

Complete Dictionary Recovery over the Sphere I: Overview and the Geometric Picture

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

Abstract

We consider the problem of recovering a complete (i.e., square and invertible) matrix \mathbf{A}_0 , from $\mathbf{Y} \in \mathbb{R}^{n \times p}$ with $\mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$, provided \mathbf{X}_0 is sufficiently sparse. This recovery problem is central to theoretical understanding of dictionary learning, which seeks a sparse representation for a collection of input signals and finds numerous applications in modern signal processing and machine learning. We give the first efficient algorithm that provably recovers \mathbf{A}_0 when \mathbf{X}_0 has $O(n)$ nonzeros per column, under suitable probability model for \mathbf{X}_0 . In contrast, prior results based on efficient algorithms either only guarantee recovery when \mathbf{X}_0 has $O(\sqrt{n})$ zeros per column, or require multiple rounds of SDP relaxation to work when \mathbf{X}_0 has $O(n^{1-\delta})$ nonzeros per column (for any constant $\delta \in (0, 1)$).

Our algorithmic pipeline centers around solving a certain nonconvex optimization problem with a spherical constraint. In this paper, we provide a geometric characterization of the objective landscape. In particular, we show that the problem is highly structured: with high probability, (1) there are no “spurious” local minimizers; and (2) around all saddle points the objective has a negative directional curvature. This distinctive structure makes the problem amenable to efficient optimization algorithms. In a companion paper [3], we design a second-order trust-region algorithm over the sphere that provably converges to a local minimizer from arbitrary initializations, despite the presence of saddle points.

Index Terms

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

I. INTRODUCTION

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 to 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 successful; the numerous theoretical advances in these fields (see, e.g., [4]–[8] for summary of relevant results) provide ever more powerful representations, ranging from the classic Fourier basis 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 [10], [11]. The discovery has spurred a flurry of algorithmic developments

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

Manuscript received xxx; revised xxx.

¹As Donoho et al [9] 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.”

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., [12], [13] for review of 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 [12], [13] and theories [14]–[17].

A. Theoretical and Algorithmic Challenges

In contrast to the above empirical successes, 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 I-E, 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:

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

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 (I.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 minimizers. These global minimizers are generally isolated, likely jeopardizing natural convex relaxation (see similar discussions in, e.g., [18] and [19]).³ This contrasts sharply with problems in sparse recovery and compressed sensing, in which simple convex relaxations are often provably effective [26]–[35]. Is there any hope to obtain global solutions to the DL problem?

B. 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 [36]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.

³Simple convex relaxations normally replace the objective function with a convex surrogate, and the constraint set with its convex hull. When there are multiple isolated global minimizers for the original nonconvex problem, any point in the convex hull of these global minimizers are necessarily feasible for the relaxed version, and such points tend to produce smaller or equal values than that of the original global minimizers by the relaxed objective function, due to convexity. This implies such relaxations are bound to be loose. Semidefinite programming (SDP) lifting may be one useful general strategy to convexify bilinear inverse problems, see, e.g., [20], [21]. However, for problems with general nonlinear constraints, it is unclear whether the lifting always yields tight relaxation; consider, e.g., [21]–[23] and the identification issue in blind deconvolution [24], [25].

⁴See Section I-C for more detailed discussions of this point. [37] also gave motivations and algorithms for learning (union of) orthobases as dictionaries.

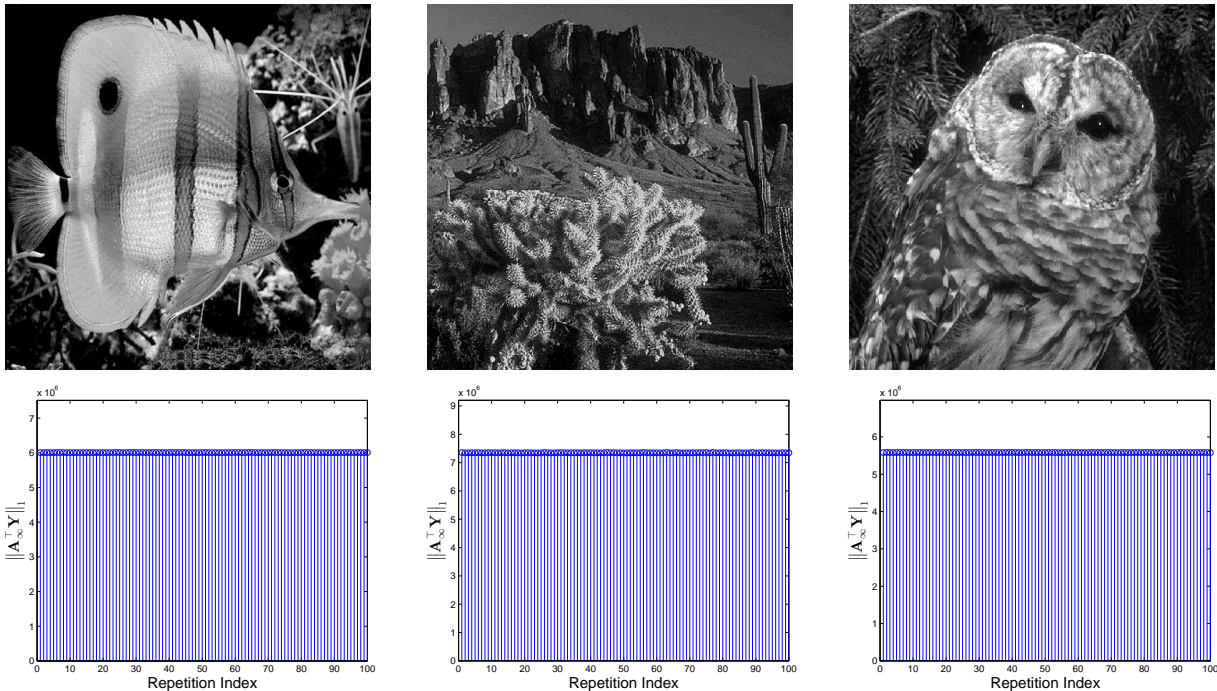


Fig. 1: **Alternating direction method for (I.2) on uncompressed real images seems to always produce the same solution!** **Top:** Each image is 512×512 in resolution and encoded in the uncompressed `pgm` format (uncompressed images to prevent possible bias towards standard bases used for compression, 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 (I.2) 100 times with independent and randomized (uniform over the orthogonal group) initialization \mathbf{A}_0 . Let \mathbf{A}_∞ denote the value of \mathbf{A} at convergence (we set the maximally allowable number of ADM iterations to be 10^4 and $\lambda = 2$). 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 (I.1) to this setting, we obtain the optimization problem:

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

where O_n is the set of order n orthogonal matrices, i.e., order- n orthogonal group. To derive a concrete algorithm for (I.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 .

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

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

⁵This method is also called alternating minimization or (block) coordinate descent method. see, e.g., [38], [39] for classic results and [40], [41] for several interesting recent developments.

This observation is consistent with the possibility that the heuristic ADM algorithm may *always converge to a 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 [42]–[47].

C. 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 [36], and admit faster algorithms for learning and encoding. ⁸
- The coefficient matrix \mathbf{X}_0 follows the Bernoulli-Gaussian (BG) model with rate θ : $[X_0]_{ij} = \Omega_{ij} V_{ij}$, with $\Omega_{ij} \sim \text{Ber}(\theta)$ and $V_{ij} \sim \mathcal{N}(0, 1)$, where all the different random variables are jointly independent. We write compactly $\mathbf{X}_0 \sim_{i.i.d.} \text{BG}(\theta)$. This BG model, or the Bernoulli-Subgaussian model as used in [42], is a reasonable first model for generic sparse coefficients: the Bernoulli process enables explicit control on the (hard) sparsity level, and the (sub)-Gaussian process seems plausible for modeling variations in magnitudes. Real signals may admit encoding coefficients with additional or different characteristics. We will focus on generic sparse encoding coefficients as a first step towards theoretical understanding.

In this paper, we provide a nonconvex formulation for the DR problem, and characterize the geometric structure of the formulation that allows development of efficient algorithms for optimization. In the companion paper [3], we derive an efficient algorithm taking advantage of the structure, and describe a complete algorithmic pipeline for efficient recovery. Together, we prove the following result:

Theorem I.1 (Informal statement of our results, a detailed version included in the companion paper [3]). *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 (up to sign, scale, and permutation) \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 constant $c > 0$.*

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.). Ensuring 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 an 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.,

⁶Technically, the convergence to global solutions is surprising because even convergence of ADM to critical points is not guaranteed in general, see, e.g., [40], [41] and references therein.

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 [11] argues for the necessity from a 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 III for relevant discussion.

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 .

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 I-E provides detailed comparison of our result with other recent recovery results for complete and overcomplete dictionaries.

D. Main Ingredients and Innovations

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

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 $p \geq \Omega(n \log n)$ [42]. Thus, recovering rows of \mathbf{X}_0 is equivalent to finding the sparsest vectors/directions (due to the scale ambiguity) in $\text{row}(\mathbf{Y})$. Since any vector in $\text{row}(\mathbf{Y})$ can be written as $\mathbf{q}^*\mathbf{Y}$ for a certain \mathbf{q} , one might try to solve

$$\text{minimize } \|\mathbf{q}^*\mathbf{Y}\|_0 \quad \text{subject to } \mathbf{q}^*\mathbf{Y} \neq \mathbf{0} \quad (\text{I.3})$$

to find the sparsest vector in $\text{row}(\mathbf{Y})$. Once the sparsest one is found, one then appropriately reduces the subspace $\text{row}(\mathbf{Y})$ by one dimension, and solves an analogous version of (I.3) to find the second sparsest vector. The process is continued recursively until all sparse vectors are obtained. The above idea of reducing the original recovery problem into finding sparsest vectors in a known subspace first appeared in [42].

The objective is discontinuous, and the domain is an open set. In particular, the homogeneous constraint is unconventional 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 [42], [48] fix the scale by setting $\|\mathbf{q}^*\mathbf{Y}\|_\infty = 1$ and use $\|\cdot\|_1$ as a surrogate to $\|\cdot\|_0$, where $\|\cdot\|_\infty$ is the elementwise ℓ^∞ norm, leading to the optimization problem

$$\text{minimize } \|\mathbf{q}^*\mathbf{Y}\|_1 \quad \text{subject to } \|\mathbf{q}^*\mathbf{Y}\|_\infty = 1. \quad (\text{I.4})$$

The constraint means at least one coordinate of $\mathbf{q}^*\mathbf{Y}$ has unit magnitude¹⁰. Thus, (I.4) reduces to a sequence of convex (linear) programs. [42] has shown that (see also [48]) solving (I.4) recovers $(\mathbf{A}_0, \mathbf{X}_0)$ for very sparse \mathbf{X}_0 , but the idea provably breaks down when θ is slightly above $O(1/\sqrt{n})$, or equivalently when each column of \mathbf{X}_0 has more than $O(\sqrt{n})$ nonzeros.

Inspired by our previous image experiment, we work with a *nonconvex* alternative¹¹:

$$\text{minimize } f(\mathbf{q}; \hat{\mathbf{Y}}) \doteq \frac{1}{p} \sum_{k=1}^p h_\mu(\mathbf{q}^*\hat{\mathbf{y}}_k), \quad \text{subject to } \|\mathbf{q}\| = 1, \quad (\text{I.5})$$

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,

$$h_\mu(z) = \mu \log \left(\frac{\exp(z/\mu) + \exp(-z/\mu)}{2} \right) = \mu \log \cosh(z/\mu), \quad (\text{I.6})$$

which is infinitely differentiable and μ controls the smoothing level.¹² An illustration of the $h_\mu(\cdot)$ function vs. the ℓ^1 function is provided in Fig. 2. The spherical constraint is nonconvex. Hence, a-priori, it is unclear whether (I.5) admits efficient algorithms that attain global optima. Surprisingly, simple descent algorithms for (I.5) exhibit very striking behavior: on many practical numerical examples¹³, they appear to produce global solutions. Our next section will uncover interesting geometrical structures underlying the phenomenon.




Fig. 2: The smooth ℓ^1 surrogate defined in (I.6) vs. the ℓ^1 function, for varying values of μ . The surrogate approximates the ℓ^1 function more closely when μ gets smaller.

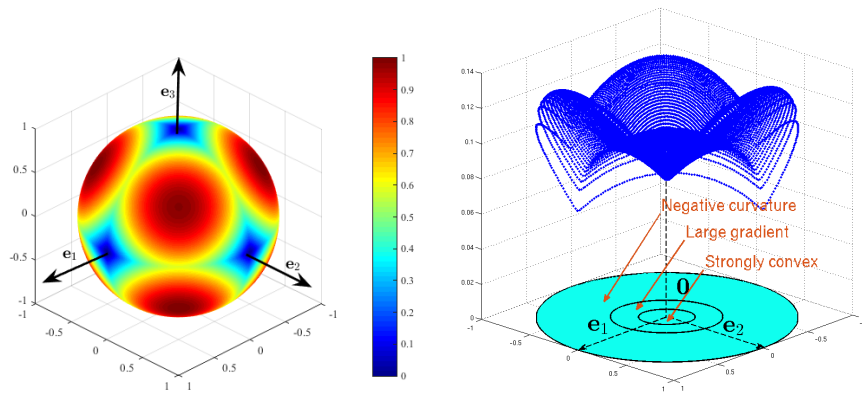


Fig. 3: **Why is dictionary learning over \mathbb{S}^{n-1} tractable?** Assume the target dictionary $\mathbf{A}_0 = \mathbf{I}$. **Left:** Large sample objective function $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q})]$. The only local minimizers are the signed basis vectors $\pm \mathbf{e}_i$. **Right:**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

2) **A Glimpse into High-dimensional Function Landscape:** For the moment, suppose $\mathbf{A}_0 = \mathbf{I}$ and take $\widehat{\mathbf{Y}} = \mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0 = \mathbf{X}_0$ in (I.5). Fig. 3 (left) plots $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q}; \mathbf{X}_0)]$ over $\mathbf{q} \in \mathbb{S}^2$ ($n = 3$). Remarkably, $\mathbb{E}_{\mathbf{X}_0} [f(\mathbf{q}; \mathbf{X}_0)]$ has no spurious local minimizers. In fact, every local minimizer $\widehat{\mathbf{q}}$ is one of the signed standard basis vectors, i.e., $\pm \mathbf{e}_i$'s where $i \in \{1, 2, 3\}$. Hence, $\widehat{\mathbf{q}}^* \mathbf{Y}$ reproduces a certain row of \mathbf{X}_0 , and all minimizers reproduce all rows of \mathbf{X}_0 .

Let \mathbf{e}_3^\perp be the equatorial section that is orthogonal to \mathbf{e}_3 , i.e., $\mathbf{e}_3^\perp \doteq \text{span}(\mathbf{e}_1, \mathbf{e}_2) \cap \mathbb{B}^3$. To better illustrate the above point, we project the upper hemisphere above \mathbf{e}_3^\perp onto \mathbf{e}_3^\perp . The projection is bijective and we equivalently define a reparameterization $g : \mathbf{e}_3^\perp \mapsto \mathbb{R}$ of f . Fig. 3 (right) plots the graph of g . Obviously the only local minimizers are $\mathbf{0}, \pm \mathbf{e}_1, \pm \mathbf{e}_2$, 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 strong gradient

¹⁰The sign ambiguity is tolerable here.

¹¹A similar formulation has been proposed in [49] in the context of blind source separation; see also [50].

¹²In 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.

¹³... not restricted to the model we assume here for \mathbf{A}_0 and \mathbf{X}_0 .

region, and a region where at each point one can always find a direction of negative curvature. 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 a 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 take advantage of this structure without knowing \mathbf{A}_0 ahead of time. In Section I-D3, we describe a Riemannian trust region method which addresses the latter challenge.

a) Geometry for orthogonal \mathbf{A}_0 .: In this case, we take $\hat{\mathbf{Y}} = \mathbf{Y} = \mathbf{A}_0 \mathbf{X}_0$. 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 exemplar. An illustration of the symmetric sections and the exemplar we choose to work with on \mathbb{S}^2 is provided in Fig. 4. 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} .

