

NEURAL NETWORKS LEARN GENERIC MULTI-INDEX MODELS NEAR INFORMATION-THEORETIC LIMIT

Bohan Zhang*
Peking University
2200010903@stu.pku.edu.cn

Zihao Wang*
Stanford University
zihao@stanford.edu

Hengyu Fu
UC Berkeley
hengyuf@berkeley.edu

Jason D. Lee
UC Berkeley
jasondlee@berkeley.edu

ABSTRACT

In deep learning, a central issue is to understand how neural networks efficiently learn high-dimensional features. To this end, we explore the gradient descent learning of a general Gaussian Multi-index model $f(\mathbf{x}) = g(\mathbf{U}\mathbf{x})$ with hidden subspace $\mathbf{U} \in \mathbb{R}^{r \times d}$, which is the canonical setup to study representation learning. We prove that under generic non-degenerate assumptions on the link function, a standard two-layer neural network trained via layer-wise gradient descent can agnostically learn the target with $o_d(1)$ test error using $\tilde{\mathcal{O}}(d)$ samples and $\tilde{\mathcal{O}}(d^2)$ time. The sample and time complexity both align with the information-theoretic limit up to leading order and are therefore optimal. During the first stage of gradient descent learning, the proof proceeds via showing that the inner weights can perform a power-iteration process. This process implicitly mimics a spectral start for the whole span of the hidden subspace and eventually eliminates finite-sample noise and recovers this span. It surprisingly indicates that optimal results can only be achieved if the first layer is trained for more than $\mathcal{O}(1)$ steps. This work demonstrates the ability of neural networks to effectively learn hierarchical functions with respect to both sample and time efficiency.

1 INTRODUCTION

Modern large language models powered by deep neural networks Krizhevsky et al. (2012); He et al. (2016); Vaswani et al. (2023) have shown strong empirical success across many domains. It is believed that much of the effectiveness of the pretraining–finetuning paradigm comes from their ability to extract useful representations from data. In practice, neural networks often first capture important hidden features, which helps them to learn complex tasks more efficiently. This adaptability offers a clear performance advantage over conventional fixed feature methods, such as kernel techniques Wei et al. (2020); Allen-Zhu & Li (2020); Bai & Lee (2020).

Despite these empirical successes, the theoretical understanding of why neural networks can efficiently and adaptively learn important features from data remains largely incomplete. Significant effort has been devoted in this direction, with a series of works Damian et al. (2022); Ba et al. (2022); Berthier et al. (2024); Lee et al. (2024); Montanari & Urbani (2025); Montanari & Wang (2026) investigating how neural networks can uncover hidden features and subsequently learn single-index and multi-index models, where the labels depend on the d -dimensional data inputs only through their projection onto a one-dimensional or low-dimensional subspace, respectively. The study of learning single-index models via neural networks is relatively well developed. Notably, Lee et al. (2024) show that neural networks can successfully learn single-index models with $\tilde{\mathcal{O}}(d)$ sample complexity, which is optimal up to leading order, while Montanari & Urbani (2025) provides a sharp DMFT theory analyzing the learning dynamics of single-index models using a two-layer neural network when the sample size and input dimension are proportional.

*Equal contribution. The order was determined by a coin flip.

However, a single-index model involves only one hidden feature, whereas in practice we often expect multiple distinct features, motivating the study of neural networks learning of multi-index models. On the contrary, positive results on neural networks learning of multi-index models remain rather limited. Note first that the information-theoretic limit for learning multi-index models is still $\Theta(d)$ samples, and there exist polynomial time algorithms Chen & Meka (2020); Troiani et al. (2025); Defilippis et al. (2025); Kovačević et al. (2025); Damian et al. (2025) that achieve this for generic generative-exponent-two multi-index models which covers almost all canonical examples. Nevertheless, in general, it is not clear whether neural networks can achieve this algorithmic limit and, if so, how. The earliest work on neural networks learning of multi-index models Damian et al. (2022) showed that standard two-layer neural networks can efficiently learn a large subclass of multi-index models with a sub-optimal $\tilde{\Theta}(d^2)$ sample size. Subsequently, Abbe et al. (2022) Abbe et al. (2023a) studied the problem of neural network learning multi-index models with a staircase structure. However, their results are far from optimal and are restricted to special classes of functions. The closest attempt to the general problem is Montanari & Wang (2026), which shows that a two-layer network trained via gradient descent can weakly recover some of the features in a multi-index model using $\Theta(d)$ samples. Nevertheless, since the features are only weakly recovered and are not guaranteed to cover the entire span, their results do not establish that these features can be directly leveraged to successfully learn the targets in later stages of training. Therefore, the following remains open.

*Can neural networks trained via standard optimizers (e.g., gradient descent) achieve **efficient, near-information-theoretic-optimal** learning of generic multi-index models?*

1.1 MAIN CONTRIBUTIONS

In this work, we work on this and conclude it positively for generative-exponent-two, non-generative-staircase multi-index models. Specifically, we consider learning the target function $f(\mathbf{x}) = g(\mathbf{U}\mathbf{x})$ over the standard d -dimensional Gaussian input distribution, where $\mathbf{U} \in \mathbb{R}^{r \times d}$ is the hidden subspace and $g : \mathbb{R}^r \rightarrow \mathbb{R}$ is a polynomial of constant degree. This corresponds to the standard multi-index model with a polynomial link function, which generalizes the single-index model as the special case $r = 1$.

Our main result, Theorem 1, shows that under generic assumptions, a two-layer neural network trained with layer-wise gradient descent can agnostically learn the target function up to $o_d(1)$ test error using $\tilde{\mathcal{O}}(d)$ samples and $\tilde{\mathcal{O}}(d^2)$ training time¹. These sample and time complexities are information-theoretically optimal up to leading order, thereby demonstrating the adaptivity and effectiveness of neural network learning. Moreover, our analysis applies to generic activation functions and a broad subclass of loss functions, including but not limited to square loss, ℓ^1 loss, the Huber loss, and the pseudo-Huber loss. We emphasize that the assumption on the link function in our work is rather mild and generic, as all polynomials possess a generative leap exponent of two Chen & Meka (2020); Damian et al. (2025) and the loss function can be interpreted as a form of data preprocessing. The idea of data preprocessing has a rich literature in spectral methods and so on, but rarely considered in theoretical analysis of neural networks learning. The only substantial assumption in this work is that the target function has no generative-staircase structure. We acknowledge that this assumption is important and it is not clear whether and how to extend our results to the case with generative-staircase structure within a neural network framework. For further discussion, please see Section 2.1.

Our proof begins by observing that near initialization, the gradient descent dynamics is nearly independent across different neurons, with each neuron evolving similarly to power method iterations applied to the local Hessian. The local Hessian exhibits large eigenvalues corresponding to the signal, namely the hidden subspace, as well as smaller eigenvalues arising from the noise because of the finite sample size. The key insight is that if this power iteration approximation works for too long, all features will align with the directions of the largest eigenvalues, that is, the most dominant directions, while ignoring the others, leading to failure of learning. However, if the number of training steps is too small, those noisy directions cannot be sufficiently eliminated, particularly when the sample size is nearly proportional to the input dimension. Therefore, the optimal strategy is to stop training the first layer at an appropriately chosen intermediate time. The features will still span the entire subspace while the noise is simultaneously eliminated.

¹We use $\tilde{\mathcal{O}}(\cdot)$ to hide those $d^{o_d(1)}$ factors.

We note that we use a standard initialization scheme for the neural network parameters, without using a spectral initialization or a hot start as in Mignacco et al. (2021); Chen & Shen (2025). The neural networks can only learn the right features from and only from gradient descent, trained for a diverging number of steps. This time scale requirement rule out all the approaches based on DMFT or state evolution, such as Agoritsas et al. (2018); Ba et al. (2022); Celentano et al. (2025); Gerbelot et al. (2023); Berthier et al. (2024) which requires either information/generative exponent one or a hot start/spectral initialization, to guarantee the learning timescale can be made constant. We also highlight that the power iteration approximation does not work if the learning rate is too small. It is important to choose a moderately large learning rate. For a more detailed discussion, please see Section 4.

We defer the detailed discussion of related works to Appendix A.

1.2 NOTATIONS

We shall use $X \lesssim Y$ to denote $X \leq CY$ for some absolute positive constant C and $X \gtrsim Y$ is defined analogously. We also use the standard big-O notations: $\Theta(\cdot)$, $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$ to only hide absolute positive constants. In addition, we use $\tilde{\mathcal{O}}$, $\tilde{\Theta}$ and $\tilde{\Omega}$ to hide all the $d^{o_d(1)}$ terms throughout the paper. Let $[k] = \{1, 2, \dots, k\}$ for $k \in \mathbb{N}^+$ and $[a; b] = \{a, \dots, b\}$ for $a, b \in \mathbb{N}^+$. For a vector \mathbf{v} , denote by $\|\mathbf{v}\|_p := (\sum_i |v_i|^p)^{1/p}$ the ℓ^p norm. When $p = 2$, we omit the subscript for simplicity. For a matrix \mathbf{A} , let $\|\mathbf{A}\|_{\text{op}}$ and $\|\mathbf{A}\|_F$ be the spectral norm and Frobenius norm, respectively. We use $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ to denote the maximal and the minimal eigenvalue of a real symmetric matrix.

2 PROBLEM SETUP

Our purpose is to learn the target $f^* : \mathbb{R}^d \rightarrow \mathbb{R}$, where \mathbb{R}^d is the input domain equipped with the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Our target follows a multi-index model $f^*(\mathbf{x}) = g^*(\mathbf{U}\mathbf{x})$ with hidden directions $\mathbf{U} \in \mathbb{R}^{r \times d}$ and the link function $g^* : \mathbb{R}^r \rightarrow \mathbb{R}$. We assume \mathbf{U} has orthogonal unit rows without loss of generality.

Assumption 1. $g^*(\cdot)$ is a polynomial of degree p . Furthermore, $\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)}[g^{*2}(\mathbf{z})] \leq C$ where C is an absolute constant.

We treat r, p as absolute constants throughout the paper. Next, we consider the model as a standard two-layer neural network

$$f_{\Theta}(\mathbf{x}) = \sum_{j=1}^m a_j \sigma(\mathbf{w}_j^T \mathbf{x} + b_j) \quad (1)$$

Here we use $\Theta := \{a_j, b_j, \mathbf{w}_j\}_{j=1}^m$ to denote all the parameters. We make the following assumption on the activation function.

Assumption 2. $\sigma(\cdot) \in C^3(\mathbb{R})$ is a smooth activation function satisfying $\sigma'(0) = 0$, $\sigma''(0) = 1$ and $|\sigma^{(3)}(\cdot)| \leq M$ where M is an absolute constant.

Let $\mathcal{D} := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ be one dataset, where $\mathbf{x}_i \in \mathbb{R}^d$ are drawn i.i.d. from distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $y_i = f^*(\mathbf{x}_i)$. The empirical and population loss is given respectively by

$$\hat{\mathcal{L}}_{\mathcal{D}}(\Theta) = \frac{1}{n} \sum_{i=1}^n \ell(f_{\Theta}(\mathbf{x}_i), y_i), \quad \mathcal{L}_{\mathcal{D}}(\Theta) = \mathbb{E}[\ell(f_{\Theta}(\mathbf{x}), y)] \quad (2)$$

where $\ell = \ell(t, y) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a general scalar loss function. We make the following assumptions on the loss function.

Assumption 3. $\ell(\cdot, \cdot) \in C^2(\mathbb{R}^2)$, $\ell(t, y)|_{t=y} = 0$ and that it is convex in its first argument. Furthermore, we assume

$$\left| \frac{\partial}{\partial t} \ell(t, y) \right| \leq L \quad \text{and} \quad \left| \frac{\partial^2}{\partial t^2} \ell(t, y) \right| \leq L \quad (3)$$

for all $t, y \in \mathbb{R}$ where L is a positive absolute constant.

We also consider a slightly relaxed version of Assumption 3 for some specific activation functions. This assumption includes square loss as a special case.

Assumption 3’. *We weaken the assumption in Assumption 3 by allowing the first derivative to grow linearly with the input*

$$\left| \frac{\partial}{\partial t} \ell(t, y) \right| \leq L(1 + |t| + |y|) \quad \text{for all } t, y \in \mathbb{R}. \quad (4)$$

Denote $\ell'(0, y_i) := \frac{\partial}{\partial t} \ell(t, y_i)|_{t=0}$ as ℓ_i .

Assumption 4. *We assume $\mathbb{E}[\ell_i] = 0$.*

We define

$$\widehat{\Sigma}_\ell := \frac{1}{n} \sum_{i=1}^n \ell_i \cdot \mathbf{x}_i \mathbf{x}_i^\top \quad (5)$$

and also its population counterpart $\Sigma_\ell := \mathbb{E}[\ell'(0, y) \cdot \mathbf{x} \mathbf{x}^\top]$. Recall that ℓ_i is a function of y_i and that y_i is itself a function of $\mathbf{U} \mathbf{x}_i$. Therefore, ℓ_i is a function of $\mathbf{U} \mathbf{x}_i$ and we conclude that $\text{rank}(\Sigma_\ell) \leq r$ via invoking Stein’s Lemma and performing integration by parts.

Assumption 5. *The rank is full, i.e., $\text{rank}(\Sigma_\ell) = r$.*

Furthermore, denote the non-zero eigenvalues as $\lambda_r \leq \dots \leq \lambda_1$. We will assume $\lambda_r > 0$ and set $\lambda_1 = 1$ without loss of generality.

In our setting, we assume $\lambda_r > 0$ to ensure that Σ_ℓ is not degenerate in its support and provides a meaningful directional signal during optimization. We will denote and track the condition number $\kappa := \frac{\lambda_1}{\lambda_r}$ but will still assume $\kappa = \Theta_d(1)$.

2.1 DISCUSSION OF ASSUMPTIONS

We note that the first four assumptions are quite standard in theory literature. Assumption 1 is fairly generic. In Assumption 2, the key requirement are $\sigma'(0) = 0$, ensuring that the activation behaves locally as a quadratic. The smoothness assumptions can be relaxed with a more refined analysis. For Assumption 3 or 3’, convexity in the first argument of the loss is essential to guarantee that the second layer weights converge to its minimizer during the second stage of training. We note that this condition holds for a large subclass of commonly used losses such as square loss, ℓ^1 loss and Huber loss. The bounds on the derivatives are not critical and can be replaced by a polynomial growth condition under a more careful analysis, since most of our data remain in a bounded regime and the outliers are rare. Finally, Assumption 4 is also standard and could be removed with either a more refined argument or an additional data pre-processing step.

The central assumption in our analysis is Assumption 5. Informally, this assumption requires that the second-order information of the data remain non-degenerate after applying the pre-processing map $y \mapsto \ell'(0, y)$. We admit that for one fixed loss function ℓ , it is straightforward to construct a target task such that the second-order information of $\ell'(0, y)$ becomes degenerate. In practice, however, the target task is typically fixed, and one has the flexibility to choose among different loss functions to learn it. Recent studies on generative exponents and learning thresholds Troiani et al. (2025); Damian et al. (2024); Kovačević et al. (2025); Damian et al. (2025); Defilippis et al. (2025) show that for a large subclass of common tasks, including all polynomial targets with non-generative-staircase structure, there exists a pre-processing function that ensures the second-order information is non-degenerate, and importantly, such a function need not be highly specialized. Consequently, we expect that $\ell'(0, y)$ exhibits non-degenerate second-order information for at least some of those commonly used loss functions in practice, and at least this already includes all targets in Damian et al. (2022). We will expand on this in Section 5.

2.2 TRAINING ALGORITHM

We introduce our layer-wise gradient descent algorithm.

At initialization, we set b_j as zero and impose a symmetric structure on a_j and w_j . Specifically, let the width m be even, and for all $j \in [m]$ we set

$$a_j = -a_{m-j}, \quad w_j^{(0)} = w_{m-j}^{(0)}. \quad (6)$$

For the first half of indices, let a_j be i.i.d. Rademacher variables: $a_j \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}\{-1, +1\}$. For the first half of the features w_j , we initialize them independently and uniformly on the sphere of radius ϵ_0 .

The network is trained via layer-wise gradient descent with sample splitting, a common approach in theoretical analysis such as Wang et al. (2023); Nichani et al. (2025). To start with, we generate two independent datasets \mathcal{D}_i where $i = 1, 2$ and each consisting of n independent samples. The training algorithm proceeds as follows. we first train \mathbf{W} via gradient descent for T_1 steps on the empirical loss $\widehat{\mathcal{L}}_{\mathcal{D}_1}(\Theta)$, then reinitialize \mathbf{b} and train \mathbf{a} via gradient descent for T_2 steps on the empirical loss $\widehat{\mathcal{L}}_{\mathcal{D}_2}(\Theta)$. The pseudocode of this procedure is given below.

Algorithm 1 Layer-wise Gradient Descent

Require: Learning rates η_1, η_2 , weight decay β_1, β_2 , initialization level ϵ_0 , number of steps T_1, T_2 .

- 1: Initialize $\mathbf{a}, \mathbf{b}, \mathbf{W}$.
 - 2: **Train \mathbf{W} on dataset \mathcal{D}_1**
 - 3: **for** $t = 1$ **to** $T_1 - 1$ **do**
 - 4: $\mathbf{W}^{(t)} \leftarrow \mathbf{W}^{(t-1)} - \eta_1 [\nabla_{\mathbf{W}} \widehat{\mathcal{L}}_{\mathcal{D}_1}(\Theta^{(t-1)}) + \beta_1 \mathbf{W}^{(t-1)}]$
 - 5: **end for**
 - 6: $\mathbf{W}^{(T_1)} \leftarrow \mathbf{W}^{(T_1-1)} - \eta_1 \epsilon_0^{-1} [\nabla_{\mathbf{W}} \widehat{\mathcal{L}}_{\mathcal{D}_1}(\Theta^{(T_1-1)}) + \beta_1 \epsilon_0 \mathbf{W}^{(T_1-1)}]$
 - 7: **Re-initialize**
 - 8: $b_i^{(T_1)} \sim_{iid} \text{Unif}([-3, 3]), i \in [m]$
 - 9: $\mathbf{a}^{(T_1)} \leftarrow \mathbf{a}^{(0)}$
 - 10: $\Theta^{(T_1)} \leftarrow (\mathbf{a}^{(T_1)}, \mathbf{W}^{(T_1)}, \mathbf{b}^{(T_1)})$
 - 11: **end**
 - 12: **Train \mathbf{a} on dataset \mathcal{D}_2**
 - 13: **for** $t = T_1 + 1$ **to** $T_1 + T_2$ **do**
 - 14: $\mathbf{a}^{(t)} \leftarrow \mathbf{a}^{(t-1)} - \eta_2 [\nabla_{\mathbf{a}} \widehat{\mathcal{L}}_{\mathcal{D}_2}(\Theta^{(t-1)}) + \beta_2 \mathbf{a}^{(t-1)}]$
 - 15: **end for**
 - 16: Note $T = T_1 + T_2$
 - 17: **return** Prediction function $f_{\Theta^{(T)}}(\cdot) : \mathbf{x} \rightarrow \langle \mathbf{a}^{(T)}, \sigma(\mathbf{W}^{(T_1)} \mathbf{x} + \mathbf{b}^{(T_1)}) \rangle$
-

3 MAIN RESULTS

Our goal is to show that the network defined in (1) trained via Algorithm 1 can efficiently learn multi-index models near the information-theoretic limit.

We use D to denote an arbitrarily large absolute constant and, for technical convenience, assume $n, m \lesssim d^{D/4}$ throughout the paper. Before stating the main results, we further assume below that finite linear combinations of the activation can approximate any one-dimensional monomial to arbitrary precision.

Assumption 6. We assume the activation function has the monomial approximation property with exponent $\beta \geq 0$. That is, for every integer $k \geq 0$ and every $\rho > 0$ there exists a bounded weight function $v_k^{(\rho)} : \{\pm 1\} \times [-3, 3] \rightarrow \mathbb{R}$ with

$$\|v_k^{(\rho)}\|_{\infty} \leq V_{k,\rho} \lesssim_{\sigma,k} \rho^{-\beta} \quad (7)$$

where $\lesssim_{\sigma,k}$ hides absolute constants depending solely on k and activations such that

$$\sup_{|z| \leq 1} \left| \mathbb{E}_{a,b} [v_k^{(\rho)}(a,b) \sigma(az + b)] - z^k \right| \leq \rho \quad (8)$$

where a is a Rademacher RV i.e. $\mathbb{P}(a = +1) = \mathbb{P}(a = -1) = \frac{1}{2}$ and $b \sim \text{Unif}[-3, 3]$ independent from a .

We treat β as an absolute constant throughout. We remark that this assumption is rather mild, which only requires that a two-layer network with one-dimensional input can approximate any monomial to arbitrary precision using an expectation form. Moreover, we conjecture that for most activations one can take $\beta = 0$ and further set $\rho = 0$, as suggested by Barron (1993); E et al. (2021). For concrete examples, Wang et al. (2023) shows that we have $\beta = 0$ for ReLU activations in their Appendix D.6 and we also show in our Appendix E.3 that for the following activation

$$\sigma(t) = \begin{cases} 2|t| - 1 & |t| \geq 1 \\ t^2 & |t| < 1 \end{cases} \quad (9)$$

we provably have $\beta = 0$ and thus can set $\rho = 0$ as well.

We also require a small initialization scale, which is our last assumption.

Assumption 7. We set the initialization level $\epsilon_0 \leq \frac{1}{C_\epsilon m n^{1/2} d^{7/2}} \left(\frac{4}{5}\right)^{T_1}$, where C_ϵ is a universal constant.

We note that ϵ_0 is only required to be polynomially small provided $T_1 = o(\log d)$. With our assumptions in place, we are ready to state our main theorem.

Theorem 1. Under assumptions 1, 2, 3, 4, 5, 6, 7, or under the additional assumption that $\beta = 0$ and assumptions 1, 2, 3', 4, 5, 6, 7, we have the following result: For any first stage training time $T_1 = o(\log d)$ and any sample size

$$n \gtrsim d \log d T_1^2 \kappa^{2T_1} + d^{1+\frac{1}{T_1}} \kappa^2 \quad (10)$$

there exists a proper choice of the hyperparameters and second stage training time T_2 such that with probability at least $1 - \mathcal{O}(d^{-D/2})$, the Algorithm 1 returns

$$\mathcal{L}(\Theta^{(T)}) \lesssim \kappa^{4pT_1} (\log d)^{4p+1} \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right)^{\frac{1}{\beta+1}} \quad (11)$$

In short, Theorem 1 states that we can learn the target up to the vanishing test error with the width $m = \tilde{\Theta}(1)$ and the sample size $n \gtrsim d \log d T_1^2 \kappa^{2T_1} + d^{1+\frac{1}{T_1}} \kappa^2$. We specify those hyperparameters later in the statement of Theorem 2 in Appendix D. We further remark that the sample size requirement $n \gtrsim d \log d T_1^2 \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ guarantees that the first-stage training learns the correct representations, whereas the error term $\kappa^{4pT_1} (\log d)^{4p+1} \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right)^{\frac{1}{\beta+1}}$ is contributed purely by the second-stage training. We note that under the first set of assumptions, the error bound can be refined to $\kappa^{2pT_1} (\log d)^{2p+1} \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right)^{\frac{1}{\beta+1}}$.

Via balancing two terms in our sample complexity requirement

$$n \gtrsim d \log d T_1^2 \kappa^{2T_1} + d^{1+\frac{1}{T_1}} \kappa^2 \quad (12)$$

the optimal dependence shall be derived from setting $T_1 = \left\lceil \sqrt{\frac{\log d}{\log \kappa}} \right\rceil$. Note that below we use $\tilde{\mathcal{O}}$, $\tilde{\Theta}$ and $\tilde{\Omega}$ to hide all the $d^{o_d(1)}$ terms.

Corollary 3.1. Consider $T_1 = \left\lceil \sqrt{\frac{\log d}{\log \kappa}} \right\rceil$. Under this, we further require

$$m = (\log d)^{8p(\beta+1)+1} d^{8p(\beta+1)} \sqrt{\frac{\log \kappa}{\log d}} \quad \text{and} \quad n = C \kappa^2 (\log d)^2 d^{1+2} \sqrt{\frac{\log \kappa}{\log d}} \quad (13)$$

where C is a large universal constant.

Then, there exists a proper choice of hyperparameters that the Algorithm 1 returns $\mathcal{L}(\Theta^{(T)}) = o_d(1)$ with $T_2 = \tilde{\Theta}(1)$ and probability $1 - \mathcal{O}(d^{-D/2})$.

We note that the overall time complexity of Algorithm 1 is $\Theta(nmdT)$. Under the setup of Corollary 3.1, we have $T = \tilde{\Theta}(1)$, $m = \tilde{\Theta}(1)$, and $n = \tilde{\Theta}(d)$, yielding an overall time of $\tilde{\Theta}(d^2)$. This is also optimal up to leading order, since any algorithm must spend at least $\Omega(nd) = \Omega(d^2)$ time to read the full dataset.

Compared with Damian et al. (2022), we note that our results exactly and completely recover the results in Damian et al. (2022) when we set $T_1 = 1$. In this regime, successful learning is guaranteed only with a larger sample size $n = \tilde{\Theta}(d^2)$, and Dandi et al. (2025) further indicates that $n = \tilde{\Theta}(d^2)$ is necessary when $T_1 = 1$. Our results improve upon Damian et al. (2022) and all previous works by managing to execute a careful analysis of the multi-step gradient-descent dynamics on the empirical loss over a diverging time horizon. We provide the details in Section 4.

4 OUTLINE OF THE PROOF

4.1 GRADIENT DESCENT CAN MIMIC POWER ITERATIONS

To begin, we compute the gradient with respect to the feature \mathbf{w}_j as follows.

$$\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}}_n(\Theta^{(t)}) = \frac{1}{n} \sum_{i=1}^n \ell'(f_{\Theta^{(t)}}(\mathbf{x}_i), y_i) a_j \sigma'(\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i \quad (14)$$

Our initialization ensures that the network output is zero at the beginning of the training. Therefore, within a moderate large time horizon, we expect $f_{\Theta^{(t)}}(\mathbf{x})$ to remain near zero. Furthermore, recall that $\|\mathbf{w}_j^{(0)}\| = o_d(1)$. For small $\mathbf{w}_j^{(t)\top} \mathbf{x}_i$ we have

$$\sigma'(\mathbf{w}_j^{(t)\top} \mathbf{x}_i) = \sigma''(0) \mathbf{w}_j^{(t)\top} \mathbf{x}_i + o_d(1) \quad (15)$$

using $\sigma'(0) = 0$ and substituting these yields

$$\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}}_n(\Theta^{(t)}) = \frac{1}{n} \sum_{i=1}^n \ell'(0, y_i) a_j \sigma''(0) \mathbf{w}_j^{(t)\top} \mathbf{x}_i \mathbf{x}_i + \text{smaller order terms.} \quad (16)$$

Letting $\ell_i := \ell'(0, y_i)$, under a proper learning rate and weight decay, the dynamics can be expressed as

$$\mathbf{w}_j^{(t+1)} \propto \left(\frac{1}{n} \sum_{i=1}^n \ell_i \mathbf{x}_i \mathbf{x}_i^\top \right) \mathbf{w}_j^{(t)} = \widehat{\Sigma}_\ell \mathbf{w}_j^{(t)}. \quad (17)$$

We emphasize that the above approximation works whenever $f_{\Theta^{(t)}}(\mathbf{x})$ and $\mathbf{w}_j^{(t)}$ both remain small. Under those approximations, the neurons decouple from each other.

Under these approximations, the dynamics can be viewed as a power method iteration with respect to $\widehat{\Sigma}_\ell$. Recall that $\widehat{\Sigma}_\ell$ is the empirical version of Σ_ℓ , satisfying $\|\widehat{\Sigma}_\ell - \Sigma_\ell\|_{\text{op}} = \Theta(\sqrt{d/n})$. Moreover, from the assumptions, the population one admits the eigendecomposition

$$\Sigma_\ell = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{u}_i^\top, \quad (18)$$

where $\{\mathbf{u}_i\}_{i=1}^r$ are orthonormal unit vectors spanning \mathbf{U} , with eigenvalues ordered as $1 = \lambda_1 \geq \dots \geq \lambda_r = 1/\kappa$, in accordance with our assumptions. We further assume a constant eigenvalue gap only in this section for convenience, namely $\min_{i \neq j} |\lambda_i - \lambda_j| \geq c$ for some universal constant c ; We shall emphasize that we manage to avoid this assumption in our final results via a more fine-grained analysis (see Appendix C).

Eigenvalues and eigenvectors of $\widehat{\Sigma}_\ell^\top$. The empirical covariance $\widehat{\Sigma}_\ell$ then admits the decomposition

$$\widehat{\Sigma}_\ell = \sum_{i=1}^d \widehat{\lambda}_i \widehat{\mathbf{u}}_i \widehat{\mathbf{u}}_i^\top, \quad (19)$$

where $\{\widehat{\mathbf{u}}_i\}_{i=1}^d$ forms an orthonormal unit basis of \mathbb{R}^d . We note that we have $|\widehat{\lambda}_i - \lambda_i| = \Theta(\sqrt{d/n})$ and $\|\widehat{\mathbf{u}}_i - \mathbf{u}_i\| = \Theta(\sqrt{d/n})$ for $i \leq r$. For the others we have $|\widehat{\lambda}_i| = \Theta(\sqrt{d/n})$. This decomposition makes it straightforward to analyze powers of $\widehat{\Sigma}_\ell$

$$\widehat{\Sigma}_\ell^\top = \sum_{i=1}^d \widehat{\lambda}_i^\top \widehat{\mathbf{u}}_i \widehat{\mathbf{u}}_i^\top, \quad (20)$$

where T denotes the number of gradient descent steps in the first stage. For a random initialization $\mathbf{w}_j^{(0)}$, the contribution from the noise components is

$$\left\| \sum_{i=r+1}^d \widehat{\lambda}_i^T \widehat{\mathbf{u}}_i \widehat{\mathbf{u}}_i^T \mathbf{w}_j^{(0)} \right\| = \Theta \left((\sqrt{d/n})^T \left\| \mathbf{w}_j^{(0)} \right\| \right), \quad (21)$$

while the signal contribution along direction $\widehat{\mathbf{u}}_i$ is

$$\widehat{\lambda}_i^T |\widehat{\mathbf{u}}_i^T \mathbf{w}_j^{(0)}| = \Theta \left(\lambda_i^T \frac{1}{\sqrt{d}} \left\| \mathbf{w}_j^{(0)} \right\| \right) \gtrsim \kappa^{-T} \frac{1}{\sqrt{d}} \left\| \mathbf{w}_j^{(0)} \right\|. \quad (22)$$

To ensure that the signal is at least comparable with the noise, we require

$$\frac{\kappa^{-T}}{\sqrt{d}} \gtrsim (\sqrt{d/n})^T \quad (23)$$

which corresponds to the condition $n \gtrsim d^{1+1/T} \kappa^2$.

Assuming the noise term is negligible, for $k \in [r]$ we now focus on the signal strength along direction \mathbf{u}_k instead of $\widehat{\mathbf{u}}_k$

$$\mathbf{u}_k^T \widehat{\Sigma}_\ell^T \mathbf{w}_j^{(0)} \approx \widehat{\lambda}_k^T \mathbf{u}_k^T \widehat{\mathbf{u}}_k \widehat{\mathbf{u}}_k^T \mathbf{w}_j^{(0)} + \sum_{\substack{1 \leq i \leq r \\ i \neq k}} \widehat{\lambda}_i^T \mathbf{u}_k^T \widehat{\mathbf{u}}_i \widehat{\mathbf{u}}_i^T \mathbf{w}_j^{(0)}. \quad (24)$$

Here we require the first term to dominate the second so that the signal strength along the direction \mathbf{u}_k is well aligned with that along $\widehat{\mathbf{u}}_k$. For a typical $\mathbf{w}_j^{(0)}$, the magnitude of the first term is at least $\Theta \left(\kappa^{-T} / \sqrt{d} \left\| \mathbf{w}_j^{(0)} \right\| \right)$, while we have

$$\left| \sum_{\substack{1 \leq i \leq r \\ i \neq k}} \widehat{\lambda}_i^T \mathbf{u}_k^T \widehat{\mathbf{u}}_i \widehat{\mathbf{u}}_i^T \mathbf{w}_j^{(0)} \right| = \mathcal{O} \left(\left\| \mathbf{w}_j^{(0)} \right\| / \sqrt{n} \right) \quad (25)$$

using the standard estimates $|\mathbf{u}_k^T \widehat{\mathbf{u}}_i| = \mathcal{O}(\sqrt{d/n})$ and $|\widehat{\mathbf{u}}_i^T \mathbf{w}_j^{(0)}| = \Theta \left(\left\| \mathbf{w}_j^{(0)} \right\| \sqrt{1/d} \right)$. Therefore, it suffices to take $n \gtrsim d \log d T^2 \kappa^{2T}$. We emphasize and highlight that it is quite interesting to observe that two conflicting factors are involved at the same time. If T is too small, the noise is not sufficiently suppressed; whereas if T is too large, the weakest signal direction vanishes exponentially and becomes undetectable. Therefore, we have to balance these two effects to find the optimal T , which is interestingly found to be of the order $\Theta(\sqrt{\log d})$.

Therefore, assuming $n \gtrsim d \log d T^2 \kappa^{2T} + d^{1+\frac{1}{T}} \kappa^2$, the learned feature approximately lies in and successfully spans the whole \mathbf{U} subspace. We can then conclude that the second stage is able to learn any multi-index model with hidden subspace \mathbf{U} from the learned features via using a standard kernel method argument. The complexity of this learning scales independent with d and polynomial with κ^T , as some compensation is required to learn the weakest direction, which is still $\tilde{\mathcal{O}}(1)$ when $T = \Theta(\sqrt{\log d})$. For a concrete explanation, we refer to the proof roadmap of the second stage in Appendix D.

4.2 COMPARISON WITH THE EXISTING APPROACHES

First of all, we note that our results cannot be recovered using a DMFT or state evolution framework, such as Agoritsas et al. (2018); Ba et al. (2022); Celentano et al. (2025); Gerbelot et al. (2023); Berthier et al. (2024) and many more. Essentially, in order to learn those generative-exponent two directions under the sample size $\tilde{\Theta}(d)$, one needs a diverging time horizon to escape the initial saddle point if there is no hot start or explicit spectral initialization. Under the DMFT/state evolution proportional limit, the neurons shall not escape the initial stationary point.

