

Neural Networks can Learn Representations with Gradient Descent

Alex Damian

Princeton University

AD27@PRINCETON.EDU

Jason D. Lee

Princeton University

JASONLEE@PRINCETON.EDU

Mahdi Soltanolkotabi

University of Southern California

SOLTANOL@USC.EDU

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

Significant theoretical work has established that in specific regimes, neural networks trained by gradient descent behave like kernel methods. However, in practice, it is known that neural networks strongly outperform their associated kernels. In this work, we explain this gap by demonstrating that there is a large class of functions which cannot be efficiently learned by kernel methods but can be easily learned with gradient descent on a two layer neural network outside the kernel regime by learning representations that are relevant to the target task. We also demonstrate that these representations allow for efficient transfer learning, which is impossible in the kernel regime.

Specifically, we consider the problem of learning polynomials which depend on only a few relevant directions, i.e. of the form $f^*(x) = g(Ux)$ where $U : \mathbb{R}^d \rightarrow \mathbb{R}^r$ with $d \gg r$. When the degree of f^* is p , it is known that $n \asymp d^p$ samples are necessary to learn f^* in the kernel regime. Our primary result is that gradient descent learns a representation of the data which depends only on the directions relevant to f^* . This results in an improved sample complexity of $n \asymp d^2$ and enables transfer learning with sample complexity independent of d .

Keywords: neural network, gradient descent, representation learning, transfer learning, kernel

1. Introduction

Crucial to the practical success of deep learning is the ability of gradient-based algorithms to learn good feature representations from the training data and learn simple functions on top of these representations. Despite significant progress towards a theoretical foundation for neural networks, a robust understanding of this unique representation learning capability of gradient descent methods has remained elusive. A major challenge is that due to the highly nonconvex loss landscape, establishing convergence to a global optimum that achieves near zero training loss is challenging. Furthermore, due to the overparameterized nature of modern neural nets (containing many more parameters than training data) the training landscape has many global optima. In fact, there are many global optima with poor generalization performance (Zhang et al., 2016; Liu et al., 2020). This paper thus focuses on answering this intriguing question:

How do gradient-based methods learn feature representations and why do these representations allow for efficient generalization and transfer learning?

The most prominent contemporary approach to understanding neural networks is the linearization or neural tangent kernel (NTK) (Soltanolkotabi et al., 2018; Jacot et al., 2018) technique. The

premise of the linearization method is that the dynamics of gradient descent are well-approximated by gradient descent on a linear regression instance with fixed feature representation. Using this linearization technique, it is possible to prove convergence to a zero training loss point (Soltanolkotabi et al., 2018; Du et al., 2018b, 2019a). However, this technique often requires unrealistic hyper-parameter choices (e.g. small learning rate, large initialization, or wide networks) that force the features to not move and the generalization error with this technique cannot be better than that of a kernel method. Indeed, precise lower bounds show that the NTK solutions do not generalize better than the polynomial kernel (Ghorbani et al., 2019a). As a result this regime of training is also sometimes referred to as the *lazy* regime Chizat et al. (2019).¹ In practice, neural networks far outperform their corresponding induced kernels (Arora et al., 2019a). Therefore, understanding the representation learning of neural networks beyond the lazy regime is of fundamental importance.

In this paper, we initiate the study of the representation learning of neural networks beyond this NTK/linear/lazy regime. To this aim, we consider the problem of learning polynomials with low-dimensional latent representation of the form $f^*(x) = g(Ux)$, where U maps from d to r dimensions with $d \gg r$ with g a multivariate polynomial of degree p . This is a natural choice as the failure of the NTK solution is in part due to its inability to learn data-dependent feature representations that adapt to the intrinsic low latent dimensionality of the ground truth function. Existing analysis based on the NTK regime provably require $n \asymp d^p$ samples (Ghorbani et al., 2019a) to learn *any* degree p polynomial, even if they only depend on a few relevant directions. In contrast we show that gradient descent from random initialization only requires $n \asymp d^2$ samples, breaking the sample complexity barrier dictated by NTK proof techniques. More specifically, our contributions are as follows:

1. **Feature Learning:** When the target function $f^* = g(Ux)$ only depends on the projection of x onto a hidden subspace $\text{span}(U)$, we show that gradient descent learns features that span $\text{span}(U)$. Leveraging these features, gradient descent can reach vanishing training loss with a very small network which guarantees good generalization performance. See Section 5.1.
2. **Improved Sample Complexity:** Using classical generalization theory, we demonstrate that when $f^* : \mathbb{R}^d \rightarrow \mathbb{R}$ is a polynomial of degree p which depends on r relevant dimensions (Assumption 1), gradient descent on a two layer neural network learns f^* with only $n \asymp d^2r + dr^p$ samples. This contrasts with the lower bound for random features/NTK methods which require d^p samples to learn *any* degree p polynomial. See Theorem 1.
3. **Transfer learning:** We show that when the target task ground truth is $f_{\text{target}}^*(x) = \tilde{g}(Ux)$, then by simply retraining the network head, gradient descent learns f_{target}^* with only $N \asymp r^p$ target samples and width $m \asymp r^p$, which is independent of the ambient dimension d . In contrast, learning from scratch would require $N \asymp d^{\Omega(p)}$ target samples.
4. **Lower Bound:** Finally, we show a lower bound that demonstrates our non-degeneracy assumption (Assumption 2) is strictly necessary. Without the non-degeneracy, there is a family of polynomials which depend on single relevant dimensions (i.e. of the form $f^*(x) = g(u \cdot x)$) which cannot be learned with fewer than $n \asymp d^{p/2}$ by any gradient descent based learner.

1. See Section 4 for a more in depth discussion of this literature and other related work.

2. Setup

2.1. Input Distribution and Target Function

In this paper we focus on learning a target function $f^*(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ over the input distribution $\mathcal{D} := N(0, I_d)$. We assume that f^* is a degree p polynomial, normalized so that $\mathbb{E}_{x \sim \mathcal{D}}[f^*(x)^2] = 1$. We will attempt to learn f^* given n i.i.d. datapoints $\{x_i, y_i\}_{i \in [n]}$ with

$$x_i \sim \mathcal{D}, \quad y_i = f^*(x_i) + \epsilon_i \quad \text{and} \quad \epsilon_i \sim \{-\varsigma, \varsigma\}$$

where ς^2 controls the strength of the label noise.

In order to make the problem of learning f^* tractable, additional assumptions are necessary. The set of degree p polynomials in d dimensions span a linear subspace of $L^2(\mathcal{D})$ of dimension $\Theta(d^p)$. Learning arbitrary degree p polynomials therefore requires $n \gtrsim d^p$ samples. We follow [Chen and Meka \(2020\)](#); [Chen et al. \(2020b\)](#) in assuming that the ground truth f^* has a special low dimensional latent structure. Specifically, we assume that f^* only depends on a small number of relevant dimensions and that the expected Hessian is non degenerate. We show in [Theorem 2](#) that this non degeneracy assumption is strictly necessary to avoid sample complexity $d^{\Omega(p)}$.

Assumption 1 *There exists a function $g : \mathbb{R}^r \rightarrow \mathbb{R}$ and linearly independent vectors u_1, \dots, u_r such that for all $x \in \mathbb{R}^d$,*

$$f^*(x) = g(\langle x, u_1 \rangle, \dots, \langle x, u_r \rangle).$$

We will call $S^* := \text{span}(u_1, \dots, u_r)$ the principal subspace of f^* . We will also denote by $\Pi^* := \Pi_{S^*}$ the orthogonal projection onto S^* .

Assumption 2 *$H := \mathbb{E}_{x \sim \mathcal{D}}[\nabla^2 f^*(x)]$ has rank r .*

We will also denote the normalized condition number of H by $\kappa := \frac{\|H^\dagger\|}{\sqrt{r}}$.

2.2. The Network and Loss

Let $\sigma(x) = \text{ReLU}(x) = \max(0, x)$, let $a \in \mathbb{R}^m$, $W \in \mathbb{R}^{m \times d}$, $b \in \mathbb{R}^m$, and let $\theta = (a, W, b)$. We define the neural network f_θ by

$$f_\theta(x) = a^T \sigma(Wx + b) = \sum_{j=1}^m a_j \sigma(w_j \cdot x + b_j),$$

where m denotes the width of the network. We use a symmetric initialization, so that $f_{\theta_0}(x) = 0$ ([Chizat and Bach, 2018a](#)). Explicitly, we will assume that m is an even number and that

$$a_j = -a_{m-j}, \quad w_j = w_{m-j} \quad \text{and} \quad b_j = b_{m-j} \quad \forall j \in [m/2].$$

We will use the following initialization:

$$a_j \sim \{-1, 1\}, \quad w_j \sim N\left(0, \frac{1}{d} I_d\right) \quad \text{and} \quad b_j = 0.$$

We note that while we focus on such symmetric initialization for clarity of exposition, our results also hold with small random initialization that is not necessarily symmetric. This holds by simple modifications in the proof accounting for the small nonzero output of the network at initialization. We will also denote the empirical and population losses by $\mathcal{L}(\theta)$ and $\mathcal{L}_{\mathcal{D}}(\theta)$ respectively:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2 \quad \text{and} \quad \mathcal{L}_{\mathcal{D}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}} [(f(x) - f^*(x))^2].$$

2.3. Notation

We use $\lesssim, O(\cdot), \Omega(\cdot)$ to denote quantities that are related by absolute constants and we treat $p, \varsigma = O(1)$. We use $\tilde{O}, \tilde{\Omega}$ to hide additional dependencies on $\text{polylog}(mnd)$. We denote the $L^1(\mathcal{D})$ and $L^2(\mathcal{D})$ losses of a function f by $\mathbb{E}_{x,y} |f(x) - y|$ and $\mathbb{E}_{x,y} (f(x) - y)^2$ respectively where $x \sim \mathcal{D}$, $y = f^*(x) + \epsilon$, and $\epsilon \sim \{\pm \varsigma\}$.

3. Main Results

Before we formally state our main result let us specify the exact form of gradient-based training we use in our theory.

Input: Learning rates η_t , weight decay λ_t , number of steps T

preprocess data

```

|    $\alpha \leftarrow \frac{1}{n} \sum_{i=1}^n y_i$ ,  $\beta \leftarrow \frac{1}{n} \sum_{i=1}^n y_i x_i$ 
|    $y_i \leftarrow y_i - \alpha - \beta \cdot x_i$  for  $i = 1, \dots, n$ 
end
```

$W^{(1)} \leftarrow W^{(0)} - \eta_1 [\nabla_W \mathcal{L}(\theta) + \lambda_1 W]$

re-initialize $b_j \sim N(0, 1)$

for $t = 2$ **to** T **do**

```

|    $a^{(t)} \leftarrow a^{(t-1)} - \eta_t [\nabla_a \mathcal{L}(\theta^{(t-1)}) + \lambda_t a^{(t-1)}]$ 
```

end

return Prediction function $x \rightarrow \alpha + \beta \cdot x + a^T \sigma(Wx + b)$

Algorithm 1: Gradient-based training

With this algorithm in place, we are now ready to state our main result.

Theorem 1 Consider the data model, network and loss per Section 2 and train the network via Algorithm 1 with parameters $\eta_1 = \tilde{O}(\sqrt{d})$, $\lambda_1 = \eta_1^{-1}$, and $\eta_t = \eta$, $\lambda_t = \lambda$ for $t \geq 2$. Furthermore, assume $n \geq \tilde{\Omega}(d^2 \kappa^2 r)$ and $d \geq \tilde{\Omega}(\kappa r^{3/2})$. Then, there exists λ such that if η is sufficiently small, $T = \tilde{\Theta}(\eta^{-1} \lambda^{-1})$ and $\theta^{(T)}$ denotes the final iterate of Algorithm 1, we have that the excess population loss in $L^1(\mathcal{D})$ is bounded with probability at least 0.99 by

$$\mathbb{E}_{x,y} |f_{\theta^{(T)}}(x) - y| - \varsigma \leq \tilde{O} \left(\sqrt{\frac{dr^p \kappa^{2p}}{n}} + \sqrt{\frac{r^p \kappa^{2p}}{m}} + \frac{1}{n^{1/4}} \right).$$

It is useful to note that the use of λ in the algorithm corresponds to the common practice of weight decay and its value is chosen in such a way that $\|a^{(T)}\| \leq B_a$, i.e. to solve a constrained

minimization problem (see Section 5.1). In practice, one simply tunes the hyperparameter λ in order to achieve the desired tradeoff between training and test loss.

An intriguing aspect of the above result is that despite the fact that f^* may be of arbitrarily high degree, learning f^* requires only $n \gtrsim dr^p + d^2r$ samples and only requires a very small network with $m \gtrsim r^p$. We note that our dependence on the latent dimension r is near optimal as the minimax sample complexity even when the principal subspace S^* is known is $\Theta(r^p)$.

We show in Theorem 3 that by resampling the data after the first step, the sample complexity can be further reduced to $d^2r + r^p$, dropping a factor of d from the second term. The extra factor of d results from the dependence between the data used in the first and second stages and we believe that a more careful analysis could remove this additional factor.

We contrast Theorem 1 with the following lower bound for learning a function class which satisfies Assumption 1 with $r = 1$ but does not satisfy Assumption 2.

Theorem 2 *For any $p \geq 0$, there exists a function class \mathcal{F}_p of polynomials of degree p , each of which depends on a single relevant dimension, such that any correlational statistical query learner using q queries requires a tolerance of at most*

$$\tau^2 \leq \frac{\log^{p/2} (qd)}{d^{p/2}}$$

in order to output a function $f \in \mathcal{F}_p$ with $L^2(\mathcal{D})$ loss at most 1.

Using the heuristic $\tau \approx \frac{1}{\sqrt{n}}$, which represents the expected scale of the concentration error, we get the immediate corollary that violating Assumption 2 allows us to construct a function class which *any neural network* with polynomially many parameters trained for polynomially many steps of gradient descent cannot learn without at least $n \gtrsim d^{p/2}$ samples. We emphasize that this is only a heuristic argument as concentration errors are random rather than adversarial.

On the other hand, Theorem 1 shows that incorporating Assumption 2 allows gradient descent to efficiently learn polynomials of arbitrarily high degree with only d^2 samples.

The difference in sample complexity between Theorem 1 and Theorem 2 is that in Theorem 1, our non-degeneracy assumption (Assumption 2) allows the network f_θ to extract useful features that aid robust learning and allowed learning high degree polynomials with $n \gtrsim d^2$ samples. Theorem 2 shows that violating this assumption allows us to construct a function class which cannot be learned without $d^{\Omega(p)}$ samples, demonstrating the necessity of Assumption 2.

