
Towards Understanding Hierarchical Learning: Benefits of Neural Representations

Minshuo Chen* Yu Bai† Jason D. Lee◊ Tuo Zhao* Huan Wang†
Caiming Xiong† Richard Socher†
*Georgia Tech †Salesforce Research ◊Princeton University
{mchen393, tourzhao}@gatech.edu jasonlee@princeton.edu
{yu.bai, huan.wang, cxiong, rsocher}@salesforce.com

Abstract

Deep neural networks can empirically perform efficient hierarchical learning, in which the layers learn useful representations of the data. However, how they make use of the intermediate representations are not explained by recent theories that relate them to “shallow learners” such as kernels. In this work, we demonstrate that intermediate *neural representations* add more flexibility to neural networks and can be advantageous over raw inputs. We consider a fixed, randomly initialized neural network as a representation function fed into another trainable network. When the trainable network is the quadratic Taylor model of a wide two-layer network, we show that neural representation can achieve improved sample complexities compared with the raw input: For learning a low-rank degree- p polynomial ($p \geq 4$) in d dimension, neural representation requires only $\tilde{O}(d^{\lceil p/2 \rceil})$ samples, while the best-known sample complexity upper bound for the raw input is $\tilde{O}(d^{p-1})$. We contrast our result with a lower bound showing that neural representations do not improve over the raw input (in the infinite width limit), when the trainable network is instead a neural tangent kernel. Our results characterize when neural representations are beneficial, and may provide a new perspective on why depth is important in deep learning.

1 Introduction

Deep neural networks have been empirically observed to be more powerful than their shallow counterparts on a variety of machine learning tasks [38]. For example, on the ImageNet classification task, a 152-layer residual network can achieve 8%-10% better top-1 accuracy than a shallower 18-layer ResNet [30]. A widely held belief on why depth helps is that deep neural networks are able to perform efficient *hierarchical learning*, in which the layers learn representations that are increasingly useful for the present task. Such a hierarchical learning ability has been further leveraged in transfer learning. For example, [28] and [19] show that by combining with additional task-specific layers, the bottom layers of pre-trained neural networks for image classification and language modeling can be naturally transferred to other related tasks and achieve significantly improved performance.

Despite significant empirical evidence, we are in the lack of practical theory for understanding the hierarchical learning abilities of deep neural networks. Classical approximation theory has established a line of “depth separation” results which show that deep networks are able to approximate certain functions with much fewer parameters than shallow networks [18, 49, 24, 54, 12]. These work often manipulates the network parameters in potentially pathological ways, and it is unclear whether the resulting networks can be efficiently found through gradient-based optimization. A more recent line of work shows that overparametrized deep networks can be provably optimized and generalize as well as the so-called Neural Tangent Kernels (NTKs) [35, 21, 22, 3, 4, 7]. However, these results

do not take the hierarchical structure of the neural networks into account, and cannot justify any advantage of deep architectures. More recently, [33] show that some NTK models of deep networks are actually degenerate, and their generalization performance are no better than those associated with shallow networks.

In this paper, we provide a new perspective for understanding hierarchical learning through studying intermediate *neural representations*—that is, feeding fixed, randomly initialized neural networks as a representation function (feature map) into another trainable model. The prototypical model we consider is a wide two-layer neural network taking a representation function \mathbf{h} as the input, that is,

$$f_{\mathbf{W}}(\mathbf{x}) := \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \phi(\mathbf{w}_r^\top \mathbf{h}(\mathbf{x})), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the feature, $\mathbf{h} : \mathbb{R}^d \rightarrow \mathbb{R}^D$ is a data-independent representation function that is held fixed during learning, and $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times D}$ is the weight matrix to be learned from the data. For example, when $\mathbf{h}(\mathbf{x}) = \sigma(\mathbf{V}\mathbf{x} + \mathbf{b})$ is another one-hidden-layer network (i.e. neural representations), the model f is a three-layer network in which we only learn the weight matrix \mathbf{W} . Studying this model will reveal how the lower-level representation affects learning in a three-layer network, a previously missing yet important aspect of hierarchical learning.

