}(t) \left(18n^{3/2}\|\mathbf{X}\|_{\infty}^2\|\tilde{\boldsymbol{\Xi}}\| + 10n^2\|\mathbf{X}\|_{\infty}^2\|\tilde{\boldsymbol{\Xi}}\|^2 \right) \\ &\quad + 3\sqrt{2}L_{\dot{h}_{\mu}}n^2\|\tilde{\boldsymbol{\Xi}}\|\|\mathbf{X}\|_{\infty}^2 + \max_t \dot{h}_{\mu}(t) 2\sqrt{2}n^2\|\tilde{\boldsymbol{\Xi}}\|\|\mathbf{X}\|_{\infty}, \end{aligned}$$

where $L_{\ddot{h}_{\mu}}$ denotes the Lipschitz constant for $\ddot{h}_{\mu}(\cdot)$. Since

$$\max_t \dot{h}_{\mu}(t) \leq 1, \quad \max_t \ddot{h}_{\mu}(t) \leq \frac{1}{\mu}, \quad L_{h_{\mu}} \leq 1, \quad L_{\dot{h}_{\mu}} \leq \frac{1}{\mu}, \quad L_{\ddot{h}_{\mu}} \leq \frac{2}{\mu^2},$$

and by Lemma 7.11, $\|\mathbf{X}\|_\infty \leq 4\sqrt{\log(np)}$ with probability at least $1 - \theta(np)^{-7} - \exp(-0.3\theta np)$, we obtain

$$\begin{aligned} \left\| \nabla_{\mathbf{w}} g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}) \right\| &\leq C_1 \frac{n}{\mu} \log(np) \left\| \tilde{\boldsymbol{\Xi}} \right\|, \\ \left\| \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}) \right\| &\leq C_2 \max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \left\| \tilde{\boldsymbol{\Xi}} \right\| \end{aligned}$$

for numerical constants $C_1, C_2 > 0$. ■

Proof [of Theorem 2.3] Assume the constant c_* as defined in Theorem 2.1. By Lemma 2.14, when

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

the magnitude of the perturbation is bounded as

$$\left\| \tilde{\boldsymbol{\Xi}} \right\| \leq C_2 c_* \theta \left(\max \left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1},$$

where C_2 can be made arbitrarily small by making C_1 large. Combining this result with Lemma 2.15, we obtain that for all $\mathbf{w} \in \Gamma$,

$$\begin{aligned} \left\| \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) - \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}) \right\| &\leq \frac{c_* \theta}{2} \\ \left\| \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) - \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}) \right\| &\leq \frac{c_* \theta}{2}, \end{aligned}$$

with probability at least $1 - p^{-8} - \theta(np)^{-7} - \exp(-0.3\theta np)$. In view of (2.11) in Theorem 2.1, we have

$$\begin{aligned} \frac{\mathbf{w}^* g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} &= \frac{\mathbf{w}^* g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} + \frac{\mathbf{w}^* g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} - \frac{\mathbf{w}^* g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \\ &\leq -c_* \theta + \left\| \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) - \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}) \right\| \leq -\frac{1}{2} c_* \theta. \end{aligned}$$

By similar arguments, we obtain (2.9) through (2.11) in Theorem 2.3.

To show the unique local minimizer over Γ is near $\mathbf{0}$, we note that (recall the last part of proof of Theorem 2.1 in Section 7.2) $g(\mathbf{w}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0)$ being $\frac{c_* \theta}{2\mu}$ strongly convex near $\mathbf{0}$ implies that

$$\|\mathbf{w}_*\| \leq \frac{4\mu}{c_* \theta} \left\| \nabla g(\mathbf{0}; \mathbf{X}_0 + \tilde{\boldsymbol{\Xi}} \mathbf{X}_0) \right\|.$$

The above perturbation analysis implies there exists $C_3 > 0$ such that when

$$p \geq \frac{C_3}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right),$$

it holds that

$$\left\| \nabla_{\mathbf{w}} g(\mathbf{0}; \mathbf{X}_0 + \tilde{\Xi} \mathbf{X}_0) - \nabla_{\mathbf{w}} g(\mathbf{0}; \mathbf{X}) \right\| \leq \frac{c_* \theta}{400},$$

which in turn implies

$$\|\mathbf{w}_*\| \leq \frac{4\mu}{c_* \theta} \|\nabla g(\mathbf{0}; \mathbf{X}_0)\| + \frac{4\mu}{c_* \theta} \frac{c_* \theta}{400} \leq \frac{\mu}{8} + \frac{\mu}{100} < \frac{\mu}{7},$$

where we have recall the result that $\frac{2\mu}{c_* \theta} \|\nabla g(\mathbf{0}; \mathbf{X}_0)\| \leq \mu/16$ from proof of Theorem 2.1. A simple union bound with careful bookkeeping gives the success probability. ■

8 Proof of Convergence for the Trust-Region Algorithm

Proof [of Lemma 3.3] Using the fact $\tanh(\cdot)$ and $1 - \tanh^2(\cdot)$ are bounded by one in magnitude, by (3.12) and (3.13) we have

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

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

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

where at the last line we have used the fact the mapping $\mathbf{q} \mapsto \mathbf{q}^* \mathbf{x}_k / \mu$ is $\|\mathbf{x}_k\| / \mu$ Lipschitz, and $x \mapsto \tanh(x)$ is 1-Lipschitz, and the composition rule in Lemma 7.5. Similar argument yields the final bound. ■

Proof [of Lemma 3.4] Suppose we can establish

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

Applying this twice we obtain

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

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}\| \\ & \quad + 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 3.3. By the integral form of Taylor's theorem in Lemma A.12 and the result above, we have

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

$$\leq t^2 \int_0^1 (1-s) st \eta_f ds = \frac{\eta_f t^3}{6},$$

with $t = \|\delta\|$ we obtain the desired result. \blacksquare

Proof [of Lemma 3.5] By the integral form of Taylor's theorem in Lemma A.12, 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 $\delta \in T_{\mathbf{q}}\mathbb{S}^{n-1}$ such that $\mathbf{q}(\mathbf{w}') = \exp_{\mathbf{q}}(\delta)$. It remains to show that $\|\delta\| \leq \Delta$. By Lemma 7.7, we know that $\|\mathbf{q}(\mathbf{w}') - \mathbf{q}(\mathbf{w})\| \leq 2\sqrt{n}\|\mathbf{w}' - \mathbf{w}\| \leq 3\Delta/\pi$. Hence,

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

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

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

as claimed. \blacksquare

Proof [of Lemma 3.6] 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.12, we have

$$g\left(\mathbf{w} - t\sigma \frac{\mathbf{w}}{\|\mathbf{w}\|}\right)$$

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

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 3.5, 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. \blacksquare

Proof [of Lemma 3.8] For any $t \in \left[0, \frac{\Delta}{\|\text{grad } f(\mathbf{q}^{(k)})\|}\right]$, it holds that $\|t \text{ grad } f(\mathbf{q}^{(k)})\| \leq \Delta$, and the quadratic approximation

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

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

$$\widehat{f}\left(\mathbf{q}^{(k)}, -t_0 \text{ grad } f\left(\mathbf{q}^{(k)}\right)\right) \leq f\left(\mathbf{q}^{(k)}\right) - \frac{1}{2} \min \left\{ \frac{\Delta}{\|\text{grad } f(\mathbf{q}^{(k)})\|}, \frac{1}{M_H} \right\} \left\|\text{grad } f\left(\mathbf{q}^{(k)}\right)\right\|^2. \quad (8.1)$$

Now let \mathbf{U} be an arbitrary orthonormal basis for $T_{\mathbf{q}^{(k)}}\mathbb{S}^{n-1}$. Since the norm constraint is active, by the optimality condition in (3.16), we have

$$\Delta \leq \left\| \left[\mathbf{U}^* \text{Hess } f\left(\mathbf{q}^{(k)}\right) \mathbf{U} \right]^{-1} \mathbf{U}^* \text{grad } f\left(\mathbf{q}^{(k)}\right) \right\|$$

$$\leq \left\| \left[\mathbf{U}^* \text{Hess } f(\mathbf{q}^{(k)}) \mathbf{U} \right]^{-1} \right\| \left\| \mathbf{U}^* \text{grad } f(\mathbf{q}^{(k)}) \right\| \leq \frac{\|\text{grad } f(\mathbf{q}^{(k)})\|}{m_H},$$

which means that $\|\text{grad } f(\mathbf{q}^{(k)})\| \geq m_H \Delta$. Substituting this into (8.1), we obtain

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

By the key comparison result established in proof of Lemma 3.4, we have

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

This completes the proof. \blacksquare

It takes certain delicate work to prove Lemma 3.9. 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 8.1 For $\tau \in [0, 1]$ and $\|\boldsymbol{\delta}\| \leq 1/2$, we have

$$\|\mathcal{P}_\gamma^{\tau \leftarrow 0} - \mathbf{I}\| \leq \frac{5}{4} \tau \|\boldsymbol{\delta}\|, \quad (8.2)$$

$$\|\mathcal{P}_\gamma^{0 \leftarrow \tau} - \mathbf{I}\| \leq \frac{3}{2} \tau \|\boldsymbol{\delta}\|. \quad (8.3)$$

Proof By (3.17), 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

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

$$\leq \frac{5}{4}\tau \|\boldsymbol{\delta}\| \frac{1}{\cos(\tau \|\boldsymbol{\delta}\|)} \leq \frac{5}{4}\tau \|\boldsymbol{\delta}\| \frac{1}{\cos(1/2)} \leq \frac{3}{2}\tau \|\boldsymbol{\delta}\|,$$

completing the proof. \blacksquare

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

Lemma 8.2 *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 $\tau \in [0, 1]$,*

$$\|\mathcal{P}_{\gamma}^{0\leftarrow\tau} \text{Hess } f(\gamma(\tau))\mathcal{P}_{\gamma}^{\tau\leftarrow 0} - \text{Hess } f(\mathbf{q})\| \leq L_H \cdot \tau \|\boldsymbol{\delta}\|, \quad (8.4)$$

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

Proof First of all, by (3.15) and using the fact that the operator norm of a projection operator is unitary bounded, we have

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

By the estimates in Lemma 3.3, we obtain

$$\begin{aligned} & \|\text{Hess } f(\gamma(\tau)) - \text{Hess } f(\mathbf{q})\| \\ & \leq \frac{2}{\mu^2}n^{3/2} \|\mathbf{X}\|_{\infty}^3 \|\gamma(\tau) - \mathbf{q}\| + \frac{n}{\mu} \|\mathbf{X}\|_{\infty}^2 \|\gamma(\tau) - \mathbf{q}\| + \sqrt{n} \|\mathbf{X}\|_{\infty} \|\gamma(\tau) - \mathbf{q}\| \\ & \quad + 2 \|\gamma(\tau) \gamma^*(\tau) - \mathbf{q}\mathbf{q}^*\| \left(\frac{n}{\mu} \|\mathbf{X}\|_{\infty}^2 + \sqrt{n} \|\mathbf{X}\|_{\infty} \right) \\ & \leq \left(\frac{5}{2\mu^2}n^{3/2} \|\mathbf{X}\|_{\infty}^3 + \frac{25n}{4\mu} \|\mathbf{X}\|_{\infty}^2 + \frac{25}{4}\sqrt{n} \|\mathbf{X}\|_{\infty} \right) \tau \|\boldsymbol{\delta}\|, \quad (8.5) \end{aligned}$$