The fact that the network f_θ extracts useful features not only allows it to learn f^* efficiently, but also allows for efficient *transfer learning*. In particular, Theorem 3 shows that we can efficiently learn any target polynomial $g^*(x)$ that depends on the same relevant dimensions as f^* with sample complexity *independent of d* by simply truncating and retraining the head of the network:

Theorem 3 *Let $g^*(x)$ be a degree p polynomial with $\mathbb{E}_{\mathcal{D}}[g^*(x)^2] = 1$ and $g(x) = g(\Pi^*x)$ for all $x \in \mathbb{R}^d$. Let $D_N = \{(x_i, y_i)\}_{i \in [N]}$ be a second dataset with $y_i = g(x_i) + \epsilon_i$. We retrain the last layer of the network f_θ in Theorem 1 with gradient descent with learning rate η and weight decay λ , i.e. we will use the function class:*

$$g_a(x) = a^T(W^{(1)}x + b)$$

where $W^{(1)}$ is the second iterate of Algorithm 1. Assume that $d = \tilde{\Omega}(\kappa r^{3/2})$. Then there exists λ such that if the network is pretrained on $n \geq \tilde{\Omega}(d^2 \kappa^2 r)$ datapoints from f^* and η is sufficiently small, the excess population loss in $L^1(\mathcal{D})$ after $T = \tilde{\Theta}(\eta^{-1} \lambda^{-1})$ steps is bounded with probability at least 0.99 by

$$\mathbb{E}_{x,y} |g_{a^{(T)}}(x) - y| - \varsigma \leq \tilde{O} \left(\sqrt{\frac{r^p \kappa^{2p}}{\min(m, N)}} + \frac{1}{N^{1/4}} \right).$$

Learning $g^*(x)$ therefore only requires $N, m \gtrsim r^p$, which is independent of the ambient dimension d . We note that this is minimax optimal for learning arbitrary degree p polynomials even when the hidden subspace S^* is known. Theorem 3 also shows that $n \geq d^2 r$ pre-training samples are necessary for gradient descent to learn the subspace S^* from the pre-training data.

4. Related work

A growing body of recent work show the connection between gradient descent on the full network and the Neural Tangent Kernel (NTK) [Jacot et al. \(2018\)](#); [Oymak and Soltanolkotabi \(2019, 2020\)](#); [Du et al. \(2019b\)](#); [Arora et al. \(2019b\)](#); [Du et al. \(2018a\)](#); [Lee et al. \(2019\)](#). Using this technique one can prove concrete results about neural network training ([Li and Liang, 2018](#); [Du et al., 2018a, 2019b](#); [Allen-Zhu et al., 2018](#); [Zou et al., 2018](#)) and generalization ([Arora et al., 2019b](#); [Oymak et al., 2019](#); [Allen-Zhu et al., 2019](#); [Cao and Gu, 2019](#); [Oymak et al., 2021](#)) in the kernel regime. The key idea is that for a large enough initialization, it suffices to consider a linearization of the neural network around the origin. This allows connecting the analysis of neural networks with the well-studied theory of kernel methods. This is also sometimes referred to as *lazy training*, as with such an initialization the parameters of the neural networks stay close to the parameters at initialization and these results can only show that neural networks are as powerful as shallow learners such as kernels. There is however growing evidence that this NTK-style analysis might not be sufficient to completely explain the success of neural networks in practice. The papers [Chizat et al. \(2019\)](#); [Woodworth et al. \(2019\)](#) provides empirical evidence that by choosing a smaller initialization the test error of the neural network decreases. A similar performance gap between the performance of the NTK and neural networks has been observed in [Ghorbani et al. \(2020\)](#). This NTK-style analysis however does not yield satisfactory results in the setting studied in this paper. In particular for learning the polynomials of the form we study in this paper, [Ghorbani et al. \(2019b\)](#) demonstrates that one needs at least d^p samples in the kernel regime. In contrast, our results only require on the order of d^2 samples.

Leveraging the fact that linearized models are not feature learners, [Ghorbani et al. \(2019b\)](#) and [Wei et al. \(2019\)](#) showed precise upper and lower bounds on the sample complexity of NTK methods. They showed that because NTK is unable to learn new features, learning any polynomial in dimension d of degree p requires $n = \Theta(d^p)$ samples, which *gives no improvement* over polynomial kernels. On the empirical front, the NTK linearization analysis is also lacking. [Arora et al. \(2019a\)](#) demonstrated that the kernel predictor loses more than 20% in test accuracy relative to a deep network trained with SGD and state-of-art regularization on CIFAR-10. Our work is motivated by the contrast between these negative theoretical results for linearized NTK models and the spectacular empirical performance of deep learning.

The gap between such shallow learners and the full neural network has been established in theory ([Wei et al., 2019](#); [Allen-Zhu and Li, 2020, 2019](#); [Yehudai and Shamir, 2019b](#); [Ghorbani et al.,](#)

2019a; Woodworth et al., 2020; Dyer and Gur-Ari, 2019; Du and Lee, 2018) and observed in practice (Arora et al., 2019a; Lee et al., 2019; Chizat and Bach, 2018a). There is an emerging literature on learning beyond the lazy/NTK regime in the small initialization setting. The papers Li et al. (2018); Stöger and Soltanolkotabi (2021) shows that for the problem of low-rank reconstruction in a non-lazy regime with small random initialization gradient descent finds globally optimal solutions with good generalization capability. This is carried out by utilizing a spectral bias phenomena exhibited by the early stages of gradient descent from small random initialization that puts the iterates on the trajectory towards generalizable models. For the problem of tensor decomposition it has also been shown that gradient descent with small initialization is able to leverage low-rank structure (Wang et al., 2020). In Li et al. (2020), it has been shown that neural networks with orthogonal weights can be learned via SGD and outperform any kernel method. One crucial element in their analysis is that the early stage of the training is connected with learning the first and second moment of the data. Higher-order approximations of the training dynamics (Bai and Lee, 2020; Bai et al., 2020) and the Neural Tangent Hierarchy (Huang and Yau, 2019) have also been recently proposed towards closing this gap. None of the above papers, however, focus on learning polynomial representations efficiently via neural networks as carried out in this paper.

Another line of work focuses on learning single activations such as the ReLU function. In this context (Yehudai and Shamir, 2019a) shows that it is hard to learn a single ReLU activation via stochastic gradient descent with random features where as learning such activations is possible in a non-NTK regime (Soltanolkotabi, 2017; Goel et al., 2017, 2019) again highlighting this important gap. In related work where the label also only depends on a single relevant direction (Daniely and Malach, 2020), the authors show that in the context of learning the parity function, gradient descent is able to efficiently learn the planted set. However, this is a result of the unbalanced data distribution which skews the gradient towards the planted set. In contrast, we consider isotropic Gaussian data so that no information can be extracted from the data distribution itself and features must be extracted from higher order correlations between the data and the labels. Chen and Meka (2020) also studied the problem of learning polynomials of few relevant dimensions. They provide an algorithm that learns polynomials of degree p in d dimensions that depends on r hidden dimensions with $n \gtrsim C(r, p)d$ samples where $C(r, p)$ is an unspecified function of r, p which is likely exponential in r . However, their algorithm is not a variant of gradient descent, and requires a clever spectral initialization. On the other hand, this work focuses on the ability of gradient descent to *automatically* extract hidden features and learn representations from the data.

There is also a line of work Mei et al. (2018); Chizat and Bach (2018b); Mei et al. (2019); Javanmard et al. (2020); Sirignano and Spiliopoulos (2020); Wei et al. (2019), which is concerned with the mean-field analysis of neural networks. The insight is that for sufficiently large width the training dynamics of the neural network can be coupled with the evolution of a probability distribution described by a PDE. These papers use a smaller initialization than in the NTK-regime and, hence, the parameters can move away from the initialization. However, these results do not provide explicit convergence rates and require an unrealistically large width of the neural network. To the extent of our knowledge such an analysis technique has not been used to show efficient learning of polynomial representations using neural networks as carried out in this paper.

A concurrent line of work studied the feature learning ability of gradient descent in the mean field regime with data sampled from the boolean cube (Abbe et al., 2022). The authors identified a necessary and sufficient condition for learning with sample complexity linear in d , dubbed the *merged staircase property*, in the special case when the hidden weights of the two layer neural

network are initialized at 0. However, the zero initialization hinders the feature learning ability of the network. For example, the boolean function XOR violates the merged staircase property, however noisy XOR is known to be learnable by two layer neural networks with sample complexity linear in d (Bai and Lee, 2020; Chen et al., 2020a). In this work we study the impact that the nonzero initialization of the hidden weights has on the feature learning ability of neural networks.

5. Proof Sketches

5.1. Proof of Theorems 1 and 3

The proofs of Theorems 1 and 3 are essentially identical so we will focus on Theorem 1. We begin by noting that the symmetric initialization implies that $f_\theta(x) = 0$ for all $x \in \mathbb{R}^d$. This implies that the population gradient of each feature w_j can be written as

$$\nabla_{w_j} \mathcal{L}_{\mathcal{D}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}} [2(f_\theta(x) - f^*(x)) \nabla_{w_j} f_\theta(x)] = -2 \mathbb{E}_{x \sim \mathcal{D}} [f^*(x) \nabla_{w_j} f_\theta(x)].$$

Using the chain rule, we can further expand this as

$$-2 \mathbb{E}_{x \sim \mathcal{D}} [f^*(x) \nabla_{w_j} f_\theta(x)] = -2a_j \mathbb{E}_x [f^*(x) x \mathbf{1}_{w_j \cdot x \geq 0}].$$

The main computation that drives Theorems 1 and 3 is that for any unit vector $w \in \mathbb{R}^d$, the expression $\mathbb{E}_x [f^*(x) x \mathbf{1}_{w \cdot x \geq 0}]$ has a natural series expansion in powers of w , which can be computed explicitly in terms of the Hermite expansions of f^* and σ' . Explicitly, if $C_k = \mathbb{E}_x [\nabla^k f^*(x)]$ is a symmetric k tensor denoting the expected k th derivative of f^* and c_k are the Hermite coefficients of $\sigma'(x) = \mathbf{1}_{x \geq 0}$,

$$\begin{aligned} \mathbb{E}_x [f^*(x) x \mathbf{1}_{w \cdot x \geq 0}] &= \sum_{k=0}^{p-1} \frac{1}{k!} \left[c_{k+1} C_{k+1}(w^{\otimes k}) + c_{k+2} w C_k(w^{\otimes k}) \right] \\ &= \underbrace{\frac{Hw}{\sqrt{2\pi}}}_{O(d^{-1/2})} + \underbrace{\frac{1}{2} [c_3 C_3(w, w) + c_4 w C_2(w, w)]}_{O(d^{-1})} + \underbrace{\frac{1}{6} [\dots]}_{O(d^{-3/2})} + \dots \end{aligned} \quad (1)$$

where we note that $C_2 = \mathbb{E}_x [\nabla^2 f^*(x)] = H$. We emphasize that because w is a unit vector, its inner product with any fixed unit vector is of order $d^{-1/2}$ so temporarily ignoring factors of r , $\|C_{k+1}(w^{\otimes k})\|, |C_k(w^{\otimes k})| = O(d^{-\frac{k}{2}})$. Therefore Equation (1) is an asymptotic series in $d^{-1/2}$. As k increases, each term in Equation (1) reveals more information about f^* . However, this information is also better hidden. A standard concentration argument shows that extracting information from the C_k term in this series requires $n \geq d^k$ samples. This paper focuses on the first term in this expansion, $\frac{Hw}{\sqrt{2\pi}}$, which requires $n \geq d^2$ samples to isolate. We directly truncate this series expansion:

Lemma 4 *With high probability over the random initialization,*

$$\nabla_{w_j} \mathcal{L}_{\mathcal{D}}(\theta) = -2a_j \frac{Hw_j}{\sqrt{2\pi}} + \tilde{O}\left(\frac{\sqrt{r}}{d}\right).$$

Note that the remainder term, of order d^{-1} , contains all higher order terms in the series expansion.

Recall that $H = \mathbb{E}_{x \sim \mathcal{D}}[\nabla^2 f^*(x)]$ is the average Hessian of f^* with respect to \mathcal{D} . Because f^* depends only on the subspace S^* , this implies that up to higher order terms, the population gradient at initialization already points each feature vector w_j towards the principal subspace S^* . In addition, Assumption 2 guarantees that the gradients at initialization span the principal subspace S^* .

However, it is also important to note that the population gradient is bounded by $\|\nabla_{w_j} \mathcal{L}_{\mathcal{D}}(\theta)\| = O(d^{-1/2})$ and we only have access to the empirical gradient $\nabla_{w_j} \mathcal{L}(\theta)$. As mentioned above, extracting the necessary subspace information from $\nabla_{w_j} \mathcal{L}_{\mathcal{D}}(\theta)$ to learn f^* therefore requires $n \gtrsim d^2$ samples, which is the dominant term in our final sample complexity result.

Once we show that the gradient at initialization contains all the relevant features, we note that after the first step of gradient descent,

$$W^{(1)} = W^{(0)} - \eta_1 [\nabla_W \mathcal{L}(\theta^{(0)}) + \eta_1^{-1} W] = -\eta_1 \nabla_W \mathcal{L}(\theta^{(0)}).$$

After the first step, the model therefore resembles a random feature model with random features $\{Hw\}_{w \in S^{d-1}} \subset S^*$. Previous results have shown that in these linearized regimes, e.g. random feature models/NTK, learning degree p polynomials requires $n \gtrsim d^p$ samples and width $m \gtrsim d^p$. As our “random features” are now constrained to the hidden subspace S^* , which has dimension r , we should expect that our sample complexity improves to $n \gtrsim r^p$.

The remainder of Algorithm 1 runs ridge regression on the network head a with fixed features $x \rightarrow \sigma(W^{(1)}x + b)$. We can directly analyze the generalization of this algorithm using standard techniques from Rademacher complexity. In particular, a high level sketch of the remainder of the proof goes as follows:

1. (Section A.2): We use the features from Lemma 4 to construct a vector $a^* \in \mathbb{R}^m$ such that

$$\mathcal{L}(a^*, W^{(1)}, b) \ll 1 \quad \text{and} \quad \|a^*\| = \tilde{O}\left(\frac{r^p}{\sqrt{m}}\right).$$

2. (Section A.3): We show the equivalence between ridge regression and norm constrained linear regression implies the existence of $\lambda > 0$ such that the T th iterate $a^{(T)}$ satisfies

$$\mathcal{L}(a^{(T)}, W^{(1)}, b) \ll 1 \quad \text{and} \quad \|a^{(T)}\| \leq \|a^*\|.$$

3. (Section A.3): A standard Rademacher generalization bound for two layer neural networks bounds the population risk $\mathbb{E}_{x,y} |f_{\theta^{(T)}}(x) - y|$ by the empirical risk $\frac{1}{n} \sum_{i=1}^n |f_{\theta^{(T)}}(x_i) - y_i|$ and $\|a^{(T)}\|$ which are small from step 2.