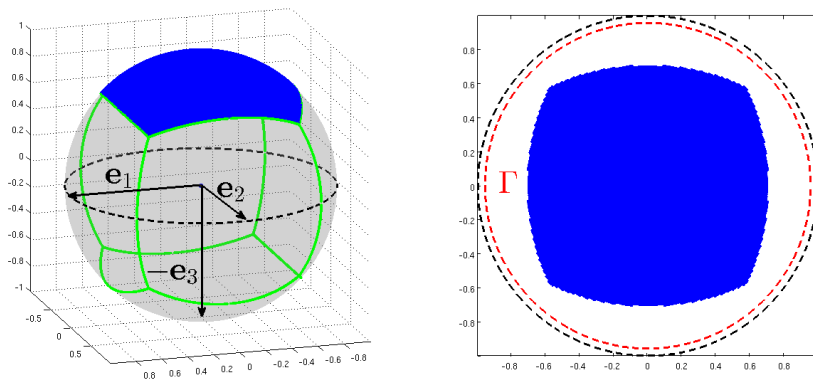


Fig. 4: Illustration of the six symmetric sections on \mathbb{S}^2 and the exemplar we work with. **Left:** The six symmetric sections on \mathbb{S}^2 , as divided by the green curves. The signed basis vectors, $\pm e_i$'s, are centers of these sections. We choose to work with the exemplar that is centered around e_3 that is shaded in blue. **Right:** Projection of the upper hemisphere onto the equatorial section e_3^\perp . The blue region is projection of the exemplar under study. The larger region enclosed by the red circle is the Γ set on which we characterize the reparametrized function g .

To study the function on this exemplar region, we again invoke the projection trick described above, this time onto the equatorial section e_n^\perp . This can be formally captured by the reparameterization mapping:

$$\mathbf{q}(\mathbf{w}) = \left(\mathbf{w}, \sqrt{1 - \|\mathbf{w}\|^2} \right), \quad \mathbf{w} \in \mathbb{B}^{n-1}, \quad (\text{I.7})$$

where \mathbf{w} is the new variable 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

$$\Gamma \doteq \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\} \subsetneq \mathbb{B}^{n-1}. \quad (\text{I.8})$$

It can be verified the exemplar we chose to work with is strictly contained in this set¹⁴. This is illustrated for the case $n = 3$ in Fig. 4 (right).

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

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

¹⁴Indeed, if $\langle \mathbf{q}, \mathbf{e}_n \rangle \geq |\langle \mathbf{q}, \mathbf{e}_i \rangle|$ for all $i \neq n$, $1 - \|\mathbf{w}\|^2 = q_n^2 \geq 1/n$, 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. We study behavior of g over this slightly larger set Γ , instead of just the projection of the chosen symmetric section, to conveniently deal with the boundary effect: if we choose to work with just projection of the chosen symmetric section, there would be considerable technical subtleties at the boundaries when we call the union argument to cover the whole sphere.

respectively over three consecutive regions moving away from the origin, corresponding to the three regions in Fig. 3 (right). In particular, through typical expectation-concentration style arguments, we show that there exists a positive constant c such that

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0) \succeq \frac{1}{\mu} c\theta \mathbf{I}, \quad \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \geq c\theta, \quad \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -c\theta \quad (\text{I.9})$$

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 $O(\mu)$ ¹⁵.

b) 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 a small perturbation. We can then argue that the perturbation does not significantly affect qualitative properties of the objective landscape. Write

$$\bar{\mathbf{Y}} = \left(\frac{1}{p\theta} \mathbf{Y} \mathbf{Y}^* \right)^{-1/2} \mathbf{Y}. \quad (\text{I.10})$$

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

$$\begin{aligned} (\mathbf{A}_0 \mathbf{A}_0^*)^{-1/2} \mathbf{A}_0 \mathbf{X}_0 &= (\mathbf{U} \Sigma \mathbf{V}^* \mathbf{V} \Sigma \mathbf{U}^*)^{-1/2} \mathbf{U} \Sigma \mathbf{V}^* \mathbf{X}_0 \\ &= \mathbf{U} \Sigma^{-1} \mathbf{U}^* \mathbf{U} \Sigma \mathbf{V}^* \mathbf{X}_0 \\ &= \mathbf{U} \mathbf{V}^* \mathbf{X}_0 \end{aligned} \quad (\text{I.11})$$

where $\text{SVD}(\mathbf{A}_0) = \mathbf{U} \Sigma \mathbf{V}^*$. 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

$$\bar{\mathbf{Y}} = \mathbf{U} \mathbf{V}^* \mathbf{X}_0 + \Xi \mathbf{X}_0, \quad (\text{I.12})$$

where Ξ is a matrix with a small magnitude. Simple perturbation argument shows that the constant c in (I.9) 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.

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 minimizers, it does have many saddle points with indefinite Hessian, which we call *ridable saddles*¹⁶ (Fig. 3). We can use second-order information to guarantee to escape from such saddle points. In the companion paper [3], we derive an algorithm based on the Riemannian trust region method (TRM) [53], [54] for this purpose. There are other algorithmic possibilities; see, e.g., [52], [55].

We provide here only the basic intuition why a local minimizer can be retrieved by the second-order trust-region method. Consider an unconstrained optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} \phi(\mathbf{x}).$$

Typical (second-order) TRM proceeds by successively forming a second-order approximation to ϕ at the current iterate,

$$\hat{\phi}(\boldsymbol{\delta}; \mathbf{x}^{(r-1)}) \doteq \phi(\mathbf{x}^{(r-1)}) + \nabla^* \phi(\mathbf{x}^{(r-1)}) \boldsymbol{\delta} + \frac{1}{2} \boldsymbol{\delta}^* \mathbf{Q}(\mathbf{x}^{(r-1)}) \boldsymbol{\delta}, \quad (\text{I.13})$$

¹⁵When $p \rightarrow \infty$, the local minimizer is exactly $\mathbf{0}$; deviation from $\mathbf{0}$ that we described is due to finite-sample perturbation. The deviation distance depends both the $h_\mu(\cdot)$ and p ; see Theorem II.1 for example.

¹⁶See [51] and [52].

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

$$\boldsymbol{\delta}^{(r)} \doteq \arg \min_{\boldsymbol{\delta} \in \mathbb{R}^n, \|\boldsymbol{\delta}\|_p \leq \Delta} \widehat{\phi}(\boldsymbol{\delta}; \mathbf{x}^{(r-1)}), \quad (\text{I.14})$$

where Δ is called the trust-region radius that controls how far the movement can be made. If we take $\mathbf{Q}(\mathbf{x}^{(r-1)}) = \nabla^2\phi(\mathbf{x}^{(r-1)})$ for all r , then whenever the gradient is nonvanishing or the Hessian is indefinite, we expect to decrease the objective function by a concrete amount provided $\|\boldsymbol{\delta}\|$ is sufficiently small. Since the domain is compact, the iterate sequence ultimately moves into the strongly convex region, where the trust-region algorithm behaves like a typical Newton algorithm. All these are generalized to our objective over the sphere and made rigorous in the companion paper [3].

E. 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 [10]. We refer the reader to Chapter 12 - 15 of the book [12] and the survey paper [13] 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.

a) Theoretical Dictionary Learning: The theoretical study of DL in the recovery setting started only very recently. [56] 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 [57]. Subsequent work [18], [19], [58]–[60] studied when the target dictionary is a local optimizer 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 optimizer and hence the dictionary.

[42] 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 (under the Bernoulli-Subgaussian model) with $O(\sqrt{n})$ nonzeros per column (and the method provably breaks down when \mathbf{X}_0 contains slightly more than $\Omega(\sqrt{n})$ nonzeros per column). [43], [45] and [44], [47] gave efficient algorithms that provably recover overcomplete ($m \geq n$), incoherent dictionaries, based on a combination of {clustering or spectral initialization} and local refinement. These algorithms again succeed when \mathbf{X}_0 has $\widetilde{O}(\sqrt{n})$ ¹⁷ nonzeros per column. Recent work [61] provided the first polynomial-time algorithm that provably recovers most “nice” overcomplete dictionaries when \mathbf{X}_0 has $O(n^{1-\delta})$ nonzeros per column for any constant $\delta \in (0, 1)$. However, the proposed algorithm runs in super-polynomial (quasipolynomial) time when the sparsity level goes up to $O(n)$. Similarly, [46] also proposed a super-polynomial time algorithm that guarantees recovery with (almost) $O(n)$ nonzeros per column. Detailed models for those methods dealing with overcomplete dictionaries are differ from one another; nevertheless, they all assume each column of \mathbf{X}_0 has bounded sparsity levels, and the nonzero coefficients have certain sub-Gaussian magnitudes¹⁸. 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, under the BG model.

Aside from efficient recovery, other theoretical work on DL includes results on identifiability [56], [57], [62], generalization bounds [14]–[17], and noise stability [63].

b) Finding Sparse Vectors in a Linear Subspace: We have followed [42] 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 model¹⁹, [48] showed that solving a sequence of linear programs similar to [42] can recover sparse vectors with sparsity up to $O(p/\sqrt{n})$, sublinear in the vector dimension. [50] improved the recovery limit to $O(p)$ by solving a nonconvex

¹⁷The \widetilde{O} suppresses some logarithm factors.

¹⁸Thus, one may anticipate that the performances of those methods do not change much qualitatively, if the BG model for the coefficients had been assumed.

¹⁹... where one sparse vector embedded in an otherwise random subspace.

sphere-constrained problem similar to (I.5)²⁰ via an ADM algorithm. The idea of seeking rows of \mathbf{X}_0 sequentially by solving the above core problem sees precursors in [49] for blind source separation, and [64] for matrix sparsification. [49] also proposed a nonconvex optimization similar to (I.5) here and that employed in [50].

c) Nonconvex Optimization Problems: For other nonconvex optimization problems of recovery of structured signals²¹, including low-rank matrix completion/recovery [73]–[82], phase retrieval [83]–[86], tensor recovery [87]–[90], mixed regression [91], [92], structured element pursuit [50], and recovery of simultaneously structured signals [92], numerical linear algebra and optimization [93], [94], the initialization plus local refinement strategy adopted in theoretical DL [43]–[47] is also crucial: nearness to the target solution enables exploiting the local property of the optimizing objective to ensure that the local refinement succeeds.²² By comparison, we provide a complete characterization of the global geometry, which admits efficient algorithms without any special initialization.

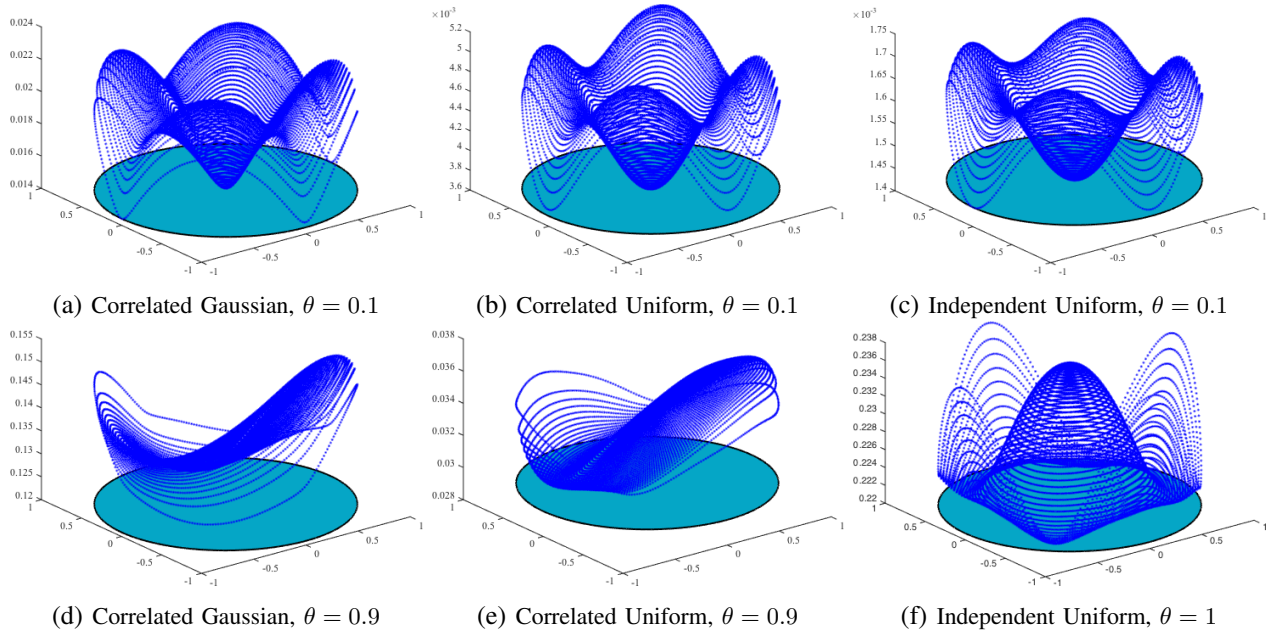


Fig. 5: Asymptotic function landscapes when rows of \mathbf{X}_0 are not independent. W.l.o.g., we again assume $\mathbf{A}_0 = \mathbf{I}$. In (a) and (d), $\mathbf{X}_0 = \mathbf{\Omega} \odot \mathbf{V}$, with $\mathbf{\Omega} \sim_{i.i.d.} \text{Ber}(\theta)$ and columns of \mathbf{X}_0 i.i.d. Gaussian vectors obeying $\mathbf{v}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}^2)$ for symmetric $\mathbf{\Sigma}$ 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), $\mathbf{X}_0 = \mathbf{\Omega} \odot \mathbf{W}$, with $\mathbf{\Omega} \sim_{i.i.d.} \text{Ber}(\theta)$ and columns of \mathbf{X}_0 i.i.d. vectors generated as $\mathbf{w}_i = \mathbf{\Sigma} \mathbf{u}^i$ with $\mathbf{u}_i \sim_{i.i.d.} \text{Uniform}[-0.5, 0.5]$. For comparison, in (c) and (f), $\mathbf{X}_0 = \mathbf{\Omega} \odot \mathbf{W}$ with $\mathbf{\Omega} \sim_{i.i.d.} \text{Ber}(\theta)$ and $\mathbf{W} \sim_{i.i.d.} \text{Uniform}[-0.5, 0.5]$. Here \odot denote the elementwise product, and the objective function is still based on the log cosh function as in (I.5).

d) 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 [98], and learning deep neural nets [99]–[102].

ICA factors a data matrix \mathbf{Y} as $\mathbf{Y} = \mathbf{A}\mathbf{X}$ such that \mathbf{A} is square and rows of \mathbf{X} achieve maximal statistical independence [103], [104]. In theoretical study of the recovery problem, it is often assumed that rows of \mathbf{X}_0 are (weakly) independent (see, e.g., [105]–[107]). Our i.i.d. probability model on \mathbf{X}_0 implies rows of \mathbf{X}_0 are independent, aligning our problem perfectly with the ICA problem. More interestingly, the log cosh objective we

²⁰The only difference is that they chose to work with the Huber function as a proxy of the $\|\cdot\|_1$ function.

²¹This is a body of recent work studying nonconvex recovery up to statistical precision, including, e.g., [65]–[72].

²²The powerful framework [40], [41] to establish local convergence of ADM algorithms to critical points applies to DL/DR also, see, e.g., [95]–[97]. However, these results do not guarantee to produce global optima.

analyze here was proposed as a general-purpose *contrast function* in ICA that has not been thoroughly analyzed [108]. Algorithms and analysis with another popular contrast function, the fourth-order cumulants, however, indeed overlap with ours considerably [106], [107]²³. While this interesting connection potentially helps port our analysis to ICA, it is a fundamental question to ask what is playing a more vital role for DR, sparsity or independence.

Fig. 5 helps shed some light in this direction, where we again plot the asymptotic objective landscape with the natural reparameterization as in Section I-D2. From the left and central panels, it is evident that even without independence, \mathbf{X}_0 with sparse columns induces the familiar geometric structures we saw in Fig. 3; 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 panel²⁴, it seems that 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.²⁵ This suggests that sparsity, instead of independence, makes our current algorithm for DR work.

e) Nonconvex Problems with Similar Geometric Structure: Besides ICA discussed above, it turns out that a handful of other practical problems arising in signal processing and machine learning induce the “no spurious minimizers, all saddles are second-order” structure under natural setting, including the eigenvalue problem, generalized phase retrieval [109], orthogonal tensor decomposition [52], low-rank matrix recovery/completion [110], [111], noisy phase synchronization and community detection [112]–[114], linear neural nets learning [115]–[117]. [51] gave a review of these problems, and discussed how the methodology developed in this and the companion paper [3] can be generalized to solve those problems.