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

$$\begin{aligned} \|\gamma(\tau) - \mathbf{q}\| & = \left\| \mathbf{q} (\cos(\tau \|\boldsymbol{\delta}\|) - 1) + \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|} \sin(\tau \|\boldsymbol{\delta}\|) \right\| \leq \frac{5}{4}\tau \|\boldsymbol{\delta}\|, \quad (\text{Proof of Lemma 8.1}) \\ \|\gamma(\tau) \gamma^*(\tau) - \mathbf{q}\mathbf{q}^*\| & \leq \left\| \left(\frac{\boldsymbol{\delta}\boldsymbol{\delta}^*}{\|\boldsymbol{\delta}\|^2} - \mathbf{q}\mathbf{q}^* \right) \sin^2(\tau \|\boldsymbol{\delta}\|) \right\| + 2 \sin(\tau \|\boldsymbol{\delta}\|) \cos(\tau \|\boldsymbol{\delta}\|) \\ & \leq \sin^2(\tau \|\boldsymbol{\delta}\|) + \sin(2\tau \|\boldsymbol{\delta}\|) \leq \frac{5}{2}\tau \|\boldsymbol{\delta}\|. \end{aligned}$$

Therefore, by Lemma 8.1, we obtain

$$\|\mathcal{P}_{\gamma}^{0\leftarrow\tau} \text{Hess } f(\gamma(\tau))\mathcal{P}_{\gamma}^{\tau\leftarrow 0} - \text{Hess } f(\mathbf{q})\|$$

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

By Lemma 3.3 and substituting the estimate in (8.5), we obtain the claimed result. \blacksquare

Proof [of Lemma 3.9] 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 $T_{\mathbf{q}}\mathbb{S}^{n-1}$. We could compare $\mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}$ with $\nabla_{\mathbf{w}}^2 g(\mathbf{w})$, and build on the known results for the latter. Instead, we present a direct proof here that yields tighter results as stated in the lemma. Again we first work with the ‘‘canonical’’ section in the vicinity of \mathbf{e}_n with the ‘‘canonical’’ reparametrization $\mathbf{q}(\mathbf{w}) = [\mathbf{w}; \sqrt{1 - \|\mathbf{w}\|^2}]$.

By definition of the Riemannian Hessian in (3.15), expressions of $\nabla^2 f$ and ∇f in (3.12) and (3.13), and exchange of differential and expectation operators (justified similarly as in Section 7.1.3), we obtain

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

Now because $\mathbf{U}^* \mathbf{q} = \mathbf{0}$, we have

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{\mu} \left\{ 1 - \tanh^2\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right) \right\} \mathbf{U}^* \mathbf{x}_k \mathbf{x}_k^* \mathbf{U}\right] &= \frac{1}{\mu} \mathbb{E}\left[1 - \tanh^2\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right)\right] \mathbb{E}[\mathbf{U}^* \mathbf{x}_k \mathbf{x}_k^* \mathbf{U}] \\
&= \frac{\theta}{\mu} \mathbb{E}\left[1 - \tanh^2\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right)\right] \mathbf{I}.
\end{aligned}$$

So it is enough to bound the scalar part:

$$\begin{aligned}
&\frac{\theta}{\mu} \mathbb{E}\left[1 - \tanh^2\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right)\right] - \mathbb{E}\left[\tanh\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right) \mathbf{q}^* \mathbf{x}_k\right] \\
&= \mathbb{E}_{\mathcal{I}} \left[\frac{\theta}{\mu} \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [1 - \tanh^2(Z/\mu)] - \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [\tanh(Z/\mu) Z] \right] \\
&= \mathbb{E}_{\mathcal{I}} \left[\left(\frac{\theta}{\mu} - \frac{\|\mathbf{q}_{\mathcal{I}}\|^2}{\mu} \right) \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{q}_{\mathcal{I}}\|^2)} [1 - \tanh^2(Z/\mu)] \right] \quad (\text{by (B.6)}) \\
&= (1 - \theta) \mathbb{E}_{\mathcal{J}} \left[\left(\frac{\theta}{\mu} - \frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\mu} \right) \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)} [1 - \tanh^2(Z/\mu)] \right] \\
&\quad + \theta \mathbb{E}_{\mathcal{J}} \left[\left(\frac{\theta}{\mu} - \frac{\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2}{\mu} \right) \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [1 - \tanh^2(Z/\mu)] \right] \\
&\geq (1 - \theta) \left(\frac{\theta}{\mu} - \frac{\mu}{32} \right) \mathbb{E}_{\mathcal{J}} \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)} [1 - \tanh^2(Z/\mu)] \\
&\quad - \frac{\theta}{\mu} \mathbb{E}_{\mathcal{J}} \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [1 - \tanh^2(Z/\mu)], \tag{8.6}
\end{aligned}$$

where to obtain the last line above we have used the fact $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$ and that $\|\mathbf{w}_{\mathcal{J}}\| + q_n^2 \leq \|\mathbf{q}\|^2 = 1$. Now by Lemma A.1,

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

where at the last inequality we have applied Gaussian tail lower bound of Type III in Lemma A.5. Also note the function $x \mapsto \sqrt{4+x^2} - x$ is monotonically decreasing and $\|\mathbf{w}_{\mathcal{J}}\| \leq \|\mathbf{w}\| \leq \mu/(4\sqrt{2})$, we obtain

$$\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)} [1 - \tanh^2(Z/\mu)] \geq \frac{1}{\sqrt{2\pi}} \left(\sqrt{\frac{1}{8} + 4} - \frac{2}{4\sqrt{2}} \right) \geq \frac{3}{2\sqrt{2\pi}}. \quad (8.7)$$

Similarly we also have

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

where at the last inequality we have applied Gaussian tail upper bound of Type II in Lemma A.5. 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

$$\mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2)} [1 - \tanh^2(Z/\mu)] \leq \frac{4}{\sqrt{2\pi}} \frac{\mu}{\sqrt{31/32}} \leq \frac{4}{\sqrt{2\pi}} \mu. \quad (8.8)$$

Substituting (8.7) and (8.8) into (8.6), we obtain

$$\begin{aligned} &\frac{\theta}{\mu} \mathbb{E} \left[1 - \tanh^2\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right) \right] - \mathbb{E} \left[\tanh\left(\frac{\mathbf{q}^* \mathbf{x}_k}{\mu}\right) \mathbf{q}^* \mathbf{x}_k \right] \\ &\geq \frac{3}{2\sqrt{2\pi}} (1 - \theta) \left(\frac{\theta}{\mu} - \frac{\mu}{32} \right) - \frac{4}{\sqrt{2\pi}} \theta \\ &\geq \frac{3}{4\sqrt{2\pi}} \left(\frac{\theta}{\mu} - \frac{1}{10} \frac{\theta}{\mu} \right) - \frac{4}{10\sqrt{2\pi}} \frac{\theta}{\mu} \geq \frac{1}{4\sqrt{2\pi}} \frac{\theta}{\mu}, \end{aligned}$$

where at the last line we have used the fact $\mu < \theta$ and that $\mu < 1/10$ to simplify. To summarize, we obtain that

$$\mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} \succeq \frac{1}{4\sqrt{2\pi}} \frac{\theta}{\mu}. \quad (8.9)$$

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

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

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

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

where we have used Lemma 7.4 to obtain the last inequality. By Lemma A.7, 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.10, we obtain

$$\mathbb{P} \left[\left\| \frac{1}{p} \sum_{k=1}^p \mathbf{W}_k - \frac{1}{p} \sum_{k=1}^p \mathbb{E}[\mathbf{W}_k] \right\| > t/2 \right] \leq 2n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right) \quad (8.10)$$

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

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

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

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

where at the first inequality we used the fact $|\tanh(\cdot)| \leq 1$, at the second we invoked Lemma 7.4, and at the third we invoked Lemma A.6. Taking $R_Z = \sigma_Z^2 = 1$, by Lemma A.9, we obtain

$$\mathbb{P} \left[\left| \frac{1}{p} \sum_{k=1}^p Z_k - \frac{1}{p} \sum_{k=1}^p \mathbb{E}[Z_k] \right| > t/2 \right] \leq 2 \exp(-pt^2/16) \quad (8.11)$$

for any $t > 0$. Gathering (8.10) and (8.11), we obtain that for any $t > 0$,

$$\begin{aligned} & \mathbb{P}[\|\mathbf{U}^* \text{Hess } \mathbb{E}[f(\mathbf{q})] \mathbf{U} - \mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U}\| > t] \\ & \leq \mathbb{P}[\|\mathbf{U}^* \nabla^2 f(\mathbf{q}) \mathbf{U} - \nabla^2 \mathbb{E}[f(\mathbf{q})]\| > t/2] + \mathbb{P}[|\langle \nabla f(\mathbf{q}), \mathbf{q} \rangle - \langle \nabla \mathbb{E}[f(\mathbf{q})], \mathbf{q} \rangle| > t/2] \\ & \leq 2n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right) + 2 \exp \left(-\frac{pt^2}{16} \right) \leq 4n \exp \left(-\frac{p\mu^2 t^2}{32n^2 + 8nt} \right). \end{aligned} \quad (8.12)$$

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 8.2, 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 8.2). 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 7.11, $\mathbb{P}[\mathcal{E}_\infty^c] \leq \theta(np)^{-7} + \exp(-0.3\theta np)$. By Lemma 8.2, 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

$$\mathbf{U}^* \text{Hess } f(\mathbf{q}) \mathbf{U} \succeq c_\# \frac{\theta}{\mu} \tag{8.13}$$

for any $c_\# \leq 1/(12\sqrt{2\pi})$. Setting $t = \frac{\theta}{12\sqrt{2\pi}\mu}$ in (8.12), 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. \blacksquare

Proof [of Lemma 3.11] 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 (3.12) and (3.14), we have

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

where to get the last line we have used (7.1). 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. \blacksquare

Proof of Lemma 3.12 combines the local Lipschitz property of Hess $f(\mathbf{q})$ in Lemma 8.2, and the Taylor's theorem (manifold version, Lemma 7.4.7 of [AMS09]).

Proof [of Lemma 3.12] Let $\gamma(t)$ be the unique geodesic that satisfies $\gamma(0) = \mathbf{q}^{(k)}$, $\gamma(1) = \mathbf{q}^{(k+1)}$, and its directional derivative $\dot{\gamma}(0) = \boldsymbol{\delta}_*$. Since the parallel translation defined by the Riemannian connection is an isometry, then $\|\text{grad } f(\mathbf{q}^{(k+1)})\| = \|\mathcal{P}_\gamma^{0 \leftarrow 1} \text{grad } f(\mathbf{q}^{(k+1)})\|$. Moreover, since $\|\boldsymbol{\delta}_*\| \leq \Delta$, the unconstrained optimality condition in (3.16) implies that $\text{grad } f(\mathbf{q}^{(k)}) + \text{Hess } f(\mathbf{q}^{(k)}) \boldsymbol{\delta}_* = \mathbf{0}_{\mathbf{q}^{(k)}}$. Thus, by using Taylor's theorem in [AMS09], we have

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

$$\begin{aligned}
&= \left\| \int_0^1 \left(\mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) \mathcal{P}_\gamma^{t \leftarrow 0} \boldsymbol{\delta}_* - \text{Hess } f(\mathbf{q}^{(k)}) \boldsymbol{\delta}_* \right) dt \right\| \\
&\leq \|\boldsymbol{\delta}_*\| \int_0^1 \left\| \mathcal{P}_\gamma^{0 \leftarrow t} \text{Hess } f(\gamma(t)) \mathcal{P}_\gamma^{t \leftarrow 0} - \text{Hess } f(\mathbf{q}^{(k)}) \right\| dt.
\end{aligned}$$

From the Lipschitz bound in Lemma 8.2 and the optimality condition in (3.16), we obtain

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

This completes the proof. ■

Proof [of Lemma 3.14] By invoking Taylor's theorem in [AMS09], 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. ■

9 Proofs of Technical Results for Section 4

We need one technical lemma to prove Lemma 4.2 and the relevant lemma for complete dictionaries.