5.2. Proof of Theorem 2

Statistical query learners are a family of learners that can query values $q(x, y)$ and receive outputs \hat{q} with $|\hat{q} - \mathbb{E}_{x,y}[q(x, y)]| \leq \tau$ where τ denotes the query tolerance (Goel et al., 2020; Diakonikolas et al., 2020). An important class of statistical query learners is that of correlational/inner product statistical queries (CSQ) of the form $q(x, y) = yh(x)$. This includes a wide class of algorithms including gradient descent with square loss. For example from Section 5.1, for a two layer neural network we have

$$\nabla_{w_j} \mathcal{L}_{\mathcal{D}}(\theta) = \mathbb{E}_{x,y}[yh(x)] \quad \text{where} \quad h(x) = -2a_j x \mathbf{1}_{w_j \cdot x + b_j \geq 0}.$$

In order to prove Theorem 2, we must construct a function class \mathcal{F}_p such that inner product queries of the form $\mathbb{E}_{x,y}[yh(x)]$ provide little to no information about the target function. The standard approach is to construct a function class with small pairwise correlations, i.e. for $f \neq g \in \mathcal{F}_p$, $|\mathbb{E}_x[f(x)g(x)]| \leq \epsilon$ (Goel et al., 2020; Diakonikolas et al., 2020). The number of functions in the function class \mathcal{F}_p and the size of the pairwise correlations ϵ directly imply a correlational statistical query lower bound:

Lemma 5 (Modified from Theorem 2 in Szörényi (2009)) *Let \mathcal{F} be a class of functions and \mathcal{D} be a data distribution such that*

$$\mathbb{E}_{x \sim \mathcal{D}}[f(x)^2] = 1 \quad \text{and} \quad |\mathbb{E}_{x \sim \mathcal{D}}[f(x)g(x)]| \leq \epsilon \quad \forall f \neq g \in \mathcal{F}.$$

Then any correlational statistical query learner requires at least $\frac{|\mathcal{F}|(\tau^2 - \epsilon)}{2}$ queries of tolerance τ to output a function in \mathcal{F} with $L^2(\mathcal{D})$ loss at most $2 - 2\epsilon$.

To construct \mathcal{F}_p , we begin by showing that there are a large number of approximately orthogonal unit vectors in S^{d-1} :

Lemma 6 *There exists an absolute constant c such that for any $\epsilon > 0$, there exists a set S of $\frac{1}{2}e^{ce^2d}$ unit vectors such that for any $v, w \in S$ such that $v \neq w$, we have $|v \cdot w| \leq \epsilon$.*

The proof bounds the probability that randomly sampled unit vectors have a large inner product and existence then follows from the probabilistic method. Therefore for any m , we can find m unit vectors in \mathbb{R}^d such that their pairwise inner products are all bounded by $d^{-1/2}\sqrt{\log m}$. We combine this with the fact that if $f_u(x) = \frac{He_k(u \cdot x)}{\sqrt{k!}}$ where He_k denotes the k th Hermite polynomial,

$$\mathbb{E}_{x \sim \mathcal{D}}[f_u(x)f_v(x)] = (u \cdot v)^k.$$

Therefore $|u \cdot v| \leq d^{-1/2}\sqrt{\log m}$ implies $|\mathbb{E}_{x \sim \mathcal{D}}[f_u(x)f_v(x)]| \leq d^{-k/2}(\log m)^{k/2}$. Theorem 2 then directly follows from Lemma 5 (see Appendix D for a more detailed proof).

6. Experiments

6.1. Sample Complexity

In this section we present a toy example that clearly demonstrates the gap between kernel methods and gradient descent on two layer networks. For $u \in S^{d-1}$, consider the target function

$$f_u^*(x) = g(u \cdot x) \quad \text{where} \quad g(x) = \frac{He_2(x)}{2} + \frac{He_p(x)}{\sqrt{2p!}}, \quad (2)$$

which satisfies $\mathbb{E}_{x \sim \mathcal{D}}[f_u^*(x)^2] = 1$. Note that f^* only depends on the projection of x onto a single relevant direction, u . We show in Section 5.1 that gradient descent is capable of isolating the subspace spanned by u and then fitting a one dimensional random feature model to g , and that this entire process requires $n \asymp d^2$ samples to generalize.

On the other hand, existing works Ghorbani et al. (2019b, 2020) have shown that $n \asymp d^p$ samples are strictly necessary in order to learn f^* in the NTK or random features regime. The theory predicts that with $n < d^2$ samples, kernel regression will return the 0 predictor and with $d^2 < n < d^p$ samples, kernel regression will return $\frac{1}{2}He_2(u \cdot x)$, incurring a $L^2(\mathcal{D})$ loss of $\frac{1}{2}$.

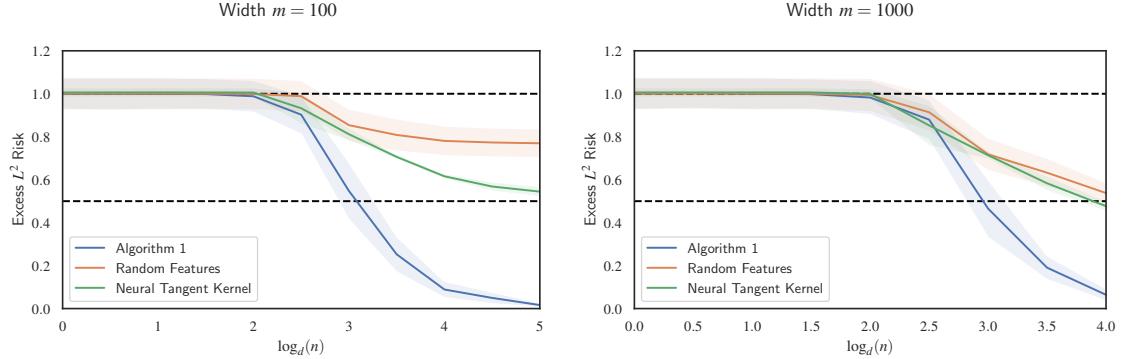


Figure 1: **Sample Complexity:** The x axis plots $\log_d(n)$ and the y axis plots the excess risk in $L^2(\mathcal{D})$ for each of three methods: Algorithm 1, random features, and the neural tangent kernel. The top dashed horizontal line is at $L^2 = 1$, which corresponds to outputting the zero predictor. The middle horizontal line is at $L^2 = \frac{1}{2}$ and corresponds to learning the optimal quadratic predictor $\frac{1}{2}He_2(x_1)$. Due to its improved sample efficiency, Algorithm 1 easily achieves near zero excess risk despite the relatively high degree of f^* , while random features and the neural tangent kernel are only able to learn the optimal quadratic predictor. See Section 6 for additional experimental details.

We empirically verify these predictions. We take $d = 10$ and $p = 4$ and consider the function $f_{e_1}^*(x) = \frac{He_2(x_1)}{2} + \frac{He_4(x_1)}{4\sqrt{3}}$. We use label noise $\sigma^2 = 1$ and attempt to learn f^* using Algorithm 1, a random feature model, and a linearized NTK model. All experiments are conducted on a two layer neural network with widths $m = 100$ and $m = 1000$. For each value of n , the weight decay parameter λ is tuned on a holdout set of size 10^5 and test accuracies are reported over a separate test set of size 10^5 . Errors bars reflect the mean and standard deviation over 10 random seeds.

We note that while Algorithm 1 easily converged to vanishing excess risk, even at width $m = 100$, both the random features model and the neural tangent kernel model only managed to fit the quadratic term $\frac{1}{2}He_2(u \cdot x)$, as predicted by the theory in Ghorbani et al. (2019b, 2020).

The key to learning a function of the form f_u^* is to use the fact that the $\frac{1}{2}He_2(u \cdot x)$ component of f_u^* gives enough information to identify u . Afterwards, any random feature or kernel method can efficiently fit any sufficiently smooth univariate function $g : \mathbb{R} \rightarrow \mathbb{R}$ ontop of $u \cdot x$. Our analysis in Section 5.1 shows that this is exactly the way that Algorithm 1 learns f_u^* and this is reflected by the steep and sudden drop from trivial risk ($L^2 = 1$) to vanishing excess risk without plateauing at $L^2 = 0.5$ in Figure 1.

6.2. Transfer Learning

The proof of Theorem 1 involves showing that Algorithm 1 learns features corresponding to S^* (see Section 5.1) and the proof of Theorem 3 shows that this implies efficient transfer learning. We again verify this empirically. We consider the function:

$$f_{\text{target}}^*(x) = g_{\text{target}}(u \cdot x) \quad \text{where} \quad g_{\text{target}} = \frac{He_p(x)}{\sqrt{p!}}.$$

Note that this was exactly the hard example in Theorem 2 that was unlearnable without $n \gtrsim d^{\frac{p}{2}}$ samples by a correlational statistical query learner (and in particular, gradient-based learners).

We pretrain with n samples on the $f^*(x)$ from Section 6.1, then train the output layer using N samples from f_{target}^* . As in Section 6.1, we use a label noise strength of $\sigma^2 = 1$. We pick $p = 3$ so that random feature methods or the neural tangent kernel will require at least $n \gtrsim d^3$ samples to learn f^* .

We note that in Figure 2, when $n = d^0, d^1$, fine tuning on N target samples gives trivial risk until $N \gtrsim d^3$, which is to be expected of a kernel method with no prior information. However, for $n \geq d^2$ pretraining samples, we can fine tune on just $N = O(1)$ target samples to reach nontrivial loss and the loss decays rapidly as a function of N . This experiment therefore fully supports the conclusion of Theorem 3.

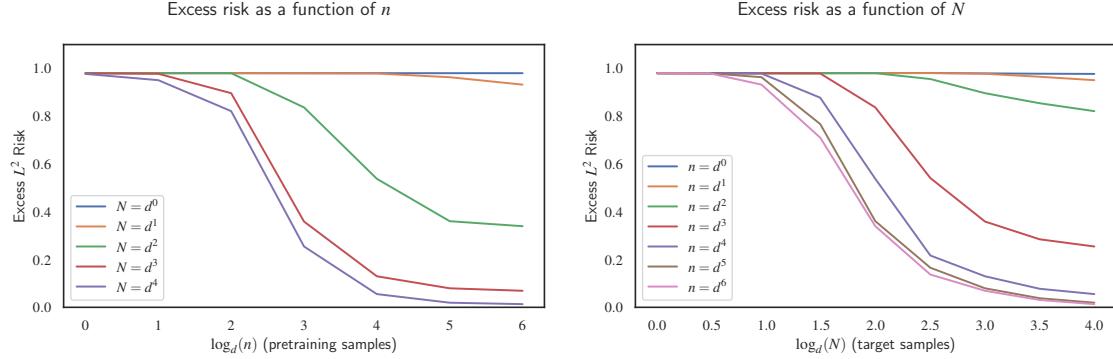


Figure 2: Transfer Learning: The x axes plot $\log_d(n)$ and $\log_d(N)$ respectively. We note that with little pretraining ($\log_d(n) = 0, 1$), Algorithm 1 is unable to extract a robust representation that enables transfer. For $n \geq d^2$, we observe that finetuning the representation from Algorithm 1 gives nontrivial loss even for $N = O(1)$, as predicted by Theorem 3. See Section 6 for additional experimental details.

7. Discussion and Future Work

In this work we provide a clear separation between gradient-based training and kernel methods. We show that there is a large family of degree p polynomials which are efficiently learnable by gradient descent with $n \asymp d^2$ samples, in contrast to the lower bound of d^p for random feature/NTK analysis. The main idea driving both our sample complexity result (Theorem 1) and our transfer learning result (Theorem 3) is that gradient descent learns useful representations of the data.

One promising direction for future work is tightening the dimension dependence of our upper bound. In particular, our $n \asymp d^2$ sample complexity is driven by the difficult in learning from a degree 2 Hermite polynomial. However, our lower bound for such functions (Theorem 2) only rules out learning with $n \leq d$ samples. In this situation the lower bound is tight as Chen et al. (2020b) show that sparse degree 2 polynomials can be efficiently learned with $n \asymp d$ samples.

Another promising direction from future work is generalizing our result to the situation in which the hidden layer and the output layer are trained together. This introduces dependencies between

the hidden and output layers which are difficult to control. However, such analysis may lead to a better understanding of learning order and inductive bias in deep learning.

Acknowledgments

AD acknowledges support from a NSF Graduate Research Fellowship. JDL and AD acknowledge support of the ARO under MURI Award W911NF-11-1-0304, the Sloan Research Fellowship, NSF CCF 2002272, NSF IIS 2107304, ONR Young Investigator Award, and NSF-CAREER under award #2144994. MS is supported by the Packard Fellowship in Science and Engineering, a Sloan Fellowship in Mathematics, an NSF-CAREER under award #1846369, DARPA Learning with Less Labels (LwLL) and FastNICS programs, and NSF-CIF awards #1813877 and #2008443.

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. *arXiv preprint arXiv:2202.08658*, 2022.
- Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels? *arXiv preprint arXiv:1905.10337*, 2019.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep learning. *arXiv preprint arXiv:2001.04413*, 2020.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6155–6166, 2019.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019a.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019b.
- Yu Bai and Jason D Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. *International Conference on Learning Representations (ICLR)*, 2020.
- Yu Bai, Ben Krause, Huan Wang, Caiming Xiong, and Richard Socher. Taylorized training: Towards better approximation of neural network training at finite width. *arXiv preprint arXiv:2002.04010*, 2020.
- Yuan Cao and Quanquan Gu. Generalization error bounds of gradient descent for learning overparameterized deep relu networks. *arXiv preprint arXiv:1902.01384*, 2019.
- Minshuo Chen, Yu Bai, J. Lee, Tuo Zhao, Huan Wang, Caiming Xiong, and Richard Socher. Towards understanding hierarchical learning: Benefits of neural representations. *ArXiv*, abs/2006.13436, 2020a.
- Minshuo Chen, Yu Bai, Jason D Lee, Tuo Zhao, Huan Wang, Caiming Xiong, and Richard Socher. Towards understanding hierarchical learning: Benefits of neural representations. *Neural Information Processing Systems (NeurIPS)*, 2020b.
- Sitan Chen and Raghu Meka. Learning polynomials of few relevant dimensions. *arXiv preprint arXiv:2004.13748*, 2020.
- Lenaic Chizat and Francis Bach. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 8, 2018a.

- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 3036–3046. Curran Associates, Inc., 2018b. URL <https://proceedings.neurips.cc/paper/2018/file/a1afc58c6ca9540d057299ec3016d726-Paper.pdf>.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32:2937–2947, 2019.
- Amit Daniely and Eran Malach. Learning parities with neural networks. *Advances in Neural Information Processing Systems*, 33:20356–20365, 2020.
- Ilias Diakonikolas, Daniel M Kane, Vasilis Kontonis, and Nikos Zarifis. Algorithms and sq lower bounds for pac learning one-hidden-layer relu networks. In *Conference on Learning Theory*, pages 1514–1539, 2020.
- Simon Du and Jason Lee. On the power of over-parametrization in neural networks with quadratic activation. In *International conference on machine learning*, pages 1329–1338. PMLR, 2018.
- Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*, 2018a.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018b.
- Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. *International Conference on Machine Learning (ICML)*, 2019a.
- Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685, 2019b.
- Ethan Dyer and Guy Gur-Ari. Asymptotics of wide networks from feynman diagrams. *arXiv preprint arXiv:1909.11304*, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Limitations of lazy training of two-layers neural network. In *Advances in Neural Information Processing Systems*, pages 9108–9118, 2019a.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *arXiv preprint arXiv:1904.12191*, 2019b.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *arXiv preprint arXiv:2006.13409*, 2020.
- S. Goel, S. Karmalkar, and A. Klivans. Time/accuracy tradeoffs for learning a relu with respect to gaussian marginals. In *NeurIPS*, 2019.
- Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. In *Conference on Learning Theory*, pages 1004–1042, 2017.

- Surbhi Goel, Aravind Gollakota, Zhihan Jin, Sushrut Karmalkar, and Adam Klivans. Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. *arXiv preprint arXiv:2006.12011*, 2020.
- Friedrich Gotze, Holger Sambale, and Arthur Sinulis. Concentration inequalities for polynomials in alpha-sub-exponential random variables. *arXiv: Probability*, 2019.
- Jiaoyang Huang and Horng-Tzer Yau. Dynamics of deep neural networks and neural tangent hierarchy. *arXiv preprint arXiv:1909.08156*, 2019.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Adel Javanmard, Marco Mondelli, Andrea Montanari, et al. Analysis of a two-layer neural network via displacement convexity. *Ann. Statist.*, 48(6):3619–3642, 2020.
- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in neural information processing systems*, pages 8570–8581, 2019.
- Y. Li, T. Ma, and H. Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pages 2–47, 2018.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.
- Yuanzhi Li, Tengyu Ma, and Hongyang R Zhang. Learning over-parametrized two-layer relu neural networks beyond ntk. *arXiv preprint arXiv:2007.04596*, 2020.
- Shengchao Liu, Dimitris Papailiopoulos, and Dimitris Achlioptas. Bad global minima exist and sgd can reach them. *Advances in Neural Information Processing Systems*, 33:8543–8552, 2020.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1806579115. URL <https://www.pnas.org/content/115/33/E7665>.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR, 2019.
- Samet Oymak and Mahdi Soltanolkotabi. Overparameterized nonlinear learning: Gradient descent takes the shortest path? In *International Conference on Machine Learning*, pages 4951–4960. PMLR, 2019.

Samet Oymak and Mahdi Soltanolkotabi. Towards moderate overparameterization: global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 2020.

Samet Oymak, Zalan Fabian, Mingchen Li, and Mahdi Soltanolkotabi. Generalization guarantees for neural networks via harnessing the low-rank structure of the jacobian. *arXiv preprint arXiv:1906.05392*, 2019.

Samet Oymak, Mingchen Li, and Mahdi Soltanolkotabi. Generalization guarantees for neural architecture search with train-validation split. In *International Conference on Machine Learning*, pages 8291–8301. PMLR, 2021.

Giuseppe Da Prato and Luciano Tubaro. Wick powers in stochastic pdes: an introduction. 2007.

Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: a central limit theorem. *Stochastic Processes Appl.*, 130(3):1820–1852, 2020. ISSN 0304-4149.

Mahdi Soltanolkotabi. Learning relus via gradient descent. In *Advances in neural information processing systems*, pages 2007–2017, 2017.

Mahdi Soltanolkotabi, Adel Javanmard, and Jason D Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 65(2):742–769, 2018.

Dominik Stöger and Mahdi Soltanolkotabi. Small random initialization is akin to spectral learning: Optimization and generalization guarantees for overparameterized low-rank matrix reconstruction. *Advances in Neural Information Processing Systems*, 34, 2021.

Balázs Szörényi. Characterizing statistical query learning: Simplified notions and proofs. In *ALT*, 2009.

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

Xiang Wang, Chenwei Wu, Jason D Lee, Tengyu Ma, and Rong Ge. Beyond lazy training for over-parameterized tensor decomposition. *Advances in Neural Information Processing Systems*, 33, 2020.

Colin Wei, Jason D Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9709–9721, 2019.

Blake Woodworth, Suriya Genesekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Kernel and deep regimes in overparametrized models. In *Conference on Learning Theory (COLT)*, 2019.

Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. *arXiv preprint arXiv:2002.09277*, 2020.

Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *CoRR*, abs/1904.00687, 2019a. URL <http://arxiv.org/abs/1904.00687>.

Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *arXiv preprint arXiv:1904.00687*, 2019b.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*, 2018.

Appendix A. Proofs

We define $\iota = C_\iota \log(nmd)$ for a sufficiently large constant C_ι . Throughout the appendix we will use $e^{-\iota}$ to track failure probabilities of various lemmas and theorems.

Definition 7 (High probability events) *We say that an event A happens with high probability if it happens with probability at least $1 - \text{poly}(n, m, d)e^{-\iota}$ where $\text{poly}(n, m, d)$ does not depend on C_ι .*

Note that high probability events are closed under taking union bounds over sets of size $\text{poly}(n, m, d)$. We will assume throughout that $\iota \leq cd$ for a sufficiently small absolute constant c .

The following lemma bounds $\|x_i\|$ and is a direct corollary of Lemma 27:

Lemma 8 *With high probability, $\|x_i\|^2 \in [\frac{d}{2}, 2d]$ for $i = 1, \dots, n$.*

All remaining proofs will be conditioned on this high probability event.

A.1. Hermite Expansions

A.1.1. HERMITE EXPANSION OF σ

Let $\sigma(x) := \text{ReLU}(x) = \max(0, x)$. Then the Hermite expansion of $\sigma(x)$ is

$$\sigma(x) = \frac{1}{\sqrt{2\pi}} + \frac{x}{2} + \frac{1}{\sqrt{2\pi}} \sum_{k \geq 1} \frac{(-1)^{k-1}}{k! 2^k (2k-1)} H e_{2k}(x).$$

Let c_k denote the Hermite coefficients of σ , i.e. $\sigma(x) = \sum_{k \geq 0} \frac{c_k}{k!} H e_k(x)$. Note that

$$\sigma'(x) = \sum_{k \geq 0} \frac{c_{k+1}}{k!} H e_k(x) = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \sum_{k \geq 0} \frac{(-1)^k}{k! 2^k (2k+1)} H e_{2k+1}(x).$$

A.1.2. HERMITE EXPANSION OF f^*

Let the Hermite expansion of f^* be

$$f^*(x) = \sum_{k=0}^p \frac{\langle C_k, H e_k(x) \rangle}{k!}$$

where $C_k \in (\mathbb{R}^d)^{\otimes k}$ is the symmetric k -tensor defined by

$$C_k := \mathbb{E}_x[\nabla^k f^*(x)].$$

Note that

$$\nabla f^*(x) = \sum_{k=0}^{p-1} \frac{C_{k+1}(H e_k(x))}{k!} \in \mathbb{R}^d.$$

Lemma 9 (Parseval's Identity)

$$1 = \mathbb{E}_x[f^*(x)^2] = \sum_{k=0}^p \frac{\|C_k\|_F^2}{k!}.$$

Note that as an immediate consequence of Lemma 9, $\|C_k\|_F^2 \leq k!$. In addition, Assumption 1 guarantees that $C_k(x^{\otimes k}) = C_k((\Pi^* x)^{\otimes k})$.

A.1.3. CONCENTRATING α, β

Lemma 10 Let $\alpha = \frac{1}{n} \sum_{i=1}^n y_i$ and $\beta = \frac{1}{n} \sum_{i=1}^n y_i x_i$. Then, with high probability,

$$|\alpha - C_0| \lesssim \sqrt{\frac{\nu^p}{n}} \quad \text{and} \quad \|C_1 - \beta\| \lesssim \sqrt{\frac{d\nu^{p+1}}{n}}.$$

Proof Let $F(x_1, \dots, x_n) = \frac{1}{n} \sum_{i=1}^n f^*(x_i) - C_0$. Note that

$$\mathbb{E}_{x_1, \dots, x_n} [F(x)^2] = \frac{1}{n} \text{Var}(f^*(x)) \leq \frac{1}{n}.$$

The bound on $|\alpha - C_0|$ therefore immediately follows from Lemma 30 applied to F . The bound on $\|\beta - C_1\|$ is a special case of Lemma 32 with $\sigma(x) = x$. \blacksquare

A.1.4. HERMITE EXPANDING THE FEATURES

Note that by the scale invariance of $\sigma(x) = \text{ReLU}(x)$, Algorithm 1 does not depend on $\|w_j\|$ for $j = 1, \dots, m$. Therefore we can assume WLOG that $\|w_j\| = 1$ for $j = 1, \dots, m$ and $w_j \sim \text{Unif}(S^{d-1})$. For the remainder of the appendix we will assume that $\|w_j\| = 1$.

Definition 11 We define $\hat{f}^*(x) := f^*(x) - \alpha - \beta \cdot x$.

The functions $g(w)$ and $g_n(w)$ capture the features that can be learned after one step of gradient descent:

Definition 12 For $\|w\| = 1$, we define

$$g(w) := \mathbb{E}_x [\hat{f}^*(x) x \sigma'(w \cdot x)] \quad \text{and} \quad g_n(w) := \frac{1}{n} \sum_{i=1}^n (\hat{f}^*(x_i) + \epsilon_i) x_i \sigma'(w \cdot x_i).$$

We note that $w_j^{(1)} = 2\eta_1 a_j g_n(w_j)$ and $g(w) = \mathbb{E}_x [g_n(w)]$. In fact, Corollary 34 shows that with probability at least $1 - 4ne^{-\nu}$,

$$\sup_w \|g(w) - g_n(w)\| \lesssim \sqrt{\frac{d\nu^{p+1}}{n}}.$$

Lemma 13 With high probability,

$$g(w) = \frac{Hw}{\sqrt{2\pi}} + O\left(\sqrt{\frac{d\nu^{p+1}}{n}} + \sqrt{\frac{r\nu^2}{d^2}}\right).$$

Proof By Stein's lemma and the orthogonality of Hermite polynomials,

$$\begin{aligned} g(x) &= \mathbb{E}_x [\hat{f}^*(x) x \sigma'(w \cdot x)] \\ &= \mathbb{E}_x [\nabla \hat{f}^*(x) \sigma'(w \cdot x) + w f^*(x) \sigma''(w \cdot x)] \\ &= \sum_{k=0}^{p-1} \frac{c_{k+1} \mathbb{E}_x [\nabla_x^{k+1} \hat{f}^*(x)] (w^{\otimes k})}{k!} + w \sum_{k=0}^p \frac{c_{k+2} \mathbb{E}_x [\nabla_x^k \hat{f}^*(x)] (w^{\otimes k})}{k!} \\ &= \left(\frac{C_1 - \beta}{2} + \frac{w(C_0 - \alpha)}{\sqrt{2\pi}} \right) + \frac{C_2 w}{\sqrt{2\pi}} + \sum_{k \geq 2} \frac{c_{2k} C_{2k} (w^{\otimes 2k-1})}{(2k-1)!} + w \sum_{k \geq 1} \frac{c_{2k+2} C_{2k} (w^{\otimes 2k})}{(2k)!}. \end{aligned}$$

Note that these sums are finite as $C_k = 0$ for $k > p$. Next, by Lemma 46 we have the high probability bounds,

$$\begin{aligned}\|C_{k+1}(w^{\otimes k})\|_F &\lesssim \sqrt{\frac{r^{\lfloor \frac{k}{2} \rfloor} \iota^k}{d^k}} \lesssim \sqrt{\frac{r\iota^3}{d^3}} \quad \text{for } k \geq 3 \\ \|wC_k(w^{\otimes k})\|_F &\lesssim \sqrt{\frac{r^{\lfloor \frac{k}{2} \rfloor} \iota^k}{d^k}} \lesssim \sqrt{\frac{r\iota^2}{d^2}} \quad \text{for } k \geq 2.\end{aligned}$$

Applying these bounds term by term and using Lemma 10 to bound $|C_0 - \alpha|$ and $\|C_1 - \beta\|$ gives the desired result. \blacksquare

Corollary 14 *With high probability,*

$$g_n(w) = \frac{Hw}{\sqrt{2\pi}} + O\left(\sqrt{\frac{d\iota^{p+1}}{n}} + \sqrt{\frac{r\iota^2}{d^2}}\right).$$

Corollary 15 *With high probability,*

$$\|g_n(w)\| \lesssim \sqrt{\frac{\iota^2}{d}} + \sqrt{\frac{d\iota^{p+1}}{n}}.$$

Furthermore, it will become necessary to bound terms of the form $g_n(w) \cdot x_i$. Note that $g_n(w)$ and x_i are dependent random variables. The following lemma handles this dependence.

Lemma 16 *Let $w \sim S^{d-1}$ and assume $n \geq d^2\iota^p$. Then with high probability,*

$$\max_{j \in [n]} \|g_n(w) \cdot x_j\| \leq \sqrt{\frac{\iota^3}{d}}.$$

Proof We can decompose

$$|g_n(w) \cdot x_j| \leq |g(w) \cdot x_j| + |[g(w) - g_n(w)] \cdot x_j|.$$

For the first term, note that $g(w)$ and x_j are independent so $g(w) \cdot x_j \sim N(0, \|g(w)\|^2)$ so with high probability,

$$|g(w) \cdot x_i| \leq \|g(w)\| \sqrt{2\iota} \lesssim \sqrt{\frac{\iota^3}{d}} + \sqrt{\frac{d\iota^{p+2}}{n}}.$$

Next,

$$\begin{aligned}&[g(w) - g_n(w)] \cdot x_i \\ &= x_j \cdot \left[\frac{1}{n} \sum_{i \neq j} \left[\widehat{f}^*(x_i) x_i \sigma'(w \cdot x_i) - g(w) \right] \right] + \frac{1}{n} \left[\widehat{f}^*(x_j) \|x_j\|^2 \sigma'(w \cdot x_j) - g(w) \right].\end{aligned}$$

Note that in the first term, the x_j and the sum are independent. Therefore by Corollary 34 the first term is bounded with high probability by $O\left(\sqrt{\frac{d\iota^{p+2}}{n}}\right)$. In addition, by Lemma 30, the second term is bounded by $O\left(\frac{\iota^{p/2}d}{n}\right)$ which completes the proof. \blacksquare

A.2. Random Feature Approximation

A.2.1. UNIVARIATE RANDOM FEATURE APPROXIMATION

This section shows that after we reinitialize the biases we can use random features to transform the activation $\sigma(x) = \text{ReLU}(x)$ into $\sigma(x) = x^p$ which is more natural for learning polynomials.

