

Recovery Guarantees for One-hidden-layer Neural Networks*

Kai Zhong [†] zhongkai@ices.utexas.edu UT-Austin	Zhao Song [‡] zhaos@utexas.edu UT-Austin	Prateek Jain prajain@microsoft.com Microsoft Research, India
Peter L. Bartlett [§] bartlett@cs.berkeley.edu UC, Berkeley	Inderjit S. Dhillon [¶] inderjit@cs.utexas.edu UT-Austin	

Abstract

In this paper, we consider regression problems with one-hidden-layer neural networks (1NNs). We distill some properties of activation functions that lead to *local strong convexity* in the neighborhood of the ground-truth parameters for the 1NN squared-loss objective and most popular nonlinear activation functions satisfy the distilled properties, including rectified linear units (ReLU), leaky ReLUs, squared ReLUs and sigmoids. For activation functions that are also smooth, we show *local linear convergence* guarantees of gradient descent under a resampling rule. For homogeneous activations, we show tensor methods are able to initialize the parameters to fall into the local strong convexity region. As a result, tensor initialization followed by gradient descent is guaranteed to recover the ground truth with sample complexity $d \cdot \log(1/\epsilon) \cdot \text{poly}(k, \lambda)$ and computational complexity $n \cdot d \cdot \text{poly}(k, \lambda)$ for smooth homogeneous activations with high probability, where d is the dimension of the input, k ($k \leq d$) is the number of hidden nodes, λ is a conditioning property of the ground-truth parameter matrix between the input layer and the hidden layer, ϵ is the targeted precision and n is the number of samples. To the best of our knowledge, this is the first work that provides recovery guarantees for 1NNs with both sample complexity and computational complexity *linear* in the input dimension and *logarithmic* in the precision.

*A preliminary version of this paper appears in Proceedings of the Thirty-fourth International Conference on Machine Learning (ICML 2017).

[†]Supported in part by NSF grants CCF-1320746, IIS-1546452 and CCF-1564000, and part of the work was done while interning in Microsoft research, India.

[‡]Supported in part by UTCS TAsip (CS361 Spring 17 Introduction to Computer Security).

[§]Supported in part by Australian Research Council through an Australian Laureate Fellowship (FL110100281) and through the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), and NSF grants IIS-1619362.

[¶]Supported in part by NSF grants CCF-1320746, IIS-1546452 and CCF-1564000.

Contents

1	Introduction	3
2	Related Work	4
2.1	Expressive Power	4
2.2	Achievability of Global Optima	4
2.3	Generalization Bound / Recovery Guarantees	5
3	Problem Formulation	6
4	Positive Definiteness of Hessian	7
5	Tensor Methods for Initialization	9
5.1	Preliminary	9
5.2	Algorithm	10
5.3	Theoretical Analysis	11
6	Global Convergence	12
7	Numerical Experiments	12
8	Conclusion	13
	References	14
A	Notation	19
B	Preliminaries	19
B.1	Useful Facts	19
B.2	Matrix Bernstein	22
C	Properties of Activation Functions	25
D	Local Positive Definiteness of Hessian	26
D.1	Main Results for Positive Definiteness of Hessian	26
D.1.1	Bounding the Spectrum of the Hessian near the Ground Truth	26
D.1.2	Local Linear Convergence of Gradient Descent	27
D.2	Positive Definiteness of Population Hessian at the Ground Truth	28
D.2.1	Lower Bound on the Eigenvalues of Hessian for the Orthogonal Case	29
D.2.2	Lower Bound on the Eigenvalues of Hessian for Non-orthogonal Case	33
D.2.3	Upper Bound on the Eigenvalues of Hessian for Non-orthogonal Case	37
D.3	Error Bound of Hessians near the Ground Truth for Smooth Activations	38
D.3.1	Second-order Smoothness near the Ground Truth for Smooth Activations	38
D.3.2	Empirical and Population Difference for Smooth Activations	42
D.4	Error Bound of Hessians near the Ground Truth for Non-smooth Activations	48
D.5	Positive Definiteness for a Small Region	51

E	Tensor Methods	55
E.1	Tensor Initialization Algorithm	55
E.2	Main Result for Tensor Methods	57
E.3	Error Bound for the Subspace Spanned by the Weight Matrix	59
	E.3.1 Error Bound for the Second-order Moment in Different Cases	59
	E.3.2 Error Bound for the Second-order Moment	63
	E.3.3 Subspace Estimation Using Power Method	63
E.4	Error Bound for the Reduced Third-order Moment	66
	E.4.1 Error Bound for the Reduced Third-order Moment in Different Cases	66
	E.4.2 Final Error Bound for the Reduced Third-order Moment	70
E.5	Error Bound for the Magnitude and Sign of the Weight Vectors	71
	E.5.1 Robustness for Solving Linear Systems	71
	E.5.2 Error Bound for the First-order Moment	71
	E.5.3 Linear System for the First-order Moment	73
	E.5.4 Linear System for the Second-order Moment	74
F	Acknowledgments	75

1 Introduction

Neural Networks (NNs) have achieved great practical success recently. Many theoretical contributions have been made very recently to understand the extraordinary performance of NNs. The remarkable results of NNs on complex tasks in computer vision and natural language processing inspired works on the expressive power of NNs [CSS16, CS16, RPK⁺16, DFS16, PLR⁺16, MPCB14, Tel16]. Indeed, several works found NNs are very powerful and the deeper the more powerful. However, due to the high non-convexity of NNs, knowing the expressivity of NNs doesn't guarantee that the targeted functions will be learned. Therefore, several other works focused on the achievability of global optima. Many of them considered the over-parameterized setting, where the global optima or local minima close to the global optima will be achieved when the number of parameters is large enough, including [FB16, HV15, LSSS14, DPG⁺14, SS16, HM17]. This, however, leads to overfitting easily and can't provide any generalization guarantees, which are actually the essential goal in most tasks.

A few works have considered generalization performance. For example, [XLS17] provide generalization bound under the Rademacher generalization analysis framework. Recently [ZBH⁺17] describe some experiments showing that NNs are complex enough that they actually memorize the training data but still generalize well. As they claim, this cannot be explained by applying generalization analysis techniques, like VC dimension and Rademacher complexity, to classification loss (although it does not rule out a margins analysis—see, for example, [Bar98]; their experiments involve the unbounded cross-entropy loss).

In this paper, we don't develop a new generalization analysis. Instead we focus on parameter recovery setting, where we assume there are underlying ground-truth parameters and we provide recovery guarantees for the ground-truth parameters up to equivalent permutations. Since the parameters are exactly recovered, the generalization performance will also be guaranteed.

Several other techniques are also provided to recover the parameters or to guarantee generalization performance, such as tensor methods [JSA15] and kernel methods [AGMR17]. These methods require sample complexity $O(d^3)$ or computational complexity $\tilde{O}(n^2)$, which can be intractable in practice.

Recently [Sha16] show that neither specific assumptions on the niceness of the input distribution or niceness of the target function alone is sufficient to guarantee learnability using gradient-based methods. In this paper, we assume data points are sampled from Gaussian distribution and the parameters of hidden neurons are linearly independent.

Our main contributions are as follows,

1. We distill some properties for activation functions, which are satisfied by a wide range of activations, including ReLU, squared ReLU, sigmoid and tanh. With these properties we show positive definiteness (PD) of the Hessian in the neighborhood of the ground-truth parameters given enough samples (Theorem 4.2). Further, for activations that are also smooth, we show local linear convergence is guaranteed using gradient descent.

2. We propose a tensor method to initialize the parameters such that the initialized parameters fall into the local positive definiteness area. Our contribution is that we reduce the sample/computational complexity from cubic dependency on dimension to linear dependency (Theorem 5.6).

3. Combining the above two results, we provide a globally converging algorithm (Algorithm 2) for smooth homogeneous activations satisfying the distilled properties. The whole procedure requires sample/computational complexity linear in dimension and logarithmic in precision (Theorem 6.1).

2 Related Work

The recent empirical success of NNs has boosted their theoretical analyses [FZK⁺16, Bal16, BMBY16, SBL16, APVZ14, AGMR17, GKKT17]. In this paper, we classify them into three main directions.

2.1 Expressive Power

Expressive power is studied to understand the remarkable performance of neural networks on complex tasks. Although one-hidden-layer neural networks with sufficiently many hidden nodes can approximate any continuous function [Hor91], shallow networks can't achieve the same performance in practice as deep networks. Theoretically, several recent works show the depth of NNs plays an essential role in the expressive power of neural networks [DFS16]. As shown in [CSS16, CS16, Tel16], functions that can be implemented by a deep network of polynomial size require exponential size in order to be implemented by a shallow network. [RPK⁺16, PLR⁺16, MPCB14, AGMR17] design some measures of expressivity that display an exponential dependence on the depth of the network. However, the increasing of the expressivity of NNs or its depth also increases the difficulty of the learning process to achieve a good enough model. In this paper, we focus on 1NNs and provide recovery guarantees using a finite number of samples.

2.2 Achievability of Global Optima

The global convergence is in general not guaranteed for NNs due to their non-convexity. It is widely believed that training deep models using gradient-based methods works so well because the error surface either has no local minima, or if they exist they need to be close in value to the global minima. [SCP16] present examples showing that for this to be true additional assumptions on the data, initialization schemes and/or the model classes have to be made. Indeed the achievability of global optima has been shown under many different types of assumptions.

In particular, [CHM⁺15] analyze the loss surface of a special random neural network through spin-glass theory and show that it has exponentially many local optima, whose loss is small and close to that of a global optimum. Later on, [Kaw16] eliminate some assumptions made by [CHM⁺15] but still require the independence of activations as [CHM⁺15], which is unrealistic. [SS16] study the geometric structure of the neural network objective function. They have shown that with high probability random initialization will fall into a basin with a small objective value when the network is over-parameterized. [LSSS14] consider polynomial networks where the activations are square functions, which are typically not used in practice. [HV15] show that when a local minimum has zero parameters related to a hidden node, a global optimum is achieved. [FB16] study the landscape of 1NN in terms of topology and geometry, and show that the level set becomes connected as the network is increasingly over-parameterized. [HM17] show that products of matrices don't have spurious local minima and that deep residual networks can represent any function on a sample, as long as the number of parameters is larger than the sample size. [SC16] consider over-specified NNs, where the number of samples is smaller than the number of weights. [DPG⁺14] propose a new approach to second-order optimization that identifies and attacks the saddle point problem in high-dimensional non-convex optimization. They apply the approach to recurrent neural networks and show practical performance. [AGMR17] use results from tropical geometry to show global optimality of an algorithm, but it requires $(2n)^k$ poly(n) computational complexity.

Almost all of these results require the number of parameters is larger than the number of points, which probably overfits the model and no generalization performance will be guaranteed. In this paper, we propose an efficient and provable algorithm for 1NNs that can achieve the underlying ground-truth parameters.

2.3 Generalization Bound / Recovery Guarantees

The achievability of global optima of the objective from the training data doesn't guarantee the learned model to be able to generalize well on unseen testing data. In the literature, we find three main approaches to generalization guarantees.

1) *Use generalization analysis frameworks*, including VC dimension/Rademacher complexity, to bound the generalization error. A few works have studied the generalization performance for NNs. [XLS17] follow [SC16] but additionally provide generalization bounds using Rademacher complexity. They assume the obtained parameters are in a regularization set so that the generalization performance is guaranteed, but this assumption can't be justified theoretically. [HRS16] apply stability analysis to the generalization analysis of SGD for convex and non-convex problems, arguing early stopping is important for generalization performance.

2) *Assume an underlying model and try to recover this model*. This direction is popular for many non-convex problems including matrix completion/sensing [JNS13, Har14, SL15, BLWZ17], mixed linear regression [ZJD16], subspace recovery [EV09] and other latent models [AGH⁺14].

Without making any assumptions, those non-convex problems are intractable [AGKM12, GV15, SWZ17a, GG11, RSW16, SR11, HM13, AGM12, YCS14]. Recovery guarantees for NNs also need assumptions. Several different approaches under different assumptions are provided to have recovery guarantees on different NN settings.

Tensor methods [AGH⁺14, WTSA15, WA16, SWZ16] are a general tool for recovering models with latent factors by assuming the data distribution is known. Some existing recovery guarantees for NNs are provided by tensor methods [SA15, JSA15]. However, [SA15] only provide guarantees to recover the subspace spanned by the weight matrix and no sample complexity is given, while [JSA15] require $O(d^3/\epsilon^2)$ sample complexity. In this paper, we use tensor methods as an initialization step so that we don't need very accurate estimation of the moments, which enables us to reduce the total sample complexity from $1/\epsilon^2$ to $\log(1/\epsilon)$.

[ABGM14] provide polynomial sample complexity and computational complexity bounds for learning deep representations in unsupervised setting, and they need to assume the weights are sparse and randomly distributed in $[-1, 1]$.

[Tia17] analyze 1NN by assuming Gaussian inputs in a supervised setting, in particular, regression and classification with a teacher. This paper also considers this setting. However, there are some key differences. a) [Tia17] require the second-layer parameters are all ones, while we can learn these parameters. b) In [Tia17], the ground-truth first-layer weight vectors are required to be orthogonal, while we only require linear independence. c) [Tia17] require a good initialization but doesn't provide initialization methods, while we show the parameters can be efficiently initialized by tensor methods. d) In [Tia17], only the population case (infinite sample size) is considered, so there is no sample complexity analysis, while we show finite sample complexity.

Recovery guarantees for convolution neural network with Gaussian inputs are provided in [BG17], where they show a globally converging guarantee of gradient descent on a one-hidden-layer no-overlap convolution neural network. However, they consider population case, so no sample complexity is provided. Also their analysis depends on ReLU activations and the no-overlap case is very unlikely to be used in practice. In this paper, we consider a large range of activation functions, but for one-hidden-layer fully-connected NNs.

3) *Improper Learning*. In the improper learning setting for NNs, the learning algorithm is not restricted to output a NN, but only should output a prediction function whose error is not much larger than the error of the best NN among all the NNs considered. [ZLJ16, ZLW16] propose kernel methods to learn the prediction function which is guaranteed to have generalization performance close to that of the NN. However, the sample complexity and computational complexity are expo-

nential. [AZS14] transform NNs to convex semi-definite programming. The works by [Bac14] and [BRV⁺05] are also in this direction. However, these methods are actually not learning the original NNs. Another work by [ZLWJ17] uses random initializations to achieve arbitrary small excess risk. However, their algorithm has exponential running time in $1/\epsilon$.

Roadmap. The paper is organized as follows. In Section 3, we present our problem setting and show three key properties of activations required for our guarantees. In Section 4, we introduce the formal theorem of local strong convexity and show local linear convergence for smooth activations. Section 5 presents a tensor method to initialize the parameters so that they fall into the basin of the local strong convexity region.

3 Problem Formulation

We consider the following regression problem. Given a set of n samples

$$S = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\} \subset \mathbb{R}^d \times \mathbb{R},$$

let \mathcal{D} denote a underlying distribution over $\mathbb{R}^d \times \mathbb{R}$ with parameters

$$\{w_1^*, w_2^*, \dots, w_k^*\} \subset \mathbb{R}^d, \text{ and } \{v_1^*, v_2^*, \dots, v_k^*\} \subset \mathbb{R}$$

such that each sample $(x, y) \in S$ is sampled i.i.d. from this distribution, with

$$\mathcal{D} : \quad x \sim \mathcal{N}(0, I), \quad y = \sum_{i=1}^k v_i^* \cdot \phi(w_i^{*\top} x), \quad (1)$$

where $\phi(z)$ is the activation function, k is the number of nodes in the hidden layer. The main question we want to answer is: How many samples are sufficient to recover the underlying parameters?

It is well-known that, training one hidden layer neural network is NP-complete [BR88]. Thus, without making any assumptions, learning deep neural network is intractable. Throughout the paper, we assume x follows a standard normal distribution; the data is noiseless; the dimension of input data is at least the number of hidden nodes; and activation function $\phi(z)$ satisfies some reasonable properties.

Actually our results can be easily extended to multivariate Gaussian distribution with positive definite covariance and zero mean since we can estimate the covariance first and then transform the input to a standard normal distribution but with some loss of accuracy. Although this paper focuses on the regression problem, we can transform classification problems to regression problems if a good teacher is provided as described in [Tia17]. Our analysis requires k to be no greater than d , since the first-layer parameters will be linearly dependent otherwise.

For activation function $\phi(z)$, we assume it is continuous and if it is non-smooth let its first derivative be left derivative. Furthermore, we assume it satisfies Property 3.1, 3.2, and 3.3. These properties are critical for the later analyses. We also observe that most activation functions actually satisfy these three properties.

Property 3.1. *The first derivative $\phi'(z)$ is nonnegative and homogeneously bounded, i.e., $0 \leq \phi'(z) \leq L_1|z|^p$ for some constants $L_1 > 0$ and $p \geq 0$.*

Property 3.2. *Let $\alpha_q(\sigma) = \mathbb{E}_{z \sim \mathcal{N}(0,1)}[\phi'(\sigma \cdot z)z^q]$, $\forall q \in \{0, 1, 2\}$, and $\beta_q(\sigma) = \mathbb{E}_{z \sim \mathcal{N}(0,1)}[\phi'^2(\sigma \cdot z)z^q]$, $\forall q \in \{0, 2\}$. Let $\rho(\sigma)$ denote $\min\{\beta_0(\sigma) - \alpha_0^2(\sigma) - \alpha_1^2(\sigma), \beta_2(\sigma) - \alpha_1^2(\sigma) - \alpha_2^2(\sigma), \alpha_0(\sigma) \cdot \alpha_2(\sigma) - \alpha_1^2(\sigma)\}$ The first derivative $\phi'(z)$ satisfies that, for all $\sigma > 0$, we have $\rho(\sigma) > 0$.*

Property 3.3. The second derivative $\phi''(z)$ is either (a) globally bounded $|\phi''(z)| \leq L_2$ for some constant L_2 , i.e., $\phi(z)$ is L_2 -smooth, or (b) $\phi''(z) = 0$ except for e (e is a finite constant) points.

Remark 3.4. The first two properties are related to the first derivative $\phi'(z)$ and the last one is about the second derivative $\phi''(z)$. At high level, Property 3.1 requires ϕ to be non-decreasing with homogeneously bounded derivative; Property 3.2 requires ϕ to be highly non-linear; Property 3.3 requires ϕ to be either smooth or piece-wise linear.

Theorem 3.5. ReLU $\phi(z) = \max\{z, 0\}$, leaky ReLU $\phi(z) = \max\{z, 0.01z\}$, squared ReLU $\phi(z) = \max\{z, 0\}^2$ and any non-linear non-decreasing smooth functions with bounded symmetric $\phi'(z)$, like the sigmoid function $\phi(z) = 1/(1 + e^{-z})$, the tanh function and the erf function $\phi(z) = \int_0^z e^{-t^2} dt$, satisfy Property 3.1, 3.2, 3.3. The linear function, $\phi(z) = z$, doesn't satisfy Property 3.2 and the quadratic function, $\phi(z) = z^2$, doesn't satisfy Property 3.1 and 3.2.

The proof can be found in Appendix C.

4 Positive Definiteness of Hessian

In this section, we study the Hessian of empirical risk near the ground truth. We consider the case when v^* is already known. Note that for homogeneous activations, we can assume $v_i^* \in \{-1, 1\}$ since $v\phi(z) = \frac{v}{|v|}\phi(|v|^{1/p}z)$, where p is the degree of homogeneity. As v_i^* only takes discrete values for homogeneous activations, in the next section, we show we can exactly recover v^* using tensor methods with finite samples.

For a set of samples S , we define the *Empirical Risk*,

$$\widehat{f}_S(W) = \frac{1}{2|S|} \sum_{(x,y) \in S} \left(\sum_{i=1}^k v_i^* \phi(w_i^\top x) - y \right)^2. \quad (2)$$

For a distribution \mathcal{D} , we define the *Expected Risk*,

$$f_{\mathcal{D}}(W) = \frac{1}{2} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\left(\sum_{i=1}^k v_i^* \phi(w_i^\top x) - y \right)^2 \right]. \quad (3)$$

Let's calculate the gradient and the Hessian of $\widehat{f}_S(W)$ and $f_{\mathcal{D}}(W)$. For each $j \in [k]$, the partial gradient of $f_{\mathcal{D}}(W)$ with respect to w_j can be represented as

$$\frac{\partial f_{\mathcal{D}}(W)}{\partial w_j} = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\left(\sum_{i=1}^k v_i^* \phi(w_i^\top x) - y \right) v_j^* \phi'(w_j^\top x) x \right].$$

For each $j, l \in [k]$ and $j \neq l$, the second partial derivative of $f_{\mathcal{D}}(W)$ for the (j, l) -th off-diagonal block is,

$$\frac{\partial^2 f_{\mathcal{D}}(W)}{\partial w_j \partial w_l} = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[v_j^* v_l^* \phi'(w_j^\top x) \phi'(w_l^\top x) x x^\top \right],$$

and for each $j \in [k]$, the second partial derivative of $f_{\mathcal{D}}(W)$ for the j -th diagonal block is

$$\frac{\partial^2 f_{\mathcal{D}}(W)}{\partial w_j^2} = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\left(\sum_{i=1}^k v_i^* \phi(w_i^\top x) - y \right) v_j^* \phi''(w_j^\top x) x x^\top + (v_j^* \phi'(w_j^\top x))^2 x x^\top \right].$$

If $\phi(z)$ is non-smooth, we use the Dirac function and its derivatives to represent $\phi''(z)$. Replacing the expectation $\mathbb{E}_{(x,y)\sim\mathcal{D}}$ by the average over the samples $|S|^{-1}\sum_{(x,y)\in S}$, we obtain the Hessian of the empirical risk.

Considering the case when $W = W^* \in \mathbb{R}^{d\times k}$, for all $j, l \in [k]$, we have,

$$\frac{\partial^2 f_{\mathcal{D}}(W^*)}{\partial w_j \partial w_l} = \mathbb{E}_{(x,y)\sim\mathcal{D}} \left[v_j^* v_l^* \phi'(w_j^{*\top} x) \phi'(w_l^{*\top} x) x x^\top \right].$$

If Property 3.3(b) is satisfied, $\phi''(z) = 0$ almost surely. So in this case the diagonal blocks of the empirical Hessian can be written as,

$$\frac{\partial^2 \widehat{f}_S(W)}{\partial w_j^2} = \frac{1}{|S|} \sum_{(x,y)\in S} (v_j^* \phi'(w_j^\top x))^2 x x^\top.$$

Now we show the Hessian of the objective near the global optimum is positive definite.

Definition 4.1. Given the ground truth matrix $W^* \in \mathbb{R}^{d\times k}$, let $\sigma_i(W^*)$ denote the i -th singular value of W^* , often abbreviated as σ_i . Let $\kappa = \sigma_1/\sigma_k$, $\lambda = (\prod_{i=1}^k \sigma_i)/\sigma_k^k$. Let v_{\max} denote $\max_{i\in[k]} |v_i^*|$ and v_{\min} denote $\min_{i\in[k]} |v_i^*|$. Let $\nu = v_{\max}/v_{\min}$. Let ρ denote $\rho(\sigma_k)$. Let $\tau = (3\sigma_1/2)^{4p}/\min_{\sigma\in[\sigma_k/2, 3\sigma_1/2]} \{\rho^2(\sigma)\}$.

Theorem 4.2 (Informal version of Theorem D.1). For any $W \in \mathbb{R}^{d\times k}$ with $\|W - W^*\| \leq \text{poly}(1/k, 1/\lambda, 1/\nu, \rho/\sigma_1^{2p}) \cdot \|W^*\|$, let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)) and let the activation function satisfy Property 3.1, 3.2, 3.3. Then for any $t \geq 1$, if $|S| \geq d \cdot \text{poly}(\log d, t, k, \nu, \tau, \lambda, \sigma_1^{2p}/\rho)$, we have with probability at least $1 - d^{-\Omega(t)}$,

$$\Omega(v_{\min}^2 \rho(\sigma_k)/(\kappa^2 \lambda)) I \preceq \nabla^2 \widehat{f}_S(W) \preceq O(kv_{\max}^2 \sigma_1^{2p}) I.$$

Remark 4.3. As we can see from Theorem 4.2, $\rho(\sigma_k)$ from Property 3.2 plays an important role for positive definite (PD) property. Interestingly, many popular activations, like ReLU, sigmoid and tanh, have $\rho(\sigma_k) > 0$, while some simple functions like linear ($\phi(z) = z$) and square ($\phi(z) = z^2$) functions have $\rho(\sigma_k) = 0$ and their Hessians are rank-deficient. Another important numbers are κ and λ , two different condition numbers of the weight matrix, which directly influences the positive definiteness. If W^* is rank deficient, $\lambda \rightarrow \infty$, $\kappa \rightarrow \infty$ and we don't have PD property. In the best case when W^* is orthogonal, $\lambda = \kappa = 1$. In the worse case, λ can be exponential in k . Also W should be close enough to W^* . In the next section, we provide tensor methods to initialize w_i^* and v_i^* such that they satisfy the conditions in Theorem 4.2.

For the PD property to hold, we need the samples to be independent of the current parameters. Therefore, we need to do resampling at each iteration to guarantee the convergence in iterative algorithms like gradient descent. The following theorem provides the linear convergence guarantee of gradient descent for smooth activations.

Theorem 4.4 (Linear convergence of gradient descent, informal version of Theorem D.2). Let W be the current iterate satisfying $\|W - W^*\| \leq \text{poly}(1/\nu, 1/k, 1/\lambda, \rho/\sigma_1^{2p}) \|W^*\|$. Let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)) with $|S| \geq d \cdot \text{poly}(\log d, t, k, \nu, \tau, \lambda, \sigma_1^{2p}/\rho)$ and let the activation function satisfy Property 3.1, 3.2 and 3.3(a). Define $m_0 := \Theta(v_{\min}^2 \rho(\sigma_k)/(\kappa^2 \lambda))$ and $M_0 := \Theta(kv_{\max}^2 \sigma_1^{2p})$. If we perform gradient descent with step size $1/M_0$ on $\widehat{f}_S(W)$ and obtain the next iterate,

$$\widetilde{W} = W - \frac{1}{M_0} \nabla \widehat{f}_S(W),$$

then with probability at least $1 - d^{-\Omega(t)}$,

$$\|\widetilde{W} - W^*\|_F^2 \leq \left(1 - \frac{m_0}{M_0}\right) \|W - W^*\|_F^2.$$

We provide the proofs in the Appendix D.1

5 Tensor Methods for Initialization

In this section, we show that Tensor methods can recover the parameters W^* to some precision and exactly recover v^* for homogeneous activations.

It is known that most tensor problems are NP-hard [Hås90, HL13] or even hard to approximate [SWZ17b]. However, by making some assumptions, tensor decomposition method becomes efficient [AGH+14, WTSA15, WA16, SWZ16]. Here we utilize the noiseless assumption and Gaussian inputs assumption to show a provable and efficient tensor methods.

5.1 Preliminary

Let's define a special outer product $\widetilde{\otimes}$ for simplification of the notation. If $v \in \mathbb{R}^d$ is a vector and I is the identity matrix, then $v \widetilde{\otimes} I = \sum_{j=1}^d [v \otimes e_j \otimes e_j + e_j \otimes v \otimes e_j + e_j \otimes e_j \otimes v]$. If M is a symmetric rank- r matrix factorized as $M = \sum_{i=1}^r s_i v_i v_i^\top$ and I is the identity matrix, then

$$M \widetilde{\otimes} I = \sum_{i=1}^r s_i \sum_{j=1}^d \sum_{l=1}^6 A_{l,i,j},$$

where $A_{1,i,j} = v_i \otimes v_i \otimes e_j \otimes e_j$, $A_{2,i,j} = v_i \otimes e_j \otimes v_i \otimes e_j$, $A_{3,i,j} = e_j \otimes v_i \otimes v_i \otimes e_j$, $A_{4,i,j} = v_i \otimes e_j \otimes e_j \otimes v_i$, $A_{5,i,j} = e_j \otimes v_i \otimes e_j \otimes v_i$ and $A_{6,i,j} = e_j \otimes e_j \otimes v_i \otimes v_i$.

Denote $\bar{w} = w/\|w\|$. Now let's calculate some moments.

Definition 5.1. We define M_1, M_2, M_3, M_4 and $m_{1,i}, m_{2,i}, m_{3,i}, m_{4,i}$ as follows :

$$M_1 = \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot x].$$

$$M_2 = \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot (x \otimes x - I)].$$

$$M_3 = \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot (x^{\otimes 3} - x \widetilde{\otimes} I)].$$

$$M_4 = \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot (x^{\otimes 4} - (x \otimes x) \widetilde{\otimes} I + I \widetilde{\otimes} I)].$$

$$\gamma_j(\sigma) = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\phi(\sigma \cdot z) z^j], \quad \forall j = 0, 1, 2, 3, 4.$$

$$m_{1,i} = \gamma_1(\|w_i^*\|).$$

$$m_{2,i} = \gamma_2(\|w_i^*\|) - \gamma_0(\|w_i^*\|).$$

$$m_{3,i} = \gamma_3(\|w_i^*\|) - 3\gamma_1(\|w_i^*\|).$$

$$m_{4,i} = \gamma_4(\|w_i^*\|) + 3\gamma_0(\|w_i^*\|) - 6\gamma_2(\|w_i^*\|).$$

According to Definition 5.1, we have the following results,

Claim 5.2. For each $j \in [4]$, $M_j = \sum_{i=1}^k v_i^* m_{j,i} \bar{w}_i^{*\otimes j}$.

Note that some $m_{j,i}$'s will be zero for specific activations. For example, for activations with symmetric first derivatives, i.e., $\phi'(z) = \phi'(-z)$, like sigmoid and erf, we have $\phi(z) + \phi(-z)$ being a constant and $M_2 = 0$ since $\gamma_0(\sigma) = \gamma_2(\sigma)$. Another example is ReLU. ReLU functions have vanishing M_3 , i.e., $M_3 = 0$, as $\gamma_3(\sigma) = 3\gamma_1(\sigma)$. To make tensor methods work, we make the following assumption.

Assumption 5.3. Assume the activation function $\phi(z)$ satisfies the following conditions:

1. If $M_j \neq 0$, then $m_{j,i} \neq 0$ for all $i \in [k]$.
2. At least one of M_3 and M_4 is non-zero.
3. If $M_1 = M_3 = 0$, then $\phi(z)$ is an even function, i.e., $\phi(z) = \phi(-z)$.
4. If $M_2 = M_4 = 0$, then $\phi(z)$ is an odd function, i.e., $\phi(z) = -\phi(-z)$.

If $\phi(z)$ is an odd function then $\phi(z) = -\phi(-z)$ and $v\phi(w^\top x) = -v\phi(-w^\top x)$. Hence we can always assume $v > 0$. If $\phi(z)$ is an even function, then $v\phi(w^\top x) = v\phi(-w^\top x)$. So if w recovers w^* then $-w$ also recovers w^* . Note that ReLU, leaky ReLU and squared ReLU satisfy Assumption 5.3. We further define the following non-zero moments.

Definition 5.4. Let $\alpha \in \mathbb{R}^d$ denote a randomly picked vector. We define P_2 and P_3 as follows: $P_2 = M_{j_2}(I, I, \alpha, \dots, \alpha)$, where $j_2 = \min\{j \geq 2 | M_j \neq 0\}$ and $P_3 = M_{j_3}(I, I, I, \alpha, \dots, \alpha)$, where $j_3 = \min\{j \geq 3 | M_j \neq 0\}$.

According to Definition 5.1 and 5.4, we have,

Claim 5.5. $P_2 = \sum_{i=1}^k v_i^* m_{j_2,i} (\alpha^\top \bar{w}_i^*)^{j_2-2} \bar{w}_i^{*\otimes 2}$ and $P_3 = \sum_{i=1}^k v_i^* m_{j_3,i} (\alpha^\top \bar{w}_i^*)^{j_3-3} \bar{w}_i^{*\otimes 3}$.

In other words for the above definition, P_2 is equal to the first non-zero matrix in the ordered sequence $\{M_2, M_3(I, I, \alpha), M_4(I, I, \alpha, \alpha)\}$. P_3 is equal to the first non-zero tensor in the ordered sequence $\{M_3, M_4(I, I, I, \alpha)\}$. Since α is randomly picked up, $w_i^{*\top} \alpha \neq 0$ and we view this number as a constant throughout this paper. So by construction and Assumption 5.3, both P_2 and P_3 are rank- k . Also, let $\hat{P}_2 \in \mathbb{R}^{d \times d}$ and $\hat{P}_3 \in \mathbb{R}^{d \times d \times d}$ denote the corresponding empirical moments of $P_2 \in \mathbb{R}^{d \times d}$ and $P_3 \in \mathbb{R}^{d \times d \times d}$ respectively.