Lemma 9.1 *There exists a positive constant C , such that 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 $1 - n_2^{-10} - \theta (n_1 n_2)^{-7} - \exp(-0.3 \theta n_1 n_2)$.

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 7.4 and Lemma A.6, 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.9, 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

$$\mathbb{P}\left[T(\mathbf{v}) < \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta\right] \leq \exp(-c_1 \theta^2 n_2) \quad (9.1)$$

for some positive constant c_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

$$\mathbb{P}[\mathcal{E}_{bg}^c] \leq \exp\left(-c_1 \theta^2 n_2 + n_1 \log\left(\frac{3}{\varepsilon}\right)\right) \leq \exp\left(-c_1 \theta^2 n_2 + c_2 n_1 \log \frac{n_1 \log n_2}{\theta}\right), \quad (9.2)$$

where $c_2 > 0$ is numerical. 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,

$$T(\mathbf{z}) = \|(\mathbf{v} + \mathbf{e})^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|(\mathbf{v} + \mathbf{e})^* \mathbf{M}_{\mathcal{I}}\|_1 \geq T(\mathbf{v}) - \|\mathbf{e}^* \mathbf{M}_{\mathcal{I}^c}\|_1 - \|\mathbf{e}^* \mathbf{M}_{\mathcal{I}}\|_1$$

$$\begin{aligned}
&= \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \|\mathbf{e}^* \mathbf{M}\|_1 = \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \sum_{k=1}^{n_2} |\mathbf{e}^* \mathbf{m}_k| \\
&\geq \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \varepsilon \sum_{k=1}^{n_2} \|\mathbf{m}_k\|.
\end{aligned}$$

By Lemma 7.11, with probability at least $1 - \theta (n_1 n_2)^{-7} - \exp(-0.3\theta n_1 n_2)$, $\|\mathbf{M}\|_\infty \leq 4\sqrt{\log(n_1 n_2)}$. Thus,

$$T(\mathbf{z}) \geq \frac{n_2}{5} \sqrt{\frac{2}{\pi}} \theta - \sqrt{\frac{2}{\pi}} \frac{\theta}{120} \frac{n_2 \sqrt{n_1} 4 \sqrt{\log(n_1 n_2)}}{\sqrt{n_1} \sqrt{\log(n_1 n_2)}} = \frac{n_2}{6} \sqrt{\frac{2}{\pi}} \theta. \quad (9.3)$$

Thus, by (9.2), it is enough to take $n_2 > C n_1 \log(n_1/\theta)/\theta^2$ for sufficiently large $C > 0$ to make the overall failure probability small enough so that the lower bound (9.3) holds. \blacksquare

Proof [Proof of Lemma 4.2] The proof is similar to that of [QSW14]. First, let us assume the dictionary $\mathbf{A}_0 = \mathbf{I}$. Wlog, 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 (4.1) takes the form:

$$\text{minimize}_{\mathbf{q}} \|\mathbf{q}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad \langle \mathbf{r}, \mathbf{q} \rangle = 1. \quad (9.4)$$

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 $\mathbf{X}_0 = [\bar{\mathbf{X}}; \mathbf{x}_n^*]$, where $\bar{\mathbf{X}} \in \mathbb{R}^{(n-1) \times p}$ and \mathbf{x}_n^* is the last row of \mathbf{X}_0 . Set $\mathbf{q} = [\bar{\mathbf{q}}, q_n]$, where $\bar{\mathbf{q}}$ 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 (9.4),

$$\text{minimize}_{\mathbf{q}} \|\mathbf{q}^* \mathbf{X}_0\|_1, \quad \text{subject to} \quad q_n r_n + \langle \bar{\mathbf{q}}, \bar{\mathbf{r}} \rangle \geq 1, \quad (9.5)$$

It is obvious that the feasible set of (9.5) contains that of (9.4). So if \mathbf{e}_n is the unique optimal solution (UOS) of (9.5), it is the UOS of (9.4). 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} . (9.5) can be further relaxed as

$$\text{minimize}_{\mathbf{q}} \|\mathbf{x}_n\|_1 |q_n| - \|\bar{\mathbf{q}}^* \bar{\mathbf{X}}_{\mathcal{I}}\|_1 + \|\bar{\mathbf{q}}^* \bar{\mathbf{X}}_{\mathcal{I}^c}\|_1, \quad \text{subject to} \quad q_n r_n + \|\bar{\mathbf{q}}\| \|\bar{\mathbf{r}}\| \geq 1. \quad (9.6)$$

The objective value of (9.6) lower bounds that of (9.5), and are equal when $\mathbf{q} = \mathbf{e}_n$. So if $\mathbf{q} = \mathbf{e}_n$ is UOS of (9.6), it is UOS of (9.4). By Lemma 9.1, we know that

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

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 (9.6) by

$$\text{minimize}_{\mathbf{q}} \|\mathbf{x}_n\|_1 |q_n| + \zeta \|\bar{\mathbf{q}}\|, \quad \text{subject to} \quad q_n r_n + \|\bar{\mathbf{q}}\| \|\bar{\mathbf{r}}\| \geq 1. \quad (9.7)$$

By similar arguments, if \mathbf{e}_n is the UOS of (9.7), it is also the UOS of (9.4). For the optimal solution of (9.7), notice that it is necessary to have $\text{sign}(q_n) = \text{sign}(r_n)$ and $q_n r_n + \|\bar{\mathbf{q}}\| \|\bar{\mathbf{r}}\| = 1$. Therefore, the problem (9.7) is equivalent to

$$\text{minimize}_{q_n} \|\mathbf{x}_n\|_1 |q_n| + \zeta \frac{1 - |r_n| |q_n|}{\|\bar{\mathbf{r}}\|}, \quad \text{subject to} \quad |q_n| \leq \frac{1}{|r_n|}. \quad (9.8)$$

Notice that the problem (9.8) 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, $\mathbf{q} = \mathbf{e}_n$ is the UOS of (9.8) if and only if

$$\frac{1}{|r_n|} \|\mathbf{x}_n\|_1 < \frac{\zeta}{\|\bar{\mathbf{r}}\|}. \quad (9.9)$$

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

$$\mathbb{P} \left[\|\mathbf{x}_n\|_1 \geq \frac{5}{4} \sqrt{\frac{2}{\pi}} \theta p \right] \leq \exp \left(-\frac{\theta^2 p}{64\pi} \right). \quad (9.10)$$

Therefore, by (9.9) and (9.10), for $\mathbf{q} = \mathbf{e}_n$ to be the UOS of (9.4) w.h.p., it is sufficient to have

$$\frac{5}{4|r_n|} \sqrt{\frac{2}{\pi}} \theta p < \frac{\theta p}{6\sqrt{1-|r_n|^2}} \sqrt{\frac{2}{\pi}}, \quad (9.11)$$

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, wlog, 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 (9.4) 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 $\langle \mathbf{A}_0^* \mathbf{r}, \mathbf{e}_n \rangle > 249/250$ to exactly recover the target, which is equivalent to $\langle \mathbf{r}, \hat{\mathbf{q}}_* \rangle > 249/250$. This implies that our rounding (4.1) is invariant to change of basis, completing the proof. ■

Proof [of Lemma 4.4] Define $\tilde{\mathbf{q}} \doteq (\mathbf{U}\mathbf{V}^* + \mathbf{\Xi})^* \mathbf{q}$. By Lemma 2.14, and in particular (2.15), when $p \geq \frac{C}{c_*^2 \theta} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8 (\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta} \right)$, $\|\mathbf{\Xi}\| \leq 1/2$ so that $\mathbf{U}\mathbf{V}^* + \mathbf{\Xi}$ is invertible. Then the LP rounding can be written as

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

By Lemma 4.2, to obtain $\tilde{\mathbf{q}} = \mathbf{e}_n$ from this LP, it is enough to have

$$\langle (\mathbf{U}\mathbf{V}^* + \mathbf{\Xi})^{-1} \mathbf{r}, \mathbf{e}_n \rangle \geq 249/250,$$

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

Proof [of Lemma 4.5] 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 $\mathbf{U}^* \left[\mathbf{q}_*^1, \dots, \mathbf{q}_*^\ell \mid \widehat{\mathbf{V}} \right] = \begin{bmatrix} \mathbf{0} & \mathbf{U}^* \widehat{\mathbf{V}} \end{bmatrix}$, we have

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \begin{bmatrix} \mathbf{0} & \mathbf{U}^* \widehat{\mathbf{V}} \end{bmatrix} \mathbf{X}_0 + \begin{bmatrix} \mathbf{0} & \mathbf{U}^* \widehat{\mathbf{V}} \end{bmatrix} \mathbf{\Delta}_1 \mathbf{X}_0 = \mathbf{U}^* \widehat{\mathbf{V}} \mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}_2 \mathbf{X}_0, \quad (9.12)$$

where $\mathbf{\Delta}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{U}^* \widehat{\mathbf{V}} \end{bmatrix} \mathbf{\Delta}_1$. Let $\delta = \|\mathbf{\Xi}\|$, so that

$$\|\mathbf{\Delta}_2\| \leq \frac{\|\mathbf{\Delta}_1\|}{\sigma_{\min}(\mathbf{Q} + \mathbf{\Xi})} \leq \frac{3 \|\mathbf{\Xi}\|}{\sigma_{\min}(\mathbf{Q} + \mathbf{\Xi})} \leq \frac{3\delta}{1 - \delta}. \quad (9.13)$$

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, $\mathbf{V} = \mathbf{UR}$ for some orthogonal matrix \mathbf{R} , so that spans the same subspace as that of \mathbf{U} . Next, we control the spectral norm of $\mathbf{\Delta}_3$ so that it is sufficiently small,

$$\|\mathbf{\Delta}_3\| = \min_{\mathbf{R} \in \mathcal{O}_\ell} \left\| \mathbf{UR} - \widehat{\mathbf{V}} \right\| \leq \min_{\mathbf{R} \in \mathcal{O}_\ell} \left\| \mathbf{UR} - \mathbf{Q}_{[n-\ell]} \right\| + \left\| \mathbf{Q}_{[n-\ell]} - \widehat{\mathbf{V}} \right\|, \quad (9.14)$$

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 [SS90]). To bound the first term, from Lemma B.4, 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 $\mathbf{Q}[e_1, \dots, e_\ell]$. Write $\mathbf{I}_{[\ell]} \doteq [e_1, \dots, e_\ell]$ for short, we bound $\sin \theta_1$ as

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

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

where in the first line we have used the fact that for any full column rank matrix \mathbf{M} , $\mathbf{M}(\mathbf{M}^*\mathbf{M})^{-1}\mathbf{M}^*$ is the orthogonal projection onto the its column span, and to obtain the fifth and six lines we have invoked the matrix inverse perturbation bound again. Use the facts that $\delta < 1/20$ 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.4,

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

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

$$\|\mathbf{\Delta}_3\| \leq \delta/(1 - \delta) + 8\sqrt{2}\delta \leq 12.5\delta. \quad (9.15)$$

Therefore, by (9.12), we obtain

$$\mathbf{U}^*(\mathbf{Q} + \mathbf{\Xi})\mathbf{X}_0 = \mathbf{U}^*\mathbf{V}\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}, \quad \text{with} \quad \mathbf{\Delta} \doteq \mathbf{U}^*\mathbf{\Delta}_3\mathbf{X}_0^{[n-\ell]} + \mathbf{\Delta}_2\mathbf{X}_0. \quad (9.16)$$