F. Notations, and Reproducible Research

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

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

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

²³Nevertheless, the objective functions are apparently different. Moreover, we have provided a complete geometric characterization of the objective, in contrast to [106], [107]. 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.

²⁴We 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.

²⁵For solving the ICA problem, this suggests the log cosh contrast function, that works well empirically [108], may not work for all distributions (rotation-invariant Gaussian excluded of course), at least when one does not process the data (say perform certain whitening or scaling).

II. THE HIGH-DIMENSIONAL FUNCTION LANDSCAPE

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

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

induced by the reparametrization

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

In particular, we focus our attention to the smaller set

$$\Gamma = \left\{ \mathbf{w} : \|\mathbf{w}\| < \sqrt{\frac{4n-1}{4n}} \right\} \subsetneq \mathbb{B}^{n-1}, \quad (\text{II.3})$$

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 similarly characterize other parts of f on \mathbb{S}^{n-1} using projection onto other equatorial sections. Note that over Γ , $q_n = \sqrt{1 - \|\mathbf{w}\|^2} \geq 1/(2\sqrt{n})$.

A. Main Geometric Theorems

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

$$p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta}, \quad (\text{II.4})$$

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

$$\nabla^2 g(\mathbf{w}; \mathbf{X}_0) \succeq \frac{c_\star \theta}{\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (\text{II.5})$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \geq c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (\text{II.6})$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (\text{II.7})$$

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

$$\|\mathbf{w}_\star - \mathbf{0}\| \leq \min \left\{ \frac{c_c \mu}{\theta} \sqrt{\frac{n \log p}{p}}, \frac{\mu}{16} \right\}. \quad (\text{II.8})$$

Here c_a through c_c are all positive 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 equatorial section \mathbf{e}_n^\perp to \mathbf{e}_i^\perp . Doing this $2n$ times (and multiplying the failure probability in Theorem II.1 by $2n$), we obtain a characterization of $f(\mathbf{q}; \mathbf{X}_0)$ over the entirety of \mathbb{S}^{n-1} .²⁶ The result is captured by the next corollary.

Corollary II.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 $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, whenever $p \geq \frac{C}{\mu^2 \theta^2} n^3 \log \frac{n}{\mu \theta}$, with probability at least $1 - c_b p^{-5}$, 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*

²⁶In fact, it is possible to pull the very detailed geometry captured in (II.5) through (II.7) 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.

between these minimizers and signed basis vectors $\{\pm e_i\}_i$, such that the corresponding local minimizer \mathbf{q}_\star and $\mathbf{b} \in \{\pm e_i\}_i$ satisfy

$$\|\mathbf{q}_\star - \mathbf{b}\| \leq \sqrt{2} \min \left\{ \frac{c_c \mu}{\theta} \sqrt{\frac{n \log p}{p}}, \frac{\mu}{16} \right\}. \quad (\text{II.9})$$

Here c_a to c_c are positive constants.

Proof. By Theorem II.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 IV.8), $g(\mathbf{w}(\mathbf{q}'); \mathbf{X}_0) \leq g(\mathbf{w}(\mathbf{q}); \mathbf{X}_0)$ 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

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

Repeating the argument $2n$ times in the vicinity of other signed basis vectors $\pm e_i$ gives $2n$ local minimizers of f . Indeed, the $2n$ symmetric sections cover the sphere with certain overlaps. We claim that none of the $2n$ local minimizers lies in the overlapped regions. This is due to the nearness of these local minimizers to standard basis vectors. To see this, w.l.o.g., suppose \mathbf{q} , which is the local minimizer next to e_n , is in the overlapped region determined by e_n and e_i for some $i \neq n$. This implies that

$$\|[w_1(\mathbf{q}), \dots, w_{i-1}(\mathbf{q}), w_{i+1}(\mathbf{q}), \dots, q_n]\|^2 < \frac{4n-1}{4n}$$

by the definition of our symmetric sections. On the other hand, we know

$$\|[w_1(\mathbf{q}), \dots, w_{i-1}(\mathbf{q}), w_{i+1}(\mathbf{q}), \dots, q_n]\|^2 \geq q_n^2 = 1 - \eta^2.$$

Thus, so long as $1 - \eta^2 \geq \frac{4n-1}{4n}$, or $\eta \leq 1/(2\sqrt{n})$, a contradiction arises. Since $\eta \in O(\mu)$ and $\mu \leq O(n^{-5/4})$ by our assumption, our claim is confirmed. There are no extra local minimizers, as any extra local minimizer must be contained in at least one of the $2n$ symmetric sections, making two different local minimizers in one section, contradicting the uniqueness result we obtained above. \square

Though the $2n$ isolated local minimizers may have different objective values, they are equally good in the sense each of them helps produce a close approximation to a certain row of \mathbf{X}_0 . As discussed in Section I-D2, 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 II.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_\star (particularly, the same constant as in Theorem II.1) and C , such that for any $\theta \in (0, 1/2)$ and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, when*

$$p \geq \frac{C}{c_\star^2 \theta^2} \max \left\{ \frac{n^4}{\mu^4}, \frac{n^5}{\mu^2} \right\} \kappa^8(\mathbf{A}_0) \log^4 \left(\frac{\kappa(\mathbf{A}_0) n}{\mu \theta} \right) \quad (\text{II.10})$$

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

$$\nabla^2 g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}}) \succeq \frac{c_\star \theta}{2\mu} \mathbf{I} \quad \forall \mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\| \leq \frac{\mu}{4\sqrt{2}}, \quad (\text{II.11})$$

$$\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}})}{\|\mathbf{w}\|} \geq \frac{1}{2} c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{\mu}{4\sqrt{2}} \leq \|\mathbf{w}\| \leq \frac{1}{20\sqrt{5}} \quad (\text{II.12})$$

$$\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{V}\mathbf{U}^*\bar{\mathbf{Y}}) \mathbf{w}}{\|\mathbf{w}\|^2} \leq -\frac{1}{2} c_\star \theta \quad \forall \mathbf{w} \quad \text{s.t.} \quad \frac{1}{20\sqrt{5}} \leq \|\mathbf{w}\| \leq \sqrt{\frac{4n-1}{4n}}, \quad (\text{II.13})$$

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

$$\|\mathbf{w}_\star - \mathbf{0}\| \leq \mu/7. \quad (\text{II.14})$$

Here c_a, a_b are both positive constants.

Corollary II.4. Suppose \mathbf{A}_0 is complete with its condition number $\kappa(\mathbf{A}_0)$. There exist positive constants c_* (particularly, the same constant as in Theorem II.1) and C , such that for any $\theta \in (0, 1/2)$ and $\mu < c_a \min\{\theta n^{-1}, n^{-5/4}\}$, when $p \geq \frac{C}{c_*^2 \theta^2} \max\left\{\frac{n^4}{\mu^4}, \frac{n^5}{\mu^2}\right\} \kappa^8(\mathbf{A}_0) \log^4\left(\frac{\kappa(\mathbf{A}_0)n}{\mu\theta}\right)$ 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 $1 - c_b p^{-5}$, 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

$$\|\mathbf{q}_* - \mathbf{b}\| \leq \sqrt{2}\mu/7. \quad (\text{II.15})$$

Here c_a, c_b are both positive constants.

We omit the proof to Corollary II.4 as it is almost identical to that of corollary II.2. From the above theorems, it is clear that for any saddle point in the \mathbf{w} space, the Hessian has at least one negative eigenvalue with an associated eigenvector $\mathbf{w}/\|\mathbf{w}\|$. Now the question is whether all saddle points of f on \mathbb{S}^{n-1} have analogous properties, since as alluded to in Section I-D3, we need to perform actual optimization in the \mathbf{q} space. This is indeed true, but we will only argue informally in the companion paper [3]. The arguments need to be put in the language of Riemannian geometry, and we can switch back and forth between \mathbf{q} and \mathbf{w} spaces in our algorithm analysis without stating this fact.

B. Useful Technical Lemmas and Proof Ideas for Orthogonal Dictionaries

Proving Theorem II.1 is conceptually straightforward: one shows that the expectation of each quantity of interest has the claimed property, and then proves that each quantity concentrates uniformly about its expectation. The detailed calculations are nontrivial.

Note that

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

The next three propositions show that in the expected function landscape, we see successively strongly convex region, large gradient region, and negative directional curvature region when moving away from zero, as depicted in Fig. 3 and sketched in Section I-D2.

Proposition II.5. For any $\theta \in (0, 1/2)$, if $\mu \leq c \min\{\theta n^{-1}, n^{-5/4}\}$, it holds for all \mathbf{w} with $1/(20\sqrt{5}) \leq \|\mathbf{w}\| \leq \sqrt{(4n-1)/(4n)}$ that

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

Here $c > 0$ is a constant.

Proof. See Page 17 under Section IV-A1. □

Proposition II.6. For any $\theta \in (0, 1/2)$, if $\mu \leq 9/50$, it holds for all \mathbf{w} with $\mu/(4\sqrt{2}) \leq \|\mathbf{w}\| \leq 1/(20\sqrt{5})$ 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 Page 23 under Section IV-A2. □

Proposition II.7. For any $\theta \in (0, 1/2)$, if $\mu \leq 1/(20\sqrt{n})$, it holds for all \mathbf{w} with $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$ that

$$\nabla_{\mathbf{w}}^2 \mathbb{E} [h_\mu(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \succeq \frac{\theta}{5\sqrt{2\pi}\mu} \mathbf{I}.$$

Proof. See Page 24 under Section IV-A3. □

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 concentrates 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 II.8. *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 28 under Section IV-A4. □

Proposition II.9. *Suppose $0 < \mu \leq 1/\sqrt{n}$. 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 28 under Section IV-A4. □

Proposition II.10. *Suppose $0 < \mu \leq 1/\sqrt{n}$. For every $\mathbf{w} \in \Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \leq 1/4\}$, it holds that for any $t > 0$,*

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

Proof. See Page 29 under Section IV-A4. □

The next three propositions provide the desired Lipschitz results.

Proposition II.11 (Hessian Lipschitz). *Fix any $r_\wedge \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_\wedge\}$, $\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w} / \|\mathbf{w}\|^2$ is L_\wedge -Lipschitz with*

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

Proof. See Page 33 under Section IV-A5. □

Proposition II.12 (Gradient Lipschitz). *Fix any $r_g \in (0, 1)$. Over the set $\Gamma \cap \{\mathbf{w} : \|\mathbf{w}\| \geq r_g\}$, $\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0) / \|\mathbf{w}\|$ is L_g -Lipschitz with*

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

Proof. See Page 33 under Section IV-A5. □

Proposition II.13 (Lipschitz for Hessian around zero). *Fix any $r_\vee \in (0, 1/2)$. 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 34 under Section IV-A5. □

Integrating the above pieces, Section IV-B provides a complete proof of Theorem II.1.

C. Extending to Complete Dictionaries

As hinted in Section I-D2, 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

$$\bar{\mathbf{Y}} \doteq \sqrt{p\theta} (\mathbf{Y}\mathbf{Y}^*)^{-1/2} \mathbf{Y} \tag{II.16}$$

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 II.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 (II.16) as

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

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 $C > 0$ is a constant.

Proof. See Page 37 under Section IV-C. □

Notice that $\mathbf{U}\mathbf{V}^*$ 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 $f(\mathbf{q}; \mathbf{X}_0 + \mathbf{V}\mathbf{U}^* \mathbf{\Xi} \mathbf{X}_0) = f(\mathbf{q}; \mathbf{X}_0 + \tilde{\mathbf{\Xi}} \mathbf{X}_0)$ with $\tilde{\mathbf{\Xi}} \doteq \mathbf{V}\mathbf{U}^* \mathbf{\Xi}$. 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 II.15. 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 $1 - \theta(np)^{-7} - \exp(-0.3\theta np)$. Here C_a, C_b are positive constants.

Proof. See Page 37 under Section IV-C. □

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 II.14), and hence changes to the gradient and Hessian caused by the perturbation are small. This gives the results presented in Theorem II.3; see Section IV-C for a detailed proof. In particular, for the p chosen in Theorem II.3, it holds that

$$\|\tilde{\mathbf{\Xi}}\| \leq c c_\star \theta \left(\max\left\{ \frac{n^{3/2}}{\mu^2}, \frac{n^2}{\mu} \right\} \log^{3/2}(np) \right)^{-1} \quad (\text{II.17})$$

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

III. DISCUSSION

The dependency of p on n and other parameters could be suboptimal due to several factors: (1) The ℓ^1 proxy. Derivatives of the logcosh function we adopted entail the tanh function, which is not amenable to effective approximation and affects the sample complexity; (2) Space of geometric characterization. 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) Dealing with the complete case. Treating the complete case directly, rather than using (pessimistic) bounds to treat it as a perturbation of the orthogonal case, is very likely to improve the sample complexity. 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 [54], [118]. 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 w.r.t. the coefficients and work with implicit functions.²⁷ For the coefficient model, as we alluded to in Section I-E, 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 I-E 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 [43]–[47], [50], [73]–[78], [83], [84], [87], [88], [91], [92], the initialization plus local refinement strategy mostly differs from practice, whereby random initializations seem to work well, and the analytic techniques developed in that line are mostly fragmented and highly specialized. The analytic and algorithmic framework we developed here holds promise to providing a coherent account of these problems, see [51]. In particular, we have intentionally separated the geometric characterization and algorithm development, hoping to making both parts modular. It is interesting to see how far we can streamline the geometric characterization. Moreover, the separation allows development of more provable and practical algorithms, say in the direction of [52].

IV. PROOFS OF TECHNICAL RESULTS

In this section, we provide complete proofs for technical results stated in Section II. Before that, let us introduce some convenient notations and common results. Since we deal with BG random variables and random vectors, it is often convenient to write such vector explicitly as $\mathbf{x} = [\Omega_1 v_1, \dots, \Omega_n v_n] = \mathbf{\Omega} \odot \mathbf{v}$, where $\Omega_1, \dots, \Omega_n$ are i.i.d. Bernoulli 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} . Naturally, x_n and $q_n(\mathbf{w})$ denote the last coordinates in \mathbf{x} and \mathbf{q} , respectively. Hence, by our notation, $\mathbf{q}^*(\mathbf{w})\mathbf{x} = \mathbf{w}^*\bar{\mathbf{x}} + q_n(\mathbf{w})x_n$. By Lemma A.1 and chain rules, the following are immediate:

$$\nabla_{\mathbf{w}} h_{\mu}(\mathbf{q}^*(\mathbf{w})\mathbf{x}) = \tanh\left(\frac{\mathbf{q}^*(\mathbf{w})\mathbf{x}}{\mu}\right) \left(\bar{\mathbf{x}} - \frac{x_n}{q_n(\mathbf{w})}\mathbf{w}\right), \quad (\text{IV.1})$$

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

A. Proofs for Section II-B

1) *Proof of Proposition II.5:* The proof involves some delicate analysis, particularly polynomial approximation of the function $f(t) = 1/(1+t)^2$ 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 IV.1. *Consider $f(t) = 1/(1+t)^2$ 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 $0 < \sum_{k=0}^{\infty} \frac{b_k}{(1+k)^3} < \sum_{k=0}^{\infty} \frac{|b_k|}{(1+k)^3} < 2$.

Proof. See page 42 under Section B. □

²⁷This recent work [47] 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.