To demonstrate the importance of the representation function \mathbf{h} , we investigate the *sample complexity* for learning certain target functions using model (1). This is a fine-grained measure of the power of \mathbf{h} compared with other notions such as approximation ability. Indeed, we expect $f_{\mathbf{W}}$ to be able to approximate any “regular” (e.g. Lipschitz) function of \mathbf{x} , whenever we use a non-degenerate \mathbf{h} and a sufficiently large width m . However, different choices of \mathbf{h} can result in different ways (for the trainable two-layer network) to approximate the same target function, thereby leading to different sample complexity guarantees. We will specifically focus on understanding when learning with the neural representation $\mathbf{h}(\mathbf{x}) = \sigma(\mathbf{V}\mathbf{x} + \mathbf{b})$ is more sample efficient than learning with the raw input $\mathbf{h}(\mathbf{x}) = \mathbf{x}$, which is a sensible baseline for capturing the benefits of representations.

As the optimization and generalization properties of a general two-layer network can be rather elusive, we consider more *optimization aware* versions of the prototype (1)—we replace the trainable two-layer network in $f_{\mathbf{W}}$ by tractable alternatives such as its **linearized model** [21] (also known as “lazy training” in [15]) or **quadratic Taylor model** [8]:

$$\text{(NTK-h)} \quad f_{\mathbf{W}}^L(\mathbf{x}) = \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \phi'(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x})),$$

$$\text{(Quad-h)} \quad f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x}))^2.$$

When \mathbf{h} is the raw input (NTK-Raw, Quad-Raw), these are models with concrete convergence and generalization guarantees, and can approximate the training of the full two-layer network in appropriate infinite-width limits (e.g. [21, 7, 4, 39, 8]). However, for learning with other representation functions, these models are less understood. The goal of this paper is to provide a quantitative understanding of these models, in particular when \mathbf{h} is a one-hidden-layer neural network (NTK-Neural, Quad-Neural), in terms of their convergence, generalization, and sample complexities of learning.

The contributions of this paper are summarized as follows:

- We show that the Quad-h model has a benign optimization landscape, and prove generalization error bounds with a precise dependence on the norm of the features and weight matrices, as well as the conditioning of the empirical covariance matrix of the features (Section 3).
- We study sample complexities of learning when the representation is chosen as a one-hidden-layer neural network (Quad-Neural model, Section 4). For achieving a small excess risk against a low-rank degree- p polynomial, we show that the Quad-Neural model requires $\tilde{O}(d^{\lceil p/2 \rceil})$ samples. When p is large, this is significantly better than the best known $\tilde{O}(d^{p-1})$ upper bound for the Quad-Raw model, demonstrating the benefits of neural representations.
- When the trainable network is instead a linearized model (or an NTK), we present a lower bound showing that neural representations are provably *not* beneficial: in a certain infinite-width limit, the NTK-Neural model requires at least $\Omega(d^p)$ samples for learning a degree- p polynomial (Section 5).

Since $O(d^p)$ samples also suffice for learning with the NTK-Raw model, this shows that neural representations are not beneficial when fed into a linearized neural network.

Additional paper organization We present the problem setup and algorithms in Section 2, review related work in Section 6, and provide conclusions as well as a broader impact statement in Section 7.

Notations We use bold lower-case letters to denote vectors, e.g., $\mathbf{x} \in \mathbb{R}^d$, and bold upper-case letters to denote matrices, e.g., $\mathbf{W} \in \mathbb{R}^{d_1 \times d_2}$. Given a matrix $\mathbf{W} \in \mathbb{R}^{d_1 \times d_2}$, we let $\|\mathbf{W}\|_{\text{op}}$ denote its operator norm, and $\|\mathbf{W}\|_{2,4}$ denote its (2, 4)-norm defined as $\|\mathbf{W}\|_{2,4}^4 = \sum_{i=1}^{d_1} \|\mathbf{W}_{i,:}\|_2^4$, where $\mathbf{W}_{i,:} \in \mathbb{R}^{d_2}$ is the i -th row of \mathbf{W} . Given a function $f(\mathbf{x})$ defined on domain \mathcal{X} with a probability measure \mathcal{D} , the L_2 norm is defined as $\|f\|_{L_2}^2 = \int_{\mathcal{X}} f^2(\mathbf{x}) \mathcal{D}(d\mathbf{x})$.