By using the results in (9.13) and (9.15), we get the desired result. \blacksquare

Appendices

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

$$\dot{h}_\mu(z) = \tanh\left(\frac{z}{\mu}\right), \quad \ddot{h}_\mu(z) = \frac{1}{\mu} \left[1 - \tanh^2\left(\frac{z}{\mu}\right)\right]. \quad (\text{A.1})$$

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

$$\frac{1}{2} \left(1 - \exp\left(-\frac{2z}{\mu}\right)\right) \leq \tanh\left(\frac{z}{\mu}\right) \leq 1 - \exp\left(-\frac{2z}{\mu}\right), \quad (\text{A.2})$$

$$\exp\left(-\frac{2z}{\mu}\right) \leq 1 - \tanh^2\left(\frac{z}{\mu}\right) \leq 4 \exp\left(-\frac{2z}{\mu}\right). \quad (\text{A.3})$$

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

$$\left|\dot{h}_\mu(z) - \dot{h}_\mu(z')\right| \leq \frac{1}{\mu} |z - z'|, \quad \left|\ddot{h}_\mu(z) - \ddot{h}_\mu(z')\right| \leq \frac{2}{\mu^2} |z - z'| \quad (\text{A.4})$$

Lemma A.2 (Chebyshev's Association Inequality) *Let X denote a real-valued random variable, and $f, g : \mathbb{R} \mapsto \mathbb{R}$ nondecreasing (nonincreasing) functions of X with $\mathbb{E}[f(X)] < \infty$ and $\mathbb{E}[g(X)] < \infty$. Then*

$$\mathbb{E}[f(X)g(X)] \geq \mathbb{E}[f(X)]\mathbb{E}[g(X)]. \quad (\text{A.5})$$

If f is nondecreasing (nonincreasing) and g is nonincreasing (nondecreasing), we have

$$\mathbb{E}[f(X)g(X)] \leq \mathbb{E}[f(X)]\mathbb{E}[g(X)]. \quad (\text{A.6})$$

Proof Consider Y , an independent copy of X . Then it is easy to see

$$\mathbb{E}[(f(X) - f(Y))(g(X) - g(Y))] \geq 0.$$

Expanding the expectation and noticing $\mathbb{E}[f(X)g(Y)] = \mathbb{E}[f(Y)g(X)] = \mathbb{E}[f(X)]\mathbb{E}[g(X)]$ and also $\mathbb{E}[f(X)g(X)] = \mathbb{E}[f(Y)g(Y)]$ yields the result. Similarly, we can prove the second one. ■

This lemma implies the following lemma.

Lemma A.3 (Harris' Inequality, [Har60], see also Theorem 2.15 of [BLM13]) *Let X_1, \dots, X_n be independent, real-valued random variables and $f, g : \mathbb{R}^n \mapsto \mathbb{R}$ be nonincreasing (nondecreasing) w.r.t. any one variable while fixing the others. Define a random vector $\mathbf{X} = (X_1, \dots, X_n) \in \mathbb{R}^n$, then we have*

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] \geq \mathbb{E}[f(\mathbf{X})]\mathbb{E}[g(\mathbf{X})]. \quad (\text{A.7})$$

Similarly, if f is nondecreasing (nonincreasing) and g is nonincreasing (nondecreasing) coordinatewise in the above sense, we have

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] \leq \mathbb{E}[f(\mathbf{X})]\mathbb{E}[g(\mathbf{X})]. \quad (\text{A.8})$$

Proof Again, it suffices to prove the first equality, which can be shown by induction. For $n = 1$, it reduces to Lemma A.2. Suppose the claim is true for any $m < n$. Since both g and f are nondecreasing functions in X_n given $\widehat{\mathbf{X}} = (X_1, \dots, X_{n-1})$, then

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] = \mathbb{E}\left[\mathbb{E}\left[f(\mathbf{X})g(\mathbf{X}) \mid \widehat{\mathbf{X}}\right]\right] \geq \mathbb{E}\left[\mathbb{E}\left[f(\mathbf{X}) \mid \widehat{\mathbf{X}}\right]\mathbb{E}\left[g(\mathbf{X}) \mid \widehat{\mathbf{X}}\right]\right]$$

Now, it follows by independence that $f'(\widehat{\mathbf{X}}) = \mathbb{E}[f(\mathbf{X}) | \widehat{\mathbf{X}}]$ and $g'(\widehat{\mathbf{X}}) = \mathbb{E}[g(\mathbf{X}) | \widehat{\mathbf{X}}]$ are both nondecreasing functions, then by the induction hypothesis, we have

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] \geq \mathbb{E}[f'(\widehat{\mathbf{X}})] \mathbb{E}[g'(\widehat{\mathbf{X}})] = \mathbb{E}[f(\mathbf{X})] \mathbb{E}[g(\mathbf{X})],$$

as desired. \blacksquare

Lemma A.4 (Differentiation under the Integral Sign) Consider a function $F : \mathbb{R}^n \times \mathbb{R} \mapsto \mathbb{R}$ such that $\frac{\partial F(\mathbf{x}, s)}{\partial s}$ is well defined and measurable over $\mathcal{U} \times (0, t_0)$ for some open subset $\mathcal{U} \subset \mathbb{R}^n$ and some $t_0 > 0$. For any probability measure μ on \mathbb{R}^n and any $t \in (0, t_0)$ such that $\int_0^t \int_{\mathcal{U}} \left| \frac{\partial F(\mathbf{x}, s)}{\partial s} \right| \mu(d\mathbf{x}) ds < \infty$, it holds that

$$\frac{d}{dt} \int_{\mathcal{U}} F(\mathbf{x}, t) \mu(d\mathbf{x}) = \int_{\mathcal{U}} \frac{\partial F(\mathbf{x}, t)}{\partial t} \mu(d\mathbf{x}), \text{ or } \frac{d}{dt} \mathbb{E}_{\mathbf{x}} [F(\mathbf{x}, t) \mathbb{1}_{\mathcal{U}}] = \mathbb{E}_{\mathbf{x}} \left[\frac{\partial F(\mathbf{x}, t)}{\partial t} \mathbb{1}_{\mathcal{U}} \right]. \quad (\text{A.9})$$

Proof We have

$$\begin{aligned} \int_{\mathcal{U}} \frac{\partial F(\mathbf{x}, t)}{\partial t} \mu(d\mathbf{x}) &= \frac{d}{dt} \int_0^t \int_{\mathcal{U}} \frac{\partial F(\mathbf{x}, s)}{\partial s} \mu(d\mathbf{x}) ds \\ &= \frac{d}{dt} \int_{\mathcal{U}} \int_0^t \frac{\partial F(\mathbf{x}, s)}{\partial s} ds \mu(d\mathbf{x}) \\ &= \frac{d}{dt} \int_{\mathcal{U}} (F(\mathbf{x}, t) - F(\mathbf{x}, 0)) \mu(d\mathbf{x}) \\ &= \frac{d}{dt} \int_{\mathcal{U}} F(\mathbf{x}, t) \mu(d\mathbf{x}), \end{aligned}$$

where we have used the fundamental theorem of calculus for the first and third equalities, and measure-theoretic Fubini's theorem (see, e.g., Theorem 2.37 of [Fol99]) for the second equality (as justified by our integrability assumption). \blacksquare

Lemma A.5 (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 $\Phi^c(x) \doteq 1 - \Phi(x)$:

$$\left(\frac{1}{x} - \frac{1}{x^3} \right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}} \leq \Phi^c(x) \leq \left(\frac{1}{x} - \frac{1}{x^3} + \frac{3}{x^5} \right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}}, \quad (\text{Type I}) \quad (\text{A.10})$$

$$\frac{x}{x^2 + 1} \frac{\exp(-x^2/2)}{\sqrt{2\pi}} \leq \Phi^c(x) \leq \frac{1}{x} \frac{\exp(-x^2/2)}{\sqrt{2\pi}}, \quad (\text{Type II}) \quad (\text{A.11})$$

$$\frac{\sqrt{x^2 + 4} - x}{2} \frac{\exp(-x^2/2)}{\sqrt{2\pi}} \leq \Phi^c(x) \leq \left(\sqrt{2 + x^2} - x \right) \frac{\exp(-x^2/2)}{\sqrt{2\pi}} \quad (\text{Type III}). \quad (\text{A.12})$$

Proof Type I bounds can be obtained by integration by parts with proper truncations. Type II upper bound can again be obtained via integration by parts, and the lower bound can be obtained via considering the function $f(x) \doteq \Phi^c(x) - \frac{x}{x^2+1} \frac{\exp(-x^2/2)}{\sqrt{2\pi}}$ and noticing it is always nonnegative. Type III bounds are mentioned in [Due10] and reproduced by the systematic approach developed therein (section 2). \blacksquare

Lemma A.6 (Moments of the Gaussian Random Variables) If $X \sim \mathcal{N}(0, \sigma^2)$, then it holds for all integer $p \geq 1$ that

$$\mathbb{E}[|X|^p] = \sigma^p (p-1)!! \left[\sqrt{\frac{2}{\pi}} \mathbb{1}_{p \text{ odd}} + \mathbb{1}_{p \text{ even}} \right] \leq \sigma^p (p-1)!!. \quad (\text{A.13})$$

Lemma A.7 (Moments of the χ^2 Random Variables) If $X \sim \chi^2(n)$, then it holds for all integer $p \geq 1$,

$$\mathbb{E}[X^p] = 2^p \frac{\Gamma(p + n/2)}{\Gamma(n/2)} = \prod_{k=1}^p (n + 2k - 2) \leq \frac{p!}{2} (2n)^p. \quad (\text{A.14})$$

Lemma A.8 (Moments of the χ Random Variables) If $X \sim \chi(n)$, then it holds for all integer $p \geq 1$,

$$\mathbb{E}[X^p] = 2^{p/2} \frac{\Gamma(p/2 + n/2)}{\Gamma(n/2)} \leq p! n^{p/2}. \quad (\text{A.15})$$

Lemma A.9 (Moment-Control Bernstein's Inequality for Scalar RVs, Theorem 2.10 of [FR13]) Let X_1, \dots, X_p be i.i.d. real-valued random variables. Suppose that there exist some positive number R and σ^2 such that

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

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

$$\mathbb{P}[|S - \mathbb{E}[S]| \geq t] \leq 2 \exp\left(-\frac{pt^2}{2\sigma^2 + 2Rt}\right). \quad (\text{A.16})$$

Lemma A.10 (Moment-Control Bernstein's Inequality for Matrix RVs, Theorem 6.2 of [Tro12]) 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

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

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

$$\mathbb{P}[\|\mathbf{S} - \mathbb{E}[\mathbf{S}]\| \geq t] \leq 2d \exp\left(-\frac{pt^2}{2\sigma^2 + 2Rt}\right). \quad (\text{A.17})$$

Proving this lemma requires some modification to the original proof of Theorem 6.2 in [Tro12]. We record it here for the sake of completeness.