Second, our methods differ from the standard power iteration in matrix/tensor PCA problems Balzano et al. (2011); Montanari & Richard (2014). If we enforce the matrix power iteration approximation for long time, the neurons shall collapse to one single direction, which prevents a successful

learning. Instead, we force a early stop (but still longer than constant) to the power iteration to ensure the neurons span the entire subspace. In practice, we believe neural networks can adaptively find the optimal stopping time when the second layer is trained together with the first layer. We provide numerical validations of our main results in Section B.

5 FURTHER DISCUSSIONS

We here further explain our Assumption 5, which essentially requires our target function has generative exponent two without generative-staircase structure. To start with, we briefly introduce the concept of generative exponents. The so-called generative exponents has a long standing literature (see related works) and is introduced in its most general form in recent work Damian et al. (2025), which sharply characterize the complexity of efficiently learning Gaussian multi-index models up to leading order. In the following, we use h_k to denote degree k Hermite polynomials and He_k to denote degree k Hermite tensors. A short introduction to these objects is provided in Appendix E.1.

We denote $\mathbf{z} = \mathbf{U}^\top \mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ as the hidden directions and y as the model output. Denote further S as a subspace of the \mathbf{U} subspace, and let $\mathbf{z}_S \in S$ denote the orthogonal projection of \mathbf{z} onto S . The generative exponent given a subspace S , $k_S^* = k_S^*(\mathbf{z}, y)$, is defined as the largest k such that $\mathbb{E}[\mathcal{T}(y, \mathbf{z}_S)p(\mathbf{z}_{S^\perp})] = 0$ for any measurable function \mathcal{T} and any polynomial p of degree lower than k . This indicates that the tensor $\mathbb{E}[\mathcal{T}^*(y, \mathbf{z}_S)\text{He}_{k_S^*}(\mathbf{z}_{S^\perp})] \neq 0$ and we will learn some new directions S_{new} . Therefore, we get a new learned subspace $S' = S \cup S_{\text{new}}$. Repeating this process, we will get an increasing sequence of subspace $\emptyset = S_0 \subsetneq S_1 \subsetneq \dots \subsetneq S_L = \text{span}(\mathbf{U})$, and the overall generative leap exponent is then formally defined as $k^* := \max_i k_{S_i}^*$.

The generative exponent is important since Damian et al. (2025) proves that the complexity of efficiently learning a multi-index model is $\Theta(d^{k^*/2})$ via providing both the LDP lower bound and a matching spectral algorithm upper bound. Further, it can be shown that almost all common multi-index models have the generative leap exponent $k^* \leq 2$. For instance, it works when the link function is 1) all the polynomials, 2) piecewise linear functions, 3) intersection of halfspace. Therefore, we can always extract information from computing $\mathbb{E}[\mathcal{T}(y, \mathbf{z}_{S_i})\text{He}_2(\mathbf{z}_{S_i^\perp})]$ with a proper pre-processing function \mathcal{T} . In addition, Lemma F.14 of Damian et al. (2024) indicates that \mathcal{T} need not to be highly specific. In fact, they prove that even random polynomials suffice in the single-index setting.

We emphasize that this observation motivates our Assumption 5. In addition to generative leap exponent two, we are essentially requiring that there exists one pre-processing function \mathcal{T} which is able to extract all the spanned directions in the subspace \mathbf{U} from the second order information. Our assumption is weaker than the previous literature of using neural networks to learn multi-index models up to vanishing test error. For instance, via using a Huber loss, we can deal with the link function $g(\mathbf{z}) = \sum_{i=1}^k h_4(z_i)$ which is not covered in previous papers like Damian et al. (2022).

6 CONCLUSIONS

In this paper, we demonstrate that under Gaussian input, a two-layer neural network trained with layer-wise gradient descent can learn generic multi-index models up to $o_d(1)$ test error using $\tilde{O}(d)$ sample and $\tilde{O}(d^2)$ time complexity, which matches the information-theoretic threshold up to the leading order. We emphasize that our analysis does not rely on any particular choice of activation, loss, or link function; rather, the assumptions imposed are generic and hold for a broad class of functions. An important and somewhat unexpected insight from the proof is that the inner layer must be trained for more than a constant number of steps to successfully learn the hidden features, in order to recover the hidden features most efficiently and thus achieve optimal performance. This observation is impossible to be observed via a DMFT or state evolution framework over bounded time horizon, and has the potential to extend to other settings. Overall, our results demonstrate that neural networks are capable of effective representation learning and can fully learn the target function on this canonical task with near-optimal efficiency.

ACKNOWLEDGEMENT

The authors thank anonymous reviewers for their valuable feedback and suggestions. Zihao Wang also thanks Alex Damian, Andrea Montanari, Eshaan Nichani, Fan Nie, Yunwei Ren, Rui Sun, Zixuan Wang, Denny Wu, Jiatong Yu for their insightful discussions and valuable feedback during the completion of this work. Jason D. Lee acknowledges support of NSF CCF 2002272, NSF IIS 2107304, NSF CIF 2212262, ONR Young Investigator Award, and NSF CAREER Award 2144994.

REFERENCES

- Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pp. 4782–4887. PMLR, 2022.
- Emmanuel Abbe, Enric Boix Adserà, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 2552–2623. PMLR, 2023a.
- Emmanuel Abbe, Enric Boix-Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics, 2023b.
- Emmanuel Abbe, Enric Boix-Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks, 2024.
- Elisabeth Agoritsas, Giulio Biroli, Pierfrancesco Urbani, and Francesco Zamponi. Out-of-equilibrium dynamical mean-field equations for the perceptron model. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085002, January 2018. ISSN 1751-8121.
- Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels?, 2020.
- Luca Arnaboldi, Yatin Dandi, Florent Krzakala, Luca Pesce, and Ludovic Stephan. Repetita iuvant: Data repetition allows sgd to learn high-dimensional multi-index functions, 2025.
- Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference, 2021.
- Benjamin Aubin, Antoine Maillard, Jean Barbier, Florent Krzakala, Nicolas Macris, and Lenka Zdeborová. The committee machine: computational to statistical gaps in learning a two-layers neural network. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124023, December 2019. ISSN 1742-5468.
- Jimmy Ba, Murat A. Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation, 2022.
- Yu Bai and Jason D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks, 2020.
- Laura Balzano, Robert Nowak, and Benjamin Recht. Online identification and tracking of subspaces from highly incomplete information, 2011.
- Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, March 2019. ISSN 1091-6490.
- Andrew Barron. Barron, a.e.: Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39:930–945. *Information Theory, IEEE Transactions on*, 39:930 – 945, 06 1993.
- William Beckner. Sobolev inequalities, the poisson semigroup, and analysis on the sphere s^n . *Proceedings of the American Mathematical Society*, 118(2):519–528, 1992.

- Raphaël Berthier, Andrea Montanari, and Kangjie Zhou. Learning time-scales in two-layers neural networks. *Foundations of Computational Mathematics*, 25(5):1627–1710, August 2024. ISSN 1615-3383.
- Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- Joan Bruna and Daniel Hsu. Survey on algorithms for multi-index models, 2025.
- Michael Celentano, Chen Cheng, and Andrea Montanari. The high-dimensional asymptotics of first order methods with random data, 2025.
- Sitan Chen and Raghu Meka. Learning polynomials of few relevant dimensions, 2020.
- Yuchen Chen and Yandi Shen. Learning single index model with gradient descent: spectral initialization and precise asymptotics, 2025.
- Alex Damian, Jason D. Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent, 2022.
- Alex Damian, Loucas Pillaud-Vivien, Jason D. Lee, and Joan Bruna. Computational-statistical gaps in gaussian single-index models, 2024.
- Alex Damian, Jason D. Lee, and Joan Bruna. The generative leap: Sharp sample complexity for efficiently learning gaussian multi-index models, 2025.
- Yatin Dandi, Emanuele Troiani, Luca Arnaboldi, Luca Pesce, Lenka Zdeborová, and Florent Krzakala. The benefits of reusing batches for gradient descent in two-layer networks: Breaking the curse of information and leap exponents, 2024.
- Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, and Ludovic Stephan. How two-layer neural networks learn, one (giant) step at a time, 2025.
- Leonardo Defilippis, Yatin Dandi, Pierre Mergny, Florent Krzakala, and Bruno Loureiro. Optimal spectral transitions in high-dimensional multi-index models, 2025.
- Ilias Diakonikolas, Giannis Iakovidis, Daniel M. Kane, and Lisheng Ren. Algorithms and sq lower bounds for robustly learning real-valued multi-index models, 2025.
- Weinan E, Chao Ma, and Lei Wu. The barron space and the flow-induced function spaces for neural network models, 2021.
- Hengyu Fu, Zihao Wang, Eshaan Nichani, and Jason D. Lee. Learning hierarchical polynomials of multiple nonlinear features with three-layer networks, 2024.
- Cedric Gerbelot, Emanuele Troiani, Francesca Mignacco, Florent Krzakala, and Lenka Zdeborova. Rigorous dynamical mean field theory for stochastic gradient descent methods, 2023.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Filip Kovačević, Yihan Zhang, and Marco Mondelli. Spectral estimators for multi-index models: Precise asymptotics and optimal weak recovery, 2025.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.
- Jason D. Lee, Kazusato Oko, Taiji Suzuki, and Denny Wu. Neural network learns low-dimensional polynomials with sgd near the information-theoretic limit, 2024.
- Yue M. Lu and Gen Li. Phase transitions of spectral initialization for high-dimensional nonconvex estimation, 2019.
- Wangyu Luo, Wael Alghamdi, and Yue M. Lu. Optimal spectral initialization for signal recovery with applications to phase retrieval. *IEEE Transactions on Signal Processing*, 67(9):2347–2356, May 2019. ISSN 1941-0476.

- Antoine Maillard, Florent Krzakala, Yue M. Lu, and Lenka Zdeborová. Construction of optimal spectral methods in phase retrieval, 2021.
- Francesca Mignacco, Pierfrancesco Urbani, and Lenka Zdeborová. Stochasticity helps to navigate rough landscapes: comparing gradient-descent-based algorithms in the phase retrieval problem. *Machine Learning: Science and Technology*, 2(3):035029, July 2021. ISSN 2632-2153.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT Press, 2 edition, 2018.
- Marco Mondelli and Andrea Montanari. Fundamental limits of weak recovery with applications to phase retrieval, 2018.
- Andrea Montanari and Emile Richard. A statistical model for tensor pca, 2014.
- Andrea Montanari and Pierfrancesco Urbani. Dynamical decoupling of generalization and overfitting in large two-layer networks, 2025.
- Andrea Montanari and Zihao Wang. Phase transitions for feature learning in neural networks, 2026.
- Eshaan Nichani, Alex Damian, and Jason D Lee. Provable guarantees for nonlinear feature learning in three-layer neural networks. *arXiv preprint arXiv:2305.06986*, 2023.
- Eshaan Nichani, Alex Damian, and Jason D. Lee. Provable guarantees for nonlinear feature learning in three-layer neural networks, 2025.
- Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- Giuseppe Da Prato and Luciano Tubaro. Wick powers in stochastic pdes: an introduction. 2007.
- Yunwei Ren, Mo Zhou, and Rong Ge. Depth separation with multilayer mean-field networks. *arXiv preprint arXiv:2304.01063*, 2023.
- Emanuele Troiani, Yatin Dandi, Leonardo Defilippis, Lenka Zdeborová, Bruno Loureiro, and Florent Krzakala. Fundamental computational limits of weak learnability in high-dimensional multi-index models, 2025.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Zihao Wang, Eshaan Nichani, and Jason D. Lee. Learning hierarchical polynomials with three-layer neural networks, 2023.
- Colin Wei, Jason D. Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets v.s. their induced kernel, 2020.

Appendix

A	Related works	13
B	Numerical illustrations	15
B.1	Results	15
B.2	Proper choice of loss function	15
C	Proof of Lemma C.1: Feature learning	17
C.1	Proof of Lemma C.1	19
C.2	Proof of supporting lemmas in Appendix C.1	21
D	Proof of Theorem 1	26
D.1	Proof roadmap of Theorem 1	26
D.2	Events of high probabilities	27
D.3	Proof of Theorem 1	28
D.4	Proof under $\beta = 0$ and Assumption 3'	36
D.5	Proof of supporting lemmas in Appendix D.3	39
E	Technical background	57
E.1	Hermite polynomials	57
E.2	Uniform generalization bounds	58
E.3	Univariate approximation	59
E.4	Convex optimization	61
E.5	Sphere lemmas	61
E.6	Supporting technical lemmas	62

A RELATED WORKS

Learning multi-index models. In a multi-index model, the output depends on a low-dimensional subspace of the input data. This framework has a rich history in statistics and ML, where it is used to model low-dimensional structures with high-dimensional data. A comprehensive survey of current algorithms for learning multi-index models is provided in Bruna & Hsu (2025). In particular, it is known that under generic conditions, the information-theoretic sample complexity required to learn such targets is proportional to the input dimension, as rigorously established in Barbier et al. (2019); Aubin et al. (2019). Nevertheless, it remains largely unclear whether and how widely used polynomial-time algorithms can efficiently recover the hidden subspace.

To start with, Barbier et al. (2019); Mondelli & Montanari (2018); Lu & Li (2019); Luo et al. (2019); Maillard et al. (2021) studies generic single-index models and sharply derives the learning threshold. In particular, Mondelli & Montanari (2018); Lu & Li (2019); Luo et al. (2019); Maillard et al. (2021) show that for a broad class of single-index models, spectral methods with suitable data pre-processing can recover the hidden direction at the optimal threshold, matching the performance

of an optimal approximate message passing (AMP) Barbier et al. (2019). In contrast, Arous et al. (2021) considers a more restrictive setting and examines the learning of single-index models when using online SGD as the optimizer without data preprocessing, and introduces the concept of an information exponent, which characterizes the learning complexity in their setup.

A substantial line of works studies how to extend these results to multi-index models such as Chen & Meka (2020); Abbe et al. (2022; 2023a); Troiani et al. (2025); Damian et al. (2025); Kovačević et al. (2025); Defilippis et al. (2025); Diakonikolas et al. (2025). On one hand, Abbe et al. (2024; 2023b) introduces the notion of leap complexity, showing that hierarchical/staircase structure in multi-index models can facilitate learning. On the other hand, Damian et al. (2024; 2025) extend this line of work by formalizing the generative exponent for single-index models and the generative leap exponent for multi-index models, extending the previous results to higher order and incorporating the benefits of data pre-processing. In particular, we highlight Damian et al. (2025) for its two key contributions.

- For multi-index models with generative exponent k^* , the authors establish up to leading order a lower bound $\tilde{\Omega}(d^{k^*/2})$ under the low-degree polynomial (LDP) framework and a matching upper bound $\tilde{\mathcal{O}}(d^{k^*/2})$ achieved by a spectral polynomial-time algorithm for the sample complexity.
- Furthermore, it shows that the generative exponent is at most two for almost all common multi-index models, indicating that $\mathcal{O}(d)$ samples are sufficient to efficiently recover the hidden subspace.

These papers are in general not tight in terms of leading order constants. We also note that Diakonikolas et al. (2025) introduces a related notion, i.e., the m -well-behaved multi-index model, which coincides with the generative exponent in many cases. While their sample complexity exhibits a worse dependence on d compared to Damian et al. (2025), their work offers two advantages: it targets agnostic PAC learning rather than subspace recovery, and it provides uniform sample complexity guarantees that apply agnostically to all multi-index models within a certain given class.

Subsequently, for the precise (weak recovery) learning threshold for multi-index models, Troiani et al. (2025) non-rigorously derives the sharp optimal sample size required for weakly learning generative-leap-exponent-two multi-index models by performing a local analysis of Bayes AMP, extending the results of Barbier et al. (2019); Mondelli & Montanari (2018) for single-index models when the generative exponent of the target equals two. Moreover, as two follow-up works, Kovačević et al. (2025); Defilippis et al. (2025) study learning multi-index models using an optimal spectral method, although their sample thresholds in Kovačević et al. (2025) are in general not always optimal, and none of those works can deal with the generative-staircase setup. We emphasize that the above works depend on the special structure of multi-index models and therefore do not use general-purpose algorithms such as learning with standard neural networks trained via gradient descent.

Feature learning in neural networks. The term feature learning refers to the following phenomenon. In practice, neural networks trained with local gradient-based algorithms such as gradient descent can adapt to the low-dimensional structure of the tasks and efficiently learn useful hidden features, which in turn facilitates successful learning and enables effective transfer learning to downstream tasks. Among the first several works in this direction, as we have mentioned previously, Damian et al. (2022) demonstrates that neural networks can learn generic multi-index models at a sub-optimal sample size of $n = \tilde{\Theta}(d^2)$ under some additional assumptions. Similarly, Abbe et al. (2024; 2023b) establish a sub-optimal rate of $n = \tilde{\mathcal{O}}(d^L)$ for neural network learning, where L denotes the leap complexity. Ba et al. (2022) studies the problem under a much more restrictive assumption of an information exponent one, and in this setting shows that neural networks can succeed with $\mathcal{O}(d)$ samples.

Furthermore, as noted earlier, Lee et al. (2024) shows that neural networks can learn generic single-index models with generative exponent at most two with sample complexity $\tilde{\mathcal{O}}(d)$ through data re-using while Montanari & Urbani (2025) provides a sharp DMFT theory analyzing the learning dynamics of single-index models using a standard two-layer neural network when the sample size and input dimension are proportional. Together, these two works essentially conclude the studies of learning single-index models with neural networks. Next, Dandi et al. (2024) demonstrates that

neural networks can learn multi-index models with generative leap exponent one using $n = \mathcal{O}(d)$ samples in constant time, leveraging DMFT theory. Finally, as noted previously, Arnaboldi et al. (2025) provides a loose result, showing that neural networks with $\mathcal{O}(d \log^2 d)$ samples can recover the hidden subspace of multi-index models but in a weak recovery sense when the generative leap exponent of the target equals two. Besides linear features, we note that there is also a substantial line of work on learning nonlinear features with deep neural networks, including Nichani et al. (2023); Ren et al. (2023); Wang et al. (2023); Fu et al. (2024).

We are aware of and highlight this recent concurrent work Montanari & Wang (2026). Our works differ in the following aspects. Montanari & Wang (2026) studies learning a multi-index model using a two-layer network trained with standard gradient descent instead of modified gradient-based algorithms. Their arguments are performed in the proportional regime, that is, $n/d \rightarrow \delta \in (0, +\infty)$, leveraging toolboxes from discrete time DMFT and random matrix theory. Besides, their results are sharp in constants, and they derive the sharp phase transition threshold at which the trained neural network can start to weakly recover the first hard direction (with generative exponent two). In contrast, Montanari & Wang (2026) does not establish full subspace recovery and therefore cannot provide an end-to-end analysis of the problem. Instead our work provides an end-to-end learning guarantee via a somehow different mechanism, albeit at the cost of a slightly worse sample complexity and the use of a modified layer-wise training scheme.

B NUMERICAL ILLUSTRATIONS

We conduct preliminary numerical experiments to validate our theoretical predictions.

We consider the following target function $f^*(\mathbf{x}) = (x_1^2 + \frac{1}{2}x_2^2)/\sqrt{5/2}$ where $\mathbf{x} \in \mathbb{R}^d$ and x_i denotes the i -th coordinate. This is a multi-index model with $r = 2$ hidden directions. We use a two-layer neural network with width $m = 4$ and quadratic activation function $\sigma(t) = t^2$. Both layers use standard Kaiming uniform initialization for weights and biases. For each dimension d and sample size exponent α , we generate $n = \lfloor d^\alpha \rfloor$ training samples from $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. We train the network using Adam optimizer with learning rate $\eta = 0.005$ for up to 1000 epochs with batch size 32. We vary the dimension d logarithmically and search the sample size exponent α in the range $[1.1, 1.8]$ with step size 0.01. For each (d, α) pair, we run 5 independent trials with different random seeds and report the average minimal sample size exponent α required to achieve test error below threshold ϵ .

B.1 RESULTS

Figure 1 shows the minimal-sample-size exponent α required to achieve test error below threshold ϵ converges to one as d increases for different thresholds $\epsilon \in \{1.0, 0.1, 0.01\}$. Each point represents the average over 5 random seeds. These experiments validate that neural networks can indeed learn multi-index models with near-optimal sample complexity. The decreasing trend in α as dimension grows provides empirical support for our theoretical finding that $n = \tilde{\mathcal{O}}(d)$ samples suffice for successful learning, achieving the information-theoretic limit up to leading order.

B.2 PROPER CHOICE OF LOSS FUNCTION

We demonstrate that the choice of loss function critically impacts learning performance for a subset of targets. We consider a challenging target $h_4(x_1) + h_4(x_2)$ where $h_4(z) = \frac{1}{\sqrt{24}}(z^4 - 6z^2 + 3)$ is the degree-4 Hermite polynomial. This target has generative exponent 2 but information exponent 4 which is larger than generative exponent and requires a properly chosen loss function rather than square loss to ensure the second-order information is non-degenerate. To this end, we train a two-layer neural network with width $m = 4$ and cosine activation $\sigma(t) = \cos(t)$ using two different loss functions: Mean Squared Error (MSE) and Huber loss. We measure the learning process by the alignment metric $\cos(\text{best})$, defined as the maximum cosine similarity between any first-layer weight vector and the true directions \mathbf{e}_1 or \mathbf{e}_2 . A value close to one indicates successful recovery of the hidden directions. For dimensions $d \in \{200, 300, 500\}$, we vary the sample size n and plot the alignment $\cos(\text{best})$ as a function of sample complexity ratio n/d . For each (d, n) pair, we run 10 independent trials with different random seeds and report the median along with 30th and 70th percentiles.

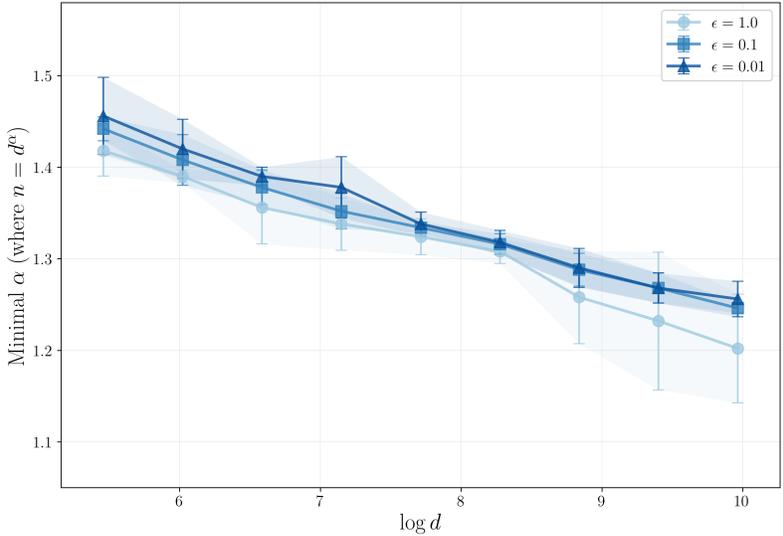


Figure 1: Minimal sample size exponent α (where $n = d^\alpha$) required to achieve test error $\leq \epsilon$ as a function of dimension d . Three thresholds are shown: $\epsilon = 1.0$ (light blue), $\epsilon = 0.1$ (medium blue), and $\epsilon = 0.01$ (dark blue). More stringent error thresholds require larger α (more samples), but all curves exhibit similar downward trends.

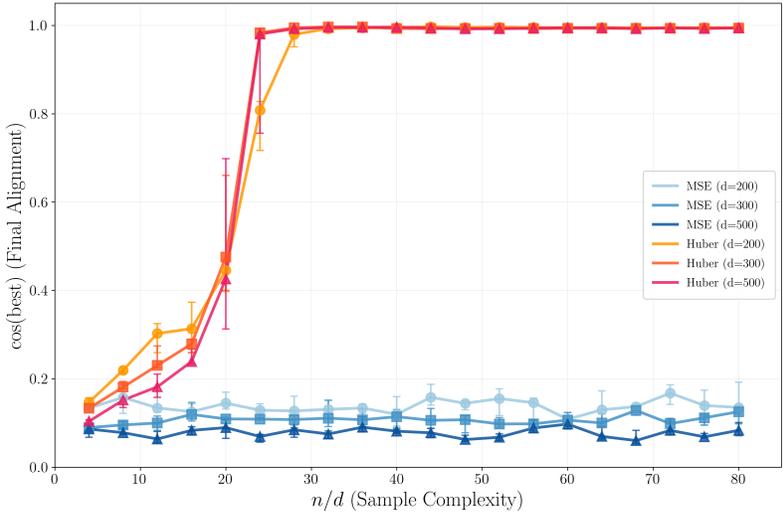


Figure 2: MSE and Huber loss for learning $h_4(x_1) + h_4(x_2)$ across different dimensions $d \in \{200, 300, 500\}$. The x -axis shows n/d , and the y -axis shows the cos similarity $\cos(\text{best})$. Lines show median with error bars indicating 30th to 70th percentiles over 10 random seeds. Blue shades represent MSE and red/orange shades represent Huber loss.

Figure 2 reveals a striking difference between the two loss functions. MSE loss consistently fails to learn the target function, with alignment remaining near zero even when $n/d = 80$. In stark contrast, Huber loss exhibits a clear phase transition: the alignment rapidly increases from near 0 to near 1 once the sample size exceeds a critical threshold around $n/d = 20$. This threshold is consistent across all three dimensions. This experiment validates our theoretical assumption (Assumption 5) that the pre-processing function $\ell'(0, y)$ must ensure non-degenerate second-order information. For this particular target, MSE loss corresponds to a degenerate pre-processing (since $\mathbb{E}[yxx^T] = \mathbf{0}$), while Huber loss provides the necessary non-degeneracy. This demonstrates that: (i) the choice of loss function is not merely a matter of optimization convenience but fundamentally influences

whether the information-theoretic limit can be achieved, and (ii) our framework correctly predicts that with proper loss selection, neural networks can learn targets beyond previous theoretical guarantees with linear sample complexity.

C PROOF OF LEMMA C.1: FEATURE LEARNING

We will use η to refer to the learning rate η_1 and use β to refer to β_1 throughout this section. We recall that we use D to refer to a large absolute constant.

We will first introduce several more notations to help us to do the proof. In the first stage, we expect the prediction $f_{\Theta^{(t)}}(\mathbf{x}_i)$ to remain small, since the norm of the features should remain small. We apply a first-order Taylor expansion of the first argument of $\ell_i^{(t)}$ around zero

$$\ell_i^{(t)} = \ell'(0, y_i) + \ell''(\zeta_i^{(t)}, y_i) f_{\Theta^{(t)}}(\mathbf{x}_i) \quad \text{for some } \left| \zeta_i^{(t)} \right| \leq |f_{\Theta^{(t)}}(\mathbf{x}_i)| \quad (26)$$

Denote the remainder term by

$$\Delta_i^{(t)} := \ell''(\zeta_i^{(t)}, y_i) f_{\Theta^{(t)}}(\mathbf{x}_i) \quad (27)$$

which satisfies the bound $|\Delta_i^{(t)}| \leq L |f_{\Theta^{(t)}}(\mathbf{x}_i)|$ and should be small in principle. Substituting the expansion into the gradient expression for the j -th neuron, we get

$$\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}}_n(\Theta^{(t)}) = \frac{1}{n} \sum_{i=1}^n (\ell'(0, y_i) + \Delta_i^{(t)}) a_j \sigma'(\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i. \quad (28)$$

Grouping terms, we have

$$\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}}_n(\Theta^{(t)}) = a_j \left[\underbrace{\frac{1}{n} \sum_{i=1}^n \ell'(0, y_i) \sigma'(\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i}_{\text{Zeroth-order term}} + \underbrace{\frac{1}{n} \sum_{i=1}^n \Delta_i^{(t)} \sigma'(\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i}_{\text{Perturbation}} \right]. \quad (29)$$

To analyze the gradient dynamics, we further expand the derivative of the activation function using a second-order Taylor approximation around zero. Since $\sigma'(0) = 0$, we obtain

$$\sigma'(z) = \sigma''(0)z + \frac{1}{2} \sigma^{(3)}(\xi) z^2, \quad \text{for some } |\xi| \leq |z|. \quad (30)$$

Substituting the expansion of $\sigma'(\cdot)$ into $\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}}$, we obtain

$$\begin{aligned} \nabla_{\mathbf{w}_j} \widehat{\mathcal{L}} &= a_j \sigma''(0) \widehat{\Sigma}_\ell \mathbf{w}_j + \underbrace{\frac{a_j}{2n} \sum_{i=1}^n \ell_i \sigma^{(3)}(\xi_i) (\mathbf{w}_j^\top \mathbf{x}_i)^2 \mathbf{x}_i + \frac{a_j}{n} \sum_{i=1}^n \Delta_i^{(t)} \sigma'(\mathbf{w}_j^\top \mathbf{x}_i) \mathbf{x}_i}_{\text{Perturbation}} \\ &= a_j \sigma''(0) \widehat{\Sigma}_\ell \mathbf{w}_j + \frac{a_j}{2n} \left(\sum_{i=1}^n \ell_i \sigma^{(3)}(\xi_i) (\mathbf{w}_j^\top \mathbf{x}_i)^2 \mathbf{x}_i \right. \\ &\quad \left. + 2\sigma''(0) \sum_{i=1}^n \Delta_i^{(t)} (\mathbf{w}_j^\top \mathbf{x}_i) \mathbf{x}_i + \sum_{i=1}^n \Delta_i^{(t)} \sigma^{(3)}(\xi_i) (\mathbf{w}_j^\top \mathbf{x}_i)^2 \mathbf{x}_i \right). \end{aligned} \quad (31)$$

The first term captures the leading-order linear dynamics that should be dominating, while the second term is a higher-order perturbation term that will be controlled using norm bounds on $\mathbf{w}_j^{(t)}$ and also bounds on $\Delta_i^{(t)}$.

Under gradient descent with a proper weight decay (recall that we set $\beta = \frac{1}{\eta}$) and the learning rate as η , the update rule for \mathbf{w}_j becomes

$$\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} - \eta (\nabla_{\mathbf{w}_j} \widehat{\mathcal{L}} + \beta \mathbf{w}_j^{(t)}) = -\eta a_j \sigma''(0) \widehat{\Sigma}_\ell \mathbf{w}_j^{(t)} + R_j^{(t)} \quad (32)$$

where the residual term $R_j^{(t)}$ is defined by

$$R_j^{(t)} := -\frac{\eta a_j}{2n} \left(\sum_{i=1}^n \ell_i \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x}_i)^2 \mathbf{x}_i + 2 \sum_{i=1}^n \Delta_i^{(t)} (\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i + \sum_{i=1}^n \Delta_i^{(t)} \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x}_i)^2 \mathbf{x}_i \right) \quad (33)$$

Since $(\mathbf{w}_j^{(t)\top} \mathbf{x}_i)^2 \mathbf{x}_i = (\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \mathbf{w}_j^{(t)}$, we can write $R_j^{(t)}$ as $R_j^{(t)} = -\eta a_j \mathbf{Q}^{(t)} \mathbf{w}_j^{(t)}$ where $\mathbf{Q}^{(t)} \in \mathbb{R}^{d \times d}$ is a symmetric matrix defined by

$$\mathbf{Q}^{(t)} := \frac{1}{2n} \left(\sum_{i=1}^n \ell_i \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top + 2 \sum_{i=1}^n \Delta_i^{(t)} \mathbf{x}_i \mathbf{x}_i^\top + \sum_{i=1}^n \Delta_i^{(t)} \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \right) \quad (34)$$

Now the entire update can be expressed as

$$\mathbf{w}_j^{(t+1)} = -\eta a_j (\widehat{\Sigma}_\ell + \mathbf{Q}^{(t)}) \mathbf{w}_j^{(t)}. \quad (35)$$

Define

$$\mathbf{M}^{(t)} := \widehat{\Sigma}_\ell + \mathbf{Q}^{(t)}, \quad \text{so that} \quad \mathbf{w}_j^{(t+1)} = -\eta a_j \mathbf{M}^{(t)} \mathbf{w}_j^{(t)}. \quad (36)$$

Through all the calculations above, we obtain a compact form of the update rule with a clear linear structure. Then by repeatedly following the updates for T_1 iterations, we achieve

$$\mathbf{w}_j^{(T_1)} = \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} \right) \mathbf{w}_j^{(0)}. \quad (37)$$

This shows that $\mathbf{w}_j^{(t)}$ lies in the image of a composition of linear operators, each consisting of a dominant direction determined by $\widehat{\Sigma}_\ell$, plus a time-dependent perturbation $\mathbf{Q}^{(t)}$.

The general picture is that $\mathbf{Q}^{(t)}$ is negligible when $t = o(\log d)$. In this case, we have $\mathbf{w}_j^{(t+1)} \propto \widehat{\Sigma}_\ell \mathbf{w}_j^{(t)}$ and the dynamics of GD will be similar to a power method iteration. For now note that our power method iteration is with respect to $\widehat{\Sigma}_\ell$, not Σ_ℓ , since the sample size is finite.

This section is devoted to a complete proof to bound the difference between $\mathbf{w}_j^{(T_1)}$ and the ideal update

$$\widehat{\mathbf{w}}_j^{(T_1)} := \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \Sigma_\ell^{T_1} \mathbf{w}_j^{(0)} \quad (38)$$

Note that we have

$$\left\| \mathbf{w}_j^{(T_1)} - \widehat{\mathbf{w}}_j^{(T_1)} \right\| = \frac{1}{\epsilon_0} \eta^{T_1} \left\| \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \Sigma_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \right\| \quad (39)$$

and therefore we just need to bound $\left\| \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \Sigma_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \right\|$.

Lemma C.1 (Total Perturbation Error). *With probability at least $1 - \mathcal{O}(d^{-D})$, assuming $n \gtrsim dT_1^2$ and $m \leq d^{D/4}$, for all $j \in [m]$, the total perturbation error is bounded by*

$$\left\| \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \Sigma_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \right\| \lesssim \epsilon_0 T_1 \gamma \sqrt{\frac{\log d}{d}} + \epsilon_0 \gamma^{T_1}, \quad \text{where} \quad \gamma = C \sqrt{\frac{d}{n}} \quad (40)$$

where C is a universal constant.

The proof of Lemma C.1 is provided in Appendix C.1.

C.1 PROOF OF LEMMA C.1

We first give a sharp characterization of $\widehat{\Sigma}_\ell$. We will bound the difference of the empirical and the population. Define

$$\mathbf{H}_\ell := \frac{1}{n} \sum_{i=1}^n \ell_i \mathbf{x}_i \mathbf{x}_i^\top - \mathbb{E}[\ell \mathbf{x} \mathbf{x}^\top]. \quad (41)$$

Lemma C.2 (Operator Norm Bound for \mathbf{H}_ℓ). *With probability $1 - \mathcal{O}(d^{-2D})$ we have $\|\mathbf{H}_\ell\|_{\text{op}} \lesssim \sqrt{\frac{d}{n}}$.*

From this lemma, we expect a sharp concentration once $n \gg d$. This is standard in literature. Next, we bound the magnitude of $f_{\Theta^{(t)}}(\mathbf{x})$.