2 Preliminaries

Problem setup We consider the standard supervised learning task, in which we receive n i.i.d. training samples $S_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ from some data distribution \mathcal{D} , where $\mathbf{x} \in \mathcal{X}$ is the input and $y \in \mathcal{Y}$ is the label. In this paper, we assume that $\mathcal{X} = \mathbb{S}^{d-1} \subset \mathbb{R}^d$ (the unit sphere) so that inputs have unit norm $\|\mathbf{x}\|_2 = 1$. Our goal is to find a predictor $f : \mathcal{X} \mapsto \mathbb{R}$ such that the population risk

$$\mathcal{R}(f) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\ell(f(\mathbf{x}), y)]$$

is low, where $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a loss function. We assume that $\ell(\cdot, y)$ is convex, twice differentiable with the first and second derivatives bounded by 1, and satisfies $|\ell(0, y)| \leq 1$ for any $y \in \mathcal{Y}$. These assumptions are standard and are satisfied by commonly used loss functions such as the logistic loss and soft hinge loss.

Given dataset S_n , we define the empirical risk of a predictor f as

$$\widehat{\mathcal{R}}(f) := \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{x}_i), y_i).$$

Model, regularization, and representation We consider the case where f is either the linearized or the quadratic Taylor model of a wide two-layer network that takes a fixed representation function as the input:

$$\text{(NTK-h)} \quad f_{\mathbf{W}}^L(\mathbf{x}) = \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \phi'(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x})),$$

$$\text{(Quad-h)} \quad f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x}))^2,$$

where $\mathbf{h} : \mathbb{R}^d \rightarrow \mathbb{R}^D$ is a fixed representation function, $\mathbf{w}_{0,r} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$ and $a_r \stackrel{\text{iid}}{\sim} \text{Unif}(\{\pm 1\})$ are randomly initialized and held fixed during the training, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times D}$ is the trainable weight matrix¹, and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear activation. These models are taken as proxies for a full two-layer network of the form $\frac{1}{\sqrt{m}} \mathbf{a}^\top \phi((\mathbf{W}_0 + \mathbf{W})\mathbf{h}(\mathbf{x}))$, so as to enable better understandings of their optimization.

For the Quad-h model, we add a regularizer to the risk so as to encourage \mathbf{W} to have low norm. We use the regularizer $\|\mathbf{W}\|_{2,4}^4 = \sum_{r=1}^m \|\mathbf{w}_r\|_2^4$, and consider minimizing the regularized empirical risk

$$\widehat{\mathcal{R}}_\lambda(f_{\mathbf{W}}^Q) := \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) + \lambda \|\mathbf{W}\|_{2,4}^4 = \frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) + \lambda \|\mathbf{W}\|_{2,4}^4. \quad (2)$$

In the majority of this paper, we will focus on the case where $\mathbf{h}(\mathbf{x})$ is a fixed, randomly initialized neural network with one hidden layer of the form $\sigma(\mathbf{V}\mathbf{x} + \mathbf{b})$, with certain pre-processing steps when necessary. However, before we make the concrete choices, we think of \mathbf{h} as a general function that maps the raw input space \mathbb{R}^d into a feature space \mathbb{R}^D without any additional assumptions.

¹Our parameterization decouples the weight matrix in a standard two-layer network into two parts: the initialization $\mathbf{W}_0 \in \mathbb{R}^{m \times D}$ that is held fixed during training, and the ‘‘weight movement matrix’’ $\mathbf{W} \in \mathbb{R}^{m \times D}$ that can be thought of as initialized at $\mathbf{0}$.