Proof Let us define $\mathbf{S}_p = \sum_{k=1}^p \mathbf{X}_k$, by Proposition 3.1 of [Tro12], we have

$$\mathbb{P}[\lambda_{\max}(\mathbf{S}_p - \mathbb{E}[\mathbf{S}_p]) \geq t] \leq \inf_{t>0} e^{-\theta t} \mathbb{E}[\text{tr exp}(\theta \mathbf{S}_p - \theta \mathbb{E}[\mathbf{S}_p])], \quad (\text{A.18})$$

To proceed, notice that

$$\mathbb{E}[\text{tr exp}(\theta \mathbf{S}_p - \theta \mathbb{E}[\mathbf{S}_p])]$$

$$\begin{aligned}
&= \mathbb{E}_{\mathbf{S}_{p-1}} \mathbb{E}_{\mathbf{X}_p} [\text{tr exp}(\theta(\mathbf{S}_{p-1} - \mathbb{E}[\mathbf{S}_{p-1}]) + \theta\mathbf{X}_p - \theta\mathbb{E}[\mathbf{X}_p])] \\
&\leq \mathbb{E}_{\mathbf{S}_{p-1}} \left[\text{tr exp} \left(\theta(\mathbf{S}_{p-1} - \mathbb{E}[\mathbf{S}_{p-1}]) + \log \left(\mathbb{E} \left[e^{\theta\mathbf{X}_p} \right] \right) - \theta\mathbb{E}[\mathbf{X}_p] \right) \right] \\
&\leq \mathbb{E}_{\mathbf{S}_{p-1}} \left[\text{tr exp} \left(\theta(\mathbf{S}_{p-1} - \mathbb{E}[\mathbf{S}_{p-1}]) + \mathbb{E} \left[e^{\theta\mathbf{X}_p} \right] - \mathbf{I} - \theta\mathbb{E}[\mathbf{X}_p] \right) \right] \\
&= \mathbb{E}_{\mathbf{S}_{p-1}} \left[\text{tr exp} \left(\theta(\mathbf{S}_{p-1} - \mathbb{E}[\mathbf{S}_{p-1}]) + \sum_{\ell=2}^{\infty} \frac{\theta^\ell \mathbb{E}[\mathbf{X}_k^\ell]}{\ell!} \right) \right]
\end{aligned}$$

where at the third line we have used the result of Corollary 3.3 of [Tro12], i.e., $\mathbb{E}[\text{tr exp}(\mathbf{H} + \mathbf{X})] \leq \text{tr exp}(\mathbf{H} + \log(\mathbb{E}[e^{\mathbf{X}}]))$ for any fixed \mathbf{H} and random, symmetric \mathbf{X} , at the fourth we have used the fact that $\log \mathbf{X} \preceq \mathbf{X} - \mathbf{I}$ for any $\mathbf{X} \succ \mathbf{0}$ (as $\log u \leq u - 1$ for any $u > 0$ and transfer rule applies here), and the last line relies on exchange of infinite summation and expectation, justified as \mathbf{X}_p has a bounded spectral radius. By repeating the argument backwards for $\mathbf{X}_{p-1}, \dots, \mathbf{X}_1$, we get

$$\begin{aligned}
&\mathbb{E}[\text{tr exp}(\theta\mathbf{S}_p - \theta\mathbb{E}[\mathbf{S}_p])] \\
&\leq \text{tr exp} \left(p \sum_{\ell=2}^{\infty} \frac{\theta^\ell \mathbb{E}[\mathbf{X}_k^\ell]}{\ell!} \right) \leq \text{tr exp} \left(p \sum_{\ell=2}^p \frac{\theta^\ell \sigma^2 R^{\ell-2}}{2} \mathbf{I} \right) \\
&\leq d \left\| \exp \left(p \sum_{\ell=2}^p \frac{\theta^\ell \sigma^2 R^{\ell-2}}{2} \mathbf{I} \right) \right\| \leq d \exp \left(\frac{p\theta^2 \sigma^2}{2(1-\theta R)} \right), \tag{A.19}
\end{aligned}$$

where we used the fact that $\mathbb{E}[\mathbf{X}_i^m] \preceq \frac{m!}{2} \sigma^2 R^{m-2} \mathbf{I}$ in (A.17) and restrict $\theta < \frac{1}{R}$. Combining the results in (A.18) and (A.19), we have

$$\mathbb{P}[\lambda_{\max}(\mathbf{S}_p - \mathbb{E}[\mathbf{S}_p]) \geq t] \leq d \inf_{\theta < 1/R} \exp \left(\frac{p\theta^2 \sigma^2}{2(1-\theta R)} - \theta t \right) \tag{A.20}$$

by taking $\theta = t/(p\sigma^2 + Rt) < 1/R$, we obtain

$$\mathbb{P}[\lambda_{\max}(\mathbf{S}_p - \mathbb{E}[\mathbf{S}_p]) \geq t] \leq d \exp \left(-\frac{t^2}{2p\sigma^2 + 2Rt} \right). \tag{A.21}$$

Considering $\mathbf{X}'_k = -\mathbf{X}_k$ and repeating the above argument, we can similarly obtain

$$\mathbb{P}[\lambda_{\min}(\mathbf{S}_p - \mathbb{E}[\mathbf{S}_p]) \leq -t] \leq d \exp \left(-\frac{t^2}{2p\sigma^2 + 2Rt} \right). \tag{A.22}$$

Putting the above bounds together, we have

$$\mathbb{P}[\|\mathbf{S}_p - \mathbb{E}[\mathbf{S}_p]\| \geq t] \leq 2d \exp \left(-\frac{t^2}{2p\sigma^2 + 2Rt} \right). \tag{A.23}$$

We obtain the claimed bound by substituting $\mathbf{S}_p = p\mathbf{S}$ and simplifying the resulting expressions. ■

Corollary A.11 (Moment-Control Bernstein's Inequality for Vector RVs) *Let $\mathbf{x}_1, \dots, \mathbf{x}_p \in \mathbb{R}^d$ be i.i.d. random vectors. Suppose there exist some positive number R and σ^2 such that*

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

Let $\mathbf{s} = \frac{1}{p} \sum_{k=1}^p \mathbf{x}_k$, then for any $t > 0$, it holds that

$$\mathbb{P} [\|\mathbf{s} - \mathbb{E}[\mathbf{s}]\| \geq t] \leq 2(d+1) \exp\left(-\frac{pt^2}{2\sigma^2 + 2Rt}\right). \quad (\text{A.24})$$

Proof To obtain the result, we apply the matrix Bernstein inequality in Lemma A.10 to a suitable embedding of the random vectors $\{\mathbf{x}_k\}_{k=1}^p$. For any $k \in [p]$, define the symmetric matrix

$$\mathbf{X}_k = \begin{bmatrix} 0 & \mathbf{x}_k^* \\ \mathbf{x}_k & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}.$$

Then it holds that

$$\mathbf{X}_k^{2\ell+1} = \|\mathbf{x}_k\|_2^{2\ell} \begin{bmatrix} 0 & \mathbf{x}_k^* \\ \mathbf{x}_k & \mathbf{0} \end{bmatrix}, \quad \mathbf{X}_k^{2\ell+2} = \|\mathbf{x}_k\|_2^{2\ell} \begin{bmatrix} \|\mathbf{x}_k\|_2^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{x}_k \mathbf{x}_k^* \end{bmatrix}, \quad \text{for all integers } \ell \geq 0.$$

Using the fact that

$$\mathbf{x}_k \mathbf{x}_k^* \preceq \|\mathbf{x}_k\|^2 \mathbf{I}, \quad \|\mathbf{X}_k\| = \sqrt{\|\mathbf{X}_k^2\|} = \|\mathbf{x}_k\| \implies -\|\mathbf{x}_k\| \mathbf{I} \preceq \mathbf{X}_k \preceq \|\mathbf{x}_k\| \mathbf{I},$$

and combining the above expressions for $\mathbf{X}_k^{2\ell+1}$ and $\mathbf{X}_k^{2\ell+2}$, we obtain

$$\mathbb{E}[\mathbf{X}_k^m], -\mathbb{E}[\mathbf{X}_k^m] \preceq \mathbb{E}[\|\mathbf{x}_k\|_2^m] \mathbf{I} \preceq \frac{m!}{2} \sigma^2 R^{m-2} \mathbf{I}, \quad \text{for all integers } m \geq 2, \quad (\text{A.25})$$

Let $\mathbf{S} = \frac{1}{p} \sum_{k=1}^p \mathbf{X}_k$, noting that

$$\|\mathbf{S} - \mathbb{E}[\mathbf{S}]\| = \|\mathbf{s} - \mathbb{E}[\mathbf{s}]\|, \quad (\text{A.26})$$

and applying Lemma A.10, we complete the proof. \blacksquare

Lemma A.12 (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

$$f(\mathbf{x} + t\mathbf{y}) = f(\mathbf{x}) + t \int_0^1 \langle \nabla f(\mathbf{x} + s\mathbf{y}), \mathbf{y} \rangle ds, \quad (\text{A.27})$$

$$f(\mathbf{x} + t\mathbf{y}) = f(\mathbf{x}) + t \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle + t^2 \int_0^1 (1-s) \langle \nabla^2 f(\mathbf{x} + s\mathbf{y}) \mathbf{y}, \mathbf{y} \rangle ds. \quad (\text{A.28})$$

Proof By the fundamental theorem of calculus, since f is continuous differentiable, it is obvious that

$$f(\mathbf{x} + t\mathbf{y}) = f(\mathbf{x}) + \int_0^t \langle \nabla f(\mathbf{x} + \tau\mathbf{y}), \mathbf{y} \rangle d\tau. \quad (\text{A.29})$$

If f is twice continuously differentiable, by using integral by parts, we obtain

$$\begin{aligned} f(\mathbf{x} + t\mathbf{y}) &= f(\mathbf{x}) + [(\tau - t) \langle \nabla f(\mathbf{x} + \tau\mathbf{y}), \mathbf{y} \rangle]_0^t - \int_0^t (\tau - t) d \langle \nabla f(\mathbf{x} + \tau\mathbf{y}), \mathbf{y} \rangle \\ &= f(\mathbf{x}) + t \langle \nabla f(\mathbf{x} + \tau\mathbf{y}), \mathbf{y} \rangle + \int_0^t (t - \tau) \langle \nabla^2 f(\mathbf{x} + \tau\mathbf{y}) \mathbf{y}, \mathbf{y} \rangle d\tau. \end{aligned} \quad (\text{A.30})$$

By a change of variable $\tau = st$ ($0 \leq s \leq 1$) for (A.29) and (A.30), we get the desired results. \blacksquare

B Auxillary Results for Proofs

Lemma B.1 Let $X \sim \mathcal{N}(0, \sigma_X^2)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$ be independent random variables and $\Phi^c(t) = \frac{1}{\sqrt{2\pi}} \int_t^\infty \exp(-x^2/2) dx$ be the complementary cumulative distribution function of the standard normal. For any $a > 0$, we have

$$\mathbb{E}[X \mathbb{1}_{X>0}] = \frac{\sigma_X}{\sqrt{2\pi}}, \quad (\text{B.1})$$

$$\mathbb{E}[\exp(-aX) X \mathbb{1}_{X>0}] = \frac{\sigma_X}{\sqrt{2\pi}} - a\sigma_X^2 \exp\left(\frac{a^2\sigma_X^2}{2}\right) \Phi^c(a\sigma_X), \quad (\text{B.2})$$

$$\mathbb{E}[\exp(-aX) \mathbb{1}_{X>0}] = \exp\left(\frac{a^2\sigma_X^2}{2}\right) \Phi^c(a\sigma_X), \quad (\text{B.3})$$

$$\begin{aligned} \mathbb{E}[\exp(-a(X+Y)) X^2 \mathbb{1}_{X+Y>0}] &= \sigma_X^2 (1 + a^2\sigma_X^2) \exp\left(\frac{a^2\sigma_X^2 + a^2\sigma_Y^2}{2}\right) \Phi^c\left(a\sqrt{\sigma_X^2 + \sigma_Y^2}\right) \\ &\quad - \frac{a\sigma_X^4}{\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}}, \end{aligned} \quad (\text{B.4})$$

$$\begin{aligned} \mathbb{E}[\exp(-a(X+Y)) XY \mathbb{1}_{X+Y>0}] &= a^2\sigma_X^2\sigma_Y^2 \exp\left(\frac{a^2\sigma_X^2 + a^2\sigma_Y^2}{2}\right) \Phi^c\left(a\sqrt{\sigma_X^2 + \sigma_Y^2}\right) \\ &\quad - \frac{a\sigma_X^2\sigma_Y^2}{\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}}, \end{aligned} \quad (\text{B.5})$$

$$\mathbb{E}[\tanh(aX) X] = a\sigma_X^2 \mathbb{E}[1 - \tanh^2(aX)], \quad (\text{B.6})$$

$$\mathbb{E}[\tanh(a(X+Y)) X] = a\sigma_X^2 \mathbb{E}[1 - \tanh^2(a(X+Y))]. \quad (\text{B.7})$$