Lemma C.3. *Let $f_{\Theta^{(t)}}(\mathbf{x}) = \sum_{j=1}^m a_j \sigma(\mathbf{w}_j^{(t)\top} \mathbf{x})$ be the output of a two-layer neural network at iteration t , where each $\mathbf{w}_j^{(t)} \in \mathbb{R}^d$ satisfies $\|\mathbf{w}_j^{(t)}\| \leq \epsilon$ and the input $\mathbf{x} \in \mathbb{R}^d$ satisfies $\|\mathbf{x}\| \leq 2\sqrt{d}$.*

Then, for all such \mathbf{x} , the network output satisfies the bound

$$|f_{\Theta^{(t)}}(\mathbf{x})| \leq m \left(2\epsilon^2 d + \frac{4}{3} M \epsilon^3 d^{3/2} \right). \quad (42)$$

Next, we start to control the operator norm of the matrix $\mathbf{Q}^{(t)}$.

Lemma C.4 (Operator Norm Bound for $\mathbf{Q}^{(t)}$). *Suppose the following holds. Each input satisfies $\|\mathbf{x}_i\| \leq 2\sqrt{d}$. Each weight vector satisfies $\|\mathbf{w}_j^{(t)}\| \leq \epsilon$. The network prediction at iteration t satisfies $|f_{\Theta^{(t)}}(\mathbf{x}_i)| \leq C_f$.*

Then the operator norm of $\mathbf{Q}^{(t)}$ satisfies

$$\|\mathbf{Q}^{(t)}\|_{\text{op}} \leq 4\epsilon d^{3/2} (LM + LC_f M) + 4LC_f d \quad (43)$$

Plugging in the concrete estimation for the constant C_f , also note that without loss of generality we can assume $L \geq 1$ and $M \geq 1$. we have the following.

Corollary C.5. *Suppose the following holds. Each input satisfies $\|\mathbf{x}_i\| \leq 2\sqrt{d}$. Each weight vector satisfies $\|\mathbf{w}_j^{(t)}\| \leq \epsilon$.*

Then $\mathbf{Q}^{(t)}$ satisfies

$$\|\mathbf{Q}^{(t)}\|_{\text{op}} \leq 4LM\epsilon d^{3/2} + 8Lm\epsilon^2 d^2 + \frac{40}{3}LMm\epsilon^3 d^{5/2} + \frac{16}{3}LM^2m\epsilon^4 d^3 \leq 32L^2M^2md^3\epsilon. \quad (44)$$

With all the lemmas and corollaries above, define the following high probability events. First, we define the event \mathcal{E}_H (finite-sample noise bound) as follows:

$$\mathcal{E}_H := \left\{ \|\mathbf{H}_\ell\|_{\text{op}} \lesssim \gamma \text{ where } \gamma = C\sqrt{\frac{d}{n}} \right\}. \quad (45)$$

By Lemma C.2, this holds with probability at least $1 - \mathcal{O}(d^{-2D})$. Furthermore, provided $n \gtrsim dT_1^2$, we know that $\gamma \leq \frac{1}{4}$ also $\gamma \leq \frac{1}{2T_1}$. Second, we define the event \mathcal{E}_D (truncated norm) as follows:

$$\mathcal{E}_D := \{\|\mathbf{x}_i\| \leq 2\sqrt{d} \text{ for } i \in [n]\}. \quad (46)$$

By Lemma E.18, this holds with probability at least $1 - n \exp(-cd)$, thus at least $1 - \mathcal{O}(d^{-2D})$.

Third, we define the event $\mathcal{E}_Q^{(t)}$ as $\mathcal{E}_Q^{(t)} := \left\{ \|\mathbf{Q}^{(t)}\|_{\text{op}} \leq \frac{1}{4} \right\}$. Note that under event \mathcal{E}_D , provided $\|\mathbf{w}_j^{(t)}\| \leq \frac{1}{128mL^2M^2d^3}$, we have $\mathcal{E}_Q^{(t)}$ happens. We notice that $\mathcal{E}_Q^{(t)}$ has no independent randomness, it is defined just for notation convenience.

In the following we analyze the gradient update conditioned on the event $\mathcal{E}_H \cap \mathcal{E}_D$, which holds with high probability. First bound the norm of $\mathbf{w}_j^{(t)}$.

Lemma C.6 (Norm Bound after T_1 Steps). *Suppose the event $\mathcal{E}_H \cap \mathcal{E}_D$ holds, and let the learning rate be set as*

$$\eta = \frac{1}{C_\eta} \left(\frac{d}{r\iota^2} \right)^{1/(2T_1)} \quad (47)$$

for a sufficiently large but fixed constant C_η where $\iota = 4D \log d$. Suppose

$$\|\mathbf{w}_j^{(0)}\| = \epsilon_0 \leq \frac{1}{128mL^2M^2d^{7/2}}. \quad (48)$$

Then for all $0 \leq t \leq T_1 - 1$, we have

$$\|\mathbf{w}_j^{(t)}\| \leq \sqrt{d}\epsilon_0 \quad \text{and} \quad \|\mathbf{w}_j^{(t)}\| \leq \frac{1}{128mL^2M^2d^3} \quad (49)$$

Next, we remove $\mathbf{Q}^{(t)}$ out of $\mathbf{M}^{(t)} = \widehat{\Sigma}_\ell + \mathbf{Q}^{(t)}$ and compute the perturbation error.

Lemma C.7. *Suppose the event $\mathcal{E}_H \cap \mathcal{E}_D$ holds. Define $\Delta := \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w}_j^{(0)}$. Assume that $\epsilon_0 \leq \frac{1}{128mL^2M^2d^{7/2}}$. Then the perturbation error is bounded by*

$$\|\Delta\| \lesssim T_1 \left(\frac{5}{4} \right)^{T_1-1} m d^{7/2} \epsilon_0^2. \quad (50)$$

Next, we deal with the difference between Σ_ℓ and $\widehat{\Sigma}_\ell$.

Lemma C.8. *Let $\mathbf{w} \sim \text{Unif}(\mathbb{S}^{d-1}(\epsilon_0))$ be a random independent vector. Then with probability at least $1 - \mathcal{O}(d^{-2D})$, we have*

$$\begin{aligned} \|(\Sigma_\ell + \mathbf{H}_\ell)^{T_1} \mathbf{w} - \Sigma_\ell^{T_1} \mathbf{w}\| &\lesssim \epsilon_0 \left((\|\Sigma_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}})^{T_1} - \|\Sigma_\ell\|_{\text{op}}^{T_1} \right) \sqrt{\frac{\log d}{d}} \\ &\quad + \epsilon_0 \|\mathbf{H}_\ell\|_{\text{op}}^{T_1}. \end{aligned} \quad (51)$$

Corollary C.9. *Under the setup of Lemma C.8 and note we assumed $\|\mathbf{H}_\ell\|_{\text{op}} \leq \gamma$ with $\gamma \leq 1/(2T_1)$. Then with probability at least $1 - \mathcal{O}(d^{-2D})$, we have*

$$\|(\Sigma_\ell + \mathbf{H}_\ell)^{T_1} \mathbf{w} - \Sigma_\ell^{T_1} \mathbf{w}\| \lesssim \epsilon_0 \left(T_1 \gamma \sqrt{\frac{\log d}{d}} + \gamma^{T_1} \right). \quad (52)$$

Proof. From the binomial expansion,

$$(1 + \gamma)^{T_1} - 1 = \sum_{j=1}^{T_1} \binom{T_1}{j} \gamma^j = T_1 \gamma + \sum_{j=2}^{T_1} \binom{T_1}{j} \gamma^j. \quad (53)$$

Since $\sum_{j=2}^{\infty} (T_1 \gamma)^j \leq (T_1 \gamma)^2 / (1 - T_1 \gamma) \leq 2(T_1 \gamma)^2 \leq T_1 \gamma$, we get $(1 + \gamma)^{T_1} - 1 \leq 2T_1 \gamma$. On the other hand, as in Lemma C.8, one shows under the same event that

$$\|(\Sigma_\ell + \mathbf{H}_\ell)^{T_1} \mathbf{w} - \Sigma_\ell^{T_1} \mathbf{w}\| \lesssim \epsilon_0 \left((1 + \gamma)^{T_1} - 1 \right) \sqrt{\frac{\log d}{d}} + \gamma^{T_1}. \quad (54)$$

Combining the two bounds yields the claimed result. \square

By putting all the things together, we obtain the final bound of Lemma C.1.

Proof of Lemma C.1. We first prove another lemma. We aim to first show that under the stated condition on ϵ_0 , the perturbation error from the time-varying iterates is smaller than the leading-order perturbation induced by \mathbf{H}_ℓ .

Lemma C.10. Suppose the event $\mathcal{E}_H \cap \mathcal{E}_D$ holds. Recall $\Delta = \left(\prod_{t=0}^{T_1-1} M^{(t)} - \widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w}_j^{(0)}$ denote the error from time-varying perturbations.

Then, provided that $\epsilon_0 \leq \frac{1}{mLM^2 d^{7/2}} \sqrt{\frac{\log d}{n}} (4/5)^{T_1}$ we have

$$\|\Delta\| \lesssim T_1 \gamma \sqrt{\frac{\log d}{d}} \cdot \epsilon_0 \quad (55)$$

Proof. Recall the bounds established in previous lemmas. We have the one-step expansion error from iterated time-varying perturbations

$$\|\Delta\| \lesssim T_1 \left(\frac{5}{4} \right)^{T_1-1} m d^{7/2} \epsilon_0^2 \quad (56)$$

We plug in one ϵ_0

$$\begin{aligned} \|\Delta\| &\lesssim T_1 \left(\frac{5}{4} \right)^{T_1-1} m d^{7/2} \epsilon_0^2 \\ &\lesssim T_1 \left(\frac{5}{4} \right)^{T_1-1} m d^{7/2} \epsilon_0 \frac{1}{m d^{7/2}} \left(\frac{4}{5} \right)^{T_1} \sqrt{\frac{\log d}{n}} \\ &\lesssim T_1 \gamma \sqrt{\frac{\log d}{d}} \epsilon_0 \end{aligned} \quad (57)$$

and the proof is complete. \square

Finally we decompose the total perturbation error into two parts

$$\begin{aligned} \left(\prod_{t=0}^{T_1-1} M^{(t)} - \Sigma_\ell^{T_1} \right) \mathbf{w}_j^{(0)} &= \left(\prod_{t=0}^{T_1-1} M^{(t)} - \widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w}_j^{(0)} + \left(\widehat{\Sigma}_\ell^{T_1} - \Sigma_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \\ &\lesssim T_1 \gamma \sqrt{\frac{\log d}{d}} \epsilon_0 + \epsilon_0 \left(T_1 \gamma \sqrt{\frac{\log d}{d}} + \gamma^{T_1} \right) \\ &\lesssim \epsilon_0 \left(T_1 \gamma \sqrt{\frac{\log d}{d}} + \gamma^{T_1} \right) \end{aligned} \quad (58)$$

The conclusion is then straightforward if we adapt the previous two lemmas to bound these two parts separately. \square

C.2 PROOF OF SUPPORTING LEMMAS IN APPENDIX C.1

C.2.1 PROOF OF LEMMA C.2

Proof. We analyze the operator norm via a net argument.

Recall that $\|\mathbf{H}_\ell\|_{\text{op}} = \sup_{\|\mathbf{u}\|=\|\mathbf{v}\|=1} \mathbf{u}^\top \mathbf{H}_\ell \mathbf{v}$. Let $M^{1/4}, N^{1/4}$ be two $\frac{1}{4}$ -nets of the unit sphere in \mathbb{R}^d . Then

$$\|\mathbf{H}_\ell\|_{\text{op}} \leq 2 \max_{\mathbf{u} \in M^{1/4}, \mathbf{v} \in N^{1/4}} \mathbf{u}^\top \mathbf{H}_\ell \mathbf{v} \quad (59)$$

Fix any \mathbf{u}, \mathbf{v} , and define

$$Z_i := (\mathbf{u}^\top \mathbf{x}_i)(\mathbf{v}^\top \mathbf{x}_i) \ell_i - \mathbb{E}[(\mathbf{u}^\top \mathbf{x})(\mathbf{v}^\top \mathbf{x}) \ell]. \quad (60)$$

Note that we have $\mathbf{u}^\top \mathbf{H}_\ell \mathbf{v} = \frac{1}{n} \sum_{i=1}^n Z_i$. Since \mathbf{x}_i is standard Gaussian and $\ell_i \in [-L, L]$, we have $\|Z_i\|_{\psi_1} \lesssim L$. Therefore we can invoke Bernstein's inequality and do a union bound over the net to obtain

$$\mathbb{P}(\|\mathbf{H}_\ell\|_{\text{op}} \geq t) \leq 81^d \exp\left(-\frac{cnt^2}{L^2}\right). \quad (61)$$

Choosing $t = CL\sqrt{\frac{d}{n}}$ for sufficiently large C , we obtain the desired conclusion. \square

C.2.2 PROOF OF LEMMA C.3

Proof. Using the Taylor expansion of $\sigma(\cdot)$ at 0 with $\sigma'(0) = 0$ and $\sigma''(0) = 1$, we have

$$\sigma(z) = \sigma(0) + \frac{1}{2}z^2 + \frac{1}{6}\sigma^{(3)}(\xi)z^3 \quad \text{for some } |\xi| \leq |z| \quad (62)$$

Then

$$f_{\Theta^{(t)}}(\mathbf{x}) = \sum_{j=1}^m a_j \sigma(\mathbf{w}_j^{(t)\top} \mathbf{x}) = \sum_{j=1}^m a_j \left[\sigma(0) + \frac{1}{2}(\mathbf{w}_j^{(t)\top} \mathbf{x})^2 + \frac{1}{6}\sigma^{(3)}(\xi_j^{(t)})(\mathbf{w}_j^{(t)\top} \mathbf{x})^3 \right] \quad (63)$$

Since $\sum_j a_j = 0$, the constant term cancels and

$$f_{\Theta^{(t)}}(\mathbf{x}) = \frac{1}{2} \sum_{j=1}^m a_j (\mathbf{w}_j^{(t)\top} \mathbf{x})^2 + \frac{1}{6} \sum_{j=1}^m a_j \sigma^{(3)}(\xi_j^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x})^3 \quad (64)$$

We then bound each term. We have

$$|\mathbf{w}_j^{(t)\top} \mathbf{x}| \leq \|\mathbf{w}_j^{(t)}\| \|\mathbf{x}\| \leq 2\epsilon\sqrt{d}. \quad (65)$$

Therefore $(\mathbf{w}_j^{(t)\top} \mathbf{x})^2 \leq 4\epsilon^2 d$ and $(\mathbf{w}_j^{(t)\top} \mathbf{x})^3 \leq 8\epsilon^3 d^{3/2}$. Then

$$\left| \sum_{j=1}^m a_j (\mathbf{w}_j^{(t)\top} \mathbf{x})^2 \right| \leq 4m\epsilon^2 d, \quad \left| \sum_{j=1}^m a_j \sigma^{(3)}(\xi_j^{(t)}) (\mathbf{w}_j^{(t)\top} \mathbf{x})^3 \right| \leq 8mM\epsilon^3 d^{3/2} \quad (66)$$

and the following is straightforward

$$|f_{\Theta^{(t)}}(\mathbf{x})| \leq \frac{1}{2}(4m\epsilon^2 d) + \frac{1}{6}(8mM\epsilon^3 d^{3/2}) = 2m\epsilon^2 d + \frac{4}{3}mM\epsilon^3 d^{3/2}. \quad (67)$$

The proof is complete. \square

C.2.3 PROOF OF LEMMA C.4

Proof. We upper bound the operator norm of $\mathbf{Q}^{(t)}$ by treating the three additive components separately.

Let

$$\mathbf{Q}_I^{(t)} := \frac{1}{2n} \sum_{i=1}^n \ell_i \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^\top \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \quad (68)$$

We bound each component $|\mathbf{w}_j^\top \mathbf{x}_i| \leq \|\mathbf{w}_j\| \|\mathbf{x}_i\| \leq 2\epsilon\sqrt{d}$, $\|\mathbf{x}_i \mathbf{x}_i^\top\|_{\text{op}} = \|\mathbf{x}_i\|^2 \leq 4d$, $|\ell_i \sigma^{(3)}(\xi_i^{(t)})| \leq LM$. Thus,

$$\|\mathbf{Q}_I^{(t)}\|_{\text{op}} \leq \frac{1}{2n} nLM(2\epsilon\sqrt{d})(4d) = 4LM\epsilon d^{3/2}. \quad (69)$$

Let

$$\mathbf{Q}_{II}^{(t)} := \frac{1}{n} \sum_{i=1}^n \Delta_i^{(t)} \mathbf{x}_i \mathbf{x}_i^\top. \quad (70)$$

Each term satisfies $|\Delta_i^{(t)}| \leq LC_f$, $\|\mathbf{x}_i \mathbf{x}_i^\top\|_{\text{op}} \leq 4d$, so

$$\|\mathbf{Q}_{II}^{(t)}\|_{\text{op}} \leq LC_f(4d). \quad (71)$$

Let

$$\mathbf{Q}_{III}^{(t)} := \frac{1}{2n} \sum_{i=1}^n \Delta_i^{(t)} \sigma^{(3)}(\xi_i^{(t)}) (\mathbf{w}_j^\top \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top. \quad (72)$$

We again use $|\Delta_i^{(t)} \sigma^{(3)}(\xi_i^{(t)})| \leq LC_f M$, $|\mathbf{w}_j^\top \mathbf{x}_i| \leq 2\epsilon\sqrt{d}$, $\|\mathbf{x}_i \mathbf{x}_i^\top\|_{\text{op}} \leq 4d$, to obtain

$$\|\mathbf{Q}_{III}^{(t)}\|_{\text{op}} \leq \frac{1}{2n} nLC_f M(2\epsilon\sqrt{d})(4d) = 4LC_f M\epsilon d^{3/2}. \quad (73)$$

We combine the operator norm bounds

$$\|\mathbf{Q}^{(t)}\|_{\text{op}} \leq \|\mathbf{Q}_I^{(t)}\|_{\text{op}} + \|\mathbf{Q}_{II}^{(t)}\|_{\text{op}} + \|\mathbf{Q}_{III}^{(t)}\|_{\text{op}} \leq 4\epsilon d^{3/2}(LM + LC_f M) + 4LC_f d. \quad (74)$$

The proof is complete. \square

C.2.4 PROOF OF LEMMA C.6

Proof. We first prove the following supporting lemma.

Lemma C.11 (One-Step Norm Control). *Suppose the event $\mathcal{E}_H \cap \mathcal{E}_D \cap \mathcal{E}_Q^{(t)}$ holds, and let the learning rate be set as*

$$\eta = \frac{1}{C_\eta} \left(\frac{d}{r\iota^2} \right)^{1/(2T_1)} \quad (75)$$

for a sufficiently large constant C_η where $\iota = 4D \log d$. Then the gradient update satisfies

$$\|\mathbf{w}_j^{(t+1)}\| \leq d^{1/(2T_1)} \|\mathbf{w}_j^{(t)}\| \quad (76)$$

Proof. Under the event $\mathcal{E}_H \cap \mathcal{E}_Q^{(t)}$, we have

$$\|\mathbf{M}^{(t)}\|_{\text{op}} \leq \|\Sigma_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}} + \|\mathbf{Q}^{(t)}\|_{\text{op}} \leq 1 + \frac{1}{4} + \frac{1}{4} = \frac{3}{2} \quad (77)$$

Then,

$$\|\mathbf{w}_j^{(t+1)}\| = \eta \|\mathbf{M}^{(t)} \mathbf{w}_j^{(t)}\| \leq \eta \|\mathbf{M}^{(t)}\|_{\text{op}} \cdot \|\mathbf{w}_j^{(t)}\| \leq \frac{3}{2} \eta \cdot \|\mathbf{w}_j^{(t)}\|. \quad (78)$$

Plug the formula for η into the above one and we obtain the desired conclusion. \square

We prove by induction on $t \in \{0, 1, \dots, T_1 - 1\}$ that $\|\mathbf{w}_j^{(t)}\| \leq d^{t/(2T_1)} \epsilon_0$ where we assume $\epsilon_0 \leq \frac{1}{128mL^2M^2d^{7/2}}$. The $t = 0$ case is trivial. Suppose for some $t \leq T_1$, we have

$$\|\mathbf{w}_j^{(t)}\| \leq d^{t/(2T_1)} \epsilon_0. \quad (79)$$

Plugging ϵ_0 in, $\|\mathbf{w}_j^{(t)}\| \leq \frac{1}{128mL^2M^2d^3}$, thus the event $\mathcal{E}_H \cap \mathcal{E}_D \cap \mathcal{E}_Q^{(t)}$ holds, apply the previous lemma C.11, we get

$$\|\mathbf{w}_j^{(t+1)}\| \leq d^{1/(2T_1)} \|\mathbf{w}_j^{(t)}\| \leq d^{(t+1)/(2T_1)} \epsilon_0. \quad (80)$$

Repeating this induction, the result is straightforward. In particular, for all $0 \leq t \leq T_1 - 1$, we have $\|\mathbf{w}_j^{(t)}\| \leq \sqrt{d} \epsilon_0$. \square

C.2.5 PROOF OF LEMMA C.7

Proof. We begin by expanding the product of perturbed operators using a standard multiplicative expansion

$$\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} = \widehat{\Sigma}_\ell^{T_1} + \sum_{s=1}^{T_1} \sum_{0 \leq T_1 < \dots < T_1} \widehat{\Sigma}_\ell^{T_1-t_s-1} \mathbf{Q}^{(t_s)} \widehat{\Sigma}_\ell^{t_s-1-1} \dots \mathbf{Q}^{(T_1)} \widehat{\Sigma}_\ell^{T_1}. \quad (81)$$

Subtracting $\widehat{\Sigma}_\ell^{T_1}$ from both sides yields

$$\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \widehat{\Sigma}_\ell^{T_1} = \sum_{s=1}^{T_1} \sum_{0 \leq T_1 < \dots < T_1} \widehat{\Sigma}_\ell^{T_1-t_s-1} \mathbf{Q}^{(t_s)} \widehat{\Sigma}_\ell^{t_s-1} \dots \mathbf{Q}^{(T_1)} \widehat{\Sigma}_\ell^{T_1}. \quad (82)$$

We now apply this expansion to $\mathbf{w}_j^{(0)}$, and estimate the resulting norm. For each term in the expansion with s perturbation matrices, we use the operator norm bound

$$\|\widehat{\Sigma}_\ell^{T_1-t_s-1} \mathbf{Q}^{(t_s)} \dots \mathbf{Q}^{(T_1)} \widehat{\Sigma}_\ell^{T_1} \mathbf{w}_j^{(0)}\| \leq \|\widehat{\Sigma}_\ell\|_{\text{op}}^{T_1-s} \cdot \prod_{k=1}^s \|\mathbf{Q}^{(t_k)}\|_{\text{op}} \cdot \|\mathbf{w}_j^{(0)}\| \quad (83)$$

Note that the event $\mathcal{E}_H \cap \mathcal{E}_D$ holds, $\|\widehat{\Sigma}_\ell\|_{\text{op}} \leq \|\Sigma_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}} \leq \frac{5}{4}$. Combining with

$$\epsilon_0 \leq \frac{1}{128mL^2M^2d^{7/2}} \quad (84)$$

by Lemma C.6, for all $0 \leq t \leq T_1 - 1$ $\|\mathbf{w}_j^{(t)}\| \leq \sqrt{d}\epsilon_0$. Moreover, using Corollary C.5, it is easy to know that

$$\left\| \mathbf{Q}^{(t)} \right\|_{\text{op}} \leq 32LM^2md^{7/2}\epsilon_0 \quad (85)$$

Hence, each term with s perturbation matrices is bounded by

$$\left(\frac{5}{4}\right)^{T_1-s} (32L^2M^2md^{7/2}\epsilon_0)^s \epsilon_0. \quad (86)$$

Summing over all such terms (total of $\binom{T_1}{s}$ per order s), we have

$$\left\| \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} - \widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \right\| \leq \sum_{s=1}^{T_1} \binom{T_1}{s} \left(\frac{5}{4}\right)^{T_1-s} (32L^2M^2md^{7/2}\epsilon_0)^s \epsilon_0. \quad (87)$$

This sum can be rewritten as

$$\epsilon_0 \left[\left(\frac{5}{4} + 32L^2M^2md^{7/2}\epsilon_0 \right)^{T_1} - \left(\frac{5}{4} \right)^{T_1} \right]. \quad (88)$$

Applying the first-order binomial expansion which is valid for the small enough ϵ_0 in our assumption and we conclude

$$\|\Delta\| \lesssim T_1 \left(\frac{5}{4}\right)^{T_1-1} mM^2d^{7/2}\epsilon_0^2. \quad (89)$$

The proof is complete. \square

C.2.6 PROOF OF LEMMA C.8

Proof. We first introduce several standard concentration results.

Lemma C.12. *Suppose \mathbf{A} is a matrix in $\mathbb{R}^{d \times d}$, and $\mathbf{w} \sim \text{Unif}(\mathbb{S}^{d-1})$. Then with probability at least $1 - \mathcal{O}(d^{-2D})$, we have*

$$|(\mathbf{A}\mathbf{w})_s| \lesssim \|\mathbf{A}\|_{\text{op}} \sqrt{\frac{\log d}{d}} \quad \text{for all } s \in [r]. \quad (90)$$

Proof. Fix $s \in [r]$, and define the function $f(\mathbf{w}) = (\mathbf{A}\mathbf{w})_s = \mathbf{e}_s^\top \mathbf{A}\mathbf{w}$.

This function is Lipschitz with

$$\|f\|_{\text{Lip}} = \|\mathbf{e}_s^\top \mathbf{A}\| \leq \|\mathbf{A}\|_{\text{op}}. \quad (91)$$

By the standard concentration results for Lipschitz function on the sphere for instance Theorem 5.1.4 in Vershynin (2018), and noting that $\mathbb{E}[f(\mathbf{w})] = 0$, we have the following concentration

$$\mathbb{P}(|f(\mathbf{w})| \geq t) \leq 2 \exp\left(-\frac{cdt^2}{\|\mathbf{A}\|^2}\right). \quad (92)$$

Set $t = 2\|\mathbf{A}\|_{\text{op}} \sqrt{\frac{D \log d}{cd}}$. Then we obtain

$$\mathbb{P}\left(|(\mathbf{A}\mathbf{w})_s| \geq 2\|\mathbf{A}\|_{\text{op}} \sqrt{\frac{D \log d}{cd}}\right) \leq d^{-2D}. \quad (93)$$

Applying a union bound over all the r coordinates and the desired result follows. \square

Lemma C.13. *Let $\mathbf{w} \sim \text{Unif}(\mathbb{S}^{d-1}(\epsilon_0))$. Let $\mathbf{H}_\ell \in \mathbb{R}^{d \times d}$ be a fixed matrix. Let $\mathbf{\Pi}^* \in \mathbb{R}^{r \times d}$ denote the projection onto the first r coordinates, that is $(\mathbf{\Pi}^* \mathbf{v})_s = v_s$ for $s \in [r]$.*

Then with probability at least $1 - \mathcal{O}(d^{-2D})$, we have

$$\|\mathbf{\Pi}^* \mathbf{H}_\ell^k \mathbf{w}\| \lesssim \epsilon_0 \cdot \|\mathbf{H}_\ell^k\|_{\text{op}} \cdot \sqrt{\frac{r \log d}{d}}, \quad \text{for all } k = 0, 1, \dots, T_1. \quad (94)$$

Proof. Fix $k \in \{0, 1, \dots, T_1\}$, and define the matrix $\mathbf{A} = \mathbf{H}_\ell^k$. By Lemma C.12, for each fixed $s \in [r]$, with probability at least $1 - \mathcal{O}(d^{-2D})$, we have

$$|(\mathbf{A}\mathbf{w})_s| = |(\mathbf{H}_\ell^k \mathbf{w})_s| \leq \epsilon_0 \|\mathbf{H}_\ell^k\|_{\text{op}} \sqrt{\frac{4D \log d}{cd}}, \quad (95)$$

for some small absolute constant c . Applying a union bound over all $s \in [r]$, we get the following bound holds

$$\forall s \in [r] \quad |(\mathbf{H}_\ell^k \mathbf{w})_s| \leq \epsilon_0 \|\mathbf{H}_\ell^k\|_{\text{op}} \sqrt{\frac{4D \log d}{cd}} \quad (96)$$

Hence

$$\|\mathbf{\Pi}^* \mathbf{H}_\ell^k \mathbf{w}\|^2 = \sum_{s=1}^r ((\mathbf{H}_\ell^k \mathbf{w})_s)^2 \leq r \left(\epsilon_0 \|\mathbf{H}_\ell^k\|_{\text{op}} \sqrt{\frac{4D \log d}{cd}} \right)^2 \quad (97)$$

and the conclusion follows with a square root calculation and a union bound over time. \square

To prove Lemma C.8, we invoke the following auxiliary results.

(i) **Operator norm bound.** $\|\mathbf{H}_\ell\|_{\text{op}} \lesssim \sqrt{\frac{d}{n}}$ with probability at least $1 - \mathcal{O}(d^{-2D})$.

(ii) **Projection bound.** For any integer $0 \leq \ell \leq T_1$, we have by Lemma C.13

$$\|\mathbf{\Pi}^* \mathbf{H}_\ell^\ell \mathbf{w}\| \lesssim \epsilon_0 \cdot \|\mathbf{H}_\ell\|_{\text{op}}^\ell \cdot \sqrt{\frac{r \log d}{d}} \quad \text{with probability at least } 1 - \mathcal{O}(d^{-2D}). \quad (98)$$

Under those two events, we expand

$$(\mathbf{\Sigma}_\ell + \mathbf{H}_\ell)^{T_1} \mathbf{w} - \mathbf{\Sigma}_\ell^{T_1} \mathbf{w} = \sum_{k=1}^{T_1} M_k(\mathbf{w}), \quad (99)$$

where $M_k(\mathbf{w})$ collects all degree- T_1 monomials with exactly k occurrences of \mathbf{H}_ℓ . Among them we have two kinds of terms.

First, the pure \mathbf{H}_ℓ term. There is exactly one such term:

$$\|\mathbf{H}_\ell^{T_1} \mathbf{w}\| \leq \epsilon_0 \|\mathbf{H}_\ell\|_{\text{op}}^{T_1}. \quad (100)$$

Second, the mixed terms. Each $M_k(\mathbf{w})$ consists of expressions of the form

$$T(\mathbf{w}) = \mathbf{A} \mathbf{\Sigma}_\ell^m \mathbf{H}_\ell^\ell \mathbf{w}, \quad (101)$$

where \mathbf{A} consists of a product of $T_1 - k - m$ $\mathbf{\Sigma}_\ell$'s and $k - \ell$ \mathbf{H}_ℓ 's. Because $\mathbf{\Sigma}_\ell$ is rank- r , and leftmost $\mathbf{\Sigma}_\ell^m$ acts as projection to top- r eigenspace, we may write

$$T(\mathbf{w}) = \mathbf{A} \mathbf{\Sigma}_\ell^m \mathbf{\Pi}^* (\mathbf{H}_\ell^\ell \mathbf{w}). \quad (102)$$

Now, we estimate each term.

$$\|\mathbf{A}\|_{\text{op}} \leq \|\mathbf{\Sigma}_\ell\|_{\text{op}}^{T_1 - k - m} \|\mathbf{H}_\ell\|_{\text{op}}^{k - \ell}; \quad (103)$$

$$\|\mathbf{\Sigma}_\ell^m\|_{\text{op}} = \|\mathbf{\Sigma}_\ell\|_{\text{op}}^m; \quad (104)$$

$$\|\mathbf{\Pi}^* (\mathbf{H}_\ell^\ell \mathbf{w})\| \lesssim \epsilon_0 \|\mathbf{H}_\ell\|_{\text{op}}^\ell \sqrt{\frac{r \log d}{d}}; \quad (105)$$

and therefore

$$\|T(\mathbf{w})\| \lesssim \epsilon_0 \|\mathbf{\Sigma}_\ell\|_{\text{op}}^{T_1 - k} \|\mathbf{H}_\ell\|_{\text{op}}^k \sqrt{\frac{r \log d}{d}}. \quad (106)$$

The number of such terms is at most $\binom{T_1}{k}$, and summing over $1 \leq k < T_1$

$$\begin{aligned} & \epsilon_0 \sum_{k=1}^{T_1-1} \binom{T_1}{k} \|\mathbf{\Sigma}_\ell\|_{\text{op}}^{T_1 - k} \|\mathbf{H}_\ell\|_{\text{op}}^k \sqrt{\frac{r \log d}{d}} \\ & \lesssim \epsilon_0 \left((\|\mathbf{\Sigma}_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}})^{T_1} - \|\mathbf{\Sigma}_\ell\|_{\text{op}}^{T_1} - \|\mathbf{H}_\ell\|_{\text{op}}^{T_1} \right) \sqrt{\frac{r \log d}{d}}. \end{aligned} \quad (107)$$

Adding the pure \mathbf{H}_ℓ term back, we conclude the final result with $1 - \mathcal{O}(d^{-2D})$ probability

$$\|(\boldsymbol{\Sigma}_\ell + \mathbf{H}_\ell)^{T_1} \mathbf{w} - \boldsymbol{\Sigma}_\ell^{T_1} \mathbf{w}\| \lesssim \epsilon_0 \left((\|\boldsymbol{\Sigma}_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}})^{T_1} - \|\boldsymbol{\Sigma}_\ell\|_{\text{op}}^{T_1} \right) \sqrt{\frac{r \log d}{d}} + \epsilon_0 \|\mathbf{H}_\ell\|_{\text{op}}^{T_1}. \quad (108)$$

The proof is complete. \square

D PROOF OF THEOREM 1

D.1 PROOF ROADMAP OF THEOREM 1

General roadmap. For notational convenience, given parameters Θ we denote the population loss and its empirical counterpart on the second-stage training set as

$$\mathcal{L}(\Theta) := \mathbb{E}_{\mathbf{x}}[\ell(f_\Theta(\mathbf{x}), f^*(\mathbf{x}))] \quad \text{and} \quad \widehat{\mathcal{L}}(\Theta) := \frac{1}{n} \sum_{i=1}^n \ell(f_\Theta(\mathbf{x}_i), f^*(\mathbf{x}_i)) \quad (109)$$

To prove Theorem 2, i.e., upper bounding the population loss $\mathcal{L}(\Theta^{(T)})$ where $\Theta^{(T)}$ is returned by Algorithm 1, we first construct an ‘‘ideal’’ parameter candidate Θ^* based on our feature learning analysis (Appendix C.1) and analysis of standard random feature models as a transition term, which lies in a proper parameter class $\Theta^* \in \mathcal{F}$ with well-controlled $\|\Theta^*\|$ and yields small empirical error. Then, we can decompose the error $\mathcal{L}(\Theta^{(T)})$ as

$$\mathcal{L}(\Theta^{(T)}) = \underbrace{\widehat{\mathcal{L}}(\Theta^*)}_{\text{Lemma D.17 (subtle issue)}} + \underbrace{\widehat{\mathcal{L}}(\Theta^{(T)}) - \widehat{\mathcal{L}}(\Theta^*)}_{\text{Lemma D.18}} + \underbrace{\mathcal{L}(\Theta^{(T)}) - \widehat{\mathcal{L}}(\Theta^{(T)})}_{\text{Lemma D.19}}. \quad (110)$$

Next, by choosing a proper weight decay hyperparameter which relates to control the complexity of \mathcal{F} and leveraging the convexity of loss function with respect to \mathbf{a} , we show in Lemma D.18 that with $\widetilde{\mathcal{O}}(1)$ second-stage training steps, the second term can be bounded appropriately. Then, given that $\Theta^{(T)} \in \mathcal{F}$, we use standard Rademacher complexity analysis in Lemma D.19 with proper truncation arguments to bound the generalization error (the third term) via uniform convergence arguments. Now it suffices to construct the ideal Θ^* with a controllable norm and a small error.