Connection to a three-layer model It is worth noticing that when \mathbf{h} is indeed a neural network, say $\mathbf{h}(\mathbf{x}) = \sigma(\mathbf{V}\mathbf{x})$ (omitting bias for simplicity), our NTK- \mathbf{h} and Quad- \mathbf{h} models are closely related to the Taylor expansion of a *three-layer network*

$$\tilde{f}_{\mathbf{W}, \mathbf{V}}(\mathbf{x}) = \frac{1}{\sqrt{m}} \mathbf{a}^\top \phi((\mathbf{W}_0 + \mathbf{W})\sigma(\mathbf{V}\mathbf{x})).$$

Indeed, the {NTK- \mathbf{h} , Quad- \mathbf{h} } models correspond to the {linear, quadratic} Taylor expansion of the above network over \mathbf{W} , and is thus a part of the full Taylor expansion of the three-layer network. By studying these Taylor models, we gain understandings about how deep networks use its intermediate representation functions, which is lacking in existing work on Taylorized models.

3 Quadratic model with representations

We begin by studying the (non-convex) optimization landscape as well as the generalization properties of the model (Quad- \mathbf{h}), providing insights on what can be a good representation \mathbf{h} for such a model.

Base case of $\mathbf{h}(\mathbf{x}) = \mathbf{x}$: a brief review When $\mathbf{h}(\mathbf{x}) = \mathbf{x}$ is the raw input, model (Quad- \mathbf{h}) becomes

$$\text{(Quad-Raw)} \quad f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{x})(\mathbf{w}_r^\top \mathbf{x})^2,$$

which is the quadratic Taylor model of a wide two-layer neural network. This model is analyzed by Bai and Lee [8] who show that (1) the (regularized) risk $\widehat{\mathcal{R}}_\lambda(f_{\mathbf{W}})$ enjoys a nice optimization landscape despite being non-convex, and (2) the generalization gap of the model $f_{\mathbf{W}}^Q$ is controlled by $\|\mathbf{W}\|_{2,4}$ as well as $\|\frac{1}{n} \sum_{i \in [n]} \mathbf{x}_i \mathbf{x}_i^\top\|_{\text{op}}$. Building on these results, [8] show that learning low-rank polynomials with (Quad-Raw) achieves a better sample complexity than with the NTK. Besides the theoretical investigation, [9] empirically show that (Quad-Raw) model also approximates the training trajectories of standard neural networks better than the linearized model.

General case We analyze optimization landscape and establish generalization guarantees when \mathbf{h} is a general representation function, extending the results in [8]. We make the following assumption:

Assumption 1 (Bounded representation and activation). There exists a constant B_h such that $\|\mathbf{h}(\mathbf{x})\|_2 \leq B_h$ almost surely for $(\mathbf{x}, y) \sim \mathcal{D}$. The activation ϕ'' is uniformly bounded: $\sup_{t \in \mathbb{R}} |\phi''(t)| \leq C$ for some absolute constant C .

Theorem 1 (Optimization landscape and generalization of Quad- \mathbf{h}). Suppose Assumption 1 holds.

(1) (Optimization) Given any $\epsilon > 0$, $\tau = \Theta(1)$, and some radius $B_{w,*} > 0$, suppose the width $m \geq \tilde{O}(B_h^4 B_{w,*}^4 \epsilon^{-1})$ and we choose a proper regularization coefficient $\lambda > 0$. Then any second-order stationary point ² (SOSP) $\widehat{\mathbf{W}}$ of the regularized risk $\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q)$ satisfies $\|\widehat{\mathbf{W}}\|_{2,4} \leq O(B_{w,*})$, and achieves

$$\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau) \min_{\|\mathbf{W}\|_{2,4} \leq B_{w,*}} \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) + \epsilon.$$

(2) (Generalization) For any radius $B_w > 0$, we have with high probability (over $(\mathbf{a}, \mathbf{W}_0)$) that

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) \right| \right] \leq \tilde{O} \left(\frac{B_h^2 B_w^2 M_{h,\text{op}}}{\sqrt{n}} + \frac{1}{\sqrt{n}} \right),$$

where $M_{h,\text{op}}^2 = B_h^{-2} \mathbb{E}_{\mathbf{x}} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right]$.

² \mathbf{W} is a second-order stationary point (SOSP) of a twice-differentiable loss $L(\mathbf{W})$ if $\nabla L(\mathbf{W}) = \mathbf{0}$ and $\nabla^2 L(\mathbf{W}) \succeq \mathbf{0}$.