Proof Equalities (B.1), (B.2), (B.3), (B.4) and (B.5) can be obtained by direct integrations. Equalities (B.6) and (B.7) can be derived using integration by part. ■

Proof [of Lemma 7.1] Indeed $\frac{1}{(1+\beta t)^2} = \sum_{k=0}^{\infty} (-1)^k (k+1) \beta^k t^k$, as

$$\sum_{k=0}^{\infty} (-1)^k (k+1) \beta^k t^k = \sum_{k=0}^{\infty} (-\beta t)^k + \sum_{k=0}^{\infty} k (-\beta t)^k = \frac{1}{1+\beta t} + \frac{-\beta t}{(1+\beta t)^2} = \frac{1}{(1+\beta t)^2}.$$

The magnitude of the coefficient vector is

$$\|\mathbf{b}\|_{\ell^1} = \sum_{k=0}^{\infty} \beta^k (1+k) = \sum_{k=0}^{\infty} \beta^k + \sum_{k=0}^{\infty} k \beta^k = \frac{1}{1-\beta} + \frac{\beta}{(1-\beta)^2} = \frac{1}{(1-\beta)^2} = T.$$

Observing that $\frac{1}{(1+\beta t)^2} > \frac{1}{(1+t)^2}$ for $t \in [0, 1]$ when $0 < \beta < 1$, we obtain

$$\|p - f\|_{L^1[0,1]} = \int_0^1 |p(t) - f(t)| dt = \int_0^1 \left[\frac{1}{(1+\beta t)^2} - \frac{1}{(1+t)^2} \right] dt = \frac{1-\beta}{2(1+\beta)} \leq \frac{1}{2\sqrt{T}}. \quad (\text{B.8})$$

Moreover, we have

$$\|f - p\|_{L^\infty[0,1]} = \max_{t \in [0,1]} p(t) - f(t) = \max_{t \in [0,1]} \frac{t(1-\beta)(2+t(1+\beta))}{(1+t)^2(1+\beta t)^2} \leq 1-\beta = \frac{1}{\sqrt{T}}. \quad (\text{B.9})$$

Finally, notice that

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{b_k}{(1+k)^3} &= \sum_{k=0}^{\infty} \frac{(-\beta)^k}{(1+k)^2} = \sum_{i=0}^{\infty} \left[\frac{\beta^{2i}}{(1+2i)^2} - \frac{\beta^{2i+1}}{(2i+2)^2} \right] \\ &= \sum_{i=0}^{\infty} \beta^{2i} \frac{(2i+2)^2 - \beta(2i+1)^2}{(2i+2)^2(2i+1)^2} > 0, \end{aligned} \quad (\text{B.10})$$

where at the second equality we have grouped consecutive even-odd pair of summands. In addition, we have

$$\sum_{k=0}^n \frac{b_k}{(1+k)^3} \leq \sum_{k=0}^n \frac{|b_k|}{(1+k)^3} = \sum_{k=0}^n \frac{\beta^k}{(1+k)^2} \leq 1 + \sum_{k=1}^n \frac{1}{(1+k)k} = 2 - \frac{1}{n+1}, \quad (\text{B.11})$$

which converges to 2 when $n \rightarrow \infty$, completing the proof. \blacksquare

Proof [of Lemma 7.4] The first inequality is obviously true for $\mathbf{v} = \mathbf{0}$. When $\mathbf{v} \neq \mathbf{0}$, we have

$$\begin{aligned} \mathbb{E} [|\mathbf{v}^* \mathbf{z}|^m] &= \sum_{\ell=0}^n \theta^\ell (1-\theta)^{n-\ell} \sum_{\mathcal{J} \in \binom{[n]}{\ell}} \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{v}_{\mathcal{J}}\|^2)} [|Z|^m] \\ &\leq \sum_{\ell=0}^n \theta^\ell (1-\theta)^{n-\ell} \sum_{\mathcal{J} \in \binom{[n]}{\ell}} \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{v}\|^2)} [|Z|^m] \\ &= \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{v}\|^2)} [|Z|^m] \sum_{\ell=0}^n \theta^\ell (1-\theta)^{n-\ell} \binom{n}{\ell} \\ &= \mathbb{E}_{Z \sim \mathcal{N}(0, \|\mathbf{v}\|^2)} [|Z|^m], \end{aligned}$$

where the second line relies on the fact $\|\mathbf{v}_{\mathcal{J}}\| \leq \|\mathbf{v}\|$ and that for a fixed order, central moment of Gaussian is monotonically increasing w.r.t. its variance. Similarly, to see the second inequality,

$$\begin{aligned} \mathbb{E} [|\mathbf{z}|^m] &= \sum_{\ell=0}^n \theta^\ell (1-\theta)^{n-\ell} \sum_{\mathcal{J} \in \binom{[n]}{\ell}} \mathbb{E} [|\mathbf{z}'_{\mathcal{J}}|^m] \\ &\leq \mathbb{E} [|\mathbf{z}'|^m] \sum_{\ell=0}^n \theta^\ell (1-\theta)^{n-\ell} \binom{n}{\ell} = \mathbb{E} [|\mathbf{z}'|^m], \end{aligned}$$

as desired. \blacksquare

Proof [of Lemma 7.11] Consider one component of \mathbf{X} , i.e., $X_{ij} = B_{ij}V_{ij}$ for $i \in [n]$ and $j \in [p]$, where $B_{ij} \sim \text{Ber}(\theta)$ and $V_{ij} \sim \mathcal{N}(0, 1)$. We have

$$\mathbb{P} [|X_{ij}| > 4\sqrt{\log(np)}] \leq \theta \mathbb{P} [|V_{ij}| > 4\sqrt{\log(np)}] \leq \theta \exp(-8 \log(np)) = \theta(np)^{-8}.$$

And also

$$\mathbb{P} [|X_{ij}| < 1] = 1 - \theta + \theta \mathbb{P} [|V_{ij}| < 1] \leq 1 - 0.3\theta.$$

Applying a union bound as

$$\mathbb{P} \left[\|\mathbf{X}\|_\infty \leq 1 \text{ or } \|\mathbf{X}\|_\infty \geq 4\sqrt{\log(np)} \right] \leq (1 - 0.3\theta)^{np} + np\theta (np)^{-8} \leq \exp(-0.3\theta np) + \theta (np)^{-7},$$

we complete the proof. \blacksquare

Lemma B.2 *Suppose $\mathbf{A} \succ \mathbf{0}$. Then for any symmetric perturbation matrix Δ with $\|\Delta\| \leq \frac{\sigma_{\min}(\mathbf{A})}{2}$, it holds that*

$$\left\| (\mathbf{A} + \Delta)^{-1/2} - \mathbf{A}^{-1/2} \right\| \leq \frac{2 \|\mathbf{A}\|^{1/2} \|\Delta\|}{\sigma_{\min}^2(\mathbf{A})}. \quad (\text{B.12})$$

Proof First note that

$$\left\| (\mathbf{A} + \Delta)^{-1/2} - \mathbf{A}^{-1/2} \right\| \leq \frac{\left\| (\mathbf{A} + \Delta)^{-1} - \mathbf{A}^{-1} \right\|}{\sigma_{\min}^{1/2}(\mathbf{A}^{-1})}$$

as by our assumption $\mathbf{A} + \Delta \succ \mathbf{0}$ and the fact (Theorem 6.2 in [Hig08]) that $\|\mathbf{X}^{1/2} - \mathbf{Y}^{1/2}\| \leq \|\mathbf{X} - \mathbf{Y}\| / \left(\sigma_{\min}^{1/2}(\mathbf{X}) + \sigma_{\min}^{1/2}(\mathbf{Y}) \right)$ for any $\mathbf{X}, \mathbf{Y} \succ \mathbf{0}$ applies. Moreover, using the fact

$$\left\| (\mathbf{X} + \Delta)^{-1} - \mathbf{X}^{-1} \right\| \leq \frac{\|\mathbf{X}^{-1}\| \|\mathbf{X}^{-1}\Delta\|}{1 - \|\mathbf{X}^{-1}\Delta\|} \leq \frac{\|\Delta\| \|\mathbf{X}^{-1}\|^2}{1 - \|\mathbf{X}^{-1}\| \|\Delta\|}$$

for nonsingular \mathbf{X} and perturbation Δ with $\|\mathbf{X}^{-1}\| \|\Delta\| < 1$ (see, e.g., Theorem 2.5 of Chapter III in [SS90]), we obtain

$$\frac{1}{\sigma_{\min}^{1/2}(\mathbf{A}^{-1})} \left\| (\mathbf{A} + \Delta)^{-1} - \mathbf{A}^{-1} \right\| \leq \|\mathbf{A}\|^{1/2} \frac{\|\Delta\| \|\mathbf{A}^{-1}\|^2}{1 - \|\mathbf{A}^{-1}\| \|\Delta\|} \leq \frac{2 \|\mathbf{A}\|^{1/2} \|\Delta\|}{\sigma_{\min}^2(\mathbf{A})},$$

where we have used the fact $\|\mathbf{A}^{-1}\| \|\Delta\| \leq 1/2$ to simplify at the last inequality. \blacksquare

Lemma B.3 *There exists a positive constant C such that for any $\theta \in (0, 1/2)$ and $n_2 > Cn_1^2 \log n_1$, the random matrix $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with $\mathbf{X} \sim_{i.i.d.} \text{BG}(\theta)$ obeys*

$$\left\| \frac{1}{n_2\theta} \mathbf{X}\mathbf{X}^* - \mathbf{I} \right\| \leq 10\sqrt{\frac{\theta n_1 \log n_2}{n_2}} \quad (\text{B.13})$$

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

Proof Observe that $\mathbb{E} \left[\frac{1}{\theta} \mathbf{x}_k \mathbf{x}_k^* \right] = \mathbf{I}$ for any column \mathbf{x}_k of \mathbf{X} and so $\frac{1}{n_2\theta} \mathbf{X}\mathbf{X}^*$ can be considered as a normalized sum of independent random matrices. Moreover, for any integer $m \geq 2$,

$$\mathbb{E} \left[\left(\frac{1}{\theta} \mathbf{x}_k \mathbf{x}_k^* \right)^m \right] = \frac{1}{\theta^m} \mathbb{E} \left[\|\mathbf{x}_k\|^{2m-2} \mathbf{x}_k \mathbf{x}_k^* \right].$$