Subtle issues in constructing Θ^* . An intricate problem arises when constructing Θ^* and estimating $\widehat{\mathcal{L}}(\Theta^*)$. We note that each neuron $w_j^{(T)}$ after training depends simultaneously on all neurons at initialization in a complicated manner. Consequently, the network after time T cannot be expressed as an average of conditionally independent terms, which hinders a direct Monte Carlo sampling analysis.

To overcome this, we introduce the so-called decoupled parameter group $\widetilde{\Theta} = (\mathbf{a}, \mathbf{b}, \widetilde{\mathbf{W}})$. In this construction, each row \widetilde{w}_j preserves its dependence on the corresponding initialization $w_j^{(0)}$ but is independent of the other neurons at initialization. Then, we can further decompose the error $\widehat{\mathcal{L}}(\Theta^*)$ as

$$\widehat{\mathcal{L}}(\Theta^*) = \underbrace{\widehat{\mathcal{L}}(\widetilde{\Theta}^*)}_{\text{Lemma D.15}} + \underbrace{\widehat{\mathcal{L}}(\Theta^*) - \widehat{\mathcal{L}}(\widetilde{\Theta}^*)}_{\text{Lemma D.16}}. \quad (111)$$

For the first term, we invoke standard Monte Carlo concentration arguments (Lemma D.14) to bound the squared loss, and then use the Lipschitz continuity of ℓ to extend the squared loss to the general loss $\widehat{\mathcal{L}}(\widetilde{\Theta}^*)$ (Lemma D.15). Moreover, we show that the empirical risks of Θ^* and $\widetilde{\Theta}^*$ differ only by a negligible amount (Lemma D.16), which is well-controlled by the Lipschitz continuity of the loss function ℓ and the small distance between $\|\mathbf{W}^{(T)} - \widetilde{\mathbf{W}}^{(T)}\|$ (Appendix C.1), ensuring that Θ^* and $\widetilde{\Theta}^*$ are sufficiently close.

Now by combining all the above analysis, we can give an upper bound for $\mathcal{L}(\Theta^{(T)})$ (Proposition D.20). Last, we manage to get a balance among the error terms with undetermined coefficients related to the monomial approximation property of the target function in Lemma D.21 and compute the time complexity explicitly in Lemma D.26, concluding our formal proof of Theorem 1².

²The formal version is stated in Theorem 2, where we explicitly specify all the hyperparameters.

D.2 EVENTS OF HIGH PROBABILITIES

We will first define the following events.

Briefly, we define \mathcal{E}_α and \mathcal{E}_β as events that the features are ensured to be well learned in the first stage, which has been proved in Appendix C.1, we then define \mathcal{E}_η and \mathcal{E}_ζ to control the range of the learned features and the second-stage training data for convenience of learning the outer polynomial link function g in the second training stage.

First, we define the events solely related to our first dataset. Recall from Lemma C.2 that the concentration error \mathbf{H}_ℓ admits the high-probability operator-norm bound $\|\mathbf{H}_\ell\|_{\text{op}} \leq \gamma := C\sqrt{d/n}$ where C is an absolute constant. We note that for sufficiently but moderately large n ($n \gtrsim dT_1^2$) the perturbation bound satisfies $\gamma \leq \frac{1}{2T_1}$. We will assume this regime throughout.

We define

$$\mathcal{E}_\alpha := \{\|\mathbf{H}_\ell\|_{\text{op}} \leq \gamma\} \cap \left(\bigcap_{i \in [n]} \left\{ \frac{1}{2} \|\mathbf{x}_i\|^2 \leq d \right\} \right) \quad (112)$$

and we know from Lemma C.2 and Lemma E.18 that $\mathbb{P}(\mathcal{E}_\alpha) \geq 1 - \mathcal{O}(d^{-2D})$. Next, for a fixed perturbation \mathbf{H}_ℓ , the projection control event (for \mathbf{w} drawn uniformly from the sphere $\epsilon_0 \mathbb{S}^{d-1}$) for a single feature is

$$\widetilde{\mathcal{E}}_\beta := \left\{ \frac{\|\Pi^*(\mathbf{H}_\ell^\ell \mathbf{w})\|}{\epsilon_0} \lesssim \|\mathbf{H}_\ell^\ell\|_{\text{op}} \sqrt{\frac{\log d}{d}} \quad \text{for all } \ell \leq T_1 \right\} \quad (113)$$

and we know from Lemma C.13 that $\mathbb{P}(\widetilde{\mathcal{E}}_\beta | \mathbf{H}_\ell) \geq 1 - \mathcal{O}(d^{-2D})$. Next, we consider the events related to the sampling of our features at initialization. Note that because we use symmetric initialization, only the first $m/2$ weight vectors are independent. For a fixed perturbation \mathbf{H}_ℓ and $1 \leq j \leq m/2$, the projection control event $\mathcal{E}_\beta^{(j)}$ is

$$\mathcal{E}_\beta^{(j)} := \left\{ \frac{\|\Pi^*(\mathbf{H}_\ell^\ell \mathbf{w}_j^{(0)})\|}{\epsilon_0} \lesssim \|\mathbf{H}_\ell^\ell\|_{\text{op}} \sqrt{\frac{\log d}{d}} \quad \text{for all } \ell \leq T_1 \right\} \quad (114)$$

and we know from Lemma C.13 that $\mathbb{P}(\mathcal{E}_\beta^{(j)} | \mathbf{H}_\ell) \geq 1 - \mathcal{O}(d^{-2D})$. Next, define the desired event as the intersection $\mathcal{E}_\beta := \bigcap_{j=1}^{m/2} \mathcal{E}_\beta^{(j)}$. Applying a union bound gives

$$\mathbb{P}(\mathcal{E}_\beta | \mathbf{H}_\ell) \geq 1 - \mathcal{O}(md^{-2D}) \quad (115)$$

From here, in this section we will write η_1 as η from time to time. Next, conditioning on the first dataset and the features, after the first stage of training, we define new events over the randomness of the second dataset. We need to introduce more notations first. From the previous section, we have the following formula for our final feature

$$\mathbf{w}_j^{(T_1)} = \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} \right) \mathbf{w}_j^{(0)} \quad (116)$$

and recall the ideal feature $\widehat{\mathbf{w}}_j^{(T_1)} := \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \Sigma_\ell^{T_1} \mathbf{w}_j^{(0)}$. For technical convenience, we next define the approximate feature to remove the influence of the interaction between neurons.

$$\widetilde{\mathbf{w}}_j^{(T_1)} = \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \left(\widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w}_j^{(0)} \quad (117)$$

To help the proof, we further define

$$\mathbf{h}_n(\mathbf{w}) := \frac{1}{\epsilon_0} \left(\widehat{\Sigma}_\ell^{T_1} \right) \mathbf{w} \quad \text{and} \quad \mathbf{h}(\mathbf{w}) := \frac{1}{\epsilon_0} \Sigma_\ell^{T_1} \mathbf{w} \quad (118)$$

and the residual $\mathbf{r}(\mathbf{w}) := \mathbf{h}_n(\mathbf{w}) - \mathbf{h}(\mathbf{w})$ for notational convenience. These approximate features are useful for our future analysis since $\widehat{\Sigma}_\ell$ does not depend on $\{\mathbf{w}_j\}_{j=1}^m$. Therefore, the approximate

dynamics then will decouple across neurons. We can bound the difference between the approximate features and true features afterwards with some additional efforts.

We next define the following event over the randomness of the second dataset conditioning on the first dataset for one feature

$$\tilde{\mathcal{E}}_\eta := \left\{ |\mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_i| \leq \|\mathbf{h}_n(\mathbf{w})\| \sqrt{2\iota} \quad \text{for all } n+1 \leq i \leq 2n \right\} \quad (119)$$

where $\iota = 4D \log d$. Via invoking the Gaussian tail bound plus a union bound over dataset, we have $\mathbb{P}(\tilde{\mathcal{E}}_\eta | \mathbf{H}_\ell, \mathbf{w}) \geq 1 - \mathcal{O}(nd^{-2D})$. Then we define

$$\mathcal{E}_\eta^{(j)} := \left\{ |\mathbf{h}_n(\mathbf{w}_j^{(0)})^\top \mathbf{x}_i| \leq \|\mathbf{h}_n(\mathbf{w}_j^{(0)})\| \sqrt{2\iota} \quad \text{for all } n+1 \leq i \leq 2n \right\} \quad (120)$$

where again $\iota = 4D \log d$ for each individual neuron. We have $\mathbb{P}(\mathcal{E}_\eta^{(j)} | \mathbf{H}_\ell, \{\mathbf{w}_j^{(0)}\}_{j=0}^m) \geq 1 - \mathcal{O}(nd^{-2D})$ due to the same reasons as above. Define the desired event as the intersection $\mathcal{E}_\eta := \bigcap_{j=1}^{m/2} \mathcal{E}_\eta^{(j)}$. Applying a union bound over these indices $j \in [m/2]$ gives $\mathbb{P}(\mathcal{E}_\eta | \mathbf{H}_\ell, \{\mathbf{w}_j^{(0)}\}_{j=0}^m) \geq 1 - \mathcal{O}(nmd^{-2D})$.

Last, we also need to define another event depending on $\{\mathbf{x}_i\}_{i=n+1}^{2n}$, that is, $\mathcal{E}_\zeta := \bigcap_{i \in [n]} \left\{ \frac{1}{2} \|\mathbf{x}_{i+n}\|^2 \leq d \right\} \bigcap_{i \in [n]} \left\{ |(\mathbf{x}_{i+n})_k| \lesssim \sqrt{\log d} \text{ for all } k \leq r \right\}$. By Lemma E.18 and Lemma E.19 we have that $\mathbb{P}(\mathcal{E}_\zeta) \geq 1 - \mathcal{O}(d^{-D})$.

Finally, we denote $\mathcal{E}_\mu := \mathcal{E}_\alpha \cap \mathcal{E}_\beta \cap \mathcal{E}_\eta \cap \mathcal{E}_\zeta$ from here. We will condition on this event throughout this section unless we particularly specify.

D.3 PROOF OF THEOREM 1

We first recall some basic facts from the results of the first stage. We have $\|\mathbf{H}_\ell\|_{\text{op}} \leq \gamma$ under \mathcal{E}_α , which gives rise to

$$\|\widehat{\Sigma}_\ell\|_{\text{op}} = \|\Sigma_\ell + \mathbf{H}_\ell\|_{\text{op}} \leq \|\Sigma_\ell\|_{\text{op}} + \|\mathbf{H}_\ell\|_{\text{op}} \leq 1 + \gamma \quad (121)$$

and taking the product over t yields $\left\| \widehat{\Sigma}_\ell^{\text{T}_1} \right\|_{\text{op}} \leq (1 + \gamma)^{\text{T}_1}$.

On those events we can obtain the crude uniform bound

$$\|\mathbf{h}_n(\mathbf{w})\| = \frac{\left\| \widehat{\Sigma}_\ell^{\text{T}_1} \mathbf{w} \right\|}{\epsilon_0} \leq \frac{(1 + \gamma)^{\text{T}_1}}{\epsilon_0} \|\mathbf{w}\| = (1 + \gamma)^{\text{T}_1} \quad (122)$$

that we might use later. Note that this estimation is uniform for all the features. We denote this upper bound by $V := (1 + \gamma)^{\text{T}_1}$.

Furthermore, we know by Corollary C.9, for \mathbf{w} drawn uniformly from the sphere $\epsilon_0 \mathbb{S}^{d-1}$, under $\mathcal{E}_\alpha \cap \tilde{\mathcal{E}}_\beta$, the residual term is dominated by

$$\|\mathbf{r}(\mathbf{w})\| \leq C \left(\gamma^{\text{T}_1} \sqrt{\frac{rD \log d}{d}} + \gamma^{\text{T}_1} \right) \quad (123)$$

where C is a universal constant. For notational convenience, we denote the right side as K . We note that this is also uniformly correct for all features under the full event \mathcal{E} .

Next, we need to do the following moment estimation for the residual term. The below is not straightforward since we need to take the expectation of \mathbf{w} .

Lemma D.1. *Under the high probability event \mathcal{E}_α , for any integer $1 \leq j \leq 4p$ we have*

$$\left(\mathbb{E}_{\mathbf{w}} \|\mathbf{\Pi}^* \mathbf{r}(\mathbf{w})\|^j \right)^{1/j} \lesssim K. \quad (124)$$

The next lemma computes the sample size needed to control K .

Lemma D.2. For each absolute constant c , in order to bound K as $d\kappa^{2T_1}K^2 \leq c^2$ it is just sufficient to require

$$n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2. \quad (125)$$

Lemma D.3. Given $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, under the high probability event of $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta$, we have

$$\|\mathbf{h}_n(\mathbf{w})\| \lesssim \sqrt{\log d} \frac{1}{\sqrt{d}}. \quad (126)$$

The proofs of the three lemmas above are provided in Appendices D.5.1, D.5.2 and D.5.3, respectively. Then the following corollary holds.

Corollary D.4. Let $\mathbf{w} \sim \epsilon_0 \text{Unif}(\mathbb{S}^{d-1})$. Given $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, under the high probability event of $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta$ we have

$$\max_{j \in [n+1, 2n]} \|\mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_j\| \lesssim \frac{\log d}{\sqrt{d}}. \quad (127)$$

Proof of Corollary D.4. From previous proof we can see under the high probability event of $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta$ we have

$$\max_{j \in [n+1, 2n]} |\mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_j| \leq \|\mathbf{h}_n(\mathbf{w})\| \sqrt{4D \log d}. \quad (128)$$

Also by Lemma D.3, as $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta$ happens, we have $\|\mathbf{h}_n(\mathbf{w})\| \lesssim \sqrt{\log d} \frac{1}{\sqrt{d}}$. Plug this in and we get our conclusion. \square

Before proving the results below, we introduce the notation used throughout. First, let $S^* = \text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_r\} \subset \mathbb{R}^d$ be the subspace spanned by the first r standard basis vectors, and recall $\Pi^* : \mathbb{R}^d \rightarrow S^*$ denote the orthogonal projector onto this subspace (i.e. Π^* preserves the first r coordinates and removes other coordinates). Then, for a symmetric tensor $\mathbf{T} \in \text{Sym}^{2k}(\mathbb{R}^d)$, write $\text{Mat}(\mathbf{T})$ for its canonical matricization as a linear operator on $\text{Sym}^k(\mathbb{R}^d)$; equivalently, $\text{Mat}(\mathbf{T})$ is the unique linear operator satisfying

$$\langle \mathbf{u}^{\otimes k}, \text{Mat}(\mathbf{T}) \mathbf{v}^{\otimes k} \rangle = \mathbf{T}(\mathbf{u}^{\otimes k}, \mathbf{v}^{\otimes k}) \quad \text{for all } \mathbf{u}, \mathbf{v} \in \mathbb{R}^d. \quad (129)$$

The next lemma will be important for our proof. We can now show that the obtained features $\mathbf{h}_n(\mathbf{w})$ are sufficiently expressive to allow us to efficiently represent any polynomial of degree p restricted to the principal subspace S^* .

Lemma D.5. Under the high probability event of \mathcal{E}_α , given the sample size $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, for any $k \leq p$, we have

$$\text{Mat}(\mathbb{E}_{\mathbf{w}}[(\Pi^* \mathbf{h}_n(\mathbf{w}))^{\otimes 2k}]) \lesssim \left(\frac{d}{\lambda_r^{2T_1}} \right)^{-k} \Pi_{\text{Sym}^k(S^*)} \quad (130)$$

where $\Pi_{\text{Sym}^k(S^*)}$ denotes the projection onto symmetric k tensors restricted to S^* .

In this lemma we will use monomial basis to construct the following approximations for $\langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle$ for any $k \leq p$ and any k tensor \mathbf{T} supported on $S^* = \{\mathbf{e}_1, \dots, \mathbf{e}_r\}$.

Lemma D.6. Assume the sample size $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$. Then under the high probability event \mathcal{E}_α , for each $k \leq p$ and any symmetric k tensor \mathbf{T} supported on S^* , there exists $z_{\mathbf{T}}(\mathbf{w})$ such that

$$\mathbb{E}_{\mathbf{w}} [z_{\mathbf{T}}(\mathbf{w})(\mathbf{h}_n(\mathbf{w})^\top \mathbf{x})^k] = \langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle \quad (131)$$

and we have the bounds

$$\mathbb{E}_{\mathbf{w}} [z_{\mathbf{T}}(\mathbf{w})^2] \lesssim \left(\frac{d}{\lambda_r^{2T_1}} \right)^k \|\mathbf{T}\|_F^2 \quad \text{and} \quad |z_{\mathbf{T}}(\mathbf{w})| \lesssim \left(\frac{d}{\lambda_r^{2T_1}} \right)^k \|\mathbf{T}\|_F \|\mathbf{h}_n(\mathbf{w})\|^k. \quad (132)$$

The proofs are provided in Appendices D.5.4 and D.5.5.

In the following, we show that bias reinitialization together with our learned features can be used to transform a wide class of activations σ into (or uniformly approximate by) the monomial x^p , thereby making it able to approximate low-dimensional polynomials.

Definition D.7. Let a be a Rademacher RV i.e. $\mathbb{P}(a = +1) = \mathbb{P}(a = -1) = \frac{1}{2}$ and let $b \sim \text{Unif}[-3, 3]$, independent from a .

We say the activation function $\sigma(\cdot)$ has the monomial approximation property with exponent $\beta \geq 0$ if for every integer $k \geq 0$ and every positive ρ there exists a bounded weight function $v_k^{(\rho)} : \{\pm 1\} \times [-3, 3] \rightarrow \mathbb{R}$ with

$$\|v_k^{(\rho)}\|_\infty \leq V_{k,\rho} \leq C_\sigma(k)\rho^{-\beta} \quad (133)$$

such that

$$\sup_{|z| \leq 1} \left| \mathbb{E}_{a,b} [v_k^{(\rho)}(a, b)\sigma(az + b)] - z^k \right| \leq \rho \quad (134)$$

The uniform approximation property in Definition D.7 is expected to hold for many standard activations, with different growth behavior of the required weights $V_{k,\rho}$. For example, we provably have $\beta = 0$ for the activation function below.

Definition D.8. The locally quadratic activation is defined as

$$\sigma(t) = \begin{cases} 2|t| - 1, & |t| \geq 1, \\ t^2, & |t| < 1. \end{cases} \quad (135)$$

In Section E.3 we show that $\beta = 0$ for locally quadratic function.

We first prove the following lemma in preparation for the proof that follows.

Lemma D.9. Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$. Then under the high probability event \mathcal{E}_μ we have for all features $j \in [m]$

$$\|\widetilde{\mathbf{w}}_j^{(T_1)}\| \leq 1 \quad \text{and} \quad \max_{i \in [n+1, 2n]} |(\widetilde{\mathbf{w}}_j^{(T_1)})^\top \mathbf{x}_i| \leq 1 \quad \text{and} \quad \|\mathbf{w}_j^{(T_1)}\| \lesssim 1 \quad (136)$$

Note that we already constructed the approximations for $\langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle$ using monomial basis, now we will utilize those bounds above to construct the following approximations for $\langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle$ for any $k \leq p$ and any k tensor \mathbf{T} supported on $S^* = \{\mathbf{e}_1, \dots, \mathbf{e}_r\}$ using activation function σ .

Lemma D.10. Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ also $T_1 = o(\log d)$. Then under the high probability event $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta$, Consider

$$f_{h_{\mathbf{T}}}(\mathbf{x}) := \mathbb{E}_{a,b,\mathbf{w}} \left[h_{\mathbf{T}}(a, b, \mathbf{w}) \sigma(\widetilde{\mathbf{w}}^{(T_1)\top} \mathbf{x} + b) \right] \quad (137)$$

where $\mathbf{w}^{(T_1)} = (-\eta a)^{T_1} \mathbf{h}_n(\mathbf{w})$. Here

$$h_{\mathbf{T}}(a, b, \mathbf{w}) := \frac{v_k^{(\epsilon_k)}(a, b) z_{\mathbf{T}}(\mathbf{w})}{(-\eta)^{kT_1}} \mathbf{1}_{\eta^{T_1} \|\mathbf{h}_n(\mathbf{w})\| \leq 1} \prod_{i=1}^n \mathbf{1}_{|\eta^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_{i+n}| \leq 1} \quad (138)$$

where $v_k^{(\epsilon_k)}$ is constructed from Definition D.7 and $z_{\mathbf{T}}(\mathbf{w})$ is the same from Corollary D.6. We further assume $\epsilon_k \geq d^{-\frac{D}{12\beta}}$.

Then, for any training point $\mathbf{x} \in \{\mathbf{x}_i\}_{i=n+1}^{2n}$, it holds that

$$|f_{h_{\mathbf{T}}}(\mathbf{x}) - \langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle| - \Delta_k \lesssim d^{-\frac{D}{8}} \|\mathbf{T}\|_F \quad (139)$$

where

$$\Delta_k := \frac{\epsilon_k}{\eta^{kT_1}} |\mathbb{E}_{\mathbf{w}} [z_{\mathbf{T}}(\mathbf{w})]| \lesssim \epsilon_k \kappa^{kT_1} \log^k d \|\mathbf{T}\|_F. \quad (140)$$

Additionally we have the following moment estimations where V_{k,ϵ_k} is from Definition D.7.

$$\mathbb{E}[h_{\mathbf{T}}^2] \lesssim V_{k,\epsilon_k}^2 (\kappa^{2T_1} \log^2 d)^k \|\mathbf{T}\|_F^2, \quad \|h_{\mathbf{T}}\|_\infty \lesssim V_{k,\epsilon_k} (\kappa^{2T_1} \log^2 d)^k \|\mathbf{T}\|_F. \quad (141)$$

The proofs of Lemmas D.9 and D.10 are provided in Appendices D.5.6 and D.5.7, respectively.

The previous lemmas establish the pointwise approximation guarantee for an arbitrary data point \mathbf{x} under the event $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta$. The next corollary will bound the empirical squared loss over the dataset $\{\mathbf{x}_i\}_{i=n+1}^{2n}$.

Corollary D.11. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ also $T_1 = o(\log d)$. Then under the event $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta$, for any $k \leq p$ and any k tensor \mathbf{T} supported on $S^* = \{\mathbf{e}_1, \dots, \mathbf{e}_r\}$, we have the empirical squared loss bound*

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{h_{\mathbf{T}}}(\mathbf{x}_i) - \langle \mathbf{T}, \mathbf{x}_i^{\otimes k} \rangle)^2 \lesssim \Delta_k^2 + d^{-D/4} \|\mathbf{T}\|_F^2. \quad (142)$$

where $f_{h_{\mathbf{T}}}$ is defined in Lemma D.10.

Proof of Corollary D.11. Let $e_i := f_{h_{\mathbf{T}}}(\mathbf{x}_i) - \langle \mathbf{T}, \mathbf{x}_i^{\otimes k} \rangle$. By Lemma D.10 we have, for every training point \mathbf{x}_i , $n+1 \leq i \leq 2n$,

$$|e_i| \lesssim \Delta_k + d^{-\frac{D}{8}} \|\mathbf{T}\|_F, \quad (143)$$

where

$$\Delta_k := \frac{\epsilon_k}{\eta^{kT_1}} |\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})]| \lesssim \epsilon_k \kappa^{kT_1} \log^k d \|\mathbf{T}\|_F. \quad (144)$$

Squaring and averaging, we get the following bound as desired.

$$\frac{1}{n} \sum_{i=n+1}^{2n} e_i^2 \lesssim \frac{1}{n} \sum_{i=1}^n (\Delta_k + d^{-\frac{D}{8}} \|\mathbf{T}\|_F)^2 \leq 2\Delta_k^2 + 2d^{-\frac{D}{4}} \|\mathbf{T}\|_F^2. \quad (145)$$

The proof is complete. \square

Definition D.12. *For accuracy parameters $\{\epsilon_k\}_{k \leq p}$ define*

$$V_{p,\epsilon} := \max_{0 \leq k \leq p} V_{k,\epsilon_k} \quad (146)$$

where V_{k,ϵ_k} is the uniform bound on $\|v_k^{(\epsilon_k)}\|_\infty$ from Definition D.7.

Notice that we already have empirical squared loss bounds for all \mathbf{T} supported on S^* in the Corollary D.11. Next we will derive empirical squared loss bounds for $f^*(\mathbf{x})$.

Lemma D.13. *We assume the event $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta$, the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ and $T_1 = o(\log d)$. Further, let $h(a, b, \mathbf{w}) := \sum_{k \leq p} h_{\mathbf{T}_k}(a, b, \mathbf{w})$, where $\{\mathbf{T}_k\}_{k=1}^p$ is defined from tensor expansion of f^* introduced in Lemma E.21 $f^*(\mathbf{x}) = \sum_{k \leq p} \langle \mathbf{T}_k, \mathbf{x}^{\otimes k} \rangle$ and $h_{\mathbf{T}_k}(\cdot)$ is given by Lemma D.10 with accuracy parameters $\{\epsilon_k\}_{k \leq p}$.*

Then we have the following empirical squared loss bound for $f^(\mathbf{x})$*

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_h(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-D/4}, \quad (147)$$

Where $f_h(\mathbf{x})$ is defined as

$$f_h(\mathbf{x}) := \mathbb{E}_{a,b,\mathbf{w}} \left[h(a, b, \mathbf{w}) \sigma(\tilde{\mathbf{w}}^{(T_1)\top} \mathbf{x} + b) \right] \quad (148)$$

Additionally we have the following moment estimation

$$\mathbb{E}[h^2] \lesssim V_{p,\epsilon}^2 (\kappa^{T_1})^{2p} (\log d)^{2p}, \quad \|h\|_\infty \lesssim V_{p,\epsilon} (\kappa^{T_1})^{2p} (\log d)^{2p}. \quad (149)$$

The proof is provided in Appendix D.5.8. We next introduce the following lemma using Monte-Carlo Approximation.

Lemma D.14. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ also $T_1 = o(\log d)$. Consider $\mathbf{a}_j^* := \frac{1}{m} h(a_j, b_j, \mathbf{w}_j^{(0)})$ where $h(\cdot)$ is defined in Lemma D.13 and $a_j, b_j, \mathbf{w}_j^{(0)}$ is initialized as in Algorithm 1. Denote $\tilde{\Theta}^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \tilde{\mathbf{W}}^{(T_1)})$.*

Then under the event \mathcal{E}_μ , we have

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-D/4} + \frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m} \quad (150)$$

and $\|\mathbf{a}^*\|^2 \lesssim \frac{1}{m} V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}$.

The proof is in Appendix D.5.9. For notational simplicity, given the parameter group $\Theta = (\mathbf{a}, \mathbf{b}, \mathbf{W})$, we denote the population loss and its empirical counterpart on the second-stage training set as

$$\mathcal{L}(\Theta) := \mathbb{E}_{\mathbf{x}}[\ell(f_{\Theta}(\mathbf{x}), f^*(\mathbf{x}))] \quad \text{and} \quad \hat{\mathcal{L}}(\Theta) := \frac{1}{n} \sum_{i=1}^n \ell(f_{\Theta}(\mathbf{x}_i), f^*(\mathbf{x}_i)) \quad (151)$$

respectively. Also write $\hat{\mathcal{L}}(\Theta)$ for $\hat{\mathcal{L}}(f_{\Theta})$.

In this lemma, we bound our general loss using the squared loss bound.

Lemma D.15. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ and under the event \mathcal{E}_μ , we have*

$$\hat{\mathcal{L}}(\tilde{\Theta}^*) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) \quad (152)$$

where Δ_k is defined from Lemma D.10.

The proof is presented in Appendix D.5.10.

Next, we consider the gap between the real features we learn $\mathbf{W}^{(T_1)}$ and the approximate features $\tilde{\mathbf{W}}^{(T_1)}$. Recall $\tilde{\Theta}^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \tilde{\mathbf{W}}^{(T_1)})$ and $\Theta^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})$.

Lemma D.16. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ and under the event \mathcal{E}_μ , we have*

$$\left| \hat{\mathcal{L}}(\Theta^*) - \hat{\mathcal{L}}(\tilde{\Theta}^*) \right| \lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}. \quad (153)$$

We provide the proof in Appendix D.5.11.

Combining the above two lemmas, we directly get the bound for $\hat{\mathcal{L}}(f_{\Theta^*})$.

Lemma D.17. *Under the same condition as the last lemma we have*

$$\hat{\mathcal{L}}(f_{\Theta^*}) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right). \quad (154)$$

The proof of this bound is provided in D.5.12.

The preceding lemmas construct a parameter Θ^* with small empirical risk $\hat{\mathcal{L}}(f_{\Theta^*})$. We now analyze the parameter output by the second-stage procedure,

$$\Theta^{(T)} = (\mathbf{a}^{(T)}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)}), \quad T = T_1 + T_2 \quad (155)$$

where $\mathbf{a}^{(T)}$ denotes the iterate after T iterations. We recall that only \mathbf{a} is updated in the second stage while \mathbf{b} and \mathbf{W} remain at their values at time T_1 . The next lemma gives an upper bound on the empirical risk $\hat{\mathcal{L}}(\Theta^{(T)})$ of this output.

For notation convenience, we define $J := L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right)$ and $U = \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}$.

Lemma D.18. We consider doing gradient descent on the ridge-regularized empirical loss

$$F(\mathbf{a}) := \widehat{\mathcal{L}}((\mathbf{a}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})) + \frac{\beta_2}{2} \|\mathbf{a}\|^2 \quad (156)$$

with the hyperparameters $\beta_2 := \frac{J}{U^2}$ and $\eta_2 = \frac{U^2}{L\sqrt{m}U^2+2J}$.

Then for any $T_2 \gtrsim \frac{U^2 L m}{J} \log(Um)$ also under the event \mathcal{E}_μ , our second stage algorithm returns $\mathbf{a}^{(T)}$ such that for $\Theta^{(T)} := (\mathbf{a}^{(T)}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})$

$$\widehat{\mathcal{L}}(\Theta^{(T)}) \lesssim J \quad \text{and} \quad \|\mathbf{a}^{(T)}\| \lesssim U. \quad (157)$$

The proof is provided in Appendix D.5.13. The following lemma bounds the gap between empirical loss and population loss for $\Theta^{(T)}$.

Lemma D.19. Under the event \mathcal{E} , where \mathcal{E} is defined in Definition D.37 and the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ we have

$$\left| \mathcal{L}(f_{\Theta^{(T)}}) - \widehat{\mathcal{L}}(f_{\Theta^{(T)}}) \right| \lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+1}}{n}}. \quad (158)$$

The proof is provided in Appendix D.5.14. Then the following proposition combines the previous lemmas and gives a quick proof map of bounding the excess population risk $\mathcal{L}(f_{\Theta^{(T)}})$.

Proposition D.20. On the event \mathcal{E} and the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ the excess population risk satisfies

$$\mathcal{L}(f_{\Theta^{(T)}}) - \mathcal{L}(f^*) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 \kappa^{4p T_1} (\log d)^{4p+1}}{n}} + \sqrt{\frac{V_{p,\epsilon}^2 \kappa^{4p T_1} (\log d)^{4p+2}}{m}} \right). \quad (159)$$

Proof of Proposition D.20. Note that to upper bound the population loss $\mathcal{L}(\Theta^{(T)})$ of the parameter group $\Theta^{(T)}$ returned by Algorithm 1, we first construct an ‘‘ideal’’ parameter group Θ^* which lies in a proper parameter class $\Theta^* \in \mathcal{F}$ (small $\|\Theta^*\|$) and yields small empirical error.

So the first several lemmas this section give an upper bound for $\widehat{\mathcal{L}}(f_{\Theta^*})$ through $\widetilde{\Theta}^*$, which is introduced to decouple the dependence

$$\begin{aligned} \underbrace{\widehat{\mathcal{L}}(f_{\Theta^*})}_{\text{Lemma D.17}} &= \underbrace{\left[\widehat{\mathcal{L}}(f_{\Theta^*}) - \widehat{\mathcal{L}}(f_{\widetilde{\Theta}^*}) \right]}_{\text{Lemma D.16}} + \underbrace{\left[\widehat{\mathcal{L}}(f_{\widetilde{\Theta}^*}) - \widehat{\mathcal{L}}(f^*) \right]}_{\text{Lemma D.15}} \\ &\lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m}} + L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) \\ &\lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right). \end{aligned} \quad (160)$$

And recall that we denote $J := L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right)$.