Efficient optimization; role of feature isotropicity Theorem 1 has two main implications: (1) With a sufficiently large width, any SOSP of the regularized risk $\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)$ achieves risk close to the optimum in a certain norm ball, and has controlled norm itself. Therefore, escaping-saddle type algorithms such as noisy SGD [36, 40] that can efficiently find SOSPs can also efficiently find these near global minima. (2) The generalization gap is controlled by $M_{h,\text{op}}$, which involves the operator norm of $\frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top$. It is thus beneficial if our representation $\mathbf{h}(\mathbf{x})$ is (approximately) *isotropic*, so that $M_{h,\text{op}} \asymp O(1/\sqrt{D})$, which is much lower than its naive upper bound 1. This will be a key insight for designing our neural representations in Section 4. The proof of Theorem 1 can be found in Appendix A.

4 Learning with neural representations

We now develop theories for learning with neural representations, where we choose \mathbf{h} to be a wide one-hidden-layer neural network.

4.1 Neural representations

We consider a fixed, randomly initialized one-hidden-layer neural network:

$$\mathbf{g}(\mathbf{x}) = \sigma(\mathbf{V}\mathbf{x} + \mathbf{b}) = [\sigma(\mathbf{v}_1^\top \mathbf{x} + b_1), \dots, \sigma(\mathbf{v}_D^\top \mathbf{x} + b_D)]^\top \in \mathbb{R}^D, \quad (3)$$

where $\mathbf{v}_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_d)$ and $b_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ are the weights. Throughout this section we will use the indicator activation $\sigma(t) = \mathbb{1}\{t \geq 0\}$. We will also choose $\phi(t) = \text{relu}(t)^2/2$ so that $\phi''(t) = \mathbb{1}\{t \geq 0\}$ as well.³

We define the representation function $\mathbf{h}(\mathbf{x})$ as the *whitened* version of $\mathbf{g}(\mathbf{x})$:

$$\mathbf{h}(\mathbf{x}) = \widehat{\Sigma}^{-1/2} \mathbf{g}(\mathbf{x}), \quad \text{where} \quad \widehat{\Sigma} = \frac{1}{n_0} \sum_{i=1}^{n_0} \mathbf{g}(\tilde{\mathbf{x}}_i) \mathbf{g}(\tilde{\mathbf{x}}_i)^\top. \quad (4)$$

Above, $\widehat{\Sigma}$ is an estimator of the population covariance matrix⁴ $\Sigma = \mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x}) \mathbf{g}(\mathbf{x})^\top] \in \mathbb{R}^{D \times D}$, and $\{\tilde{\mathbf{x}}_i\}_{i \in [n_0]} =: \tilde{S}_{n_0}$ is an additional set of unlabeled training examples of size n_0 (or a split from the existing training data). Such a whitening step makes $\mathbf{h}(\mathbf{x})$ more isotropic than the original features $\mathbf{g}(\mathbf{x})$, which according to Theorem 1 item (2) reduces the sample complexity for achieving low test error. We will discuss this more in Section 4.2.

We summarize our overall learning algorithm (with the neural representation) in Algorithm 1.

Algorithm 1 Learning with Neural Representations (Quad-Neural method)

Input: Labeled data S_n , unlabeled data \tilde{S}_{n_0} , initializations $\mathbf{V} \in \mathbb{R}^{D \times d}$, $\mathbf{b} \in \mathbb{R}^D$, $\mathbf{W}_0 \in \mathbb{R}^{m \times D}$, parameters (λ, ϵ) .

Step 1: Construct model $f_{\mathbf{W}}^Q$ as

$$\text{(Quad-Neural)} \quad f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x}))^2,$$

where $\mathbf{h}(\mathbf{x}) = \widehat{\Sigma}^{-1/2} \mathbf{g}(\mathbf{x})$ is the neural representation (4) (using \tilde{S}_{n_0} to estimate the covariance).

Step 2: Find a second-order stationary point $\widehat{\mathbf{W}}$ of the regularized empirical risk (on the data S_n):

$$\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q) = \frac{1}{n} \sum_{i=1}^n \ell(f_{\widehat{\mathbf{W}}}^Q(\mathbf{x}_i), y_i) + \lambda \|\widehat{\mathbf{W}}\|_{2,4}^4.$$

³We can use a non-smooth σ since (\mathbf{V}, \mathbf{b}) are not trained. Our results can be extended to the situation where σ or ϕ'' is the relu activation as well.