Now $\mathbb{E} \left[\|\mathbf{x}_k\|^{2m-2} \mathbf{x}_k \mathbf{x}_k^* \right]$ is a diagonal matrix (as $\mathbb{E} \left[\|\mathbf{x}_k\|^2 x_k(i) x_k(j) \right] = -\mathbb{E} \left[\|\mathbf{x}_k\|^2 x_k(i) x_k(j) \right]$ for any $i \neq j$ by symmetry of the distribution) in the form $\mathbb{E} \left[\|\mathbf{x}_k\|^{2m-2} \mathbf{x}_k \mathbf{x}_k^* \right] = \mathbb{E} \left[\|\mathbf{x}\|^{2m-2} x(1)^2 \right] \mathbf{I}$ for $\mathbf{x} \sim_{i.i.d.} \text{BG}(\theta)$ with $\mathbf{x} \in \mathbb{R}^{n_1}$. Let $t^2(\mathbf{x}) = \|\mathbf{x}\|^2 - x(1)^2$. Then if $m = 2$,

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{x}\|^2 x(1)^2 \right] &= \mathbb{E} [x(1)^4] + \mathbb{E} [t^2(\mathbf{x})] \mathbb{E} [x(1)^2] \\ &= \mathbb{E} [x(1)^4] + (n_1 - 1) (\mathbb{E} [x(1)^2])^2 = 3\theta + (n_1 - 1) \theta^2 \leq 3n_1\theta, \end{aligned}$$

where for the last simplification we use the assumption $\theta \leq 1/2$. For $m \geq 3$,

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{x}\|^{2m-2} x(1)^2 \right] &= \sum_{k=0}^{m-1} \binom{m-1}{k} \mathbb{E} \left[t^{2k}(\mathbf{x}) x(1)^{2m-2k} \right] = \sum_{k=0}^{m-1} \binom{m-1}{k} \mathbb{E} \left[t^{2k}(\mathbf{x}) \right] \mathbb{E} \left[x(1)^{2m-2k} \right] \\ &\leq \sum_{k=0}^{m-1} \binom{m-1}{k} \mathbb{E}_{Z \sim \chi^2(n_1-1)} [Z^k] \theta \mathbb{E}_{W \sim \mathcal{N}(0,1)} [W^{2m-2k}] \\ &\leq \theta \sum_{k=0}^{m-1} \binom{m-1}{k} \frac{k!}{2} (2n_1 - 2)^k (2m - 2k)!! \\ &\leq \theta 2^m \frac{m!}{2} \sum_{k=0}^{m-1} \binom{m-1}{k} (n_1 - 1)^k \\ &\leq \frac{m!}{2} n_1^{m-1} 2^{m-1}, \end{aligned}$$

where we have used the moment estimates for Gaussian and χ^2 random variables from Lemma A.6 and Lemma A.7, and also $\theta \leq 1/2$. Taking $\sigma^2 = 3n_1\theta$ and $R = 2n_1$, and invoking the matrix Bernstein in Lemma A.10, we obtain

$$\mathbb{E} \left[\left\| \frac{1}{p\theta} \sum_{k=1}^p \mathbf{x}_k \mathbf{x}_k^* - \mathbf{I} \right\| > t \right] \leq \exp \left(-\frac{n_2 t^2}{6n_1\theta + 4n_1 t} + 2 \log n_1 \right) \quad (\text{B.14})$$

for any $t \geq 0$. Taking $t = 10\sqrt{\theta n_1 \log(n_2)/n_2}$ gives the claimed result. \blacksquare

Lemma B.4 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 Proof to i) is similar to that of II. Theorem 4.11 in [SS90]. For $2k \leq n$, w.l.o.g., we can assume \mathbf{U} and \mathbf{V} are the canonical bases for \mathcal{U} and \mathcal{V} , respectively. Then

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

Now by definition

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

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

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

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

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

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

References

- [AAJ⁺13] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, Praneeth Netrapalli, and Rashish Tandon, *Learning sparsely used overcomplete dictionaries via alternating minimization*, arXiv preprint arXiv:1310.7991 (2013).
- [AAN13] Alekh Agarwal, Animashree Anandkumar, and Praneeth Netrapalli, *Exact recovery of sparsely used overcomplete dictionaries*, arXiv preprint arXiv:1309.1952 (2013).
- [ABG07] P-A Absil, Christopher G Baker, and Kyle A Gallivan, *Trust-region methods on riemannian manifolds*, Foundations of Computational Mathematics **7** (2007), no. 3, 303–330.
- [ABGM13] Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma, *Provable bounds for learning some deep representations*, arXiv preprint arXiv:1310.6343 (2013).
- [ABGM14] ———, *More algorithms for provable dictionary learning*, arXiv preprint arXiv:1401.0579 (2014).
- [ABRS10] Hedy Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran, *Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the kurdyka-lojasiewicz inequality*, Mathematics of Operations Research **35** (2010), no. 2, 438–457.
- [AEB06] Michal Aharon, Michael Elad, and Alfred M Bruckstein, *On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them*, Linear algebra and its applications **416** (2006), no. 1, 48–67.
- [AGJ14] Animashree Anandkumar, Rong Ge, and Majid Janzamin, *Guaranteed non-orthogonal tensor decomposition via alternating rank-1 updates*, arXiv preprint arXiv:1402.5180 (2014).
- [AGKM12] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra, *Computing a nonnegative matrix factorization—provably*, Proceedings of the forty-fourth annual ACM symposium on Theory of computing, ACM, 2012, pp. 145–162.

- [AGM13] Sanjeev Arora, Rong Ge, and Ankur Moitra, *New algorithms for learning incoherent and overcomplete dictionaries*, arXiv preprint arXiv:1308.6273 (2013).
- [AGMM15] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra, *Simple, efficient, and neural algorithms for sparse coding*, arXiv preprint arXiv:1503.00778 (2015).
- [AGMS12] Sanjeev Arora, Rong Ge, Ankur Moitra, and Sushant Sachdeva, *Provable ica with unknown gaussian noise, with implications for gaussian mixtures and autoencoders*, Advances in Neural Information Processing Systems, 2012, pp. 2375–2383.
- [ALMT14] Dennis Amelunxen, Martin Lotz, Michael B McCoy, and Joel A Tropp, *Living on the edge: Phase transitions in convex programs with random data*, Information and Inference (2014), iau005.
- [AMS09] Pierre-Antoine Absil, Robert Mahoney, and Rodolphe Sepulchre, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2009.
- [ARR14] Ali Ahmed, Benjamin Recht, and Justin Romberg, *Blind deconvolution using convex programming*, Information Theory, IEEE Transactions on **60** (2014), no. 3, 1711–1732.
- [BC13] Chenglong Bao, Jian-Feng Cai, and Hui Ji, *Fast sparsity-based orthogonal dictionary learning for image restoration*, Computer Vision (ICCV), 2013 IEEE International Conference on, IEEE, 2013, pp. 3384–3391.
- [BJS14] Chenglong Bao, Hui Ji, Yuhui Quan, and Zuowei Shen, *L0 norm based dictionary learning by proximal methods with global convergence*, Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE, 2014, pp. 3858–3865.
- [BJS14] Chenglong Bao, Hui Ji, and Zuowei Shen, *Convergence analysis for iterative data-driven tight frame construction scheme*, Applied and Computational Harmonic Analysis (2014).
- [BKS13] Afonso S Bandeira, Christopher Kennedy, and Amit Singer, *Approximating the little grothendieck problem over the orthogonal and unitary groups*, arXiv preprint arXiv:1308.5207 (2013).
- [BKS14] Boaz Barak, Jonathan A Kelner, and David Steurer, *Dictionary learning and tensor decomposition via the sum-of-squares method*, arXiv preprint arXiv:1407.1543 (2014).
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford University Press, 2013.
- [BQJ14] Chenglong Bao, Yuhui Quan, and Hui Ji, *A convergent incoherent dictionary learning algorithm for sparse coding*, Computer Vision–ECCV 2014, Springer, 2014, pp. 302–316.
- [BR14] Jop Briët and Oded Regev, *Tight hardness of the non-commutative grothendieck problem*, arXiv preprint arXiv:1412.4413 (2014).
- [BST14] Jérôme Bolte, Shoham Sabach, and Marc Teboulle, *Proximal alternating linearized minimization for nonconvex and nonsmooth problems*, Mathematical Programming **146** (2014), no. 1-2, 459–494.
- [BT89] Dimitri P Bertsekas and John N Tsitsiklis, *Parallel and distributed computation: numerical methods*, vol. 23, Prentice hall Englewood Cliffs, NJ, 1989.
- [BV04] Stephen Boyd and Lieven Vandenberghe, *Convex optimization*, Cambridge University Press, New York, NY, USA, 2004.
- [BWY14] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu, *Statistical guarantees for the em algorithm: From population to sample-based analysis*, arXiv preprint arXiv:1408.2156 (2014).
- [Can02] Emmanuel J. Candès, *New ties between computational harmonic analysis and approximation theory*, Approximation Theory X (2002), 87–153.
- [Can14] Emmanuel J. Candès, *Mathematics of sparsity (and few other things)*, Proceedings of the International Congress of Mathematicians, Seoul, South Korea, 2014.

- [CGT00a] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint, *Trust-region methods*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [CGT00b] Andrew R Conn, Nicholas IM Gould, and Ph L Toint, *Trust region methods*, vol. 1, Siam, 2000.
- [CLMW11] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright, *Robust principal component analysis?*, Journal of the ACM (JACM) **58** (2011), no. 3, 11.
- [CLS14] Emmanuel Candès, Xiaodong Li, and Mahdi Soltanolkotabi, *Phase retrieval via wirtinger flow: Theory and algorithms*, arXiv preprint arXiv:1407.1065 (2014).
- [CM14] Sunav Choudhary and Urbashi Mitra, *Identifiability scaling laws in bilinear inverse problems*, arXiv preprint arXiv:1402.2637 (2014).
- [Com94] Pierre Comon, *Independent component analysis, a new concept?*, Signal processing **36** (1994), no. 3, 287–314.
- [CRPW12] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky, *The convex geometry of linear inverse problems*, Foundations of Computational mathematics **12** (2012), no. 6, 805–849.
- [CSV13] Emmanuel J Candès, Thomas Strohmer, and Vladislav Voroninski, *Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming*, Communications on Pure and Applied Mathematics **66** (2013), no. 8, 1241–1274.
- [DeV98] Ronald A. DeVore, *Nonlinear approximation*, Acta numerica **7** (1998), 51–150.
- [DeV09] Ronald A DeVore, *Nonlinear approximation and its applications*, Multiscale, Nonlinear and Adaptive Approximation, Springer, 2009, pp. 169–201.
- [DGM13] David L Donoho, Matan Gavish, and Andrea Montanari, *The phase transition of matrix recovery from gaussian measurements matches the minimax mse of matrix denoising*, Proceedings of the National Academy of Sciences **110** (2013), no. 21, 8405–8410.
- [DH14] Laurent Demanet and Paul Hand, *Scaling law for recovering the sparsest element in a subspace*, Information and Inference **3** (2014), no. 4, 295–309.
- [DT09] David Donoho and Jared Tanner, *Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **367** (2009), no. 1906, 4273–4293.
- [Due10] Lutz Duembgen, *Bounding standard gaussian tail probabilities*, arXiv preprint arXiv:1012.2063 (2010).
- [DVDD98] David L. Donoho, Martin Vetterli, Ronald A. DeVore, and Ingrid Daubechies, *Data compression and harmonic analysis*, Information Theory, IEEE Transactions on **44** (1998), no. 6, 2435–2476.
- [EAS98] Alan Edelman, Tomás A Arias, and Steven T Smith, *The geometry of algorithms with orthogonality constraints*, SIAM journal on Matrix Analysis and Applications **20** (1998), no. 2, 303–353.
- [Ela10] Michael Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*, Springer, 2010.
- [FJK96] Alan Frieze, Mark Jerrum, and Ravi Kannan, *Learning linear transformations*, 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, IEEE Computer Society, 1996, pp. 359–359.
- [Fol99] Gerald B Folland, *Real analysis: Modern techniques and their applications*, 2nd ed., John Wiley & Sons, 1999.