Lemma 17 *Let $a \sim \text{Unif}(\{-1, 1\})$, and $b \sim \text{Unif}([-1, 1])$. Then for any $k \geq 0$ there exists $v_k(a, b)$ such that for $|x| \leq 1$,*

$$\mathbb{E}[v_k(a, b)\sigma(ax + b)] = x^k \quad \text{and} \quad \sup_{a, b} |v_k(a, b)| \lesssim 1.$$

Proof First, for $k = 0$ we can take $v_0(a, b) := 6b$. Then,

$$\begin{aligned} \mathbb{E}[v_0(a, b)\sigma(ax + b)] &= \frac{3}{2} \int_{-1}^1 b[\sigma(x + b) + \sigma(-x + b)]db \\ &= \frac{3}{2} \left[\int_{-x}^1 b(x + b)db + \int_x^1 b(-x + b)db \right] \\ &= 1 \end{aligned}$$

and $\sup_{a, b} |v_0(a, b)| = 6$. Next, for $k = 1$ we can take $v_1(a, b) := 2a$. Then,

$$\begin{aligned} \mathbb{E}[v_1(a, b)\sigma(ax + b)] &= \frac{1}{2} \int_{-1}^1 [\sigma(x + b) - \sigma(-x + b)]db \\ &= \frac{1}{2} \left[\int_{-x}^1 (x + b)db - \int_x^1 (-x + b)db \right] \\ &= x \end{aligned}$$

and we have $\sup_{a, b} |v_1(a, b)| = 2$. Next, note that by integration by parts we have for any function f ,

$$\begin{aligned} \mathbb{E}[2(1 - a)f''(b)\sigma(ax + b)] &= \int_x^1 f''(b)(-x + b)db \\ &= f'(1)(-x + 1) - \int_x^1 f'(b)db \\ &= f(x) + f'(1)(-x + 1) - f(1) \\ &= f(x) + [f'(1) - f(1)] - f'(1)x. \end{aligned}$$

Therefore for $k \geq 2$ if $f(x) = x^k$ and

$$v_k(a, b) := 2(1 - a)f''(b) - [f'(1) - f(1)]v_0(a, b) + f'(1)v_1(a, b)$$

we have

$$\mathbb{E}[v_k(a, b)\sigma(ax + b)] = x^k \quad \text{and} \quad \sup_{a, b} |v_k(a, b)| \lesssim 1.$$

■

Corollary 18 *Let $a \sim \text{Unif}(\{-1, 1\})$, and $b \sim N(0, 1)$. Then for any $k \geq 0$ there exists $v_k(a, b)$ such that for $|x| \leq 1$,*

$$\mathbb{E}[v_k(a, b)\sigma(ax + b)] = x^k \quad \text{and} \quad \sup_{a, b} |v_k(a, b)| \lesssim 1.$$

Proof Let \bar{v}_k be the function constructed in Lemma 17 and let

$$v_k(a, b) = \mathbf{1}_{|b| \leq 1} \frac{\bar{v}_k(a, b)}{2\mu(b)}$$

where $\mu(b) := \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}$ denotes the density of b . Then,

$$\begin{aligned} \mathbb{E}_{a, b}[v_k(a, b)\sigma(ax + b)] &= \mathbb{E}_a \left[\int_b v_k(a, b)\sigma(ax + b)\mu(b)db \right] \\ &= \mathbb{E}_{a, b \sim \text{Unif}([-1, 1])} [\bar{v}_k(a, b)\sigma(ax + b)] \\ &= x^k \end{aligned}$$

and

$$\sup_{a, b} |v_k(a, b)| = \sup_{a, b \in [-1, 1]} \left[\frac{\bar{v}_k(a, b)}{2\mu(b)} \right] \lesssim 1.$$

■

A.2.2. MULTIVARIABLE RANDOM FEATURE APPROXIMATION

Definition 19 *For $\|w\| = 1$, we define*

$$r(w) := g_n(w) - \frac{Hw}{\sqrt{2\pi}}.$$

Recall that Corollary 14 shows that with high probability, $\|r(w)\| \lesssim \tilde{O}\left(\sqrt{\frac{d}{n}} + \sqrt{\frac{r}{d^2}}\right)$.

Lemma 20 *With high probability over the data $\{x_i\}_{i \in [n]}$, we have for $j \leq 4p$,*

$$\mathbb{E}_w \left[\|\Pi^\star r(w)\|^j \right]^{1/j} \lesssim \sqrt{\frac{d\mu^{p+1}}{n}} + \sqrt{\frac{r^2}{d^3}}.$$

Proof We can decompose $r(w) = [g_n(w) - g(w)] + \left[g(w) - \frac{Hw}{\sqrt{2\pi}}\right]$ and note that

$$\begin{aligned} \mathbb{E}_w \left[\|\Pi^\star r(w)\|^j \right]^{1/j} &\leq \mathbb{E}_w \left[\|\Pi^\star [g_n(w) - g(w)]\|^j \right]^{1/j} + \mathbb{E}_w \left[\left\| \Pi^\star \left[g(w) - \frac{Hw}{\sqrt{2\pi}} \right] \right\|^j \right]^{1/j} \\ &\leq \mathbb{E}_w \left[\left\| \Pi^\star \left[g(w) - \frac{Hw}{\sqrt{2\pi}} \right] \right\|^j \right]^{1/j} + O\left(\sqrt{\frac{d\mu^{p+1}}{n}}\right). \end{aligned}$$

Recall that

$$\begin{aligned} g(w) - \frac{Hw}{\sqrt{2\pi}} \\ = \left(\frac{C_1 - \beta}{2} + \frac{w(C_0 - \alpha)}{\sqrt{2\pi}} \right) + \sum_{k \geq 2} \frac{c_{2k} C_{2k}(w^{\otimes 2k-1})}{(2k-1)!} + w \sum_{k \geq 1} \frac{c_{2k+2} C_{2k}(w^{\otimes 2k})}{(2k)!}. \end{aligned}$$

Therefore,

$$\Pi^* \left[g(w) - \frac{Hw}{\sqrt{2\pi}} \right] = \sum_{k \geq 2} \frac{c_{2k} C_{2k}(w^{\otimes 2k-1})}{(2k-1)!} + \Pi^* w \sum_{k \geq 1} \frac{c_{2k+2} C_{2k}(w^{\otimes 2k})}{(2k)!} + O\left(\frac{d\ell^{p+1}}{n}\right).$$

We can bound the j th moment term by term. We have by Corollary 39 and Lemma 43 that for $k \geq 2$,

$$\left(\mathbb{E}_w \left\| C_{2k}(w^{\otimes 2k-1}) \right\|^j \right)^{1/j} \lesssim \sqrt{\frac{r^{\lfloor \frac{2k-1}{2} \rfloor}}{d^{2k-1}}} \lesssim \sqrt{\frac{r}{d^3}}$$

and for $k \geq 1$,

$$\begin{aligned} \left(\mathbb{E}_w \left[\left\| \Pi^* w \right\| \left| C_{2k}(w^{\otimes 2k}) \right| \right]^j \right)^{1/j} &\lesssim \left(\mathbb{E}_w \left\| \Pi^* w \right\|^{2j} \right)^{\frac{1}{2j}} \left(\mathbb{E}_w \left| C_{2k}(w^{\otimes 2k}) \right|^{2j} \right)^{\frac{1}{2j}} \\ &\lesssim \sqrt{\frac{r}{d}} \sqrt{\frac{r^k}{d^{2k}}} \\ &\lesssim \sqrt{\frac{r^2}{d^3}}. \end{aligned}$$

■

We can now show that the random features $g_n(w)$ are sufficiently expressive to allow us to efficiently represent any polynomial of degree p restricted to the principal subspace S^* .

Lemma 21 *For any $k \leq p$, there exists an absolute constant C such that if $n \geq Cd^2r\kappa^2\ell^{p+1}$ and $d \geq C\kappa r^{3/2}$,*

$$\text{Mat} \left(\mathbb{E} \left[(\Pi^* g_n(w))^{\otimes 2k} \right] \right) \gtrsim (rd\kappa^2)^{-k} \Pi_{\text{Sym}^k(S^*)}$$

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

Proof Note that because every vector in $\text{span} \text{Mat} \left(\mathbb{E} \left[(\Pi^* g_n(w))^{\otimes 2k} \right] \right)$ is a vectorized symmetric k tensor, it suffices to show that

$$\mathbb{E}_w [(\Pi^* g_n(w))^{\otimes 2k}](T, T) \gtrsim (rd\kappa^2)^{-k}$$

for all symmetric k tensor T with $\|T\|_F^2 = 1$. Recall that $g_n(w) = \frac{Hw}{\sqrt{2\pi}} + r(w)$. Therefore by the binomial theorem,

$$\left\langle T, (\Pi^* g_n(w))^{\otimes k} \right\rangle = \left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle + \delta(w)$$

where $|\delta(w)| \lesssim \sum_{i=1}^k \|T((Hw)^{\otimes k-i})\|_F \|\Pi^\star r(w)\|^i$. Therefore by Young's inequality,

$$\mathbb{E} \left[\left\langle T, (\Pi^\star g_n(w))^{\otimes k} \right\rangle^2 \right] \geq \mathbb{E} \left[\frac{1}{2} \left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right] - \mathbb{E}[\delta(w)^2].$$

Next by Cauchy-Schwarz,

$$\begin{aligned} \mathbb{E}[\delta(w)^2] &\lesssim \sum_{i=1}^k \sqrt{\mathbb{E}_w \left[\|T((Hw)^{\otimes k-i})\|_F^4 \right] \mathbb{E}_w \left[\|\Pi^\star r(w)\|^{4i} \right]} \\ &\lesssim \sum_{i=1}^k \mathbb{E}_w \left[\left\| T \left((Hw)^{\otimes k-i} \right) \right\|_F^2 \right] \sqrt{\mathbb{E}_w \left[\|\Pi^\star r(w)\|^{4i} \right]}. \end{aligned}$$

Let \hat{T} be the symmetric k tensor defined by $\hat{T}(v_1, \dots, v_k) = T(Hv_1, \dots, Hv_k)$. Then by Lemma 47,

$$\begin{aligned} \mathbb{E} \left[\left\| T \left((Hw)^{\otimes k-i} \right) \right\|_F^2 \right] &\leq \frac{1}{\lambda_{\min}(H)^{2i}} \mathbb{E} \left[\left\| \hat{T} \left(w^{\otimes k-i} \right) \right\|^2 \right] \\ &\lesssim \left(\frac{d}{\lambda_{\min}(H)^2} \right)^i \mathbb{E} \left[\left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right] \\ &= (rd\kappa^2)^i \mathbb{E} \left[\left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right]. \end{aligned}$$

Therefore,

$$\mathbb{E}_w[\delta(w)^2] \lesssim \mathbb{E} \left[\left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right] \sum_{i=1}^k \left((rd\kappa^2) \left(\frac{dt^{p+1}}{n} + \frac{r^2}{d^3} \right) \right)^i.$$

Because we assumed $n \geq Cd^2r\kappa^2t^{p+1}$ and $d \geq C\kappa r^{3/2}$ for a sufficiently large constant C , we have

$$\mathbb{E}_w[\delta(w)^2] \lesssim \frac{1}{4} \mathbb{E} \left[\left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right].$$

Combining everything gives

$$\begin{aligned} \mathbb{E} \left[\left\langle T, (\Pi^\star g(w))^{\otimes k} \right\rangle^2 \right] &\geq \mathbb{E} \left[\frac{1}{2} \left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right] - \mathbb{E}[\delta(w)^2] \\ &\geq \frac{1}{4} \mathbb{E} \left[\left\langle T, \left(\frac{Hw}{\sqrt{2\pi}} \right)^{\otimes k} \right\rangle^2 \right] \\ &\gtrsim d^{-k} \left\| \hat{T} \right\|_F^2 \\ &\geq d^{-k} \lambda_{\min}(H)^{2k} \\ &= (rd\kappa^2)^{-k}. \end{aligned}$$

■

Corollary 22 Assume $n \geq Cd^2r\kappa^2\iota^{p+1}$ and $d \geq C\kappa r^{3/2}$ for a sufficiently large constant C . Then for any $k \leq p$ and any symmetric k tensor T supported on S^* , there exists $z_T(w)$ such that

$$\mathbb{E}_w[z_T(w)(g_n(w) \cdot x)^p] = \langle T, x^{\otimes k} \rangle$$

and we have the bounds

$$\mathbb{E}_w[z_T(w)^2] \lesssim (rd\kappa^2)^k \|T\|_F^2 \quad \text{and} \quad |z_T(w)| \lesssim (rd\kappa^2)^k \|T\|_F \|g_n(w)\|^k.$$

Proof Let

$$z_T(w) := \text{Vec}(T)^T \text{Mat}(\mathbb{E}[g_n(w)^{\otimes 2k}])^\dagger \text{Vec}(g_n(w)^{\otimes k}).$$

Note that $\text{Vec}(T) \in \text{span}(\text{Mat}(\mathbb{E}[g(w)^{\otimes 2k}]))$ by Lemma 21. Therefore,

$$\begin{aligned} & \mathbb{E}_w[z_T(w)(g_n(w) \cdot x)^k] \\ &= \mathbb{E}_w \left[\text{Vec}(T)^T \text{Mat}(\mathbb{E}[g_n(w)^{\otimes 2k}])^\dagger \text{Vec}(g(w)^{\otimes k}) \text{Vec}(g_n(w)^{\otimes k})^T \text{Vec}(x^{\otimes k}) \right] \\ &= \langle T, x^{\otimes k} \rangle. \end{aligned}$$

For the bounds on z we have

$$\begin{aligned} & \mathbb{E}_w[z_T(w)^2] \\ &= \mathbb{E}_w[\text{Vec}(T)^T \text{Mat}(\mathbb{E}[g_n(w)^{\otimes 2k}])^\dagger \text{Vec}(g(w)^{\otimes k}) \text{Vec}(g_n(w)^{\otimes k})^T \text{Mat}(\mathbb{E}[g(w)^{\otimes 2k}])^\dagger \text{Vec}(T)] \\ &= \text{Vec}(T)^T \text{Mat}(\mathbb{E}[g_n(w)^{\otimes 2k}])^\dagger \text{Vec}(T) \\ &\lesssim (rd\kappa^2)^k \|T\|_F^2 \end{aligned}$$

and

$$\begin{aligned} |z_T(w)| &= \left| \text{Vec}(T)^T \text{Mat}(\mathbb{E}[g_n(w)^{\otimes 2k}])^\dagger \text{Vec}(g_n(w)^{\otimes k}) \right| \\ &\lesssim (rd\kappa^2)^k \|T\|_F \|g_n(w)\|^k. \end{aligned}$$