5.2 Algorithm

Now we briefly introduce how to use a set of samples with size linear in dimension to recover the ground truth parameters to some precision. As shown in the previous section, we have a rank- k 3rd-order moment P_3 that has tensor decomposition formed by $\{\bar{w}_1^*, \bar{w}_2^*, \dots, \bar{w}_k^*\}$. Therefore, we can use the non-orthogonal decomposition method [KCL15] to decompose the corresponding estimated tensor \hat{P}_3 and obtain an approximation of the parameters. The precision of the obtained parameters depends on the estimation error of P_3 , which requires $\Omega(d^3/\epsilon^2)$ samples to achieve ϵ error. Also, the time complexity for tensor decomposition on a $d \times d \times d$ tensor is $\Omega(d^3)$.

In this paper, we reduce the cubic dependency of sample/computational complexity in dimension [JSA15] to linear dependency. Our idea follows the techniques used in [ZJD16], where they first used a 2nd-order moment P_2 to approximate the subspace spanned by $\{\bar{w}_1^*, \bar{w}_2^*, \dots, \bar{w}_k^*\}$, denoted as V , then use V to reduce a higher-dimensional third-order tensor $P_3 \in \mathbb{R}^{d \times d \times d}$ to a lower-dimensional tensor $R_3 := P_3(V, V, V) \in \mathbb{R}^{k \times k \times k}$. Since the tensor decomposition and the tensor estimation are conducted on a lower-dimensional $\mathbb{R}^{k \times k \times k}$ space, the sample complexity and computational complexity are reduced.

The detailed algorithm is shown in Algorithm 1. First, we randomly partition the dataset into three subsets each with size $\tilde{O}(d)$. Then apply the power method on \hat{P}_2 , which is the estimation of P_2 from S_2 , to estimate V . After that, the non-orthogonal tensor decomposition (KCL)[KCL15] on \hat{R}_3 outputs \hat{u}_i which estimates $s_i V^\top \bar{w}_i^*$ for $i \in [k]$ with unknown sign $s_i \in \{-1, 1\}$. Hence \bar{w}_i^* can be estimated by $s_i V \hat{u}_i$. Finally we estimate the magnitude of w_i^* and the signs s_i, v_i^* in the RECMAGSIGN function for homogeneous activations. We discuss the details of each procedure and provide POWERMETHOD and RECMAGSIGN algorithms in Appendix E.

Algorithm 1 Initialization via Tensor Method

```

1: procedure INITIALIZATION( $S$ ) ▷ Theorem 5.6
2:    $S_2, S_3, S_4 \leftarrow \text{PARTITION}(S, 3)$ 
3:    $\widehat{P}_2 \leftarrow \mathbb{E}_{S_2}[P_2]$ 
4:    $V \leftarrow \text{POWERMETHOD}(\widehat{P}_2, k)$ 
5:    $\widehat{R}_3 \leftarrow \mathbb{E}_{S_3}[P_3(V, V, V)]$ 
6:    $\{\widehat{u}_i\}_{i \in [k]} \leftarrow \text{KCL}(\widehat{R}_3)$ 
7:    $\{w_i^{(0)}, v_i^{(0)}\}_{i \in [k]} \leftarrow \text{RECMAGSIGN}(V, \{\widehat{u}_i\}_{i \in [k]}, S_4)$ 
8:   Return  $\{w_i^{(0)}, v_i^{(0)}\}_{i \in [k]}$ 
9: end procedure


```

Algorithm 2 Globally Converging Algorithm

```

1: procedure LEARNING1NN( $S, d, k, \epsilon$ ) ▷ Theorem 6.1
2:    $T \leftarrow \log(1/\epsilon) \cdot \text{poly}(k, \nu, \lambda, \sigma_1^{2p}/\rho)$ .
3:    $\eta \leftarrow 1/(k\nu_{\max}^2 \sigma_1^{2p})$ .
4:    $S_0, S_1, \dots, S_q \leftarrow \text{PARTITION}(S, q+1)$ .
5:    $W^{(0)}, v^{(0)} \leftarrow \text{INITIALIZATION}(S_0)$ .
6:   Set  $v_i^* \leftarrow v_i^{(0)}$  in Eq. (2) for all  $\widehat{f}_{S_q}(W)$ ,  $q \in [T]$ 
7:   for  $q = 0, 1, 2, \dots, T-1$  do
8:      $W^{(q+1)} = W^{(q)} - \eta \nabla \widehat{f}_{S_{q+1}}(W^{(q)})$ 
9:   end for
10:  Return  $\{w_i^{(T)}, v_i^{(0)}\}_{i \in [k]}$ 
11: end procedure

```



(a) Sample complexity for recovery (b) Tensor initialization error (c) Objective v.s. iterations

Figure 1: Numerical Experiments

5.3 Theoretical Analysis

We formally present our theorem for Algorithm 1, and provide the proof in the Appendix E.2.

Theorem 5.6. *Let the activation function be homogeneous satisfying Assumption 5.3. For any $0 < \epsilon < 1$ and $t \geq 1$, if $|S| \geq \epsilon^{-2} \cdot d \cdot \text{poly}(t, k, \kappa, \log d)$, then there exists an algorithm (Algorithm 1) that takes $|S|k \cdot \widetilde{O}(d)$ time and outputs a matrix $W^{(0)} \in \mathbb{R}^{d \times k}$ and a vector $v^{(0)} \in \mathbb{R}^k$ such that, with probability at least $1 - d^{-\Omega(t)}$,*

$$\|W^{(0)} - W^*\|_F \leq \epsilon \cdot \text{poly}(k, \kappa) \|W^*\|_F, \text{ and } v_i^{(0)} = v_i^*.$$

6 Global Convergence

Combining the positive definiteness of the Hessian near the global optimal in Section 4 and the tensor initialization methods in Section 5, we come up with the overall globally converging algorithm Algorithm 2 and its guarantee Theorem 6.1.

Theorem 6.1 (Global convergence guarantees). *Let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)) and let the activation function be homogeneous satisfying Property 3.1, 3.2, 3.3(a) and Assumption 5.3. Then for any $t \geq 1$ and any $\epsilon > 0$, if $|S| \geq d \log(1/\epsilon) \cdot \text{poly}(\log d, t, k, \lambda)$, $T \geq \log(1/\epsilon) \cdot \text{poly}(k, \nu, \lambda, \sigma_1^{2p}/\rho)$ and $0 < \eta \leq 1/(k\nu_{\max}^2 \sigma_1^{2p})$, then there is an Algorithm (procedure LEARNING1NN in Algorithm 2) taking $|S| \cdot d \cdot \text{poly}(\log d, k, \lambda)$ time and outputting a matrix $W^{(T)} \in \mathbb{R}^{d \times k}$ and a vector $v^{(0)} \in \mathbb{R}^k$ satisfying*

$$\|W^{(T)} - W^*\|_F \leq \epsilon \|W^*\|_F, \text{ and } v_i^{(0)} = v_i^*.$$

with probability at least $1 - d^{-\Omega(t)}$.

This follows by combining Theorem 4.4 and Theorem 5.6.

7 Numerical Experiments

In this section we use synthetic data to verify our theoretical results. We generate data points $\{x_i, y_i\}_{i=1,2,\dots,n}$ from Distribution \mathcal{D} (defined in Eq. (1)). We set $W^* = U\Sigma V^\top$, where $U \in \mathbb{R}^{d \times k}$ and $V \in \mathbb{R}^{k \times k}$ are orthogonal matrices generated from QR decomposition of Gaussian matrices, Σ is a diagonal matrix whose diagonal elements are $1, 1 + \frac{\kappa-1}{k-1}, 1 + \frac{2(\kappa-1)}{k-1}, \dots, \kappa$. In this experiment, we set $\kappa = 2$ and $k = 5$. We set v_i^* to be randomly picked from $\{-1, 1\}$ with equal chance. We use squared ReLU $\phi(z) = \max\{z, 0\}^2$, which is a smooth homogeneous function. For non-orthogonal tensor methods, we directly use the code provided by [KCL15] with the number of random projections fixed as $L = 100$. We pick the stepsize $\eta = 0.02$ for gradient descent. In the experiments, we don't do the resampling since the algorithm still works well without resampling.

First we show the number of samples required to recover the parameters for different dimensions. We fix $k = 5$, change d for $d = 10, 20, \dots, 100$ and n for $n = 1000, 2000, \dots, 10000$. For each pair of d and n , we run 10 trials. We say a trial successfully recovers the parameters if there exists a permutation $\pi : [k] \rightarrow [k]$, such that the returned parameters W and v satisfy

$$\max_{j \in [k]} \{\|w_j^* - w_{\pi(j)}\| / \|w_j^*\|\} \leq 0.01 \text{ and } v_{\pi(j)} = v_j^*.$$

We record the recovery rates and represent them as grey scale in Fig. 1(a). As we can see from Fig. 1(a), the least number of samples required to have 100% recovery rate is about proportional to the dimension.

Next we test the tensor initialization. We show the error between the output of the tensor method and the ground truth parameters against the number of samples under different dimensions in Fig 1(b). The pure dark blocks indicate, in at least one of the 10 trials, $\sum_{i=1}^k v_i^{(0)} \neq \sum_{i=1}^k v_i^*$, which means $v_i^{(0)}$ is not correctly initialized. Let $\Pi(k)$ denote the set of all possible permutations $\pi : [k] \rightarrow [k]$. The grey scale represents the averaged error,

$$\min_{\pi \in \Pi(k)} \max_{j \in [k]} \{\|w_j^* - w_{\pi(j)}^{(0)}\| / \|w_j^*\|\},$$

over 10 trials. As we can see, with a fixed dimension, the more samples we have the better initialization we obtain. We can also see that to achieve the same initialization error, the sample complexity required is about proportional to the dimension.

We also compare different initialization methods for gradient descent in Fig. 1(c). We fix $d = 10, k = 5, n = 10000$ and compare three different initialization approaches, (I) Let both v and W be initialized from tensor methods, and then do gradient descent for W while v is fixed; (II) Let both v and W be initialized from random Gaussian, and then do gradient descent for both W and v ; (III) Let $v = v^*$ and W be initialized from random Gaussian, and then do gradient descent for W while v is fixed. As we can see from Fig 1(c), Approach (I) is the fastest and Approach (II) doesn't converge even if more iterations are allowed. Both Approach (I) and (III) have linear convergence rate when the objective value is small enough, which verifies our local linear convergence claim.

8 Conclusion

As shown in Theorem 6.1, the tensor initialization followed by gradient descent will provide a globally converging algorithm with linear time/sample complexity in dimension, logarithmic in precision and polynomial in other factors for smooth homogeneous activation functions. Our distilled properties for activation functions include a wide range of non-linear functions and hopefully provide an intuition to understand the role of non-linear activations played in optimization. Deeper neural networks and convergence for SGD will be considered in the future.

References

- [ABGM14] Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma. Provable bounds for learning some deep representations. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 584–592. <https://arxiv.org/pdf/1310.6343.pdf>, 2014.
- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *JMLR*, 15:2773–2832, 2014.
- [AGKM12] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a non-negative matrix factorization—provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing (STOC)*, pages 145–162. ACM, 2012.
- [AGM12] Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models—going beyond svd. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 1–10. IEEE, 2012.
- [AGMR17] Sanjeev Arora, Rong Ge, Tengyu Ma, and Andrej Risteski. Provable learning of noisy-or networks. In *Proceedings of the 49th Annual Symposium on the Theory of Computing (STOC)*. <https://arxiv.org/pdf/1612.08795.pdf>, 2017.
- [AKZ12] Peter Arbenz, Daniel Kressner, and D-MATH ETH Zürich. Lecture notes on solving large scale eigenvalue problems. <http://people.inf.ethz.ch/arbenz/ewp/Lnotes/lsevp2010.pdf>, 2012.
- [APVZ14] Alexandr Andoni, Rina Panigrahy, Gregory Valiant, and Li Zhang. Learning polynomials with neural networks. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1908–1916, 2014.
- [AZS14] Özlem Aslan, Xinhua Zhang, and Dale Schuurmans. Convex deep learning via normalized kernels. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3275–3283, 2014.
- [Bac14] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *arXiv preprint arXiv:1412.8690*, 2014.
- [Bal16] David Balduzzi. Deep online convex optimization with gated games. *arXiv preprint arXiv:1604.01952*, 2016.
- [Bar98] Peter L. Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Transactions on Information Theory*, 44(2):525–536, 1998.
- [BG17] Alon Brutzkus and Amir Globerson. Globally optimal gradient descent for a convnet with gaussian inputs. *arXiv preprint arXiv:1702.07966*, 2017.
- [BLWZ17] Maria-Florina Balcan, Yingyu Liang, David P. Woodruff, and Hongyang Zhang. Optimal sample complexity for matrix completion and related problems via ℓ_2 -regularization. *arXiv preprint arXiv:1704.08683*, 2017.

- [BMBY16] David Balduzzi, Brian McWilliams, and Tony Butler-Yeoman. Neural taylor approximations: Convergence and exploration in rectifier networks. *arXiv preprint arXiv:1611.02345*, 2016.
- [BR88] Avrim Blum and Ronald L Rivest. Training a 3-node neural network is np-complete. In *Proceedings of the 1st International Conference on Neural Information Processing Systems (NIPS)*, pages 494–501. MIT Press, 1988.
- [BRV⁺05] Yoshua Bengio, Nicolas L Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marquette. Convex neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 123–130, 2005.
- [CHM⁺15] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gerard Ben Arous, and Yann LeCun. The loss surfaces of multilayer networks. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 192–204, 2015.
- [CS16] Nadav Cohen and Amnon Shashua. Convolutional rectifier networks as generalized tensor decompositions. In *International Conference on Machine Learning (ICML)*, 2016.
- [CSS16] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *29th Annual Conference on Learning Theory (COLT)*, pages 698–728, 2016.
- [DFS16] Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In *Advances in neural information processing systems (NIPS)*, pages 2253–2261, 2016.
- [DPG⁺14] Yann N Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in neural information processing systems (NIPS)*, pages 2933–2941, 2014.
- [EV09] Ehsan Elhamifar and René Vidal. Sparse subspace clustering. In *CVPR*, pages 2790–2797, 2009.
- [FB16] C Daniel Freeman and Joan Bruna. Topology and geometry of half-rectified network optimization. In *arXiv preprint <https://arxiv.org/pdf/1611.01540.pdf>*, 2016.
- [FZK⁺16] Jiashi Feng, Tom Zahavy, Bingyi Kang, Huan Xu, and Shie Mannor. Ensemble robustness of deep learning algorithms. *arXiv preprint arXiv:1602.02389*, 2016.
- [GG11] Nicolas Gillis and François Glineur. Low-rank matrix approximation with weights or missing data is np-hard. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1149–1165, 2011.
- [GKKT17] Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. In *30th Annual Conference on Learning Theory (COLT)*. <https://arxiv.org/pdf/1611.10258.pdf>, 2017.
- [GV15] Nicolas Gillis and Stephen A Vavasis. On the complexity of robust pca and ℓ_1 -norm low-rank matrix approximation. *arXiv preprint arXiv:1509.09236*, 2015.

- [Har14] Moritz Hardt. Understanding alternating minimization for matrix completion. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 651–660. IEEE, 2014.
- [Hås90] Johan Håstad. Tensor rank is np-complete. *Journal of Algorithms*, 11(4):644–654, 1990.
- [HK13] Daniel Hsu and Sham M Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *ITCS*, pages 11–20. ACM, 2013.
- [HKZ12] Daniel Hsu, Sham M Kakade, and Tong Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electron. Commun. Probab.*, 17(52):1–6, 2012.
- [HL13] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are np-hard. In *Journal of the ACM (JACM)*, volume 60(6), page 45. <https://arxiv.org/pdf/0911.1393.pdf>, 2013.
- [HM13] Moritz Hardt and Ankur Moitra. Algorithms and hardness for robust subspace recovery. In *COLT*, volume 30, pages 354–375, 2013.
- [HM17] Moritz Hardt and Tengyu Ma. Identity matters in deep learning. *ICLR*, 2017.
- [Hor91] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [HRS16] Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *ICML*, pages 1225–1234, 2016.
- [HV15] Benjamin D Haeffele and René Vidal. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing (STOC)*, 2013.
- [JSA15] Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv preprint arXiv:1506.08473*, 2015.
- [Kaw16] Kenji Kawaguchi. Deep learning without poor local minima. *arXiv preprint arXiv:1605.07110*, 2016.
- [KCL15] Volodymyr Kuleshov, Arun Chaganty, and Percy Liang. Tensor factorization via matrix factorization. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 507–516, 2015.
- [LSS14] Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In *Advances in neural information processing systems (NIPS)*, pages 855–863, 2014.
- [MPCB14] Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In *Advances in neural information processing systems (NIPS)*, pages 2924–2932, 2014.

- [PLR⁺16] Ben Poole, Subhaneil Lahiri, Maithreyi Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Advances In Neural Information Processing Systems (NIPS)*, pages 3360–3368, 2016.
- [RPK⁺16] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. *arXiv preprint arXiv:1606.05336*, 2016.
- [RSW16] Ilya P Razenshteyn, Zhao Song, and David P. Woodruff. Weighted low rank approximations with provable guarantees. In *Proceedings of the 48th Annual Symposium on the Theory of Computing (STOC)*, pages 250–263, 2016.
- [SA15] Hanie Sedghi and Anima Anandkumar. Provable methods for training neural networks with sparse connectivity. In *International Conference on Learning Representation (ICLR)*, 2015.
- [SBL16] Levent Sagun, Léon Bottou, and Yann LeCun. Singularity of the Hessian in deep learning. *arXiv preprint arXiv:1611.07476*, 2016.
- [SC16] Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016.
- [SCP16] Grzegorz Swirszcz, Wojciech Marian Czarnecki, and Razvan Pascanu. Local minima in training of deep networks. *arXiv preprint arXiv:1611.06310*, 2016.
- [Sha16] Ohad Shamir. Distribution-specific hardness of learning neural networks. *arXiv preprint arXiv:1609.01037*, 2016.
- [SL15] Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via non-convex factorization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 270–289. IEEE, 2015.
- [SR11] David Sontag and Dan Roy. Complexity of inference in latent dirichlet allocation. In *Advances in neural information processing systems*, pages 1008–1016, 2011.
- [SS16] Itay Safran and Ohad Shamir. On the quality of the initial basin in overspecified neural networks. In *International Conference on Machine Learning (ICML)*, 2016.
- [SWZ16] Zhao Song, David P. Woodruff, and Huan Zhang. Sublinear time orthogonal tensor decomposition. In *Advances in Neural Information Processing Systems(NIPS)*, pages 793–801, 2016.
- [SWZ17a] Zhao Song, David P. Woodruff, and Peilin Zhong. Low rank approximation with entrywise ℓ_1 -norm error. In *Proceedings of the 49th Annual Symposium on the Theory of Computing (STOC)*. ACM, <https://arxiv.org/pdf/1611.00898.pdf>, 2017.
- [SWZ17b] Zhao Song, David P. Woodruff, and Peilin Zhong. Relative error tensor low rank approximation. In *arXiv preprint*. <https://arxiv.org/pdf/1704.08246.pdf>, 2017.
- [Tel16] Matus Telgarsky. Benefits of depth in neural networks. In *29th Annual Conference on Learning Theory (COLT)*, pages 1517–1539, 2016.

- [Tia17] Yuandong Tian. Symmetry-breaking convergence analysis of certain two-layered neural networks with ReLU nonlinearity. In *Workshop at International Conference on Learning Representation*, 2017.
- [Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- [WA16] Yining Wang and Anima Anandkumar. Online and differentially-private tensor decomposition. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3531–3539, 2016.
- [WTSA15] Yining Wang, Hsiao-Yu Tung, Alexander J Smola, and Anima Anandkumar. Fast and guaranteed tensor decomposition via sketching. In *Advances in Neural Information Processing Systems (NIPS)*, pages 991–999, 2015.
- [XLS17] Bo Xie, Yingyu Liang, and Le Song. Diversity leads to generalization in neural networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- [YCS14] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi. Alternating minimization for mixed linear regression. In *ICML*, pages 613–621, 2014.
- [ZBH⁺17] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.
- [ZJD16] Kai Zhong, Prateek Jain, and Inderjit S Dhillon. Mixed linear regression with multiple components. In *Advances in neural information processing systems (NIPS)*, pages 2190–2198, 2016.
- [ZLJ16] Yuchen Zhang, Jason D Lee, and Michael I Jordan. L1-regularized neural networks are improperly learnable in polynomial time. In *Proceedings of The 33rd International Conference on Machine Learning (ICML)*, pages 993–1001, 2016.
- [ZLW16] Yuchen Zhang, Percy Liang, and Martin J Wainwright. Convexified convolutional neural networks. *arXiv preprint arXiv:1609.01000*, 2016.
- [ZLWJ17] Yuchen Zhang, Jason D. Lee, Martin J. Wainwright, and Michael I. Jordan. On the learnability of fully-connected neural networks. In *International Conference on Artificial Intelligence and Statistics*, 2017.

A Notation

For any positive integer n , we use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For random variable X , let $\mathbb{E}[X]$ denote the expectation of X (if this quantity exists). For any vector $x \in \mathbb{R}^n$, we use $\|x\|$ to denote its ℓ_2 norm.

We provide several definitions related to matrix A . Let $\det(A)$ denote the determinant of a square matrix A . Let A^\top denote the transpose of A . Let A^\dagger denote the Moore-Penrose pseudoinverse of A . Let A^{-1} denote the inverse of a full rank square matrix. Let $\|A\|_F$ denote the Frobenius norm of matrix A . Let $\|A\|$ denote the spectral norm of matrix A . Let $\sigma_i(A)$ denote the i -th largest singular value of A . We often use capital letter to denote the stack of corresponding small letter vectors, e.g., $W = [w_1 \ w_2 \ \dots \ w_k]$. For two same-size matrices, $A, B \in \mathbb{R}^{d_1 \times d_2}$, we use $A \circ B \in \mathbb{R}^{d_1 \times d_2}$ to denote element-wise multiplication of these two matrices.

We use \otimes to denote outer product and \cdot to denote dot product. Given two column vectors $u, v \in \mathbb{R}^n$, then $u \otimes v \in \mathbb{R}^{n \times n}$ and $(u \otimes v)_{i,j} = u_i \cdot v_j$, and $u^\top v = \sum_{i=1}^n u_i v_i \in \mathbb{R}$. Given three column vectors $u, v, w \in \mathbb{R}^n$, then $u \otimes v \otimes w \in \mathbb{R}^{n \times n \times n}$ and $(u \otimes v \otimes w)_{i,j,k} = u_i \cdot v_j \cdot w_k$. We use $u^{\otimes r} \in \mathbb{R}^{n^r}$ to denote the vector u outer product with itself $r - 1$ times.

Tensor $T \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is symmetric if and only if for any i, j, k , $T_{i,j,k} = T_{i,k,j} = T_{j,i,k} = T_{j,k,i} = T_{k,i,j} = T_{k,j,i}$. Given a third order tensor $T \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and three matrices $A \in \mathbb{R}^{n_1 \times d_1}, B \in \mathbb{R}^{n_2 \times d_2}, C \in \mathbb{R}^{n_3 \times d_3}$, we use $T(A, B, C)$ to denote a $d_1 \times d_2 \times d_3$ tensor where the (i, j, k) -th entry is,

$$\sum_{i'=1}^{n_1} \sum_{j'=1}^{n_2} \sum_{k'=1}^{n_3} T_{i',j',k'} A_{i',i} B_{j',j} C_{k',k}.$$

We use $\|T\|$ to denote the operator norm of the tensor T , i.e.,

$$\|T\| = \max_{\|a\|=1} |T(a, a, a)|.$$

For tensor $T \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, we use matrix $T^{(1)} \in \mathbb{R}^{n_1 \times n_2 n_3}$ to denote the flattening of tensor T along the first dimension, i.e., $[T^{(1)}]_{i,(j-1)n_3+k} = T_{i,j,k}, \forall i \in [n_1], j \in [n_2], k \in [n_3]$. Similarly for matrices $T^{(2)} \in \mathbb{R}^{n_2 \times n_3 n_1}$ and $T^{(3)} \in \mathbb{R}^{n_3 \times n_1 n_2}$.

We use $\mathbf{1}_f$ to denote the indicator function, which is 1 if f holds and 0 otherwise. Let $I_d \in \mathbb{R}^{d \times d}$ denote the identity matrix. We use $\phi(z)$ to denote an activation function. We define $(z)_+ := \max\{0, z\}$. We use \mathcal{D} to denote a Gaussian distribution $\mathcal{N}(0, I_d)$ or to denote a joint distribution of $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$, where the marginal distribution of X is $\mathcal{N}(0, I_d)$.

For any function f , we define $\tilde{O}(f)$ to be $f \cdot \log^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions f, g , we use the shorthand $f \lesssim g$ (resp. \gtrsim) to indicate that $f \leq Cg$ (resp. \geq) for an absolute constant C . We use $f \approx g$ to mean $cf \leq g \leq Cf$ for constants c, C .

B Preliminaries

In this section, we introduce some lemmata and corollaries that will be used in the proofs.

B.1 Useful Facts

We provide some facts that will be used in the later proofs.

Fact B.1. Let z denote a fixed d -dimensional vector, then for any $C \geq 1$ and $n \geq 1$, we have

$$\Pr_{x \sim \mathcal{N}(0, I_d)} [|\langle x, z \rangle|^2 \leq 5C \|z\|^2 \log n] \geq 1 - 1/(nd^C).$$

Proof. This follows by Proposition 1.1 in [HKZ12]. \square

Fact B.2. For any $C \geq 1$ and $n \geq 1$, we have

$$\Pr_{x \sim \mathcal{N}(0, I_d)} [\|x\|^2 \leq 5Cd \log n] \geq 1 - 1/(nd^C).$$

Proof. This follows by Proposition 1.1 in [HKZ12]. \square

Fact B.3. Given a full column-rank matrix $W = [w_1, w_2, \dots, w_k] \in \mathbb{R}^{d \times k}$, let $\bar{W} = [\frac{w_1}{\|w_1\|}, \frac{w_2}{\|w_2\|}, \dots, \frac{w_k}{\|w_k\|}]$. Then, we have: (I) for any $i \in [k]$, $\sigma_k(W) \leq \|w_i\| \leq \sigma_1(W)$; (II) $1/\kappa(W) \leq \sigma_k(\bar{W}) \leq \sigma_1(\bar{W}) \leq \sqrt{k}$.

Proof. Part (I). We have,

$$\sigma_k(W) \leq \|W e_i\| = \|w_i\| \leq \sigma_1(W)$$

Part (II). We first show how to lower bound $\sigma_k(\bar{W})$,

$$\begin{aligned} \sigma_k(\bar{W}) &= \min_{\|s\|=1} \|\bar{W} s\| \\ &= \min_{\|s\|=1} \left\| \sum_{i=1}^k \frac{s_i}{\|w_i\|} w_i \right\| && \text{by definition of } \bar{W} \\ &\geq \min_{\|s\|=1} \sigma_k(W) \left(\sum_{i=1}^k \left(\frac{s_i}{\|w_i\|} \right)^2 \right)^{\frac{1}{2}} && \text{by } \|w_i\| \geq \sigma_k(W) \\ &\geq \min_{\|s\|_2=1} \sigma_k(W) \left(\sum_{i=1}^k \left(\frac{s_i}{\max_{j \in [k]} \|w_j\|} \right)^2 \right)^{\frac{1}{2}} && \text{by } \max_{j \in [k]} \|w_j\| \geq \|w_i\| \\ &= \sigma_k(W) / \max_{j \in [k]} \|w_j\| && \text{by } \|s\| = 1 \\ &\geq \sigma_k(W) / \sigma_1(W). && \text{by } \max_{j \in [k]} \|w_j\| \leq \sigma_1(W) \\ &= 1/\kappa(W). \end{aligned}$$

It remains to upper bound $\sigma_1(\bar{W})$,

$$\sigma_1(\bar{W}) \leq \left(\sum_{i=1}^k \sigma_i^2(\bar{W}) \right)^{\frac{1}{2}} = \|\bar{W}\|_F \leq \sqrt{k}.$$

\square

Fact B.4. Let $U \in \mathbb{R}^{d \times k}$ and $V \in \mathbb{R}^{d \times k}$ ($k \leq d$) denote two orthogonal matrices. Then $\|UU^\top - VV^\top\| = \|(I - UU^\top)V\| = \|(I - VV^\top)U\| = \sqrt{1 - \sigma_k^2(U^\top V)}$.

Proof. Let $U_\perp \in \mathbb{R}^{d \times (d-k)}$ and $V_\perp \in \mathbb{R}^{d \times (d-k)}$ be the orthogonal complementary matrices of $U, V \in \mathbb{R}^{d \times k}$ respectively.

$$\begin{aligned} \|UU^\top - VV^\top\| &= \|(I - VV^\top)UU^\top - VV^\top(I - UU^\top)\| \\ &= \|V_\perp V_\perp^\top UU^\top - VV^\top U_\perp U_\perp^\top\| \\ &= \left\| \begin{bmatrix} V_\perp & V \end{bmatrix} \begin{bmatrix} V_\perp^\top U & 0 \\ 0 & -V^\top U_\perp \end{bmatrix} \begin{bmatrix} U^\top \\ U_\perp^\top \end{bmatrix} \right\| \\ &= \max(\|V_\perp^\top U\|, \|V^\top U_\perp\|). \end{aligned}$$

We show how to simplify $\|V_\perp^\top U\|$,

$$\|V_\perp^\top U\| = \|(I - VV^\top)U\| = \sqrt{\|U^\top(I - VV^\top)U\|} = \max_{\|a\|=1} \sqrt{1 - \|V^\top Ua\|^2} = \sqrt{1 - \sigma_k^2(V^\top U)}.$$

Similarly we can simplify $\|U_\perp^\top V\|$,

$$\|U_\perp^\top V\| = \sqrt{1 - \sigma_k^2(U^\top V)} = \sqrt{1 - \sigma_k^2(V^\top U)}.$$

□

Fact B.5. Let $C \in \mathbb{R}^{d_1 \times d_2}, B \in \mathbb{R}^{d_2 \times d_3}$ be two matrices. Then $\|CB\| \leq \|C\| \|B\|_F$ and $\|CB\| \geq \sigma_{\min}(C) \|B\|_F$.

Proof. For each $i \in [d_3]$, let b_i denote the i -th column of B . We can upper bound $\|CB\|$,

$$\|CB\|_F = \left(\sum_{i=1}^{d_3} \|Cb_i\|^2 \right)^{1/2} \leq \left(\sum_{i=1}^{d_3} \|C\|^2 \|b_i\|^2 \right)^{1/2} = \|C\| \|B\|_F.$$

We show how to lower bound $\|CB\|$,

$$\|CB\| = \left(\sum_{i=1}^{d_3} \|Cb_i\|^2 \right)^{1/2} \geq \left(\sum_{i=1}^{d_3} \sigma_k^2(C) \|b_i\|^2 \right)^{1/2} = \sigma_{\min}(C) \|B\|_F.$$

□

Fact B.6. Let $a, b, c \geq 0$ denote three constants, let $u, v, w \in \mathbb{R}^d$ denote three vectors, let \mathcal{D}_d denote Gaussian distribution $\mathcal{N}(0, I_d)$ then

$$\mathbb{E}_{x \sim \mathcal{D}_d} [|u^\top x|^a |v^\top x|^b |w^\top x|^c] \approx \|u\|^a \|v\|^b \|w\|^c.$$

Proof.

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{D}_d} [|u^\top x|^a |v^\top x|^b |w^\top x|^c] &\leq \left(\mathbb{E}_{x \sim \mathcal{D}_d} [|u^\top x|^{2a}] \right)^{1/2} \cdot \left(\mathbb{E}_{x \sim \mathcal{D}_d} [|v^\top x|^{4b}] \right)^{1/4} \cdot \left(\mathbb{E}_{x \sim \mathcal{D}_d} [|w^\top x|^{4c}] \right)^{1/4} \\ &\lesssim \|u\|^a \|v\|^b \|w\|^c, \end{aligned}$$

where the first step follows by Hölder's inequality, i.e., $\mathbb{E}[|XYZ|] \leq (\mathbb{E}[|X|^2])^{1/2} \cdot (\mathbb{E}[|Y|^4])^{1/4} \cdot (\mathbb{E}[|Z|^4])^{1/4}$, the third step follows by calculating the expectation and a, b, c are constants.