Also note that we can decompose the error $\mathcal{L}(\Theta^{(T)})$ in the equation that follows

$$\begin{aligned} \mathcal{L}(f_{\Theta^{(T)}}) - \widehat{\mathcal{L}}(f_{\Theta^*}) &= \underbrace{\left[\mathcal{L}(f_{\Theta^{(T)}}) - \widehat{\mathcal{L}}(f_{\Theta^{(T)}}) \right]}_{\text{Lemma D.19}} + \underbrace{\left[\widehat{\mathcal{L}}(f_{\Theta^{(T)}}) - \widehat{\mathcal{L}}(f_{\Theta^*}) \right]}_{\text{Lemma D.18}} \\ &\lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+1}}{n}} + J \\ &\lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 \kappa^{4p T_1} (\log d)^{4p+1}}{n}} + \sqrt{\frac{V_{p,\epsilon}^2 \kappa^{4p T_1} (\log d)^{4p+2}}{m}} \right). \end{aligned} \quad (161)$$

Combining this two yields

$$\mathcal{L}(f_{\Theta(\tau)}) - \mathcal{L}(f^*) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+1}}{n}} + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) \quad (162)$$

which is precisely the claimed bound. This completes the proof. \square

Back to the main proof, in the previous proposition, we have the generalization bound

$$\mathcal{L}(f_{\Theta(\tau)}) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+1}}{n}} + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) \quad (163)$$

We note that the bound is depending on both $\{\Delta_k\}_{k=0}^p$ and $V_{p,\epsilon}$, which are left unspecified. Fixing $\beta > 0$, we notice that Δ_k has the bound relying on ϵ_k which is

$$\begin{aligned} \Delta_k &= \frac{\epsilon_k}{\eta^{kT_1}} |\mathbb{E}_{\mathbf{w}} [z_{T_k}(\mathbf{w})]| \\ &\lesssim \epsilon_k \kappa^{kT_1} \log^k d \|\mathbf{T}_k\|_F \lesssim \epsilon_k \kappa^{pT_1} \log^p d. \end{aligned} \quad (164)$$

On the other hand, we observe that $V_{k,\epsilon_k} \lesssim \epsilon_k^{-\beta}$. Thus we need to choose proper $\{\epsilon_k\}_{k=0}^p$ to keep the balance of V_{k,ϵ_k} and Δ_k . For notation convenience, let

$$A := \kappa^{2pT_1} (\log d)^{2p+1}, \quad C_m := \sqrt{\frac{1}{m}}, \quad C_n := \sqrt{\frac{1}{n}}. \quad (165)$$

Thus the ϵ -dependent contribution is upper bounded (up to absolute constants) by

$$\Phi(\{\epsilon_k\}_{k \leq p}) := A \left(\sum_{k \leq p} \epsilon_k + (C_m + C_n) V_{p,\epsilon} \right) \quad (166)$$

Specifically, when $\beta = 0$, from Definition D.7, we know that the approximation precision ρ can be zero, which indicates that $\epsilon_k = 0$ for all $\{\epsilon_k\}_{k=0}^p$ and also $V_{p,\epsilon} \lesssim 1$. Thus

$$\Phi = A((C_m + C_n) V_{p,\epsilon}) \lesssim A(C_m + C_n) \quad (167)$$

Furthermore, for $\beta > 0$, plugging in $V_{k,\epsilon_k} \lesssim \epsilon_k^{-\beta}$, we have

$$\Phi(\{\epsilon_k\}_{k \leq p}) := A \left(\sum_{k \leq p} \epsilon_k + (C_m + C_n) \max_{k \leq p} \{\epsilon_k^{-\beta}\} \right) \quad (168)$$

The following lemma gives the minimum value of Φ .

Lemma D.21. Fix $\beta > 0$. Consider Φ as a function of $\{\epsilon_k\}_{k=0}^p$. Then the minimizer is unique and attained at the symmetric point

$$\epsilon_0^* = \dots = \epsilon_p^* = \left(\frac{\beta(C_m + C_n)}{p+1} \right)^{\frac{1}{\beta+1}}. \quad (169)$$

Moreover, it holds that

$$\min_{\{\epsilon_k > 0\}} \Phi = A \cdot (\beta + 1) \beta^{-\frac{\beta}{\beta+1}} (p+1)^{\frac{\beta}{\beta+1}} (C_m + C_n)^{\frac{1}{\beta+1}} \lesssim A(C_m + C_n)^{\frac{1}{\beta+1}}. \quad (170)$$

The proof is provided in Appendix D.5.15. The following corollary is direct from the previous statement.

Corollary D.22. When $\beta = 0$, we have $\inf \Phi = A(C_m + C_n)$.

Now we state the main theorem. For the theorem stated below, we recall that we assume $T_1 = o(\log d)$, $n \gtrsim d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ in addition to the assumptions in Section 2 and Section 3.

Theorem 2 (Main Theorem). *Under the assumptions above, there exists a set of hyperparameters*

$$\eta_1 = \frac{1}{C_\eta} \left(\frac{d}{4D^2 \log^2 d} \right)^{\frac{1}{2T_1}}, \quad \beta_1 = \frac{1}{\eta_1}, \quad \eta_2 = \frac{U^2}{L\sqrt{m}U^2 + 2J}, \quad \beta_2 = \frac{J}{U^2} \quad (171)$$

such that for any $T_2 \gtrsim \frac{U^2 L m}{J} \log(Um)$, the population risk satisfies

$$\mathbb{E}_{\mathbf{x}}[\ell(f_{\Theta(T)}(\mathbf{x}), y)] \lesssim \kappa^{2pT_1} (\log d)^{2p+1} \left(\sqrt{\frac{1}{m}} + \sqrt{\frac{1}{n}} \right)^{\frac{1}{\beta+1}} \quad (172)$$

under the high probability event \mathcal{E} , which happens with probability at least $1 - \mathcal{O}(d^{-D/2})$, where C_η is an absolute constant from Lemma C.6 and J, U are quantities defined in D.3.

We further estimate sample and time complexity in the next lemmas, showing that they are both nearly optimal.

Lemma D.23. *Under the setting $T_1 := \lceil \sqrt{\frac{\log d}{\log \kappa}} \rceil$ we have*

$$T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2 \lesssim \kappa^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}} (\log d)^2 \quad (173)$$

The proof is provided in Appendix D.5.16. Note that from Theorem 2, for $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ the main result applies. Since our sample-complexity upper bound coincides with that threshold, the corollary follows immediately.

Corollary D.24 (Sample Complexity). *The sample complexity is $\tilde{\Theta}(\kappa^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}})$, where $d^{2\sqrt{\frac{\log \kappa}{\log d}}}$ is subpolynomial in d .*

Next we first estimate T_2 under particular m, n choices and other hyperparameters.

Lemma D.25. *Fix the set of hyperparameters*

$$T_1 := \lceil \sqrt{\frac{\log d}{\log \kappa}} \rceil, \quad \beta_2 := \frac{J}{U^2}, \quad \eta_2 = \frac{U^2}{L\sqrt{m}U^2 + 2J}, \quad T_2 = \frac{U^2 L m}{J} \log(Um) \quad (174)$$

where J and U is defined as

$$J := L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) \text{ and } U = \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}. \quad (175)$$

$\{\Delta_k\}_{k \leq p}$ is defined in Lemma D.10 and $V_{p,\epsilon}$ is defined in Definition D.12. Note that $\{\epsilon_k\}_{k=0}^p$ was fixed in Lemma D.21, conclusively J and U are also fixed.

Furthermore, fix the sample size and net width as $m = (\log d)^{4p(\beta+1)+1} d^{4p(\beta+1)\sqrt{\frac{\log \kappa}{\log d}}}$, $n = C\kappa^2 (\log d)^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}}$ where C is a large universal constant. Then under the above setting,

$$T_2 \lesssim \kappa^{5p} (\log d)^{8p\beta+13p+2} d^{(8p\beta+13p)\sqrt{\frac{\log \kappa}{\log d}}}. \quad (176)$$

By ignoring the logarithmic terms, we can get $T_2 = \tilde{\mathcal{O}}(\kappa^{5p} d^{(8p\beta+13p)\sqrt{\frac{\log \kappa}{\log d}}})$.

The proof is provided in Appendix D.5.17.

We note that the total time complexity is upper bound by $nmd \cdot T$ up to absolute constants, so plugging n and m , the next lemma follows directly.

Lemma D.26 (Time complexity). *Under the same condition as the last lemma, we have $T_{\text{complexity}} \lesssim \kappa^{5p+2} (\log d)^{12p\beta+17p+5} d^{2+(12p\beta+17p+2)\sqrt{\frac{\log \kappa}{\log d}}}$*

The proof is provided in Appendix D.5.18.

D.4 PROOF UNDER $\beta = 0$ AND ASSUMPTION 3'

We now extend our analysis to a more general loss class (Assumption 3') under an additional assumption $\beta = 0$. For clarity, we provide a complete proof here.

Under Assumption 3' and $\beta = 0$, Lemma D.5.1 to Lemma D.14 still hold. In particular, given that $\beta = 0$ we can conclude that $\Delta_k = 0$ for all $k \leq p$ and $V_{p,\epsilon}$ is bounded by universal constant. We restate this as the following lemma, which is parallel to Lemma D.14 in the previous section.

Lemma D.27. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$ also $T_1 = o(\log d)$. Consider $a_j^* := \frac{1}{m} h(a_j, b_j, \mathbf{w}_j^{(0)})$ where $h(\cdot)$ is defined in Lemma D.13 and $a_j, b_j, \mathbf{w}_j^{(0)}$ is initialized as in Algorithm 1. Denote $\tilde{\Theta}^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \tilde{\mathbf{W}}^{(T_1)})$.*

Then under the event \mathcal{E}_μ , we have

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \lesssim d^{-D/4} + \frac{(\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m} \quad (177)$$

and $\|\mathbf{a}^*\|^2 \lesssim \frac{1}{m} (\kappa^{T_1})^{4p} (\log d)^{4p}$.

We sequentially present the parallel versions of the other theorems in the previous section. In the following lemma, we bound our general loss using the squared loss bound. For notation convenience, we denote $L_{\text{aug}} := L \kappa^{2pT_1} (\log d)^{2p}$. It can be observed that the original conclusion remains almost unchanged, except that L is replaced by the new L_{aug} .

Lemma D.28. *Under $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ and under the event \mathcal{E}_μ , we have*

$$\hat{\mathcal{L}}(\tilde{\Theta}^*) \lesssim L_{\text{aug}} \sqrt{\frac{\kappa^{4pT_1} (\log d)^{4p+2}}{m}} \quad (178)$$

Proof. Under Assumption 3', for any $z, z', y \in \mathbb{R}$,

$$|\ell(z, y) - \ell(z', y)| \leq L(1 + |z| + |z'| + |y|)|z - z'|. \quad (179)$$

Combining with the above Lemma, we have

$$\begin{aligned} \hat{\mathcal{L}}(\tilde{\Theta}^*) &= \frac{1}{n} \sum_{i=1}^n \ell(f_{\tilde{\Theta}^*}(\mathbf{x}_i), f^*(\mathbf{x}_i)) = \frac{1}{n} \sum_{i=1}^n (\ell(f_{\tilde{\Theta}^*}(\mathbf{x}_i), f^*(\mathbf{x}_i)) - \ell(f^*(\mathbf{x}_i), f^*(\mathbf{x}_i))) \\ &\lesssim \frac{1}{n} \sum_{i=1}^n L(1 + |f_{\tilde{\Theta}^*}(\mathbf{x}_i)| + |f^*(\mathbf{x}_i)|) |f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i)| \end{aligned} \quad (180)$$

Furthermore, by the same argument as in Lemma D.9, we have $|(\mathbf{w}_j^{(T)})^\top \mathbf{x}| \lesssim 1$ for all j , thus we have

$$\left\| \sigma(\tilde{\mathbf{W}}^{(T_1)} \mathbf{x} + \mathbf{b}) \right\| \lesssim \sqrt{m} \quad (181)$$

$$|f_{\tilde{\Theta}^*}(\mathbf{x}_i)| = \mathbf{a}^* \sigma(\tilde{\mathbf{W}}^{(T_1)} \mathbf{x} + \mathbf{b}) \lesssim U \sqrt{m}. \quad (182)$$

Also note that $U \lesssim \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}$ and $|f^*(\mathbf{x}_i)| \lesssim (\log d)^p$. Thus by Jensen's inequality and the Lipschitz property of ℓ , we have

$$\hat{\mathcal{L}}(\tilde{\Theta}^*) \leq \frac{1}{n} \sum_{i=1}^n L(1 + |f_{\tilde{\Theta}^*}(\mathbf{x}_i)| + |f^*(\mathbf{x}_i)|) |f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i)| \quad (183)$$

$$\lesssim L(\kappa^{T_1})^{2p} (\log d)^{2p} \frac{1}{n} \left(\sum_{i=1}^n |f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i)| \right) \lesssim L_{\text{aug}} \sqrt{\frac{\kappa^{4pT_1} (\log d)^{4p+2}}{m}}. \quad (184)$$

The proof is complete. \square

Next, we consider the gap between the real features we learn $\mathbf{W}^{(T_1)}$ and the approximate features $\widehat{\mathbf{W}}^{(T_1)}$. Recall $\widehat{\Theta}^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \widehat{\mathbf{W}}^{(T_1)})$ and $\Theta^* = (\mathbf{a}^*, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})$.

Lemma D.29. *Under the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ and under the event \mathcal{E}_μ , we have*

$$\left| \widehat{\mathcal{L}}(\Theta^*) - \widehat{\mathcal{L}}(\widehat{\Theta}^*) \right| \lesssim L_{\text{aug}} \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}. \quad (185)$$

We omit the proof here, since the only change is to change L to L_{aug} .

Combining the above two lemmas, we directly get the bound for $\widehat{\mathcal{L}}(f_{\Theta^*})$.

Lemma D.30. *Under the same condition as Lemma D.28, we have $\widehat{\mathcal{L}}(f_{\Theta^*}) \lesssim L_{\text{aug}} \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}}$.*

We now analyze the parameter output by the second-stage procedure,

$$\Theta^{(T)} = (\mathbf{a}^{(T)}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)}) \quad T = T_1 + T_2 \quad (186)$$

in the new setting. Define $J' = L_{\text{aug}} \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}}$ and $U' = \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}$. This is similar to the above section, but simpler under the new setting.

Lemma D.31. *We consider doing gradient descent on the ridge-regularized empirical loss*

$$F(\mathbf{a}) := \widehat{\mathcal{L}}((\mathbf{a}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})) + \frac{\beta_2}{2} \|\mathbf{a}\|^2 \quad (187)$$

with the hyperparameters $\beta_2 := \frac{J'}{U'^2}$ and $\eta_2 = \frac{U'^2}{LmU'^2 + J'}$.

Then for any $T_2 \gtrsim \frac{m}{\log d} \log(m\kappa^T \log d)$ also under the event \mathcal{E}_μ , our second stage algorithm returns $\mathbf{a}^{(T)}$ such that for $\Theta^{(T)} := (\mathbf{a}^{(T)}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})$

$$\widehat{\mathcal{L}}(\Theta^{(T)}) \lesssim J' \quad \text{and} \quad \|\mathbf{a}^{(T)}\| \lesssim U'. \quad (188)$$

Proof. Consider the training objective $F(\mathbf{a}) = \widehat{\mathcal{L}}((\mathbf{a}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})) + \frac{1}{2}\beta_2\|\mathbf{a}\|^2$ with $\beta_2 = \frac{J'}{U'^2}$.

Let

$$\mathbf{a}^{(\infty)} := \arg \min_{\mathbf{a} \in \mathbb{R}^m} F(\mathbf{a}) \quad (189)$$

be the minimizer of F .

We further denote ϵ_0 is the prescribed optimization tolerance $\|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \leq \epsilon_0$. Furthermore, let $\epsilon_0 = \frac{1}{mL_{\text{aug}}}$.

By optimality of $\mathbf{a}^{(\infty)}$ and the *constructed* comparator \mathbf{a}^* we have

$$F(\mathbf{a}^{(\infty)}) \leq F(\mathbf{a}^*) = \widehat{\mathcal{L}}(\mathbf{a}^*) + \frac{\beta_2}{2} \|\mathbf{a}^*\|_2^2. \quad (190)$$

Here we only use that \mathbf{a}^* is a feasible point with $\|\mathbf{a}^*\| \lesssim U'$ and $\widehat{\mathcal{L}}(\mathbf{a}^*) \lesssim J'$ (Lemma D.30).

On the other hand,

$$F(\mathbf{a}^{(\infty)}) = \widehat{\mathcal{L}}(\mathbf{a}^{(\infty)}) + \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|^2 \geq \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|^2, \quad (191)$$

so combining with (331) gives

$$\|\mathbf{a}^{(\infty)}\|^2 \leq \|\mathbf{a}^*\|^2 + \frac{2}{\beta_2} \widehat{\mathcal{L}}(\mathbf{a}^*) \lesssim U'^2 + \frac{2J'}{\beta_2} = U'^2 + \frac{2J'}{J'/U'^2} \lesssim U'^2. \quad (192)$$

Thus

$$\|\mathbf{a}^{(\infty)}\| \lesssim U'. \quad (193)$$

and also

$$\|\mathbf{a}^{(\text{T})}\| \leq \|\mathbf{a}^{(\infty)}\| + \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \leq \epsilon_0 + U' \lesssim U' \quad (194)$$

Moreover, from (331) and (334),

$$\widehat{\mathcal{L}}(\mathbf{a}^{(\infty)}) = F(\mathbf{a}^{(\infty)}) - \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|_2^2 \leq F(\mathbf{a}^*) \lesssim J' + \frac{\beta_2}{2} U'^2 \lesssim J'. \quad (195)$$

Thus we have

$$\begin{aligned} |F(\mathbf{a}^{(\text{T})}) - F(\mathbf{a}^{(\infty)})| &\lesssim \sup_{i \in [n+1, 2n]} \left\{ L_{\text{aug}} |\mathbf{a}^{(\text{T})\top} \sigma(\mathbf{W}^{(\text{T})} \mathbf{x}_i + \mathbf{b}^{(\text{T})}) - \mathbf{a}^{(\infty)\top} \sigma(\mathbf{W}^{(\text{T})} \mathbf{x}_i + \mathbf{b}^{(\text{T})})| \right\} \\ &\quad + \beta_2 (\|\mathbf{a}^{(\text{T})}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \\ &\lesssim L_{\text{aug}} \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \cdot \sup_{i \in [n+1, 2n]} \|\sigma(\mathbf{W}^{(\text{T})} \mathbf{x}_i + \mathbf{b}^{(\text{T})})\| \\ &\quad + \beta_2 (\|\mathbf{a}^{(\text{T})}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \end{aligned} \quad (196)$$

By previous proof $\|\sigma(\mathbf{W}^{(\text{T})} \mathbf{x}_i + \mathbf{b}^{(\text{T})})\|_\infty$ can be bounded by an absolute constant for all i under event \mathcal{E} . Thus we can have a loose ℓ^2 norm bound and

$$\begin{aligned} |F(\mathbf{a}^{(\text{T})}) - F(\mathbf{a}^{(\infty)})| &\lesssim L_{\text{aug}} \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \cdot \sup_{i \in [n+1, 2n]} \|\sigma(\mathbf{W}^{(\text{T})} \mathbf{x}_i + \mathbf{b}^{(\text{T})})\| \\ &\quad + \beta_2 (\|\mathbf{a}^{(\text{T})}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \\ &\lesssim L_{\text{aug}} \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \sqrt{m} + \beta_2 U' \|\mathbf{a}^{(\text{T})} - \mathbf{a}^{(\infty)}\| \\ &\lesssim (L_{\text{aug}} \sqrt{m} + \beta_2 U') \epsilon_0 \lesssim J' \end{aligned} \quad (197)$$

Combining with the same argument as above,

$$\widehat{\mathcal{L}}(\mathbf{a}^{(\text{T})}) \lesssim J'. \quad (198)$$

By Assumption 3' on the loss function ℓ , the function $F(\mathbf{a}) - \frac{\beta_2}{2} \|\mathbf{a}\|^2$ is convex and satisfies the gradient Lipschitz condition

By the similar arguments as in the proof of Lemma D.18, the lip constant for F satisfy $L_{\text{lip}} \lesssim Lm + J'/U'^2$.

Therefore, by Lemma E.13, if the step size $\eta_2 \leq \frac{1}{L_{\text{lip}}}$, gradient descent will approximate $\mathbf{a}^{(\infty)}$ to arbitrary accuracy in

$$T_2 \gtrsim \frac{1}{\eta_2 \beta_2} \log \left(\frac{\|\mathbf{a}^{(\text{T}_1)} - \mathbf{a}^{(\infty)}\|^2}{\epsilon_0^2} \right) \quad (199)$$

steps. Note that when the step size $\eta_2 = \frac{1}{Lm + J'/U'^2}$

$$\frac{1}{\eta_2 \beta_2} \lesssim \frac{U'^2 Lm}{J'} \quad (200)$$

Using the assumption that $m \lesssim d^{\frac{D}{4}}$ and $\|\mathbf{a}^{(\infty)}\| \lesssim U$, $\|\mathbf{a}^{(\text{T}_1)}\| \lesssim 1$, we can further simplify:

$$\frac{1}{\eta_2 \beta_2} \log \left(\frac{\|\mathbf{a}^{(\text{T}_1)} - \mathbf{a}^{(\infty)}\|^2}{\epsilon_0^2} \right) \lesssim \frac{U'^2 Lm}{J'} \log(U' m L_{\text{aug}}) \lesssim \frac{m}{\log d} \log(m \kappa^T \log d) \quad (201)$$

Thus, we conclude that gradient descent will approximate $\mathbf{a}^{(\infty)}$ to arbitrary accuracy in at most

$$T_2 \gtrsim \frac{m}{\log d} \log(m \kappa^T \log d) \quad (202)$$

steps. The proof is complete. \square

The following lemma bounds the gap between empirical loss and population loss.

Lemma D.32. *Under the event \mathcal{E} , where \mathcal{E} is defined in Definition D.37 and the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ we have*

$$\left| \mathcal{L}(f_{\Theta(T)}) - \widehat{\mathcal{L}}(f_{\Theta(T)}) \right| \lesssim L_{\text{aug}} \sqrt{\frac{(\kappa^{T_1})^{4p} (\log d)^{4p+1}}{n}}. \quad (203)$$

Proof. Use the same argument as the proof of Lemma D.19 in Appendix D.5.14 and use the same truncation, we can get

$$\sup_{f \in \mathcal{F}} \left| \widehat{\mathcal{L}}^{\mathcal{M}}(f) - \mathcal{L}^{\mathcal{M}}(f) \right| \lesssim L_{\text{aug}} U' \sqrt{\frac{m \log d}{n}}. \quad (204)$$

We notice that under the event \mathcal{E}_μ , \mathcal{M} holds for all \mathbf{x}_i , $i \in [n+1, 2n]$, thus $\widehat{\mathcal{L}}^{\mathcal{M}}(f) = \widehat{\mathcal{L}}(f)$.

For the other term, we have

$$\begin{aligned} \mathcal{L}^{\mathcal{M}}(f) - \mathcal{L}(f) &= \mathbb{E}[\ell(f(\mathbf{x}), y)] - \mathbb{E}[\ell^{\mathcal{M}}(f(\mathbf{x}), y; \mathbf{x})] = \mathbb{E}[\ell(f(\mathbf{x}), y) \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] \\ &\leq L \mathbb{E}[(1 + |f(\mathbf{x})| + |f^*(\mathbf{x})|) |f(\mathbf{x}) - f^*(\mathbf{x})| \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] \\ &\lesssim L \left(\mathbb{E}[|f(\mathbf{x})|^2 \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] + \mathbb{E}[|f^*(\mathbf{x})|^2 \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] \right) \\ &\lesssim L \left(\sqrt{\mathbb{E}[f(\mathbf{x})^4 \mathbb{P}\{\mathbf{x} \notin \mathcal{M}\}] + \sqrt{\mathbb{E}[(f^*(\mathbf{x}))^4 \mathbb{P}\{\mathbf{x} \notin \mathcal{M}\}]} \right). \end{aligned} \quad (205)$$

Recall that it has been assumed that the activation function has bounded third derivatives as in Assumption 2. Therefore, there exists an absolute constant $C_\sigma > 0$ such that

$$|\sigma(u)| \leq C_\sigma (1 + |u|^3) \quad \text{for all } u \in \mathbb{R}. \quad (206)$$

Via a crude moment estimation bound (plus some hypercontractivity arguments), we can see the right side is $\lesssim LU' d^{-\frac{D}{8}}$. Thus, we have the following final estimation

$$\sup_{f \in \mathcal{F}} \left| \widehat{\mathcal{L}}(f) - \mathcal{L}(f) \right| \lesssim L_{\text{aug}} U' \sqrt{\frac{m \log d}{n}} + LU' d^{-D/8}. \quad (207)$$

Since we have assume $n \leq d^{D/4}$, we have $\sup_{f \in \mathcal{F}} \left| \widehat{\mathcal{L}}(f) - \mathcal{L}(f) \right| \lesssim L_{\text{aug}} U' \sqrt{\frac{m \log d}{n}} + LU' d^{-\frac{D}{8}}$. Plug in the formula for U and we get our desired result. \square

We then have the final proposition.

Proposition D.33. *On the event \mathcal{E} and the sample-size condition $n \gtrsim T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, also $T_1 = o(\log d)$ the excess population risk satisfies*

$$\mathcal{L}(f_{\Theta(T)}) - \mathcal{L}(f^*) \lesssim L_{\text{aug}} \kappa^{2pT_1} (\log d)^{2p+1} \left(\sqrt{\frac{1}{n}} + \sqrt{\frac{1}{m}} \right). \quad (208)$$

Roadmap to remove $\beta = 0$ assumption. We note that $\beta = 0$ assumption is not essential and can be removed via a careful analysis. To be concrete, one can keep the Δ_k and $V_{p,\epsilon}$ terms in final generalization bound under the weakened assumption 3'. Next, we perform another round of balancing those terms. The minimizer shall be slightly different and a bit technical troublesome, but the final result remains the same up to logarithmic factors.

D.5 PROOF OF SUPPORTING LEMMAS IN APPENDIX D.3

D.5.1 PROOF OF LEMMA D.1

Proof. We split the expectation over \mathbf{w} into the good event $\widetilde{\mathcal{E}}_\beta$ and its complement

$$\mathbb{E}[\|\mathbf{\Pi}^* \mathbf{r}\|^j] = \mathbb{E}[\|\mathbf{\Pi}^* \mathbf{r}\|^j \mathbf{1}_{\widetilde{\mathcal{E}}_\beta}] + \mathbb{E}[\|\mathbf{\Pi}^* \mathbf{r}\|^j \mathbf{1}_{\neg \widetilde{\mathcal{E}}_\beta}]. \quad (209)$$

We first deal with the first term.

By the high-probability bound, we have $\|\mathbf{\Pi}^* \mathbf{r}(\mathbf{w})\| \leq K$ for all \mathbf{w} , so

$$\mathbb{E} \left[\|\mathbf{\Pi}^* \mathbf{r}\|^j \mathbf{1}_{\widetilde{\mathcal{E}}_\beta} \right] \leq K^j \mathbb{P}(\widetilde{\mathcal{E}}_\beta) \leq K^j. \quad (210)$$

Next for the second one, we only need to invoke the trivial operator norm bound. By definition we have $\|\widehat{\Sigma}_\ell^{\mathbf{T}_1}\| \leq (1 + \gamma)^{\mathbf{T}_1}$. Thus

$$\|\mathbf{r}(\mathbf{w})\| \leq \|\widehat{\Sigma}_\ell^{\mathbf{T}_1}\| + \|\Sigma_\ell\|^{\mathbf{T}_1} \leq 2M \quad (211)$$

Hence

$$\mathbb{E} \left[\|\mathbf{\Pi}^* \mathbf{r}\|^j \mathbf{1}_{-\widetilde{\mathcal{E}}_\beta} \right] \leq (2M)^j \mathbb{P}(-\widetilde{\mathcal{E}}_\beta) \lesssim (2M)^j r \mathbf{T}_1 d^{-D} \quad (212)$$

Therefore we have

$$\mathbb{E} \|\mathbf{\Pi}^* \mathbf{r}\|^j \lesssim K^j + (2M)^j r \mathbf{T}_1 d^{-D} \quad \text{and} \quad [\mathbb{E} \|\mathbf{\Pi}^* \mathbf{r}\|^j]^{1/j} \lesssim K \left[1 + \left(\frac{2M}{K} \right)^j r \mathbf{T}_1 d^{-D} \right]^{1/j} \quad (213)$$

To bound the right side, further note

$$\frac{2M}{K} = \frac{2(1 + \gamma)^{\mathbf{T}_1}}{\mathbf{T}_1 \gamma \sqrt{2rD \log d/d}} = \frac{2(1 + \gamma)^{\mathbf{T}_1}}{\mathbf{T}_1 \gamma \sqrt{2rD \log d/d}} \quad \gamma \leq 1/(2\mathbf{T}_1) \quad (214)$$

Since $(1 + \gamma)^{\mathbf{T}_1} \leq e^{\mathbf{T}_1 \gamma} \leq \sqrt{e}$, we get

$$\frac{2M}{K} \lesssim \frac{1}{\mathbf{T}_1 \gamma} \sqrt{\frac{d}{rD \log d}} \lesssim \frac{1}{\mathbf{T}_1 C_H} \sqrt{\frac{n}{rD \log d}} \lesssim \sqrt{n} \quad (215)$$

Thus via choosing a large enough universal constant D , we have for all $j \leq 4p$

$$\left(\frac{2M}{K} \right)^j r \mathbf{T}_1 d^{-D} \lesssim n^{2p} d^{-D} \lesssim 1 \quad (216)$$

and we finally obtain the desired conclusion via combining two terms. \square

D.5.2 PROOF OF LEMMA D.2

Proof. We want to ensure $\frac{d}{\lambda_r^{2\mathbf{T}_1}} K^2 \leq c^2$. Taking square roots on both sides, we note that this is equivalent to

$$\frac{\sqrt{d}}{\lambda_r^{\mathbf{T}_1}} K \leq c \quad \text{and} \quad K \leq c \frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}}. \quad (217)$$

To get a sufficient condition, we bound the two terms in K separately. Recall that we set

$$K = C \left(\mathbf{T}_1 \gamma \sqrt{\frac{rD \log d}{d}} + \gamma^{\mathbf{T}_1} \right) \quad (218)$$

therefore we need

$$C \mathbf{T}_1 \gamma \sqrt{\frac{rD \log d}{d}} \leq c \frac{\lambda_r^{\mathbf{T}_1}}{2\sqrt{d}} \quad \text{and} \quad C \gamma^{\mathbf{T}_1} \leq c \frac{\lambda_r^{\mathbf{T}_1}}{2\sqrt{d}}. \quad (219)$$

Under these, we get $K \leq c \frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}}$ as desired. Note that

$$\gamma = C_H \sqrt{\frac{d}{n}} \quad \text{and} \quad C C_H \mathbf{T}_1 \sqrt{\frac{d}{n}} \sqrt{\frac{rD \log d}{d}} \lesssim \mathbf{T}_1 \sqrt{\frac{rD \log d}{n}} \quad (220)$$

Then, to satisfy the first one, we require $\mathbf{T}_1 \sqrt{\frac{rD \log d}{n}} \lesssim \frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}}$, which is equivalent to

$$n \gtrsim \frac{\mathbf{T}_1^2 r D d \log d}{\lambda_r^{2\mathbf{T}_1}} = \mathbf{T}_1^2 r D d \log d \kappa^{2\mathbf{T}_1} \quad (221)$$

For the second one, we require $\sqrt{\frac{d}{n}} \lesssim \left(\frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}} \right)^{1/\mathbf{T}_1} = \frac{\lambda_r}{(\sqrt{d})^{1/\mathbf{T}_1}}$ which is equivalent to

$$n \gtrsim \frac{d^{1+1/\mathbf{T}_1}}{\lambda_r^2} = d^{1+1/\mathbf{T}_1} \kappa^2 \quad (222)$$

Combining the two parts above, the final conclusion directly follows. \square

D.5.3 PROOF OF LEMMA D.3

Proof. From Lemma D.2 we know that $K \lesssim \frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}}$. Then on the intersection of event $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta$ the residual satisfies

$$\|\mathbf{r}(\mathbf{w})\| \lesssim K \lesssim \frac{\lambda_r^{\mathbf{T}_1}}{\sqrt{d}} \leq \sqrt{\log d/d}. \quad (223)$$

Moreover $\mathbf{h}_n(\mathbf{w}) = \mathbf{h}(\mathbf{w}) + \mathbf{r}(\mathbf{w})$ and $\mathbf{h}(\mathbf{w}) = \Sigma_\ell^{\mathbf{T}_1} \mathbf{w}$. Under $\widetilde{\mathcal{E}}_\beta$, we have $\|\mathbf{h}(\mathbf{w})\| \lesssim \sqrt{\log d/d}$. Therefore, under the event $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta$ we have

$$\|\mathbf{h}_n(\mathbf{w})\| \leq \|\mathbf{h}(\mathbf{w})\| + \|\mathbf{r}(\mathbf{w})\| \lesssim \sqrt{\log d/d}. \quad (224)$$

□

D.5.4 PROOF OF LEMMA D.5

Proof. Note that since every vector in $\text{span}(\text{Mat}(\mathbb{E}_\mathbf{w}[(\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes 2k}]))$ is a vectorized symmetric k tensor, it suffices to show that

$$\mathbb{E}[(\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes 2k}](\mathbf{T}, \mathbf{T}) \gtrsim \left(\frac{d}{\lambda_r^{\mathbf{T}_1}}\right)^{-k} \quad (225)$$

for all symmetric k tensors \mathbf{T} supported on S^* with $\|\mathbf{T}\|_F^2 = 1$.

For notation convenience, we denote $\mathbf{w}_\epsilon := \mathbf{w}/\epsilon_0$ where we know $\mathbf{w}_\epsilon \sim \text{Unif}(\mathbb{S}^{d-1})$. Recalling that $\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}) = \Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon + \mathbf{\Pi}^* \mathbf{r}(\mathbf{w})$ and the tensor binomial expansion, we have

$$(\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes k} = \sum_{i=0}^k \binom{k}{i} \left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \otimes (\mathbf{\Pi}^* \mathbf{r}(\mathbf{w}))^{\otimes i} \quad (226)$$

Thus for any symmetric k -tensor \mathbf{T} , we have

$$\langle \mathbf{T}, (\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes k} \rangle = \left\langle \mathbf{T}, \left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes k} \right\rangle + \delta(\mathbf{w}) \quad (227)$$

where

$$\delta(\mathbf{w}) := \sum_{i=1}^k \binom{k}{i} \langle \mathbf{T}, \left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \otimes (\mathbf{\Pi}^* \mathbf{r}(\mathbf{w}))^{\otimes i} \rangle \quad (228)$$

Note that by Young's inequality,

$$\mathbb{E} \left[\langle \mathbf{T}, (\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes k} \rangle^2 \right] \geq \mathbb{E} \left[\frac{1}{2} \left\langle \mathbf{T}, \left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes k} \right\rangle^2 \right] - \mathbb{E} [\delta(\mathbf{w})^2] \quad (229)$$

Therefore we need to upper bound $\mathbb{E}[\delta(\mathbf{w})^2]$ and lower bound $\mathbb{E}[\frac{1}{2} \langle \mathbf{T}, (\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2]$. Intuitively the $\delta(\mathbf{w})$ term should be negligible since $\mathbf{r}(\mathbf{w})$ is of small norm.