⁴Strictly speaking, Σ is the second moment matrix of $\mathbf{g}(\mathbf{x})$.

4.2 Learning low-rank polynomials with neural representations

We now study the sample complexity of Algorithm 1 to achieve low excess test risk compared with the best *low-rank degree- p polynomial*, that is, sum of polynomials of the form $(\boldsymbol{\beta}^\top \mathbf{x})^p$. This setting has been considered in a variety of prior work on learning polynomials [47, 13] as well as analyses of wide neural networks [7, 8].

We need the following additional assumption on the random features.

Assumption 2 (Lower Bounded Covariance). For any k and $D \leq O(d^k)$, with high probability over \mathbf{V}, \mathbf{b} (as $d \rightarrow \infty$), we have the minimum eigenvalue $\lambda_{\min}(\boldsymbol{\Sigma}) \geq \lambda_k$ for some constant $\lambda_k > 0$ that only depends on k but not d , where $\boldsymbol{\Sigma} = \mathbb{E}_{\mathbf{x}}[\sigma(\mathbf{V}\mathbf{x} + \mathbf{b})\sigma(\mathbf{V}\mathbf{x} + \mathbf{b})^\top]$.

Assumption 2 states the features $\{\sigma(\mathbf{v}_i^\top \mathbf{x} + \mathbf{b}_i)\}$ to be not too correlated, which roughly requires the distribution of \mathbf{x} to span all directions in \mathbb{R}^d . For example, when $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$ (and with our choice of $\sigma(t) = \mathbb{1}\{t \geq 0\}$), we show that this assumption is satisfied with

$$\lambda_k = \Theta\left(\min_{\deg(q) \leq k} \mathbb{E}_{z \sim \mathcal{N}(0,1)}[(\sigma(z) - q(z))^2]\right) \asymp k^{-1/2},$$

where $q(z)$ denotes a polynomial in z and its degree is denoted as $\deg(q)$. For general distributions of \mathbf{x} , we show Assumption 2 still holds under certain moment conditions on the distribution of \mathbf{x} (see the formal statement and proof of both results in Appendix B).

Sample complexity for learning polynomials We focus on low-rank polynomials of the form

$$f_\star(\mathbf{x}) = \sum_{s=1}^{r_\star} \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}, \quad \text{where } |\alpha_s| \leq 1, \|\boldsymbol{\beta}_s^\top \mathbf{x}\|_{L_2} \leq 1, p_s \leq p \text{ for all } s. \quad (5)$$

We state our main result for the Quad-Neural model to achieve low excess risk over such functions.

Theorem 2 (Sample complexity of learning with Quad-Neural). Suppose Assumption 2 holds, and there exists some f_\star of the form (5) that achieves low risk: $\mathcal{R}(f_\star) \leq \text{OPT}$. Then for any $\epsilon, \delta \in (0, 1)$ and $\tau = \Theta(1)$, choosing

$$D = \Theta\left(\text{poly}(r_\star, p) \sum_s \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil} \epsilon^{-2} \delta^{-1}\right), \quad m \geq \tilde{O}(\text{poly}(r_\star, D) \epsilon^{-2} \delta^{-1}), \quad (6)$$

$n_0 = \tilde{O}(D\delta^{-2})$, and a proper $\lambda > 0$, Algorithm 1 achieves the following guarantee: with probability at least $1 - \delta$ over the randomness of data and initialization, any second-order stationary point $\widehat{\mathbf{W}}$ of $\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q)$ satisfies

$$\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau)\text{OPT} + \underbrace{\epsilon}_{\substack{\text{approx.,} \\ \text{requires large } D}} + \underbrace{\tilde{O}\left(\sqrt{\frac{\text{poly}(r_\star, p, \delta^{-1}) \lambda_{\lceil p/2 \rceil}^{-1} \epsilon^{-2} \sum_{s=1}^{r_\star} \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil}}{n}}\right)}_{\text{generalization, requires large } n \text{ (given } \epsilon)}.$$