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

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

Proof. For $X + Y \geq 0$, let $Z = \exp(-2(X + Y)/\mu) \in [0, 1]$, then

$$1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) = \frac{4Z}{(1+Z)^2}.$$

First fix any $T > 1$. By Lemma IV.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 due to

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

and the dominated convergence theorem (see, e.g., theorem 2.24 and 2.25 of [119]). 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 rearranged the terms to obtain the last line. Using the following estimates (see Lemma IV.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

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

Since the above holds for any $T > 1$, we obtain the claimed result by letting $T \rightarrow \infty$ such that $\beta \rightarrow 1$. \square

Lemma IV.3. Let $X \sim \mathcal{N}(0, \sigma_X^2)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$ be independent. We have

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

Proof. For $X + Y > 0$, let $z = \exp(-2(X+Y)/\mu) \in [0, 1]$, then

$$1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) = \frac{4Z}{(1+Z)^2}.$$

First fix any $T > 1$. By Lemma IV.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

$$\begin{aligned} & \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) XY \mathbb{1}_{X+Y>0} \right] \\ &= 4\mathbb{E} [Z f(Z) XY \mathbb{1}_{X+Y>0}] \\ &= 4\mathbb{E} [Z p_\beta(Z) XY \mathbb{1}_{X+Y>0}] - 4\mathbb{E} [Z (p_\beta(Z) - f(Z)) XY \mathbb{1}_{X+Y>0}]. \end{aligned}$$

Now the first term can be rewritten as

$$4\mathbb{E} [Z p_\beta(Z) XY \mathbb{1}_{X+Y>0}] = 4 \sum_{k=0}^{\infty} \left\{ b_k \mathbb{E} \left[Z^{k+1} X^2 \mathbb{1}_{X+Y>0} \right] \right\},$$

where $b_k = (-1)^k (k+1)\beta^k$, and exchange of infinite summation and expectation is justified, due to

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

and the dominated convergence theorem (see, e.g., theorem 2.24 and 2.25 of [119]). By Lemma B.1, we have

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

where we have applied Type I lower and upper bounds for $\Phi^c(\cdot)$ to even k and odd k respectively and rearranged the terms to obtain the last line. Using the following estimates (see Lemma IV.1)

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

we obtain

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

For the second term, by Lemma IV.1, we have

$$\begin{aligned} \mathbb{E} [(p_\beta(Z) - f(Z)) Z XY \mathbb{1}_{X+Y>0}] &\leq \|p - f\|_{L^\infty[0,1]} \mathbb{E} [Z |XY| \mathbb{1}_{X+Y>0}] \\ &\leq \|p - f\|_{L^\infty[0,1]} \mathbb{E} [|X|] \mathbb{E} [|Y|] \quad (Z \leq 1, X, Y \text{ independent}) \\ &\leq \frac{2}{\pi\sqrt{T}} \sigma_X \sigma_Y. \end{aligned}$$

Thus,

$$\mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{X+Y}{\mu} \right) \right) XY \mathbb{1}_{X+Y>0} \right] \geq -\frac{2}{\sqrt{2\pi}} \frac{\mu \sigma_X^2 \sigma_Y^2}{(\sigma_X^2 + \sigma_Y^2)^{3/2}} \frac{1}{1+\beta} - \frac{3}{\sqrt{2\pi}} \frac{\sigma_X^2 \sigma_Y^2 \mu^3}{(\sigma_X^2 + \sigma_Y^2)^{5/2}} - \frac{8}{\pi \sqrt{T}} \sigma_X \sigma_Y.$$

Since the above bound holds for any $T > 1$, we obtain the claimed result by letting $T \rightarrow \infty$ such that $\beta \rightarrow 1$ and $1/\sqrt{T} \rightarrow 0$. \square

Lemma IV.4. *Let $X \sim \mathcal{N}(0, \sigma_X^2)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$ be independent. We have*

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

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 $Z \doteq \exp\left(-2\frac{X+Y}{\mu}\right)$ and $f(Z) \doteq \frac{1}{(1+Z)^2}$. First fix 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

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

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 [119]) applicable. To proceed, from Lemma B.1, we obtain

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

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

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

For the second term, by Lemma B.1 and Lemma IV.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

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

Since the above holds for any $T > 1$, we obtain the claimed result by letting $T \rightarrow \infty$, such that $\beta \rightarrow 1$ and $1/\sqrt{T} \rightarrow 0$. \square

Proof. (of Proposition II.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.3 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.3 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

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

1) **An upper bound for (A).** When x_n is not in support set of \mathbf{x} , the term reduces to

$$\begin{aligned} &\frac{2}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) \right) (\mathbf{w}^* \bar{\mathbf{x}})^2 \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} > 0} \right] \\ &\leq \frac{8}{\mu} \mathbb{E} \left[\exp \left(-2 \frac{\mathbf{w}^* \bar{\mathbf{x}}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}})^2 \mathbb{1}_{\mathbf{w}^* \bar{\mathbf{x}} > 0} \right] \quad (\text{Lemma A.1}) \\ &\leq 8 \exp(-2)\mu, \end{aligned}$$

where to obtain the last line we used that $t \mapsto \exp(-2t/\mu)t^2$ for $t > 0$ is maximized at μ .

When x_n is in the support set, we expand the square term inside the expectation and obtain

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

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 $Y \doteq \mathbf{w}_{\mathcal{J}}^* \bar{\mathbf{v}} \sim \mathcal{N}(0, \|\mathbf{w}_{\mathcal{J}}\|^2)$. An upper bound for the above is obtained by calling the estimates in Lemma IV.2 and Lemma IV.3:

$$\begin{aligned}
 (\mathcal{A})_{x_n \neq 0} &\leq \frac{2}{\mu} \mathbb{E}_{\mathcal{J}} \left[\frac{1}{\sqrt{2\pi}} \frac{\mu \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{3}{4\sqrt{2\pi}} \frac{\|\mathbf{w}_{\mathcal{J}}\|^2 \mu^3}{\|\mathbf{q}_{\mathcal{I}}\|^5} (3\mu^2 + 4\|\mathbf{w}_{\mathcal{J}}\|^2) \right] \\
 &\quad + \frac{2}{\mu} \frac{\|\mathbf{w}\|^4}{q_n^4(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{1}{\sqrt{2\pi}} \frac{\mu \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{3}{4\sqrt{2\pi}} \frac{q_n^2(\mathbf{w}) \mu^3}{\|\mathbf{q}_{\mathcal{I}}\|^5} (3\mu^2 + 4q_n^2(\mathbf{w})) \right] \\
 &\quad + \frac{4}{\mu} \frac{\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{1}{\sqrt{2\pi}} \frac{\mu \|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w})}{\|\mathbf{q}_{\mathcal{I}}\|^3} + \frac{3}{\sqrt{2\pi}} \frac{\|\mathbf{w}_{\mathcal{J}}\|^2 q_n^2(\mathbf{w}) \mu^3}{\|\mathbf{q}_{\mathcal{I}}\|^5} \right] \\
 &\leq \sqrt{\frac{2}{\pi}} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2 q_n^4(\mathbf{w}) + \|\mathbf{w}_{\mathcal{J}}\|^2 \|\mathbf{w}\|^4 + 2\|\mathbf{w}_{\mathcal{J}}\|^2 \|\mathbf{w}\|^2 q_n^2(\mathbf{w})}{q_n^2(\mathbf{w}) \|\mathbf{q}_{\mathcal{I}}\|^3} \right] \\
 &\quad + \frac{12\mu^2}{\sqrt{2\pi} q_n^5(\mathbf{w})} + \frac{9\mu^4}{2\sqrt{2\pi} q_n^5(\mathbf{w})} \\
 &\leq \frac{1}{q_n^2(\mathbf{w})} \sqrt{\frac{2}{\pi}} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + \frac{12\mu^2}{\sqrt{2\pi} q_n^5(\mathbf{w})} + \frac{9\mu^4}{2\sqrt{2\pi} q_n^5(\mathbf{w})},
 \end{aligned}$$

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.

Thus, we obtain that

$$\begin{aligned}
 (\mathcal{A}) &\leq (1 - \theta) \cdot 8 \exp(-2)\mu + \frac{\theta}{q_n^2(\mathbf{w})} \sqrt{\frac{2}{\pi}} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + \frac{12\theta\mu^2}{\sqrt{2\pi} q_n^5(\mathbf{w})} + \frac{9\theta\mu^4}{2\sqrt{2\pi} q_n^5(\mathbf{w})} \\
 &\leq \frac{\theta}{q_n^2(\mathbf{w})} \sqrt{\frac{2}{\pi}} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + 2\mu + \frac{12\theta\mu^2}{\sqrt{2\pi}} \left(\frac{1}{q_n^3(\mathbf{w})} + \frac{1}{q_n^5(\mathbf{w})} \right). \tag{IV.3}
 \end{aligned}$$

2) **A lower bound for (\mathcal{B}) .** Similarly, we obtain

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

$$\geq \sqrt{\frac{2}{\pi}} \frac{\theta}{q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|} \right] - \frac{4\theta\mu^2}{\sqrt{2\pi}q_n^5(\mathbf{w})}.$$

Collecting the above estimates, we obtain

$$\begin{aligned} & \mathbf{w}^* \nabla_{\mathbf{w}}^2 \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \mathbf{w} \\ & \leq \sqrt{\frac{2}{\pi}} \frac{\theta}{q_n^2(\mathbf{w})} \mathbb{E}_{\mathcal{J}} \left[\frac{\|\mathbf{w}_{\mathcal{J}}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} - \frac{\|\mathbf{w}\|^2 (\|\mathbf{w}_{\mathcal{J}}\|^2 + q_n^2(\mathbf{w}))}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + 2\mu + \frac{4\theta\mu^2}{\sqrt{2\pi}} \left(\frac{3}{q_n^3(\mathbf{w})} + \frac{4}{q_n^5(\mathbf{w})} \right) \\ & = -\sqrt{\frac{2}{\pi}} \theta \mathbb{E} \left[\frac{\|\mathbf{w}_{\mathcal{J}^c}\|^2}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + 2\mu + \frac{4\theta\mu^2}{\sqrt{2\pi}} \left(\frac{3}{q_n^3(\mathbf{w})} + \frac{4}{q_n^5(\mathbf{w})} \right) \\ & \leq -\sqrt{\frac{2}{\pi}} \theta (1-\theta) \|\mathbf{w}\|^2 \mathbb{E} \left[\frac{1}{\|\mathbf{q}_{\mathcal{I}}\|^3} \right] + 2\mu + \frac{4\theta\mu^2}{\sqrt{2\pi}} \left(\frac{3}{q_n^3(\mathbf{w})} + \frac{4}{q_n^5(\mathbf{w})} \right), \end{aligned} \quad (\text{IV.4})$$

where to obtain the last line we have invoked the association inequality in Lemma A.2, 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 (IV.4) and noting $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ (implied by the assumption $\|\mathbf{w}\| \leq \sqrt{(4n-1)/(4n)}$), we obtain the claimed result. \square

2) Proof of Proposition II.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

$$\begin{aligned} \mathbf{w}^* \nabla \mathbb{E} [h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] &= \nabla \mathbb{E} [\mathbf{w}^* h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \\ &= \underbrace{\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) (\mathbf{w}^* \bar{\mathbf{x}}) \right]}_{(A)} - \frac{\|\mathbf{w}\|^2}{q_n} \underbrace{\mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) x_n \right]}_{(B)}. \end{aligned} \quad (\text{IV.5})$$

1) A lower bound for (A). We have

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

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.2 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 $\tanh(z) \geq (1 - \exp(-2z))/2$ for all $z \geq 0$ and integral results in Lemma B.1, we obtain

$$\begin{aligned} \mathbb{E}\left[\tanh\left(\frac{Y}{\mu}\right)Y\right] &= 2\mathbb{E}\left[\tanh\left(\frac{Y}{\mu}\right)Y\mathbb{1}_{Y>0}\right] \\ &\geq \mathbb{E}\left[\left(1 - \exp\left(-\frac{2Y}{\mu}\right)\right)Y\mathbb{1}_{Y>0}\right] \\ &= \frac{2\sigma_Y^2}{\mu} \exp\left(\frac{2\sigma_Y^2}{\mu^2}\right) \Phi^c\left(\frac{2\sigma_Y}{\mu}\right) \\ &\geq \frac{2\sigma_Y^2}{\mu\sqrt{2\pi}} \left(\sqrt{1 + \frac{\sigma_Y^2}{\mu^2}} - \frac{\sigma_Y}{\mu}\right) \quad (\text{Type III lower bound for } \Phi^c(\cdot), \text{ Lemma A.4}) \\ &\geq \frac{2\sigma_Y^2}{\mu\sqrt{2\pi}} \left(\sqrt{1 + \frac{\|\mathbf{w}\|^2}{\mu^2}} - \frac{\|\mathbf{w}\|}{\mu}\right). \quad (t \mapsto \sqrt{1+t^2} - t \text{ decreasing over } t > 0) \end{aligned}$$

Collecting the above estimates, we have

$$\begin{aligned} \mathbb{E}\left[\tanh\left(\frac{\mathbf{q}^*(\mathbf{w})\mathbf{x}}{\mu}\right)(\mathbf{w}^*\bar{\mathbf{x}})\right] &\geq (1-\theta)\mathbb{E}_{\mathcal{J}}\left[\frac{2\|\mathbf{w}_{\mathcal{J}}\|^2}{\mu\sqrt{2\pi}}\left(\sqrt{1 + \frac{\|\mathbf{w}\|_2^2}{\mu^2}} - \frac{\|\mathbf{w}\|}{\mu}\right)\right] \\ &\geq (1-\theta)\mathbb{E}_{\mathcal{J}}\left[\frac{2\|\mathbf{w}_{\mathcal{J}}\|^2}{\mu\sqrt{2\pi}}\frac{\mu}{10\|\mathbf{w}\|}\right] \\ &\geq \frac{\theta(1-\theta)\|\mathbf{w}\|}{5\sqrt{2\pi}}, \end{aligned} \quad (\text{IV.6})$$

where at the second line we have used the assumption that $\|\mathbf{w}\| \geq \mu/(4\sqrt{2})$ and also the fact that $\sqrt{1+x^2} \geq x + \frac{1}{10x}$ for $x \geq 1/(4\sqrt{2})$.

2) **An upper bound for (B).** We have

$$\mathbb{E}\left[\tanh\left(\frac{\mathbf{q}^*(\mathbf{w})\mathbf{x}}{\mu}\right)x_n\right] \leq \theta\mathbb{E}\left[\left|\tanh\left(\frac{\mathbf{q}^*(\mathbf{w})\mathbf{x}}{\mu}\right)\right||v_n|\right] \leq \theta\sqrt{\frac{2}{\pi}}, \quad (\text{IV.7})$$

as $\tanh(\cdot)$ is bounded by one in magnitude

Plugging the results of (IV.6) and (IV.7) into (IV.5) 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 used $\frac{2\|\mathbf{w}\|}{\sqrt{1-\|\mathbf{w}\|^2}} \leq \frac{1}{10}(1-\theta)$ when $\|\mathbf{w}\| \leq 1/(20\sqrt{5})$ and $\theta \leq 1/2$, completing the proof. \square

3) *Proof of Proposition II.7:*

Proof. By consideration similar to proof of Proposition II.5, we can exchange the Hessian and expectation, i.e.,

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

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 $0 \leq \|\mathbf{w}\| \leq \mu/(4\sqrt{2})$.

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

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

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 $\mu \leq 1/(20\sqrt{n}) \leq 1/20$ and $\theta < 1/2$, we obtain

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

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

$$\begin{aligned}
 & \mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \\
 &= \underbrace{\frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) \bar{\mathbf{x}} \bar{\mathbf{x}}^* \right]}_{(\mathcal{A})} - \underbrace{\frac{1}{q_n^2(\mathbf{w})} \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n(\mathbf{w}) x_n \right] \mathbf{I}}_{(\mathcal{B})} \\
 &\quad - \underbrace{\frac{1}{\mu q_n^2(\mathbf{w})} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n (\mathbf{w} \bar{\mathbf{x}}^* + \bar{\mathbf{x}} \mathbf{w}^*) \right]}_{(\mathcal{C})} \\
 &\quad + \underbrace{\frac{1}{q_n^4(\mathbf{w})} \left\{ \frac{1}{\mu} \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) (q_n(\mathbf{w}) x_n)^2 \right] - \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n(\mathbf{w}) x_n \right] \right\}}_{(\mathcal{D})} \mathbf{w} \mathbf{w}^*.
 \end{aligned} \tag{IV.8}$$

We will first provide bounds for (C) and (D), which are relatively simple. Then we will bound (A) and (B), which are slightly more tricky.