■

Lemma 23 Assume $n \geq Cd^2r\kappa^2\iota^{p+1}$ and $d \geq C\kappa r^{3/2}$ for a sufficiently large constant C . Let $\eta_1 = \sqrt{\frac{d}{C^2\iota^3}}$, let $k \leq p$ and let T be a k tensor. Then with high probability, there exists $h_T(a, w, b)$ such that if

$$f_{h_T}(x) := \mathbb{E}_{a,w,b}[h_T(a, w, b)\sigma(w^{(1)} \cdot x + b)]$$

we have

$$\frac{1}{n} \sum_{i=1}^n (f_{h_T}(x_i) - \langle T, x_i^{\otimes p} \rangle)^2 \lesssim \frac{1}{n}$$

and the moment bounds

$$\begin{aligned} \mathbb{E}_{w,a,b}[h_T(a, w, b)^2] &\lesssim r^k \kappa^{2k} \iota^{3k} \|T\|_F^2 \\ \sup_w |h_T(a, w, b)| &\lesssim r^k \kappa^{2k} \iota^{6k} \|T\|_F. \end{aligned}$$

Proof We define

$$h_T(a, w, b) := \frac{v_k(a, b)z_T(w)}{(2\eta)^k} \mathbf{1}_{\eta_1 \|g_n(w)\| \leq 1} \prod_{i=1}^n \mathbf{1}_{|g_n(w) \cdot x_i| \leq 1}.$$

where $v_k(a, b)$ and $z_T(w)$ are constructed in Corollary 18 and Corollary 22 respectively. Recall that $w^{(1)} = 2\eta_1 a g_n(w)$. Then for $x \in \{x_1, \dots, x_n\}$,

$$\begin{aligned} & f_{h_T}(x) \\ &= \frac{1}{(2\eta_1)^k} \mathbb{E}_{a,w,b} [v_k(a, b)z_T(w)\sigma(2\eta_1 a g_n(w) \cdot x + b)] \\ &= \mathbb{E}_w \left[z_T(w) \left[(g_n(w) \cdot x)^k + O\left(|g_n(w) \cdot x|^k \left[1 - \mathbf{1}_{\eta_1 \|g_n(w)\| \leq 1} \prod_{i=1}^n \mathbf{1}_{|g_n(w) \cdot x_i| \leq 1} \right] \right) \right] \right] \\ &= \langle T, x^{\otimes k} \rangle + \text{poly}(d) \left[\mathbb{P}_w[\eta_1 \|g_n(w)\| \geq 1] + \sum_{i=1}^n \mathbb{P}_w[|2\eta_1 g_n(w) \cdot x_i| \geq 1] \right] \\ &= \langle T, x^{\otimes k} \rangle + \text{poly}(n, d) e^{-\iota} \\ &= \langle T, x^{\otimes k} \rangle + O\left(\frac{1}{n}\right) \end{aligned}$$

where the second to last line followed from Lemma 16. The first part of the lemma now follows from a union bound over x_1, \dots, x_n . For the bounds on h , we have

$$\begin{aligned} & \mathbb{E}_{a,w,b}[h(a, w, b)^2] \\ &= \frac{1}{(2\eta_1)^{2k}} \mathbb{E}_{a,w,b} [v_k(a, b)^2 z_T(w)^2] \\ &\lesssim \eta_1^{-2k} (rd\kappa^2)^k \|T\|_F^2 \\ &= r^k \kappa^{2k} \iota^{3k} \|T\|_F^2. \end{aligned}$$

and

$$\begin{aligned} \sup_w |h(a, w, b)| &= \sup_w \left| \frac{v_k(a, b)z_T(w)}{(2\eta)^k} \mathbf{1}_{\eta_1 \|g_n(w)\| \leq 1} \right| \\ &\lesssim \eta_1^{-k} (rd\kappa^2)^k \|g_n(w)\|^k \|T\|_F \\ &= \eta_1^{-2k} (rd\kappa^2)^k \|T\|_F \\ &\lesssim r^k \kappa^{2k} \iota^{6k} \|T\|_F. \end{aligned}$$

■

Corollary 24 Assume $n \geq Cd^2r\kappa^2\iota^{p+1}$ and $d \geq C\kappa r^{3/2}$ for a sufficiently large constant C and let $\eta_1 = \sqrt{\frac{d}{\iota^3}}$. Then with high probability, there exists $h(a, w, b)$ such that if

$$f_h(x) := \mathbb{E}_{a,w,b}[h(a, w, b)\sigma(w^{(1)} \cdot x + b)]$$

we have

$$\frac{1}{n} \sum_{i=1}^n (f_h(x_i) - f^*(x_i))^2 \lesssim \frac{1}{n}$$

and the moment bounds

$$\begin{aligned} \mathbb{E}_{w,a,b}[h_T(a, w, b)^2] &\lesssim r^p \kappa^{2p} \iota^{3p} \\ \sup_w |h_T(a, w, b)| &\lesssim r^p \kappa^{2p} \iota^{6p}. \end{aligned}$$

Proof We know from Lemma 35 that

$$f^*(x) = \sum_{k \leq p} \langle T_k, x^{\otimes k} \rangle$$

with $\|T_k\|_F \lesssim r^{\frac{p-k}{4}}$. Let

$$h(a, w, b) := \sum_{k \leq p} h_{T_k}(a, w, b).$$

Then $\frac{1}{n} \sum_{i=1}^n (f_h(x_i) - f^*(x_i))^2 \lesssim \frac{1}{n}$ is immediate from Lemma 23 and

$$\mathbb{E}_{a,w,b}[h(a, w, b)^2] \lesssim \sum_{k \leq p} \mathbb{E}_{a,w,b}[h_{T_k}(a, w, b)^2] \lesssim \sum_{k \leq p} r^k \kappa^{2k} \iota^{3k} r^{\frac{p-k}{2}} \lesssim r^p \kappa^{2p} \iota^{3p}.$$

and

$$\sup_{a,w,b} |h(a, w, b)| \leq \sum_{k \leq p} \sup_{a,w,b} |h_{T_k}(a, w, b)| \lesssim \sum_{k \leq p} r^k \kappa^{2k} \iota^{6k} r^{\frac{p-k}{2}} \lesssim r^p \kappa^{2p} \iota^{6p}.$$

complete the proof. ■

Lemma 25 Assume $n \geq Cd^2 r \kappa^2 \iota^{p+1}$, $d \geq C \kappa r^{3/2}$, and $m \geq r^p \kappa^{2p} \iota^{6p+1}$ for a sufficiently large constant C and let $\eta_1 = \sqrt{\frac{d}{\iota^3}}$. Then with high probability, there exists $a^* \in \mathbb{R}^m$ such that if $\theta^* = (a^*, W^{(1)}, b^{(1)})$,

$$\frac{1}{n} \sum_{i=1}^n (f_{\theta^*}(x_i) - f^*(x_i))^2 \lesssim \frac{1}{n} + \frac{r^p \kappa^{2p} \iota^{6p+1}}{m} \quad \text{and} \quad \|a^*\|^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p}}{m}.$$

Proof Let $a_j^* := \frac{1}{m} h(a_j, w_j, b_j)$ where h is the function constructed in Lemma 24. Then,

$$\begin{aligned} \mathbb{E}_{i \in [n]} [(f_{\theta^*}(x_i) - f^*(x_i))^2] &\lesssim \mathbb{E}_{i \in [n]} [(f_{\theta^*}(x_i) - f_h(x_i))^2] + \mathbb{E}_{i \in [n]} [(f_h(x_i) - f^*(x_i))^2] \\ &= \mathbb{E}_{i \in [n]} [(f_{\theta^*}(x_i) - f_h(x_i))^2] + \frac{1}{n}. \end{aligned}$$

For $j \leq m/2$, let

$$Z_j(x) := a_j^* \sigma(w_j^{(1)} \cdot x + b_j) + a_{m-j}^* \sigma(w_{m-j}^{(1)} \cdot x + b_{m-j}).$$

Note that

$$f_{\theta^*}(x) - f_h(x) = \sum_{j \leq \frac{m}{2}} (Z_j(x) - \mathbb{E}[Z_j(x)])$$

and the $Z_j(x)$ are all i.i.d.. Let

$$\bar{Z}_j(x) := Z_j(x) \mathbf{1}_{|w_j^{(1)} \cdot x| \leq 1} \mathbf{1}_{|w_{m-j}^{(1)} \cdot x| \leq 1}.$$

Then with probability $1 - \text{poly}(n, m, d)e^{-\ell}$ we have that $Z_j(x_i) = \bar{Z}_j(x_i)$ for $i = 1, \dots, n$. Therefore,

$$f_{\theta^*}(x) - f_h(x) = \sum_{j \leq \frac{m}{2}} (\bar{Z}_j(x) - \mathbb{E}[\bar{Z}_j(x)]) + \frac{m}{2} [\mathbb{E} \bar{Z}_j(x) - \mathbb{E} Z_j(x)].$$

For the first term, by Bernstein's inequality we have with probability at least $1 - 2e^{-\ell}$,

$$\sum_{j \leq \frac{m}{2}} \bar{Z}_j(x) - \mathbb{E}[\bar{Z}_j(x)] \lesssim \sqrt{\frac{\ell \mathbb{E}_{a,w,b}[h(a, w, b)^2]}{m}} + \frac{\ell r^p \kappa^{2p} \ell^{3p}}{m} \lesssim \sqrt{\frac{r^p \kappa^{2p} \ell^{6p+1}}{m}}.$$

The second term is bounded as in the proof of Lemma 23 by $\text{poly}(n, d)e^{-\ell} \leq \frac{1}{m}$ because $\mathbb{P}[w^{(1)} \cdot x > 1] \leq e^{-\ell}$ from the choice of η_1 . Therefore for any fixed x , with high probability we have

$$f_{\theta^*}(x) = f^*(x) + O\left(\frac{1}{n} + \sqrt{\frac{r^p \kappa^{2p} \ell^{6p+1}}{m}}\right)$$

and the first part of the lemma follows from a union bound.

We will now turn to the bound on $\|a^*\|^2$. Let $z_i = (a_i^*)^2 + (a_{m-i}^*)^2$. Note that $\{z_i\}_{i \leq m/2}$ are positive, i.i.d., and bounded by $O(m^{-2} r^{2p} \kappa^{4p} \ell^{12p})$. In addition, they have expectation $O(m^{-2} r^p \kappa^{2p} \ell^{3p})$. Therefore by Popoviciu's inequality they have variance bounded by

$$O([m^{-1} r^p \kappa^{2p} \ell^{3p}] [m^{-2} r^{2p} \kappa^{4p} \ell^{12p}]) = O(m^{-3} r^{3p} \kappa^{6p} \ell^{15p}).$$

Therefore by Bernstein's inequality we have that with high probability,

$$\|a^*\|^2 = \mathbb{E}[\|a^*\|^2] + O\left(\frac{1}{m} \sqrt{\frac{r^{3p} \kappa^{6p} \ell^{15p}}{m}} + \frac{r^{2p} \kappa^{3p} \ell^{6p}}{m^2}\right) \lesssim \frac{r^p \kappa^{2p} \ell^{6p}}{m}.$$

■

A.3. Proof of Theorem 1

We will define

$$\tilde{\mathcal{L}}(\theta)(\theta) := \frac{1}{n} \sum_{i=1}^n (f_\theta(x_i) - f^*(x_i))^2.$$

to be the empirical L^2 losses with respect to the true labels (recall $y_i = f^*(x_i) + \epsilon_i$, $\epsilon_i \sim \{-\sigma, \sigma\}$).

Lemma 26 *Assume $n \geq Cd^2r\kappa^2\iota^{p+1}$ and $d \geq C\kappa r^{3/2}$ for a sufficiently large constant C and let $\eta_1 = \sqrt{\frac{d}{\iota^3}}$. Let a^* be the vector constructed in the proof of Lemma 25 and let $\theta = (a^*, W^{(1)}, b^{(1)})$. Then with high probability,*

$$\mathcal{L}(\theta) - \zeta^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p+1}}{m} + \sqrt{\frac{\iota}{n}}.$$

Proof Let $\delta_i = f_\theta(x_i) - f^*(x_i)$. Then,

$$\frac{1}{n} \|\delta + \epsilon\|_2^2 = \frac{1}{n} \left(\|\delta\|^2 + 2\langle \delta, \epsilon \rangle + \|\epsilon\|_2^2 \right).$$

First, by Hoeffding's inequality, we have with high probability,

$$\frac{\|\epsilon\|^2}{n} \leq \zeta^2 + \frac{C\zeta^2 \sqrt{\iota}}{\sqrt{n}} = \zeta^2 + O\left(\sqrt{\frac{\iota}{n}}\right).$$

Similarly, by Hoeffding's inequality we have with high probability, $\frac{1}{n} \langle \delta, \epsilon \rangle \leq \zeta \sqrt{\frac{2\iota \tilde{\mathcal{L}}(\theta)}{n}} = O\left(\sqrt{\frac{\iota}{n}}\right)$. ■

We are now ready to directly prove Theorem 1.

Proof [Proof of Theorem 1] Note that we can assume that there is an absolute constant C such that $n \geq Cd^2r\kappa^2\iota^{p+1}$, and $m \geq r^p \kappa^{2p} \iota^{6p+1}$. Otherwise, we can simply take $\lambda \rightarrow \infty$ and return the zero predictor.