Since all the three components $|u^\top x|, |v^\top x|, |w^\top x|$ are positive and related to a common random vector x , we can show a lower bound,

$$\mathbb{E}_{x \sim \mathcal{D}_d} [|u^\top x|^a |v^\top x|^b |w^\top x|^c] \gtrsim \|u\|^a \|v\|^b \|w\|^c.$$

□

B.2 Matrix Bernstein

In many proofs we need to bound the difference between some population matrices/tensors and their empirical versions. Typically, the classic matrix Bernstein inequality requires the norm of the random matrix be bounded *almost surely* (e.g., Theorem 6.1 in [Tro12]) or the random matrix satisfies subexponential property (Theorem 6.2 in [Tro12]). However, in our cases, most of the random matrices don't satisfy these conditions. So we derive the following lemmata that can deal with random matrices that are not bounded almost surely or follow subexponential distribution, but are bounded with high probability.

Lemma B.7 (Matrix Bernstein for unbounded case (A modified version of bounded case, Theorem 6.1 in [Tro12])). *Let \mathcal{B} denote a distribution over $\mathbb{R}^{d_1 \times d_2}$. Let $d = d_1 + d_2$. Let B_1, B_2, \dots, B_n be i.i.d. random matrices sampled from \mathcal{B} . Let $\bar{B} = \mathbb{E}_{B \sim \mathcal{B}}[B]$ and $\hat{B} = \frac{1}{n} \sum_{i=1}^n B_i$. For parameters $m \geq 0, \gamma \in (0, 1), \nu > 0, L > 0$, if the distribution \mathcal{B} satisfies the following four properties,*

- (I) $\Pr_{B \sim \mathcal{B}} [\|B\| \leq m] \geq 1 - \gamma;$
- (II) $\left\| \mathbb{E}_{B \sim \mathcal{B}} [B] \right\| > 0;$
- (III) $\max \left(\left\| \mathbb{E}_{B \sim \mathcal{B}} [BB^\top] \right\|, \left\| \mathbb{E}_{B \sim \mathcal{B}} [B^\top B] \right\| \right) \leq \nu;$
- (IV) $\max_{\|a\|=\|b\|=1} \left(\mathbb{E}_{B \sim \mathcal{B}} \left[(a^\top B b)^2 \right] \right)^{1/2} \leq L.$

Then we have for any $0 < \epsilon < 1$ and $t \geq 1$, if

$$n \geq (18t \log d) \cdot (\nu + \|\bar{B}\|^2 + m\|\bar{B}\|\epsilon) / (\epsilon^2 \|\bar{B}\|^2) \quad \text{and} \quad \gamma \leq (\epsilon \|\bar{B}\| / (2L))^2$$

with probability at least $1 - 1/d^{2t} - n\gamma$,

$$\|\hat{B} - \bar{B}\| \leq \epsilon \|\bar{B}\|.$$

Proof. Define the event

$$\xi_i = \{\|B_i\| \leq m\}, \forall i \in [n].$$

Define $M_i = \mathbf{1}_{\|B_i\| \leq m} B_i$. Let $\bar{M} = \mathbb{E}_{B \sim \mathcal{B}}[\mathbf{1}_{\|B\| \leq m} B]$ and $\widehat{M} = \frac{1}{n} \sum_{i=1}^n M_i$. By triangle inequality, we have

$$\|\hat{B} - \bar{B}\| \leq \|\hat{B} - \widehat{M}\| + \|\widehat{M} - \bar{M}\| + \|\bar{M} - \bar{B}\|. \quad (4)$$

In the next a few paragraphs, we will upper bound the above three terms.

The first term in Eq. (4). Denote ξ^c as the complementary set of ξ , thus $\Pr[\xi_i^c] \leq \gamma$. By a union bound over $i \in [n]$, with probability $1 - n\gamma$, $\|B_i\| \leq m$ for all $i \in [n]$. Thus $\widehat{M} = \hat{B}$.

The second term in Eq. (4). For a matrix B sampled from \mathcal{B} , we use ξ to denote the event

that $\xi = \{\|B\| \leq m\}$. Then, we can upper bound $\|\overline{M} - \overline{B}\|$ in the following way,

$$\begin{aligned}
& \|\overline{M} - \overline{B}\| \\
&= \left\| \mathbb{E}_{B \sim \mathcal{B}}[\mathbf{1}_{\|B\| \leq m} \cdot B] - \mathbb{E}_{B \sim \mathcal{B}}[B] \right\| \\
&= \left\| \mathbb{E}_{B \sim \mathcal{B}}[B \cdot \mathbf{1}_{\xi^c}] \right\| \\
&= \max_{\|a\|=\|b\|=1} \mathbb{E}_{B \sim \mathcal{B}}[a^\top B b \mathbf{1}_{\xi^c}] \\
&\leq \max_{\|a\|=\|b\|=1} \mathbb{E}_{B \sim \mathcal{B}}[(a^\top B b)^2]^{1/2} \cdot \mathbb{E}_{B \sim \mathcal{B}}[\mathbf{1}_{\xi^c}]^{1/2} && \text{by Hölder's inequality} \\
&\leq L \mathbb{E}_{B \sim \mathcal{B}}[\mathbf{1}_{\xi^c}]^{1/2} && \text{by Property (IV)} \\
&\leq L\gamma^{1/2}, && \text{by } \Pr[\xi^c] \leq \gamma \\
&\leq \frac{1}{2}\epsilon\|\overline{B}\|, && \text{by } \gamma \leq (\epsilon\|\overline{B}\|/(2L))^2
\end{aligned}$$

which implies

$$\|\overline{M} - \overline{B}\| \leq \frac{\epsilon}{2}\|\overline{B}\|.$$

Since $\epsilon < 1$, we also have $\|\overline{M} - \overline{B}\| \leq \frac{1}{2}\|\overline{B}\|$ and $\frac{3}{2}\|\overline{B}\| \geq \|\overline{M}\| \geq \frac{1}{2}\|\overline{B}\|$.

The third term in Eq. (4). We can bound $\|\widehat{M} - \overline{M}\|$ by Matrix Bernstein's inequality [Tro12].

We define $Z_i = M_i - \overline{M}$. Thus we have $\mathbb{E}_{B_i \sim \mathcal{B}}[Z_i] = 0$, $\|Z_i\| \leq 2m$, and

$$\left\| \mathbb{E}_{B_i \sim \mathcal{B}}[Z_i Z_i^\top] \right\| = \left\| \mathbb{E}_{B_i \sim \mathcal{B}}[M_i M_i^\top] - \overline{M} \cdot \overline{M}^\top \right\| \leq \nu + \|\overline{M}\|^2 \leq \nu + 3\|\overline{B}\|^2.$$

Similarly, we have $\left\| \mathbb{E}_{B_i \sim \mathcal{B}}[Z_i^\top Z_i] \right\| \leq \nu + 3\|\overline{B}\|^2$. Using matrix Bernstein's inequality, for any $\epsilon > 0$,

$$\Pr_{B_1, \dots, B_n \sim \mathcal{B}} \left[\frac{1}{n} \left\| \sum_{i=1}^n Z_i \right\| \geq \epsilon\|\overline{B}\| \right] \leq d \exp \left(- \frac{\epsilon^2 \|\overline{B}\|^2 n / 2}{\nu + 3\|\overline{B}\|^2 + 2m\|\overline{B}\|\epsilon / 3} \right).$$

By choosing

$$n \geq (3t \log d) \cdot \frac{\nu + 3\|\overline{B}\|^2 + 2m\|\overline{B}\|\epsilon / 3}{\epsilon^2 \|\overline{B}\|^2 / 2},$$

for $t \geq 1$, we have with probability at least $1 - 1/d^{2t}$,

$$\left\| \frac{1}{n} \sum_{i=1}^n M_i - \overline{M} \right\| \leq \frac{\epsilon}{2}\|\overline{B}\|$$

Putting it all together, we have for $0 < \epsilon < 1$, if

$$n \geq (18t \log d) \cdot (\nu + \|\overline{B}\|^2 + m\|\overline{B}\|\epsilon) / (\epsilon^2 \|\overline{B}\|^2) \quad \text{and} \quad \gamma \leq (\epsilon\|\overline{B}\|/(2L))^2$$

with probability at least $1 - 1/d^{2t} - n\gamma$,

$$\left\| \frac{1}{n} \sum_{i=1}^n B_i - \mathbb{E}_{B \sim \mathcal{B}}[B] \right\| \leq \epsilon \left\| \mathbb{E}_{B \sim \mathcal{B}}[B] \right\|.$$

□

Corollary B.8 (Error bound for symmetric rank-one random matrices). *Let x_1, x_2, \dots, x_n denote n i.i.d. samples drawn from Gaussian distribution $\mathcal{N}(0, I_d)$. Let $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function satisfying the following properties (I), (II) and (III).*

- (I) $\Pr_{x \sim \mathcal{N}(0, I_d)} [|h(x)| \leq m] \geq 1 - \gamma$
- (II) $\left\| \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h(x)xx^\top] \right\| > 0;$
- (III) $\left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h^4(x)] \right)^{1/4} \leq L.$

Define function $B(x) = h(x)xx^\top \in \mathbb{R}^{d \times d}$, $\forall i \in [n]$. Let $\bar{B} = \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h(x)xx^\top]$. For any $0 < \epsilon < 1$ and $t \geq 1$, if

$$n \gtrsim (t \log d) \cdot (L^2 d + \|\bar{B}\|^2 + (mtd \log n) \|\bar{B}\| \epsilon) / (\epsilon^2 \|\bar{B}\|^2), \text{ and } \gamma + 1/(nd^{2t}) \lesssim (\epsilon \|\bar{B}\| / L)^2$$

then

$$\Pr_{x_1, \dots, x_n \sim \mathcal{N}(0, I_d)} \left[\left\| \bar{B} - \frac{1}{n} \sum_{i=1}^n B(x_i) \right\| \leq \epsilon \|\bar{B}\| \right] \geq 1 - 2/(d^{2t}) - n\gamma.$$

Proof. We show that the four Properties in Lemma B.7 are satisfied. Define function $B(x) = h(x)xx^\top$.

$$(I) \|B(x)\| = \|h(x)xx^\top\| = |h(x)| \|x\|^2.$$

By using Fact B.2, we have

$$\Pr_{x \sim \mathcal{N}(0, I_d)} [\|x\|^2 \leq 10td \log n] \geq 1 - 1/(nd^{2t})$$

Therefore,

$$\Pr_{x \sim \mathcal{N}(0, I_d)} [\|B(x)\| \leq m \cdot 10td \log(n)] \geq 1 - \gamma - 1/(nd^{2t}).$$

$$(II) \left\| \mathbb{E}_{B \sim \mathcal{B}} [B] \right\| = \left\| \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h(x)xx^\top] \right\| > 0.$$

(III)

$$\begin{aligned} & \max \left(\left\| \mathbb{E}_{B \sim \mathcal{B}} [BB^\top] \right\|, \left\| \mathbb{E}_{B \sim \mathcal{B}} [B^\top B] \right\| \right) \\ &= \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [(h(x))^2 \|x\|^2 (a^\top x)^2] \\ &\leq \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [(h(x))^4] \right)^{1/2} \cdot \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [\|x\|^8] \right)^{1/4} \cdot \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [(a^\top x)^8] \right)^{1/4} \\ &\lesssim L^2 d. \end{aligned}$$

(IV)

$$\begin{aligned}
& \max_{\|a\|=\|b\|=1} \left(\mathbb{E}_{B \sim \mathcal{B}} [(a^\top B b)^2] \right)^{1/2} \\
&= \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h^2(x) (a^\top x)^4] \right)^{1/2} \\
&\leq \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [h^4(x)] \right)^{1/4} \cdot \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} [(a^\top x)^8] \right)^{1/4} \\
&\lesssim L.
\end{aligned}$$

Applying Lemma B.7, we obtain, for any $0 < \epsilon < 1$ and $t \geq 1$, if

$$n \gtrsim (t \log d) \cdot (L^2 d + \|\bar{B}\|^2 + (mtd \log n) \|\bar{B}\| \epsilon) / (\epsilon^2 \|\bar{B}\|^2), \text{ and } \gamma + 1/(nd^{2t}) \lesssim (\epsilon \|\bar{B}\| / L)^2$$

then

$$\Pr_{x_1, \dots, x_n \sim \mathcal{N}(0, I_d)} \left[\left\| \bar{B} - \frac{1}{n} \sum_{i=1}^n B(x_i) \right\| \leq \epsilon \|\bar{B}\| \right] \geq 1 - 2/(d^{2t}) - n\gamma.$$

□

C Properties of Activation Functions

Theorem 3.5. ReLU $\phi(z) = \max\{z, 0\}$, leaky ReLU $\phi(z) = \max\{z, 0.01z\}$, squared ReLU $\phi(z) = \max\{z, 0\}^2$ and any non-linear non-decreasing smooth functions with bounded symmetric $\phi'(z)$, like the sigmoid function $\phi(z) = 1/(1 + e^{-z})$, the tanh function and the erf function $\phi(z) = \int_0^z e^{-t^2} dt$, satisfy Property 3.1, 3.2, 3.3. The linear function, $\phi(z) = z$, doesn't satisfy Property 3.2 and the quadratic function, $\phi(z) = z^2$, doesn't satisfy Property 3.1 and 3.2.

Proof. We can easily verify that ReLU, leaky ReLU and squared ReLU satisfy Property 3.2 by calculating $\rho(\sigma)$ in Property 3.2, which is shown in Table 1. Property 3.1 for ReLU, leaky ReLU and squared ReLU can be verified since they are non-decreasing with bounded first derivative. ReLU and leaky ReLU are piece-wise linear, so they satisfy Property 3.3(b). Squared ReLU is smooth so it satisfies Property 3.3(a).

Activations	ReLU	Leaky ReLU	squared ReLU	erf	sigmoid ($\sigma = 0.1$)	sigmoid ($\sigma = 1$)	sigmoid ($\sigma = 10$)
$\alpha_0(\sigma)$	$\frac{1}{2}$	$\frac{1.01}{2}$	$\sigma \sqrt{\frac{2}{\pi}}$	$\frac{1}{(2\sigma^2+1)^{1/2}}$	0.99	0.605706	0.079
$\alpha_1(\sigma)$	$\frac{1}{\sqrt{2\pi}}$	$\frac{0.99}{\sqrt{2\pi}}$	σ	0	0	0	0
$\alpha_2(\sigma)$	$\frac{1}{2}$	$\frac{1.01}{2}$	$2\sigma \sqrt{\frac{2}{\pi}}$	$\frac{1}{(2\sigma^2+1)^{3/2}}$	0.97	0.24	0.00065
$\beta_0(\sigma)$	$\frac{1}{2}$	$\frac{1.0001}{2}$	2σ	$\frac{1}{(4\sigma^2+1)^{1/2}}$	0.98	0.46	0.053
$\beta_2(\sigma)$	$\frac{1}{2}$	$\frac{1.0001}{2}$	6σ	$\frac{1}{(4\sigma^2+1)^{3/2}}$	0.94	0.11	0.00017
$\rho(\sigma)$	0.091	0.089	0.27σ	$\rho_{\text{erf}}(\sigma)^1$	1.8E-4	4.9E-2	5.1E-5

Table 1: $\rho(\sigma)$ values for different activation functions. Note that we can calculate the exact values for ReLU, Leaky ReLU, squared ReLU and erf. We can't find a closed-form value for sigmoid or tanh, but we calculate the numerical values of $\rho(\sigma)$ for $\sigma = 0.1, 1, 10$. ¹ $\rho_{\text{erf}}(\sigma) = \min\{(4\sigma^2 + 1)^{-1/2} - (2\sigma^2 + 1)^{-1}, (4\sigma^2 + 1)^{-3/2} - (2\sigma^2 + 1)^{-3}, (2\sigma^2 + 1)^{-2}\}$

Smooth non-decreasing activations with bounded first derivatives automatically satisfy Property 3.1 and 3.3. For Property 3.2, since their first derivatives are symmetric, we have $\mathbb{E}[\phi'(\sigma \cdot z)z] = 0$. Then by Hölder's inequality and $\phi'(z) \geq 0$, we have

$$\begin{aligned} \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'^2(\sigma \cdot z)] &\geq \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma \cdot z)] \right)^2, \\ \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'^2(\sigma \cdot z)z^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [z^2] &\geq \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma \cdot z)z^2] \right)^2, \\ \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma \cdot z)z^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma \cdot z)] &= \mathbb{E}_{z \sim \mathcal{D}_1} [(\sqrt{\phi'(\sigma \cdot z)z^2})^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [(\sqrt{\phi'(\sigma \cdot z)})^2] \geq \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma \cdot z)z] \right)^2. \end{aligned}$$

The equality in the first inequality happens when $\phi'(\sigma \cdot z)$ is a constant a.e.. The equality in the second inequality happens when $|\phi'(\sigma \cdot z)|$ is a constant a.e., which is invalidated by the non-linearity and smoothness condition. The equality in the third inequality holds only when $\phi'(z) = 0$ a.e., which leads to a constant function under non-decreasing condition. Therefore, $\rho(\sigma) > 0$ for any smooth non-decreasing non-linear activations with bounded symmetric first derivatives. The statements about linear activations and quadratic activation follow direct calculations. \square

D Local Positive Definiteness of Hessian

D.1 Main Results for Positive Definiteness of Hessian

D.1.1 Bounding the Spectrum of the Hessian near the Ground Truth

Theorem D.1 (Bounding the spectrum of the Hessian near the ground truth). *For any $W \in \mathbb{R}^{d \times k}$ with $\|W - W^*\| \lesssim v_{\min}^4 \rho^2(\sigma_k) / (k^2 \kappa^5 \lambda^2 v_{\max}^4 \sigma_1^{4p}) \cdot \|W^*\|$, let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)) and let the activation function satisfy Property 3.1, 3.2, 3.3. Then for any $t \geq 1$, if $|S| \geq d \cdot \text{poly}(\log d, t) \cdot k^2 v_{\max}^4 \tau \kappa^8 \lambda^2 \sigma_1^{4p} / (v_{\min}^4 \rho^2(\sigma_k))$, we have with probability at least $1 - d^{-\Omega(t)}$,*

$$\Omega(v_{\min}^2 \rho(\sigma_k) / (\kappa^2 \lambda)) I \preceq \nabla^2 \widehat{f}_S(W) \preceq O(kv_{\max}^2 \sigma_1^{2p}) I.$$

Proof. The main idea of the proof follows the following inequalities,

$$\nabla^2 f_{\mathcal{D}}(W^*) - \|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| I \preceq \nabla^2 \widehat{f}_S(W) \preceq \nabla^2 f_{\mathcal{D}}(W^*) + \|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| I$$

The proof sketch is first to bound the range of the eigenvalues of $\nabla^2 f_{\mathcal{D}}(W^*)$ (Lemma D.3) and then bound the spectral norm of the remaining error, $\|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\|$. $\|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\|$ can be further decomposed into two parts, $\|\nabla^2 \widehat{f}_S(W) - H\|$ and $\|H - \nabla^2 f_{\mathcal{D}}(W^*)\|$, where H is $\nabla^2 f_{\mathcal{D}}(W)$ if ϕ is smooth, otherwise H is a specially designed matrix. We can upper bound them when W is close enough to W^* and there are enough samples. In particular, if the activation satisfies Property 3.3(a), see Lemma D.10 for bounding $\|H - \nabla^2 f_{\mathcal{D}}(W^*)\|$ and Lemma D.11 for bounding $\|H - \nabla^2 \widehat{f}_S(W)\|$. If the activation satisfies Property 3.3(b), see Lemma D.15.

Finally we can complete the proof by setting $\delta = O(v_{\min}^2 \rho(\sigma_1) / (kv_{\max}^2 \kappa^2 \lambda \sigma_1^{2p}))$ in Lemma D.11 and Lemma D.15, setting $\|W - W^*\| \lesssim v_{\min}^2 \rho(\sigma_k) / (k^2 \kappa^2 \lambda v_{\max}^2 \sigma_1^p)$ in Lemma D.10 and setting $\|W - W^*\| \leq v_{\min}^4 \rho^2(\sigma_k) \sigma_k / (k^2 \kappa^4 \lambda^2 v_{\max}^4 \sigma_1^{4p})$ in Lemma D.15. \square

D.1.2 Local Linear Convergence of Gradient Descent

Although Theorem D.1 gives upper and lower bounds for the spectrum of the Hessian w.h.p., it only holds when the current set of parameters W are independent of samples. When we use iterative methods, like gradient descent, to optimize this objective, the next iterate calculated from the current set of samples will depend on this set of samples. Therefore, we need to do resampling at each iteration. Here we show that for activations that satisfies Properties 3.1, 3.2 and 3.3(a), linear convergence of gradient descent is guaranteed. To the best of our knowledge, there is no linear convergence guarantees for general non-smooth objective. So the following proposition also applies to smooth objectives only, which excludes ReLU.

Theorem D.2 (Linear convergence of gradient descent, formal version of Theorem 4.4). *Let $W^c \in \mathbb{R}^{d \times k}$ be the current iterate satisfying*

$$\|W^c - W^*\| \lesssim v_{\min}^4 \rho^2(\sigma_k) / (k^2 \kappa^5 \lambda^2 v_{\max}^4 \sigma_1^{4p}) \|W^*\|.$$

Let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)) Let the activation function satisfy Property 3.1, 3.2 and 3.3(a). Define

$$m_0 = \Theta(v_{\min}^2 \rho(\sigma_k) / (\kappa^2 \lambda)) \text{ and } M_0 = \Theta(kv_{\max}^2 \sigma_1^{2p}).$$

For any $t \geq 1$, if we choose

$$|S| \geq d \cdot \text{poly}(\log d, t) \cdot k^2 v_{\max}^4 \tau \kappa^8 \lambda^2 \sigma_1^{4p} / (v_{\min}^4 \rho^2(\sigma_k)) \quad (5)$$

and perform gradient descent with step size $1/M_0$ on $\widehat{f}_S(W^c)$ and obtain the next iterate,

$$\widetilde{W} = W^c - \frac{1}{M_0} \nabla \widehat{f}_S(W^c),$$

then with probability at least $1 - d^{-\Omega(t)}$,

$$\|\widetilde{W} - W^*\|_F^2 \leq (1 - \frac{m_0}{M_0}) \|W^c - W^*\|_F^2.$$

Proof. To prove Theorem D.2, we need to show the positive definite properties on the entire line between the current iterate and the optimum by constructing a set of anchor points, which are independent of the samples. Then we apply traditional analysis for the linear convergence of gradient descent.

In particular, given a current iterate W^c , we set $d^{(p+1)/2}$ anchor points $\{W^a\}_{a=1,2,\dots,d^{(p+1)/2}}$ equally along the line $\xi W^* + (1 - \xi)W^c$ for $\xi \in [0, 1]$.

According to Theorem D.1, by setting $t \leftarrow t + (p + 1)/2$, we have with probability at least $1 - d^{-(t+(p+1)/2)}$ for each anchor point $\{W^a\}$,

$$m_0 I \preceq \nabla^2 \widehat{f}_S(W^a) \preceq M_0 I.$$

Then given an anchor point W^a , according to Lemma D.16, we have with probability $1 - 2d^{-(t+(p+1)/2)}$, for any points W between $(W^{a-1} + W^a)/2$ and $(W^a + W^{a+1})/2$,

$$m_0 I \preceq \nabla^2 \widehat{f}_S(W) \preceq M_0 I. \quad (6)$$

Finally by applying union bound over these $d^{(p+1)/2}$ small intervals, we have with probability at least $1 - d^{-t}$ for any points W on the line between W^c and W^* ,

$$m_0 I \preceq \nabla^2 \widehat{f}_S(W) \preceq M_0 I.$$

Now we can apply traditional analysis for linear convergence of gradient descent. Let η denote the stepsize.

$$\begin{aligned} & \|\widetilde{W} - W^*\|_F^2 \\ &= \|W^c - \eta \nabla \widehat{f}_S(W^c) - W^*\|_F^2 \\ &= \|W^c - W^*\|_F^2 - 2\eta \langle \nabla \widehat{f}_S(W^c), (W^c - W^*) \rangle + \eta^2 \|\nabla \widehat{f}_S(W^c)\|_F^2 \end{aligned}$$

We can rewrite $\widehat{f}_S(W^c)$,

$$\nabla \widehat{f}_S(W^c) = \left(\int_0^1 \nabla^2 \widehat{f}_S(W^* + \gamma(W^c - W^*)) d\gamma \right) \text{vec}(W^c - W^*).$$

We define function $\widehat{H}_S : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}^{dk \times dk}$ such that

$$\widehat{H}_S(W^c - W^*) = \left(\int_0^1 \nabla^2 \widehat{f}_S(W^* + \gamma(W^c - W^*)) d\gamma \right).$$

According to Eq. (6),

$$m_0 I \preceq \widehat{H} \preceq M_0 I. \quad (7)$$

We can upper bound $\|\nabla \widehat{f}_S(W^c)\|_F^2$,

$$\|\nabla \widehat{f}_S(W^c)\|_F^2 = \langle \widehat{H}_S(W^c - W^*), \widehat{H}_S(W^c - W^*) \rangle \leq M_0 \langle W^c - W^*, \widehat{H}_S(W^c - W^*) \rangle.$$

Therefore,

$$\begin{aligned} & \|\widetilde{W} - W^*\|_F^2 \\ & \leq \|W^c - W^*\|_F^2 - (-\eta^2 M_0 + 2\eta) \langle W^c - W^*, \widehat{H}(W^c - W^*) \rangle \\ & \leq \|W^c - W^*\|_F^2 - (-\eta^2 M_0 + 2\eta) m_0 \|W^c - W^*\|_F^2 \\ & = \|W^c - W^*\|_F^2 - \frac{m_0}{M_0} \|W^c - W^*\|_F^2 \\ & \leq \left(1 - \frac{m_0}{M_0}\right) \|W^c - W^*\|_F^2 \end{aligned}$$

where the third equality holds by setting $\eta = \frac{1}{M_0}$. \square

D.2 Positive Definiteness of Population Hessian at the Ground Truth

The goal of this Section is to prove Lemma D.3.

Lemma D.3 (Positive definiteness of population Hessian at the ground truth). *If $\phi(z)$ satisfies Property 3.1, 3.2 and 3.3, we have the following property for the second derivative of function $f_{\mathcal{D}}(W)$ at W^* ,*

$$\Omega(v_{\min}^2 \rho(\sigma_k) / (\kappa^2 \lambda)) I \preceq \nabla^2 f_{\mathcal{D}}(W^*) \preceq O(kv_{\max}^2 \sigma_1^{2p}) I.$$

Proof. The proof directly follows Lemma D.6 (Section D.2.2) and Lemma D.7 (Section D.2.3). \square

D.2.1 Lower Bound on the Eigenvalues of Hessian for the Orthogonal Case

Lemma D.4. Let \mathcal{D}_1 denote Gaussian distribution $\mathcal{N}(0, 1)$. Let $\alpha_0 = \mathbb{E}_{z \sim \mathcal{D}_1}[\phi'(z)]$, $\alpha_1 = \mathbb{E}_{z \sim \mathcal{D}_1}[\phi'(z)z]$, $\alpha_2 = \mathbb{E}_{z \sim \mathcal{D}_1}[\phi'(z)z^2]$, $\beta_0 = \mathbb{E}_{z \sim \mathcal{D}_1}[\phi'^2(z)]$, $\beta_2 = \mathbb{E}_{z \sim \mathcal{D}_1}[\phi'^2(z)z^2]$. Let ρ denote $\min\{(\beta_0 - \alpha_0^2 - \alpha_1^2), (\beta_2 - \alpha_1^2 - \alpha_2^2)\}$. Let $P = [p_1 \ p_2 \ \cdots \ p_k] \in \mathbb{R}^{k \times k}$. Then we have,

$$\mathbb{E}_{u \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k p_i^\top u \cdot \phi'(u_i) \right)^2 \right] \geq \rho \|P\|_F^2 \quad (8)$$

Proof. The main idea is to explicitly calculate the LHS of Eq (8), then reformulate the equation and find a lower bound represented by $\alpha_0, \alpha_1, \alpha_2, \beta_0, \beta_2$.

$$\begin{aligned} & \mathbb{E}_{u \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k p_i^\top u \cdot \phi'(u_i) \right)^2 \right] \\ &= \sum_{i=1}^k \sum_{l=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} [p_i^\top (\phi'(u_i)\phi'(u_l) \cdot uu^\top) p_l] \\ &= \underbrace{\sum_{i=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} [p_i^\top (\phi'(u_i)^2 \cdot uu^\top) p_i]}_A + \underbrace{\sum_{i \neq l} \mathbb{E}_{u \sim \mathcal{D}_k} [p_i^\top (\phi'(u_i)\phi'(u_l) \cdot uu^\top) p_l]}_B \end{aligned}$$

Further, we can rewrite the diagonal term in the following way,

$$\begin{aligned} A &= \sum_{i=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} [p_i^\top (\phi'(u_i)^2 \cdot uu^\top) p_i] \\ &= \sum_{i=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} \left[p_i^\top \left(\phi'(u_i)^2 \cdot \left(u_i^2 e_i e_i^\top + \sum_{j \neq i} u_i u_j (e_i e_j^\top + e_j e_i^\top) + \sum_{j \neq i} \sum_{l \neq i} u_j u_l e_j e_l^\top \right) \right) p_i \right] \\ &= \sum_{i=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} \left[p_i^\top \left(\phi'(u_i)^2 \cdot \left(u_i^2 e_i e_i^\top + \sum_{j \neq i} u_j^2 e_j e_j^\top \right) \right) p_i \right] \\ &= \sum_{i=1}^k \left[p_i^\top \left(\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_i^2] e_i e_i^\top + \sum_{j \neq i} \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_j^2] e_j e_j^\top \right) p_i \right] \\ &= \sum_{i=1}^k \left[p_i^\top \left(\beta_2 e_i e_i^\top + \sum_{j \neq i} \beta_0 e_j e_j^\top \right) p_i \right] \\ &= \sum_{i=1}^k p_i^\top ((\beta_2 - \beta_0) e_i e_i^\top + \beta_0 I_k) p_i \\ &= (\beta_2 - \beta_0) \sum_{i=1}^k p_i^\top e_i e_i^\top p_i + \beta_0 \sum_{i=1}^k p_i^\top p_i \\ &= (\beta_2 - \beta_0) \|\text{diag}(P)\|^2 + \beta_0 \|P\|_F^2, \end{aligned}$$

where the second step follows by rewriting $uu^\top = \sum_{i=1}^k \sum_{j=1}^k u_i u_j e_i e_j^\top$, the third step follows by

$\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_i u_j] = 0, \forall j \neq i$ and $\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_j u_l] = 0, \forall j \neq l$, the fourth step follows by pushing expectation, the fifth step follows by $\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_i^2] = \beta_2$ and $\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2 u_j^2] = \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_i)^2] = \beta_0$,

and the last step follows by $\sum_{i=1}^k p_{i,i}^2 = \|\text{diag}(P)\|^2$ and $\sum_{i=1}^k p_i^\top p_i = \sum_{i=1}^k \|p_i\|^2 = \|P\|_F^2$.