1) **An upper bound for (C).** We have

$$\begin{aligned}
 (\mathcal{C}) &\leq \frac{1}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n (\mathbf{w} \bar{\mathbf{x}}^* + \bar{\mathbf{x}} \mathbf{w}^*) \right] \right\| \\
 &\leq \frac{2}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n \bar{\mathbf{x}} \right] \mathbf{w}^* \right\| \\
 &\leq \frac{2}{\mu q_n^2(\mathbf{w})} \left\| \mathbb{E} \left[\left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n \bar{\mathbf{x}} \right] \right\| \|\mathbf{w}\| \\
 &\leq \frac{2}{\mu q_n^2(\mathbf{w})} \mathbb{E} \left\| \left(1 - \tanh^2 \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) \right) q_n(\mathbf{w}) x_n \bar{\mathbf{x}} \right\| \|\mathbf{w}\| \quad (\text{Jensen's inequality}) \\
 &\leq \frac{2}{\mu q_n(\mathbf{w})} \theta^2 \mathbb{E} [|v_n|] \mathbb{E} [\|\bar{\mathbf{v}}\|] \|\mathbf{w}\| \quad (1 - \tanh^2(\cdot) \leq 1, x_n \text{ and } \bar{\mathbf{x}} \text{ independent}) \\
 &\leq \frac{4\theta^2}{\pi \mu q_n(\mathbf{w})} \sqrt{n} \|\mathbf{w}\| \leq \frac{\theta}{\mu} \frac{4\theta \sqrt{n} \|\mathbf{w}\|}{\pi \sqrt{1 - \|\mathbf{w}\|^2}} \leq \frac{\theta}{\mu} \frac{1}{40\pi},
 \end{aligned}$$

where to obtain the final bound we have invoked the assumptions: $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$, $\mu \leq 1/(20\sqrt{n})$, and $\theta \leq 1/2$.

- 2) **A lower bound for (\mathcal{D}) .** We directly drop the first expectation which is positive, and derive an upper for the second expectation term as:

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

Thus,

$$(\mathcal{D}) \succeq -\frac{1}{q_n^4(\mathbf{w})} \sqrt{\frac{2}{\pi}} \theta q_n(\mathbf{w}) \mathbf{w} \mathbf{w}^* \succeq -\frac{\theta}{q_n^3(\mathbf{w})} \sqrt{\frac{2}{\pi}} \|\mathbf{w}\|^2 \mathbf{I} \succeq -\frac{\theta}{32000\mu} \sqrt{\frac{2}{\pi}} \mathbf{I},$$

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

- 3) **An upper bound for (\mathcal{B}) .** Similar to the way we bound (\mathcal{D}) ,

$$\frac{1}{q_n^2(\mathbf{w})} \mathbb{E} \left[\tanh \left(\frac{\mathbf{q}^*(\mathbf{w}) \mathbf{x}}{\mu} \right) q_n x_n \right] \leq \frac{1}{q_n(\mathbf{w})} \sqrt{\frac{2}{\pi}} \theta \leq \frac{\theta}{10\mu} \sqrt{\frac{2}{\pi}}.$$

- 4) **A lower bound for (\mathcal{A}) .** First note that

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

Thus, 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 $\beta \in (0, 1)$, as $\mathbb{E}_{\bar{\mathbf{x}}} [\bar{\mathbf{x}} \bar{\mathbf{x}}^*] = \theta \mathbf{I}$. 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 $\text{diag}(\boldsymbol{\Sigma}; \alpha \theta \mathbf{I}_{n-1-k})$, where

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

So in order to derive the $\theta \beta \mathbf{I}$ lower bound as desired, it is sufficient to show $\boldsymbol{\Sigma} \succeq \theta \beta \mathbf{I}$ and $\beta \leq \alpha$, 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 above implication, suppose the latter claimed holds, then for any \mathbf{z} with unit norm,

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

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] \\
 &\quad + \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] \quad (\text{cross-term vanishes due to independence}) \\
 &= \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] \\
 &\quad + \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] \quad (\text{exp. factorizes due to independence}) \\
 &\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, the above lower bound is further bounded as:

$$\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 2 \frac{(\tilde{\mathbf{w}}_{\mathcal{S}}^* \mathbf{z})^2}{\|\mathbf{w}_{\mathcal{S}}\|^4} \left[\left(\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^2 + \frac{4}{\mu^2} \|\tilde{\mathbf{w}}_{\mathcal{S}}\|^4 \right) \exp \left(\frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^2}{\mu^2} \right) \Phi^c \left(\frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|}{\mu} \right) - \frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^3}{\mu\sqrt{2\pi}} \right] \\
 &\quad + 2 \exp \left(\frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^2}{\mu^2} \right) \Phi^c \left(\frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|}{\mu} \right) \|(\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z}\|^2 \\
 &\geq \underbrace{\left(2 \frac{(\tilde{\mathbf{w}}_{\mathcal{S}}^* \mathbf{z})^2}{\|\mathbf{w}_{\mathcal{S}}\|^2} + 2 \|(\mathbf{I} - \mathcal{P}_{\tilde{\mathbf{w}}_{\mathcal{S}}}) \mathbf{z}\|^2 \right)}_{=2\|\mathbf{z}\|^2=2} \exp \left(\frac{2\|\tilde{\mathbf{w}}_{\mathcal{S}}\|^2}{\mu^2} \right) \Phi^c \left(\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}}\|} \\
 &\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}}\|} \quad (\text{Type III lower bound for } \Phi^c(\cdot)) \\
 &\geq \frac{1}{\sqrt{2\pi}} \left(\sqrt{4 + \frac{4\|\mathbf{w}\|^2}{\mu^2}} - \frac{2\|\mathbf{w}\|}{\mu} \right) - \frac{4\|\mathbf{w}\|}{\mu\sqrt{2\pi}} \quad (t \mapsto \sqrt{4+t^2} - t \text{ nonincreasing and Cauchy-Schwarz}) \\
 &\geq \frac{1}{\sqrt{2\pi}} \left(2 - \frac{3}{4}\sqrt{2} \right),
 \end{aligned}$$

where to obtain the last line we have used $\|\mathbf{w}\| \leq \mu/(4\sqrt{2})$. On the other hand, we similarly obtain

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

So we can take $\beta = \frac{1}{\sqrt{2\pi}} \left(2 - \frac{3}{4}\sqrt{2} \right) < 1$.

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

$$\mathbb{E} [\nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}^*(\mathbf{w}) \mathbf{x})] \succeq \frac{\theta}{\mu\sqrt{2\pi}} \left(1 - \frac{3}{8}\sqrt{2} - \frac{\sqrt{2\pi}}{40\pi} - \frac{1}{16000} - \frac{1}{5} \right) \mathbf{I} \succeq \frac{1}{5\sqrt{2\pi}} \frac{\theta}{\mu} \mathbf{I}.$$

Hence for all \mathbf{w} , we can take the $\frac{1}{5\sqrt{2\pi}} \frac{\theta}{\mu}$ as the lower bound, completing the proof. \square

4) *Proof of Pointwise Concentration Results:* We first establish a useful comparison lemma between random i.i.d. Bernoulli random vectors random i.i.d. normal random vectors.

Lemma IV.5. *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 $m \geq 1$.

Proof. See page 42 under Section B. □

Now, we are ready to prove Proposition II.8 to Proposition II.10 as follows.

Proof. (of Proposition II.8) Let

$$X_k = \frac{\mathbf{w}^*}{\|\mathbf{w}\|_2} \nabla h_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k),$$

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

$$|X_k| = \left| \tanh\left(\frac{\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k}{\mu}\right) \left(\frac{\mathbf{w}^* \bar{\mathbf{x}}_{0k}}{\|\mathbf{w}\|} - \frac{\|\mathbf{w}\|_2 x_{0k}(n)}{q_n(\mathbf{w})}\right) \right| \leq \left| \frac{\mathbf{w}^* \bar{\mathbf{x}}_{0k}}{\|\mathbf{w}\|} - \frac{\|\mathbf{w}\|_2 x_{0k}(n)}{q_n(\mathbf{w})} \right|,$$

as the magnitude of $\tanh(\cdot)$ is bounded by one. Because

$$\frac{\mathbf{w}^* \bar{\mathbf{x}}_{0k}}{\|\mathbf{w}\|_2} - \frac{\|\mathbf{w}\|_2 x_{0k}(n)}{q_n(\mathbf{w})} = \left(\frac{\mathbf{w}}{\|\mathbf{w}\|}, -\frac{\|\mathbf{w}\|}{q_n(\mathbf{w})}\right)^* (\mathbf{x}_0)_k \quad \text{and} \quad (\mathbf{x}_0)_k \sim_{i.i.d.} \text{BG}(\theta),$$

invoking Lemma IV.5, 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.5 and we have used that $(m-1)!! \leq m!/2$ and the assumption that $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ 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.8. □

Proof. (of Proposition II.9) Let

$$Y_k = \frac{1}{\|\mathbf{w}\|^2} \mathbf{w}^* \nabla^2 h_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) \mathbf{w},$$

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

$$Y_k = \underbrace{\frac{1}{\mu} \left(1 - \tanh^2\left(\frac{\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k}{\mu}\right)\right)}_{\doteq W_k} \underbrace{\left(\frac{\mathbf{w}^* \bar{\mathbf{x}}_{0k}}{\|\mathbf{w}\|} - \frac{x_{0k}(n) \|\mathbf{w}\|}{q_n(\mathbf{w})}\right)^2}_{\doteq V_k} - \tanh\left(\frac{\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k}{\mu}\right) \frac{x_{0k}(n)}{q_n^3(\mathbf{w})}.$$

Then by similar argument as in proof to Proposition II.8, we have for all integers $m \geq 2$ that

$$\begin{aligned} \mathbb{E} [|W_k|^m] &\leq \frac{1}{\mu^m} \mathbb{E} \left[\left| \frac{\mathbf{w}^* \bar{\mathbf{x}}_k}{\|\mathbf{w}\|} - \frac{x_k(n) \|\mathbf{w}\|}{q_n(\mathbf{w})} \right|^{2m} \right] \leq \frac{1}{\mu^m} \mathbb{E}_{Z \sim \mathcal{N}(0, 1/q_n^2(\mathbf{w}))} [|Z|^{2m}] \\ &\leq \frac{1}{\mu^m} (2m-1)!! (4n)^m \leq \frac{m!}{2} \left(\frac{4n}{\mu}\right)^m, \\ \mathbb{E} [|V_k|^m] &\leq \frac{1}{q_n^{3m}(\mathbf{w})} \mathbb{E} [|v_k(n)|^m] \leq (2\sqrt{n})^{3m} (m-1)!! \leq \frac{m!}{2} (8n\sqrt{n})^m, \end{aligned}$$

where we have again used the assumption that $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ 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.8, 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

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

provided that $\mu \leq 1/\sqrt{n}$, as desired. \square

Proof. (of Proposition II.10) Let $\mathbf{Z}_k = \nabla_{\mathbf{w}}^2 h_{\mu}(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k)$, then $\nabla_{\mathbf{w}}^2 g(\mathbf{w}) = \frac{1}{p} \sum_{k=1}^p \mathbf{Z}_k$. From (IV.2), we know that

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

where

$$\begin{aligned} \mathbf{W}_k &= \frac{1}{\mu} \left(1 - \tanh^2 \left(\frac{\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k}{\mu} \right) \right) \left(\bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right) \left(\bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right)^* \\ \mathbf{V}_k &= -\tanh \left(\frac{\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k}{\mu} \right) \left(\frac{x_{0k}(n)}{q_n(\mathbf{w})} \mathbf{I} + \frac{x_{0k}(n)\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right). \end{aligned}$$

For \mathbf{W}_k , we have

$$\begin{aligned} \mathbf{0} \preceq \mathbb{E}[\mathbf{W}_k^m] &\preceq \frac{1}{\mu^m} \mathbb{E} \left[\left\| \bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right\|^{2m-2} \left(\bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right) \left(\bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right)^* \right] \\ &\preceq \frac{1}{\mu^m} \mathbb{E} \left[\left\| \bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)\mathbf{w}}{q_n(\mathbf{w})} \right\|^{2m} \right] \mathbf{I} \\ &\preceq \frac{2^m}{\mu^m} \mathbb{E} \left[\left(\|\bar{\mathbf{x}}_{0k}\|^2 + \frac{x_{0k}^2(n)\|\mathbf{w}\|^2}{q_n^2(\mathbf{w})} \right)^m \right] \mathbf{I} \\ &\preceq \frac{2^m}{\mu^m} \mathbb{E} \left[\|(\mathbf{x}_0)_k\|^{2m} \right] \mathbf{I} \preceq \frac{2^m}{\mu^m} \mathbb{E}_{Z \sim \chi^2(n)} [Z^m] \mathbf{I}, \end{aligned}$$

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 $\|\mathbf{w}\|_2 \leq 1/4$ and Lemma IV.5 to obtain the last line. By Lemma A.6, 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 $R_W = 4n/\mu$ and $\sigma_W^2 = 16n^2/\mu^2 \geq \mathbb{E}[\mathbf{W}_k^2]$, and letting $\mathbf{S}_W = \frac{1}{p} \sum_{k=1}^p \mathbf{W}_k$, by Lemma A.9, 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

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

where we have used the fact $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ to simplify the result. Similar argument also shows $-\mathbb{E}[\mathbf{V}_k^m] \leq m!(8n\sqrt{n})^m \mathbf{I}/2$. 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.9, 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 using $\mu \leq 1/\sqrt{n}$. \square

5) *Proof of Lipschitz Results:* We need the following lemmas to prove the Lipschitz results.

Lemma IV.6. *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 IV.7. *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 $\mathbf{x} \in \mathcal{D}$ 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 IV.8. *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

$$\begin{aligned} \|\mathbf{q}(\mathbf{w}) - \mathbf{q}(\mathbf{w}')\|^2 &= \|\mathbf{w} - \mathbf{w}'\|^2 + |q_n(\mathbf{w}) - q_n(\mathbf{w}')|^2 \leq \left(1 + \frac{\max(\|\mathbf{w}\|^2, \|\mathbf{w}'\|^2)}{\min(q_n^2(\mathbf{w}), q_n^2(\mathbf{w}'))}\right) \|\mathbf{w} - \mathbf{w}'\|^2 \\ &= \frac{1}{\min(q_n^2(\mathbf{w}), q_n^2(\mathbf{w}'))} \|\mathbf{w} - \mathbf{w}'\|^2 \leq 4n \|\mathbf{w} - \mathbf{w}'\|^2, \end{aligned}$$

where we have used the fact $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ 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 IV.6, we obtain the desired claims. \square

Lemma IV.9. 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 any 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 $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ 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

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

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 $q_n(\mathbf{w}) \geq 1/(2\sqrt{n})$ to simplify the last result. Collecting the above estimates, we obtain

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

as desired. For the last one, we have

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

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

Lemma IV.10. *For any fixed \mathbf{x} , consider the function*

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

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 any constant $r \in (0, 1/2)$, it holds that

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

Proof. Simple calculation shows

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

For the second one, we have

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

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 IV.8 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

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

Combining the above estimates gives the claimed result. \square

Lemma IV.11. *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 any constant $r \in (0, 1/2)$, 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. \square

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

Proof. (of Proposition II.11) Let

$$F_k(\mathbf{w}) = \ddot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) t_{(\mathbf{x}_0)_k}^2(\mathbf{w}) + \dot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) \frac{x_{0k}(n)}{q_n^3(\mathbf{w})}.$$

Then, $\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w} / \|\mathbf{w}\|^2 = \frac{1}{p} \sum_{k=1}^p F_k(\mathbf{w})$. Noticing that $\ddot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k)$ is bounded by $1/\mu$ and $\dot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k)$ is bounded by 1, both in magnitude. Applying Lemma IV.7, Lemma IV.8 and Lemma IV.9, we can see $F_k(\mathbf{w})$ is L_\wedge^k -Lipschitz with

$$\begin{aligned}
 L_\wedge^k &= 4n \|\mathbf{x}_0\|_k^2 \frac{4\sqrt{n}}{\mu^2} \|\mathbf{x}_0\|_k + \frac{1}{\mu} 8\sqrt{n} \|\mathbf{x}_0\|_k \left(\frac{\|\mathbf{x}_0\|_k}{r_\wedge} + 4n^{3/2} \|\mathbf{x}_0\|_\infty \right) \\
 &\quad + (2\sqrt{n})^3 \|\mathbf{x}_0\|_\infty \frac{2\sqrt{n}}{\mu} \|\mathbf{x}_0\|_k + \sup_{r_\wedge < a < \sqrt{\frac{2n-1}{2n}}} \frac{3}{(1-a^2)^{5/2}} \|\mathbf{x}_0\|_\infty \\
 &= \frac{16n^{3/2}}{\mu^2} \|\mathbf{x}_0\|_k^3 + \frac{8\sqrt{n}}{\mu r_\wedge} \|\mathbf{x}_0\|_k^2 + \frac{48n^2}{\mu} \|\mathbf{x}_0\|_k \|\mathbf{x}_0\|_\infty + 96n^{5/2} \|\mathbf{x}_0\|_\infty.
 \end{aligned}$$