From Lemma 26 we know that with high probability, there exists a^* such that if $\theta = (a^*, W^{(1)}, b^{(1)})$,

$$\mathcal{L}(\theta) - \zeta^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p+1}}{m} + \sqrt{\frac{\iota}{n}}.$$

and $\|a^*\|_2^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p+1}}{m}$. Therefore by equality of norm constrained linear regression and ridge regression, there exists $\lambda > 0$ such that if

$$a^{(\infty)} = \min_a \mathcal{L}\left((a, W^{(1)}, b^{(1)})\right) + \lambda \frac{\|a\|^2}{2},$$

$$\mathcal{L}\left((a^{(\infty)}, W^{(1)}, b^{(1)})\right) \leq \mathcal{L}\left((a^*, W^{(1)}, b^{(1)})\right) \quad \text{and} \quad \|a_\infty\| \leq \|a^*\|.$$

Note that we can approximate $a^{(\infty)}$ by $a^{(T)}$ to within arbitrary accuracy within $T = \tilde{\Theta}(\eta^{-1}\lambda^{-1})$ steps. Let

$$\mathcal{F} = \{f_\theta : \|a\|_2 \leq \|a^*\|, \|w_j\| \leq 1\}.$$

Then with high probability, $f_{(a^{(T)}, W^{(1)}, b^{(1)})} \in \mathcal{F}$. In addition, from Lemma 49,

$$\begin{aligned} \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n |f(x_i) - y_i| - \mathbb{E}_{x,y} |f(x) - y| \right| &\lesssim \sqrt{\frac{\|a^*\|^2 md}{n}} + \sqrt{\frac{\iota}{n}} \\ &\lesssim \sqrt{\frac{dr^p \kappa^{2p} \iota^{6p}}{n}}. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}_{x,y} |f_{\theta^{(T)}}(x) - y| - \mathbb{E}_{x,y} |f^*(x) - y| \\ \lesssim \sqrt{\frac{dr^p \kappa^{2p} \iota^{6p}}{n}} + \sqrt{\frac{r^p \kappa^{2p} \iota^{6p+1}}{m}} + \left(\frac{\iota}{n}\right)^{1/4}. \end{aligned}$$

which completes the proof. ■

Appendix B. Transfer Learning

Proof [Proof of Theorem 3] The proof of Theorem 3 is virtually identical to that of Theorem 1. We can use Lemma 25 to construct a^* such that if $\theta^* = (a^*, W^{(1)}, b^{(1)})$ then with high probability,

$$L(\theta^*) - \varsigma^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p+1}}{m} + \sqrt{\frac{\iota}{n}} \quad \text{and} \quad \|a^*\|^2 \lesssim \frac{r^p \kappa^{2p} \iota^{6p}}{m}.$$

In addition, there exists λ such that if $T \geq \Theta(\eta^{-1}\lambda^{-1})$,

$$L(\theta^{(T)}) \leq L(\theta^*) \quad \text{and} \quad \|a^{(T)}\| \leq \|a^*\|.$$

Now let $\mathcal{F} = \{f_{(a, W, b)} : \|a\|_2 \leq \|a^*\|\}$. Then by Lemma 48,

$$\mathbb{E}_{x,y} |g_{a^{(T)}}(x) - y| - \varsigma \leq \tilde{O} \left(\sqrt{\frac{d^2 r}{n}} + \sqrt{\frac{r^p \kappa^{2p}}{\min(n, m)}} \right).$$
■

Appendix C. Concentration Lemmas

Lemma 27 (Corollary of Lemma 1 in (??)) *Let $X \sim \chi^2(d)$. Then, for any $t \geq 0$,*

$$\begin{aligned} \mathbb{P}[X \geq d + 2\sqrt{dt} + 2t] &\leq \exp(-t) \\ \mathbb{P}[X \leq d - 2\sqrt{dt}] &\leq \exp(-t). \end{aligned}$$

Corollary 28 *Let $w \sim N(0, I_d)$. Then for some constant C ,*

$$\mathbb{P} \left[\frac{\|\Pi^* w\|^2}{\|w\|^2} \in \left[\frac{r}{Cd}, \frac{Cr}{d} \right] \right] \gtrsim 1.$$

Lemma 29 (Corollary 5.35 in Vershynin (2018)) *Let $X \in \mathbb{R}^{n \times d}$ with $X_{ij} \sim N(0, 1)$. Then with probability at least $1 - 2e^{-\ell}$,*

$$\|X\|_2 \leq \sqrt{n} + \sqrt{d} + \sqrt{2\ell}.$$

C.1. Polynomial Concentration

Lemma 30 *Let g be a polynomial of degree p . Then there exists an absolute constant C_p depending only on p such that for any δ ,*

$$\mathbb{P}[|g(x) - \mathbb{E}[g(x)]| \geq \delta \sqrt{\mathbb{E}[g(x)^2]}] \leq 2 \exp\left(-C_p \min(\delta^2, \delta^{2/p})\right).$$

Proof Note that by Lemma 9,

$$\left\| \mathbb{E}[\nabla^k g(x)] \right\|_{HS} \leq \sqrt{k!}.$$

Therefore by Theorem 1.2 of (Gotze et al., 2019), there exists an absolute constant C_p such that

$$\mathbb{P}[|g(x) - \mathbb{E}[g(x)]| \geq \delta \sqrt{\mathbb{E}[g(x)^2]}] \leq 2 \exp\left(-C_p \min_{1 \leq s \leq p} \delta^{2/s}\right) = 2 \exp\left(-C_p \min(\delta^2, \delta^{2/p})\right).$$

■

Lemma 31 *Let $\sigma(x) \in \{x, \text{ReLU}(x)\}$. There exists an absolute constant C such that for any $x_1, \dots, x_n \in \mathbb{R}^d$, there exists $\mathcal{N}_\epsilon^x, \pi$ with $|\mathcal{N}_\epsilon^x| \leq e^{Cd \log(n/\epsilon)}$ such that for every $w \in S^{d-1}$, $\pi(w) \in \mathcal{N}_\epsilon^x, \sigma'(w \cdot x_i) = \sigma'(\pi(w) \cdot x_i)$ for $i = 1, \dots, n$ and $\|w - \pi(x)\| \leq \epsilon$.*

Proof Note that the planes $w \cdot x_1 = 0, \dots, w \cdot x_n = 0$ divides the sphere S^{d-1} into at most $\sum_{i=0}^d \binom{n}{i} \lesssim n^d$ convex regions. For each region there exists an ϵ net of size $(\frac{3}{\epsilon})^d$. Therefore we can take the union of these nets over each region which has size at most $(\frac{3n}{\epsilon})^d = e^{Cd \log(n/\epsilon)}$. ■

Lemma 32 *Let $f(x)$ be a polynomial of degree p and let $\sigma(x) \in \{x, \text{ReLU}(x)\}$. Then there exists an absolute constant C_p depending only on p such that for any $\ell > 0$, with probability at least $1 - 2ne^{-\ell}$, we have*

$$\sup_{w \in S^{d-1}} \left\| \frac{1}{n} \sum_{i=1}^n f(x_i) x_i \sigma'(w \cdot x_i) - \mathbb{E}[f(x) x \sigma'(w \cdot x)] \right\| \leq C_p \sqrt{\mathbb{E}[g(x)^2]} \sqrt{\frac{d\ell^{p+1}}{n}}.$$

Proof Note that we may assume $\iota \geq \log(2n)$ otherwise there is nothing to prove. Let C be a sufficiently large absolute constant. We fix a truncation radius $R := (C\iota)^{p/2}$ and WLOG assume that $\mathbb{E}[f(x)^2] = 1$. Let

$$Y(w) := \frac{1}{n} \sum_{i=1}^n f(x_i) x_i \sigma'(w \cdot x_i) \quad \text{and} \quad \tilde{Y}(w) := \frac{1}{n} \sum_{i=1}^n f(x_i) x_i \sigma'(w \cdot x_i) \mathbf{1}_{\{|f(x_i)| \leq R\}}.$$

First, note that by Lemma 30, with probability at least $1 - 2e^{-2\iota}$, we have $|f(x)| \leq R$. Therefore by a union bound we have with probability at least $1 - 2ne^{-2\iota}$ we have $|f(x_i)| \leq R$ for $i = 1, \dots, n$. Conditioned on this event, $Y(w) = \tilde{Y}(w)$ uniformly over all $w \in S^{d-1}$. Next, we will bound $\sup_w \|\mathbb{E}_x[Y(w)] - \mathbb{E}_x[\tilde{Y}(w)]\|$:

$$\begin{aligned} \sup_w \|\mathbb{E}_x[Y(w)] - \mathbb{E}_x[\tilde{Y}(w)]\| &= \sup_w \|\mathbb{E}_x[g(x)x\sigma'(w \cdot x)\mathbf{1}_{\{|g(x_i)| > R\}}]\| \\ &\leq \mathbb{E}_x [|f(x)| \|x\| \mathbf{1}_{\{|g(x_i)| > R\}}] \\ &\leq \mathbb{E}[g(x)^2]^{1/2} \mathbb{E}[\|x\|^4]^{1/4} \mathbb{P}[|g(x_i)| > R]^{1/4} \\ &\leq 2\sqrt{2d} \exp(\iota/2) \\ &\lesssim \sqrt{\frac{d}{n}}. \end{aligned}$$

Finally, we concentrate $\sup_w \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\|$. Let $\epsilon = \sqrt{\frac{d}{n}}$, let $\mathcal{N}_{1/4}$ be a minimal $1/4$ -net of S^{d-1} with $|\mathcal{N}_{1/4}| \leq e^{Cd}$ and let \mathcal{N}_ϵ^x be the net defined in Lemma 31 with $|\mathcal{N}_\epsilon^x| \leq e^{Cd\log(n/\epsilon)}$ and let $\pi(w)$ be the projection function defined in Lemma 31. Then because $\tilde{Y}(w) = \tilde{Y}(\pi(w))$,

$$\begin{aligned} &\sup_w \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| \\ &\leq \sup_{w \in \mathcal{N}_\epsilon^x} \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| + \sup_w \|\mathbb{E}_x[\tilde{Y}(w)] - \mathbb{E}_x[\tilde{Y}(\pi(w))]\| \\ &\leq \sup_{w \in \mathcal{N}_\epsilon^x} \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| + \sup_w \|\mathbb{E}_x[Y(w)] - \mathbb{E}_x[Y(\pi(w))]\| + O\left(\sqrt{\frac{d}{n}}\right). \end{aligned}$$

Next, because $w \rightarrow \mathbb{E}_x[Y(w)]$ is $O(1)$ Lipschitz (see Section A.1.4), we can bound this by

$$\sup_w \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| \leq \sup_{w \in \mathcal{N}_\epsilon^x} \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| + O\left(\epsilon + \sqrt{\frac{d}{n}}\right).$$

Therefore it remains to bound $\sup_{w \in \mathcal{N}_\epsilon^x} \|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\|$. First, for fixed w we have

$$\|\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]\| = \sup_{u \in S^{d-1}} u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]] \leq 2 \sup_{u \in \mathcal{N}_{1/4}} u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]].$$

Let $Z_i(w) := g(x_i)(u \cdot x_i)\sigma'(w \cdot x_i)\mathbf{1}_{\{|g(x)| < R\}}$ so that

$$u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]] = \frac{1}{n} \sum_{i=1}^n Z_i(w) - \mathbb{E}_x[Z_i(w)].$$

Then note that for fixed w , $Z_i(w)$ is R -sub Gaussian so for each $u \in \mathcal{N}_{1/4}$, with probability $1 - 2e^{-z}$ we have

$$u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]] \leq R \sqrt{\frac{2z}{n}}.$$

so by a union bound we have with probability $1 - 2e^{Cd\log(n/\epsilon)}e^{-z}$,

$$2 \sup_{u \in \mathcal{N}_{1/4}, w \in \mathcal{N}_\epsilon^x} u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]] \leq 2R \sqrt{\frac{2z}{n}}.$$

so setting $z = Cd\log(n/\epsilon) + \iota$ we have with probability $1 - 2e^\iota$,

$$2 \sup_{u \in \mathcal{N}_{1/4}, w \in \mathcal{N}_\epsilon^x} u \cdot [\tilde{Y}(w) - \mathbb{E}_x[\tilde{Y}(w)]] \lesssim R \sqrt{\frac{d\log(n/\epsilon) + \iota}{n}}.$$

Using $\epsilon = \sqrt{\frac{d}{n}}$ and putting everything together gives with probability $1 - 2ne^{-\iota}$,

$$\sup_w \|Y(w) - \mathbb{E}[Y(w)]\| \lesssim \sqrt{\frac{(d\log n + \iota)\iota^p}{n}} \lesssim \sqrt{\frac{d\iota^{p+1}}{n}}.$$

■

Lemma 33 *Let $\epsilon_i \sim \{-\varsigma, \varsigma\}$. Then with high probability,*

$$\sup_w \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i x_i \sigma'(w \cdot x_i) \right\| \lesssim \varsigma \sqrt{\frac{d\iota}{n}}.$$

Proof Note that

$$\sup_w \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i x_i \sigma'(w \cdot x_i) \right\| = \sup_{u,w} \left[\frac{1}{n} \sum_{i=1}^n \epsilon_i (u \cdot x_i) \sigma'(w \cdot x_i) \right].$$

Next, note that for fixed u, w , $\epsilon_i(u \cdot x_i) \sigma'(w \cdot x_i)$ is ς^2 sub-Gaussian so for any $\iota > 0$, with probability $1 - 2e^{-\iota}$,

$$\frac{1}{n} \sum_{i=1}^n \epsilon_i (u \cdot x_i) \sigma'(w \cdot x_i) \leq \varsigma \sqrt{\frac{\iota}{n}}.$$

Therefore,

$$\sup_{u,w} \left[\frac{1}{n} \sum_{i=1}^n \epsilon_i (u \cdot x_i) \sigma'(w \cdot x_i) \right] \lesssim \sup_{u \in \mathcal{N}_{1/4}, w \in \mathcal{N}_{1/4}^x} \sum_{i=1}^n \epsilon_i (u \cdot x_i) \sigma'(w \cdot x_i).$$

By a union bound, with probability at least $1 - 2e^\iota$,

$$\sup_{u \in \mathcal{N}_{1/4}, w \in \mathcal{N}_{1/4}^x} \sum_{i=1}^n \epsilon_i (u \cdot x_i) \sigma'(w \cdot x_i) \lesssim \varsigma \sqrt{\frac{d\log n + \iota}{n}} \lesssim \varsigma \sqrt{\frac{d\iota}{n}}$$

which completes the proof. ■

Corollary 34 *With high probability,*

$$\sup_w \|g(w) - g_n(w)\| \lesssim \sqrt{\frac{dt^{p+1}}{n}}.$$

Appendix D. CSQ Lower Bound

Proof [Proof of Lemma 5] The proof is a modified version of the proof in [Szörényi \(2009\)](#). Let $\langle \cdot, \cdot \rangle_{\mathcal{D}}$ denote the L^2 inner product with respect to \mathcal{D} . We will show that there are at least two functions $f, g \in \mathcal{F}$ such that for each query h_k , $|\langle f, h_k \rangle_{\mathcal{D}}| \leq \tau$ and $|\langle g, h_k \rangle_{\mathcal{D}}| \leq \tau$. Therefore, we can simply respond to each query adversarially with 0 and it is impossible for the learner to distinguish between f, g . Note that failing to do so will result in a loss of $\|f - g\|_{\mathcal{D}}^2 \geq 2 - 2\epsilon$. Let the k th query be h_k and let

$$A_k^+ = \{f \in \mathcal{F} : \langle f, h_k \rangle_{\mathcal{D}} \geq \tau\} \quad \text{and} \quad A_k^- = \{f \in \mathcal{F} : \langle f, h_k \rangle_{\mathcal{D}} \leq -\tau\}$$

Then by Cauchy-Schwarz we have

$$|A_k^+|^2 \tau^2 \leq \left\langle h_k, \sum_{f \in A_k^+} f \right\rangle_{\mathcal{D}}^2 \leq \left\| \sum_{f \in A_k^+} f \right\|_{\mathcal{D}}^2 = \sum_{f, g \in A_k^+} \langle f, g \rangle_{\mathcal{D}} \leq |A_k^+| + \epsilon (|A_k^+|^2 - |A_k^+|)$$

which implies

$$|A_k^+| \leq \frac{1 - \epsilon}{\tau^2 - \epsilon} \leq \frac{1}{\tau^2 - \epsilon}.$$

Similarly, we have that $|A_k^-| \leq \frac{1}{\tau^2 - \epsilon}$ so the number of functions that are eliminated from the k th query is at most $\frac{2}{\tau^2 - \epsilon}$. We can continue this process for at most $\frac{|F|(\tau^2 - \epsilon)}{2}$ iterations. \blacksquare

Proof [Proof of Lemma 6] Let $v_1, \dots, v_k \sim S^{d-1}$. Then for every pair $i \neq j$, $v_i \cdot v_j$ is $O(d^{-1})$ subgaussian so for an absolute constant c , with probability $1 - 2e^{-2c\epsilon^2 d}$, $|v_i \cdot v_j| \leq \epsilon$. Therefore with probability $1 - k^2 e^{-2c\epsilon^2 d} > 0$ this holds for all $i \neq j$ so there must exist at least one collection of such points. \blacksquare

Proof [Proof of Theorem 2] Let S be the set constructed in Lemma 6. Let

$$\mathcal{F} = \left\{ x \rightarrow \frac{He_p(v \cdot x)}{\sqrt{k!}} : v \in S \right\}$$

and note that for all $f \in \mathcal{F}$, $\|f\|_{\mathcal{D}} = 1$. Then for $v, w \in S$ and $v \neq w$,

$$\left| \left\langle \frac{He_k(v \cdot x)}{\sqrt{k!}}, \frac{He_k(w \cdot x)}{\sqrt{k!}} \right\rangle_{\mathcal{D}} \right| = |(v \cdot w)^k| \leq \epsilon^k.$$

Therefore, by Lemma 5 we have for any ϵ ,

$$4q \geq e^{c\epsilon^2 d} (\tau^2 - \epsilon^k)$$

In particular if we take $\epsilon = \sqrt{\frac{\log(4q(cd)^{k/2})}{cd}}$ we get

$$\tau^2 \leq \frac{1 + \log^{k/2}(4q(cd)^{k/2})}{(cd)^{k/2}} \lesssim \frac{\log^{k/2}(qd)}{d^{k/2}}.$$

■

Appendix E. Additional Technical Lemmas

For a k tensor T , let $\text{Sym}(T)$ denote the symmetrization of T along all $k!$ permutations of indices.