We can rewrite the off-diagonal term in the following way,

$$\begin{aligned}
B &= \sum_{i \neq l} \mathbb{E}_{u \sim \mathcal{D}_k} [p_i^\top (\phi'(u_l) \phi'(u_i) \cdot uu^\top) p_l] \\
&= \sum_{i \neq l} \mathbb{E}_{u \sim \mathcal{D}_k} \left[p_i^\top \left(\phi'(u_l) \phi'(u_i) \cdot \left(u_i^2 e_i e_i^\top + u_l^2 e_l e_l^\top + u_i u_l (e_i e_l^\top + e_l e_i^\top) + \sum_{j \neq l} u_i u_j e_i e_j^\top \right. \right. \right. \\
&\quad \left. \left. + \sum_{j \neq i} u_j u_l e_j e_l^\top + \sum_{j \neq i, l} \sum_{j' \neq i, l} u_j u_{j'} e_j e_{j'}^\top \right) \right] p_l \\
&= \sum_{i \neq l} \mathbb{E}_{u \sim \mathcal{D}_k} \left[p_i^\top \left(\phi'(u_l) \phi'(u_i) \cdot \left(u_i^2 e_i e_i^\top + u_l^2 e_l e_l^\top + u_i u_l (e_i e_l^\top + e_l e_i^\top) + \sum_{j \neq i, l} u_j^2 e_j e_j^\top \right) \right) p_l \right] \\
&= \sum_{i \neq l} \left[p_i^\top \left(\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_i^2] e_i e_i^\top + \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_l^2] e_l e_l^\top \right. \right. \\
&\quad \left. \left. + \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_i u_l] (e_i e_l^\top + e_l e_i^\top) + \sum_{j \neq i, l} \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_j^2] e_j e_j^\top \right) p_l \right] \\
&= \sum_{i \neq l} \left[p_i^\top \left(\alpha_0 \alpha_2 (e_i e_i^\top + e_l e_l^\top) + \alpha_1^2 (e_i e_l^\top + e_l e_i^\top) + \sum_{j \neq i, l} \alpha_0^2 e_j e_j^\top \right) p_l \right] \\
&= \sum_{i \neq l} \left[p_i^\top \left((\alpha_0 \alpha_2 - \alpha_0^2) (e_i e_i^\top + e_l e_l^\top) + \alpha_1^2 (e_i e_l^\top + e_l e_i^\top) + \alpha_0^2 I_k \right) p_l \right] \\
&= \underbrace{(\alpha_0 \alpha_2 - \alpha_0^2) \sum_{i \neq l} p_i^\top (e_i e_i^\top + e_l e_l^\top) p_l}_{B_1} + \underbrace{\alpha_1^2 \sum_{i \neq l} p_i^\top (e_i e_l^\top + e_l e_i^\top) p_l}_{B_2} + \underbrace{\alpha_0^2 \sum_{i \neq l} p_i^\top p_l}_{B_3},
\end{aligned}$$

where the third step follows by $\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_i u_j] = 0$ and $\mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(u_l) \phi'(u_i) u_j u_{j'}] = 0$ for $j' \neq j$.

For the term B_1 , we have

$$\begin{aligned}
B_1 &= (\alpha_0\alpha_2 - \alpha_0^2) \sum_{i \neq l} p_i^\top (e_i e_i^\top + e_l e_l^\top) p_l \\
&= 2(\alpha_0\alpha_2 - \alpha_0^2) \sum_{i \neq l} p_i^\top e_i e_i^\top p_l \\
&= 2(\alpha_0\alpha_2 - \alpha_0^2) \sum_{i=1}^k p_i^\top e_i e_i^\top \left(\sum_{l=1}^k p_l - p_i \right) \\
&= 2(\alpha_0\alpha_2 - \alpha_0^2) \left(\sum_{i=1}^k p_i^\top e_i e_i^\top \sum_{l=1}^k p_l - \sum_{i=1}^k p_i^\top e_i e_i^\top p_i \right) \\
&= 2(\alpha_0\alpha_2 - \alpha_0^2) (\text{diag}(P)^\top \cdot P \cdot \mathbf{1} - \|\text{diag}(P)\|^2)
\end{aligned}$$

For the term B_2 , we have

$$\begin{aligned}
B_2 &= \alpha_1^2 \sum_{i \neq l} p_i^\top (e_i e_l^\top + e_l e_i^\top) p_l \\
&= \alpha_1^2 \left(\sum_{i \neq l} p_i^\top e_i e_l^\top p_l + \sum_{i \neq l} p_i^\top e_l e_i^\top p_l \right) \\
&= \alpha_1^2 \left(\sum_{i=1}^k \sum_{l=1}^k p_i^\top e_i e_l^\top p_l - \sum_{j=1}^k p_j^\top e_j e_j^\top p_j + \sum_{i=1}^k \sum_{l=1}^k p_i^\top e_l e_i^\top p_l - \sum_{j=1}^k p_j^\top e_j e_j^\top p_j \right) \\
&= \alpha_1^2 ((\text{diag}(P)^\top \mathbf{1})^2 - \|\text{diag}(P)\|^2 + \langle P, P^\top \rangle - \|\text{diag}(P)\|^2)
\end{aligned}$$

For the term B_3 , we have

$$\begin{aligned}
B_3 &= \alpha_0^2 \sum_{i \neq l} p_i^\top p_l \\
&= \alpha_0^2 \left(\sum_{i=1}^k p_i^\top \sum_{l=1}^k p_l - \sum_{i=1}^k p_i^\top p_i \right) \\
&= \alpha_0^2 \left(\left\| \sum_{i=1}^k p_i \right\|^2 - \sum_{i=1}^k \|p_i\|^2 \right) \\
&= \alpha_0^2 (\|P \cdot \mathbf{1}\|^2 - \|P\|_F^2)
\end{aligned}$$

Let $\text{diag}(P)$ denote a length k column vector where the i -th entry is the (i, i) -th entry of

$P \in \mathbb{R}^{k \times k}$. Furthermore, we can show $A + B$ is,

$$\begin{aligned}
& A + B \\
&= A + B_1 + B_2 + B_3 \\
&= \underbrace{(\beta_2 - \beta_0)\|\text{diag}(P)\|^2 + \beta_0\|P\|_F^2}_{A} + \underbrace{2(\alpha_0\alpha_2 - \alpha_0^2)(\text{diag}(P)^\top \cdot P \cdot \mathbf{1} - \|\text{diag}(P)\|^2)}_{B_1} \\
&+ \underbrace{\alpha_1^2((\text{diag}(P)^\top \cdot \mathbf{1})^2 - \|\text{diag}(P)\|^2 + \langle P, P^\top \rangle - \|\text{diag}(P)\|^2)}_{B_2} + \underbrace{\alpha_0^2(\|P \cdot \mathbf{1}\|^2 - \|P\|_F^2)}_{B_3} \\
&= \underbrace{\|\alpha_0 P \cdot \mathbf{1} + (\alpha_2 - \alpha_0)\text{diag}(P)\|^2}_{C_1} + \underbrace{\alpha_1^2(\text{diag}(P)^\top \cdot \mathbf{1})^2}_{C_2} + \underbrace{\frac{\alpha_1^2}{2}\|P + P^\top - 2\text{diag}(\text{diag}(P))\|_F^2}_{C_3} \\
&+ \underbrace{(\beta_0 - \alpha_0^2 - \alpha_1^2)\|P - \text{diag}(\text{diag}(P))\|_F^2}_{C_4} + \underbrace{(\beta_2 - \alpha_1^2 - \alpha_2^2)\|\text{diag}(P)\|^2}_{C_5} \\
&\geq (\beta_0 - \alpha_0^2 - \alpha_1^2)\|P - \text{diag}(\text{diag}(P))\|_F^2 + (\beta_2 - \alpha_1^2 - \alpha_2^2)\|\text{diag}(P)\|^2 \\
&\geq \min\{(\beta_0 - \alpha_0^2 - \alpha_1^2), (\beta_2 - \alpha_1^2 - \alpha_2^2)\} \cdot (\|P - \text{diag}(\text{diag}(P))\|_F^2 + \|\text{diag}(P)\|^2) \\
&= \min\{(\beta_0 - \alpha_0^2 - \alpha_1^2), (\beta_2 - \alpha_1^2 - \alpha_2^2)\} \cdot (\|P - \text{diag}(\text{diag}(P))\|_F^2 + \|\text{diag}(\text{diag}(P))\|^2) \\
&\geq \min\{(\beta_0 - \alpha_0^2 - \alpha_1^2), (\beta_2 - \alpha_1^2 - \alpha_2^2)\} \cdot \|P\|_F^2 \\
&= \rho\|P\|_F^2,
\end{aligned}$$

where the first step follows by $B = B_1 + B_2 + B_3$, and the second step follows by the definition of A, B_1, B_2, B_3 the third step follows by $A + B_1 + B_2 + B_3 = C_1 + C_2 + C_3 + C_4 + C_5$, the fourth step follows by $C_1, C_2, C_3 \geq 0$, the fifth step follows $a \geq \min(a, b)$, the sixth step follows by $\|\text{diag}(P)\|^2 = \|\text{diag}(\text{diag}(P))\|^2$, the seventh step follows by triangle inequality, and the last step follows the definition of ρ . \square

Claim D.5. $A + B_1 + B_2 + B_3 = C_1 + C_2 + C_3 + C_4 + C_5$.

Proof. The key properties we need are, for two vectors a, b , $\|a + b\|^2 = \|a\|^2 + 2\langle a, b \rangle + \|b\|^2$; for two

matrices A, B , $\|A + B\|_F^2 = \|A\|_F^2 + 2\langle A, B \rangle + \|B\|_F^2$. Then, we have

$$\begin{aligned}
& C_1 + C_2 + C_3 + C_4 + C_5 \\
&= \underbrace{(\|\alpha_0 P \cdot \mathbf{1}\|^2 + 2(\alpha_0 \alpha_2 - \alpha_0^2) \langle P \cdot \mathbf{1}, \text{diag}(P) \rangle + (\alpha_2 - \alpha_0)^2 \|\text{diag}(P)\|^2)}_{C_1} + \underbrace{\alpha_1^2 (\text{diag}(P)^\top \cdot \mathbf{1})^2}_{C_2} \\
&+ \underbrace{\frac{\alpha_1^2}{2} (2\|P\|_F^2 + 4\|\text{diag}(\text{diag}(P))\|_F^2 + 2\langle P, P^\top \rangle - 4\langle P, \text{diag}(\text{diag}(P)) \rangle - 4\langle P^\top, \text{diag}(\text{diag}(P)) \rangle)}_{C_3} \\
&+ \underbrace{(\beta_0 - \alpha_0^2 - \alpha_1^2) (\|P\|_F^2 - 2\langle P, \text{diag}(\text{diag}(P)) \rangle) + \|\text{diag}(\text{diag}(P))\|_F^2}_{C_4} + \underbrace{(\beta_2 - \alpha_1^2 - \alpha_2^2) \|\text{diag}(P)\|^2}_{C_5} \\
&= \underbrace{\alpha_0^2 \|P \cdot \mathbf{1}\|^2 + 2(\alpha_0 \alpha_2 - \alpha_0^2) \langle P \cdot \mathbf{1}, \text{diag}(P) \rangle + (\alpha_2 - \alpha_0)^2 \|\text{diag}(P)\|^2}_{C_1} + \underbrace{\alpha_1^2 (\text{diag}(P)^\top \cdot \mathbf{1})^2}_{C_2} \\
&+ \underbrace{\frac{\alpha_1^2}{2} (2\|P\|_F^2 + 4\|\text{diag}(P)\|^2 + 2\langle P, P^\top \rangle - 8\|\text{diag}(P)\|^2)}_{C_3} \\
&+ \underbrace{(\beta_0 - \alpha_0^2 - \alpha_1^2) (\|P\|_F^2 - 2\|\text{diag}(P)\|^2 + \|\text{diag}(P)\|^2)}_{C_4} + \underbrace{(\beta_2 - \alpha_1^2 - \alpha_2^2) \|\text{diag}(P)\|^2}_{C_5} \\
&= \alpha_0^2 \|P \cdot \mathbf{1}\|^2 + 2(\alpha_0 \alpha_2 - \alpha_0^2) \text{diag}(P)^\top \cdot P \cdot \mathbf{1} + \alpha_1^2 (\text{diag}(P)^\top \cdot \mathbf{1})^2 + \alpha_1^2 \langle P, P^\top \rangle \\
&+ (\beta_0 - \alpha_0^2) \|P\|_F^2 + \underbrace{((\alpha_2 - \alpha_0)^2 - 2\alpha_1^2 - \beta_0 + \alpha_0^2 + \alpha_1^2 + \beta_2 - \alpha_1^2 - \alpha_2^2) \|\text{diag}(P)\|^2}_{\beta_2 - \beta_0 - 2(\alpha_2 \alpha_0 - \alpha_0^2 + \alpha_1^2)} \\
&= \underbrace{0}_{\text{part of } A} + \underbrace{2(\alpha_2 \alpha_0 - \alpha_0^2) \cdot \text{diag}(P)^\top P \cdot \mathbf{1}}_{\text{part of } B_1} + \underbrace{\alpha_1^2 \cdot ((\text{diag}(P)^\top \mathbf{1})^2 + \langle P, P^\top \rangle)}_{\text{part of } B_2} + \underbrace{\alpha_0^2 \cdot \|P \cdot \mathbf{1}\|^2}_{\text{part of } B_3} \\
&+ \underbrace{(\beta_0 - \alpha_0^2) \cdot \|P\|_F^2}_{\text{proportional to } \|P\|_F^2} + \underbrace{(\beta_2 - \beta_0 - 2(\alpha_2 \alpha_0 - \alpha_0^2 + \alpha_1^2)) \cdot \|\text{diag}(P)\|^2}_{\text{proportional to } \|\text{diag}(P)\|^2} \\
&= \underbrace{(\beta_2 - \beta_0) \|\text{diag}(P)\|^2 + \beta_0 \|P\|_F^2}_A + \underbrace{2(\alpha_0 \alpha_2 - \alpha_0^2) (\text{diag}(P)^\top \cdot P \cdot \mathbf{1} - \|\text{diag}(P)\|^2)}_{B_1} \\
&+ \underbrace{\alpha_1^2 ((\text{diag}(P)^\top \cdot \mathbf{1})^2 - \|\text{diag}(P)\|^2 + \langle P, P^\top \rangle - \|\text{diag}(P)\|^2)}_{B_2} + \underbrace{\alpha_0^2 (\|P \cdot \mathbf{1}\|^2 - \|P\|_F^2)}_{B_3} \\
&= A + B_1 + B_2 + B_3
\end{aligned}$$

where the second step follows by $\langle P, \text{diag}(\text{diag}(P)) \rangle = \|\text{diag}(P)\|^2$ and $\|\text{diag}(\text{diag}(P))\|_F^2 = \|\text{diag}(P)\|^2$. \square

D.2.2 Lower Bound on the Eigenvalues of Hessian for Non-orthogonal Case

First we show the lower bound of the eigenvalues. The main idea is to reduce the problem to a k -by- k problem and then lower bound the eigenvalues using orthogonal weight matrices.

Lemma D.6 (Lower bound). *If $\phi(z)$ satisfies Property 3.1, 3.2 and 3.3, we have the following property for the second derivative of function $f_{\mathcal{D}}(W)$ at W^* ,*

$$\Omega(v_{\min}^2 \rho(\sigma_k) / (\kappa^2 \lambda)) I \preceq \nabla^2 f_{\mathcal{D}}(W^*).$$

Proof. Let $a \in \mathbb{R}^{dk}$ denote vector $[a_1^\top \ a_2^\top \ \cdots \ a_k^\top]^\top$, let $b \in \mathbb{R}^{dk}$ denote vector $[b_1^\top \ b_2^\top \ \cdots \ b_k^\top]^\top$ and let $c \in \mathbb{R}^{dk}$ denote vector $[c_1^\top \ c_2^\top \ \cdots \ c_k^\top]^\top$. The smallest eigenvalue of the Hessian can be calculated by

$$\nabla^2 f(W^*) \succeq \min_{\|a\|=1} a^\top \nabla^2 f(W^*) a I_{dk} = \min_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k v_i^* a_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] I_{dk} \quad (9)$$

Note that

$$\begin{aligned} & \min_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k (v_i^* a_i)^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] \\ &= \min_{\|a\| \neq 0} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k (v_i^* a_i)^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] / \|a\|^2 \\ &= \min_{\sum_i \|b_i/v_i^*\|^2 \neq 0} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k b_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] / \left(\sum_{i=1}^k \|b_i/v_i^*\|^2 \right) \quad \text{by } a_i = b_i/v_i^* \\ &= \min_{\sum_i \|b_i\|^2 \neq 0} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k b_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] / \left(\sum_{i=1}^k \|b_i/v_i^*\|^2 \right) \\ &\geq v_{\min}^2 \min_{\sum_i \|b_i\|^2 \neq 0} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k b_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] / \left(\sum_{i=1}^k \|b_i\|^2 \right) \quad \text{by } v_{\min} = \min_{i \in [k]} |v_i^*| \\ &= v_{\min}^2 \min_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k a_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] \end{aligned} \quad (10)$$

Let $U \in \mathbb{R}^{d \times k}$ be the orthonormal basis of W^* and let $V = [v_1 \ v_2 \ \cdots \ v_k] = U^\top W^* \in \mathbb{R}^{k \times k}$. Also note that V and W^* have same singular values and $W^* = UV$. We use $U_\perp \in \mathbb{R}^{d \times (d-k)}$ to denote the complement of U . For any vector $a_j \in \mathbb{R}^d$, there exist two vectors $b_j \in \mathbb{R}^k$ and $c_j \in \mathbb{R}^{d-k}$ such that

$$a_j = Ub_j + U_\perp c_j.$$

We use \mathcal{D}_d to denote Gaussian distribution $\mathcal{N}(0, I_d)$, \mathcal{D}_{d-k} to denote Gaussian distribution $\mathcal{N}(0, I_{d-k})$, and \mathcal{D}_k to denote Gaussian distribution $\mathcal{N}(0, I_k)$. Then we can rewrite formulation (10) (removing v_{\min}^2) as

$$\mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k a_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] = \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k (b_i^\top U^\top + c_i^\top U_\perp^\top) x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] = A + B + C$$

where

$$\begin{aligned}
A &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k b_i^\top U^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right], \\
B &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k c_i^\top U_\perp^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right], \\
C &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[2 \left(\sum_{i=1}^k b_i^\top U^\top x \cdot \phi'(w_i^{*\top} x) \right) \cdot \left(\sum_{i=1}^k c_i^\top U_\perp^\top x \cdot \phi'(w_i^{*\top} x) \right) \right].
\end{aligned}$$

We calculate A, B, C separately. First, we can show

$$A = \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k b_i^\top U^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] = \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k b_i^\top z \cdot \phi'(v_i^\top z) \right)^2 \right].$$

Second, we can show

$$\begin{aligned}
B &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k c_i^\top U_\perp^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] \\
&= \mathbb{E}_{s \sim \mathcal{D}_{d-k}, z \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k c_i^\top s \cdot \phi'(v_i^{*\top} z) \right)^2 \right] \\
&= \mathbb{E}_{s \sim \mathcal{D}_{d-k}, z \sim \mathcal{D}_k} [(y^\top s)^2] && \text{by defining } y = \sum_{i=1}^k \phi'(v_i^{*\top} z) c_i \in \mathbb{R}^{d-k} \\
&= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\mathbb{E}_{s \sim \mathcal{D}_{d-k}} [(y^\top s)^2] \right] \\
&= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\mathbb{E}_{s \sim \mathcal{D}_{d-k}} \left[\sum_{j=1}^{d-k} s_j^2 y_j^2 \right] \right] && \text{by } \mathbb{E}[s_j s_{j'}] = 0 \\
&= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\sum_{j=1}^{d-k} y_j^2 \right] && \text{by } s_j \sim \mathcal{N}(0, 1) \\
&= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left\| \sum_{i=1}^k \phi'(v_i^{*\top} z) c_i \right\|^2 \right] && \text{by definition of } y
\end{aligned}$$

Third, we have $C = 0$ since $U_\perp^\top x$ is independent of $w_i^{*\top} x$ and $U^\top x$. Thus, putting them all together,

$$\mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{i=1}^k a_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] = \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k b_i^\top z \cdot \phi'(v_i^\top z) \right)^2 \right] + \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left\| \sum_{i=1}^k \phi'(v_i^\top z) c_i \right\|^2 \right]$$

Let us lower bound A ,

$$\begin{aligned}
A &= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k b_i^\top z \cdot \phi'(v_i^\top z) \right)^2 \right] \\
&= \int (2\pi)^{-k/2} \left(\sum_{i=1}^k b_i^\top z \cdot \phi'(v_i^\top z) \right)^2 e^{-\|z\|^2/2} dz \\
&= \int_{\xi_1} (2\pi)^{-k/2} \left(\sum_{i=1}^k b_i^\top V^{\dagger\top} s \cdot \phi'(s_i) \right)^2 e^{-\|V^{\dagger\top} s\|^2/2} \cdot |\det(V^\dagger)| ds \\
&\geq \int_{\xi_2} (2\pi)^{-k/2} \left(\sum_{i=1}^k b_i^\top V^{\dagger\top} s \cdot \phi'(s_i) \right)^2 e^{-\sigma_1^2(V^\dagger) \|s\|^2/2} \cdot |\det(V^\dagger)| ds \\
&= \int_{\xi_3} (2\pi)^{-k/2} \left(\sum_{i=1}^k b_i^\top V^{\dagger\top} u / \sigma_1(V^\dagger) \cdot \phi'(u_i / \sigma_1(V^\dagger)) \right)^2 e^{-\|u\|^2/2} |\det(V^\dagger)| / \sigma_1^k(V^\dagger) du \\
&= \int (2\pi)^{-k/2} \left(\sum_{i=1}^k p_i^\top u \cdot \phi'(\sigma_k \cdot u_i) \right)^2 e^{-\|u\|^2/2} \frac{1}{\lambda} du \\
&= \frac{1}{\lambda} \mathbb{E}_{u \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k p_i^\top u \cdot \phi'(\sigma_k \cdot u_i) \right)^2 \right]
\end{aligned}$$

where $V^\dagger \in \mathbb{R}^{k \times k}$ is the inverse of $V \in \mathbb{R}^{k \times k}$, i.e., $V^\dagger V = I$, $p_i^\top = b_i^\top V^{\dagger\top} / \sigma_1(V^\dagger)$ and $\sigma_k = \sigma_k(W^*)$. ξ_1 replaces z by $z = V^{\dagger\top} s$, so $v_i^{\top} z = s_i$. ξ_2 uses the fact $\|V^{\dagger\top} s\| \leq \sigma_1(V^\dagger) \|s\|$. ξ_3 replaces s by $s = u / \sigma_1(V^\dagger)$. Note that $\phi'(\sigma_k \cdot u_i)$'s are independent of each other, so we can simplify the analysis. In particular, Lemma D.4 gives a lower bound in this case in terms of p_i . Note that $\|p_i\| \geq \|b_i\| / \kappa$. Therefore,

$$\mathbb{E}_{z \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k b_i^\top z \cdot \phi'(v_i^\top z) \right)^2 \right] \geq \rho(\sigma_k) \frac{1}{\kappa^2 \lambda} \|b\|^2.$$

For B , similar to the proof of Lemma D.4, we have,

$$\begin{aligned}
B &= \mathbb{E}_{z \sim \mathcal{D}_k} \left[\left\| \sum_{i=1}^k \phi'(v_i^\top z) c_i \right\|^2 \right] \\
&= \int (2\pi)^{-k/2} \left\| \sum_{i=1}^k \phi'(v_i^\top z) c_i \right\|^2 e^{-\|z\|^2/2} dz \\
&= \int (2\pi)^{-k/2} \left\| \sum_{i=1}^k \phi'(\sigma_k \cdot u_i) c_i \right\|^2 e^{-\|V^\dagger u / \sigma_1(V^\dagger)\|^2/2} \cdot \det(V^\dagger / \sigma_1(V^\dagger)) du \\
&= \int (2\pi)^{-k/2} \left\| \sum_{i=1}^k \phi'(\sigma_k \cdot u_i) c_i \right\|^2 e^{-\|V^\dagger u / \sigma_1(V^\dagger)\|^2/2} \cdot \frac{1}{\lambda} du \\
&\geq \int (2\pi)^{-k/2} \left\| \sum_{i=1}^k \phi'(\sigma_k \cdot u_i) c_i \right\|^2 e^{-\|u\|^2/2} \cdot \frac{1}{\lambda} du \\
&= \frac{1}{\lambda} \mathbb{E}_{u \sim \mathcal{D}_k} \left[\left\| \sum_{i=1}^k \phi'(\sigma_k \cdot u_i) c_i \right\|^2 \right] \\
&= \frac{1}{\lambda} \left(\sum_{i=1}^k \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(\sigma_k \cdot u_i) \phi'(\sigma_k \cdot u_i) c_i^\top c_i] + \sum_{i \neq l} \mathbb{E}_{u \sim \mathcal{D}_k} [\phi'(\sigma_k \cdot u_i) \phi'(\sigma_k \cdot u_l) c_i^\top c_l] \right) \\
&= \frac{1}{\lambda} \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot u_i)^2] \sum_{i=1}^k \|c_i\|^2 + \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)] \right)^2 \sum_{i \neq l} c_i^\top c_l \right) \\
&= \frac{1}{\lambda} \left(\left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)] \right)^2 \left\| \sum_{i=1}^k c_i \right\|_2^2 + \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)^2] - \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)] \right)^2 \right) \|c\|^2 \right) \\
&\geq \frac{1}{\lambda} \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)^2] - \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_k \cdot z)] \right)^2 \right) \|c\|^2 \\
&\geq \rho(\sigma_k) \frac{1}{\lambda} \|c\|^2,
\end{aligned}$$

where the first step follows by definition of Gaussian distribution, the second step follows by replacing z by $z = V^\dagger u / \sigma_1(V^\dagger)$, and then $v_i^\top z = u_i / \sigma_1(V^\dagger) = u_i \sigma_k(W^*)$, the third step follows by $\|u\|^2 \geq \left\| \frac{1}{\sigma_1(V^\dagger)} V^\dagger u \right\|^2$, the fourth step follows by $\det(V^\dagger / \sigma_1(V^\dagger)) = \det(V^\dagger) / \sigma_1^k(V^\dagger) = 1/\lambda$, the fifth step follows by definition of Gaussian distribution, the ninth step follows by $x^2 \geq 0$ for any $x \in \mathbb{R}$, and the last step follows by Property 3.2.

Note that $1 = \|a\|^2 = \|b\|^2 + \|c\|^2$. Thus, we finish the proof for the lower bound. \square

D.2.3 Upper Bound on the Eigenvalues of Hessian for Non-orthogonal Case

Lemma D.7 (Upper bound). *If $\phi(z)$ satisfies Property 3.1, 3.2 and 3.3, we have the following property for the second derivative of function $f_{\mathcal{D}}(W)$ at W^* ,*

$$\nabla^2 f_{\mathcal{D}}(W^*) \preceq O(kv_{\max}^2 \sigma_1^{2p}) I.$$

Proof. Similarly, we can calculate the upper bound of the eigenvalues by

$$\begin{aligned}
& \|\nabla^2 f(W^*)\| \\
&= \max_{\|a\|=1} a^\top \nabla^2 f(W^*) a \\
&= v_{\max}^2 \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_k} \left[\left(\sum_{i=1}^k a_i^\top x \cdot \phi'(w_i^{*\top} x) \right)^2 \right] \\
&= v_{\max}^2 \max_{\|a\|=1} \sum_{i=1}^k \sum_{l=1}^k \mathbb{E}_{x \sim \mathcal{D}_k} [a_i^\top x \cdot \phi'(w_i^{*\top} x) \cdot a_l^\top x \cdot \phi'(w_l^{*\top} x)] \\
&\leq v_{\max}^2 \max_{\|a\|=1} \sum_{i=1}^k \sum_{l=1}^k \left(\mathbb{E}_{x \sim \mathcal{D}_k} [(a_i^\top x)^4] \cdot \mathbb{E}_{x \sim \mathcal{D}_k} [(\phi'(w_i^{*\top} x))^4] \cdot \mathbb{E}_{x \sim \mathcal{D}_k} [(a_l^\top x)^4] \cdot \mathbb{E}_{x \sim \mathcal{D}_k} [(\phi'(w_l^{*\top} x))^4] \right)^{1/4} \\
&\lesssim v_{\max}^2 \max_{\|a\|=1} \sum_{i=1}^k \sum_{l=1}^k \|a_i\| \cdot \|a_l\| \cdot \|w_i^*\|^p \cdot \|w_l^*\|^p \\
&\leq v_{\max}^2 \max_{\|a\|=1} \sum_{i=1}^k \sum_{l=1}^k \|a_i\| \cdot \|a_l\| \cdot \sigma_1^{2p} \\
&\leq k v_{\max}^2 \sigma_1^{2p},
\end{aligned}$$

where the first inequality follows Hölder's inequality, the second inequality follows by Property 3.1, the third inequality follows by $\|w_i^*\| \leq \sigma_1(W^*)$, and the last inequality by Cauchy-Schwarz inequality. \square

D.3 Error Bound of Hessians near the Ground Truth for Smooth Activations

The goal of this Section is to prove Lemma D.8

Lemma D.8 (Error bound of Hessians near the ground truth for smooth activations). *Let $\phi(z)$ satisfy Property 3.1, Property 3.2 and Property 3.3(a). Let W satisfy $\|W - W^*\| \leq \sigma_k/2$. Let S denote a set of i.i.d. samples from the distribution defined in (1). Then for any $t \geq 1$ and $0 < \epsilon < 1/2$, if*

$$|S| \geq \epsilon^{-2} d \kappa^2 \tau \cdot \text{poly}(\log d, t)$$

then we have, with probability at least $1 - 1/d^{\Omega(t)}$,

$$\|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| \lesssim v_{\max}^2 k \sigma_1^p (\epsilon \sigma_1^p + \|W - W^*\|).$$

Proof. This follows by combining Lemma D.10 and Lemma D.11 directly. \square

D.3.1 Second-order Smoothness near the Ground Truth for Smooth Activations

The goal of this Section is to prove Lemma D.10.

Fact D.9. *Let w_i denote the i -th column of $W \in \mathbb{R}^{d \times k}$, and w_i^* denote the i -th column of $W^* \in \mathbb{R}^{d \times k}$. If $\|W - W^*\| \leq \sigma_k(W^*)/2$, then for all $i \in [k]$,*

$$\frac{1}{2} \|w_i^*\| \leq \|w_i\| \leq \frac{3}{2} \|w_i^*\|.$$

Proof. Note that if $\|W - W^*\| \leq \sigma_k(W^*)/2$, we have $\sigma_k(W^*)/2 \leq \sigma_i(W) \leq \frac{3}{2}\sigma_1(W^*)$ for all $i \in [k]$ by Weyl's inequality. By definition of singular value, we have $\sigma_k(W^*) \leq \|w_i^*\| \leq \sigma_1(W^*)$. By definition of spectral norm, we have $\|w_i - w_i^*\| \leq \|W - W^*\|$. Thus, we can lower bound $\|w_i\|$,

$$\|w_i\| \leq \|w_i^*\| + \|w_i - w_i^*\| \leq \|w_i^*\| + \|W - W^*\| \leq \|w_i^*\| + \sigma_k/2 \leq \frac{3}{2}\|w_i^*\|.$$

Similarly, we have $\|w_i\| \geq \frac{1}{2}\|w_i^*\|$. □

Lemma D.10 (Second-order smoothness near the ground truth for smooth activations). *If $\phi(z)$ satisfies Property 3.1, Property 3.2 and Property 3.3(a), then for any W with $\|W - W^*\| \leq \sigma_k/2$, we have*

$$\|\nabla^2 f_{\mathcal{D}}(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| \lesssim k^2 v_{\max}^2 \sigma_1^p \|W - W^*\|.$$

Proof. Let $\Delta = \nabla^2 f_{\mathcal{D}}(W) - \nabla^2 f_{\mathcal{D}}(W^*)$. For each $(i, l) \in [k] \times [k]$, let $\Delta_{i,l} \in \mathbb{R}^{d \times d}$ denote the (i, l) -th block of Δ . Then, for $i \neq l$, we have

$$\Delta_{i,l} = \mathbb{E}_{x \sim \mathcal{D}_d} \left[v_i^* v_l^* \left(\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x) \right) x x^\top \right],$$

and for $i = l$, we have

$$\begin{aligned} & \Delta_{i,i} \\ &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* \phi(w_r^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top + v_i^{*2} \left(\phi'^2(w_i^\top x) - \phi'^2(w_i^{*\top} x) \right) x x^\top \right] \\ &= \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* \phi(w_r^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top \right] + \mathbb{E}_{x \sim \mathcal{D}_d} \left[v_i^{*2} \left(\phi'^2(w_i^\top x) - \phi'^2(w_i^{*\top} x) \right) x x^\top \right]. \quad (11) \end{aligned}$$

In the next a few paragraphs, we first show how to bound the off-diagonal term, and then show how to bound the diagonal term.