Note that there is a crude upper bound for $\delta(\mathbf{w})$. Contracting \mathbf{T} in its first $k-i$ slots and applying the Frobenius Cauchy-Schwarz inequality gives

$$\left| \langle \mathbf{T}, \left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \otimes (\mathbf{\Pi}^* \mathbf{r}(\mathbf{w}))^{\otimes i} \rangle \right| \leq \left\| \mathbf{T} \left(\left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \right) \right\|_F \|\mathbf{\Pi}^* \mathbf{r}(\mathbf{w})\|^i \quad (230)$$

Hence

$$\begin{aligned} |\delta(\mathbf{w})| &\leq \sum_{i=1}^k \binom{k}{i} \left\| \mathbf{T} \left(\left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \right) \right\|_F \|\mathbf{\Pi}^* \mathbf{r}(\mathbf{w})\|^i \\ &\lesssim \sum_{i=1}^k \left\| \mathbf{T} \left(\left(\Sigma_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon\right)^{\otimes (k-i)} \right) \right\|_F \|\mathbf{\Pi}^* \mathbf{r}(\mathbf{w})\|^i. \end{aligned} \quad (231)$$

Via Cauchy–Schwarz again, we have

$$\begin{aligned}\mathbb{E}[\delta(\mathbf{w})^2] &\lesssim \sum_{i=1}^k \sqrt{\mathbb{E}[\|\mathbf{T}((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)})\|_F^4] \mathbb{E}[\|\boldsymbol{\Pi}^* \mathbf{r}(\mathbf{w})\|^{4i}]} \\ &\lesssim \sum_{i=1}^k \mathbb{E}[\|\mathbf{T}((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)})\|_F^2] \sqrt{\mathbb{E}[\|\boldsymbol{\Pi}^* \mathbf{r}(\mathbf{w})\|^{4i}]} \end{aligned} \quad (232)$$

where the second line is derived via the Gaussian hypercontractivity lemma that we will state at Lemma E.22 in Section E.6 .

Next, let $\widehat{\mathbf{T}}$ be the symmetric k tensor defined by $\widehat{\mathbf{T}}(\mathbf{v}_1, \dots, \mathbf{v}_k) = \mathbf{T}(\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{v}_1, \dots, \boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{v}_k)$. For each i and each unit-Frobenius norm tensor \mathbf{V} supported on S^* we have

$$\langle \mathbf{T}((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)}), \mathbf{V} \rangle = \widehat{\mathbf{T}}(\mathbf{w}_\epsilon^{\otimes(k-i)}, (\boldsymbol{\Sigma}_\ell^{\text{T}_1})^\dagger \mathbf{V}) \quad (233)$$

i.e. by moving $\boldsymbol{\Sigma}_\ell^{\text{T}_1}$ from the arguments of \mathbf{T} onto the test tensor one applies $(\boldsymbol{\Sigma}_\ell^{\text{T}_1})^\dagger$ to each slot of \mathbf{V} . By Cauchy-Schwarz again,

$$\left| \mathbf{T} \left((\boldsymbol{\Sigma}_\ell^{\text{T}_1} (\mathbf{w}/\epsilon_0))^{\otimes(k-i)} \right), \mathbf{V} \right| \leq \|\widehat{\mathbf{T}}(\mathbf{w}_\epsilon^{\otimes(k-i)})\|_F \|(\boldsymbol{\Sigma}_\ell^{\text{T}_1})^\dagger \mathbf{V}\|_F. \quad (234)$$

Since $\|(\boldsymbol{\Sigma}_\ell^{\text{T}_1})^\dagger \mathbf{V}\|_F \leq 1/\lambda_r^{iT_1}$ for $\|\mathbf{V}\|_F = 1$, taking the supremum over unit V gives the pointwise bound

$$\|\mathbf{T} \left((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)} \right)\|_F \leq 1/\lambda_r^{iT_1} \|\widehat{\mathbf{T}}(\mathbf{w}_\epsilon^{\otimes(k-i)})\|_F \quad (235)$$

Squaring and taking expectation over \mathbf{w} yields

$$\mathbb{E}[\|\mathbf{T}((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)})\|_F^2] \leq \frac{1}{\lambda_r^{2iT_1}} \mathbb{E}[\|\widehat{\mathbf{T}}(\mathbf{w}_\epsilon^{\otimes(k-i)})\|_F^2] \quad (236)$$

Combining with previous calculation, we have the estimation below

$$\begin{aligned}\mathbb{E}[\|\mathbf{T}((\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes(k-i)})\|_F^2] &\leq \frac{1}{\lambda_r^{2iT_1}} \mathbb{E}[\|\widehat{\mathbf{T}}(\mathbf{w}_\epsilon^{\otimes(k-i)})\|_F^2] \\ &\lesssim \left(\frac{d}{\lambda_r^{2T_1}} \right)^i \mathbb{E}[\langle \mathbf{T}, (\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2] \end{aligned} \quad (237)$$

where the second inequality is deduced from the lemma below.

Lemma D.34. *Let $\mathbf{w} \sim \text{Unif}(\mathbb{S}^{d-1})$ and for each tensor \mathbf{T} and each $k \leq p$ we have*

$$\mathbb{E}\|\mathbf{T}(\mathbf{w}^{\otimes k})\|_F^2 \lesssim d^{p-k} \mathbb{E}\langle \mathbf{T}, \mathbf{w}^{\otimes p} \rangle^2. \quad (238)$$

The proof of Lemma D.34 is provided in Appendix E.5

We also know from Lemma D.1 that for any integer $1 \leq j \leq 4p$ we have under the high probability event of \mathcal{E}_α , which happens with probability at least $1 - \mathcal{O}(d^{-D})$ that

$$[\mathbb{E}\|\boldsymbol{\Pi}^* \mathbf{r}(\mathbf{w})\|^j]^{1/j} \lesssim K \quad (239)$$

Therefore, we have

$$\mathbb{E}[\delta(\mathbf{w})^2] \lesssim \mathbb{E}[\langle \mathbf{T}, (\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2] \sum_{i=1}^k \left(\frac{d}{\lambda_r^{2T_1}} K^2 \right)^i \quad (240)$$

Because we have assumed $n \gtrsim rDd \log d\kappa^{2T_1} + d^{1+1/T_1} \kappa^2$, we have the following from Lemma D.2 that

$$\frac{d}{\lambda_r^{2T_1}} K^2 \leq c^2 \quad (241)$$

with some prechosen absolute constant c . Thus via picking a small c and plugging the above bound into the previous equation we have

$$\mathbb{E}_{\mathbf{w}}[\delta(\mathbf{w})^2] \leq \frac{1}{4} \mathbb{E}[\langle \mathbf{T}, (\boldsymbol{\Sigma}_\ell^{\text{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2]. \quad (242)$$

Combining everything gives

$$\begin{aligned} \mathbb{E}[\langle \mathbf{T}, (\mathbf{\Pi}^* \mathbf{h}_n(\mathbf{w}))^{\otimes k} \rangle^2] &\geq \mathbb{E}\left[\frac{1}{2} \langle \mathbf{T}, (\mathbf{\Sigma}_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2\right] - \mathbb{E}_{\mathbf{w}}[\delta(\mathbf{w})^2] \\ &\geq \frac{1}{4} \mathbb{E}[\langle \mathbf{T}, (\mathbf{\Sigma}_\ell^{\mathbf{T}_1} \mathbf{w}_\epsilon)^{\otimes k} \rangle^2] \end{aligned} \quad (243)$$

$$\begin{aligned} &\gtrsim d^{-k} \|\widehat{\mathbf{T}}\|_F^2 \\ &\geq \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}}\right)^{-k} \end{aligned} \quad (244)$$

where the third line is via taking the expectation over \mathbf{w} , the detailed proof for this can be also found in Appendix E.5. The proof is finished. \square

D.5.5 PROOF OF LEMMA D.6

Proof. Let

$$z_{\mathbf{T}}(\mathbf{w}) := \text{Vec}(\mathbf{T})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k}). \quad (245)$$

Note that $\text{Vec}(\mathbf{T}) \in \text{span}(\text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}]))$ by Lemma D.5. Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})(\mathbf{h}_n(\mathbf{w})^\top \mathbf{x})^k] &= \mathbb{E}_{\mathbf{w}}[\text{Vec}(\mathbf{T})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k}) \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k})^\top \text{Vec}(\mathbf{x}^{\otimes k})] \\ &= \langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle. \end{aligned} \quad (246)$$

For the bounds on $z_{\mathbf{T}}(\mathbf{w})$ we have the following estimation

$$\begin{aligned} \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})^2] &= \mathbb{E}_{\mathbf{w}}[\text{Vec}(\mathbf{T})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k}) \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{T})] \\ &= \text{Vec}(\mathbf{T})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{T}) \\ &\lesssim \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}}\right)^k \|\mathbf{T}\|_F^2, \end{aligned} \quad (247)$$

and

$$|z_{\mathbf{T}}(\mathbf{w})| = |\text{Vec}(\mathbf{T})^\top \text{Mat}(\mathbb{E}[\mathbf{h}_n(\mathbf{w})^{\otimes 2k}])^\dagger \text{Vec}(\mathbf{h}_n(\mathbf{w})^{\otimes k})| \lesssim \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}}\right)^k \|\mathbf{T}\|_F \|\mathbf{h}_n(\mathbf{w})\|^k. \quad (248)$$

The proof is complete. \square

D.5.6 PROOF OF LEMMA D.9

Proof. Recall that from Lemma D.3 and D.4, under the high probability event of $\mathcal{E}_\alpha \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta$, which happens with probability at least $1 - \mathcal{O}(d^{-D})$ the following bounds hold

$$\max_{j \in [n+1, 2n]} \|\mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_j\| \lesssim \sqrt{\frac{\log^2 d}{d}} \quad \text{and} \quad \|\mathbf{h}_n(\mathbf{w})\| \lesssim \sqrt{\frac{\log d}{d}}. \quad (249)$$

Those bounds will be helpful for us to construct the low-dimensional approximation in the second stage.

Recall that we have set $\eta = \eta_1 = \frac{1}{C} \left(\frac{d}{\log^2 d}\right)^{1/(2\mathbf{T}_1)}$ where the universal constant C is chosen sufficiently large so we can absorb the implicit constants hidden in the \lesssim bounds. This means we choose C to make sure on the same high-probability event the following controls are valid

$$\eta^{\mathbf{T}_1} \|\mathbf{h}_n(\mathbf{w})\| \leq 1 \quad \text{and} \quad \max_{j \in [n+1, 2n]} |\eta^{\mathbf{T}_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_j| \leq 1 \quad (250)$$

which indicates that

$$\|\widetilde{\mathbf{w}}^{(\mathbf{T}_1)}\| = \left\| \frac{1}{\epsilon_0} (-a_j)^{\mathbf{T}_1} \eta^{\mathbf{T}_1} \widehat{\mathbf{\Sigma}}_\ell^{\mathbf{T}_1} \mathbf{w} \right\| \leq 1 \quad (251)$$

Note that $\mathbf{w}^{(T_1)} = \frac{1}{\epsilon_0} (-a_j)^{T_1} \eta^{T_1} \left(\prod_{t=0}^{T_1-1} \mathbf{M}^{(t)} \right) \mathbf{w}$ and by Lemma C.1, we have

$$\|\mathbf{w}^{(T_1)} - \widehat{\mathbf{w}}^{(T_1)}\| \lesssim K \lesssim \eta^{T_1} \sqrt{\frac{\log d}{d}} \quad (252)$$

Thus $\|\mathbf{w}^{(T_1)}\| \leq 2\|\widehat{\mathbf{w}}^{(T_1)}\| \lesssim 1$ and we complete our proof. \square

D.5.7 PROOF OF LEMMA D.10

Proof. For any training point \mathbf{x} and given initialization \mathbf{w} , on events $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta$ we have the following control

$$z := (-\eta)^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} \in [-1, 1]. \quad (253)$$

By our previous assumption on the activation function, we have the following low dimensional approximation result uniformly for $z \in [-1, 1]$

$$\left| \mathbb{E}_{a,b}[\mathbf{v}_k^{(\epsilon_k)}(a, b) \sigma(az + b)] - z^k \right| \leq \epsilon_k. \quad (254)$$

Divide both sides by $(-\eta)^{kT_1}$ and substitute the features back we obtain

$$\left| \frac{1}{(-\eta)^{kT_1}} \mathbb{E}_{a,b}[\mathbf{v}_k^{(\epsilon_k)}(a, b) \sigma((-\eta a)^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b)] - (\mathbf{h}_n(\mathbf{w})^\top \mathbf{x})^k \right| \leq \frac{\epsilon_k}{\eta^{kT_1}}. \quad (255)$$

Recall that from Corollary D.6 under the high probability event \mathcal{E}_α we have

$$\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})(\mathbf{h}_n(\mathbf{w})^\top \mathbf{x})^k] = \langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle, \quad (256)$$

And also recall from Lemma D.4 and the statement before the lemma we know that under the high probability event $\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta$

$$|\eta^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x}_{i+n}| \leq 1 \quad (257)$$

$$\eta^{T_1} \|\mathbf{h}_n(\mathbf{w})\| \leq 1 \quad (258)$$

thus

$$\begin{aligned} & |f_{h_{\mathbf{T}}}(\mathbf{x}) - \langle \mathbf{T}, \mathbf{x}^{\otimes k} \rangle| \\ & \leq \frac{\epsilon_k}{\eta^{kT_1}} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap \widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta}] \right| + \frac{1}{\eta^{kT_1}} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \\ & \quad + \frac{1}{\eta^{kT_1}} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{v}_k^{(\epsilon_k)}(a, b) \sigma((-\eta a)^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \\ & \leq \frac{\epsilon_k}{\eta^{kT_1}} |\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})]| + \frac{1}{\eta^{kT_1}} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \\ & \quad + \eta^{-kT_1} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{v}_k^{(\epsilon_k)}(a, b) \sigma((-\eta a)^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \end{aligned} \quad (259)$$

We bound the third term first. Denote

$$X(a, b, \mathbf{w}) := \frac{1}{\eta^{kT_1}} \mathbf{v}_k^{(\epsilon_k)}(a, b) z_{\mathbf{T}}(\mathbf{w}) \sigma((-\eta a)^{T_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b). \quad (260)$$

By Cauchy-Schwarz,

$$\left| \mathbb{E}_{a,b,\mathbf{w}}[X(a, b, \mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \leq \sqrt{\mathbb{E}_{a,b,\mathbf{w}}[X(a, b, \mathbf{w})^2]} \sqrt{\mathbb{P}(\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c)} \quad (261)$$

We bound the two factors on the right. Recall that it has been assumed that the activation function has bounded third derivatives as in Assumption 2. Therefore, there exists a constant $C_\sigma > 0$ such that

$$|\sigma(u)| \leq C_\sigma (1 + |u|^3) \quad \text{for all } u \in \mathbb{R}. \quad (262)$$

and also

$$|X(a, b, \mathbf{w})| \leq V_{k, \epsilon_k} |z_{\mathbf{T}}(\mathbf{w})| C_\sigma (1 + |\eta^{\mathbf{T}_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b|^3). \quad (263)$$

We recall that the following crude bound works for all \mathbf{w}

$$\|\mathbf{h}_n(\mathbf{w})\| \leq (1 + \gamma)^{\mathbf{T}_1} \leq 2^{\mathbf{T}_1} \quad (264)$$

where γ is defined in via directly bounding the operator norm of $\prod_{t=0}^{\mathbf{T}_1-1} \mathbf{M}^{(t)}$. Notice that under the event \mathcal{E}_ζ , we have $\|\mathbf{x}\| \leq 2\sqrt{d}$.

Thus via another crude bound we have $|\eta^{\mathbf{T}_1} \mathbf{h}_n(\mathbf{w})^\top \mathbf{x} + b| \lesssim d$ and

$$|X| \lesssim V_{k, \epsilon_k} |z_{\mathbf{T}}(\mathbf{w})| d^3. \quad (265)$$

In addition, from Definition D.7 we also have $V_{k, \epsilon_k} = \|\mathbf{v}_k^{(\epsilon_k)}\|_\infty \lesssim \epsilon_k^{-\beta} \lesssim d^{D/12}$. Also invoking $\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}^2] \lesssim d^k \kappa^{2\mathbf{T}_1 k} \|\mathbf{T}\|_F^2$ from D.6, note that for $\mathbf{T}_1 = o(\log d)$, we have $\kappa^{\mathbf{T}_1} \lesssim d$, thus it shows that we obtain

$$\sqrt{\mathbb{E}[X(a, b, \mathbf{w})^2]} \leq V_{k, \epsilon_k} d^3 \sqrt{\mathbb{E}[z_{\mathbf{T}}^2]} \lesssim d^{D/12+3k/2+3} \|\mathbf{T}\|_F. \quad (266)$$

Second, Lemma D.3 and D.4 give

$$\mathbb{P}(\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c) \leq \mathbb{P}((\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c) \lesssim d^{-D/2} \quad (267)$$

when $\iota = 2D \log d$. Hence

$$\left| \mathbb{E}_{a, b, \mathbf{w}}[X(a, b, \mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| \lesssim d^{-D/8} \|\mathbf{T}\|_F \quad (268)$$

for sufficiently large universal constant D .

The above argument shows that the third term is negligible.

Then we bound the second term. Using essentially the same method and argument as above, we obtain

$$\begin{aligned} \frac{1}{\eta^{k\mathbf{T}_1}} \left| \mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w}) \mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c}] \right| &\lesssim \frac{1}{\eta^{k\mathbf{T}_1}} \sqrt{\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}^2]} \sqrt{\mathbb{P}(\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\widetilde{\mathcal{E}}_\beta \cap \widetilde{\mathcal{E}}_\eta)^c)} \\ &\lesssim d^{3k/2} \|\mathbf{T}\|_F d^{-D/4} \\ &\lesssim d^{-D/8} \|\mathbf{T}\|_F \end{aligned} \quad (269)$$

For the first one, we define the approximation error

$$\Delta_k := \frac{\epsilon_k}{\eta^{k\mathbf{T}_1}} \|\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})]\|. \quad (270)$$

By Cauchy–Schwarz,

$$\Delta_k \leq \frac{\epsilon_k}{\eta^{k\mathbf{T}_1}} \sqrt{\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})^2]}. \quad (271)$$

Substitute the second-moment bound mentioned in Corollary D.6 and we obtain

$$\mathbb{E}_{\mathbf{w}}[z_{\mathbf{T}}(\mathbf{w})^2] \lesssim \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}} \right)^k \|\mathbf{T}\|_F^2 \quad (272)$$

thus

$$\Delta_k \lesssim \epsilon_k \frac{1}{\eta^{k\mathbf{T}_1}} \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}} \right)^{k/2} \|\mathbf{T}\|_F \lesssim \epsilon_k \kappa^{k\mathbf{T}_1} \log^k d \|\mathbf{T}\|_F. \quad (273)$$

For the other moments, from Definition D.7 we have $\|\mathbf{v}_k^{(\epsilon_k)}\|_\infty \leq V_{k, \epsilon_k}$ and

$$\mathbb{E}[h_{\mathbf{T}}^2] \leq \frac{V_{k, \epsilon_k}^2}{\eta^{2k\mathbf{T}_1}} \mathbb{E}[z_{\mathbf{T}}(\mathbf{w})^2] \lesssim V_{k, \epsilon_k}^2 \frac{1}{\eta^{2k\mathbf{T}_1}} \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}} \right)^k \|\mathbf{T}\|_F^2. \quad (274)$$

Substitute η back and we can get

$$\mathbb{E}[h_{\mathbf{T}}^2] \lesssim V_{k, \epsilon_k}^2 (\kappa^{2\mathbf{T}_1} \log^2 d)^k \|\mathbf{T}\|_F^2. \quad (275)$$

Similarly, using $|z_{\mathbf{T}}(\mathbf{w})| \lesssim \left(\frac{d}{\lambda_r^{2\mathbf{T}_1}} \right)^k \|\mathbf{T}\|_F \|\mathbf{h}_n(\mathbf{w})\|^k$ again and noticing the truncation terms in the definition of $h_{\mathbf{T}}(a, b, \mathbf{w})$, we have

$$\|h_{\mathbf{T}}\|_\infty \lesssim V_{k, \epsilon_k} (\kappa^{2\mathbf{T}_1} \log^2 d)^k \|\mathbf{T}\|_F \quad (276)$$

and the lemma is proved. \square

D.5.8 PROOF OF LEMMA D.13

Proof. Since $h(a, b, \mathbf{w}) := \sum_{k \leq p} h_{\mathbf{T}_k}(a, b, \mathbf{w})$, combining with linearity of expectation, we have

$$f_h(\mathbf{x}) = \sum_{k \leq p} f_{h_{\mathbf{T}_k}}(\mathbf{x}) \quad (277)$$

For each k , define the empirical error vector $\mathbf{o}_k \in \mathbb{R}^n$ by

$$\mathbf{o}_k(i) := f_{h_{\mathbf{T}_k}}(\mathbf{x}_i) - \langle \mathbf{T}_k, \mathbf{x}_i^{\otimes k} \rangle, \quad i = n+1, \dots, 2n. \quad (278)$$

Then for the aggregated function we have, for all \mathbf{x}_i from the second stage

$$f_h(\mathbf{x}_i) - f^*(\mathbf{x}_i) = \sum_{k \leq p} \mathbf{o}_k(i), \quad (279)$$

For now throughout this proof we equip \mathbb{R}^n with the empirical ℓ_2 norm, defined for any vector $\mathbf{v} = (v(1), \dots, v(n)) \in \mathbb{R}^n$ by

$$\|\mathbf{v}\|_n^2 := \frac{1}{n} \sum_{i=1}^n v(i)^2. \quad (280)$$

Hence we can write the empirical squared loss by

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_h(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 = \left\| \sum_{k \leq p} \mathbf{o}_k \right\|_n^2 \quad (281)$$

Note that we have

$$\left\| \sum_{k \leq p} \mathbf{o}_k \right\|_n^2 \leq \left(\sum_{k \leq p} \|\mathbf{o}_k\|_n \right)^2 \quad (282)$$

By Corollary D.11, each single-term error satisfies $\|\mathbf{o}_k\|_n^2 \lesssim \Delta_k^2 + d^{-D/4} \|\mathbf{T}_k\|_F^2$. Thus

$$\|\mathbf{o}_k\|_n \lesssim \Delta_k + d^{-D/8} \|\mathbf{T}_k\|_F. \quad (283)$$

Notice that by Lemma E.21 we have $\|\mathbf{T}_k\|_F \lesssim r^{(p-k)/4} \lesssim 1$. Therefore

$$\begin{aligned} \left(\sum_{k \leq p} \|\mathbf{o}_k\|_n \right)^2 &\lesssim \left(\sum_{k \leq p} \Delta_k + d^{-D/8} \sum_{k \leq p} \|\mathbf{T}_k\|_F \right)^2 \\ &\lesssim \left(\sum_{k \leq p} \Delta_k + d^{-D/8} p \right)^2 \\ &\lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + (d^{-D/8})^2 \\ &\lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-D/4} \end{aligned} \quad (284)$$

where the last step absorbs r, p into the \lesssim constant. This proves equation (147).

For the moment bounds, we have by Cauchy–Schwarz

$$\mathbb{E}[h^2] = \mathbb{E} \left[\left(\sum_{k \leq p} h_{\mathbf{T}_k} \right)^2 \right] = \sum_{k, \ell \leq p} \mathbb{E}[h_{\mathbf{T}_k} h_{\mathbf{T}_\ell}] \leq \left(\sum_{k \leq p} \sqrt{\mathbb{E}[h_{\mathbf{T}_k}^2]} \right)^2 \quad (285)$$

Substituting the second-moment bound mentioned in Lemma D.10, we obtain for each k ,

$$\mathbb{E}[h_{\mathbf{T}_k}^2] \lesssim V_{k, \epsilon_k}^2 (\kappa^{2\mathbf{T}_1} \log^2 d)^k \|\mathbf{T}_k\|_F^2 \quad (286)$$

hence

$$\sqrt{\mathbb{E}[h_{\mathbf{T}_k}^2]} \lesssim V_{p,\epsilon}(\kappa^{\mathbf{T}_1} \log d)^k \|\mathbf{T}_k\|_F. \quad (287)$$

Using the bound $\|\mathbf{T}_k\|_F \lesssim r^{\frac{p-k}{4}} \lesssim 1$ and summing over $k \leq p$, we obtain

$$\sum_{k \leq p} \sqrt{\mathbb{E}[h_{\mathbf{T}_k}^2]} \lesssim V_{p,\epsilon} \sum_{k \leq p} (\kappa^{\mathbf{T}_1})^k (\log d)^k \lesssim V_{p,\epsilon} (\kappa^{\mathbf{T}_1})^p (\log d)^p. \quad (288)$$

where the last step absorbs the polynomial degree p into the \lesssim constant. Further, squaring both sides yields

$$\mathbb{E}[h^2] \lesssim V_{p,\epsilon}^2 (\kappa^{\mathbf{T}_1})^{2p} (\log d)^{2p} \quad (289)$$

which is the first inequality of Equation (293). For the sup-norm, substitute the sup-norm bound mentioned in Lemma D.10 we know, for each k ,

$$\|h_{\mathbf{T}_k}\|_\infty \lesssim V_{k,\epsilon_k} (\kappa^{2\mathbf{T}_1} \log^2 d)^k \|\mathbf{T}_k\|_F. \quad (290)$$

Combining the k bounds together we have

$$\|h\|_\infty \leq \sum_{k \leq p} \|h_{\mathbf{T}_k}\|_\infty \lesssim \sum_{k \leq p} V_{k,\epsilon_k} (\kappa^{2\mathbf{T}_1} (\log d)^2)^k \|\mathbf{T}_k\|_F \lesssim V_{p,\epsilon} (\kappa^{2\mathbf{T}_1} (\log d)^2)^p \quad (291)$$

This proves the second one of Equation (293), concluding our proof. \square

D.5.9 PROOF OF LEMMA D.14

Proof. We first bound $\|\mathbf{a}^*\|$ where $\mathbf{a}^* = (a_1^*, \dots, a_m^*)$. Since

$$a_j^* := \frac{1}{m} h(a_j, b_j, \mathbf{w}_j) \quad (292)$$

and via Lemma D.13 we have the following moment estimation

$$\|h\|_\infty \lesssim V_{p,\epsilon} (\kappa^{\mathbf{T}_1})^{2p} (\log d)^{2p}. \quad (293)$$

Thus we directly get as desired

$$\|\mathbf{a}^*\|^2 \lesssim \frac{V_{p,\epsilon}^2 (\kappa^{\mathbf{T}_1})^{4p} (\log d)^{4p}}{m} \quad (294)$$

To bound the square loss for finite width, take the basic decomposition

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{\widehat{\Theta}}(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \leq 2 \frac{1}{n} \sum_{i=n+1}^{2n} (f_{\widehat{\Theta}^*}(\mathbf{x}_i) - f_h(\mathbf{x}_i))^2 + 2 \frac{1}{n} \sum_{i=n+1}^{2n} (f_h(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2. \quad (295)$$

We next bound these two terms separately.

For the second term, by Lemma D.13 we have

$$2 \frac{1}{n} \sum_{i=n+1}^{2n} (f_h(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-D/4} \quad (296)$$

It thus suffices to control the Monte Carlo fluctuation term (the first term). Note that the parameters are initialized symmetrically and this symmetry is preserved by first stage training, that is, the first half neurons are free. Therefore we should group neurons into symmetric pairs to help the proof.

For $j \leq m/2$, define

$$Z_j(\mathbf{x}) := a_j^* \sigma((\widetilde{\mathbf{w}}_j^{(\mathbf{T}_1)})^\top \mathbf{x} + b_j) + a_{m-j}^* \sigma((\widetilde{\mathbf{w}}_{m-j}^{(\mathbf{T}_1)})^\top \mathbf{x} + b_{m-j}), \quad \widetilde{\mathbf{w}}_j^{(\mathbf{T}_1)} = (-\eta a_j)^{\mathbf{T}_1} \mathbf{h}_n(\mathbf{w}_j) \quad (297)$$

so that $f_{\widehat{\Theta}^*}(\mathbf{x}) - f_h(\mathbf{x}) = \sum_{j \leq m/2} (Z_j(\mathbf{x}) - \mathbb{E}[Z_j(\mathbf{x})])$. Introduce the truncated version

$$\overline{Z}_j(\mathbf{x}) := Z_j(\mathbf{x}) \mathbf{1}_{\{|\widetilde{\mathbf{w}}_j^{(\mathbf{T}_1)\top} \mathbf{x}| \leq 1, |\widetilde{\mathbf{w}}_{m-j}^{(\mathbf{T}_1)\top} \mathbf{x}| \leq 1\}} \quad (298)$$

By the concentration event \mathcal{E}_μ , which occurs with probability at least $1 - O(d^{-D/2})$ we have $\bar{Z}_j(\mathbf{x}_i) = Z_j(\mathbf{x}_i)$ for all $n+1 \leq i \leq 2n$ and all $j \leq m/2$. Hence, on \mathcal{E} , we have the following decomposition

$$f_{\bar{\Theta}^*}(\mathbf{x}_i) - f_h(\mathbf{x}_i) = \sum_{j \leq m/2} (\bar{Z}_j(\mathbf{x}_i) - \mathbb{E}\bar{Z}_j(\mathbf{x}_i)) + \underbrace{\frac{m}{2} (\mathbb{E}\bar{Z}_j(\mathbf{x}_i) - \mathbb{E}Z_j(\mathbf{x}_i))}_{:=B_i}, \quad (299)$$

We bound the second term B_i first.

Recall

$$B_i := \frac{m}{2} (\mathbb{E}\bar{Z}_j(\mathbf{x}_i) - \mathbb{E}Z_j(\mathbf{x}_i)) = \frac{m}{2} \mathbb{E}[(\bar{Z}_j(\mathbf{x}_i) - Z_j(\mathbf{x}_i))\mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c}], \quad (300)$$

because $\bar{Z}_j = Z_j$ on \mathcal{E} . By Cauchy-Schwarz,

$$\begin{aligned} |B_i| &\leq \frac{m}{2} \mathbb{E}[|Z_j(\mathbf{x}_i)|\mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c}] + \frac{m}{2} \mathbb{E}[|\bar{Z}_j(\mathbf{x}_i)|\mathbf{1}_{\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c}] \\ &\leq \frac{m}{2} \left(\sqrt{\mathbb{E}Z_j(\mathbf{x}_i)^2} + \sqrt{\mathbb{E}\bar{Z}_j(\mathbf{x}_i)^2} \right) \sqrt{\mathbb{P}(\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c)} \\ &\leq m \sqrt{\mathbb{E}Z_j(\mathbf{x}_i)^2} \sqrt{\mathbb{P}(\mathcal{E}_\alpha \cap \mathcal{E}_\zeta \cap (\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c)}. \end{aligned} \quad (301)$$

The final inequality follows from the fact that truncation does not increase magnitude, that is, $|\bar{Z}_j(\mathbf{x}_i)| \leq |Z_j(\mathbf{x}_i)|$. For a symmetric pair

$$Z_j(\mathbf{x}_i) = a_j^* \sigma(\widetilde{\mathbf{w}}_j^{(\Gamma_1)\top} \mathbf{x}_i + b_j) + a_{m-j}^* \sigma(\widetilde{\mathbf{w}}_{m-j}^{(\Gamma_1)\top} \mathbf{x}_i + b_{m-j}) \quad (302)$$

Recall that it has been assumed that the activation function has bounded third derivatives as in Assumption 2. Therefore, there exists an absolute constant $C_\sigma > 0$ such that

$$|\sigma(u)| \leq C_\sigma(1 + |u|^3) \quad \text{for all } u \in \mathbb{R}. \quad (303)$$

Then, for any fixed i, j ,

$$\mathbb{E}[Z_j(\mathbf{x}_i)^2] \lesssim (a_j^*)^2 (1 + |\eta^{\Gamma_1} \mathbf{h}_n(\mathbf{w}_j^{(0)})^\top \mathbf{x}_i + b_j|^3)^2 + (a_{m-j}^*)^2 (1 + |\eta^{\Gamma_1} \mathbf{h}_n(\mathbf{w}_{m-j}^{(0)})^\top \mathbf{x}_i + b_{m-j}|^3)^2. \quad (304)$$

Notice that under the event \mathcal{E}_ζ , we have \mathbf{x}_i is bounded for all $i \in [n+1, 2n]$ by $\|\mathbf{x}_i\| \leq 2\sqrt{d}$. Thus via another crude bound, since we have $|\eta^{\Gamma_1} \mathbf{h}_n(\mathbf{w}_j^{(0)})^\top \mathbf{x}_i + b_j| \lesssim d$ and $|\eta^{\Gamma_1} \mathbf{h}_n(\mathbf{w}_{m-j}^{(0)})^\top \mathbf{x}_i + b_{m-j}| \lesssim d$ and we have

$$\begin{aligned} \sqrt{\mathbb{E}Z_j(\mathbf{x}_i)^2} &\lesssim d^3 \sqrt{(a_j^*)^2 + (a_{m-j}^*)^2} \\ &\lesssim d^3 \frac{V_{p,\epsilon}(\kappa^{\Gamma_1})^{4p} (\log d)^{4p}}{m} \end{aligned} \quad (305)$$

In addition, from Definition D.7 we also have $V_{k,\epsilon_k} = \|\mathbf{v}_k^{(\epsilon_k)}\|_\infty \lesssim \epsilon_k^{-\beta} \lesssim d^{\frac{D}{12}}$, thus $V_{p,\epsilon}^2 \lesssim d^{\frac{D}{6}}$

Also note that for $\Gamma_1 = o(\log d)$, we have $\kappa^{\Gamma_1} \lesssim d$, thus it shows that

$$\sqrt{\mathbb{E}Z_j(\mathbf{x}_i)^2} \lesssim \frac{d^{4+4p+\frac{D}{12}}}{m} \quad (306)$$

This gives the intermediate bound

$$|B_i| \lesssim m \frac{d^{4+4p+\frac{D}{12}}}{m} \sqrt{\mathbb{P}((\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c)} \lesssim d^{4+4p+\frac{D}{12}} \sqrt{\mathbb{P}((\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c)}. \quad (307)$$

By the concentration estimates proved above we have $\Pr((\mathcal{E}_\beta \cap \mathcal{E}_\eta)^c) \lesssim d^{-D/2}$ (union bound over the $m/2$ independent draws yields the extra factor m), hence

$$|B_i| \lesssim d^{4+4p+\frac{D}{12}} d^{-D/4} \lesssim d^{-D/8}. \quad (308)$$

for sufficiently large universal constant D .