Thus, $\frac{1}{\|\mathbf{w}\|_2} \mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}$ is L_\wedge -Lipschitz with

$$L_\wedge \leq \frac{1}{p} \sum_{k=1}^p L_\wedge^k \leq \frac{16n^3}{\mu^2} \|\mathbf{X}_0\|_\infty^3 + \frac{8n^{3/2}}{\mu r_\wedge} \|\mathbf{X}_0\|_\infty^2 + \frac{48n^{5/2}}{\mu} \|\mathbf{X}_0\|_\infty^2 + 96n^{5/2} \|\mathbf{X}_0\|_\infty,$$

as desired. \square

Proof. (of Proposition II.12) We have

$$\left\| \frac{\mathbf{w}^*}{\|\mathbf{w}\|} \nabla g(\mathbf{w}; \mathbf{X}_0) - \frac{\mathbf{w}'^*}{\|\mathbf{w}'\|} \nabla g(\mathbf{w}'; \mathbf{X}_0) \right\| \leq \frac{1}{p} \sum_{k=1}^p \left\| \dot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) t_{(\mathbf{x}_0)_k}(\mathbf{w}) - \dot{h}_\mu(\mathbf{q}(\mathbf{w}')^*(\mathbf{x}_0)_k) t_{(\mathbf{x}_0)_k}(\mathbf{w}') \right\|$$

where $\dot{h}_\mu(t) = \tanh(t/\mu)$ is bounded by one in magnitude, and $t_{(\mathbf{x}_0)_k}(\mathbf{w})$ and $t_{(\mathbf{x}_0)_k}(\mathbf{w}')$ is defined as in Lemma IV.9. By Lemma IV.7, Lemma IV.8 and Lemma IV.9, we know that $\dot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) t_{(\mathbf{x}_0)_k}(\mathbf{w})$ is L_k -Lipschitz with constant

$$L_k = \frac{2\|\mathbf{x}_0\|_k}{r_g} + 8n^{3/2} \|\mathbf{x}_0\|_\infty + \frac{4n}{\mu} \|\mathbf{x}_0\|_k^2.$$

Therefore, we have

$$\begin{aligned} \left\| \frac{\mathbf{w}^*}{\|\mathbf{w}\|} \nabla g(\mathbf{w}) - \frac{\mathbf{w}^*}{\|\mathbf{w}\|} \nabla g(\mathbf{w}') \right\| &\leq \frac{1}{p} \sum_{k=1}^p \left(\frac{2 \|(\mathbf{x}_0)_k\|}{r_g} + 8n^{3/2} \|(\mathbf{x}_0)_k\|_\infty + \frac{4n}{\mu} \|(\mathbf{x}_0)_k\|^2 \right) \|\mathbf{w} - \mathbf{w}'\| \\ &\leq \left(\frac{2\sqrt{n}}{r_g} \|\mathbf{X}_0\|_\infty + 8n^{3/2} \|\mathbf{X}_0\|_\infty + \frac{4n^2}{\mu} \|\mathbf{X}_0\|_\infty^2 \right) \|\mathbf{w} - \mathbf{w}'\|, \end{aligned}$$

as desired. \square

Proof. (of Proposition II.13) Let

$$\mathbf{F}_k(\mathbf{w}) = \ddot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) \zeta_k(\mathbf{w}) \zeta_k(\mathbf{w})^* - \dot{h}_\mu(\mathbf{q}(\mathbf{w})^*(\mathbf{x}_0)_k) \Phi_k(\mathbf{w})$$

with $\zeta_k(\mathbf{w}) = \bar{\mathbf{x}}_{0k} - \frac{x_{0k}(n)}{q_n(\mathbf{w})} \mathbf{w}$ and $\Phi_k(\mathbf{w}) = \frac{x_{0k}(n)}{q_n(\mathbf{w})} \mathbf{I} + \frac{x_{0k}(n)}{q_n^3(\mathbf{w})} \mathbf{w} \mathbf{w}^*$. Then, $\nabla^2 g(\mathbf{w}) = \frac{1}{p} \sum_{k=1}^p \mathbf{F}_k(\mathbf{w})$. Using Lemma IV.7, Lemma IV.8, Lemma IV.10 and Lemma IV.11, and the facts that $\ddot{h}_\mu(t)$ is bounded by $1/\mu$ and that $\dot{h}_\mu(t)$ is bounded by 1 in magnitude, we can see $\mathbf{F}_k(\mathbf{w})$ is L_\cup^k -Lipschitz continuous with

$$\begin{aligned} L_\cup^k &= \frac{1}{\mu} \times 8\sqrt{2}\sqrt{n} \|(\mathbf{x}_0)_k\|_\infty^2 + \frac{2\sqrt{n}}{\mu^2} \|(\mathbf{x}_0)_k\| \times 2n \|(\mathbf{x}_0)_k\|_\infty^2 + 4 \|(\mathbf{x}_0)_k\|_\infty + \frac{2\sqrt{n}}{\mu} \|(\mathbf{x}_0)_k\| \times 2 \|(\mathbf{x}_0)_k\|_\infty \\ &\leq \frac{4n^{3/2}}{\mu^2} \|(\mathbf{x}_0)_k\| \|(\mathbf{x}_0)_k\|_\infty^2 + \frac{4\sqrt{n}}{\mu} \|(\mathbf{x}_0)_k\| \|(\mathbf{x}_0)_k\|_\infty + \frac{8\sqrt{2}\sqrt{n}}{\mu} \|(\mathbf{x}_0)_k\|_\infty^2 + 4 \|(\mathbf{x}_0)_k\|_\infty. \end{aligned}$$

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. \square

B. Proofs of Theorem II.1

Before proving Theorem II.1, we record one useful lemma.

Lemma IV.12. *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*

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

Proof. See page 43 under Section B. \square

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

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

Proof. (of Theorem II.1)

a) *Strong convexity in region R_1 .*: Proposition II.7 shows that for any $\mathbf{w} \in R_1$, $\mathbb{E} [\nabla^2 g(\mathbf{w}; \mathbf{X}_0)] \succeq \frac{c_1 \theta}{\mu} \mathbf{I}$. 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 II.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

$$\mathcal{E}_1 = \left\{ \max_{\mathbf{w} \in N_1} \left\| \nabla^2 g(\mathbf{w}; \mathbf{X}_0) - \mathbb{E} [\nabla^2 g(\mathbf{w}; \mathbf{X}_0)] \right\| \leq \frac{c_1 \theta}{3\mu} \right\}.$$

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

$$\sup_{\|\mathbf{w}\| \leq \mu / (4\sqrt{2})} \left\| \nabla^2 g(\mathbf{w}; \mathbf{X}_0) - \mathbb{E} [\nabla^2 g(\mathbf{w}; \mathbf{X}_0)] \right\| \leq \frac{2c_1 \theta}{3\mu},$$

and so on $\mathcal{E}_1 \cap \mathcal{E}_\infty$, (II.5) holds for any constant $c_* \leq c_1/3$. Setting $t = c_1 \theta / 3\mu$ in Proposition II.10, we obtain that for any fixed \mathbf{w} ,

$$\mathbb{P} \left[\left\| \nabla^2 g(\mathbf{w}; \mathbf{X}_0) - \mathbb{E} [\nabla^2 g(\mathbf{w}; \mathbf{X}_0)] \right\| \geq \frac{c_1 \theta}{3\mu} \right] \leq 4n \exp \left(-\frac{c_4 p \theta^2}{n^2} \right).$$

Taking a union bound, we obtain that

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

b) *Large gradient in region R_2 .*: Similarly, for the gradient quantity, for $\mathbf{w} \in R_2$, Proposition II.6 shows that

$$\mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \right] \geq c_6 \theta.$$

Moreover, on \mathcal{E}_∞ , $\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0) / \|\mathbf{w}\|$ is

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

Lipschitz by Proposition II.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

$$\mathcal{E}_2 = \left\{ \max_{\mathbf{w} \in N_2} \left| \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \right] \right| \leq \frac{c_6 \theta}{3} \right\}.$$

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

$$\sup_{\mathbf{w} \in R_2} \left| \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \right] \right| \leq \frac{2c_6 \theta}{3}, \quad (\text{IV.9})$$

and so on $\mathcal{E}_2 \cap \mathcal{E}_\infty$, (II.6) holds for any constant $c_* \leq c_6/3$. Setting $t = c_6 \theta / 3$ in Proposition II.8, we obtain that for any fixed $\mathbf{w} \in R_2$,

$$\mathbb{P} \left[\left| \frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla g(\mathbf{w}; \mathbf{X}_0)}{\|\mathbf{w}\|} \right] \right| \geq \frac{c_9 p \theta^2}{n} \right] \leq 2 \exp \left(-\frac{c_9 p \theta^2}{n} \right),$$

and so

$$\mathbb{P} [\mathcal{E}_2^c] \leq 2 \exp \left(-\frac{c_9 p \theta^2}{n} + n \log \left(\frac{C_8 n^2 \log(np)}{\theta \mu} \right) \right). \quad (\text{IV.10})$$

c) *Existence of negative curvature direction in R_3 .*: Finally, for any $\mathbf{w} \in R_3$, Proposition II.5 shows that

$$\mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \leq -c_9 \theta.$$

On \mathcal{E}_∞ , $\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w} / \|\mathbf{w}\|^2$ is

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

Lipschitz by Proposition II.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

$$\mathcal{E}_3 = \left\{ \max_{\mathbf{w} \in N_3} \left| \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \right| \leq \frac{c_9 \theta}{3} \right\}$$

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

$$\sup_{\mathbf{w} \in R_3} \left| \frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} - \mathbb{E} \left[\frac{\mathbf{w}^* \nabla^2 g(\mathbf{w}; \mathbf{X}_0) \mathbf{w}}{\|\mathbf{w}\|^2} \right] \right| \leq \frac{2c_9 \theta}{3},$$

and (II.7) holds with any constant $c_* < c_9/3$. Setting $t = c_9 \theta / 3$ in Proposition II.9 and taking a union bound, we obtain

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

d) *The unique local minimizer located near $\mathbf{0}$.*: Let \mathcal{E}_g be the event that the bounds (II.5)-(II.7) hold. On \mathcal{E}_g , the function g is $\frac{c_* \theta}{\mu}$ -strongly convex over $R_1 = \{\mathbf{w} : \|\mathbf{w}\| \leq \mu / (4\sqrt{2})\}$. This implies that f has at most one local minimum on R_1 . It also implies that for any $\mathbf{w} \in R_1$,

$$g(\mathbf{w}; \mathbf{X}_0) \geq g(\mathbf{0}; \mathbf{X}_0) + \langle \nabla g(\mathbf{0}; \mathbf{X}_0), \mathbf{w} \rangle + \frac{c\theta}{2\mu} \|\mathbf{w}\|^2 \geq g(\mathbf{0}; \mathbf{X}_0) - \|\mathbf{w}\| \|\nabla g(\mathbf{0}; \mathbf{X}_0)\| + \frac{c_* \theta}{2\mu} \|\mathbf{w}\|^2.$$

So, if $g(\mathbf{w}; \mathbf{X}_0) \leq g(\mathbf{0}; \mathbf{X}_0)$, we necessarily have

$$\|\mathbf{w}\| \leq \frac{2\mu}{c_* \theta} \|\nabla g(\mathbf{0}; \mathbf{X}_0)\|.$$

Suppose that

$$\|\nabla g(\mathbf{0}; \mathbf{X}_0)\| \leq \frac{c_* \theta}{32}. \quad (\text{IV.11})$$

Then $g(\mathbf{w}; \mathbf{X}_0) \leq g(\mathbf{0}; \mathbf{X}_0)$ implies that $\|\mathbf{w}\| \leq \mu/16$. By Wierstrass's theorem, $g(\mathbf{w}; \mathbf{X}_0)$ has at least one minimizer \mathbf{w}_* over the compact set $S = \{\mathbf{w} : \|\mathbf{w}\| \leq \mu/10\}$. By the above reasoning, $\|\mathbf{w}_*\| \leq \mu/16$, and hence \mathbf{w}_* does not lie on the boundary of S . This implies that \mathbf{w}_* is a local minimizer of g . Moreover, as above,

$$\|\mathbf{w}_*\| \leq \frac{2\mu}{c_* \theta} \|\nabla g(\mathbf{0}; \mathbf{X}_0)\|.$$

We now use the vector Bernstein inequality to show that with our choice of p , (IV.11) is satisfied w.h.p. Notice that

$$\nabla g(\mathbf{0}; \mathbf{X}_0) = \frac{1}{p} \sum_{k=1}^p \dot{h}_\mu(x_{0k}(n)) \bar{\mathbf{x}}_{0k},$$

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

$$\mathbb{E} \left[\left\| \dot{h}_\mu(x_{0k}(n)) \bar{\mathbf{x}}_{0k} \right\|^m \right] \leq \mathbb{E} [\|(\mathbf{x}_0)_k\|^m] \leq \mathbb{E}_{Z \sim \chi(n)} [Z^m] \leq m! n^{m/2},$$

where we have applied the moment estimate for the $\chi(n)$ distribution shown in Lemma A.7. Applying the vector Bernstein inequality in Corollary A.10 with $R = \sqrt{n}$ and $\sigma^2 = 2n$, we obtain

$$\mathbb{P}[\|\nabla g(\mathbf{0}; \mathbf{X}_0)\| \geq t] \leq 2(n+1) \exp\left(-\frac{pt^2}{4n+2\sqrt{nt}}\right)$$

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 $1 - 4np^{-10}$,

$$\|\nabla g(\mathbf{0}; \mathbf{X}_0)\| \leq C_3 \sqrt{\frac{n \log p}{p}}. \quad (\text{IV.12})$$

When $\frac{p}{\log p} \geq \frac{C_{14}n}{\theta^2}$, for appropriately large C_{14} , (IV.12) implies (IV.11). Summing up failure probabilities completes the proof. \square

C. Proofs for Section II-C and Theorem II.3

Proof. (of Lemma II.14) By the generative model,

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

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

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

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 $p \geq C_1 n^2 \log n$, $\left\| \frac{1}{p\theta} \mathbf{X}_0 \mathbf{X}_0^* - \mathbf{I} \right\| \leq 10 \sqrt{\frac{\theta n \log p}{p}}$ with probability at least $1 - p^{-8}$. Thus, when $p \geq C_2 \kappa^4(\mathbf{A}_0) \theta n^2 \log(n\theta\kappa(\mathbf{A}_0))$,

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

as desired. \square

Proof. (of Lemma II.15) Let $\tilde{\mathbf{Y}} \doteq \mathbf{X}_0 + \tilde{\Xi} \mathbf{X}_0$. Note the Jacobian matrix for the mapping $\mathbf{q}(\mathbf{w})$ is $\nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) = \begin{bmatrix} \mathbf{I}, -\mathbf{w}/\sqrt{1-\|\mathbf{w}\|^2} \end{bmatrix}$. 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

$$\begin{aligned}
 & \left\| \nabla_{\mathbf{w}} g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}} g(\mathbf{w}; \mathbf{X}_0) \right\| \\
 &= \left\| \frac{1}{p} \sum_{k=1}^p \dot{h}_\mu(\mathbf{q}^*(\mathbf{w}) \tilde{\mathbf{y}}_k) \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) \tilde{\mathbf{y}}_k - \frac{1}{p} \sum_{k=1}^p \dot{h}_\mu(\mathbf{q}^*(\mathbf{w}) (\mathbf{x}_0)_k) \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) (\mathbf{x}_0)_k \right\| \\
 &\leq \left\| \frac{1}{p} \sum_{k=1}^p \dot{h}_\mu(\mathbf{q}^* \tilde{\mathbf{y}}_k) \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) (\tilde{\mathbf{y}}_k - (\mathbf{x}_0)_k) \right\| \\
 &\quad + \left\| \frac{1}{p} \sum_{k=1}^p \left[\dot{h}_\mu(\mathbf{q}^*(\mathbf{w}) \tilde{\mathbf{y}}_k) - \dot{h}_\mu(\mathbf{q}^*(\mathbf{w}) (\mathbf{x}_0)_k) \right] \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) (\mathbf{x}_0)_k \right\| \\
 &\leq \left\| \tilde{\boldsymbol{\Xi}} \right\| \left(\max_t \dot{h}_\mu(t) 3n \|\mathbf{X}_0\|_\infty + L_{\dot{h}_\mu} 3n \|\mathbf{X}_0\|_\infty^2 \right),
 \end{aligned}$$