First, we consider off-diagonal terms.

$$\begin{aligned}
& \|\Delta_{i,l}\| \\
&= \left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[v_i^* v_l^* \left(\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x) \right) x x^\top \right] \right\| \\
&\leq v_{\max}^2 \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left| \phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x) \right| \cdot (x^\top a)^2 \right] \\
&\leq v_{\max}^2 \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| + |\phi'(w_i^{*\top} x)| |\phi'(w_l^\top x) - \phi'(w_l^{*\top} x)| \right) (x^\top a)^2 \right] \\
&= v_{\max}^2 \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top a)^2 \right] \right. \\
&\quad \left. + \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^{*\top} x)| |\phi'(w_l^\top x) - \phi'(w_l^{*\top} x)| (x^\top a)^2 \right] \right) \\
&\leq v_{\max}^2 \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[L_2 |(w_i - w_i^*)^\top x| |\phi'(w_l^\top x)| (x^\top a)^2 \right] + \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^{*\top} x)| L_2 |(w_l - w_l^*)^\top x| (x^\top a)^2 \right] \right) \\
&\leq v_{\max}^2 \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[L_2 |(w_i - w_i^*)^\top x| L_1 |w_l^\top x|^p (x^\top a)^2 \right] + \mathbb{E}_{x \sim \mathcal{D}_d} \left[L_1 |w_i^{*\top} x|^p L_2 |(w_l - w_l^*)^\top x| (x^\top a)^2 \right] \right) \\
&\leq v_{\max}^2 L_1 L_2 \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[|(w_i - w_i^*)^\top x| |w_l^\top x|^p (x^\top a)^2 \right] + \mathbb{E}_{x \sim \mathcal{D}_d} \left[|(w_l - w_l^*)^\top x| |w_i^{*\top} x|^p (x^\top a)^2 \right] \right) \\
&\lesssim v_{\max}^2 L_1 L_2 \max_{\|a\|=1} (\|w_i - w_i^*\| \|w_l\|^p \|a\|^2 + \|w_l - w_l^*\| \|w_i^*\|^p \|a\|^2) \\
&\lesssim v_{\max}^2 L_1 L_2 \sigma_1^p(W^*) \|W - W^*\| \tag{12}
\end{aligned}$$

where the first step follows by definition of $\Delta_{i,l}$, the second step follows by definition of spectral norm and $v_i^* v_l^* \leq v_{\max}^2$, the third step follows by triangle inequality, the fourth step follows by linearity of expectation, the fifth step follows by Property 3.3(a), i.e., $|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| \leq L_2 |(w_i - w_i^*)^\top x|$, the sixth step follows by Property 3.1, i.e., $\phi'(z) \leq L_1 |z|^p$, the seventh step follows by Fact B.6, and the last step follows by $\|a\|^2 = 1$, $\|w_i - w_i^*\| \leq \|W - W^*\|$, $\|w_i\| \leq \frac{3}{2} \|w_i^*\|$, and $\|w_i^*\| \leq \sigma_1(W^*)$.

Note that the proof for the off-diagonal terms also applies to bounding the second-term in the diagonal block Eq. (11). Thus we only need to show how to bound the first term in the diagonal

block Eq. (11).

$$\begin{aligned}
& \left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* \phi(w_r^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top \right] \right\| \\
&= \left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* (\phi(w_r^\top x) - \phi(w_r^{*\top} x)) \right) v_i^* \phi''(w_i^\top x) x x^\top \right] \right\| \\
&\leq v_{\max}^2 \sum_{r=1}^k \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [|\phi(w_r^\top x) - \phi(w_r^{*\top} x)| \phi''(w_i^\top x) |x^\top a|^2] \\
&\leq v_{\max}^2 \sum_{r=1}^k \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [|\phi(w_r^\top x) - \phi(w_r^{*\top} x)| L_2 (x^\top a)^2] \\
&\leq v_{\max}^2 L_2 \sum_{r=1}^k \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\max_{z \in [w_r^\top x, w_r^{*\top} x]} |\phi'(z)| \cdot |(w_r - w_r^*)^\top x| \cdot (x^\top a)^2 \right] \\
&\leq v_{\max}^2 L_2 \sum_{r=1}^k \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\max_{z \in [w_r^\top x, w_r^{*\top} x]} L_1 |z|^p \cdot |(w_r - w_r^*)^\top x| \cdot (x^\top a)^2 \right] \\
&\leq v_{\max}^2 L_1 L_2 \sum_{r=1}^k \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [(|w_r^\top x|^p + |w_r^{*\top} x|^p) \cdot |(w_r - w_r^*)^\top x| \cdot (x^\top a)^2] \\
&\lesssim v_{\max}^2 L_1 L_2 \sum_{r=1}^k [(\|w_r\|^p + \|w_r^*\|^p) \|w_r - w_r^*\|] \\
&\lesssim k v_{\max}^2 L_1 L_2 \sigma_1^p(W^*) \|W - W^*\|, \tag{13}
\end{aligned}$$

where the first step follows by $y = \sum_{r=1}^k v_r^* \phi(w_r^{*\top} x)$, the second step follows by definition of spectral norm and $v_r^* v_i^* \leq |v_{\max}|^2$, the third step follows by Property 3.3(a), i.e., $|\phi''(w_i^\top x)| \leq L_2$, the fourth step follows by $|\phi(w_r^\top x) - \phi(w_r^{*\top} x)| \leq \max_{z \in [w_r^\top x, w_r^{*\top} x]} |\phi'(z)| |(w_r - w_r^*)^\top x|$, the fifth step follows Property 3.1, i.e., $|\phi'(z)| \leq L_1 |z|^p$, the sixth step follows by $\max_{z \in [w_r^\top x, w_r^{*\top} x]} |z|^p \leq (|w_r^\top x|^p + |w_r^{*\top} x|^p)$, the seventh step follows by Fact B.6.

Putting it all together, we can bound the error by

$$\begin{aligned}
& \|\nabla^2 f_{\mathcal{D}}(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| \\
&= \max_{\|a\|=1} a^\top (\nabla^2 f_{\mathcal{D}}(W) - \nabla^2 f_{\mathcal{D}}(W^*))a \\
&= \max_{\|a\|=1} \sum_{i=1}^k \sum_{l=1}^k a_i^\top \Delta_{i,l} a_l \\
&= \max_{\|a\|=1} \left(\sum_{i=1}^k a_i^\top \Delta_{i,i} a_i + \sum_{i \neq l} a_i^\top \Delta_{i,l} a_l \right) \\
&\leq \max_{\|a\|=1} \left(\sum_{i=1}^k \|\Delta_{i,i}\| \|a_i\|^2 + \sum_{i \neq l} \|\Delta_{i,l}\| \|a_i\| \|a_l\| \right) \\
&\leq \max_{\|a\|=1} \left(\sum_{i=1}^k C_1 \|a_i\|^2 + \sum_{i \neq l} C_2 \|a_i\| \|a_l\| \right) \\
&= \max_{\|a\|=1} \left(C_1 \sum_{i=1}^k \|a_i\|^2 + C_2 \left(\left(\sum_{i=1}^k \|a_i\| \right)^2 - \sum_{i=1}^k \|a_i\|^2 \right) \right) \\
&\leq \max_{\|a\|=1} \left(C_1 \sum_{i=1}^k \|a_i\|^2 + C_2 \left(k \sum_{i=1}^k \|a_i\|^2 - \sum_{i=1}^k \|a_i\|^2 \right) \right) \\
&= \max_{\|a\|=1} (C_1 + C_2(k-1)) \\
&\lesssim k^2 v_{\max}^2 L_1 L_2 \sigma_1^p(W^*) \|W - W^*\|.
\end{aligned}$$

where the first step follows by definition of spectral norm and a denotes a vector $\in \mathbb{R}^{dk}$, the first inequality follows by $\|A\| = \max_{\|x\| \neq 0, \|y\| \neq 0} \frac{x^\top A y}{\|x\| \|y\|}$, the second inequality follows by $\|\Delta_{i,i}\| \leq C_1$ and $\|\Delta_{i,l}\| \leq C_2$, the third inequality follows by Cauchy-Scharwz inequality, the eighth step follows by $\sum_{i=1}^k \|a_i\|^2 = 1$, and the last step follows by Eq (12) and (13). \square

D.3.2 Empirical and Population Difference for Smooth Activations

The goal of this Section is to prove Lemma D.11. For each $i \in [k]$, let σ_i denote the i -th largest singular value of $W^* \in \mathbb{R}^{d \times k}$.

Note that Bernstein inequality requires the spectral norm of each random matrix to be bounded almost surely. However, since we assume Gaussian distribution for x , $\|x\|^2$ is not bounded almost surely. The main idea is to do truncation and then use Matrix Bernstein inequality. Details can be found in Lemma B.7 and Corollary B.8.

Lemma D.11 (Empirical and population difference for smooth activations). *Let $\phi(z)$ satisfy Property 3.1, 3.2 and 3.3(a). Let W satisfy $\|W - W^*\| \leq \sigma_k/2$. Let S denote a set of i.i.d. samples from distribution \mathcal{D} (defined in (1)). Then for any $t \geq 1$ and $0 < \epsilon < 1/2$, if*

$$|S| \geq \epsilon^{-2} d \kappa^2 \tau \cdot \text{poly}(\log d, t)$$

then we have, with probability at least $1 - d^{-\Omega(t)}$,

$$\|\nabla^2 \hat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W)\| \lesssim v_{\max}^2 k \sigma_1^p(\epsilon \sigma_1^p + \|W - W^*\|).$$

Proof. Define $\Delta = \nabla^2 f_{\mathcal{D}}(W) - \nabla^2 \widehat{f}_S(W)$. Let's first consider the diagonal blocks. Define

$$\begin{aligned} \Delta_{i,i} = & \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* \phi(w_r^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top + v_i^{*2} \phi'^2(w_i^\top x) x x^\top \right] \\ & - \left(\frac{1}{n} \sum_{j=1}^n \left(\sum_{r=1}^k v_r^* \phi(w_r^\top x_j) - y \right) v_i^* \phi''(w_i^\top x_j) x_j x_j^\top + v_i^{*2} \phi'^2(w_i^\top x_j) x_j x_j^\top \right). \end{aligned}$$

Let's further decompose $\Delta_{i,i}$ into $\Delta_{i,i} = \Delta_{i,i}^{(1)} + \Delta_{i,i}^{(2)}$, where

$$\begin{aligned} & \Delta_{i,i}^{(1)} \\ = & \mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(\sum_{r=1}^k v_r^* \phi(w_r^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top \right] - \frac{1}{n} \sum_{j=1}^n \left(\sum_{r=1}^k v_r^* \phi(w_r^\top x_j) - y \right) v_i^* \phi''(w_i^\top x_j) x_j x_j^\top \\ = & v_i^* \sum_{r=1}^k \left(v_r^* \mathbb{E}_{x \sim \mathcal{D}_d} \left[(\phi(w_r^\top x) - \phi(w_r^{*\top} x)) \phi''(w_i^\top x) x x^\top \right] \right. \\ & \left. - \frac{1}{n} \sum_{j=1}^n (\phi(w_r^\top x_j) - \phi(w_r^{*\top} x_j)) \phi''(w_i^\top x_j) x_j x_j^\top \right), \end{aligned}$$

and

$$\Delta_{i,i}^{(2)} = \mathbb{E}_{x \sim \mathcal{D}_d} [v_i^{*2} \phi'^2(w_i^\top x) x x^\top] - \frac{1}{n} \sum_{j=1}^n [v_i^{*2} \phi'^2(w_i^\top x_j) x_j x_j^\top]. \quad (14)$$

The off-diagonal block is

$$\Delta_{i,l} = v_i^* v_l^* \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top] - \frac{1}{n} \sum_{j=1}^n \phi'(w_i^\top x_j) \phi'(w_l^\top x_j) x_j x_j^\top \right)$$

Combining Claims D.12, D.13 and D.14, and taking a union bound over k^2 different $\Delta_{i,j}$, we obtain if $n \geq \epsilon^{-2} \kappa^2(W^*) \tau d \cdot \text{poly}(t, \log d)$ for any $\epsilon \in (0, 1/2)$, with probability at least $1 - 1/d^t$,

$$\|\nabla^2 \widehat{f}_S(W) - \nabla^2 f(W)\| \lesssim v_{\max}^2 k \sigma_1^p(W^*) (\epsilon \sigma_1^p(W^*) + \|W - W^*\|)$$

□

Claim D.12. For each $i \in [k]$, if $n \geq d \text{poly}(\log d, t)$

$$\|\Delta_{i,i}^{(1)}\| \lesssim k v_{\max}^2 \sigma_1^p(W^*) \|W - W^*\|$$

holds with probability $1 - 1/d^{4t}$.

Proof. For each $r \in [k]$, we define function $\widehat{B}_r : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$,

$$\widehat{B}_r(x) = L_1 L_2 \cdot (|w_r^\top x|^p + |w_r^{*\top} x|^p) \cdot |(w_r - w_r^*)^\top x| \cdot x x^\top.$$

According to Properties 3.1,3.2 and 3.3(a), we have for each $x \in S$,

$$-\widehat{B}_r(x) \preceq (\phi(w_r^\top x) - \phi(w_r^{*\top} x))\phi''(w_i^\top x)xx^\top \preceq \widehat{B}_r(x)$$

Therefore,

$$\Delta_{i,i}^{(1)} \preceq v_{\max}^2 \sum_{r=1}^k \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\widehat{B}_r(x)] + \frac{1}{|S|} \sum_{x \in S} \widehat{B}_r(x) \right).$$

Let $h_r(x) = L_1 L_2 |w_r^\top x|^p \cdot |(w_r - w_r^*)^\top x|$. Let $\overline{B}_r = \mathbb{E}_{x \sim \mathcal{D}_d} [h_r(x)xx^\top]$. Define function $B_r(x) = h_r(x)xx^\top$.

(I) Bounding $|h_r(x)|$.

According to Fact B.1, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^{4t})$,

$$|h_r(x)| = L_1 L_2 |w_r^\top x|^p |(w_r - w_r^*)^\top x| \lesssim L_1 L_2 \|w_r\|^p \|w_r - w_r^*\| (t \log n)^{(p+1)/2}.$$

(II) Bounding $\|\overline{B}_r\|$.

$$\|\overline{B}_r\| \geq \mathbb{E}_{x \sim \mathcal{D}_d} \left[L_1 L_2 |w_r^\top x|^p |(w_r - w_r^*)^\top x| \left(\frac{(w_r - w_r^*)^\top x}{\|w_r - w_r^*\|} \right)^2 \right] \gtrsim L_1 L_2 \|w_r\|^p \|w_r - w_r^*\|,$$

where the first step follows by definition of spectral norm, and last step follows by Fact B.6. Using Fact B.6, we can also prove an upper bound $\|\overline{B}_r\|, \|\overline{B}_r\| \lesssim L_1 L_2 \|w_r\|^p \|w_r - w_r^*\|$.

(III) Bounding $(\mathbb{E}_{x \sim \mathcal{D}_d} [h^4(x)])^{1/4}$

Using Fact B.6, we have

$$\left(\mathbb{E}_{x \sim \mathcal{D}_d} [h^4(x)] \right)^{1/4} = L_1 L_2 \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[\left(|w_r^\top x|^p |(w_r - w_r^*)^\top x| \right)^4 \right] \right)^{1/4} \lesssim L_1 L_2 \|w_r\|^p \|w_r - w_r^*\|.$$

By applying Corollary B.8, if $n \geq \epsilon^{-2} d \text{poly}(\log d, t)$, then with probability $1 - 1/d^{4t}$,

$$\begin{aligned} & \left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[|w_r^\top x|^p \cdot |(w_r - w_r^*)^\top x| \cdot xx^\top \right] - \frac{1}{|S|} \left(\sum_{x \in S} |w_r^\top x|^p \cdot |(w_r - w_r^*)^\top x| \cdot xx^\top \right) \right\| \\ &= \left\| \overline{B}_r - \frac{1}{|S|} \sum_{x \in S} B_r(x) \right\| \\ &\leq \epsilon \|\overline{B}_r\| \\ &\lesssim \epsilon \|w_r\|^p \|w_r - w_r^*\|. \end{aligned} \tag{15}$$

If $\epsilon \leq 1/2$, we have

$$\|\Delta_{i,i}^{(1)}\| \lesssim \sum_{i=1}^k v_{\max}^2 \|\overline{B}_r\| \lesssim k v_{\max}^2 \sigma_1^p(W^*) \|W - W^*\|$$

□

Claim D.13. For each $i \in [k]$, if $n \geq \epsilon^{-2} d \tau \text{poly}(\log d, t)$, then

$$\|\Delta_{i,i}^{(2)}\| \lesssim \epsilon v_{\max}^2 \sigma_1^{2p}$$

holds with probability $1 - 1/d^{4t}$.

Proof. Recall the definition of $\Delta_{i,i}^{(2)}$.

$$\Delta_{i,i}^{(2)} = \mathbb{E}_{x \sim \mathcal{D}_d} [v_i^{*2} \phi'^2(w_i^\top x) x x^\top] - \frac{1}{|S|} \sum_{x \in S} [v_i^{*2} \phi'^2(w_i^\top x) x x^\top]$$

Let $h_i(x) = \phi'^2(w_i^\top x)$. Let $\bar{B}_i = \mathbb{E}_{x \sim \mathcal{D}_d} [h_i(x) x x^\top]$. Define function $B_i(x) = h_i(x) x x^\top$.

(I) Bounding $|h_i(x)|$.

For any constant $t \geq 1$, $(\phi'(w_i^\top x))^2 \leq L_1^2 |w_i^\top x|^{2p} \lesssim L_1^2 \|w_i\|^{2p} t^p \log^p(n)$ with probability $1 - 1/(nd^{4t})$ according to Fact B.1.

(II) Bounding $\|\bar{B}_i\|$.

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'^2(w_i^\top x) x x^\top] \right\| &= \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'^2(w_i^\top x) (x^\top a)^2] \\ &= \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\phi'^2(w_i^\top x) \left(\alpha \bar{w}_i^\top x + \beta x^\top v \right)^2 \right] \\ &= \max_{\alpha^2 + \beta^2 = 1, \|v\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[\phi'^2(w_i^\top x) \left((\alpha \bar{w}_i^\top x)^2 + (\beta x^\top v)^2 \right) \right] \\ &= \max_{\alpha^2 + \beta^2 = 1} \left(\alpha^2 \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'^2(\|w_i\|z) z^2] + \beta^2 \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'^2(\|w_i\|z)] \right) \\ &= \max \left(\mathbb{E}_{x \sim \mathcal{D}_1} [\phi'^2(\|w_i\|z) z^2], \mathbb{E}_{x \sim \mathcal{D}_1} [\phi'^2(\|w_i\|z)] \right) \end{aligned}$$

where $\bar{w}_i = w_i / \|w_i\|$ and v is a unit vector orthogonal to w_i such that $a = \alpha \bar{w}_i + \beta v$. Now from Property 3.2, we have

$$\rho(\|w_i\|) \leq \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'^2(w_i^\top x) x x^\top] \right\| \lesssim L_1^2 \|w_i\|^{2p}.$$

(III) Bounding $(\mathbb{E}_{x \sim \mathcal{D}_d} [h_i^4(x)])^{1/4}$.

$$\left(\mathbb{E}_{x \sim \mathcal{D}_d} [h_i^4(x)] \right)^{1/4} = \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi'^8(w_i^\top x)] \right)^{1/4} \lesssim L_1^2 \|w_i\|^{2p}.$$

By applying Corollary B.8, we have, for any $0 < \epsilon < 1$, if $n \geq \epsilon^{-2} d \frac{\|w_i\|^{4p}}{\rho^2(\|w_i\|)} \text{poly}(\log d, t)$ the following bound holds

$$\left\| \bar{B}_i - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon \|\bar{B}_i\|,$$

with probability at least $1 - 1/d^{4t}$.

Therefore, if $n \geq \epsilon^{-2} d \tau \text{poly}(\log d, t)$, where $\tau = \frac{(3\sigma_1/2)^{4p}}{\min_{\sigma \in [\sigma_k/2, 3\sigma_1/2]} \rho^2(\sigma)}$, we have with probability $1 - 1/d^{4t}$

$$\|\Delta_{i,i}^{(2)}\| \lesssim \epsilon v_{\max}^2 \sigma_1^{2p}$$

□

Claim D.14. For each $i \in [k], j \in [k], i \neq j$, if $n \geq \epsilon^{-2} \kappa^2 \tau d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,j}\| \leq \epsilon v_{\max}^2 \sigma_1^{2p}(W^*)$$

holds with probability $1 - 1/d^{4t}$.

Proof. Recall the definition of off-diagonal blocks $\Delta_{i,l}$,

$$\Delta_{i,l} = v_i^* v_l^* \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top] - \frac{1}{|S|} \sum_{x \in S} \phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top \right) \quad (16)$$

Let $h_{i,l}(x) = \phi'(w_i^\top x) \phi'(w_l^\top x)$. Define function $B_{i,l}(x) = h_{i,l}(x) x x^\top$. Let $\bar{B}_{i,l} = \mathbb{E}_{x \sim \mathcal{D}_d} [h_{i,l}(x) x x^\top]$.

(I) Bounding $|h_{i,l}(x)|$.

For any constant $t \geq 1$, we have with probability $1 - 1/(nd^{4t})$

$$\begin{aligned} |h_{i,l}(x)| &= |\phi'(w_i^\top x) \phi'(w_l^\top x)| \\ &\leq L_1^2 \|w_i^\top x\|^p \|w_l^\top x\|^p \\ &\leq L_1^2 \|w_i\|^p \|w_l\|^p (t \log n)^p \\ &\lesssim L_1^2 \sigma_1^{2p} (t \log n)^p \end{aligned}$$

where the third step follows by Fact B.1.

(II) Bounding $\|\bar{B}_{i,l}\|$.

Let $U \in \mathbb{R}^{d \times 2}$ be the orthogonal basis of $\text{span}\{w_i, w_l\}$ and $U_\perp \in \mathbb{R}^{d \times (d-2)}$ be the complementary matrix of U . Let matrix $V \in \mathbb{R}^{2 \times 2}$ denote $U^\top [w_i \ w_l]$, then $UV = [w_i \ w_l] \in \mathbb{R}^{d \times 2}$. Given any vector $a \in \mathbb{R}^d$, there exist vectors $b \in \mathbb{R}^2$ and $c \in \mathbb{R}^{d-2}$ such that $a = Ub + U_\perp c$. We can simplify $\|\bar{B}_{i,l}\|$ in the following way,

$$\begin{aligned} \|\bar{B}_{i,l}\| &= \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top] \right\| \\ &= \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) (x^\top a)^2] \\ &= \max_{\|b\|^2 + \|c\|^2 = 1} \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) (b^\top U^\top x + c^\top U_\perp^\top x)^2] \\ &= \max_{\|b\|^2 + \|c\|^2 = 1} \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) ((b^\top U^\top x)^2 + (c^\top U_\perp^\top x)^2)] \\ &= \max_{\|b\|^2 + \|c\|^2 = 1} \left(\underbrace{\mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z) \phi'(v_2^\top z) (b^\top z)^2]}_{A_1} + \underbrace{\mathbb{E}_{z \sim \mathcal{D}_2, s \sim \mathcal{D}_{d-2}} [\phi'(v_1^\top z) \phi'(v_2^\top z) (c^\top s)^2]}_{A_2} \right) \end{aligned}$$

We can lower bound the term A_1 ,

$$\begin{aligned}
A_1 &= \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z) \phi'(v_2^\top z) (b^\top z)^2] \\
&= \int (2\pi)^{-1} \phi'(v_1^\top z) \phi'(v_2^\top z) (b^\top z)^2 e^{-\|z\|^2/2} dz \\
&= \int (2\pi)^{-1} \phi'(s_1) \phi'(s_2) (b^\top V^{\dagger\top} s)^2 e^{-\|V^{\dagger\top} s\|^2/2} \cdot |\det(V^\dagger)| ds \\
&\geq \int (2\pi)^{-1} (\phi'(s_1) \phi'(s_2)) (b^\top V^{\dagger\top} s)^2 e^{-\sigma_1^2(V^\dagger) \|s\|^2/2} \cdot |\det(V^\dagger)| ds \\
&= \int (2\pi)^{-1} (\phi'(u_1/\sigma_1(V^\dagger)) \phi'(u_2/\sigma_1(V^\dagger))) \cdot (b^\top V^{\dagger\top} u/\sigma_1(V^\dagger))^2 e^{-\|u\|^2/2} |\det(V^\dagger)|/\sigma_1^2(V^\dagger) du \\
&= \frac{\sigma_2(V)}{\sigma_1(V)} \mathbb{E}_{u \sim \mathcal{D}_2} [(p^\top u)^2 \phi'(\sigma_2(V) \cdot u_1) \phi'(\sigma_2(V) \cdot u_2)] \\
&= \frac{\sigma_2(V)}{\sigma_1(V)} \mathbb{E}_{u \sim \mathcal{D}_2} [(p_1 u_1)^2 + (p_2 u_2)^2 + 2p_1 p_2 u_1 u_2] \phi'(\sigma_2(V) \cdot u_1) \phi'(\sigma_2(V) \cdot u_2) \\
&= \frac{\sigma_2(V)}{\sigma_1(V)} \left(\|p\|^2 \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z)] \right. \\
&\quad \left. + ((p^\top \mathbf{1})^2 - \|p\|^2) \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z]^2 \right)
\end{aligned}$$

where $p = V^\dagger b \cdot \sigma_2(V) \in \mathbb{R}^2$.

Since we are maximizing over $b \in \mathbb{R}^2$, we can choose b such that $\|p\| = \|b\|$. Then

$$\begin{aligned}
A_1 &= \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z) \phi'(v_2^\top z) (b^\top z)^2] \\
&\geq \frac{\sigma_2(V)}{\sigma_1(V)} \|b\|^2 \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z)] - \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z]^2 \right)
\end{aligned}$$

For the term A_2 , similarly we have

$$\begin{aligned}
A_2 &= \mathbb{E}_{z \sim \mathcal{D}_2, s \sim \mathcal{D}_{d-2}} [\phi'(v_1^\top z) \phi'(v_2^\top z) (c^\top s)^2] \\
&= \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z) \phi'(v_2^\top z)] \mathbb{E}_{s \sim \mathcal{D}_{d-2}} [(c^\top s)^2] \\
&= \|c\|^2 \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z) \phi'(v_2^\top z)] \\
&\geq \|c\|^2 \frac{\sigma_2(V)}{\sigma_1(V)} \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z)] \right)^2
\end{aligned}$$

For simplicity, we just set $\|b\| = 1$ and $\|c\| = 0$ to lower bound,

$$\begin{aligned}
&\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top] \right\| \\
&\geq \frac{\sigma_2(V)}{\sigma_1(V)} \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z^2] \cdot \mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z)] - \left(\mathbb{E}_{z \sim \mathcal{D}_1} [\phi'(\sigma_2(V) \cdot z) z] \right)^2 \right) \\
&\geq \frac{\sigma_2(V)}{\sigma_1(V)} \rho(\sigma_2(V)) \\
&\geq \frac{1}{\kappa(W^*)} \rho(\sigma_2(V))
\end{aligned}$$

where the second step is from Property 3.2 and the fact that $\sigma_k/2 \leq \sigma_2(V) \leq \sigma_1(V) \leq 3\sigma_1/2$.

For the upper bound of $\|\mathbb{E}_{x \sim \mathcal{D}_d}[\phi'(w_i^\top x)\phi'(w_l^\top x)xx^\top]\|$, we have

$$\begin{aligned} \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z)\phi'(v_2^\top z)(b^\top z)^2] &\leq L_1^2 \mathbb{E}_{z \sim \mathcal{D}_2} [|v_1^\top z|^p \cdot |v_2^\top z|^p \cdot |b^\top z|^2] \\ &\lesssim L_1^2 \|v_1\|^p \|v_2\|^p \|b\|^2 \\ &\lesssim L_1^2 \sigma_1^{2p} \end{aligned}$$

where the first step follows by Property 3.1, the second step follows by Fact B.6, and the last step follows by $\|v_1\| \leq \sigma_1$, $\|v_2\| \leq \sigma_1$ and $\|b\| \leq 1$. Similarly, we can upper bound,

$$\mathbb{E}_{z \sim \mathcal{D}_2, s \sim \mathcal{D}_{d-2}} [\phi'(v_1^\top z)\phi'(v_2^\top z)(c^\top s)^2] = \|c\|^2 \mathbb{E}_{z \sim \mathcal{D}_2} [\phi'(v_1^\top z)\phi'(v_2^\top z)] \lesssim L_1^2 \sigma_1^{2p}$$

Thus, we have

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [\phi'(w_i^\top x)\phi'(w_l^\top x)xx^\top] \right\| \lesssim L_1^2 \sigma_1^{2p} \lesssim \sigma_1^{2p}.$$

(III) Bounding $(\mathbb{E}_{x \sim \mathcal{D}_d}[h_{i,l}^4(x)])^{1/4}$.

$$\left(\mathbb{E}_{x \sim \mathcal{D}_d} [h_{i,l}^4(x)] \right)^{1/4} = \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi'^4(w_i^\top x) \cdot \phi'^4(w_l^\top x)] \right)^{1/4} \lesssim L_1^2 \|w_i\|^p \|w_l\|^p \lesssim L_1^2 \sigma_1^{2p}.$$

Therefore, applying Corollary B.8, we have, if $n \geq \epsilon^{-2} \kappa^2(W^*) \tau d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,j}\| \leq \epsilon v_{\max}^2 \sigma_1^{2p}(W^*).$$

holds with probability at least $1 - 1/d^{4t}$. \square

D.4 Error Bound of Hessians near the Ground Truth for Non-smooth Activations

The goal of this Section is to prove Lemma D.15,

Lemma D.15 (Error bound of Hessians near the ground truth for non-smooth activations). *Let $\phi(z)$ satisfy Property 3.1, 3.2 and 3.3(b). Let W satisfy $\|W - W^*\| \leq \sigma_k/2$. Let S denote a set of i.i.d. samples from the distribution defined in (1). Then for any $t \geq 1$ and $0 < \epsilon < 1/2$, if*

$$|S| \geq \epsilon^{-2} \kappa^2 \tau d \text{poly}(\log d, t)$$

with probability at least $1 - d^{-\Omega(t)}$,

$$\|\nabla^2 \widehat{f}_S(W) - \nabla^2 f_{\mathcal{D}}(W^*)\| \lesssim v_{\max}^2 k \sigma_1^{2p} (\epsilon + (\sigma_k^{-1} \cdot \|W - W^*\|)^{1/2}).$$

Proof. As we noted previously, when Property 3.3(b) holds, the diagonal blocks of the empirical Hessian can be written as, with probability 1,

$$\frac{\partial^2 \widehat{f}_S(W)}{\partial w_i^2} = \frac{1}{|S|} \sum_{(x,y) \in S} (v_i^* \phi'(w_i^\top x))^2 xx^\top$$

We construct a matrix $H \in \mathbb{R}^{dk \times dk}$ with i, l -th block as

$$H_{i,l} = v_i^* v_l^* \mathbb{E}_{x \sim \mathcal{D}_d} \left[\phi'(w_i^\top x) \phi'(w_l^\top x) x x^\top \right] \in \mathbb{R}^{d \times d}, \quad \forall i \in [k], l \in [k].$$

Note that $H \neq \nabla^2 f_{\mathcal{D}}(W)$. However we can still bound $\|H - \nabla^2 \widehat{f}_S(W)\|$ and $\|H - \nabla^2 f_{\mathcal{D}}(W^*)\|$ when we have enough samples and $\|W - W^*\|$ is small enough. The proof for $\|H - \nabla^2 \widehat{f}_S(W)\|$ basically follows the proof of Lemma D.11 as $\Delta_{ii}^{(2)}$ in Eq. (14) and Δ_{il} in Eq. (16) forms the blocks of $H - \nabla^2 \widehat{f}_S(W)$ and we can bound them without smoothness of $\phi(\cdot)$.

Now we focus on $H - \nabla^2 f_{\mathcal{D}}(W^*)$. We again consider each block.

$$\Delta_{i,l} = \mathbb{E}_{x \sim \mathcal{D}_d} \left[v_i^* v_l^* (\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x)) x x^\top \right].$$

We used the boundedness of $\phi''(z)$ to prove Lemma D.10. Here we can't use this condition. Without smoothness, we will stop at the following position.

$$\begin{aligned} & \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [v_i^* v_l^* (\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x)) x x^\top] \right\| \\ & \leq |v_i^* v_l^*| \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{*\top} x) \phi'(w_l^{*\top} x)| (x^\top a)^2 \right] \\ & \leq |v_i^* v_l^*| \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| \right. \\ & \quad \left. + |\phi'(w_i^{*\top} x)| |\phi'(w_l^\top x) - \phi'(w_l^{*\top} x)| (x^\top a)^2 \right] \\ & = |v_i^* v_l^*| \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top a)^2 \right] \right. \\ & \quad \left. + \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^{*\top} x)| |\phi'(w_l^\top x) - \phi'(w_l^{*\top} x)| (x^\top a)^2 \right] \right). \end{aligned} \tag{17}$$

where the first step follows by definition of spectral norm, the second step follows by triangle inequality, and the last step follows by linearity of expectation.