Lemma 35 *There exist T_0, \dots, T_p such that*

$$f^*(x) = \sum_{k \leq p} \langle T_k, x^{\otimes k} \rangle$$

and $\|T_k\|_F \lesssim r^{\frac{p-k}{4}}$ for $k \leq p$.

Proof Note that from the Taylor series of $f^*(x)$ we have

$$T_k = \frac{\nabla^k f^*(0)}{k!} = \sum_{j \leq p-k} \frac{C_{j+k}(He_j(0))}{k!j!} = \sum_{2j \leq p-k} \frac{(-1)^j (2j-1)!! C_{2j+k}(I^{\otimes j})}{k!(2j)!}.$$

Therefore,

$$\|T_k\|_F \lesssim \sum_{2j \leq p-k} \|C_{2j+k}(I^{\otimes j})\| \lesssim r^{\frac{p-k}{4}}.$$

■

E.1. Gaussian Lemmas

Lemma 36

$$\mathbb{E}_{w \sim N(0, I_d)}[w^{\otimes 2k}] = (2k-1)!! \text{Sym}(I_d^{\otimes k})$$

Proof We will show equality for each coordinate. Let i_1, \dots, i_{2k} be an index set and let c_1, \dots, c_d be defined by $c_j = |\{k : i_k = j\}|$. First we will consider the case there is an odd c_j . Then, $\mathbb{E}_{w \sim N(0, I_d)}[w^{\otimes 2k}]_{i_1, \dots, i_{2k}} = 0$ and $\left[(2k-1)!! \text{Sym}(I_d^{\otimes k})\right]_{i_1, \dots, i_{2k}} = 0$ because in order for this to be nonzero there must exist a pairing of i_1, \dots, i_{2k} such the numbers in each pair are identical.

Next, assume that each c_j is even. Then, $\mathbb{E}_{w \sim N(0, I_d)}[w^{\otimes 2k}]_{i_1, \dots, i_{2k}} = \prod_{j=1}^d (c_j - 1)!!$ by the standard formula for Gaussian moments. Finally, consider

$$\left[(2k-1)!! \text{Sym}(I_d^{\otimes k})\right]_{i_1, \dots, i_{2k}} = \frac{(2k-1)!!}{2k!} \sum_{\sigma} \mathbf{1}_{i_{\sigma_1} = i_{\sigma_2}} \cdots \mathbf{1}_{i_{\sigma_{2k-1}} = i_{\sigma_{2k}}}.$$

Note that by a simple counting argument, the number of permutations such that this product of indicators is nonzero is exactly $k! \prod_{i=1}^d \frac{c_j!}{(c_j/2)!}$ as you can first order the indices corresponding to each c_j , then split them into groups of two, then shuffle these groups of two. Therefore,

$$\left[(2k-1)!! I_d^k \right]_{i_1, \dots, i_{2k}} = \frac{k!(2k-1)!!}{2k!} \prod_{i=1}^d \frac{c_j!}{(c_j/2)!} = \frac{1}{2^k} \prod_{i=1}^d (c_j - 1)!! 2^{c_j/2} = \prod_{i=1}^d (c_j - 1)!!$$

because $\sum_j c_j = 2k$, which completes the proof. \blacksquare

Definition 37 Let $\{h_{kl}\}$ and $\{h_{kl}^{-1}\}$ denote the change of basis matrices between Hermite polynomials and monomials, i.e.

$$He_k(x) = \sum_{l \leq k} h_{kl} x^l \quad \text{and} \quad x^k = \sum_{l \leq k} h_{kl}^{-1} He_l(x).$$

Note that

$$h_{kl} = \begin{cases} (-1)^{\frac{k-l}{2}} (k-l-1)!! \binom{k}{l} & 2 \mid k-l \\ 0 & 2 \nmid k-l \end{cases} \quad \text{and} \quad h_{kl}^{-1} = \begin{cases} (k-l-1)!! \binom{k}{l} & 2 \mid k-l \\ 0 & 2 \nmid k-l \end{cases}.$$

Lemma 38 Let T be a symmetric p -tensor and let $w \sim N(0, I_d)$. Then for $k \leq p$,

$$\mathbb{E} \|T(w^{\otimes k})\|_F^2 = \sum_{2l \leq k} (k-2l)((2l-1)!!)^2 \binom{k}{2l}^2 \|T(I^{\otimes l})\|_F^2.$$

Proof Let $T = \sum_i c_i v_i^p$ with $\|v_i\| = 1$. Using the change of basis $x^k \rightarrow \sum_{l \leq k} h_{kl}^{-1} He_l(x)$,

$$\begin{aligned} \mathbb{E} \|T(w^{\otimes k})\|_F^2 &= \sum_{ij} c_i c_j \mathbb{E}[(w \cdot v_i)^k (w \cdot v_j)^k] (v_i \cdot v_j)^{p-k} \\ &= \sum_{l \leq k} l! (h_{kl}^{-1})^2 \sum_{ij} c_i c_j (v_i \cdot v_j)^{p-k+l} \\ &= \sum_{2l \leq k} (k-2l)((2l-1)!!)^2 \binom{k}{2l}^2 \|T(I^{\otimes l})\|_F^2. \end{aligned}$$

\blacksquare

Corollary 39 Let T be a symmetric p -tensor with $\dim(\text{span}(T)) = r$. For $k \leq p$,

$$\mathbb{E} \|T(w^{\otimes k})\|_F^2 \lesssim r^{\lfloor \frac{k}{2} \rfloor} \|T\|_F^2.$$

Proof The proof follows directly from Lemma 38 and the inequality $\|T(I^{\otimes l})\|_F = \|T(\Pi_{\text{span}(T)}^{\otimes l})\|_F \leq \|T\|_F^2 \left\| \Pi_{\text{span}(T)}^{\otimes l} \right\|_F^2 = r^l \|T\|_F^2$ for $2l \leq k$. \blacksquare

Corollary 40 *Let T be a symmetric p -tensor with $\dim(\text{span}(T)) = r$. With probability at least $1 - 2e^{-\ell}$,*

$$\|T(w^{\otimes k})\|_F \lesssim \|T\|_F \sqrt{r^{\lfloor \frac{k}{2} \rfloor} \ell^k}.$$

Proof Note that $F(w) = \|T(w^{\otimes k})\|_F^2$ is a polynomial of degree $2k$. For $k \leq p$, let \tilde{T}_k be the (k, k) tensor which comes from contracting the last $d - k$ indices of $T \otimes T$, i.e.

$$(\tilde{T}_k)_{i_1, \dots, i_k}^{j_1, \dots, j_k} = T_{i_1, \dots, i_k, i_{k+1}, \dots, i_p} T^{j_1, \dots, j_k, i_{k+1}, \dots, i_p}.$$

Note that $F(w) = \mathbb{E}_w[\tilde{T}_k(w^{\otimes 2k})]$ and $\|\tilde{T}_k\|_F \leq \|T\|_F^2$. Then by Theorem 38,

$$\mathbb{E}_w[F(w)^2] \lesssim \sum_{l \leq k} \left\| \text{Sym}(\tilde{T}_k)(I^l) \right\|_F^2 \lesssim \sum_{l \leq \frac{k}{2}} \left\| T(I^l) \right\|_F^4 \leq \|T\|_F^4 r^{2\lfloor \frac{k}{2} \rfloor}.$$

Therefore by Theorem 30, with probability at least $1 - 2e^{-\ell}$, $F(w) \lesssim \|T\|_F^2 r^{\lfloor \frac{k}{2} \rfloor} \ell^k$ and taking square roots completes the proof. \blacksquare

Corollary 41 *For $k \leq p$,*

$$\mathbb{E} \|T(w^{\otimes k})\|_F^2 \leq \mathbb{E} \langle T, w^{\otimes p} \rangle^2.$$

Proof This follows immediately from Lemma 38 and $(k - 2l)! \binom{k}{2l}^2 \leq (p - 2l)! \binom{p}{2l}^2$. \blacksquare

Corollary 42 *Let $w \sim N(0, I_d)$. Then,*

$$\mathbb{E}[w^{\otimes 2k}] \succeq k! \Pi_{\text{Sym}^k(\mathbb{R}^d)}$$

where $\Pi_{\text{Sym}^k(\mathbb{R}^d)}$ denotes the projection onto symmetric k -tensors.

Proof Considering only the $l = 0$ term in the above expansion of $\mathbb{E}[w^{\otimes 2k}](T, T)$ gives

$$\mathbb{E}[w^{\otimes 2k}](T, T) \geq k! \|T\|_F^2.$$

\blacksquare

Lemma 43 (Theorem 4.3 in (Prato and Tubaro, 2007)) *Let f be a polynomial of degree p . Then*

$$\mathbb{E}_{w \sim N(0, I_d)}[f(w)^k] \leq O_{k,p}(1) (\mathbb{E}_{w \sim N(0, I_d)}[f(w)^2])^{k/2}.$$

E.2. Sphere Lemmas

Lemma 44 *Let $\nu \sim \chi(d)$. Then,*

$$\mathbb{E}[\nu^{2k}] = \prod_{j=0}^{k-1} (d+2j) = d(d+2) \cdots (d+2k-2) = \Theta(d^k).$$

Lemma 45 *Let $\bar{w} \sim S^{d-1}$. Then,*

$$\mathbb{E}[\bar{w}^{\otimes 2k}] = \frac{\mathbb{E}_{w \sim N(0, I_d)}[w^{2k}]}{\mathbb{E}_{\nu \sim \chi(d)}[\nu^{2k}]}.$$

Proof This follows from the decomposition $w = \nu \bar{w}$ with $\nu \sim \chi(d)$, $\bar{w} \sim S^{d-1}$ independent. \blacksquare

Corollary 46 *Let T be a symmetric p -tensor with $\dim(\text{span}(T)) = r$. With probability at least $1 - 2e^{-t}$,*

$$\|T(\bar{w}^{\otimes k})\|_F \lesssim \|T\|_F \sqrt{\frac{r^{\lfloor \frac{k}{2} \rfloor} t^k}{d^k}}.$$

Corollary 47 *Let $\bar{w} \sim S^{d-1}$. For $k \leq p$,*

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

E.3. Rademacher Complexity Bounds

Lemma 48 *Let $f = a^T \sigma(W^* x + b)$ be a two layer neural network. For fixed W, b , Let*

$$\mathcal{F} = \{f_{(a, W, b)} : \|a\|_2 \leq B_a\}.$$

Then,

$$\mathfrak{R}_n(\mathcal{F}) \leq \sqrt{\frac{B_a^2(\|W\|_F^2 + \|b\|^2)}{n}}.$$

Proof

$$\begin{aligned} \mathfrak{R}_n(\mathcal{F}) &= \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \left[\frac{1}{n} \sum_i \sigma_i (a^T \sigma(W x_i + b)) \right] \right] \\ &= \frac{B_a}{n} \mathbb{E}_{x, \sigma} \left[\left\| \sum_i \sigma_i \sigma(W x_i + b) \right\|_2 \right] \\ &\leq \frac{B_a}{n} \sqrt{\mathbb{E}_{x, \sigma} \left[\left\| \sum_i \sigma_i \sigma(W x_i + b) \right\|_2^2 \right]} \\ &= \frac{B_a}{\sqrt{n}} \sqrt{\mathbb{E}_x \|\sigma(W x_1 + b)\|_2^2} \\ &\leq \sqrt{\frac{B_a^2(\|W\|_F^2 + \|b\|^2)}{n}}. \end{aligned}$$

\blacksquare

Lemma 49 Let $\theta = (a, W, b)$ and let $f = a^T \sigma(Wx + b)$ be a two layer neural network. Let

$$\mathcal{F} = \{f_\theta : \|a\|_2 \leq B_a, \|w_j\| \leq B_w\}.$$

Then,

$$\mathfrak{R}_n(\mathcal{F}) \leq 2B_a B_w \sqrt{\frac{md}{n}}.$$

Proof

$$\begin{aligned} \mathfrak{R}_n(\mathcal{F}) &= \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \left[\frac{1}{n} \sum_i \sigma_i (a^T \sigma(Wx_i + b)) \right] \right] \\ &= \frac{B_a}{n} \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \left\| \sum_i \sigma_i \sigma(Wx_i + b) \right\|_2 \right] \\ &\leq \frac{B_a \sqrt{m}}{n} \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \left\| \sum_i \sigma_i \sigma(Wx_i + b) \right\|_\infty \right] \\ &= \frac{B_a \sqrt{m}}{n} \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \left| \sum_i \sigma_i \sigma(w_j \cdot x_i + b_j) \right| \right] \\ &\leq \frac{2B_a \sqrt{m}}{n} \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \sum_i \sigma_i \sigma(w_j \cdot x_i + b_j) \right] \\ &\leq \frac{2B_a \sqrt{m}}{n} \mathbb{E}_{x, \sigma} \left[\sup_{f \in \mathcal{F}} \sum_i \sigma_i (w_j \cdot x_i) \right] \\ &\leq 2B_a B_w \sqrt{\frac{md}{n}}. \end{aligned}$$

■