Fix any $\mathbf{x} \in \{\mathbf{x}_i\}_{i=n+1}^{2n}$. The variables $\{\bar{Z}_j(\mathbf{x}) - \mathbb{E}\bar{Z}_j(\mathbf{x})\}_{j \leq m/2}$ are i.i.d. and centered, so we can use Bernstein to bound the sum. Notice that $(\mathbf{w}_j^{(T_1)})^\top \mathbf{x}$ and $\mathbf{w}_{m-j}^{(T_1)\top} \mathbf{x}$ is truncated by the definition of $Z_j(\mathbf{x})$, we have $\sigma((\mathbf{w}_j^{(T_1)})^\top \mathbf{x} + b_j)$ and $\sigma((\mathbf{w}_{m-j}^{(T_1)})^\top \mathbf{x} + b_{m-j})$ also bounded by a universal constant.

Combining with $a_j^* = \frac{1}{m}h(a_j, b_j, \mathbf{w}_j)$, we have for a universal constant C

$$|\bar{Z}_j(\mathbf{x}) - \mathbb{E}\bar{Z}_j(\mathbf{x})| \leq \frac{C}{m} \|h\|_\infty. \quad (309)$$

Lemma D.35 (Hoeffding's inequality). *Let X_1, X_2, \dots, X_n be i.i.d. random variables with mean μ and $X_i \in [a, b]$. Then,*

$$\mathbb{P} \left\{ \left| \frac{1}{n} \sum_{i=1}^n X_i - \mu \right| \geq t \right\} \leq 2 \exp \left(-\frac{2nt^2}{(b-a)^2} \right). \quad (310)$$

Due to the fact that $|\bar{Z}_j(x)| \lesssim \frac{1}{m} \|h\|_\infty$, Hoeffding inequality then yields with probability at least $1 - \mathcal{O}(d^{-D})$

$$\left| \sum_{j \leq m/2} (\bar{Z}_j(\mathbf{x}) - \mathbb{E}\bar{Z}_j(\mathbf{x})) \right| \lesssim \frac{\log d \cdot \|h\|_\infty}{\sqrt{m}}. \quad (311)$$

Recall from Lemma D.13 we have $\|h\|_\infty \lesssim V_{p,\epsilon}(\kappa^{T_1})^{2p}(\log d)^{2p}$. Plugging these bounds into the above one, it shows that

$$\left| \sum_{j \leq m/2} (\bar{Z}_j(\mathbf{x}) - \mathbb{E}\bar{Z}_j(\mathbf{x})) \right| \lesssim \frac{V_{p,\epsilon}(\kappa^{T_1})^{2p}(\log d)^{2p+1}}{\sqrt{m}}. \quad (312)$$

Therefore

$$\begin{aligned} |f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f_h(\mathbf{x}_i)| &\leq \left| \sum_{j \leq m/2} (\bar{Z}_j(\mathbf{x}_i) - \mathbb{E}\bar{Z}_j(\mathbf{x}_i)) \right| + \frac{m}{2} |\mathbb{E}\bar{Z}_j(\mathbf{x}_i) - \mathbb{E}Z_j(\mathbf{x}_i)| \\ &\lesssim \sqrt{\frac{V_{p,\epsilon}^2(\kappa^{T_1})^{4p}(\log d)^{4p+2}}{m}} + d^{-D/8} \end{aligned} \quad (313)$$

Squaring both sides and applying a union bound over the n training points, we obtain the following.

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f_h(\mathbf{x}_i))^2 \lesssim \frac{V_{p,\epsilon}^2(\kappa^{T_1})^{4p}(\log d)^{4p+2}}{m} + d^{-D/4} \quad (314)$$

By the initial decomposition and Lemma D.13,

$$\frac{1}{n} \sum_{i=n+1}^{2n} (f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 \lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-\frac{D}{4}} + \frac{V_{p,\epsilon}^2(\kappa^{T_1})^{4p}(\log d)^{4p+2}}{m}, \quad (315)$$

which is (177). This completes the proof. \square

D.5.10 PROOF OF LEMMA D.15

Proof. From Lemma D.14 and Lemma D.13, we know that with high probability at least $1 - \mathcal{O}(d^{-D/2})$ there exists $\mathbf{a}^* \in \mathbb{R}^m$ such that for $\tilde{\Theta}^* = (\mathbf{a}^*, \mathbf{b}^{(\mathbb{T}_1)}, \tilde{\mathbf{W}}^{(\mathbb{T}_1)})$,

$$\begin{aligned} \frac{1}{n} \sum_{i=n+1}^{2n} (f_{\tilde{\Theta}^*}(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 &\lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + d^{-D/4} + \frac{V_{p,\epsilon}^2(\kappa^{\mathbb{T}_1})^{4p}(\log d)^{4p+2}}{m} \\ &\lesssim \left(\sum_{k \leq p} \Delta_k \right)^2 + \frac{V_{p,\epsilon}^2(\kappa^{\mathbb{T}_1})^{4p}(\log d)^{4p+2}}{m} \end{aligned} \quad (316)$$

The last inequality is because $m \leq d^{D/4}$, so $d^{-D/4} \lesssim \frac{1}{m}$, thus $d^{-D/4}$ can be absorbed by $\frac{V_{p,\epsilon}^2(\kappa^{\mathbb{T}_1})^{4p}(\log d)^{4p+2}}{m}$.

By Jensen's inequality and the Lipschitz property of ℓ , we have

$$\begin{aligned} \widehat{\mathcal{L}}(f_{\tilde{\Theta}^*}) - \widehat{\mathcal{L}}(f^*) &= \frac{1}{n} \sum_{i=1}^n [\ell(f_{\tilde{\Theta}^*}(\mathbf{x}_{i+n}), y_i) - \ell(f^*(\mathbf{x}_{i+n}), y_i)] \\ &\leq \mathbb{L} \frac{1}{n} \sum_{i=1}^n |f_{\tilde{\Theta}^*}(\mathbf{x}_{i+n}) - f^*(\mathbf{x}_{i+n})| \\ &\leq \mathbb{L} \sqrt{\frac{1}{n} \sum_{i=1}^n (f_{\tilde{\Theta}^*}(\mathbf{x}_{i+n}) - f^*(\mathbf{x}_{i+n}))^2} \\ &\lesssim \mathbb{L} \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2(\kappa^{\mathbb{T}_1})^{4p}(\log d)^{4p+2}}{m}} \right). \end{aligned} \quad (317)$$

The proof is complete. \square

D.5.11 PROOF OF LEMMA D.16

Proof. Note that

$$\begin{aligned} \left| \widehat{\mathcal{L}}(f_{\tilde{\Theta}^*}) - \widehat{\mathcal{L}}(f_{\Theta^*}) \right| &= \left| \frac{1}{n} \sum_{i=1}^n [\ell(f_{\tilde{\Theta}^*}(\mathbf{x}_{i+n}), y_i) - \ell(f_{\Theta^*}(\mathbf{x}_{i+n}), y_i)] \right| \\ &\leq \mathbb{L} \sup_{i \in [n]} |f_{\tilde{\Theta}^*}(\mathbf{x}_{i+n}) - f_{\Theta^*}(\mathbf{x}_{i+n})| \\ &\leq \mathbb{L} m \sup_{i \in [n], j \in [m]} \left| \mathbf{a}_j^* \sigma((\tilde{\mathbf{w}}_j^{(\mathbb{T}_1)})^\top \mathbf{x}_{i+n} + b_j) - \mathbf{a}_j^* \sigma((\mathbf{w}_j^{(\mathbb{T}_1)})^\top \mathbf{x}_{i+n} + b_j) \right| \\ &\leq \mathbb{L} m \|\mathbf{a}^*\| \sup_{i \in [n], j \in [m]} \left| \sigma((\tilde{\mathbf{w}}_j^{(\mathbb{T}_1)})^\top \mathbf{x}_{i+n} + b_j) - \sigma((\mathbf{w}_j^{(\mathbb{T}_1)})^\top \mathbf{x}_{i+n} + b_j) \right|. \end{aligned} \quad (318)$$

By Assumption 2 we know that $\sigma'(0) = 0$, $\sigma''(0) = 1$ and $|\sigma^{(3)}(\cdot)| \leq M$. So we can bound the gap of $\sigma((\tilde{\mathbf{w}}_j^{(\mathbb{T}_1)})^\top \mathbf{x} + b_j) - \sigma((\mathbf{w}_j^{(\mathbb{T}_1)})^\top \mathbf{x} + b_j)$ by the following lemma.

Lemma D.36. *If an activation function σ satisfies $\sigma'(0) = 0$, $\sigma''(0) = 1$ and $|\sigma^{(3)}(\cdot)| \leq M$, then we have for all $x, y \in \mathbb{R}$:*

$$|\sigma(x) - \sigma(y)| \leq \left[\frac{1}{2}(|x| + |y|) + \frac{M}{2}(|x|^2 + |y|^2) \right] \cdot |x - y|. \quad (319)$$

Proof of Lemma D.36. Define $r(t) := \sigma(t) - \frac{1}{2}t^2 - \sigma(0)$. Since $\sigma''(0) = 1$ and $\|\sigma^{(3)}\|_\infty \leq M$, we have that σ'' is M -Lipschitz. Hence

$$|r'(t)| = |\sigma'(t) - t| = \left| \int_0^t (\sigma''(s) - 1) ds \right| \leq \frac{M}{2}|t|^2. \quad (320)$$

By the mean value theorem,

$$|r(x) - r(y)| \leq \sup_{\xi \in [x, y]} |r'(\xi)| \cdot |x - y| \leq \frac{M}{2} \max\{|x|, |y|\}^2 |x - y|. \quad (321)$$

Therefore,

$$|\sigma(x) - \sigma(y)| \leq \frac{1}{2}|x^2 - y^2| + |r(x) - r(y)| \leq \frac{1}{2}(|x| + |y|)|x - y| + \frac{M}{2} \max\{|x|, |y|\}^2 |x - y|. \quad (322)$$

Finally, since $\max\{a, b\}^2 \leq a^2 + b^2$, we get the conclusion. \square

Back to the proof of Lemma D.16, under event \mathcal{E}_μ , from Lemma D.9 notice that for all \mathbf{x} in $\{\mathbf{x}_i\}_{i=n+1}^{2n}$, we have $(\mathbf{w}_j^{(\text{T}_1)})^\top \mathbf{x}$ and $\mathbf{w}_{m-j}^{(\text{T}_1)\top} \mathbf{x}$ are truncated, so $\sigma((\mathbf{w}_j^{(\text{T}_1)})^\top \mathbf{x} + b_j)$ and $\sigma((\mathbf{w}_{m-j}^{(\text{T}_1)})^\top \mathbf{x} + b_{m-j})$ are bounded by a universal constant. Respectively, we have $(\widetilde{\mathbf{w}}_j^{(\text{T}_1)})^\top \mathbf{x}$ and $\widetilde{\mathbf{w}}_{m-j}^{(\text{T}_1)\top} \mathbf{x}$ are truncated, so $\sigma((\widetilde{\mathbf{w}}_j^{(\text{T}_1)})^\top \mathbf{x} + b_j)$ and $\sigma((\widetilde{\mathbf{w}}_{m-j}^{(\text{T}_1)})^\top \mathbf{x} + b_{m-j})$ are also bounded by a universal constant. Thus we have

$$\begin{aligned} |\widehat{\mathcal{L}}(f_{\Theta^*}) - \widehat{\mathcal{L}}(f_{\widetilde{\Theta}^*})| &\leq Lm \|\mathbf{a}^*\| \sup_{i \in [n], j \in [m]} |\sigma((\widetilde{\mathbf{w}}_j^{(\text{T}_1)})^\top \mathbf{x}_{i+n} + b_j) - \sigma((\mathbf{w}_j^{(\text{T}_1)})^\top \mathbf{x}_{i+n} + b_j)| \\ &\lesssim Lm \|\mathbf{a}^*\| \sup_{i \in [n], j \in [m]} |(\widetilde{\mathbf{w}}_j^{(\text{T}_1)})^\top \mathbf{x}_{i+n} - (\mathbf{w}_j^{(\text{T}_1)})^\top \mathbf{x}_{i+n}| \\ &\lesssim Lm \|\mathbf{a}^*\| \sqrt{d} \sup_{i \in [n], j \in [m]} \|\widetilde{\mathbf{w}}_j^{(\text{T}_1)} - \mathbf{w}_j^{(\text{T}_1)}\| \\ &\lesssim Lm \frac{1}{\epsilon_0} \|\mathbf{a}^*\| \sqrt{d} \eta^{\text{T}_1} \|\Delta\| \end{aligned} \quad (323)$$

where Δ is defined in Lemma C.7, and from this lemma, we have the bound

$$\|\Delta\| \lesssim \left(\frac{5}{4}\right)^{\text{T}_1 - 1} \cdot Md^2 \cdot \epsilon_0^2 \quad (324)$$

By Assumption 7, we know that $\epsilon_0 \leq \frac{1}{C_\epsilon m n^{1/2} d^{7/2}} \left(\frac{4}{5}\right)^{\text{T}_1}$, where C_ϵ is a universal constant. Moreover, we have $\|\mathbf{a}^*\|^2 \lesssim \frac{V_{p,\epsilon}^2 (\kappa^{\text{T}_1})^{4p} (\log d)^{4p}}{m}$. Thus, it holds that

$$\begin{aligned} |\widehat{\mathcal{L}}(f_{\widetilde{\Theta}^*}) - \widehat{\mathcal{L}}(f_{\Theta^*})| &\lesssim Lm \frac{1}{\epsilon_0} \|\mathbf{a}^*\| \sqrt{d} \eta^{\text{T}_1} \|\Delta\| \\ &\lesssim Lm \|\mathbf{a}^*\| d \text{T}_1 \cdot \left(\frac{5}{4}\right)^{\text{T}_1 - 1} \cdot Md^2 \cdot \epsilon_0 \\ &\lesssim Lm \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{\text{T}_1})^{4p} (\log d)^{4p}}{m}} d^3 \log d M \frac{1}{C_\epsilon m \sqrt{nd}^{\frac{7}{2}}} \\ &\lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{\text{T}_1})^{4p} (\log d)^{4p}}{m}} \log d \frac{1}{\sqrt{n}} \lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{\text{T}_1})^{4p} (\log d)^{4p}}{m}}. \end{aligned} \quad (325)$$

The proof is complete. \square

D.5.12 PROOF OF LEMMA D.17

Proof. Lemma D.15 implies that

$$\widehat{\mathcal{L}}(f_{\widetilde{\Theta}^*}) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{\text{T}_1})^{4p} (\log d)^{4p+2}}{m}} \right). \quad (326)$$

Lemma D.16 further suggests that

$$\left| \widehat{\mathcal{L}}(f_{\Theta^*}) - \widehat{\mathcal{L}}(f_{\Theta^*}) \right| \lesssim L \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m}}. \quad (327)$$

By combining these two bounds, we easily obtain

$$\left| \widehat{\mathcal{L}}(f_{\Theta^*}) - \widehat{\mathcal{L}}(f^*) \right| \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right). \quad (328)$$

The proof is complete. \square

D.5.13 PROOF OF LEMMA D.18

Proof. From Lemma D.17, we have the following bound

$$\widehat{\mathcal{L}}(f_{\Theta^*}) \lesssim L \left(\sum_{k \leq p} \Delta_k + \sqrt{\frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p+2}}{m}} \right) = J \quad (329)$$

Consider the training objective $F(\mathbf{a}) = \widehat{\mathcal{L}}((\mathbf{a}, \mathbf{b}^{(T_1)}, \mathbf{W}^{(T_1)})) + \frac{1}{2} \beta_2 \|\mathbf{a}\|^2$ with $\beta_2 = \frac{J}{U^2}$.

Let

$$\mathbf{a}^{(\infty)} := \arg \min_{\mathbf{a} \in \mathbb{R}^m} F(\mathbf{a}) \quad (330)$$

be the minimizer of F .

We further denote ϵ_0 is the prescribed optimization tolerance $\|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \leq \epsilon_0$. Furthermore, let $\epsilon_0 = \frac{1}{mL}$.

By optimality of $\mathbf{a}^{(\infty)}$ and the *constructed* comparator \mathbf{a}^* we have

$$F(\mathbf{a}^{(\infty)}) \leq F(\mathbf{a}^*) = \widehat{\mathcal{L}}(\mathbf{a}^*) + \frac{\beta_2}{2} \|\mathbf{a}^*\|_2^2. \quad (331)$$

Here we only use that \mathbf{a}^* is a feasible point with $\|\mathbf{a}^*\| \lesssim U$ and $\widehat{\mathcal{L}}(\mathbf{a}^*) \lesssim J$ (Lemma D.17).

On the other hand,

$$F(\mathbf{a}^{(\infty)}) = \widehat{\mathcal{L}}(\mathbf{a}^{(\infty)}) + \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|^2 \geq \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|^2, \quad (332)$$

so combining with (331) gives

$$\|\mathbf{a}^{(\infty)}\|^2 \leq \|\mathbf{a}^*\|^2 + \frac{2}{\beta_2} \widehat{\mathcal{L}}(\mathbf{a}^*) \lesssim U^2 + \frac{2J}{\beta_2} = U^2 + \frac{2J}{J/U^2} \lesssim U^2. \quad (333)$$

Thus

$$\|\mathbf{a}^{(\infty)}\| \lesssim U. \quad (334)$$

and also

$$\|\mathbf{a}^{(T)}\| \leq \|\mathbf{a}^{(\infty)}\| + \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \leq \epsilon_0 + U \lesssim U \quad (335)$$

$$\widehat{\mathcal{L}}(\mathbf{a}^{(\infty)}) = F(\mathbf{a}^{(\infty)}) - \frac{\beta_2}{2} \|\mathbf{a}^{(\infty)}\|^2 \leq F(\mathbf{a}^*) \lesssim J + \frac{\beta_2}{2} U^2 \lesssim J. \quad (336)$$

Thus we have

$$\begin{aligned} |F(\mathbf{a}^{(T)}) - F(\mathbf{a}^{(\infty)})| &\lesssim \sup_{i \in [n+1, 2n]} \left\{ L \left| \mathbf{a}^{(T)\top} \sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)}) - \mathbf{a}^{(\infty)\top} \sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)}) \right| \right\} \\ &\quad + \beta_2 (\|\mathbf{a}^{(T)}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \\ &\lesssim L \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \cdot \sup_{i \in [n+1, 2n]} \|\sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)})\| \\ &\quad + \beta_2 (\|\mathbf{a}^{(T)}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \end{aligned} \quad (337)$$

By previous proof $\|\sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)})\|_\infty$ can be bounded by an absolute constant for all i under event \mathcal{E} . Thus we can have a loose ℓ^2 norm bound and

$$\begin{aligned} |F(\mathbf{a}^{(T)}) - F(\mathbf{a}^{(\infty)})| &\lesssim L \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \cdot \sup_{i \in [n+1, 2n]} \|\sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)})\| \\ &\quad + \beta_2 (\|\mathbf{a}^{(T)}\| + \|\mathbf{a}^{(\infty)}\|) \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \\ &\lesssim L \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \sqrt{m} + \beta_2 U \|\mathbf{a}^{(T)} - \mathbf{a}^{(\infty)}\| \\ &\lesssim (L\sqrt{m} + \beta_2 U) \epsilon_0 \end{aligned} \quad (338)$$

Denote $\epsilon_{\text{opt}} = (L\sqrt{m} + \beta_2 U) \epsilon_0 = \frac{L\sqrt{m} + \beta_2 U}{mL}$. It is easy to know that $\epsilon_{\text{opt}} = \frac{L\sqrt{m} + \beta_2 U}{mL} = \frac{1}{\sqrt{m}} + \frac{J}{UmL} \lesssim J$. Thus, we obtain

$$\widehat{\mathcal{L}}(\mathbf{a}^{(T)}) \lesssim J. \quad (339)$$

Denote

$$\phi_i := \sigma(\mathbf{W}^{(T_1)} \mathbf{x}_i + \mathbf{b}^{(T_1)}) \in \mathbb{R}^m, \quad i = n+1, \dots, 2n. \quad (340)$$

By Assumption 3 on the loss function ℓ , the function $F(\mathbf{a}) - \frac{\beta_2}{2} \|\mathbf{a}\|^2$ is convex and satisfies the gradient Lipschitz condition

$$\begin{aligned} \|\nabla \widehat{\mathcal{L}}(\mathbf{a}_1) - \nabla \widehat{\mathcal{L}}(\mathbf{a}_2)\| &= \left\| \frac{1}{n} \sum_{i=n+1}^{2n} (\ell'(\mathbf{a}_1^\top \phi_i) - \ell'(\mathbf{a}_2^\top \phi_i)) \phi_i \right\| \\ &\leq \frac{1}{n} \sum_{i=n+1}^{2n} |\ell'(\mathbf{a}_1^\top \phi_i) - \ell'(\mathbf{a}_2^\top \phi_i)| \|\phi_i\| \\ &\leq \frac{L}{n} \sum_{i=n+1}^{2n} |(\mathbf{a}_1 - \mathbf{a}_2)^\top \phi_i| \|\phi_i\| \\ &= L \left(\frac{1}{n} \sum_{i=n+1}^{2n} \|\phi_i\|^2 \right) \|\mathbf{a}_1 - \mathbf{a}_2\|. \end{aligned} \quad (341)$$

By previous arguments, $\|\sigma(\mathbf{W}^{(T)} \mathbf{x}_i + \mathbf{b}^{(T)})\|_\infty$ can be bounded by an absolute constant for all $i \in [n+1, 2n]$ under event \mathcal{E} . In addition, we have $\frac{1}{n} \sum_i \|\phi_i\|^2 = \text{Tr}(\frac{1}{n} \sum_i \phi_i \phi_i^\top) \leq \sup_i \|\phi_i\|^2$. Thus the Lip constant for F satisfies $L_{\text{lip}} \lesssim Lm + J/U^2$. Therefore, by Lemma E.13, if the step size $\eta_2 \leq \frac{1}{L_{\text{lip}}}$, gradient descent will approximate $\mathbf{a}^{(\infty)}$ to arbitrary accuracy in

$$T_2 \gtrsim \frac{1}{\eta_2 \beta_2} \log \left(\frac{\|\mathbf{a}^{(T_1)} - \mathbf{a}^{(\infty)}\|^2}{\epsilon_0^2} \right) \quad (342)$$

steps.

Note that when the step size $\eta_2 = \frac{1}{Lm + J/U^2}$

$$\frac{1}{\eta_2 \beta_2} \lesssim \frac{Lm + J/U^2}{\beta_2} \lesssim \frac{U^2 Lm}{J} \quad (343)$$

Using the assumption that $m \lesssim d^{\frac{D}{4}}$ and $\|\mathbf{a}^{(\infty)}\| \lesssim U, \|\mathbf{a}^{(T_1)}\| \lesssim 1$, we can further simplify:

$$\frac{1}{\eta_2 \beta_2} \log \left(\frac{\|\mathbf{a}^{(T_1)} - \mathbf{a}^{(\infty)}\|^2}{\epsilon_0^2} \right) \lesssim \frac{U^2 Lm}{J} \log(Um) \quad (344)$$

Thus, we conclude that gradient descent will approximate $\mathbf{a}^{(\infty)}$ to arbitrary accuracy in at most

$$T_2 \gtrsim \frac{U^2 Lm}{J} \log(Um) \quad (345)$$

steps. The proof is complete. \square

D.5.14 PROOF OF LEMMA D.19

Proof. Let $\mathcal{F}(\mathbf{b}^{(T)}, \mathbf{W}^{(T)}) := \{f(\mathbf{x}) = \mathbf{a}^\top \sigma(\mathbf{W}^{(T)} \mathbf{x} + \mathbf{b}^{(T)}), \Theta : \|\mathbf{a}\| \lesssim U\}$. Note that under \mathcal{E} we have $f_{\Theta^{(T)}} \in \mathcal{F}$. We first define a new event for $\mathbf{x} \sim \mathcal{N}(0, I)$ under fixed $\mathbf{W}^{(T)}$ and $\mathbf{b}^{(T)}$

$$\mathcal{M}_\alpha^{(j)} := \left\{ |\mathbf{h}_n(\mathbf{w}_j^{(0)})^\top \mathbf{x}| \leq \|\mathbf{h}_n(\mathbf{w}_j^{(0)})\| \sqrt{4D \log d} \right\} \quad (346)$$

Define the finite-width event for the initialization as the intersection $\mathcal{M}_\alpha := \bigcap_{j=1}^{m/2} \mathcal{M}_\alpha^{(j)}$. It is straightforward to have $\mathbb{P}(\mathcal{M}_\alpha | \mathbf{H}_\ell, \{\mathbf{w}_j^{(0)}\}_{j=0}^m) \geq 1 - \mathcal{O}(d^{-D})$. We also define another event for \mathbf{x} in order to bound its first r components

$$\mathcal{M}_\beta := \left\{ \sup_{k \in [r]} |(\mathbf{x})_k| \leq \sqrt{4D \log d} \right\} \quad (347)$$

Applying a union bound gives immediately $\mathbb{P}(\mathcal{M}_\beta) \geq 1 - \mathcal{O}(d^{-D})$.

Denote $\mathcal{M} = \mathcal{M}_\alpha \cap \mathcal{M}_\beta$. Define the truncated loss

$$\ell^{\mathcal{M}}(t, y) := \ell(t, y) \mathbf{1}\{\mathbf{x} \in \mathcal{M}\}. \quad (348)$$

$$\widehat{\mathcal{L}}^{\mathcal{M}}(f) := \frac{1}{n} \sum_{i=1}^n \ell^{\mathcal{M}}(f(\mathbf{x}_{i+n}), y_{i+n}), \quad \mathcal{L}^{\mathcal{M}}(f) := \mathbb{E}[\ell^{\mathcal{M}}(f(\mathbf{x}), y)]. \quad (349)$$

By the symmetrization lemma,

$$\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}^{\mathcal{M}}(f) - \mathcal{L}^{\mathcal{M}}(f)| \leq 2\mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \xi_i \ell^{\mathcal{M}}(f(\mathbf{x}_{i+n}), y_{i+n}) \right] \quad (350)$$

where $\{\xi_i\}_{i=1}^n$ are i.i.d. Rademacher signs independent of $\{(\mathbf{x}_{i+n}, y_{i+n})\}_{i=1}^n$. Denote $S := \{(\mathbf{x}_{i+n}, y_{i+n})\}_{i=1}^n$. For fixed $(\mathbf{x}_{i+n}, y_{i+n})$, the mapping $t \mapsto \ell^{\mathcal{M}}(t, y_{i+n})$ has Lipschitz constant in t satisfying

$$|\ell^{\mathcal{M}}(t, y) - \ell^{\mathcal{M}}(s, y)| = |\ell(t, y) - \ell(s, y)| \cdot \mathbf{1}\{\mathbf{x} \in \mathcal{M}\} \leq L|t - s|. \quad (351)$$

Thus, by the Ledoux–Talagrand contraction inequality Lemma E.5 applied coordinatewise to the maps $t \mapsto \ell^{\mathcal{M}}(t, y_{i+n})$,

$$\mathbb{E}_\xi \left[\sup_f \frac{1}{n} \sum_{i=1}^n \xi_i \ell^{\mathcal{M}}(f(\mathbf{x}_{i+n}), y_{i+n}) \right] \leq L\mathbb{E}_\xi \left[\sup_f \frac{1}{n} \sum_{i=1}^n \xi_i f(\mathbf{x}_{i+n}) \right] = L\widehat{\text{Rad}}_n(\mathcal{F}). \quad (352)$$

Recalling $\text{Rad}_n(\mathcal{F}) = \mathbb{E}[\widehat{\text{Rad}}_n(\mathcal{F})]$, we obtain the in-expectation bound

$$\mathbb{E}_S \left[\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}^{\mathcal{M}}(f) - \mathcal{L}^{\mathcal{M}}(f)| \right] \leq 2L \text{Rad}_n(\mathcal{F}). \quad (353)$$

Furthermore, for \mathbf{x} under the event \mathcal{M} , by the same argument as in Lemma D.9, we have $|(\mathbf{w}_j^{(T)})^\top \mathbf{x}| \lesssim 1$ for all j , thus we have

$$\left\| \sigma(\mathbf{W}^{(T)} \mathbf{x} + \mathbf{b}^{(T)}) \right\| \lesssim \sqrt{m} \quad (354)$$

$$|f(\mathbf{x})| = \mathbf{a}^\top \sigma(\mathbf{W}^{(T)} \mathbf{x} + \mathbf{b}^{(T)}) \lesssim U\sqrt{m}. \quad (355)$$

$$\begin{aligned} \ell^{\mathcal{M}}(f(\mathbf{x}), y; \mathbf{x}) &= \ell(f(\mathbf{x}), y) \mathbf{1}\{\mathbf{x} \in \mathcal{M}\} \\ &\leq (\ell(f(\mathbf{x}), y) - \ell(f^*(\mathbf{x}), y)) \mathbf{1}\{\mathbf{x} \in \mathcal{M}\} \\ &\leq L|f(\mathbf{x}) - f^*(\mathbf{x})| \mathbf{1}\{\mathbf{x} \in \mathcal{M}\} \\ &\lesssim LU\sqrt{m} + L(\log d)^p \lesssim LU\sqrt{m}. \end{aligned} \quad (356)$$

We now give the following two high probability events in order to bound the gap between empirical loss and population loss.

Definition D.37 (Event \mathcal{E}_τ and \mathcal{E}). *We say the event \mathcal{E}_τ holds if*

$$\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}^{\mathcal{M}}(f) - \mathcal{L}^{\mathcal{M}}(f)| \lesssim \left(\text{L Rad}_n(\mathcal{F}) + \text{LU} \sqrt{m} \sqrt{\frac{\log d}{n}} \right), \quad (357)$$

Using the standard empirical-to-population deviation bound by Lemma E.4, we have \mathcal{E}_τ holds with probability at least $1 - \mathcal{O}(d^{-D})$.

Furthermore, we define event

$$\mathcal{E} := \mathcal{E}_\mu \cap \mathcal{E}_\tau. \quad (358)$$

Hence, by the preceding probability bounds (and a union bound), \mathcal{E} occurs with probability at least $1 - \mathcal{O}(d^{-D/2})$.

By Lemma E.6 we have $\text{Rad}_n(\mathcal{F}) \lesssim U \sqrt{\frac{m}{n}}$. Thus

$$\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}^{\mathcal{M}}(f) - \mathcal{L}^{\mathcal{M}}(f)| \lesssim \text{LU} \sqrt{\frac{m \log d}{n}}. \quad (359)$$

We notice that under the event \mathcal{E}_μ , \mathcal{M} holds for all \mathbf{x}_i , $i \in [n+1, 2n]$, thus $\widehat{\mathcal{L}}^{\mathcal{M}}(f) = \widehat{\mathcal{L}}(f)$. For the other term, we have

$$\begin{aligned} \mathcal{L}^{\mathcal{M}}(f) - \mathcal{L}(f) &= \mathbb{E}[\ell(f(\mathbf{x}), y)] - \mathbb{E}[\ell^{\mathcal{M}}(f(\mathbf{x}), y; \mathbf{x})] \\ &= \mathbb{E}[\ell(f(\mathbf{x}), y) \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] \\ &\leq \text{L} \mathbb{E}[|f(\mathbf{x}) - f^*(\mathbf{x})| \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] \\ &\leq \text{L} (\mathbb{E}[|f(\mathbf{x})| \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}] + \mathbb{E}[|f^*(\mathbf{x})| \mathbf{1}\{\mathbf{x} \notin \mathcal{M}\}]) \\ &\lesssim \text{L} \left(\sqrt{\mathbb{E}[f(\mathbf{x})^2 \mathbb{P}\{\mathbf{x} \notin \mathcal{M}\}] + \sqrt{\mathbb{E}[(f^*(\mathbf{x}))^2 \mathbb{P}\{\mathbf{x} \notin \mathcal{M}\}]} \right). \end{aligned} \quad (360)$$

Recall that it has been assumed that the activation function has bounded third derivatives as in Assumption 2. Therefore, there exists an absolute constant $C_\sigma > 0$ such that

$$|\sigma(u)| \leq C_\sigma(1 + |u|^3) \quad \text{for all } u \in \mathbb{R}. \quad (361)$$

Via a crude moment estimation bound, we can see the right side is $\lesssim \text{LU} d^{-\frac{D}{8}}$. Thus, we have the following final estimation

$$\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}(f) - \mathcal{L}(f)| \lesssim \text{LU} \sqrt{\frac{m \log d}{n}} + \text{LU} d^{-D/8}. \quad (362)$$

Since we have assume $n \leq d^{D/4}$, we have $\sup_{f \in \mathcal{F}} |\widehat{\mathcal{L}}(f) - \mathcal{L}(f)| \lesssim \text{LU} \sqrt{\frac{m \log d}{n}}$. Plug in the formula for U and we get our desired result. \square

D.5.15 PROOF OF LEMMA D.21

Proof. Let $\epsilon_{\min} := \min_{0 \leq k \leq p} \epsilon_k$. Because $x \mapsto x^{-\beta}$ is strictly decreasing on $(0, \infty)$, we have $\max_{0 \leq k \leq p} \epsilon_k^{-\beta} = \epsilon_{\min}^{-\beta}$. Also note that $\sum_{k=0}^p \epsilon_k \geq (p+1)\epsilon_{\min}$, with equality if and only if all ϵ_k are equal. Consequently,

$$\Phi(\{\epsilon_k\}) \geq A((p+1)\epsilon_{\min} + (C_m + C_n)\epsilon_{\min}^{-\beta}) \quad (363)$$

and equality holds precisely at symmetric choices $\epsilon_0 = \dots = \epsilon_p$. Therefore the problem reduces to minimizing the strictly convex one-variable function

$$g(\epsilon) := (p+1)\epsilon + (C_m + C_n)\epsilon^{-\beta} \quad (\epsilon > 0). \quad (364)$$

Since

$$g'(\epsilon) = (p+1) - \beta(C_m + C_n)\epsilon^{-(\beta+1)} \quad (365)$$

and

$$g''(\epsilon) = \beta(\beta + 1)(C_m + C_n)\epsilon^{-(\beta+2)} > 0, \quad (366)$$

there is a unique critical point, which is the global minimizer

$$\epsilon^* = \left(\frac{\beta(C_m + C_n)}{p + 1} \right)^{\frac{1}{\beta+1}} \quad (367)$$

Substituting back and using $(C_m + C_n)(\epsilon^*)^{-\beta} = \frac{p+1}{\beta}\epsilon^*$ yields

$$\min g = (\beta + 1)\beta^{-\frac{\beta}{\beta+1}}(p + 1)^{\frac{\beta}{\beta+1}}(C_m + C_n)^{\frac{1}{\beta+1}}. \quad (368)$$

Multiplying by LA gives the stated expression for $\min \Phi$.