Without loss of generality, we just bound the first term in the above formulation. Let U be the orthogonal basis of $\text{span}(w_i, w_i^*, w_l)$. If w_i, w_i^*, w_l are independent, U is d -by-3. Otherwise it can be d -by- $\text{rank}(\text{span}(w_i, w_i^*, w_l))$. Without loss of generality, we assume $U = \text{span}(w_i, w_i^*, w_l)$ is d -by-3. Let $[v_i \ v_i^* \ v_l] = U^\top [w_i \ w_i^* \ w_l] \in \mathbb{R}^{3 \times 3}$, and $[u_i \ u_i^* \ u_l] = U_\perp^\top [w_i \ w_i^* \ w_l] \in \mathbb{R}^{(d-3) \times 3}$. Let $a = Ub + U_\perp c$, where $U_\perp \in \mathbb{R}^{d \times (d-3)}$ is the complementary matrix of U .

$$\begin{aligned} & \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top a)^2 \right] \\ & = \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top (Ub + U_\perp c))^2 \right] \\ & \lesssim \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| \left((x^\top Ub)^2 + (x^\top U_\perp c)^2 \right) \right] \\ & = \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top Ub)^2 \right] \\ & \quad + \mathbb{E}_{x \sim \mathcal{D}_d} \left[|\phi'(w_i^\top x) - \phi'(w_i^{*\top} x)| |\phi'(w_l^\top x)| (x^\top U_\perp c)^2 \right] \\ & = \mathbb{E}_{z \sim \mathcal{D}_3} \left[|\phi'(v_i^\top z) - \phi'(v_i^{*\top} z)| |\phi'(v_l^\top z)| (z^\top b)^2 \right] \\ & \quad + \mathbb{E}_{y \sim \mathcal{D}_{d-3}} \left[|\phi'(u_i^\top y) - \phi'(u_i^{*\top} y)| |\phi'(u_l^\top y)| (y^\top c)^2 \right] \end{aligned} \tag{18}$$

where the first step follows by $a = Ub + U_\perp c$, the last step follows by $(a + b)^2 \leq 2a^2 + 2b^2$.

By Property 3.3(b), we have e exceptional points which have $\phi''(z) \neq 0$. Let these e points be p_1, p_2, \dots, p_e . Note that if $v_i^\top z$ and $v_i^{*\top} z$ are not separated by any of these exceptional points, i.e., there exists no $j \in [e]$ such that $v_i^\top z \leq p_j \leq v_i^{*\top} z$ or $v_i^{*\top} z \leq p_j \leq v_i^\top z$, then we have $\phi'(v_i^\top z) = \phi'(v_i^{*\top} z)$ since $\phi''(s)$ are zeros except for $\{p_j\}_{j=1,2,\dots,e}$. So we consider the probability that $v_i^\top z, v_i^{*\top} z$ are separated by any exception point. We use ξ_j to denote the event that $v_i^\top z, v_i^{*\top} z$ are separated by an exceptional point p_j . By union bound, $1 - \sum_{j=1}^e \Pr[\xi_j]$ is the probability that $v_i^\top z, v_i^{*\top} z$ are not separated by any exceptional point. The first term of Equation (18) can be bounded as,

$$\begin{aligned}
& \mathbb{E}_{z \sim \mathcal{D}_3} \left[|\phi'(v_i^\top z) - \phi'(v_i^{*\top} z)| |\phi'(v_l^\top z)| (z^\top b)^2 \right] \\
&= \mathbb{E}_{z \sim \mathcal{D}_3} \left[\mathbf{1}_{\cup_{j=1}^e \xi_j} |\phi'(v_i^\top z) + \phi'(v_i^{*\top} z)| |\phi'(v_l^\top z)| (z^\top b)^2 \right] \\
&\leq \left(\mathbb{E}_{z \sim \mathcal{D}_3} \left[\mathbf{1}_{\cup_{j=1}^e \xi_j} \right] \right)^{1/2} \left(\mathbb{E}_{z \sim \mathcal{D}_3} \left[(\phi'(v_i^\top z) + \phi'(v_i^{*\top} z))^2 \phi'(v_l^\top z)^2 (z^\top b)^4 \right] \right)^{1/2} \\
&\leq \left(\sum_{j=1}^e \Pr_{z \sim \mathcal{D}_3} [\xi_j] \right)^{1/2} \left(\mathbb{E}_{z \sim \mathcal{D}_3} \left[(\phi'(v_i^\top z) + \phi'(v_i^{*\top} z))^2 \phi'(v_l^\top z)^2 (z^\top b)^4 \right] \right)^{1/2} \\
&\lesssim \left(\sum_{j=1}^e \Pr_{z \sim \mathcal{D}_3} [\xi_j] \right)^{1/2} (\|v_i\|^p + \|v_i^*\|^p) \|v_l\|^p \|b\|^2
\end{aligned}$$

where the first step follows by if $v_i^\top z, v_i^{*\top} z$ are not separated by any exceptional point then $\phi'(v_i^\top z) = \phi'(v_i^{*\top} z)$ and the last step follows by Hölder's inequality and Property 3.1.

It remains to upper bound $\Pr_{z \sim \mathcal{D}_3}[\xi_j]$. First note that if $v_i^\top z, v_i^{*\top} z$ are separated by an exceptional point, p_j , then $|v_i^{*\top} z - p_j| \leq |v_i^\top z - v_i^{*\top} z| \leq \|v_i - v_i^*\| \|z\|$. Therefore,

$$\Pr_{z \sim \mathcal{D}_3} [\xi_j] \leq \Pr_{z \sim \mathcal{D}_3} \left[\frac{|v_i^\top z - p_j|}{\|z\|} \leq \|v_i - v_i^*\| \right].$$

Note that $(\frac{v_i^{*\top} z}{\|z\| \|v_i^*\|} + 1)/2$ follows Beta(1,1) distribution which is uniform distribution on $[0, 1]$.

$$\Pr_{z \sim \mathcal{D}_3} \left[\frac{|v_i^\top z - p_j|}{\|z\| \|v_i^*\|} \leq \frac{\|v_i - v_i^*\|}{\|v_i^*\|} \right] \leq \Pr_{z \sim \mathcal{D}_3} \left[\frac{|v_i^\top z|}{\|z\| \|v_i^*\|} \leq \frac{\|v_i - v_i^*\|}{\|v_i^*\|} \right] \lesssim \frac{\|v_i - v_i^*\|}{\|v_i^*\|} \lesssim \frac{\|W - W^*\|}{\sigma_k(W^*)},$$

where the first step is because we can view $\frac{v_i^{*\top} z}{\|z\|}$ and $\frac{p_j}{\|z\|}$ as two independent random variables: the former is about the direction of z and the later is related to the magnitude of z . Thus, we have

$$\mathbb{E}_{z \in \mathcal{D}_3} [|\phi'(v_i^\top z) - \phi'(v_i^{*\top} z)| |\phi'(v_l^\top z)| (z^\top b)^2] \lesssim (e \|W - W^*\| / \sigma_k(W^*))^{1/2} \sigma_1^{2p}(W^*) \|b\|^2. \quad (19)$$

Similarly we have

$$\mathbb{E}_{y \in \mathcal{D}_{d-3}} [|\phi'(u_i^\top y) - \phi'(u_i^{*\top} y)| |\phi'(u_l^\top y)| (y^\top c)^2] \lesssim (e \|W - W^*\| / \sigma_k(W^*))^{1/2} \sigma_1^{2p}(W^*) \|c\|^2. \quad (20)$$

Finally combining Eq. (17), Eq. (19) and Eq. (20), we have

$$\|H - \nabla^2 f_{\mathcal{D}}(W^*)\| \lesssim kv_{\max}^2 (e \|W - W^*\| / \sigma_k(W^*))^{1/2} \sigma_1^{2p}(W^*),$$

which completes the proof. \square

D.5 Positive Definiteness for a Small Region

Here we introduce a lemma which shows that the Hessian of any W , which may be dependent on the samples but is very close to an anchor point, is close to the Hessian of this anchor point.

Lemma D.16. *Let S denote a set of samples from Distribution \mathcal{D} defined in Eq. (1). Let $W^a \in \mathbb{R}^{d \times k}$ be a point (respect to function $\widehat{f}_S(W)$), which is independent of the samples S , satisfying $\|W^a - W^*\| \leq \sigma_k/2$. Assume ϕ satisfies Property 3.1, 3.2 and 3.3(a). Then for any $t \geq 1$, if*

$$|S| \geq d \text{ poly}(\log d, t),$$

with probability at least $1 - d^{-t}$, for any W (which is not necessarily to be independent of samples) satisfying $\|W^a - W\| \leq \sigma_k/4$, we have

$$\|\nabla^2 \widehat{f}_S(W) - \nabla^2 \widehat{f}_S(W^a)\| \leq kv_{\max}^2 \sigma_1^p (\|W^a - W^*\| + \|W - W^a\| d^{(p+1)/2}).$$

Proof. Let $\Delta = \nabla^2 \widehat{f}_S(W) - \nabla^2 \widehat{f}_S(W^a) \in \mathbb{R}^{dk \times dk}$, then Δ can be thought of as k^2 blocks, and each block has size $d \times d$. The off-diagonal blocks are,

$$\Delta_{i,l} = v_i^* v_l^* \frac{1}{|S|} \sum_{x \in S} \left(\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^{a\top} x) \phi'(w_l^{a\top} x) \right) x x^\top$$

For diagonal blocks,

$$\begin{aligned} \Delta_{i,i} = & \frac{1}{|S|} \sum_{x \in S} \left(\left(\sum_{q=1}^k v_q^* \phi(w_q^\top x) - y \right) v_i^* \phi''(w_i^\top x) x x^\top + v_i^{*2} \phi'^2(w_i^\top x) x x^\top \right) \\ & - \frac{1}{|S|} \sum_{x \in S} \left(\left(\sum_{q=1}^k v_q^* \phi(w_q^{a\top} x) - y \right) v_i^* \phi''(w_i^{a\top} x) x x^\top + v_i^{*2} \phi'^2(w_i^{a\top} x) x x^\top \right) \end{aligned}$$

We further decompose $\Delta_{i,i}$ into $\Delta_{i,i} = \Delta_{i,i}^{(1)} + \Delta_{i,i}^{(2)}$, where

$$\Delta_{i,i}^{(1)} = v_i^* \frac{1}{|S|} \sum_{x \in S} \left(\left(\sum_{q=1}^k v_q^* \phi(w_q^\top x) - y \right) \phi''(w_i^\top x) - \left(\sum_{q=1}^k v_q^* \phi(w_q^{a\top} x) - y \right) \phi''(w_i^{a\top} x) \right) x x^\top,$$

and

$$\Delta_{i,i}^{(2)} = v_i^{*2} \frac{1}{|S|} \sum_{(x,y) \in S} \phi'^2(w_i^\top x) x x^\top - \phi'^2(w_i^{a\top} x) x x^\top. \quad (21)$$

We can further decompose $\Delta_{i,i}^{(1)}$ into $\Delta_{i,i}^{(1,1)}$ and $\Delta_{i,i}^{(1,2)}$,

$$\begin{aligned}
\Delta_{i,i}^{(1)} &= v_i^* \frac{1}{|S|} \sum_{x \in S} \left(\sum_{q=1}^k v_q^* \phi(w_q^\top x) - y \right) \phi''(w_i^\top x) - \left(\sum_{q=1}^k v_q^* \phi(w_q^{a\top} x) - y \right) \phi''(w_i^{a\top} x) x x^\top \\
&= v_i^* \frac{1}{|S|} \sum_{x \in S} \left(\sum_{q=1}^k v_q^* \phi(w_q^\top x) - \sum_{q=1}^k v_q^* \phi(w_q^{a\top} x) \right) \phi''(w_i^\top x) x x^\top \\
&\quad + v_i^* \frac{1}{|S|} \sum_{x \in S} \left(\sum_{q=1}^k v_q^* \phi(w_q^{a\top} x) - y \right) (\phi''(w_i^\top x) - \phi''(w_i^{a\top} x)) x x^\top \\
&= v_i^* \frac{1}{|S|} \sum_{x \in S} \sum_{q=1}^k v_q^* (\phi(w_q^\top x) - \phi(w_q^{a\top} x)) \phi''(w_i^\top x) x x^\top \\
&\quad + v_i^* \frac{1}{|S|} \sum_{x \in S} \sum_{q=1}^k v_q^* (\phi(w_q^{a\top} x) - \phi(w_q^{*\top} x)) (\phi''(w_i^\top x) - \phi''(w_i^{a\top} x)) x x^\top \\
&= \Delta_{i,i}^{(1,1)} + \Delta_{i,i}^{(1,2)}.
\end{aligned}$$

Combining Claim D.17 and Claim D.18, we have if

$$n \geq d \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^{4t}$,

$$\|\Delta_{i,i}^{(1)}\| \lesssim k v_{\max}^2 \sigma_1^p (\|W^a - W^*\| + \|W^a - W\| d^{(p+1)/2}). \quad (22)$$

Therefore, combining Eq. (22), Claim D.19 and Claim D.20, we complete the proof. \square

Claim D.17. For each $i \in [k]$, if $n \geq d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,i}^{(1,1)}\| \lesssim k v_{\max}^2 \sigma_1^p \|W^a - W\| d^{(p+1)/2}$$

Proof. Define function $h_1(x) = \|x\|^{p+1}$ and $h_2(x) = |w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x|$. Note that h_1 and h_2 don't contain W which maybe depend on the samples. Therefore, we can use the modified matrix Bernstein inequality (Corollary B.8) to bound $\Delta_{i,i}^{(1)}$.

(I) Bounding $|h_1(x)|$.

By Fact B.2, we have $h_1(x) \lesssim (td \log n)^{(p+1)/2}$ with probability at least $1 - 1/(nd^{4t})$.

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[\|x\|^{p+1} x x^\top]\|$.

Let $g(x) = (2\pi)^{-d/2}e^{-\|x\|^2/2}$. Note that $xg(x)dx = -dg(x)$.

$$\begin{aligned}
\mathbb{E}_{x \sim \mathcal{D}_d} \left[\|x\|^{p+1} xx^\top \right] &= \int \|x\|^{p+1} g(x) xx^\top dx \\
&= - \int \|x\|^{p+1} d(g(x)) x^\top \\
&= - \int \|x\|^{p+1} d(g(x) x^\top) + \int \|x\|^{p+1} g(x) I_d dx \\
&= \int d(\|x\|^{p+1}) g(x) x^\top + \int \|x\|^{p+1} g(x) I_d dx \\
&= \int (p+1) \|x\|^{p-1} g(x) xx^\top dx + \int \|x\|^{p+1} g(x) I_d dx \\
&\preceq \int \|x\|^{p+1} g(x) I_d dx \\
&= \mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^{p+1}] I_d.
\end{aligned}$$

Since $\|x\|^2$ follows χ^2 distribution with degree d , $\mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^q] = 2^{q/2} \frac{\Gamma((d+q)/2)}{\Gamma(d/2)}$ for any $q \geq 0$. So, $d^{q/2} \lesssim \mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^q] \lesssim d^{q/2}$. Hence, $\|\mathbb{E}_{x \sim \mathcal{D}_d} [h_1(x) xx^\top]\| \gtrsim d^{(p+1)/2}$. Also

$$\begin{aligned}
\left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[h_1(x) xx^\top \right] \right\| &\leq \max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} \left[h_1(x) (x^\top a)^2 \right] \\
&\leq \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [h_1^2(x)] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(x^\top a)^4] \right)^{1/2} \\
&\lesssim d^{(p+1)/2}.
\end{aligned}$$

(III) Bounding $(\mathbb{E}_{x \sim \mathcal{D}_d} [h_1^4(x)])^{1/4}$.

$$\left(\mathbb{E}_{x \sim \mathcal{D}_d} [h_1^4(x)] \right)^{1/4} \lesssim d^{(p+1)/2}.$$

Define function $B(x) = h(x) xx^\top \in \mathbb{R}^{d \times d}$, $\forall i \in [n]$. Let $\bar{B} = \mathbb{E}_{x \sim \mathcal{D}_d} [h(x) xx^\top]$. Therefore by applying Corollary B.8, we obtain for any $0 < \epsilon < 1$, if

$$n \geq \epsilon^{-2} d \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^t$,

$$\left\| \frac{1}{|S|} \sum_{x \in S} \|x\|^{p+1} xx^\top - \mathbb{E}_{x \sim \mathcal{D}_d} \left[\|x\|^{p+1} xx^\top \right] \right\| \lesssim \delta d^{(p+1)/2}.$$

Therefore we have with probability at least $1 - 1/d^t$,

$$\left\| \frac{1}{|S|} \sum_{x \in S} \|x\|^{p+1} xx^\top \right\| \lesssim d^{(p+1)/2}. \tag{23}$$

□

Claim D.18. For each $i \in [k]$, if $n \geq d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,i}^{(1,2)}\| \lesssim k\sigma_1^p \|W^a - W^*\|,$$

holds with probability at least $1 - 1/d^{4t}$.

Proof. Recall the definition of $\Delta_{i,i}^{(1,2)}$,

$$\Delta_{i,i}^{(1,2)} = v_i^* \frac{1}{|S|} \sum_{x \in S} \sum_{q=1}^k v_q^* (\phi(w_q^{a\top} x) - \phi(w_q^{*\top} x)) (\phi''(w_i^\top x) - \phi''(w_i^{a\top} x)) x x^\top$$

In order to upper bound the $\|\Delta_{i,i}^{(1,2)}\|$, it suffices to upper bound the spectral norm of this quantity,

$$\frac{1}{|S|} \sum_{x \sim S} (\phi(w_q^{a\top} x) - \phi(w_q^{*\top} x)) (\phi''(w_i^\top x) - \phi''(w_i^{a\top} x)) x x^\top.$$

By Property 3.1, we have

$$|\phi(w_q^{a\top} x) - \phi(w_q^{*\top} x)| \lesssim L_1 (|w_q^{a\top} x|^p + |w_q^{*\top} x|) |(w_q^* - w_q^a)^\top x|.$$

By Property 3.3, we have $|\phi''(w_i^\top x) - \phi''(w_i^{a\top} x)| \leq 2L_2$.

For the second part $|w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x| x x^\top$, according to Eq. (15), we have with probability $1 - d^{-t}$, if $n \geq d \text{poly}(\log d, t)$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[|w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x| x x^\top \right] - \frac{1}{|S|} \sum_{x \in S} |w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x| x x^\top \right\| \lesssim \delta \|w_q^*\|^p \|w_q^* - w_q^a\|.$$

Also, note that

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} \left[|w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x| x x^\top \right] \right\| \lesssim \|w_q^*\|^p \|w_q^* - w_q^a\|.$$

Thus, we obtain

$$\left\| \frac{1}{|S|} \sum_{x \in S} |w_q^{*\top} x|^p |(w_q^* - w_q^a)^\top x| x x^\top \right\| \lesssim \|W^a - W^*\| \sigma_1^p. \quad (24)$$

□

Claim D.19. For each $i \in [k]$, if $n \geq d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,i}^{(2)}\| \lesssim k v_{\max}^2 \sigma_1^p \|W - W^a\| d^{(p+1)/2}$$

holds with probability $1 - 1/d^{4t}$.

Proof. We have

$$\begin{aligned} \|\Delta_{i,i}^{(2)}\| &\leq v_i^{*2} \left\| \frac{1}{|S|} \sum_{x \in S} \left((\phi'(w_i^\top x) - \phi'(w_i^{a\top} x)) \cdot (\phi'(w_i^\top x) + \phi'(w_i^{a\top} x)) \right) x x^\top \right\| \\ &\leq v_i^{*2} \left\| \frac{1}{|S|} \sum_{x \in S} \left(L_2 |(w_i - w_i^a)^\top x| \cdot L_1 (|w_i^\top x|^p + |w_i^{a\top} x|^p) \right) x x^\top \right\| \\ &\leq v_i^{*2} \|W - W^a\| \left\| \frac{1}{|S|} \sum_{x \in S} \left(L_2 \|x\| \cdot L_1 (\|w_i\|^p \|x\|^p + |w_i^{a\top} x|^p) \right) x x^\top \right\|. \end{aligned}$$

Applying Corollary B.8 finishes the proof.

□

Claim D.20. For each $i \in [k], j \in [k], i \neq j$, if $n \geq d \text{poly}(\log d, t)$, then

$$\|\Delta_{i,l}\| \lesssim v_{\max}^2 \sigma_1^p \|W^a - W\|$$

holds with probability $1 - d^{4t}$.

Proof. Recall the definition of $\Delta_{i,l}$,

$$\begin{aligned} \Delta_{i,l} &= v_i^* v_l^* \frac{1}{|S|} \sum_{x \in S} \left(\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^\top x) \phi'(w_l^{a\top} x) \right. \\ &\quad \left. + \phi'(w_i^\top x) \phi'(w_l^{a\top} x) - \phi'(w_i^{a\top} x) \phi'(w_l^{a\top} x) \right) x x^\top \\ &= v_i^* v_l^* \frac{1}{|S|} \sum_{x \in S} \left(\phi'(w_i^\top x) \phi'(w_l^\top x) - \phi'(w_i^\top x) \phi'(w_l^{a\top} x) \right) \\ &\quad + v_i^* v_l^* \frac{1}{|S|} \sum_{x \in S} \left(\phi'(w_i^\top x) \phi'(w_l^{a\top} x) - \phi'(w_i^{a\top} x) \phi'(w_l^{a\top} x) \right) x x^\top \\ &\preceq |v_i^* v_l^*| \frac{1}{|S|} \sum_{x \in S} \left(L_1 \|w_i\|^p L_2 \|w_l - w_l^a\| \|x\|^{p+1} + L_2 \|w_i - w_i^a\| \|x\| L_1 \|w_l^{a\top} x\|^p \right) x x^\top \end{aligned}$$

Applying Corollary B.8 completes the proof. \square

E Tensor Methods

E.1 Tensor Initialization Algorithm

We describe the details of each procedure in Algorithm 1 in this section.

a) Compute the subspace estimation from \widehat{P}_2 (Algorithm 3). Note that the eigenvalues of P_2 and \widehat{P}_2 are not necessarily nonnegative. However, only k of the eigenvalues will have large magnitude. So we can first compute the top- k eigenvectors/eigenvalues of both $C \cdot I + \widehat{P}_2$ and $C \cdot I - \widehat{P}_2$, where C is large enough such that $C \geq 2\|P_2\|$. Then from the $2k$ eigen-pairs, we pick the top- k eigenvectors with the largest eigenvalues in magnitude, which is executed in TOPK in Algorithm 3. For the outputs of TOPK, k_1, k_2 are the numbers of picked largest eigenvalues from $C \cdot I + \widehat{P}_2$ and $C \cdot I - \widehat{P}_2$ respectively and $\pi_1(i)$ returns the original index of i -th largest eigenvalue from $C \cdot I + \widehat{P}_2$ and similarly π_2 is for $C \cdot I - \widehat{P}_2$. Finally orthogonalizing the picked eigenvectors leads to an estimation of the subspace spanned by $\{w_1^* w_2^* \cdots w_k^*\}$. Also note that forming \widehat{P}_2 takes $O(n \cdot d^2)$ time and each step of the power method doing a multiplication between a $d \times d$ matrix and a $d \times k$ matrix takes $k \cdot d^2$ time by a naive implementation. Here we reduce this complexity from $O((k+n)d^2)$ to $O(knd)$. The idea is to compute each step of the power method without explicitly forming \widehat{P}_2 . We take $P_2 = M_2$ as an example; other cases are similar. In Algorithm 3, let the step $\widehat{P}_2 V$ be calculated as $\widehat{P}_2 V = \frac{1}{|S|} \sum_{(x,y) \in S} y(x x^\top V) - V$. Now each iteration only needs $O(knd)$ time. Furthermore, the number of iterations required will be a small number, since the power method has a linear convergence rate and as an initialization method we don't need a very accurate solution. The detailed algorithm is shown in Algorithm 3. The approximation error bound of \widehat{P}_2 to P_2 is provided in Lemma E.5. Lemma E.6 provides the theoretical bound for Algorithm 3.

b) Form and decompose the 3rd-moment \widehat{R}_3 (Algorithm 1 in [KCL15]). We apply the non-orthogonal tensor factorization algorithm, Algorithm 1 in [KCL15], to decompose \widehat{R}_3 . According to Theorem 3 in [KCL15], when \widehat{R}_3 is close enough to R_3 , the output of the algorithm, \widehat{u}_i will

close enough to $s_i V^\top \bar{w}_i^*$, where s_i is an unknown sign. Lemma E.10 provides the error bound for $\|\widehat{R}_3 - R_3\|$.

c) Recover the magnitude of w_i^* and the signs s_i, v_i^* (Algorithm 4). For Algorithm 4, we only consider homogeneous functions. Hence we can assume $v_i^* \in \{-1, 1\}$ and there exist some universal constants c_j such that $m_{j,i} = c_j \|w_i^*\|^{p+1}$ for $j = 1, 2, 3, 4$, where $p + 1$ is the degree of homogeneity. Note that different activations have different situations even under Assumption 5.3. In particular, if $M_4 = M_2 = 0$, $\phi(\cdot)$ is an odd function and we only need to know $s_i v_i^*$. If $M_3 = M_1 = 0$, $\phi(\cdot)$ is an even function, so we don't care about what s_i is.

Let's describe the details for Algorithm 4. First define two quantities Q_1 and Q_2 ,

$$Q_1 = M_{l_1}(I, \underbrace{\alpha, \dots, \alpha}_{(l_1-1) \text{ } \alpha\text{'s}}) = \sum_{i=1}^k v_i^* c_{l_1} \|w_i^*\|^{p+1} (\alpha^\top \bar{w}_i^*)^{l_1-1} \bar{w}_i^*, \quad (25)$$

$$Q_2 = M_{l_2}(V, V, \underbrace{\alpha, \dots, \alpha}_{(l_2-2) \text{ } \alpha\text{'s}}) = \sum_{i=1}^k v_i^* c_{l_2} \|w_i^*\|^{p+1} (\alpha^\top \bar{w}_i^*)^{l_2-2} (V^\top \bar{w}_i^*) (V^\top \bar{w}_i^*)^\top, \quad (26)$$

where $l_1 \geq 1$ such that $M_{l_1} \neq 0$ and $l_2 \geq 2$ such that $M_{l_2} \neq 0$. There are possibly multiple choices for l_1 and l_2 . We will discuss later on how to choose them. Now we solve two linear systems.

$$z^* = \operatorname{argmin}_{z \in \mathbb{R}^k} \left\| \sum_{i=1}^k z_i s_i \bar{w}_i^* - Q_1 \right\|, \text{ and } r^* = \operatorname{argmin}_{r \in \mathbb{R}^k} \left\| \sum_{i=1}^k r_i V^\top \bar{w}_i^* (V^\top \bar{w}_i^*)^\top - Q_2 \right\|_F. \quad (27)$$

The solutions of the above linear systems are

$$z_i^* = v_i^* s_i^{l_1} c_{l_1} \|w_i^*\|^{p+1} (\alpha^\top s_i \bar{w}_i^*)^{l_1-1}, \text{ and } r_i^* = v_i^* s_i^{l_2} c_{l_2} \|w_i^*\|^{p+1} (\alpha^\top s_i \bar{w}_i^*)^{l_2-2}.$$

We can approximate $s_i \bar{w}_i^*$ by $V \hat{u}_i$ and approximate Q_1 and Q_2 by their empirical versions \widehat{Q}_1 and \widehat{Q}_2 respectively. Hence, in practice, we solve

$$\widehat{z} = \operatorname{argmin}_{z \in \mathbb{R}^k} \left\| \sum_{i=1}^k z_i V \hat{u}_i - \widehat{Q}_1 \right\|, \text{ and } \widehat{r} = \operatorname{argmin}_{r \in \mathbb{R}^k} \left\| \sum_{i=1}^k r_i \hat{u}_i \hat{u}_i^\top - \widehat{Q}_2 \right\|_F \quad (28)$$

So we have the following approximations,

$$\widehat{z}_i \approx v_i^* s_i^{l_1} c_{l_1} \|w_i^*\|^{p+1} (\alpha^\top V \hat{u}_i)^{l_1-1}, \text{ and } \widehat{r}_i \approx v_i^* s_i^{l_2} c_{l_2} \|w_i^*\|^{p+1} (\alpha^\top V \hat{u}_i)^{l_2-2}, \forall i \in [k].$$

In Lemma E.13 and Lemma E.14, we provide robustness of the above two linear systems, i.e., the solution errors, $\|\widehat{z} - z^*\|$ and $\|\widehat{r} - r^*\|$, are bounded under small perturbations of \bar{w}_i^* , Q_1 and Q_2 . Recall that the final goal is to approximate $\|w_i^*\|$ and the signs v_i^*, s_i . Now we can approximate $\|w_i^*\|$ by $(|\widehat{z}_i / (c_{l_1} (\alpha^\top V \hat{u}_i)^{l_1-1})|)^{1/(p+1)}$. To recover v_i^*, s_i , we need to note that if l_1 and l_2 are both odd or both even, we can't recover both v_i^* and s_i . So we consider the following situations,

1. If $M_1 = M_3 = 0$, we choose $l_1 = l_2 = \min\{j \in \{2, 4\} | M_j \neq 0\}$. Return $v_i^{(0)} = \operatorname{sign}(\widehat{r}_i c_{l_2})$ and $s_i^{(0)}$ being -1 or 1 .
2. If $M_2 = M_4 = 0$, we choose $l_1 = \min\{j \in \{1, 3\} | M_j \neq 0\}$, $l_2 = 3$. Return $v_i^{(0)}$ being -1 or 1 and $s_i^{(0)} = \operatorname{sign}(v_i^{(0)} \widehat{z}_i c_{l_1})$.

Algorithm 3 Power Method via Implicit Matrix-Vector Multiplication

```

1: procedure POWERMETHOD( $\widehat{P}_2, k$ )
2:    $C \leftarrow 3\|\widehat{P}_2\|$ ,  $T \leftarrow$  a large enough constant.
3:   Initial guess  $\widehat{V}_1^{(0)} \in \mathbb{R}^{d \times k}$ ,  $\widehat{V}_2^{(0)} \in \mathbb{R}^{d \times k}$ 
4:   for  $t = 1 \rightarrow T$  do
5:      $\widehat{V}_1^{(t)} \leftarrow \text{QR}(C\widehat{V}_1^{(t-1)} + \widehat{P}_2\widehat{V}_1^{(t-1)})$   $\triangleright \widehat{P}_2\widehat{V}_1^{(t-1)}$  is not calculated directly, see Sec. E.1(a)
6:      $\widehat{V}_2^{(t)} \leftarrow \text{QR}(C\widehat{V}_2^{(t-1)} - \widehat{P}_2\widehat{V}_2^{(t-1)})$ 
7:   end for
8:   for  $j = 1, 2$  do
9:      $\widehat{V}_j^{(T)} \leftarrow [\widehat{v}_{j,1} \ \widehat{v}_{j,2} \ \cdots \ \widehat{v}_{j,k}]$ 
10:    for  $i = 1 \rightarrow k$  do
11:       $\lambda_{j,i} \leftarrow |\widehat{v}_{j,i}^\top \widehat{P}_2 \widehat{v}_{j,i}|$   $\triangleright$  Calculate the absolute of eigenvalues
12:    end for
13:  end for
14:   $\pi_1, \pi_2, k_1, k_2 \leftarrow \text{TOPK}(\lambda, k)$   $\triangleright \pi_j : [k_j] \rightarrow [k]$  and  $k_1 + k_2 = k$ , see Sec. E.1(a)
15:  for  $j = 1, 2$  do
16:     $V_j \leftarrow [\widehat{v}_{j,\pi_j(1)} \ \widehat{v}_{j,\pi_j(2)} \ \cdots \ \widehat{v}_{j,\pi_j(k_j)}]$ 
17:  end for
18:   $\widetilde{V}_2 \leftarrow \text{QR}((I - V_1 V_1^\top) V_2)$ 
19:   $V \leftarrow [V_1, \widetilde{V}_2]$ 
20:  return  $V$ 
21: end procedure

```

3. Otherwise, we choose $l_1 = \min\{j \in \{1, 3\} | M_j \neq 0\}$, $l_2 = \min\{j \in \{2, 4\} | M_j \neq 0\}$. Return $v_i^{(0)} = \text{sign}(\widehat{r}_i c_{l_2})$ and $s_i^{(0)} = \text{sign}(v_i^{(0)} \widehat{z}_i c_{l_1})$.