- [FR13] Simon Foucart and Holger Rauhut, *A mathematical introduction to compressive sensing*, Springer, 2013.
- [FW04] Charles Fortin and Henry Wolkowicz, *The trust region subproblem and semidefinite programming**, *Optimization methods and software* **19** (2004), no. 1, 41–67.
- [GHJY15] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan, *Escaping from saddle points—online stochastic gradient for tensor decomposition*, arXiv preprint arXiv:1503.02101 (2015).
- [GJB⁺13] Remi Gribonval, Rodolphe Jenatton, Francis Bach, Martin Kleinsteuber, and Matthias Seibert, *Sample complexity of dictionary learning and other matrix factorizations*, arXiv preprint arXiv:1312.3790 (2013).
- [GJB14] Rémi Gribonval, Rodolphe Jenatton, and Francis Bach, *Sparse and spurious: dictionary learning with noise and outliers*, arXiv preprint arXiv:1407.5155 (2014).
- [GN10] Lee-Ad Gottlieb and Tyler Neylon, *Matrix sparsification and the sparse null space problem*, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, Springer, 2010, pp. 205–218.
- [GS10] Rémi Gribonval and Karin Schnass, *Dictionary identification - sparse matrix-factorization via ℓ^1 -minimization*, *IEEE Transactions on Information Theory* **56** (2010), no. 7, 3523–3539.
- [GW11] Quan Geng and John Wright, *On the local correctness of ℓ^1 -minimization for dictionary learning*, *Submitted to IEEE Transactions on Information Theory* (2011), Preprint: <http://www.columbia.edu/~jw2966>.
- [Har60] Theodore E. Harris, *A lower bound for the critical probability in a certain percolation process*, *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 56, Cambridge Univ Press, 1960, pp. 13–20.
- [Har14] Moritz Hardt, *Understanding alternating minimization for matrix completion*, *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, IEEE, 2014, pp. 651–660.
- [Hig08] Nicholas J. Higham, *Functions of matrices*, Society for Industrial and Applied Mathematics, 2008.
- [HK14] Elad Hazan and Tomer Koren, *A linear-time algorithm for trust region problems*, arXiv preprint arXiv:1401.6757 (2014).
- [HMG94] Uwe Helmke, John B Moore, and Würzburg Germany, *Optimization and dynamical systems*.
- [HO00] Aapo Hyvärinen and Erkki Oja, *Independent component analysis: algorithms and applications*, *Neural networks* **13** (2000), no. 4, 411–430.
- [HO01] Karhunen J. Hyvärinen, A. and E. Oja, *Independent component analysis*, John Wiley and Sons., 2001.
- [HS11] Christopher Hillar and Friedrich T Sommer, *When can dictionary learning uniquely recover sparse data from subsamples?*, arXiv preprint arXiv:1106.3616 (2011).
- [HW14] Moritz Hardt and Mary Wootters, *Fast matrix completion without the condition number*, *Proceedings of The 27th Conference on Learning Theory*, 2014, pp. 638–678.
- [Hyv99] Aapo Hyvarinen, *Fast and robust fixed-point algorithms for independent component analysis*, *IEEE Trans. Neural Networks* **10** (1999), no. 3, 626–634.
- [JN14] Prateek Jain and Praneeth Netrapalli, *Fast exact matrix completion with finite samples*, arXiv preprint arXiv:1411.1087 (2014).
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi, *Low-rank matrix completion using alternating minimization*, *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*, ACM, 2013, pp. 665–674.

- [JO14] Prateek Jain and Sewoong Oh, *Provable tensor factorization with missing data*, Advances in Neural Information Processing Systems, 2014, pp. 1431–1439.
- [KMO10] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh, *Matrix completion from a few entries*, Information Theory, IEEE Transactions on **56** (2010), no. 6, 2980–2998.
- [LGBB05] Sylvain Lesage, Rémi Gribonval, Frédéric Bimbot, and Laurent Benaroya, *Learning unions of orthonormal bases with thresholded singular value decomposition*, Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing, vol. 5, IEEE, 2005, pp. v–293.
- [Loh15] Po-Ling Loh, *Statistical consistency and asymptotic normality for high-dimensional robust m -estimators*, arXiv preprint arXiv:1501.00312 (2015).
- [LSSS14] Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir, *On the computational efficiency of training neural networks*, Advances in Neural Information Processing Systems, 2014, pp. 855–863.
- [LV15] Kyle Luh and Van Vu, *Dictionary learning with few samples and matrix concentration*, arXiv preprint arXiv:1503.08854 (2015).
- [LW11] Po-Ling Loh and Martin J Wainwright, *High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity*, Advances in Neural Information Processing Systems, 2011, pp. 2726–2734.
- [LW13] ———, *Regularized m -estimators with nonconvexity: Statistical and algorithmic theory for local optima*, Advances in Neural Information Processing Systems, 2013, pp. 476–484.
- [LW14] ———, *Support recovery without incoherence: A case for nonconvex regularization*, arXiv preprint arXiv:1412.5632 (2014).
- [LWB13] Kiryung Lee, Yihong Wu, and Yoram Bresler, *Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization*, arXiv preprint arXiv:1312.0525 (2013).
- [MBP14] Julien Mairal, Francis Bach, and Jean Ponce, *Sparse modeling for image and vision processing*, Foundations and Trends in Computer Graphics and Vision **8** (2014), no. 2-3, 85–283.
- [MG13] Nishant Mehta and Alexander G. Gray, *Sparsity-based generalization bounds for predictive sparse coding*, Proceedings of the 30th International Conference on Machine Learning (ICML-13) **28** (2013), no. 1, 36–44.
- [MHWG13] Cun Mu, Bo Huang, John Wright, and Donald Goldfarb, *Square deal: Lower bounds and improved relaxations for tensor recovery*, arXiv preprint arXiv:1307.5870 (2013).
- [MK87] Katta G Murty and Santosh N Kabadi, *Some np -complete problems in quadratic and nonlinear programming*, Mathematical programming **39** (1987), no. 2, 117–129.
- [MP10a] Jianwei Ma and Gerlind Plonka, *A review of curvelets and recent applications*, IEEE Signal Processing Magazine **27** (2010), no. 2, 118–133.
- [MP10b] Andreas Maurer and Massimiliano Pontil, *K -dimensional coding schemes in hilbert spaces*, Information Theory, IEEE Transactions on **56** (2010), no. 11, 5839–5846.
- [MS83] Jorge J Moré and Danny C Sorensen, *Computing a trust region step*, SIAM Journal on Scientific and Statistical Computing **4** (1983), no. 3, 553–572.
- [MT14] Michael B McCoy and Joel A Tropp, *Sharp recovery bounds for convex demixing, with applications*, Foundations of Computational Mathematics **14** (2014), no. 3, 503–567.
- [NJS13] Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi, *Phase retrieval using alternating minimization*, Advances in Neural Information Processing Systems, 2013, pp. 2796–2804.
- [NNS⁺14] Praneeth Netrapalli, UN Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain, *Non-convex robust pca* , Advances in Neural Information Processing Systems, 2014, pp. 1107–1115.

- [NP13] Behnam Neyshabur and Rina Panigrahy, *Sparse matrix factorization*, arXiv preprint arXiv:1311.3315 (2013).
- [NW06] Jorge Nocedal and Stephen Wright, *Numerical optimization*, Springer, 2006.
- [OF96] Bruno A. Olshausen and David J. Field, *Emergence of simple-cell receptive field properties by learning a sparse code for natural images*, *Nature* **381** (1996), no. 6583, 607–609.
- [OF97] ———, *Sparse coding with an overcomplete basis set: A strategy employed by v1?*, *Vision research* **37** (1997), no. 23, 3311–3325.
- [OH10] Samet Oymak and Babak Hassibi, *New null space results and recovery thresholds for matrix rank minimization*, arXiv preprint arXiv:1011.6326 (2010).
- [QSW14] Qing Qu, Ju Sun, and John Wright, *Finding a sparse vector in a subspace: Linear sparsity using alternating directions*, *Advances in Neural Information Processing Systems*, 2014, pp. 3401–3409.
- [RW97] Franz Rendl and Henry Wolkowicz, *A semidefinite framework for trust region subproblems with applications to large scale minimization*, *Mathematical Programming* **77** (1997), no. 1, 273–299.
- [SA14] Hanie Sedghi and Anima Anandkumar, *Provable methods for training neural networks with sparse connectivity*, arXiv preprint arXiv:1412.2693 (2014).
- [Sch14a] Karin Schnass, *Local identification of overcomplete dictionaries*, arXiv preprint arXiv:1401.6354 (2014).
- [Sch14b] ———, *On the identifiability of overcomplete dictionaries via the minimisation principle underlying k -svd*, *Applied and Computational Harmonic Analysis* **37** (2014), no. 3, 464–491.
- [Sch15] ———, *Convergence radius and sample complexity of itkm algorithms for dictionary learning*, arXiv preprint arXiv:1503.07027 (2015).
- [SLLC15] Wei Sun, Junwei Lu, Han Liu, and Guang Cheng, *Provable sparse tensor decomposition*, arXiv preprint arXiv:1502.01425 (2015).
- [SS90] Gilbert W Stewart and Ji-guang Sun, *Matrix perturbation theory*, Academic press, 1990.
- [SWW12] Daniel A Spielman, Huan Wang, and John Wright, *Exact recovery of sparsely-used dictionaries*, *Proceedings of the 25th Annual Conference on Learning Theory*, 2012.
- [Tem03] Vladimir N Temlyakov, *Nonlinear methods of approximation*, *Foundations of Computational Mathematics* **3** (2003), no. 1, 33–107.
- [Tro12] Joel A. Tropp, *User-friendly tail bounds for sums of random matrices*, *Foundations of Computational Mathematics* **12** (2012), no. 4, 389–434.
- [Tse01] Paul Tseng, *Convergence of a block coordinate descent method for nondifferentiable minimization*, *Journal of optimization theory and applications* **109** (2001), no. 3, 475–494.
- [Udr94] Constantin Udriste, *Convex functions and optimization methods on riemannian manifolds*, vol. 297, Springer Science & Business Media, 1994.
- [VMB11] Daniel Vainsencher, Shie Mannor, and Alfred M. Bruckstein, *The sample complexity of dictionary learning*, *Journal of Machine Learning Research* **12** (2011), no. 23, 3259–3281.
- [WGNL14] Zhaoran Wang, Quanquan Gu, Yang Ning, and Han Liu, *High dimensional expectation-maximization algorithm: Statistical optimization and asymptotic normality*, arXiv preprint arXiv:1412.8729 (2014).
- [WLL14] Zhaoran Wang, Huanran Lu, and Han Liu, *Nonconvex statistical optimization: minimax-optimal sparse pca in polynomial time*, arXiv preprint arXiv:1408.5352 (2014).

- [YCS13] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi, *Alternating minimization for mixed linear regression*, arXiv preprint arXiv:1310.3745 (2013).
- [YZ03] Yinyu Ye and Shuzhong Zhang, *New results on quadratic minimization*, SIAM Journal on Optimization **14** (2003), no. 1, 245–267.
- [ZP01] Michael Zibulevsky and Barak Pearlmutter, *Blind source separation by sparse decomposition in a signal dictionary*, Neural computation **13** (2001), no. 4, 863–882.