Finally, as p is fixed, its contribution can be absorbed into the \lesssim constant, also note that $1 \leq (\beta + 1)\beta^{-\frac{\beta}{\beta+1}} \leq 2$ thus we get

$$\begin{aligned} \min_{\{\epsilon_k > 0\}} \Phi &= A \cdot (\beta + 1)\beta^{-\frac{\beta}{\beta+1}}(p + 1)^{\frac{\beta}{\beta+1}}(C_m + C_n)^{\frac{1}{\beta+1}} \\ &\lesssim A(C_m + C_n)^{\frac{1}{\beta+1}} \\ &\lesssim \kappa^{2pT_1} \log^{2p+1} d \left(\sqrt{\frac{1}{m}} + \sqrt{\frac{1}{n}} \right)^{\frac{1}{\beta+1}}. \end{aligned} \quad (369)$$

The proof is complete. □

D.5.16 PROOF OF LEMMA D.23

Proof. Let $\tau := \sqrt{\frac{\log d}{\log \kappa}}$, so $T_1 = \lceil \tau \rceil \in [\tau, \tau + 1]$. Then

$$\kappa^{2T_1} = \exp(2T_1 \log \kappa) \in \left[d^{2\sqrt{\frac{\log \kappa}{\log d}}}, \kappa^2 d^{2\sqrt{\frac{\log \kappa}{\log d}}} \right], \quad (370)$$

and therefore

$$T_1^2 d \log d \cdot \kappa^{2T_1} \lesssim \kappa^2 (\log d)^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}} \quad (371)$$

Moreover, since $\frac{1}{T_1} \leq \frac{1}{\tau} = \sqrt{\frac{\log \kappa}{\log d}}$, we have

$$d^{1+1/T_1} \kappa^2 \leq \kappa^2 d^{1+\sqrt{\frac{\log \kappa}{\log d}}} \quad (372)$$

and combining the two bounds yields

$$T_1^2 d \log d \kappa^{2T_1} + d^{1+1/T_1} \kappa^2 \lesssim \kappa^2 \log(d) d^{1+2\sqrt{\frac{\log \kappa}{\log d}}} \lesssim d^{1+2\sqrt{\frac{\log \kappa}{\log d}}} (\log d)^2. \quad (373)$$

The proof is complete. □

D.5.17 PROOF OF LEMMA D.25

Proof. Plugging in U and J we get

$$\begin{aligned} T_2 &= \frac{U^2 L m}{J} \log(Um) \\ &\lesssim \frac{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p}}{m} L m \log \left(\sqrt{V_{p,\epsilon}^2 (\kappa^{T_1})^{4p} (\log d)^{4p} m} \right) \lesssim V_{p,\epsilon}^2 \kappa^{5pT_1} (\log d)^{5p} m \end{aligned} \quad (374)$$

Note that we have $V_{k,\epsilon_k} \lesssim \epsilon_k^{-\beta}$ by Definition D.7, and also note that $\epsilon_k = \left(\frac{\beta(C_m+C_n)}{p+1}\right)^{\frac{1}{\beta+1}}$ from Lemma D.21, thus

$$\begin{aligned} T_2 &\lesssim V_{p,\epsilon}^2 \kappa^{5pT_1} (\log d)^{5p} m \lesssim \epsilon_k^{-\beta} \kappa^{5pT_1} (\log d)^{5p} m \\ &\lesssim \left(\frac{p+1}{\beta(\sqrt{\frac{1}{m}} + \sqrt{\frac{1}{n}})}\right)^{\frac{\beta}{\beta+1}} \kappa^{5pT_1} (\log d)^{5p} m \\ &\lesssim \min(m, n)^{\frac{\beta}{2(\beta+1)}} \kappa^{5pT_1} (\log d)^{5p} m \lesssim \kappa^{5pT_1} (\log d)^{5p} m^2 \end{aligned} \quad (375)$$

Moreover plugging $m = (\log d)^{4p(\beta+1)+1} d^{4p(\beta+1)\sqrt{\frac{\log \kappa}{\log d}}}$ and $T_1 := \lceil \sqrt{\frac{\log d}{\log \kappa}} \rceil$ in, we get

$$T_2 \lesssim \kappa^{5pT_1} (\log d)^{5p} m^2 \lesssim \kappa^{5p} (\log d)^{8p\beta+13p+2} d^{(8p\beta+13p)\sqrt{\frac{\log \kappa}{\log d}}}. \quad (376)$$

The proof is complete. \square

D.5.18 PROOF OF LEMMA D.26

Proof. By Lemma D.25 we have

$$T_2 \lesssim \kappa^{5p} (\log d)^{8p\beta+13p+2} d^{(8p\beta+13p)\sqrt{\frac{\log \kappa}{\log d}}}. \quad (377)$$

Plugging $m = (\log d)^{4p(\beta+1)+1} d^{4p(\beta+1)\sqrt{\frac{\log \kappa}{\log d}}}$, $n = C\kappa^2 (\log d)^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}}$ in we get

$$\begin{aligned} nmd \cdot T &= nmd \cdot (T_1 + T_2) \\ &\lesssim \kappa^2 (\log d)^2 d^{1+2\sqrt{\frac{\log \kappa}{\log d}}} (\log d)^{4p(\beta+1)+1} d^{4p(\beta+1)\sqrt{\frac{\log \kappa}{\log d}}} \kappa^{5p} (\log d)^{8p\beta+13p+2} d^{(8p\beta+13p)\sqrt{\frac{\log \kappa}{\log d}}} \\ &\lesssim \kappa^{5p+2} (\log d)^{12p\beta+17p+5} d^{2+(12p\beta+17p+2)\sqrt{\frac{\log \kappa}{\log d}}} \end{aligned} \quad (378)$$

$$(379)$$

The proof is complete. \square

E TECHNICAL BACKGROUND

E.1 HERMITE POLYNOMIALS

We briefly introduce Hermite Polynomials and Hermite Tensors.

Definition E.1 (1D Hermite polynomials). *The k -th normalized probabilist's Hermite polynomial, $h_k : \mathbb{R} \rightarrow \mathbb{R}$, is the degree k polynomial defined as*

$$h_k(x) = \frac{(-1)^k}{\sqrt{k!}} \frac{d^k \mu_\beta}{dx^k}(x), \quad (380)$$

where $\mu_\beta(x) = \exp(-x^2/2)/\sqrt{2\pi}$ is the density of the standard Gaussian.

The first such Hermite polynomials are

$$h_0(z) = 1, h_1(z) = z, h_2(z) = \frac{z^2 - 1}{\sqrt{2}}, h_3(z) = \frac{z^3 - 3z}{\sqrt{6}}, \dots \quad (381)$$

Denote $\beta = \mathcal{N}(0, 1)$ to be the standard Gaussian in 1D. A key fact is that the normalized Hermite polynomials form an orthonormal basis of $L^2(\beta)$; that is $\mathbb{E}_{x \sim \beta}[h_j(x)h_k(x)] = \delta_{jk}$.

The multidimensional analogs of the Hermite polynomials are Hermite tensors.

Definition E.2 (Hermite tensors). *The k -th Hermite tensor $He_k : \mathbb{R}^d \rightarrow (\mathbb{R}^d)^{\otimes k}$ is*

$$He_k(\mathbf{x}) := \frac{(-1)^k}{\sqrt{k!}} \frac{\nabla^k \mu_\gamma(\mathbf{x})}{\mu_\gamma(\mathbf{x})}, \quad (382)$$

where $\mu_\gamma(\mathbf{x}) = \exp(-\frac{1}{2}\|\mathbf{x}\|^2)/(2\pi)^{d/2}$ is the density of the d -dimensional Gaussian.

The Hermite tensors form an orthonormal basis of $L^2(\gamma)$; that is, for any $f \in L^2(\gamma)$, one can write the Hermite expansion

$$f(\mathbf{x}) = \sum_{k \geq 0} \langle \mathbf{C}_k(f), H e_k(\mathbf{x}) \rangle \quad \text{where} \quad \mathbf{C}_k(f) := \mathbb{E}_{\mathbf{x} \sim \gamma} [f(\mathbf{x}) H e_k(\mathbf{x})] \in (\mathbb{R}^d)^{\otimes k}. \quad (383)$$

E.2 UNIFORM GENERALIZATION BOUNDS

Definition E.3. *The empirical Rademacher complexity of a function class \mathcal{F} on finite samples is defined as*

$$\widehat{\text{Rad}}_n(\mathcal{F}) = \mathbb{E}_\xi \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \xi_i f(X_i) \right] \quad (384)$$

where $\xi_1, \xi_2, \dots, \xi_n$ are i.i.d. Rademacher random variables $\mathbb{P}(\xi_i = 1) = \mathbb{P}(\xi_i = -1) = \frac{1}{2}$. Let $\text{Rad}_n(\mathcal{F}) = \mathbb{E}[\widehat{\text{Rad}}_n(\mathcal{F})]$ be the population Rademacher complexity.

Then we recall the uniform law of large number via Rademacher complexity, which can be found in Theorem 3.3 in Mohri et al. (2018)

Lemma E.4. *Assume that f ranges in $[0, R]$ for all $f \in \mathcal{F}$. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the i.i.d. training set $S = \{X_1, \dots, X_n\}$ we have*

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(X_i) - \mathbb{E}f(X) \right| \leq 2 \text{Rad}_n(\mathcal{F}) + R \sqrt{\frac{\log(4/\delta)}{n}} \quad (385)$$

Then we recall the contraction Lemma in (Vershynin, 2018, Exercise 6.7.7) to compute Rademacher complexity.

Lemma E.5 (Contraction Lemma). *Let $\varphi_i : \mathbb{R} \mapsto \mathbb{R}$ with $i = 1, \dots, n$ be β -Lispchitz continuous. Then*

$$\frac{1}{n} \mathbb{E}_\xi \sup_{f \in \mathcal{F}} \sum_{i=1}^n \xi_i \varphi_i \circ f(x_i) \leq \beta \widehat{\text{Rad}}_n(\mathcal{F}) \quad (386)$$

The next Lemma estimates the Rademacher complexity.

Lemma E.6. *Let $f(\mathbf{x}) = \mathbf{a}^\top \sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$ be our two layer neural network with our previous assumptions on the activation functions in Assumption 2. Let*

$$\mathcal{F} = \{f_\Theta : \|\mathbf{a}\| \leq B_a, \|\mathbf{w}_j\| \leq C, \|\mathbf{b}\|_\infty \leq C\} \quad (387)$$

where C is a universal constant. Then we have the following estimation $\text{Rad}_n(\mathcal{F}) \lesssim \frac{B_a \sqrt{m}}{\sqrt{n}}$.

Proof.

$$\begin{aligned} \text{Rad}_n(\mathcal{F}) &= \mathbb{E}_{\mathbf{x}, \xi} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_i \xi_i (\mathbf{a}^\top \sigma(\mathbf{W}\mathbf{x}_i + \mathbf{b})) \right| \right] \\ &\leq \frac{B_a}{n} \mathbb{E}_{\mathbf{x}, \xi} \left[\sup_{f \in \mathcal{F}} \left\| \sum_i \xi_i \sigma(\mathbf{W}\mathbf{x}_i + \mathbf{b}) \right\| \right] \\ &\leq \frac{B_a}{n} \sqrt{\mathbb{E}_{\mathbf{x}, \xi} \left[\sup_{f \in \mathcal{F}} \left\| \sum_i \xi_i \sigma(\mathbf{W}\mathbf{x}_i + \mathbf{b}) \right\|^2 \right]} \\ &= \frac{B_a}{\sqrt{n}} \sqrt{\mathbb{E}_{\mathbf{x}} \left[\sum_{j=1}^m \sigma(\mathbf{w}_j^\top \mathbf{x} + b_j)^2 \right]} \lesssim \frac{B_a \sqrt{m}}{\sqrt{n}} \end{aligned} \quad (388)$$

□

E.3 UNIVARIATE APPROXIMATION

We will consider the following activation function

$$\sigma(t) = \begin{cases} 2|t| - 1, & |t| \geq 1, \\ t^2, & |t| < 1. \end{cases} \quad (389)$$

Lemma E.7. *There exists $v_0(a, b)$, supported on $\{\pm 1\} \times [2, 3]$, such that for any $|z| \leq 1$*

$$\mathbb{E}_{a,b}[v_0(a, b)\sigma(az + b)] = 1, \quad \sup_{a,b} |v(a, b)| \lesssim 1. \quad (390)$$

Proof. Let $v_0(a, b) = 12 \cdot \mathbf{1}_{a=1}(b - \frac{5}{2}) \cdot \frac{\mathbf{1}_{b \in [2, 3]}}{\mu(b)}$. Then, since $z + b \geq 1$,

$$\mathbb{E}_{a,b}[v_0(a, b)\sigma(az + b)] = 6 \int_2^3 (b - \frac{5}{2})\sigma(z + b)db \quad (391)$$

$$= 6 \int_2^3 (b - \frac{5}{2})(2z + 2b - 1)db \quad (392)$$

$$= z \cdot 6 \int_2^3 (b - \frac{5}{2})db + 6 \int_2^3 (b - \frac{5}{2})(2b - 1)db \quad (393)$$

$$= 1. \quad (394)$$

The proof is complete. \square

Lemma E.8. *There exists $v_1(a, b)$, supported on $\{\pm 1\} \times [2, 3]$, such that for any $|z| \leq 1$*

$$\mathbb{E}_{a,b}[v_1(a, b)\sigma(az + b)] = z, \quad \sup_{a,b} |v(a, b)| \lesssim 1. \quad (395)$$

Proof. Let $v_0(a, b) = \mathbf{1}_{a=1}(-24b + 61) \cdot \frac{\mathbf{1}_{b \in [2, 3]}}{\mu(b)}$. Then, since $z + b \geq 1$,

$$\mathbb{E}_{a,b}[v_1(a, b)\sigma(az + b)] = \frac{1}{2} \int_2^3 (-24b + 61)\sigma(z + b)db \quad (396)$$

$$= \frac{1}{2} \int_2^3 (-24b + 61)(2z + 2b - 1)db \quad (397)$$

$$= z \int_2^3 (-24b + 61)db + \frac{1}{2} \int_2^3 (-24b + 61)(2b - 1)db \quad (398)$$

$$= z. \quad (399)$$

The proof is complete. \square

Lemma E.9. *There exists $v_2(a, b)$, supported on $\{\pm 1\} \times [-2, 3]$, such that for any $|z| \leq 1$*

$$\mathbb{E}_{a,b}[v_2(a, b)\sigma(az + b)] = z^2, \quad \sup_{a,b} |v(a, b)| \lesssim 1. \quad (400)$$

Proof. First, see that

$$\int_{-2}^2 \sigma(z + b)db = \int_{-2+z}^{2+z} \sigma(b)db \quad (401)$$

$$= \int_{-1}^{-1} (-2b - 1)db + \int_{-1}^1 b^2 db + \int_1^{2+z} (2b - 1)db \quad (402)$$

$$= [-b^2 - b]_{-2+z}^{-1} + \frac{2}{3} + [b^2 - b]_1^{2+z} \quad (403)$$

$$= (z - 2)^2 + (z - 2) + \frac{2}{3} + (z + 2)^2 - (z + 2) \quad (404)$$

$$= 2z^2 + \frac{14}{3}. \quad (405)$$

Let $v_2(a, b) = \mathbf{1}_{a=1} \frac{\mathbf{1}_{b \in [-2, 2]}}{\mu(b)} - \frac{7}{3}v_0(a, b)$ Then

$$\mathbb{E}_{a,b}[v_2(a, b)\sigma(az + b)] = \frac{1}{2} \int_{-2}^2 \sigma(z + b)db - \frac{7}{3} \quad (406)$$

$$= z^2 + \frac{7}{3} - \frac{7}{3} = z^2. \quad (407)$$

The proof is complete. \square

Lemma E.10. Let $v(b) = -\frac{1}{2}k(k-1)(k-2)(1-b)^{k-3} \cdot \frac{\mathbf{1}_{b \in [0, 1]}}{\mu(b)}$. Then

$$\mathbb{E}_b[v_k(b)\sigma(z + b)] = z^k \cdot \mathbf{1}_{z > 0} - \frac{k(k-1)}{2}z^2 - kz - 1. \quad (408)$$

Proof. Plugging in $v_k(b)$ and applying integration by parts yields

$$\mathbb{E}_b[v_k(b)\sigma(z + b)] = \int_0^1 -\frac{1}{2}k(k-1)(k-2)(1-b)^{k-3}\sigma(z + b)db \quad (409)$$

$$= \left[\frac{1}{2}k(k-1)(1-b)^{k-2}\sigma(z + b)\right]_0^1 - \int_0^1 \frac{1}{2}k(k-1)(1-b)^{k-2}\sigma'(z + b)db \quad (410)$$

$$= -\frac{1}{2}k(k-1)\sigma(z) + \left[\frac{1}{2}k(1-b)^{k-1}\sigma'(z + b)\right]_0^1 - \int_0^1 \frac{1}{2}k(1-b)^{k-1}\sigma''(z + b)db \quad (411)$$

$$= -\frac{1}{2}k(k-1)\sigma(z) - \frac{1}{2}k\sigma'(z) - \int_0^1 k(1-b)^{k-1}\mathbf{1}_{|z+b| \leq 1}db \quad (412)$$

When $1 \geq z > 0$, we have

$$-\int_0^1 k(1-b)^{k-1}\mathbf{1}_{|z+b| \leq 1}db = -\int_0^{1-z} k(1-b)^{k-1}db = [(1-b)^k]_0^{1-z} = z^k - 1. \quad (413)$$

When $-1 \leq z \leq 0$, we have

$$-\int_0^1 k(1-b)^{k-1}\mathbf{1}_{|z+b| \leq 1}db = -\int_0^1 k(1-b)^{k-1}db = -1. \quad (414)$$

Since $z \in [-1, 1]$, we have that $\sigma(z) = z^2$ and $\sigma'(z) = 2z$. Therefore for $z \in [-1, 1]$

$$\mathbb{E}_b[v_k(b)\sigma(z + b)] = z^k \cdot \mathbf{1}_{z > 0} - \frac{k(k-1)}{2}z^2 - kz - 1. \quad (415)$$

The proof is complete. \square

Lemma E.11. There exists $v_k(a, b)$, supported on $\{\pm 1\} \times [-2, 3]$, such that for any $|z| \leq 1$

$$\mathbb{E}_{a,b}[v_k(a, b)\sigma(az + b)] = z^k, \quad \sup_{a,b} |v_k(a, b)| \lesssim \text{poly}(k). \quad (416)$$

Proof. We focus on $k \geq 3$. We have that

$$\mathbb{E}_b[v_k(b)\sigma(z + b)] = z^k \cdot \mathbf{1}_{z > 0} - \frac{k(k-1)}{2}z^2 - kz - 1. \quad (417)$$

$$\mathbb{E}_b[v_k(b)\sigma(-z + b)] = (-z)^k \cdot \mathbf{1}_{z < 0} - \frac{k(k-1)}{2}z^2 + kz - 1. \quad (418)$$

Therefore if k is even

$$\mathbb{E}_b[v(b)\sigma(z + b) + v(b)\sigma(-z + b)] = z^k - k(k-1)z^2 - 2. \quad (419)$$

Let $v_k(a, b) = 2av_k(b) + k(k-1)v_2(a, b) + 2$. Then

$$\mathbb{E}_{a,b}[v_k(a, b)\sigma(az + b)] = \mathbb{E}_b[v_k(b)\sigma(z + b) + v_k(b)\sigma(z - b)] + k(k-1)z^2 + 2 = z^k. \quad (420)$$

If k is odd,

$$\mathbb{E}_b[v(b)\sigma(z + b) - v(b)\sigma(-z + b)] = z^k - 2kz. \quad (421)$$

Let $v_k(a, b) = 2av_k(b) + 2av_1(a, b)$. Then

$$\mathbb{E}_{a,b}[v_k(a, b)\sigma(az + b)] = \mathbb{E}_b[v_k(b)\sigma(z + b) - v_k(b)\sigma(z - b)] + 2kz = z^k. \quad (422)$$

The proof is complete. \square

E.4 CONVEX OPTIMIZATION

Denote $f(\mathbf{x})$ as a C^1 function defined in \mathbb{R}^d . Assume that

- There exists $m > 0$ such that $f(\mathbf{x}) - \frac{m}{2} \|\mathbf{x}\|^2$ is convex.
- $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$.

The following result is standard and can be found in most convex optimization textbooks like Boyd & Vandenberghe (2004).

Lemma E.12. *There exists a unique \mathbf{x}^* such that $f(\mathbf{x}^*) = \inf_{\mathbf{x}} f(\mathbf{x})$. And if we start at the point \mathbf{x}^0 and do gradient descent with learning rate η , if $\eta \leq \frac{1}{m+L}$, then we will get*

$$\|\mathbf{x}^k - \mathbf{x}^*\|^2 \leq c^k \|\mathbf{x}^0 - \mathbf{x}^*\|^2 \quad (423)$$

where $c = 1 - \eta \frac{2mL}{m+L}$.

We compute the minimal number of the required iterations.

Corollary E.13. *Let $\epsilon \in (0, \|\mathbf{x}^0 - \mathbf{x}^*\|)$ be a target squared-distance accuracy. In order to have $\|\mathbf{x}^k - \mathbf{x}^*\|^2 \leq \epsilon^2$ it is enough to have*

$$k \gtrsim \frac{m+L}{\eta mL} \log\left(\|\mathbf{x}^0 - \mathbf{x}^*\|^2 / \epsilon^2\right) \quad (424)$$

The above corollary can be verified from a direct computation.

E.5 SPHERE LEMMAS

Lemma E.14. *Let $\nu \sim \chi(d)$. Then,*

$$\mathbb{E}[\nu^{2k}] = \prod_{j=0}^{k-1} (d+2j) = d(d+2) \cdots (d+2k-2) = \Theta(d^k). \quad (425)$$

Proof. Write $\nu^2 \sim \chi^2(d)$ with density proportional to $x^{\frac{d}{2}-1} e^{-x/2}$ on $x > 0$. Using the Gamma function identity,

$$\mathbb{E}[\nu^{2k}] = \mathbb{E}[(\nu^2)^k] = 2^k \frac{\Gamma(k + \frac{d}{2})}{\Gamma(\frac{d}{2})} = \prod_{j=0}^{k-1} (d+2j). \quad (426)$$

Since $\prod_{j=0}^{k-1} (d+2j) \asymp d^k$ for fixed k (e.g., by Stirling or monotonicity of the ratio), we obtain $\Theta(d^k)$. \square

Lemma E.15. *Let $\bar{\mathbf{w}} \sim \text{Unif}(\mathbb{S}^{d-1})$. Then,*

$$\mathbb{E}[\bar{\mathbf{w}}^{\otimes 2k}] = \frac{\mathbb{E}_{\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)}[\mathbf{w}^{2k}]}{\mathbb{E}_{\nu \sim \chi(d)}[\nu^{2k}]} \quad (427)$$

Proof. This follows from $\mathbf{w} = \nu \bar{\mathbf{w}}$ with $\nu \sim \chi(d)$, $\bar{\mathbf{w}} \sim \text{Unif}(\mathbb{S}^{d-1})$ independent. \square

Corollary E.16. *By Wick/Isserlis's formula for Gaussian moments,*

$$\mathbb{E}_{\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)}[\mathbf{w}^{\otimes 2k}] = (2k-1)!! \text{Sym}(\mathbf{I}^{\otimes k}) \quad (428)$$

Meanwhile, Lemma E.14 shows $\mathbb{E}_{\nu \sim \chi(d)}[\nu^{2k}] = \Theta(d^k)$. Hence overall we have

$$\mathbb{E}[\bar{\mathbf{w}}^{\otimes 2k}] \asymp d^{-k} \text{Sym}(\mathbf{I}^{\otimes k}). \quad (429)$$

Corollary E.17. *Let \mathbf{T} be a symmetric p -tensor with $\dim(\text{span}(\mathbf{T})) = r$. With probability at least $1 - 2d^{-D}$*

$$\|\mathbf{T}(\bar{\mathbf{w}}^{\otimes k})\|_F \lesssim \|\mathbf{T}\|_F \sqrt{\frac{r^{\lfloor \frac{k}{2} \rfloor} (\log d)^k}{d^k}}. \quad (430)$$

Proof. Denote $Y := \|\mathbf{T}(\bar{\mathbf{w}}^{\otimes k})\|_F$. We first bound the second moment. Flatten \mathbf{T} as a matrix $\mathbf{M} \in \mathbb{R}^{d^{p-k} \times d^k}$ so that $\mathbf{T}(\bar{\mathbf{w}}^{\otimes k}) = \mathbf{M}(\bar{\mathbf{w}}^{\otimes k})$. Then

$$\mathbb{E}[Y^2] = \mathbb{E}\|\mathbf{M}(\bar{\mathbf{w}}^{\otimes k})\|_F^2 = \langle \mathbf{M}, \mathbf{M}C_k \rangle = \|\mathbf{M}C_k^{1/2}\|_F^2 \leq \|\mathbf{M}\|_F^2 \|C_k\|_{\text{op}}, \quad (431)$$

where $C_k := \mathbb{E}[\bar{\mathbf{w}}^{\otimes k}(\bar{\mathbf{w}}^{\otimes k})^\top] = \mathbb{E}[\bar{\mathbf{w}}^{\otimes 2k}]$ on the d^k -dimensional k -mode space. Next, from Lemma E.15 and Wick/Isserlis' formula for Gaussian even moments, $C_k = c_{k,d}\text{Sym}(\mathbf{I}^{\otimes k})$ with $c_{k,d} \asymp d^{-k}$. Restricting to the r -dimensional span of \mathbf{T} (on each paired contraction) yields an extra factor $r^{\lfloor k/2 \rfloor}$ from pairing counts. Hence

$$\mathbb{E}[Y^2] \lesssim \frac{r^{\lfloor k/2 \rfloor}}{d^k} \|\mathbf{M}\|_F^2 = \frac{r^{\lfloor k/2 \rfloor}}{d^k} \|\mathbf{T}\|_F^2. \quad (432)$$

We then bound Y with high probability. The map $\bar{\mathbf{w}} \mapsto \mathbf{T}(\bar{\mathbf{w}}^{\otimes k})$ is a homogeneous polynomial of degree k restricted to \mathbb{S}^{d-1} . By (spherical) hypercontractivity for degree- k polynomials on \mathbb{S}^{d-1} , Lemma E.22, for all $q \geq 2$,

$$\|Y\|_{L_{2q}(\mathbb{S}^{d-1})} \leq (C\sqrt{q})^k \|Y\|_{L_2(\mathbb{S}^{d-1})}. \quad (433)$$

Combining with equation (432) gives $\|Y\|_{L_{2q}} \lesssim (C\sqrt{q})^k \|\mathbf{T}\|_F \sqrt{\frac{r^{\lfloor k/2 \rfloor}}{d^k}}$ and applying Markov's inequality to Y^{2q} with $q \asymp D \log d$ yields

$$\mathbb{P}\left(Y \gtrsim \|\mathbf{T}\|_F \sqrt{\frac{r^{\lfloor k/2 \rfloor} (\log d)^k}{d^k}}\right) \leq 2d^{-D} \quad (434)$$

which concludes the proof. \square

E.5.1 PROOF OF LEMMA D.34

Proof. As in the previous proof, flatten \mathbf{T} into $\mathbf{M} \in \mathbb{R}^{d^{p-k} \times d^k}$. Then

$$\mathbb{E}\|\mathbf{T}(\bar{\mathbf{w}}^{\otimes k})\|_F^2 = \|\mathbf{M}C_k^{1/2}\|_F^2 \leq \|\mathbf{M}\|_F^2 \|C_k\|_{\text{op}}, \quad C_k = \mathbb{E}[\bar{\mathbf{w}}^{\otimes 2k}] = c_{k,d}\text{Sym}(\mathbf{I}^{\otimes k}), \quad (435)$$

so $\|C_k\|_{\text{op}} \asymp c_{k,d} \asymp d^{-k}$. On the other hand,

$$\mathbb{E}\langle \mathbf{T}, \bar{\mathbf{w}}^{\otimes p} \rangle^2 = \langle \mathbf{T}, \mathbb{E}[\bar{\mathbf{w}}^{\otimes 2p}] \mathbf{T} \rangle = c_{p,d} \|\mathbf{T}\|_F^2 \quad \text{with} \quad c_{p,d} \asymp d^{-p}. \quad (436)$$

Therefore

$$\mathbb{E}\|\mathbf{T}(\bar{\mathbf{w}}^{\otimes k})\|_F^2 \lesssim \|\mathbf{T}\|_F^2 d^{-k} \lesssim d^{p-k} c_{p,d} \|\mathbf{T}\|_F^2 = d^{p-k} \mathbb{E}\langle \mathbf{T}, \bar{\mathbf{w}}^{\otimes p} \rangle^2 \quad (437)$$

which proves the claim. \square

E.6 SUPPORTING TECHNICAL LEMMAS

Lemma E.18 (Norm Concentration for Gaussian Samples). *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \sim \mathcal{N}(0, \mathbf{I}_d)$ be i.i.d. Gaussian vectors. Then with probability at least $1 - n \exp(-cd)$, we have*

$$\|\mathbf{x}_i\| \leq 2\sqrt{d} \quad \forall i \in [n] \quad (438)$$

where c is some positive universal constant.

Proof. Note that $\|\mathbf{x}\|^2$ follows a chi-squared distribution with freedom d and standard concentration for the chi-squared distribution, for instance Theorem 2.8.2 in Vershynin (2018) gives us

$$\mathbb{P}(\|\mathbf{x}_i\| \geq 2\sqrt{d}) \leq \exp(-cd) \quad (439)$$

with some universal positive constant c . The result follows with a union bound argument. \square

Lemma E.19 (Coordinate concentration parameterized by D). *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathbf{I}_d)$ and fix $r \in [d]$. For any $D > 0$, with probability at least $1 - 2nrD^{-D}$,*

$$\max_{i \in [n], k \leq r} |(\mathbf{x}_i)_k| \leq 2\sqrt{D \log d}. \quad (440)$$

Proof. For $Z \sim \mathcal{N}(0, 1)$, $\mathbb{P}(|Z| \geq t) \leq 2e^{-t^2/2}$. Setting $t = \sqrt{4D \log d}$ gives $\mathbb{P}(|Z| \geq \sqrt{4D \log d}) \leq 2d^{-2D}$. Then we do a union bound over all $i \in [n]$ and $k \leq r$. \square

Let the Hermite expansion of our target be $f(\mathbf{x}) = \sum_{k=0}^p \frac{1}{k!} \langle \mathbf{C}_k, H e_k(\mathbf{x}) \rangle$ where $\mathbf{C}_k \in (\mathbb{R}^d)^{\otimes k}$ is the k -tensor defined by $\mathbf{C}_k := \mathbb{E}_{\mathbf{x}}[\nabla^k f(\mathbf{x})]$ and supported on S^* . Here $H e_k$ is the unnormalized k -th Hermite tensor.

Lemma E.20 (Parseval's Identity).

$$1 = \mathbb{E}_{\mathbf{x}}[f(\mathbf{x})^2] = \sum_{k=0}^p \frac{1}{k!} \|\mathbf{C}_k\|_F^2 \quad (441)$$

As an immediate consequence of Lemma E.20 we have $\|\mathbf{C}_k\|_F^2 \leq k!$. Note that $f(\mathbf{x})$ also has the following tensor expansion

$$f(\mathbf{x}) = \sum_{k \leq p} \langle \mathbf{T}_k, \mathbf{x}^{\otimes k} \rangle \quad (442)$$

The following lemma shows $\|\mathbf{T}_k\|_F \lesssim r^{\frac{p-k}{4}} \lesssim 1$.

Lemma E.21. *There exist $\mathbf{T}_0, \dots, \mathbf{T}_p$ such that*

$$f(\mathbf{x}) = \sum_{k \leq p} \langle \mathbf{T}_k, \mathbf{x}^{\otimes k} \rangle \quad (443)$$

and $\|\mathbf{T}_k\|_F \lesssim r^{\frac{p-k}{4}} \lesssim 1$ for $k \leq p$.

Proof. Note that from the Taylor expansion of $f(\mathbf{x})$ we have

$$\mathbf{T}_k = \frac{\nabla^k f(0)}{k!} = \sum_{j \leq p-k} \frac{\mathbf{C}_{j+k}(H e_j(0))}{k!j!} = \sum_{2j \leq p-k} \frac{(-1)^j (2j-1)! \mathbf{C}_{2j+k}(I^{\otimes j})}{k!(2j)!}. \quad (444)$$

We recall from the consequence of Lemma E.20 that $\|\mathbf{C}_k\|_F^2 \leq k!$ and therefore,

$$\|\mathbf{T}_k\|_F \lesssim \sum_{2j \leq p-k} \|\mathbf{C}_{2j+k}(I^{\otimes j})\| \lesssim r^{\frac{p-k}{4}} \lesssim 1. \quad (445)$$

\square

We shall use the next Gaussian hypercontractivity lemma from Theorem 4.3, Prato & Tubaro (2007).

Lemma E.22. *For any $\ell \in \mathbb{N}$ and $f \in L^2(\gamma)$ to be a degree ℓ polynomial where the input distribution $\gamma = \mathcal{N}(0, \mathbf{I}_d)$, for any $q \geq 2$, we have*

$$\mathbb{E}_{\mathbf{z} \sim \gamma} [f(\mathbf{z})^q] \leq C_{q,\ell} \left(\mathbb{E}_{\mathbf{z} \sim \gamma} [f(\mathbf{z})^2] \right)^{q/2} \quad (446)$$

where we use $C_{q,\ell}$ to denote some universal positive constant that only depends on q, ℓ .

We also need this Spherical hypercontractivity lemma from Beckner (1992); O'Donnell (2014).

Lemma E.23. *Let Y_m be a spherical harmonic of degree m on \mathbb{S}^{d-1} . Then for all $1 < p \leq q < \infty$ we have*

$$\|Y_m\|_{L^q(\mathbb{S}^{d-1})} \leq \left(\frac{q-1}{p-1} \right)^{m/2} \|Y_m\|_{L^p(\mathbb{S}^{d-1})}. \quad (447)$$

In particular, for any polynomial P of degree at most m and $q \geq 2$ we have $\|P\|_{L^q(\mathbb{S}^{d-1})} \leq (q-1)^{\frac{m}{2}} \|P\|_{L^2(\mathbb{S}^{d-1})}$.