The 1st situation corresponds to part 3 of Assumption 5.3, where s_i doesn't matter, and the 2nd situation corresponds to part 4 of Assumption 5.3, where only $s_i v_i^*$ matters. So we recover $\|w_i^*\|$ to some precision and v_i^*, s_i exactly provided enough samples. The recovery of w_i^* and v_i^* then follows.

Sample complexity: We use matrix Bernstein inequality to bound the error between \widehat{P}_2 and P_2 , which requires $\widetilde{\Omega}(d)$ samples (Lemma E.5). To bound the estimation error between R_3 and \widehat{R}_3 , we flatten the tensor to a matrix and then use matrix Bernstein inequality to bound the error, which requires $\widetilde{\Omega}(k^3)$ samples (Lemma E.10). In Algorithm 4, we also need to approximate a \mathbb{R}^d vector and a $\mathbb{R}^{k \times k}$ matrix, which also requires $\widetilde{\Omega}(d)$. Thus, taking $\widetilde{O}(d) + \widetilde{O}(k^3)$ samples is sufficient.

Time complexity: In Part a), by using a specially designed power method, we only need $O(knd)$ time to compute the subspace estimation V . Part b) needs $O(knd)$ to form \widehat{R}_3 and the tensor factorization needs $O(k^3)$ time. Part c) requires calculation of $d \times k$ and $k^2 \times k$ linear systems in Eq. (28), which takes at most $O(knd)$ running time. Hence, the total time complexity is $O(knd)$.

E.2 Main Result for Tensor Methods

The goal of this Section is to prove Theorem 5.6.

Theorem 5.6. *Let the activation function be homogeneous satisfying Assumption 5.3. For any $0 < \epsilon < 1$ and $t \geq 1$, if $|S| \geq \epsilon^{-2} \cdot d \cdot \text{poly}(t, k, \kappa, \log d)$, then there exists an algorithm (Algorithm 1) that takes $|S|k \cdot \widetilde{O}(d)$ time and outputs a matrix $W^{(0)} \in \mathbb{R}^{d \times k}$ and a vector $v^{(0)} \in \mathbb{R}^k$ such that, with*

Algorithm 4 Recovery of the Ground Truth Parameters of the Neural Network, i.e., w_i^* and v_i^*

```

1: procedure RECMAGSIGN( $V, \{\hat{u}_i\}_{i \in [k]}, S$ )
2:   if  $M_1 = M_3 = 0$  then
3:      $l_1 \leftarrow l_2 \leftarrow \min\{j \in \{2, 4\} | M_j \neq 0\}$ 
4:   else if  $M_2 = M_4 = 0$  then
5:      $l_1 \leftarrow \min\{j \in \{1, 3\} | M_j \neq 0\}, l_2 \leftarrow 3$ 
6:   else
7:      $l_1 \leftarrow \min\{j \in \{1, 3\} | M_j \neq 0\}, l_2 \leftarrow \min\{j \in \{2, 4\} | M_j \neq 0\}.$ 
8:   end if
9:    $S_1, S_2 \leftarrow \text{PARTITION}(S, 2)$   $\triangleright |S_1|, |S_2| = \tilde{\Omega}(d)$ 
10:  Choose  $\alpha$  to be a random unit vector
11:   $\hat{Q}_1 \leftarrow \mathbb{E}_{S_1}[Q_1]$   $\triangleright \hat{Q}_1$  is the empirical version of  $Q_1$  (defined in Eq.(25))
12:   $\hat{Q}_2 \leftarrow \mathbb{E}_{S_2}[Q_2]$   $\triangleright \hat{Q}_2$  is the empirical version of  $Q_2$  (defined in Eq.(26))
13:   $\hat{z} \leftarrow \operatorname{argmin}_z \left\| \sum_{i=1}^k z_i V \hat{u}_i - \hat{Q}_1 \right\|$ 
14:   $\hat{r} \leftarrow \operatorname{argmin}_r \left\| \sum_{i=1}^k r_i \hat{u}_i \hat{u}_i^\top - \hat{Q}_2 \right\|_F$ 
15:   $v_i^{(0)} \leftarrow \operatorname{sign}(\hat{r}_i c_{l_2})$ 
16:   $s_i^{(0)} \leftarrow \operatorname{sign}(v_i^{(0)} \hat{z}_i c_{l_1})$ 
17:   $w_i^{(0)} \leftarrow s_i^{(0)} (|\hat{z}_i / (c_{l_1} (\alpha^\top V \hat{u}_i)^{l_1-1})|)^{1/(p+1)} V \hat{u}_i$ 
18:  return  $v_i^{(0)}, w_i^{(0)}$ 
19: end procedure

```

probability at least $1 - d^{-\Omega(t)}$,

$$\|W^{(0)} - W^*\|_F \leq \epsilon \cdot \operatorname{poly}(k, \kappa) \|W^*\|_F, \text{ and } v_i^{(0)} = v_i^*.$$

Proof. The success of Algorithm 1 depends on two approximations. The first is the estimation of the normalized w_i^* up to some unknown sign flip, i.e., the error $\|\bar{w}_i^* - s_i V \hat{u}_i\|$ for some $s_i \in \{-1, 1\}$. The second is the estimation of the magnitude of w_i^* and the signs v_i^*, s_i which is conducted in Algorithm 4.

For the first one,

$$\begin{aligned} \|\bar{w}_i^* - s_i V \hat{u}_i\| &\leq \|V V^\top \bar{w}_i^* - \bar{w}_i^*\| + \|V V^\top \bar{w}_i^* - V s_i \hat{u}_i\| \\ &= \|V V^\top \bar{w}_i^* - \bar{w}_i^*\| + \|V^\top \bar{w}_i^* - s_i \hat{u}_i\|, \end{aligned} \quad (29)$$

where the first step follows by triangle inequality, the second step follows by $V^\top V = I$.

We can upper bound $\|V V^\top \bar{w}_i^* - \bar{w}_i^*\|$,

$$\begin{aligned} \|V V^\top \bar{w}_i^* - \bar{w}_i^*\| &\leq (\|\hat{P}_2 - P_2\| / \sigma_k(P_2) + \epsilon) \\ &\leq (\operatorname{poly}(k, \kappa) \|\hat{P}_2 - P_2\| + \epsilon) \\ &\leq \operatorname{poly}(k, \kappa) \epsilon, \end{aligned} \quad (30)$$

where the first step follows by Lemma E.6, the second step follows by $\sigma_k(P_2) \geq 1/\operatorname{poly}(k, \kappa)$, and the last step follows by $\|\hat{P}_2 - P_2\| \leq \epsilon \operatorname{poly}(k, \kappa)$ if the number of samples is proportional to $\tilde{O}(d/\epsilon^2)$ as shown in Lemma E.5.

We can upper bound $\|V^\top \bar{w}_i^* - s_i \hat{u}_i\|$,

$$\|V^\top \bar{w}_i^* - s_i \hat{u}_i\| \leq \operatorname{poly}(k, \kappa) \|\hat{R}_3 - R_3\| \leq \epsilon \operatorname{poly}(k, \kappa), \quad (31)$$

where the first step follows by Theorem 3 in [KCL15], and the last step follows by $\|\widehat{R}_3 - R_3\| \leq \epsilon \text{poly}(k, \kappa)$ if the number of samples is proportional to $\widetilde{O}(k^2/\epsilon^2)$ as shown in Lemma E.10.

Combining Eq. (29), (30) and (31) together,

$$\|\overline{w}_i^* - s_i V \widehat{u}_i\| \leq \epsilon \text{poly}(k, \kappa).$$

For the second one, we can bound the error of the estimation of moments, Q_1 and Q_2 , using number of samples proportional to $\widetilde{O}(d)$ by Lemma E.12 and Lemma E.5 respectively. The error of the solutions of the linear systems Eq.(28) can be bounded by $\|Q_1 - \widehat{Q}_1\|, \|Q_2 - \widehat{Q}_2\|, \|\widehat{u}_i - V^\top \overline{w}_i^*\|$ and $\|(I - VV^\top)\overline{w}_i^*\|$ according to Lemma E.13 and Lemma E.14. Then we can bound the error of the output of Algorithm 4. Furthermore, since v_i^* 's are discrete values, they can be exactly recovered. All the sample complexities mentioned in the above lemmata are linear in dimension and polynomial in other factors to achieve a constant error. So accumulating all these errors we complete our proof. \square

Remark E.1. *The proofs of these lemmata for Theorem 1 can be found in the following sections. Note that these lemmata also hold for any activations satisfying Property 3.1 and Assumption 5.3. However, since we are unclear how to implement the last step of Algorithm 1 (Algorithm 4) for general non-homogeneous activations, we restrict our theorem to homogeneous activations only.*

E.3 Error Bound for the Subspace Spanned by the Weight Matrix

E.3.1 Error Bound for the Second-order Moment in Different Cases

Lemma E.2. *Let M_2 be defined as in Definition 5.1. Let \widehat{M}_2 be the empirical version of M_2 , i.e.,*

$$\widehat{M}_2 = \frac{1}{|S|} \sum_{(x,y) \in S} y \cdot (x \otimes x - I_d),$$

where S denote a set of samples from Distribution \mathcal{D} defined in Eq. (1). Assume $M_2 \neq 0$, i.e., $m_i^{(2)} \neq 0$ for any i . Then for any $0 < \epsilon < 1, t \geq 1$, if

$$|S| \geq \max_{i \in [k]} (\|w_i^*\|^{p+1} / |m_{2,i}| + 1) \cdot \epsilon^{-2} d \text{poly}(\log d, t)$$

with probability at least $1 - d^{-t}$,

$$\|M_2 - \widehat{M}_2\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{2,i}|.$$

Proof. Recall that, for each sample (x, y) , $y = \sum_{i=1}^k v_i^* \phi(w_i^{\top} x)$. We consider each component $i \in [k]$. Define function $B_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ such that

$$B_i(x) = \phi(w_i^{\top} x) \cdot (x \otimes x - I_d).$$

Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = |\int_0^z \phi'(s) ds| \leq L_1 / (p+1) |z|^{p+1}$, which follows Property 3.1.

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned} \|B_i(x)\| &\lesssim \left(\frac{L_1}{p+1} |w_i^{\top} x|^{p+1} + |\phi(0)| \right) (\|x_j\|^2 + 1) \\ &\lesssim \left(\frac{L_1}{p+1} \|w_i^*\|^{p+1} + |\phi(0)| \right) d \text{poly}(\log d, t) \end{aligned}$$

where the last step follows by Fact B.1 and Fact B.2.

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{2,i} \bar{w}_i^* \bar{w}_i^{*\top}$. Therefore, $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| = |m_{2,i}|$.

(III) Bounding $\max(\mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x) B_i(x)^\top\|, \mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x)^\top B_i(x)\|)$.

Note that $B_i(x)$ is a symmetric matrix, thus it suffices to only bound one of them.

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i^2(x)] \right\| \lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^4] \right)^{1/2} \lesssim \left(\frac{L_1}{p+1} \|w_i^*\|^{p+1} + |\phi(0)| \right)^2 d.$$

(IV) Bounding $\max_{\|a\|=\|b\|=1} (\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) b)^2])$.

Note that $B_i(x)$ is a symmetric matrix, thus it suffices to consider the case where $a = b$.

$$\max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) a)^2] \right)^{1/2} \lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi^4(w_i^{*\top} x)] \right)^{1/4} \lesssim \frac{L_1}{p+1} \|w_i^*\|^{p+1} + |\phi(0)|.$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$n \gtrsim \frac{L^2 d + |m_{2,i}|^2 + L |m_{2,i}| d \cdot \text{poly}(\log d, t) \epsilon}{\epsilon^2 |m_{2,i}|^2} t \log d$$

with probability at least $1 - 1/d^t$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{2,i}|.$$

□

Lemma E.3. Let M_3 be defined as in Definition 5.1. Let \widehat{M}_3 be the empirical version of M_3 , i.e.,

$$\widehat{M}_3 = \frac{1}{|S|} \sum_{(x,y) \in S} y \cdot (x^{\otimes 3} - x \otimes I),$$

where S denote a set of samples (each sample is i.i.d. sampled from Distribution \mathcal{D} defined in Eq. (1)). Assume $M_3 \neq 0$, i.e., $m_{3,i} \neq 0$ for any i . Let α be a fixed unit vector. Then for any $0 < \epsilon < 1, t \geq 1$, if

$$|S| \geq \max_{i \in [k]} (\|w_i^*\|^{p+1} / |m_{3,i} (\bar{w}_i^{*\top} \alpha)|^2 + 1) \cdot \epsilon^{-2} d \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^t$,

$$\|M_3(I, I, \alpha) - \widehat{M}_3(I, I, \alpha)\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{3,i} (\bar{w}_i^{*\top} \alpha)|.$$

Proof. Since $y = \sum_{i=1}^k v_i^* \phi(w_i^{*\top} x)$. We consider each component $i \in [k]$.

Define function $B_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ such that

$$B_i(x) = [\phi(w_i^{*\top} x) \cdot (x^{\otimes 3} - x \otimes I)](I, I, \alpha) = \phi(w_i^{*\top} x) \cdot ((x^\top \alpha) x^{\otimes 2} - \alpha^\top x I - \alpha x^\top - x \alpha^\top).$$

Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = |\int_0^z \phi'(s) ds| \leq \frac{L_1}{p+1} |z|^{p+1} \lesssim |z|^{p+1}$, which follows Property 3.1. In order to apply Lemma B.7, we need to bound the following four quantities,

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned}
\|B_i(x)\| &= \|\phi(w_i^{*\top} x) \cdot ((x^\top \alpha)x^{\otimes 2} - \alpha^\top x I_d - \alpha x^\top - x \alpha^\top)\| \\
&\leq |\phi(w_i^{*\top} x)| \cdot \|(x^\top \alpha)x^{\otimes 2} - \alpha^\top x I_d - \alpha x^\top - x \alpha^\top\| \\
&\lesssim (|w_i^{*\top} x|^{p+1} + |\phi(0)|) \|(x^\top \alpha)x^{\otimes 2} - \alpha^\top x I_d - \alpha x^\top - x \alpha^\top\| \\
&\lesssim (|w_i^{*\top} x|^{p+1} + |\phi(0)|)(|x^\top \alpha| \|x\|^2 + 3|\alpha^\top x|),
\end{aligned}$$

where the third step follows by definition of $g(z)$, and last step follows by definition of spectral norm and triangle inequality.

Using Fact B.1 and Fact B.2, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^{4t})$,

$$\|B_i(x)\| \lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)d \text{poly}(\log d, t).$$

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{3,i}(\bar{w}_i^{*\top} \alpha) \bar{w}_i^* \bar{w}_i^{*\top}$. Therefore, $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| = |m_{3,i}(\bar{w}_i^{*\top} \alpha)|$.

(III) Bounding $\max(\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)B_i(x)^\top]\|, \|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)^\top B_i(x)]\|)$.

Because matrix $B_i(x)$ is symmetric, thus it suffices to bound one of them,

$$\begin{aligned}
\left\| \mathbb{E}_{x \sim \mathcal{D}_d}[B_i^2(x)] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(x^\top \alpha)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\|x\|^4] \right)^{1/2} \\
&\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 d.
\end{aligned}$$

(IV) Bounding $\max_{\|a\|=\|b\|=1}(\mathbb{E}_{x \sim \mathcal{D}_d}[(a^\top B_i(x)b)^2])^{1/2}$.

$$\max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(a^\top B_i(x)a)^2] \right)^{1/2} \lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\phi^4(w_i^{*\top} x)] \right)^{1/4} \lesssim \|w_i^*\|^{p+1} + |\phi(0)|.$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$|S| \gtrsim \frac{L^2 d + |m_{3,i}(\bar{w}_i^{*\top} \alpha)|^2 + L|m_{3,i}(\bar{w}_i^{*\top} \alpha)|d \cdot \text{poly}(\log d, t)\epsilon}{\epsilon^2 |m_{3,i}(\bar{w}_i^{*\top} \alpha)|^2} \cdot t \log d$$

with probability at least $1 - d^{-t}$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{3,i}(\bar{w}_i^{*\top} \alpha)|.$$

□

Lemma E.4. Let M_4 be defined as in Definition 5.1. Let \widehat{M}_4 be the empirical version of M_4 , i.e.,

$$\widehat{M}_4 = \frac{1}{|S|} \sum_{(x,y) \in S} y \cdot (x^{\otimes 4} - (x \otimes x) \widetilde{\otimes} I + I \widetilde{\otimes} I),$$

where S denote a set of samples (where each sample is i.i.d. sampled from Distribution \mathcal{D} defined in Eq. (1)). Assume $M_4 \neq 0$, i.e., $m_{4,i} \neq 0$ for any i . Let α be a fixed unit vector. Then for any $0 < \epsilon < 1, t \geq 1$, if

$$|S| \geq \max_{i \in [k]} (\|w_i^*\|^{p+1} / |m_{4,i}(\bar{w}_i^{*\top} \alpha)|^2 + 1)^2 \cdot \epsilon^{-2} \cdot d \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^t$,

$$\|M_4(I, I, \alpha, \alpha) - \widehat{M}_4(I, I, \alpha, \alpha)\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{4,i}| (\bar{w}_i^{*\top} \alpha)^2.$$

Proof. Since $y = \sum_{i=1}^k v_i^* \phi(w_i^{*\top} x)$. We consider each component $i \in [k]$.

Define function $B_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ such that

$$\begin{aligned} B_i(x) &= [\phi(w_i^{*\top} x) \cdot (x^{\otimes 4} - (x \otimes x) \widetilde{\otimes} I + I \widetilde{\otimes} I)](I, I, \alpha, \alpha) \\ &= \phi(w_i^{*\top} x) \cdot ((x^\top \alpha)^2 x^{\otimes 2} - (\alpha^\top x)^2 I - 2(\alpha^\top x)(x \alpha^\top + \alpha x^\top) - x x^\top + 2\alpha \alpha^\top + I). \end{aligned}$$

Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = |\int_0^z \phi'(s) ds| \leq L_1/(p+1)|z|^{p+1} \lesssim |z|^{p+1}$, which follows Property 3.1.

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned} \|B_i(x)\| &\lesssim |\phi(w_i^{*\top} x)| \cdot ((x^\top \alpha)^2 \|x\|^2 + 1 + \|x\|^2 + (\alpha^\top x)^2) \\ &\lesssim (|w_i^{*\top} x|^{p+1} + |\phi(0)|) \cdot ((x^\top \alpha)^2 \|x\|^2 + 1 + \|x\|^2 + (\alpha^\top x)^2) \end{aligned}$$

Using Fact B.1 and Fact B.2, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^{4t})$,

$$\|B_i(x)\| \lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) d \text{poly}(\log d, t).$$

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{4,i}(\bar{w}_i^{*\top} \alpha)^2 \bar{w}_i^* \bar{w}_i^{*\top}$. Therefore, $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| = |m_{4,i}| (\bar{w}_i^{*\top} \alpha)^2$.

(III) Bounding $\max(\mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x) B_i(x)^\top\|, \mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x)^\top B_i(x)\|)$.

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)^2] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(x^\top \alpha)^8] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 d. \end{aligned}$$

(IV) Bounding $\max_{\|a\|=\|b\|=1} (\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) b)^2])^{1/2}$.

$$\max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) a)^2] \right)^{1/2} \lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi^4(w_i^{*\top} x)] \right)^{1/4} \lesssim \|w_i^*\|^{p+1} + |\phi(0)|.$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$n \gtrsim \frac{L^2 d + |m_{4,i}|^2 (\bar{w}_i^{*\top} \alpha)^4 + L |m_{4,i}| (\bar{w}_i^{*\top} \alpha)^2 d \text{poly}(\log d, t) \epsilon}{\epsilon^2 |m_{4,i}|^2 (\bar{w}_i^{*\top} \alpha)^4} \cdot t \log d$$

with probability at least $1 - d^{-t}$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{4,i}| (\bar{w}_i^{*\top} \alpha)^2.$$

□

E.3.2 Error Bound for the Second-order Moment

The goal of this section is to prove Lemma E.5, which shows we can approximate the second order moments up to some precision by using linear sample complexity in d .

Lemma E.5 (Estimation of the second order moment). *Let P_2 and j_2 be defined in Definition 5.4. Let S denote a set of i.i.d. samples generated from distribution \mathcal{D} (defined in (1)). Let \widehat{P}_2 be the empirical version of P_2 using dataset S , i.e., $\widehat{P}_2 = \mathbb{E}_S[P_2]$. Assume the activation function satisfies Property 3.1 and Assumption 5.3. Then for any $0 < \epsilon < 1$ and $t \geq 1$, and $m_0 = \min_{i \in [k]} \{|m_{j_2, i}|^2 (\bar{w}_i^{*\top} \alpha)^{2(j_2-2)}\}$, if*

$$|S| \gtrsim \sigma_1^{2p+2} \cdot d \cdot \text{poly}(t, \log d) / (\epsilon^2 m_0),$$

then with probability at least $1 - d^{-\Omega(t)}$,

$$\|P_2 - \widehat{P}_2\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{j_2, i} (\bar{w}_i^{*\top} \alpha)^{j_2-2}|.$$

Proof. We have shown the bound for $j_2 = 2, 3, 4$ in Lemma E.2, Lemma E.3 and Lemma E.4 respectively. To summarize, for any $0 < \epsilon < 1$ we have if

$$|S| \geq \max_{i \in [k]} \left\{ \frac{(|w_i^*|^{p+1} + |\phi(0)| + |m_{j_2, i} (\bar{w}_i^{*\top} \alpha)^{(j_2-2)}|)^2}{|m_{j_2, i}|^2 (\bar{w}_i^{*\top} \alpha)^{2(j_2-2)}} \right\} \cdot \epsilon^{-2} d \text{poly}(\log d, t)$$

with probability at least $1 - d^{-t}$,

$$\|P_2 - \widehat{P}_2\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{j_2, i} (\bar{w}_i^{*\top} \alpha)^{j_2-2}|.$$

□

E.3.3 Subspace Estimation Using Power Method

Lemma E.6 shows a small number of power iterations can estimate the subspace of $\{w_i^*\}_{i \in [k]}$ to some precision, which provides guarantees for Algorithm 3.

Lemma E.6 (Bound on subspace estimation). *Let P_2 be defined as in Definition 5.4 and \widehat{P}_2 be its empirical version. Let $U \in \mathbb{R}^{d \times k}$ be the orthogonal column span of $W^* \in \mathbb{R}^{d \times k}$. Assume $\|\widehat{P}_2 - P_2\| \leq \sigma_k(P_2)/10$. Let C be a large enough positive number such that $C > 2\|P_2\|$. Then after $T = O(\log(1/\epsilon))$ iterations, the output of Algorithm 3, $V \in \mathbb{R}^{d \times k}$, will satisfy*

$$\|UU^\top - VV^\top\| \lesssim \|\widehat{P}_2 - P_2\| / \sigma_k(P_2) + \epsilon,$$

which implies

$$\|(I - VV^\top)w_i^*\| \lesssim (\|\widehat{P}_2 - P_2\| / \sigma_k(P_2) + \epsilon) \|w_i^*\|.$$

Proof. According to Weyl's inequality, we are able to pick up the correct numbers of positive eigenvalues and negative eigenvalues in Algorithm 3 as long as \widehat{P}_2 and P_2 are close enough.

Let $U = [U_1 \ U_2] \in \mathbb{R}^{d \times k}$ be the eigenspace of $\text{span}\{w_1^* \ w_2^* \ \dots \ w_k^*\}$, where $U_1 \in \mathbb{R}^{d \times k_1}$ corresponds to positive eigenvalues of P_2 and $U_2 \in \mathbb{R}^{d \times k_2}$ is for negatives.

Let \bar{V}_1 be the top- k_1 eigenvectors of $CI + \hat{P}_2$. Let \bar{V}_2 be the top- k_2 eigenvectors of $CI - \hat{P}_2$. Let $\bar{V} = [\bar{V}_1 \bar{V}_2] \in \mathbb{R}^{d \times k}$.

According to Lemma 9 in [HK13], we have $\|(I - U_1 U_1^\top) \bar{V}_1\| \lesssim \|\hat{P}_2 - P_2\| / \sigma_k(P_2)$, $\|(I - U_2 U_2^\top) \bar{V}_2\| \lesssim \|\hat{P}_2 - P_2\| / \sigma_k(P_2)$. Using Fact B.4, we have $\|(I - UU^\top) \bar{V}\| = \|UU^\top - \bar{V} \bar{V}^\top\|$.

Let ϵ be the precision we want to achieve using power method. Let V_1 be the top- k_1 eigenvectors returned after $O(\log(1/\epsilon))$ iterations of power methods on $CI + \hat{P}_2$ and $V_2 \in \mathbb{R}^{d \times k_2}$ for $CI - \hat{P}_2$ similarly.

According to Theorem 7.2 in [AKZ12], we have $\|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \leq \epsilon$ and $\|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \leq \epsilon$. Let U_\perp be the complementary matrix of U . Then we have,

$$\begin{aligned}
\|(I - U_1 U_1^\top) \bar{V}_1\| &= \max_{\|a\|=1} \|(I - U_1 U_1^\top) \bar{V}_1 a\| \\
&= \max_{\|a\|=1} \|(U_\perp U_\perp^\top + U_2 U_2^\top) \bar{V}_1 a\| \\
&= \max_{\|a\|=1} \sqrt{\|U_\perp U_\perp^\top \bar{V}_1 a\|^2 + \|U_2 U_2^\top \bar{V}_1 a\|^2} \\
&\geq \max_{\|a\|=1} \|U_2 U_2^\top \bar{V}_1 a\| \\
&= \|U_2 U_2^\top \bar{V}_1\|, \tag{32}
\end{aligned}$$

where the first step follows by definition of spectral norm, the second step follows by $I = U_1 U_1^\top + U_2 U_2^\top + U_\perp U_\perp^\top$, the third step follows by $U_2^\top U_\perp = 0$, and last step follows by definition of spectral norm.

We can upper bound $\|(I - UU^\top) \bar{V}\|$,

$$\begin{aligned}
&\|(I - UU^\top) \bar{V}\| \\
&\leq (\|(I - U_1 U_1^\top) \bar{V}_1\| + \|(I - U_2 U_2^\top) \bar{V}_2\| + \|U_2 U_2^\top \bar{V}_1\| + \|U_1 U_1^\top \bar{V}_2\|) \\
&\leq 2(\|(I - U_1 U_1^\top) \bar{V}_1\| + \|(I - U_2 U_2^\top) \bar{V}_2\|) \\
&\lesssim \|\hat{P}_2 - P_2\| / \sigma_k(P_2), \tag{33}
\end{aligned}$$

where the first step follows by triangle inequality, the second step follows by Eq. (32), and the last step follows by Lemma 9 in [HK13].

We define matrix R such that $\bar{V}_2 R = (I - V_1 V_1^\top) V_2$ is the QR decomposition of $(I - V_1 V_1^\top) V_2$, then we have

$$\begin{aligned}
&\|(I - \bar{V}_2 \bar{V}_2^\top) \bar{V}_2\| \\
&= \|(I - \bar{V}_2 \bar{V}_2^\top) (I - V_1 V_1^\top) V_2 R^{-1}\| \\
&= \|(I - \bar{V}_2 \bar{V}_2^\top) (I - \bar{V}_1 \bar{V}_1^\top + \bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top) V_2 R^{-1}\| \\
&\leq \underbrace{\|(I - \bar{V}_2 \bar{V}_2^\top) (I - \bar{V}_1 \bar{V}_1^\top) V_2 R^{-1}\|}_\alpha + \underbrace{\|(I - \bar{V}_2 \bar{V}_2^\top) \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\|}_\beta,
\end{aligned}$$

where the first step follows by $\bar{V}_2 = (I - V_1 V_1^\top) V_2 R^{-1}$, and the last step follows by triangle inequality.

Furthermore, we have,

$$\begin{aligned}
& \alpha + \beta \\
&= \|(I - \bar{V}_2 \bar{V}_2^\top - \bar{V}_1 \bar{V}_1^\top) V_2 R^{-1}\| + \|(I - \bar{V}_2 \bar{V}_2^\top)\| \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \\
&\leq \|(I - \bar{V}_2 \bar{V}_2^\top) V_2 R^{-1}\| + \|\bar{V}_1 \bar{V}_1^\top V_2 R^{-1}\| + \|(I - \bar{V}_2 \bar{V}_2^\top)\| \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \\
&\leq \|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \|R^{-1}\| + \|\bar{V}_1 \bar{V}_1^\top V_2\| \|R^{-1}\| + \|(I - \bar{V}_2 \bar{V}_2^\top)\| \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \\
&= \|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \|R^{-1}\| + \|\bar{V}_1 \bar{V}_1^\top V_2\| \|R^{-1}\| + \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \\
&\leq \|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \|R^{-1}\| + \|(I - \bar{V}_2 \bar{V}_2^\top) V_2\| \|R^{-1}\| + \|R^{-1}\| \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \\
&\leq (2\|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| + \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\|) \|R^{-1}\| \\
&\leq 3\epsilon \|R^{-1}\| \\
&\leq 6\epsilon,
\end{aligned}$$

where the first step follows by definition of α, β , the second step follows by triangle inequality, the third step follows by $\|AB\| \leq \|A\| \|B\|$, the fourth step follows by $\|(I - \bar{V}_2 \bar{V}_2^\top)\| = 1$, the fifth step follows by Eq. (32), the sixth step follows by Fact B.4, the seventh step follows by $\|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \leq \epsilon$ and $\|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \leq \epsilon$, and the last step follows by $\|R^{-1}\| \leq 2$ (Claim E.7).

Finally,

$$\begin{aligned}
\|UU^\top - VV^\top\| &\leq \|UU^\top - \bar{V}\bar{V}^\top\| + \|\bar{V}\bar{V}^\top - VV^\top\| \\
&= \|(I - UU^\top)\bar{V}\| + \|\bar{V}\bar{V}^\top - VV^\top\| \\
&\leq \|\hat{P}_2 - P_2\| / \sigma_k(P_2) + \|\bar{V}\bar{V}^\top - VV^\top\| \\
&\leq \|\hat{P}_2 - P_2\| / \sigma_k(P_2) + \|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| + \|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \\
&\leq \|\hat{P}_2 - P_2\| / \sigma_k(P_2) + 2\epsilon,
\end{aligned}$$

where the first step follows by triangle inequality, the second step follows by Fact B.4, the third step follows by Eq. (33), the fourth step follows by triangle inequality, and the last step follows by $\|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \leq \epsilon$ and $\|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \leq \epsilon$.

Therefore we finish the proof. \square

It remains to prove Claim E.7.

Claim E.7. $\sigma_k(R) \geq 1/2$.

Proof. First, we can rewrite $R^\top R$ in the follow way,

$$R^\top R = V_2^\top (I - V_1 V_1^\top) V_2 = I - V_2^\top V_1 V_1^\top V_2$$

Second, we can upper bound $\|V_2^\top V_1\|$ by $1/4$,

$$\begin{aligned}
\|V_2^\top V_1\| &= \|V_2 V_2^\top V_1\| \\
&\leq \|(V_2 V_2^\top - \bar{V}_2 \bar{V}_2^\top) V_1\| + \|\bar{V}_2 \bar{V}_2^\top V_1\| \\
&\leq \|(V_2 V_2^\top - \bar{V}_2 \bar{V}_2^\top) V_1\| + \|\bar{V}_2^\top (V_1 V_1^\top - \bar{V}_1 \bar{V}_1^\top)\| + \|\bar{V}_2^\top \bar{V}_1 \bar{V}_1^\top\| \\
&= \|(V_2 V_2^\top - \bar{V}_2 \bar{V}_2^\top) V_1\| + \|\bar{V}_2^\top (V_1 V_1^\top - \bar{V}_1 \bar{V}_1^\top)\| \\
&\leq \|V_2 V_2^\top - \bar{V}_2 \bar{V}_2^\top\| \cdot \|V_1\| + \|\bar{V}_2^\top\| \cdot \|V_1 V_1^\top - \bar{V}_1 \bar{V}_1^\top\| \\
&\leq \epsilon + \epsilon \\
&\leq 1/4,
\end{aligned}$$

where the first step follows by $V_2^\top V_2 = I$, the second step follows by triangle inequality, the third step follows by triangle inequality, the fourth step follows by $\|\bar{V}_2^\top \bar{V}_1 \bar{V}_1^\top\| = 0$, the fifth step follows by $\|AB\| \leq \|A\| \cdot \|B\|$, and the last step follows by $\|\bar{V}_1 \bar{V}_1^\top - V_1 V_1^\top\| \leq \epsilon$, $\|V_1\| = 1$, $\|\bar{V}_2 \bar{V}_2^\top - V_2 V_2^\top\| \leq \epsilon$ and $\|\bar{V}_2^\top\| = 1$, and the last step follows by $\epsilon < 1/8$.