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

$$\begin{aligned}
 & \left\| \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \tilde{\mathbf{Y}}) - \nabla_{\mathbf{w}}^2 g(\mathbf{w}; \mathbf{X}_0) \right\| \\
 &\leq \left\| \frac{1}{p} \sum_{k=1}^p \left[\ddot{h}(\mathbf{q}^*(\mathbf{w}) \tilde{\mathbf{y}}_k) \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) \tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k^* (\nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}))^* - \ddot{h}(\mathbf{q}^*(\mathbf{w}) (\mathbf{x}_0)_k) \nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}) (\mathbf{x}_0)_k (\mathbf{x}_0)_k^* (\nabla_{\mathbf{w}} \mathbf{q}(\mathbf{w}))^* \right] \right\| \\
 &\quad + \left\| \frac{1}{p} \sum_{k=1}^p \left[\dot{h}(\mathbf{q}^*(\mathbf{w}) \tilde{\mathbf{y}}_k) \left(\frac{\mathbf{I}}{q_n(\mathbf{w})} + \frac{\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right) \tilde{\mathbf{y}}_k(n) - \dot{h}(\mathbf{q}^*(\mathbf{w}) (\mathbf{x}_0)_k) \left(\frac{\mathbf{I}}{q_n(\mathbf{w})} + \frac{\mathbf{w}\mathbf{w}^*}{q_n^3(\mathbf{w})} \right) (\mathbf{x}_0)_k(n) \right] \right\| \\
 &\leq \frac{45}{2} L_{\ddot{h}_\mu} n^{3/2} \|\mathbf{X}\|_\infty^3 \left\| \tilde{\boldsymbol{\Xi}} \right\| + \max_t \ddot{h}_\mu(t) \left(18n^{3/2} \|\mathbf{X}\|_\infty^2 \left\| \tilde{\boldsymbol{\Xi}} \right\| + 10n^2 \|\mathbf{X}\|_\infty^2 \left\| \tilde{\boldsymbol{\Xi}} \right\|^2 \right) \\
 &\quad + 3\sqrt{2} L_{\dot{h}_\mu} n^2 \left\| \tilde{\boldsymbol{\Xi}} \right\| \|\mathbf{X}\|_\infty^2 + \max_t \dot{h}_\mu(t) 2\sqrt{2}n^2 \left\| \tilde{\boldsymbol{\Xi}} \right\| \|\mathbf{X}\|_\infty,
 \end{aligned}$$

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 IV.12, $\|\mathbf{X}\|_\infty \leq 4\sqrt{\log(np)}$ with probability at least $1 - \theta(np)^{-7} - \exp(-0.3\theta np)$, we obtain

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

completing the proof. \square

Proof. (of Theorem II.3) Here c_* is as defined in Theorem II.1. By Lemma II.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 II.15, we obtain that for all $\mathbf{w} \in \Gamma$,

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

with probability at least $1 - p^{-8} - \theta (np)^{-7} - \exp(-0.3\theta np)$. In view of (II.13) in Theorem II.1, we have

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

By similar arguments, we obtain (II.11) through (II.13) in Theorem II.3.

To show the unique local minimizer over Γ is near $\mathbf{0}$, we note that (recall the last part of proof of Theorem II.1 in Section IV-B) $g(\mathbf{w}; \mathbf{X}_0 + \tilde{\Xi} \mathbf{X}_0)$ being $\frac{c_* \theta}{2\mu}$ strongly convex near $\mathbf{0}$ implies that

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

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 II.1. A simple union bound with careful bookkeeping gives the success probability. \square

APPENDIX A

TECHNICAL TOOLS AND BASIC FACTS USED IN PROOFS

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

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

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

the first two derivatives are

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

Also, for any $z \geq 0$, we have

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

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

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

Lemma A.2 (Harris' Inequality, [120], see also Theorem 2.15 of [121]). *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*

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] \geq \mathbb{E}[f(\mathbf{X})]\mathbb{E}[g(\mathbf{X})].$$

Similarly, if f is nondecreasing (nonincreasing) and g is nonincreasing (nondecreasing) coordinatewise in the above sense, we have

$$\mathbb{E}[f(\mathbf{X})g(\mathbf{X})] \leq \mathbb{E}[f(\mathbf{X})]\mathbb{E}[g(\mathbf{X})].$$

Lemma A.3 (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 ρ on \mathbb{R}^n and any $t \in (0, t_0)$ such that $\int_0^t \int_{\mathcal{U}} \left| \frac{\partial F(\mathbf{x}, s)}{\partial s} \right| \rho(d\mathbf{x}) ds < \infty$, it holds that*

$$\frac{d}{dt} \int_{\mathcal{U}} F(\mathbf{x}, t) \rho(d\mathbf{x}) = \int_{\mathcal{U}} \frac{\partial F(\mathbf{x}, t)}{\partial t} \rho(d\mathbf{x}), \text{ or } \frac{d}{dt} \mathbb{E}_{\mathbf{x}} [F(\mathbf{x}, t) \mathbb{1}_{\mathcal{U}}] = \mathbb{E}_{\mathbf{x}} \left[\frac{\partial F(\mathbf{x}, t)}{\partial t} \mathbb{1}_{\mathcal{U}} \right].$$

Proof. See proof of Lemma A.4 in the technical report [2]. □

Lemma A.4 (Gaussian Tail Estimates). *Let $X \sim \mathcal{N}(0, 1)$ and $\Phi(x)$ be CDF of X . For any $x \geq 0$, we have the following estimates for $\Phi^c(x) \doteq 1 - \Phi(x)$:*

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

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

Lemma A.5 (Moments of the Gaussian RV). *If $X \sim \mathcal{N}(0, \sigma^2)$, then it holds for all integer $p \geq 1$ that*

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

Lemma A.6 (Moments of the χ^2 RV). *If $X \sim \chi^2(n)$, then it holds for all integer $p \geq 1$ that*

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

Lemma A.7 (Moments of the χ RV). *If $X \sim \chi(n)$, then it holds for all integer $p \geq 1$ that*

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

Lemma A.8 (Moment-Control Bernstein's Inequality for Scalar RVs, Theorem 2.10 of [122]). *Let X_1, \dots, X_p be i.i.d. real-valued random variables. Suppose that there exist some positive numbers R and σ^2 such that*

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

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

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

Lemma A.9 (Moment-Control Bernstein's Inequality for Matrix RVs, Theorem 6.2 of [123]). *Let $\mathbf{X}_1, \dots, \mathbf{X}_p \in \mathbb{R}^{d \times d}$ be i.i.d. random, symmetric matrices. Suppose there exist some positive number R and σ^2 such that*

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

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

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

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

Corollary A.10 (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*

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

Let $\mathbf{s} = \frac{1}{p} \sum_{k=1}^p \mathbf{x}_k$, then for any $t > 0$, it holds that

$$\mathbb{P} [\|\mathbf{s} - \mathbb{E}[\mathbf{s}]\| \geq t] \leq 2(d+1) \exp\left(-\frac{pt^2}{2\sigma^2 + 2Rt}\right).$$

Proof. See proof of Lemma A.11 in the technical report [2]. \square

Lemma A.11 (Integral Form of Taylor's Theorem). *Let $f(\mathbf{x}) : \mathbb{R}^n \mapsto \mathbb{R}$ be a twice continuously differentiable function, then for any direction $\mathbf{y} \in \mathbb{R}^n$, we have*

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

APPENDIX 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*

$$\mathbb{E}[X \mathbb{1}_{X>0}] = \frac{\sigma_X}{\sqrt{2\pi}}, \tag{B.1}$$

$$\mathbb{E}[\exp(-aX) X \mathbb{1}_{X>0}] = \frac{\sigma_X}{\sqrt{2\pi}} - a\sigma_X^2 \exp\left(\frac{a^2\sigma_X^2}{2}\right) \Phi^c(a\sigma_X), \tag{B.2}$$

$$\mathbb{E}[\exp(-aX) \mathbb{1}_{X>0}] = \exp\left(\frac{a^2\sigma_X^2}{2}\right) \Phi^c(a\sigma_X), \tag{B.3}$$

$$\begin{aligned} \mathbb{E}[\exp(-a(X+Y)) X^2 \mathbb{1}_{X+Y>0}] &= \sigma_X^2 (1 + a^2\sigma_X^2) \exp\left(\frac{a^2\sigma_X^2 + a^2\sigma_Y^2}{2}\right) \Phi^c\left(a\sqrt{\sigma_X^2 + \sigma_Y^2}\right) \\ &\quad - \frac{a\sigma_X^4}{\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}}, \end{aligned} \tag{B.4}$$

$$\begin{aligned} \mathbb{E}[\exp(-a(X+Y)) XY \mathbb{1}_{X+Y>0}] &= a^2\sigma_X^2\sigma_Y^2 \exp\left(\frac{a^2\sigma_X^2 + a^2\sigma_Y^2}{2}\right) \Phi^c\left(a\sqrt{\sigma_X^2 + \sigma_Y^2}\right) \\ &\quad - \frac{a\sigma_X^2\sigma_Y^2}{\sqrt{2\pi}\sqrt{\sigma_X^2 + \sigma_Y^2}}, \end{aligned} \tag{B.5}$$

$$\mathbb{E}[\tanh(aX) X] = a\sigma_X^2 \mathbb{E}[1 - \tanh^2(aX)], \tag{B.6}$$

$$\mathbb{E}[\tanh(a(X+Y)) X] = a\sigma_X^2 \mathbb{E}[1 - \tanh^2(a(X+Y))]. \tag{B.7}$$

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. \square

Proof. (of Lemma IV.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

$$\|p - f\|_{L^1[0,1]} = \int_0^1 |p(t) - f(t)| dt = \int_0^1 \left[\frac{1}{(1+\beta t)^2} - \frac{1}{(1+t)^2} \right] dt = \frac{1-\beta}{2(1+\beta)} \leq \frac{1}{2\sqrt{T}}.$$

Moreover, we have

$$\|f - p\|_{L^\infty[0,1]} = \max_{t \in [0,1]} p(t) - f(t) = \max_{t \in [0,1]} \frac{t(1-\beta)(2+t(1+\beta))}{(1+t)^2(1+\beta t)^2} \leq 1-\beta = \frac{1}{\sqrt{T}}.$$

Finally, notice that

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{b_k}{(1+k)^3} &= \sum_{k=0}^{\infty} \frac{(-\beta)^k}{(1+k)^2} = \sum_{i=0}^{\infty} \left[\frac{\beta^{2i}}{(1+2i)^2} - \frac{\beta^{2i+1}}{(2i+2)^2} \right] \\ &= \sum_{i=0}^{\infty} \beta^{2i} \frac{(2i+2)^2 - \beta(2i+1)^2}{(2i+2)^2(2i+1)^2} > 0, \end{aligned}$$

where at the second equality we have grouped consecutive even-odd pair of summands. In addition, we have

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

which converges to 2 when $n \rightarrow \infty$, completing the proof. \square

Proof. (of Lemma IV.5) 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. \square

Proof. (of Lemma IV.12) 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

$$\mathbb{P} \left[|X_{ij}| > 4\sqrt{\log(np)} \right] \leq \theta \mathbb{P} \left[|V_{ij}| > 4\sqrt{\log(np)} \right] \leq \theta \exp(-8 \log(np)) = \theta(np)^{-8}.$$

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. \square

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*

$$\left\| (\mathbf{A} + \Delta)^{-1/2} - \mathbf{A}^{-1/2} \right\| \leq \frac{2 \|\mathbf{A}\|^{1/2} \|\Delta\|}{\sigma_{\min}^2(\mathbf{A})}. \quad (\text{B.8})$$

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

Lemma B.3. *For any $\theta \in (0, 1/2)$, $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with $\mathbf{X} \sim_{i.i.d.} \text{BG}(\theta)$ obeys*

$$\left\| \frac{1}{n_2\theta} \mathbf{X} \mathbf{X}^* - \mathbf{I} \right\| \leq 10 \sqrt{\frac{\theta n_1 \log n_2}{n_2}} \quad (\text{B.9})$$

with probability at least $1 - n_2^{-8}$, provided $n_2 > Cn_1^2 \log n_1$. Here $C > 0$ is a constant.

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.5 and Lemma A.6, and also $\theta \leq 1/2$. Taking $\sigma^2 = 3n_1\theta$ and $R = 2n_1$, and invoking the matrix Bernstein in Lemma A.9, we obtain

$$\mathbb{E} \left[\left\| \frac{1}{p\theta} \sum_{k=1}^p \mathbf{x}_k \mathbf{x}_k^* - \mathbf{I} \right\| > t \right] \leq \exp \left(-\frac{n_2 t^2}{6n_1\theta + 4n_1 t} + 2 \log n_1 \right) \quad (\text{B.10})$$

for any $t \geq 0$. Taking $t = 10\sqrt{\theta n_1 \log(n_2)/n_2}$ gives the claimed result. \square

ACKNOWLEDGMENT

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

REFERENCES

- [1] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery using nonconvex optimization," in *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015, pp. 2351–2360.
- [2] —, "Complete dictionary recovery over the sphere," *arXiv preprint arXiv:1504.06785*, 2015.
- [3] —, "Complete dictionary recovery over the sphere II: Recovery by Riemannian trust-region method," *arXiv preprint arXiv:1511.04777*, 2015.
- [4] R. A. DeVore, "Nonlinear approximation," *Acta numerica*, vol. 7, pp. 51–150, 1998.
- [5] V. N. Temlyakov, "Nonlinear methods of approximation," *Foundations of Computational Mathematics*, vol. 3, no. 1, pp. 33–107, 2003.
- [6] R. A. DeVore, "Nonlinear approximation and its applications," in *Multiscale, Nonlinear and Adaptive Approximation*. Springer, 2009, pp. 169–201.
- [7] E. J. Candès, "New ties between computational harmonic analysis and approximation theory," *Approximation Theory X*, pp. 87–153, 2002.
- [8] J. Ma and G. Plonka, "A review of curvelets and recent applications," *IEEE Signal Processing Magazine*, vol. 27, no. 2, pp. 118–133, 2010.
- [9] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *Information Theory, IEEE Transactions on*, vol. 44, no. 6, pp. 2435–2476, 1998.
- [10] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.
- [11] —, "Sparse coding with an overcomplete basis set: A strategy employed by v1?" *Vision research*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [12] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer, 2010.
- [13] J. Mairal, F. Bach, and J. Ponce, "Sparse modeling for image and vision processing," *Foundations and Trends in Computer Graphics and Vision*, vol. 8, no. 2-3, pp. 85–283, 2014. [Online]. Available: <http://dx.doi.org/10.1561/06000000058>
- [14] A. Maurer and M. Pontil, "K-dimensional coding schemes in hilbert spaces," *Information Theory, IEEE Transactions on*, vol. 56, no. 11, pp. 5839–5846, 2010.
- [15] D. Vainsencher, S. Mannor, and A. M. Bruckstein, "The sample complexity of dictionary learning," *Journal of Machine Learning Research*, vol. 12, no. 23, pp. 3259–3281, Nov. 2011.
- [16] N. Mehta and A. G. Gray, "Sparsity-based generalization bounds for predictive sparse coding," *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, vol. 28, no. 1, pp. 36–44, 2013.
- [17] R. Gribonval, R. Jenatton, F. Bach, M. Kleinstuber, and M. Seibert, "Sample complexity of dictionary learning and other matrix factorizations," *arXiv preprint arXiv:1312.3790*, 2013.
- [18] R. Gribonval and K. Schnass, "Dictionary identification - sparse matrix-factorization via ℓ^1 -minimization," *IEEE Transactions on Information Theory*, vol. 56, no. 7, pp. 3523–3539, 2010.
- [19] Q. Geng and J. 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>.
- [20] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *Information Theory, IEEE Transactions on*, vol. 60, no. 3, pp. 1711–1732, 2014.
- [21] S. Choudhary and U. Mitra, "Identifiability scaling laws in bilinear inverse problems," *arXiv preprint arXiv:1402.2637*, 2014.
- [22] A. S. Bandeira, C. Kennedy, and A. Singer, "Approximating the little grothendieck problem over the orthogonal and unitary groups," *arXiv preprint arXiv:1308.5207*, 2013.
- [23] J. Briët and O. Regev, "Tight hardness of the non-commutative grothendieck problem," *arXiv preprint arXiv:1412.4413*, 2014.
- [24] Y. Li, K. Lee, and Y. Bresler, "Identifiability and stability in blind deconvolution under minimal assumptions," *arXiv preprint arXiv*, vol. 1507, 2015.
- [25] M. Kech and F. Krahmer, "Optimal injectivity conditions for bilinear inverse problems with applications to identifiability of deconvolution problems," *arXiv preprint arXiv:1603.07316*, 2016.

- [26] D. Donoho and J. 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*, vol. 367, no. 1906, pp. 4273–4293, 2009.
- [27] S. Oymak and B. Hassibi, "New null space results and recovery thresholds for matrix rank minimization," *arXiv preprint arXiv:1011.6326*, 2010.
- [28] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM (JACM)*, vol. 58, no. 3, p. 11, 2011.
- [29] D. L. Donoho, M. Gavish, and A. Montanari, "The phase transition of matrix recovery from gaussian measurements matches the minimax mse of matrix denoising," *Proceedings of the National Academy of Sciences*, vol. 110, no. 21, pp. 8405–8410, 2013.
- [30] M. B. McCoy and J. A. Tropp, "Sharp recovery bounds for convex demixing, with applications," *Foundations of Computational Mathematics*, vol. 14, no. 3, pp. 503–567, 2014.
- [31] C. Mu, B. Huang, J. Wright, and D. Goldfarb, "Square deal: Lower bounds and improved relaxations for tensor recovery," *arXiv preprint arXiv:1307.5870*, 2013.
- [32] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational mathematics*, vol. 12, no. 6, pp. 805–849, 2012.
- [33] E. J. Candès, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [34] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, "Living on the edge: Phase transitions in convex programs with random data," *Information and Inference*, p. iau005, 2014.
- [35] E. J. Candès, "Mathematics of sparsity (and few other things)," in *Proceedings of the International Congress of Mathematicians, Seoul, South Korea*, 2014.
- [36] C. Bao, J.-F. Cai, and H. Ji, "Fast sparsity-based orthogonal dictionary learning for image restoration," in *Computer Vision (ICCV), 2013 IEEE International Conference on*. IEEE, 2013, pp. 3384–3391.
- [37] S. Lesage, R. Gribonval, F. Bimbot, and L. Benaroya, "Learning unions of orthonormal bases with thresholded singular value decomposition," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5. IEEE, 2005, pp. v–293.
- [38] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice hall Englewood Cliffs, NJ, 1989, vol. 23.
- [39] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of optimization theory and applications*, vol. 109, no. 3, pp. 475–494, 2001.
- [40] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran, "Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the kurdyka-lojasiewicz inequality," *Mathematics of Operations Research*, vol. 35, no. 2, pp. 438–457, 2010.
- [41] J. Bolte, S. Sabach, and M. Teboulle, "Proximal alternating linearized minimization for nonconvex and nonsmooth problems," *Mathematical Programming*, vol. 146, no. 1–2, pp. 459–494, 2014.
- [42] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proceedings of the 25th Annual Conference on Learning Theory*, 2012.
- [43] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning sparsely used overcomplete dictionaries via alternating minimization," *arXiv preprint arXiv:1310.7991*, 2013.
- [44] S. Arora, R. Ge, and A. Moitra, "New algorithms for learning incoherent and overcomplete dictionaries," *arXiv preprint arXiv:1308.6273*, 2013.
- [45] A. Agarwal, A. Anandkumar, and P. Netrapalli, "Exact recovery of sparsely used overcomplete dictionaries," *arXiv preprint arXiv:1309.1952*, 2013.
- [46] S. Arora, A. Bhaskara, R. Ge, and T. Ma, "More algorithms for provable dictionary learning," *arXiv preprint arXiv:1401.0579*, 2014.
- [47] S. Arora, R. Ge, T. Ma, and A. Moitra, "Simple, efficient, and neural algorithms for sparse coding," *arXiv preprint arXiv:1503.00778*, 2015.
- [48] L. Demanet and P. Hand, "Scaling law for recovering the sparsest element in a subspace," *Information and Inference*, vol. 3, no. 4, pp. 295–309, 2014.
- [49] M. Zibulevsky and B. Pearlmutter, "Blind source separation by sparse decomposition in a signal dictionary," *Neural computation*, vol. 13, no. 4, pp. 863–882, 2001.
- [50] Q. Qu, J. Sun, and J. Wright, "Finding a sparse vector in a subspace: Linear sparsity using alternating directions," in *Advances in Neural Information Processing Systems*, 2014, pp. 3401–3409.
- [51] J. Sun, Q. Qu, and J. Wright, "When are nonconvex problems not scary?" *arXiv preprint arXiv:1510.06096*, 2015.
- [52] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—online stochastic gradient for tensor decomposition," in *Proceedings of The 28th Conference on Learning Theory*, 2015, pp. 797–842.
- [53] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on Riemannian manifolds," *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 303–330, 2007.
- [54] P.-A. Absil, R. Mahoney, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [55] D. Goldfarb, "Curvilinear path steplength algorithms for minimization which use directions of negative curvature," *Mathematical programming*, vol. 18, no. 1, pp. 31–40, 1980.
- [56] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear algebra and its applications*, vol. 416, no. 1, pp. 48–67, 2006.
- [57] C. Hillar and F. T. Sommer, "When can dictionary learning uniquely recover sparse data from subsamples?" *arXiv preprint arXiv:1106.3616*, 2011.
- [58] K. Schnass, "Local identification of overcomplete dictionaries," *arXiv preprint arXiv:1401.6354*, 2014.
- [59] —, "On the identifiability of overcomplete dictionaries via the minimisation principle underlying k-svd," *Applied and Computational Harmonic Analysis*, vol. 37, no. 3, pp. 464–491, 2014.

- [60] —, “Convergence radius and sample complexity of itkm algorithms for dictionary learning,” *arXiv preprint arXiv:1503.07027*, 2015.
- [61] B. Barak, J. A. Kelner, and D. Steurer, “Dictionary learning and tensor decomposition via the sum-of-squares method,” *arXiv preprint arXiv:1407.1543*, 2014.
- [62] S. Wu and B. Yu, “Local identifiability of ℓ_1 -minimization dictionary learning: a sufficient and almost necessary condition,” *arXiv preprint arXiv:1505.04363*, 2015.
- [63] R. Gribonval, R. Jenatton, and F. Bach, “Sparse and spurious: dictionary learning with noise and outliers,” *arXiv preprint arXiv:1407.5155*, 2014.
- [64] L.-A. Gottlieb and T. Neylon, “Matrix sparsification and the sparse null space problem,” in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2010, pp. 205–218.
- [65] P.-L. Loh and M. J. Wainwright, “High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity,” in *Advances in Neural Information Processing Systems*, 2011, pp. 2726–2734.
- [66] —, “Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima,” in *Advances in Neural Information Processing Systems*, 2013, pp. 476–484.
- [67] Z. Wang, H. Lu, and H. Liu, “Nonconvex statistical optimization: minimax-optimal sparse pca in polynomial time,” *arXiv preprint arXiv:1408.5352*, 2014.
- [68] S. Balakrishnan, M. J. Wainwright, and B. Yu, “Statistical guarantees for the em algorithm: From population to sample-based analysis,” *arXiv preprint arXiv:1408.2156*, 2014.
- [69] Z. Wang, Q. Gu, Y. Ning, and H. Liu, “High dimensional expectation-maximization algorithm: Statistical optimization and asymptotic normality,” *arXiv preprint arXiv:1412.8729*, 2014.
- [70] P.-L. Loh and M. J. Wainwright, “Support recovery without incoherence: A case for nonconvex regularization,” *arXiv preprint arXiv:1412.5632*, 2014.
- [71] P.-L. Loh, “Statistical consistency and asymptotic normality for high-dimensional robust m-estimators,” *arXiv preprint arXiv:1501.00312*, 2015.
- [72] W. Sun, J. Lu, H. Liu, and G. Cheng, “Provable sparse tensor decomposition,” *arXiv preprint arXiv:1502.01425*, 2015.
- [73] R. H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” *Information Theory, IEEE Transactions on*, vol. 56, no. 6, pp. 2980–2998, 2010.
- [74] P. Jain, P. Netrapalli, and S. Sanghavi, “Low-rank matrix completion using alternating minimization,” in *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*. ACM, 2013, pp. 665–674.
- [75] M. Hardt, “Understanding alternating minimization for matrix completion,” in *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 2014, pp. 651–660.
- [76] M. Hardt and M. Wootters, “Fast matrix completion without the condition number,” in *Proceedings of The 27th Conference on Learning Theory*, 2014, pp. 638–678.
- [77] P. Netrapalli, U. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain, “Non-convex robust pca,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1107–1115.
- [78] P. Jain and P. Netrapalli, “Fast exact matrix completion with finite samples,” *arXiv preprint arXiv:1411.1087*, 2014.
- [79] R. Sun and Z.-Q. Luo, “Guaranteed matrix completion via non-convex factorization,” *arXiv preprint arXiv:1411.8003*, 2014.
- [80] Q. Zheng and J. Lafferty, “A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements,” *arXiv preprint arXiv:1506.06081*, 2015.
- [81] S. Tu, R. Boczar, M. Soltanolkotabi, and B. Recht, “Low-rank solutions of linear matrix equations via procrustes flow,” *arXiv preprint arXiv:1507.03566*, 2015.
- [82] Y. Chen and M. J. Wainwright, “Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees,” *arXiv preprint arXiv:1509.03025*, 2015.
- [83] P. Netrapalli, P. Jain, and S. Sanghavi, “Phase retrieval using alternating minimization,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2796–2804.
- [84] E. Candès, X. Li, and M. Soltanolkotabi, “Phase retrieval via wirtinger flow: Theory and algorithms,” *Information Theory, IEEE Transactions on*, vol. 61, no. 4, pp. 1985–2007, April 2015.
- [85] Y. Chen and E. J. Candès, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” *arXiv preprint arXiv:1505.05114*, 2015.
- [86] C. D. White, R. Ward, and S. Sanghavi, “The local convexity of solving quadratic equations,” *arXiv preprint arXiv:1506.07868*, 2015.
- [87] P. Jain and S. Oh, “Provable tensor factorization with missing data,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1431–1439.
- [88] A. Anandkumar, R. Ge, and M. Janzamin, “Guaranteed non-orthogonal tensor decomposition via alternating rank-1 updates,” *arXiv preprint arXiv:1402.5180*, 2014.
- [89] —, “Analyzing tensor power method dynamics: Applications to learning overcomplete latent variable models,” *arXiv preprint arXiv:1411.1488*, 2014.
- [90] A. Anandkumar, P. Jain, Y. Shi, and U. Niranjan, “Tensor vs matrix methods: Robust tensor decomposition under block sparse perturbations,” *arXiv preprint arXiv:1510.04747*, 2015.
- [91] X. Yi, C. Caramanis, and S. Sanghavi, “Alternating minimization for mixed linear regression,” *arXiv preprint arXiv:1310.3745*, 2013.
- [92] K. Lee, Y. Wu, and Y. Bresler, “Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization,” *arXiv preprint arXiv:1312.0525*, 2013.
- [93] P. Jain, C. Jin, S. M. Kakade, and P. Netrapalli, “Computing matrix squareroot via non convex local search,” *arXiv preprint arXiv:1507.05854*, 2015.
- [94] S. Bhojanapalli, A. Kyrillidis, and S. Sanghavi, “Dropping convexity for faster semi-definite optimization,” *arXiv preprint arXiv:1509.03917*, 2015.
- [95] C. Bao, H. Ji, Y. Quan, and Z. Shen, “L0 norm based dictionary learning by proximal methods with global convergence,” in *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. IEEE, 2014, pp. 3858–3865.

- [96] C. Bao, Y. Quan, and H. Ji, "A convergent incoherent dictionary learning algorithm for sparse coding," in *Computer Vision–ECCV 2014*. Springer, 2014, pp. 302–316.
- [97] C. Bao, H. Ji, and Z. Shen, "Convergence analysis for iterative data-driven tight frame construction scheme," *Applied and Computational Harmonic Analysis*, 2014.
- [98] S. Arora, R. Ge, R. Kannan, and A. Moitra, "Computing a nonnegative matrix factorization—provably," in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 2012, pp. 145–162.
- [99] S. Arora, A. Bhaskara, R. Ge, and T. Ma, "Provable bounds for learning some deep representations," *arXiv preprint arXiv:1310.6343*, 2013.
- [100] H. Sedghi and A. Anandkumar, "Provable tensor methods for learning mixtures of classifiers," *arXiv preprint arXiv:1412.3046*, 2014.
- [101] B. Neyshabur and R. Panigrahy, "Sparse matrix factorization," *arXiv preprint arXiv:1311.3315*, 2013.
- [102] R. Livni, S. Shalev-Shwartz, and O. Shamir, "On the computational efficiency of training neural networks," in *Advances in Neural Information Processing Systems*, 2014, pp. 855–863.
- [103] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," *Neural networks*, vol. 13, no. 4, pp. 411–430, 2000.
- [104] A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis*. John Wiley and Sons., 2001.
- [105] P. Comon, "Independent component analysis, a new concept?" *Signal processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [106] A. Frieze, M. Jerrum, and R. Kannan, "Learning linear transformations," in *focs*. IEEE, 1996, p. 359.
- [107] S. Arora, R. Ge, A. Moitra, and S. Sachdeva, "Provable ICA with unknown gaussian noise, with implications for gaussian mixtures and autoencoders," in *Advances in Neural Information Processing Systems*, 2012, pp. 2375–2383.
- [108] A. Hyvarinen, "Fast and robust fixed-point algorithms for independent component analysis," *IEEE Trans. Neural Networks*, vol. 10, no. 3, pp. 626–634, 1999.
- [109] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," *arXiv preprint arXiv:1602.06664*, 2016.
- [110] S. Bhojanapalli, B. Neyshabur, and N. Srebro, "Global optimality of local search for low rank matrix recovery," *arXiv preprint arXiv:1605.07221*, 2016.
- [111] R. Ge, J. D. Lee, and T. Ma, "Matrix completion has no spurious local minimum," *arXiv preprint arXiv:1605.07272*, 2016.
- [112] N. Boumal, V. Voroninski, and A. S. Bandeira, "The non-convex burer-monteiro approach works on smooth semidefinite programs," *arXiv preprint arXiv:1606.04970*, 2016.
- [113] N. Boumal, "Nonconvex phase synchronization," *arXiv preprint arXiv:1601.06114*, 2016.
- [114] A. S. Bandeira, N. Boumal, and V. Voroninski, "On the low-rank approach for semidefinite programs arising in synchronization and community detection," *arXiv preprint arXiv:1602.04426*, 2016.
- [115] P. Baldi and K. Hornik, "Neural networks and principal component analysis: Learning from examples without local minima," *Neural networks*, vol. 2, no. 1, pp. 53–58, 1989.
- [116] K. Kawaguchi, "Deep learning without poor local minima," *arXiv preprint arXiv:1605.07110*, 2016.
- [117] D. Soudry and Y. Carmon, "No bad local minima: Data independent training error guarantees for multilayer neural networks," *arXiv preprint arXiv:1605.08361*, 2016.
- [118] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998.
- [119] G. B. Folland, *Real Analysis: Modern Techniques and Their Applications*, 2nd ed. John Wiley & Sons, 1999.
- [120] T. E. Harris, "A lower bound for the critical probability in a certain percolation process," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 56, no. 01. Cambridge Univ Press, 1960, pp. 13–20.
- [121] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [122] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Springer, 2013.
- [123] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Foundations of Computational Mathematics*, vol. 12, no. 4, pp. 389–434, 2012.