Thus, we can lower bound $\sigma_k^2(R)$,

$$\begin{aligned}
\sigma_k^2(R) &= \lambda_{\min}(R^\top R) \\
&= \min_{\|a\|=1} a^\top R^\top R a \\
&= \min_{\|a\|=1} a^\top I a - \|V_2^\top V_1 a\|^2 \\
&= 1 - \max_{\|a\|=1} \|V_2^\top V_1 a\|^2 \\
&= 1 - \|V_2^\top V_1\|^2 \\
&\geq 3/4
\end{aligned}$$

which implies $\sigma_k(R) \geq 1/2$. □

E.4 Error Bound for the Reduced Third-order Moment

E.4.1 Error Bound for the Reduced Third-order Moment in Different Cases

Lemma E.8. *Let M_3 be defined as in Definition 5.1. Let \widehat{M}_3 be the empirical version of M_3 , i.e.,*

$$\widehat{M}_3 = \frac{1}{|S|} \sum_{(x,y) \in S} y \cdot (x^{\otimes 3} - x \otimes I),$$

where S denote a set of samples (where each sample is i.i.d. sampled from Distribution \mathcal{D} defined in Eq. (1)). Assume $M_3 \neq 0$, i.e., $m_{3,i} \neq 0$ for any i . Let $V \in \mathbb{R}^{d \times k}$ be an orthogonal matrix satisfying $\|UU^\top - VV^\top\| \leq 1/4$, where U is the orthogonal basis of $\text{span}\{w_1^*, w_2^*, \dots, w_k^*\}$. Then for any $0 < \epsilon < 1, t \geq 1$, if

$$|S| \geq \max_{i \in [k]} (\|w_i^*\|^{p+1} / |m_{3,i}|^2 + 1) \cdot \epsilon^{-2} \cdot k^2 \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^t$,

$$\|M_3(V, V, V) - \widehat{M}_3(V, V, V)\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{3,i}|.$$

Proof. Since $y = \sum_{i=1}^k v_i^* \phi(w_i^{*\top} x)$. We consider each component $i \in [k]$. We define function $T_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{k \times k \times k}$ such that,

$$T_i(x) = \phi(w_i^{*\top} x) \cdot ((V^\top x)^{\otimes 3} - (V^\top x) \tilde{\otimes} I).$$

We flatten tensor $T_i(x)$ along the first dimension into matrix $B_i(x) \in \mathbb{R}^{k \times k^2}$. Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = |\int_0^z \phi'(s) ds| \leq L_1/(p+1)|z|^{p+1}$, which follows Property 3.1.

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned} \|B_i(x)\| &\leq |\phi(w_i^{*\top} x)| \cdot (\|V^\top x\|^3 + 3k\|V^\top x\|) \\ &\lesssim (|w_i^{*\top} x|^{p+1} + |\phi(0)|) \cdot (\|V^\top x\|^3 + 3k\|V^\top x_j\|) \end{aligned}$$

Note that $V^\top x \sim \mathcal{N}(0, I_k)$. According to Fact B.1 and Fact B.2, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^t)$,

$$\|B_i(x)\| \lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) k^{3/2} \text{poly}(\log d, t)$$

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{3,i}(V^\top \bar{w}_i^*) \text{vec}((V^\top \bar{w}_i^*)(V^\top \bar{w}_i^*)^\top)^\top$. Therefore, $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| = |m_{3,i}| \|V^\top \bar{w}_i^*\|^3$. Since $\|VV^\top - UU^\top\| \leq 1/4$, $\|VV^\top \bar{w}_i^* - \bar{w}_i^*\| \leq 1/4$ and $3/4 \leq \|V^\top \bar{w}_i^*\| \leq 5/4$. So $\frac{1}{4}|m_{3,i}| \leq \|B\| \leq 2|m_{3,i}|$.

(III) Bounding $\max(\mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x)B_i(x)^\top\|, \mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x)^\top B_i(x)\|)$.

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)B_i(x)^\top] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\|V^\top x\|^6] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 k^{3/2}. \end{aligned}$$

$$\begin{aligned} &\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)^\top B_i(x)] \right\| \\ &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\|V^\top x\|^4] \right)^{1/2} \max_{\|A\|_F=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\langle A, (V^\top x)(V^\top x)^\top \rangle^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 k^2. \end{aligned}$$

(IV) Bounding $\max_{\|a\|=\|b\|=1} (\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x)b)^2])$.

$$\begin{aligned} &\max_{\|a\|=\|b\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x)b)^2] \right)^{1/2} \\ &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(\phi(w_i^{*\top} x))^4] \right)^{1/4} \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top V^\top x)^4] \right)^{1/2} \max_{\|A\|_F=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\langle A, (V^\top x)(V^\top x)^\top \rangle^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) k \end{aligned}$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$|S| \gtrsim \frac{L^2 k^2 + |m_{3,i}|^2 + k^{3/2} \text{poly}(\log d, t) |m_{3,i}| \epsilon}{\epsilon^2 |m_{3,i}|^2} t \log(k)$$

with probability at least $1 - k^{-t}$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{3,i}|.$$

We can set $t = T \log(d) / \log(k)$, then if

$$|S| \geq \epsilon^{-2} (1 + 1/|m_{3,i}|^2) \text{poly}(T, \log d)$$

with probability at least $1 - d^{-T}$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{3,i}|.$$

Also note that for any symmetric 3rd-order tensor R , the operator norm of R ,

$$\|R\| = \max_{\|a\|=1} |R(a, a, a)| \leq \max_{\|a\|=1} \|R(a, I, I)\|_F = \|R^{(1)}\|.$$

□

Lemma E.9. *Let M_4 be defined as in Definition 5.1. Let \widehat{M}_4 be the empirical version of M_4 , i.e.,*

$$\widehat{M}_4 = \frac{1}{|S|} \sum_{(x,y) \in S} y \cdot (x^{\otimes 4} - (x \otimes x) \widetilde{\otimes} I + I \widetilde{\otimes} I),$$

where S is a set of samples (where each sample is i.i.d. sampled from Distribution \mathcal{D} defined in Eq. (1)). Assume $M_4 \neq 0$, i.e., $m_{4,i} \neq 0$ for any i . Let α be a fixed unit vector. Let $V \in \mathbb{R}^{d \times k}$ be an orthogonal matrix satisfying $\|UU^\top - VV^\top\| \leq 1/4$, where U is the orthogonal basis of $\text{span}\{w_1^*, w_2^*, \dots, w_k^*\}$. Then for any $0 < \epsilon < 1, t \geq 1$, if

$$|S| \geq \max_{i \in [k]} (1 + \|w_i^*\|^{p+1} / |m_{4,i}(\alpha^\top \bar{w}_i^*)|^2) \cdot \epsilon^{-2} \cdot k^2 \text{poly}(\log d, t)$$

with probability at least $1 - d^{-t}$,

$$\|M_4(V, V, V, \alpha) - \widehat{M}_4(V, V, V, \alpha)\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{4,i}(\alpha^\top \bar{w}_i^*)|.$$

Proof. Recall that for each $(x, y) \in S$, we have $y = \sum_{i=1}^k v_i^* \phi(w_i^{*\top} x)$. We consider each component $i \in [k]$. We define function $r(x) : \mathbb{R}^d \rightarrow \mathbb{R}^k$ such that

$$r(x) = V^\top x.$$

Define function $T_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{k \times k \times k}$ such that

$$T_i(x) = \phi(w_i^{*\top} x) \left(x^\top \alpha \cdot r(x) \otimes r(x) \otimes r(x) - (V^\top \alpha) \widetilde{\otimes} (r(x) \otimes r(x)) - \alpha^\top x \cdot r(x) \widetilde{\otimes} I + (V^\top \alpha) \widetilde{\otimes} I \right).$$

We flatten $T_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{k \times k \times k}$ along the first dimension to obtain function $B_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^{k \times k^2}$. Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = \left| \int_0^z \phi'(s) ds \right| \leq L_1 / (p+1) |z|^{p+1}$, which follows Property 3.1.

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned} \|B_i(x)\| &\lesssim (|w_i^{*\top} x|^{p+1} + |\phi(0)|) \cdot (|(x^\top \alpha)| \|V^\top x\|^3 + 3\|V^\top \alpha\| \|V^\top x\|^2 \\ &\quad + 3|(x^\top \alpha)| \|V^\top x_j\| \sqrt{k} + 3\|V^\top \alpha\| \sqrt{k}) \end{aligned}$$

Note that $V^\top x \sim \mathcal{N}(0, I_k)$. According to Fact B.1 and Fact B.2, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^t)$,

$$\|B_i(x)\| \lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) k^{3/2} \text{poly}(\log d, t)$$

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{4,i}(\alpha^\top \bar{w}_i^*)(V^\top \bar{w}_i^*) \text{vec}((V^\top \bar{w}_i^*)(V^\top \bar{w}_i^*)^\top)^\top$. Therefore,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] \right\| = |m_{4,i}(\alpha^\top \bar{w}_i^*)| \|V^\top \bar{w}_i^*\|^3.$$

Since $\|VV^\top - UU^\top\| \leq 1/4$, $\|VV^\top \bar{w}_i^* - \bar{w}_i^*\| \leq 1/4$ and $3/4 \leq \|V^\top \bar{w}_i^*\| \leq 5/4$. So $\frac{1}{4}|m_{4,i}(\alpha^\top \bar{w}_i^*)| \leq \|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| \leq 2|m_{4,i}(\alpha^\top \bar{w}_i^*)|$.

(III) Bounding $\max(\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)B_i(x)^\top], \mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)^\top B_i(x)])$.

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)B_i(x)^\top] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(\alpha^\top x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\|V^\top x\|^6] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 k^{3/2}. \end{aligned}$$

$$\begin{aligned} &\left\| \mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)^\top B_i(x)] \right\| \\ &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(\alpha^\top x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\|V^\top x\|^4] \right)^{1/2} \\ &\quad \cdot \left(\max_{\|A\|_F=1} \mathbb{E}_{x \sim \mathcal{D}_d}[\langle A, (V^\top x)(V^\top x)^\top \rangle^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 k^2. \end{aligned}$$

(IV) Bounding $\max_{\|a\|=\|b\|=1} (\mathbb{E}_{x \sim \mathcal{D}_d}[(a^\top B_i(x)b)^2])^{1/2}$.

$$\begin{aligned} &\max_{\|a\|=\|b\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(a^\top B_i(x)b)^2] \right)^{1/2} \\ &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\phi^4(w_i^{*\top} x)] \right)^{1/4} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(\alpha^\top x)^4] \right)^{1/4} \max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[(a^\top V^\top x)^4] \right)^{1/2} \\ &\quad \cdot \max_{\|A\|_F=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d}[\langle A, (V^\top x)(V^\top x)^\top \rangle^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) k. \end{aligned}$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$|S| \geq \frac{L^2 k^2 + |m_{4,i}(\alpha^\top \bar{w}_i^*)|^2 + k^{3/2} \text{poly}(t, \log d) |m_{4,i}(\alpha^\top \bar{w}_i^*)| \epsilon}{\epsilon^2 (m_{4,i}(\alpha^\top \bar{w}_i^*))^2} \cdot t \log k$$

with probability at least $1 - k^{-t}$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \in S} B_i(x) \right\| \leq \epsilon |m_{4,i}(\alpha^\top \bar{w}_i^*)|. \quad (34)$$

We can set $t = T \log(d)/\log(k)$, then if

$$|S| \geq \frac{(L + |m_{4,i}(\alpha^\top \bar{w}_i^*)|)^2 k^2 \text{poly}(T, \log d)}{\epsilon^2 |m_{4,i}(\alpha^\top \bar{w}_i^*)|^2} \cdot T \log^2 d$$

with probability at least $1 - d^{-T}$, Eq. (34) holds. Also note that for any symmetric 3rd-order tensor R , the operator norm of R ,

$$\|R\| = \max_{\|a\|=1} |R(a, a, a)| \leq \max_{\|a\|=1} \|R(a, I, I)\|_F = \|R^{(1)}\|.$$

□

E.4.2 Final Error Bound for the Reduced Third-order Moment

Lemma E.10 shows \widehat{R}_3 can approximate R_3 to some small precision with $\text{poly}(k)$ samples.

Lemma E.10 (Estimation of the reduced third order moment). *Let $U \in \mathbb{R}^{d \times k}$ denote the orthogonal column span of W^* . Let α be a fixed unit vector and $V \in \mathbb{R}^{d \times k}$ denote an orthogonal matrix satisfying $\|VV^\top - UU^\top\| \leq 1/4$. Define $R_3 := P_3(V, V, V)$, where P_3 is defined as in Definition 5.4 using α . Let \widehat{R}_3 be the empirical version of R_3 using dataset S , where each sample of S is i.i.d. sampled from distribution \mathcal{D} (defined in (1)). Assume the activation function satisfies Property 3.1 and Assumption 5.3. Then for any $0 < \epsilon < 1$ and $t \geq 1$, define $j_3 = \min\{j \geq 3 | M_j \neq 0\}$ and $m_0 = \min_{i \in [k]} \{(m_i^{(j_3)}(\alpha^\top \bar{w}_i^*)^{j_3-3})^2\}$, if*

$$|S| \geq \sigma_1^{2p+2} \cdot k^2 \cdot \text{poly}(\log d, t) / (\epsilon^2 m_0)$$

then we have ,

$$\|R_3 - \widehat{R}_3\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{j_3,i}(\bar{w}_i^{*\top} \alpha)^{j_3-3}|,$$

holds with probability at least $1 - d^{-\Omega(t)}$.

Proof. The main idea is to use matrix Bernstein bound after matricizing the third-order tensor. Similar to the proof of Lemma E.5, we consider each node component individually and then sum up the errors and apply union bound.

We have shown the bound for $j_3 = 3, 4$ in Lemma E.8 and Lemma E.9 respectively. To summarize, for any $0 < \epsilon < 1$ we have if

$$|S| \geq \max_{i \in [k]} \left(1 + \|w_i^*\|^{p+1} / |m_{j_3,i}(\bar{w}_i^{*\top} \alpha)^{(j_3-3)}|^2 \right) \cdot \epsilon^{-2} \cdot k^2 \text{poly}(\log d, t)$$

with probability at least $1 - d^{-t}$,

$$\|R_3 - \widehat{R}_3\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{j_3,i}(\bar{w}_i^{*\top} \alpha)^{j_3-3}|.$$

□

E.5 Error Bound for the Magnitude and Sign of the Weight Vectors

The lemmata in this section together with Lemma E.5 provide guarantees for Algorithm 4. In particular, Lemma E.12 shows with linear sample complexity in d , we can approximate the 1st-order moment to some precision. And Lemma E.13 and Lemma E.14 provide the error bounds of linear systems Eq. (28) under some perturbations.

E.5.1 Robustness for Solving Linear Systems

Lemma E.11 (Robustness of linear system). *Given two matrices $A, \tilde{A} \in \mathbb{R}^{d \times k}$, and two vectors $b, \tilde{b} \in \mathbb{R}^d$. Let $z^* = \operatorname{argmin}_{z \in \mathbb{R}^k} \|Az - b\|$ and $\hat{z} = \operatorname{argmin}_{z \in \mathbb{R}^k} \|(A + \tilde{A})z - (b + \tilde{b})\|$. If $\|\tilde{A}\| \leq \frac{1}{4\kappa} \sigma_k(A)$ and $\|\tilde{b}\| \leq \frac{1}{4}\|b\|$, then, we have*

$$\|z^* - \hat{z}\| \lesssim (\sigma_k^{-4}(A)\sigma_1^2(A) + \sigma_k^{-2}(A))\|b\|\|\tilde{A}\| + \sigma_k^{-2}(A)\sigma_1(A)\|\tilde{b}\|.$$

Proof. By definition of z and \hat{z} , we can rewrite z and \hat{z} ,

$$\begin{aligned} z &= A^\dagger b = (A^\top A)^{-1} A^\top b \\ \hat{z} &= (A + \tilde{A})^\dagger (b + \tilde{b}) = ((A + \tilde{A})^\top (A + \tilde{A}))^{-1} (A + \tilde{A})^\top (b + \tilde{b}). \end{aligned}$$

As $\|\tilde{A}\| \leq \frac{1}{4\kappa} \sigma_k(A)$, we have $\|\tilde{A}^\top A + A^\top \tilde{A}\| \|(A^\top A)^{-1}\| \leq 1/4$. Together with $\|\tilde{b}\| \leq \frac{1}{4}\|b\|$, we can ignore the high-order errors. So we have

$$\begin{aligned} &\|\hat{z} - z^*\| \\ &\lesssim \|(A^\top A)^{-1}(\tilde{A}^\top b + A^\top \tilde{b}) + (A^\top A)^{-1}(A^\top \tilde{A} + \tilde{A}^\top A)(A^\top A)^{-1} A^\top b\| \\ &\lesssim \|(A^\top A)^{-1}(\|\tilde{A}\|\|b\| + \|A\|\|\tilde{b}\|) + \|(A^\top A)^{-2}\| \cdot \|A\|\|\tilde{A}\|\|A\|\|b\| \\ &\lesssim \sigma_k^{-2}(A)(\|\tilde{A}\|\|b\| + \sigma_1(A)\|\tilde{b}\|) + \sigma_k^{-4}(A) \cdot \sigma_1^2(A)\|\tilde{A}\|\|b\|. \end{aligned}$$

□

E.5.2 Error Bound for the First-order Moment

Lemma E.12 (Error bound for the first-order moment). *Let Q_1 be defined as in Eq. (25) and \hat{Q}_1 be the empirical version of Q_1 using dataset S , where each sample of S is i.i.d. sampled from distribution \mathcal{D} (defined in (1)). Assume the activation function satisfies Property 3.1 and Assumption 5.3. Then for any $0 < \epsilon < 1$ and $t \geq 1$, define $j_1 = \min\{j \geq 1 | M_j \neq 0\}$ and $m_0 = \min_{i \in [k]} (m_{j_1, i} (\bar{w}_i^* \alpha)^{j_1 - 1})^2$ if*

$$|S| \geq \sigma_1^{2p+2} d \operatorname{poly}(t, \log d) / (\epsilon^2 m_0)$$

we have with probability at least $1 - d^{-\Omega(t)}$,

$$\|Q_1 - \hat{Q}_1\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{j_1, i} (\bar{w}_i^* \alpha)^{j_1 - 1}|.$$

Proof. We consider the case when $l_1 = 3$, i.e.,

$$Q_1 = M_3(I, \alpha, \alpha) = \sum_{i=1}^k v_i^* m_{3, i} (\alpha^\top \bar{w}_i^*)^3 \bar{w}_i^*.$$

And other cases are similar.

Since $y = \sum_{i=1}^k v_i^* \phi(w_i^{*\top} x)$. We consider each component $i \in [k]$.

Define function $B_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$B_i(x) = [\phi(w_i^{*\top} x) \cdot (x^{\otimes 3} - x \tilde{\otimes} I)](I, \alpha, \alpha) = \phi(w_i^{*\top} x) \cdot ((x^\top \alpha)^2 x - 2(x^\top \alpha) \alpha - x).$$

Define $g(z) = \phi(z) - \phi(0)$, then $|g(z)| = |\int_0^z \phi'(s) ds| \leq L_1/(p+1)|z|^{p+1}$, which follows Property 3.1.

(I) Bounding $\|B_i(x)\|$.

$$\begin{aligned} \|B_i(x)\| &\leq |\phi(w_i^{*\top} x)| \cdot \|(x^\top \alpha)^2 x - 2\alpha^\top x \alpha - x\| \\ &\leq (\|w_i^{*\top} x\|^{p+1} + |\phi(0)|) ((x^\top \alpha)^2 + 1) \|x\| + 2|\alpha^\top x| \end{aligned}$$

According to Fact B.1 and Fact B.2, we have for any constant $t \geq 1$, with probability $1 - 1/(nd^t)$,

$$\|B_i(x)\| \lesssim (\|w_i^*\|^{p+1} + |\phi(0)|) \sqrt{d} \text{poly}(\log d, t)$$

(II) Bounding $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\|$.

Note that $\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)] = m_{3,i}(\bar{w}_i^{*\top} \alpha)^2 \bar{w}_i^*$. Therefore, $\|\mathbb{E}_{x \sim \mathcal{D}_d}[B_i(x)]\| = |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|$.

(III) Bounding $\max(\mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x) B_i(x)^\top\|, \mathbb{E}_{x \sim \mathcal{D}_d} \|B_i(x)^\top B_i(x)\|)$.

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)^\top B_i(x)] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(x^\top \alpha)^8] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\|x\|^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2 d. \end{aligned}$$

$$\begin{aligned} \left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x) B_i(x)^\top] \right\| &\lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi(w_i^{*\top} x)^4] \right)^{1/2} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(x^\top \alpha)^8] \right)^{1/2} \left(\max_{\|a\|=1} \mathbb{E}_{x \sim \mathcal{D}_d} [(x^\top a)^4] \right)^{1/2} \\ &\lesssim (\|w_i^*\|^{p+1} + |\phi(0)|)^2. \end{aligned}$$

(IV) Bounding $\max_{\|a\|=\|b\|=1} (\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) b)^2])$.

$$\max_{\|a\|=1} \left(\mathbb{E}_{x \sim \mathcal{D}_d} [(a^\top B_i(x) a)^2] \right)^{1/2} \lesssim \left(\mathbb{E}_{x \sim \mathcal{D}_d} [\phi^4(w_i^{*\top} x)] \right)^{1/4} \lesssim \|w_i^*\|^{p+1} + |\phi(0)|.$$

Define $L = \|w_i^*\|^{p+1} + |\phi(0)|$. Then we have for any $0 < \epsilon < 1$, if

$$|S| \gtrsim \frac{L^2 d + |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|^2 + L |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2| \sqrt{d} \text{poly}(\log d, t) \epsilon}{\epsilon^2 |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|^2} \cdot t \log d$$

with probability at least $1 - 1/d^t$,

$$\left\| \mathbb{E}_{x \sim \mathcal{D}_d} [B_i(x)] - \frac{1}{|S|} \sum_{x \sim S} B_i(x) \right\| \leq \epsilon |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|.$$

Summing up all k components, we obtain if

$$|S| \geq \max_{i \in [k]} \left\{ \frac{(\|w_i^*\|^{p+1} + |\phi(0)| + |m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|)^2}{|m_{3,i}(\bar{w}_i^{*\top} \alpha)^2|^2} \right\} \cdot \epsilon^{-2} d \text{poly}(\log d, t)$$

with probability at least $1 - 1/d^t$,

$$\|M_3(I, \alpha, \alpha) - \widehat{M}_3(I, \alpha, \alpha)\| \leq \epsilon \sum_{i=1}^k |v_i^* m_{3,i} (\bar{w}_i^{*\top} \alpha)^2|.$$

Other cases ($j_1 = 1, 2, 4$) are similar, so we complete the proof. \square

E.5.3 Linear System for the First-order Moment

The following lemma provides estimation error bound for the first linear system in Eq. (28).

Lemma E.13 (Solution of linear system for the first order moment). *Let $U \in \mathbb{R}^{d \times k}$ be the orthogonal column span of W^* . Let $V \in \mathbb{R}^{d \times k}$ denote an orthogonal matrix satisfying that $\|VV^\top - UU^\top\| \leq \widehat{\delta}_2 \lesssim 1/(\kappa^2 \sqrt{k})$. For each $i \in [k]$, let \widehat{u}_i denote the vector satisfying $\|\widehat{u}_i - V^\top \bar{w}_i^*\| \leq \widehat{\delta}_3 \lesssim 1/(\kappa^2 \sqrt{k})$. Let Q_1 be defined as in Eq. (25) and \widehat{Q}_1 be the empirical version of Q_1 such that $\|Q_1 - \widehat{Q}_1\| \leq \widehat{\delta}_4 \|Q_1\| \leq \frac{1}{4} \|Q_1\|$. Let $z^* \in \mathbb{R}^k$ and $\widehat{z} \in \mathbb{R}^k$ be defined as in Eq. (27) and Eq. (28). Then*

$$\|\widehat{z}_i - z_i^*\| \leq (\kappa^4 k^{3/2} (\widehat{\delta}_2 + \widehat{\delta}_3) + \kappa^2 k^{1/2} \widehat{\delta}_4) \|z^*\|_1.$$

Proof. Let $A \in \mathbb{R}^{k \times k}$ denote the matrix where the i -th column is $s_i \bar{w}_i^*$. Let $\widetilde{A} \in \mathbb{R}^{k \times k}$ denote the matrix where the i -th column is $V \widehat{u}_i$. Let $b \in \mathbb{R}^k$ denote the vector Q_1 , let \widetilde{b} denote the vector $\widehat{Q}_1 - Q_1$. Then we have

$$\|A\| \leq \sqrt{k}.$$

Using Fact B.3, we can lower bound $\sigma_k(A)$,

$$\sigma_k(A) \geq 1/\kappa.$$

We can upper bound $\|\widetilde{A}\|$ in the following way,

$$\begin{aligned} \|\widetilde{A}\| &\leq \sqrt{k} \max_{i \in [k]} \{\|V \widehat{u}_i - s_i \bar{w}_i^*\|\} \\ &\leq \sqrt{k} \max_{i \in [k]} \{\|V \widehat{u}_i - s_i V V^\top \bar{w}_i^* + s_i V V^\top \bar{w}_i^* - s_i U U^\top \bar{w}_i^*\|\} \\ &\leq \sqrt{k} (\widehat{\delta}_3 + \widehat{\delta}_2). \end{aligned}$$

We can upper bound $\|b\|$ and $\|\widetilde{b}\|$,

$$\|b\| = \|Q_1\| \leq k \sum_{i=1}^k |z_i^*|, \text{ and } \|\widetilde{b}\| \leq \widehat{\delta}_4 \|Q_1\|.$$

To apply Lemma E.11, we need $\widehat{\delta}_4 \leq 1/4$ and $\widehat{\delta}_2 \lesssim 1/(\sqrt{k} \kappa^2)$, $\widehat{\delta}_3 \lesssim 1/(\sqrt{k} \kappa^2)$. So we have,

$$\begin{aligned} \|\widehat{z}_i - z_i^*\| &\leq (\kappa^4 k^{3/2} (\widehat{\delta}_2 + \widehat{\delta}_3) + \kappa^2 k^{1/2} \widehat{\delta}_4) \|Q_1\| \\ &\leq (\kappa^4 k^{3/2} (\widehat{\delta}_2 + \widehat{\delta}_3) + \kappa^2 k^{1/2} \widehat{\delta}_4) \sum_{i=1}^k |z_i^*|. \end{aligned}$$

\square

E.5.4 Linear System for the Second-order Moment

The following lemma provides estimation error bound for the second linear system in Eq. (28).

Lemma E.14 (Solution of linear system for the second order moment). *Let $U \in \mathbb{R}^{d \times k}$ be the orthogonal column span of W^* denote an orthogonal matrix satisfying that $\|VV^\top - UU^\top\| \leq \widehat{\delta}_2 \lesssim 1/(\kappa\sqrt{k})$. For each $i \in [k]$, let \widehat{u}_i denote the vector satisfying $\|\widehat{u}_i - V^\top \overline{w}_i^*\| \leq \widehat{\delta}_3 \lesssim 1/(\sqrt{k}\kappa^3)$. Let Q_2 be defined as in Eq. (26) and \widehat{Q}_2 be the estimation of Q_2 such that $\|Q_2 - \widehat{Q}_2\|_F \leq \widehat{\delta}_4 \|Q_2\|_F \leq \frac{1}{4} \|Q_2\|_F$. Let $r^* \in \mathbb{R}^k$ and $\widehat{r} \in \mathbb{R}^k$ be defined as in Eq. (27) and Eq. (28). Then*

$$|\widehat{r}_i - r_i^*| \leq (k^3 \kappa^8 \widehat{\delta}_3 + \kappa^2 k^2 \widehat{\delta}_4) \|r^*\|.$$

Proof. For each $i \in [k]$, let $u_i = V^\top \overline{w}_i^*$. Let $A \in \mathbb{R}^{k^2 \times k}$ denote the matrix where the i -th column is $\text{vec}(u_i u_i^\top)$. Let $\widetilde{A} \in \mathbb{R}^{k^2 \times k}$ denote the matrix where the i -th column is $\text{vec}(u_i u_i^\top - \widehat{u}_i \widehat{u}_i^\top)$. Let $b \in \mathbb{R}^{k^2}$ denote the vector $\text{vec}(Q_2)$, let $\widetilde{b} \in \mathbb{R}^{k^2}$ denote the vector $\text{vec}(Q_2 - \widehat{Q}_2)$.

Let \circ be the element-wise matrix product (a.k.a. Hadamard product), $\overline{W} = [\overline{w}_1^* \ \overline{w}_2^* \ \cdots \ \overline{w}_k^*]$ and $U = [u_1 \ u_2 \ \cdots \ u_k] = V^\top \overline{W}$. We can upper bound $\|A\|$ and $\|\widetilde{A}\|$ as follows,

$$\begin{aligned} \|A\| &= \max_{\|x\|=1} \left\| \sum_{i=1}^k x_i \text{vec}(u_i u_i^\top) \right\| \\ &= \max_{\|x\|=1} \|U \text{diag}(x) U^\top\|_F \\ &\leq \|U\|^2 \\ &\leq \sigma_1^2(V^\top \overline{W}), \end{aligned}$$

and

$$\begin{aligned} \|\widetilde{A}\| &= \sqrt{k} \max_{i \in [k]} \|\widetilde{A}_i\| \\ &\leq \sqrt{k} \max_{i \in [k]} \|u_i u_i^\top - \widehat{u}_i \widehat{u}_i^\top\|_F \\ &\leq \sqrt{k} \max_{i \in [k]} 2 \|u_i - \widehat{u}_i\|_2 \\ &\leq 2\sqrt{k} \widehat{\delta}_3. \end{aligned}$$

We can lower bound $\sigma_k(A)$,

$$\begin{aligned} \sigma_k(A) &= \sqrt{\sigma_k(A^\top A)} \\ &= \sqrt{\sigma_k((U^\top U) \circ (U^\top U))} \\ &= \min_{\|x\|=1} \sqrt{x^\top ((U^\top U) \circ (U^\top U)) x} \\ &= \min_{\|x\|=1} \|(U^\top U)^{1/2} \text{diag}(x) (U^\top U)^{1/2}\|_F \\ &\geq \sigma_k^2(V^\top \overline{W}) \end{aligned}$$

where fourth step follows Schur product theorem, the last step follows by the fact that $\|CB\|_F \geq \sigma_{\min}(C) \|B\|_F$ and \circ is the element-wise multiplication of two matrices.

We can upper bound $\|b\|$ and $\|\tilde{b}\|$,

$$\|b\| \leq \|Q_2\|_F \leq \|r^*\|,$$

$$\|\tilde{b}\| = \|Q_2 - \hat{Q}_2\|_F \leq \hat{\delta}_4 \|r^*\|.$$

Since $\|VV^\top \bar{W} - \bar{W}\| \leq \sqrt{k} \hat{\delta}_2$, we have for any $x \in \mathbb{R}^k$,

$$\begin{aligned} \|VV^\top \bar{W}x\| &\geq \|\bar{W}x\| - \|(VV^\top \bar{W} - \bar{W})x\| \\ &\geq \sigma_k(\bar{W})\|x\| - \hat{\delta}_2 \sqrt{k}\|x\| \end{aligned}$$

Note that according to Fact B.3, $\sigma_k(\bar{W}) \geq 1/\kappa$. Therefore, if $\hat{\delta}_2 \leq 1/(2\kappa\sqrt{k})$, we will have $\sigma_k(V^\top \bar{W}) \geq 1/(2\kappa)$. Similarly, we have $\sigma_1(V^\top \bar{W}) \leq \|V\| \|\bar{W}\| \leq \sqrt{k}$. Then applying Lemma E.11 and setting $\hat{\delta}_2 \lesssim \frac{1}{\sqrt{k}\kappa^3}$, we complete the proof. \square

F Acknowledgments

The authors would like to thank Surbhi Goel, Adam Klivans, Qi Lei, Eric Price, David P. Woodruff, Peilin Zhong, Hongyang Zhang and Jiong Zhang for useful discussions.