In particular, for any $\epsilon > 0$, we can achieve $\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau)\text{OPT} + 2\epsilon$ with sample complexity

$$n_0 + n \leq \tilde{O}\left(\text{poly}(r_\star, p, \lambda_{\lceil p/2 \rceil}^{-1}, \epsilon^{-1}, \delta^{-1}) \sum_{s=1}^{r_\star} \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil}\right). \quad (7)$$

According to Theorem 2, Quad-Neural can learn polynomials of any degree by doing the following: (1) Choose a sufficiently large D , so that the neural representations are expressive enough; (2) Choose a large width m in the quadratic model so as to enable a nice optimization landscape, where such m only appears logarithmically in generalization error (Theorem 1).

Improved dimension dependence over Quad-Raw and NTK-Raw We parse the sample complexity bound in Theorem 2 in the following important case: \mathbf{x} is relatively uniform (e.g. $\text{Unif}(\mathbb{S}^{d-1})$), $\|\beta_s\|_2 = O(\sqrt{d})$, and the data is noiseless and realized by f_* (so that $\text{OPT} = 0$). In this case we have $\|(\beta_s^\top \mathbf{x})^{p_s}\|_{L_2} = O(1)$, and Assumption 2 holds with $\lambda_{\lceil p/2 \rceil} \geq p^{-O(1)}$. Thus, when only highlighting the d dependence⁵, the sample complexity required to achieve ϵ test risk with the Quad-Neural is (reading from (18))

$$N_{\text{quad-neural}} = \tilde{O}\left(d^{\lceil p/2 \rceil}\right).$$

In comparison, the sample complexity for learning with the Quad-Raw (quadratic neural network with the raw input) is

$$N_{\text{quad-raw}} = \tilde{O}\left(d^{p-1}\right)$$

(see, e.g. [8, Thm 7]). Therefore, Theorem 2 shows that neural representations can significantly improve the sample complexity over the raw input, when fed into a quadratic Taylor model.

Overview of techniques At a high level, the improved sample complexity achieved in Theorem 2 is due to the flexibility of the neural representation: the Quad-h model can express polynomials *hierarchically*, using weight matrices with much smaller norms than that of a shallow learner such as the Quad-Raw model. This lower norm in turn translates to a better generalization bound (according to Theorem 1) and an improved sample complexity. We sketch the main arguments here, and leave the complete proof to Appendix D.

(1) **Expressing functions using hierarchical structure:** We prove the existence of some $\mathbf{W}^* \in \mathbb{R}^{m \times D}$ such that $f_{\mathbf{W}^*}^Q \approx f_*$ by showing the following: (1) As soon as $D \geq \tilde{O}(d^k)$, $\mathbf{h}(\mathbf{x})$ can linearly express certain degree- k polynomials as “bases”; (2) For large m , the top quadratic Taylor model can further express degree $\lceil p_s/k \rceil$ polynomials of the bases, thereby expressing f_* . This is an explicit way of utilizing the hierarchical structure of the model. We note that our proof used $k = \lceil p/2 \rceil$, but the argument can be generalized to other k as well.

(2) **Making representations isotropic:** We used a whitened version of a one-hidden-layer network as our representation function \mathbf{h} (cf. (4)). The whitening operation does not affect the expressivity argument in part (1) above, but helps improve the conditioning of the feature covariance matrix (cf. the quantity $M_{h,\text{op}}$ in Theorem 1). Applying whitening, we obtain nearly isotropic features: $\mathbb{E}_{\mathbf{x}}[\mathbf{h}(\mathbf{x})\mathbf{h}(\mathbf{x})^\top] \approx \mathbf{I}_D$, which is key to the sample complexity gain over the Quad-Raw model as discussed above. We note that well-trained deep networks with BatchNorm may have been implicitly performing such whitening operations in practice [46]. We also remark that the whitening step in Algorithm 1 may be replaced with using unwhitened representations with a *data-dependent* regularizer, e.g., $\sum_{r=1}^m \|\tilde{\Sigma}^{1/2} \mathbf{w}_r\|_2^4$, which achieves similar sample complexity guarantees (see Appendix D).

5 NTK with neural representations: a lower bound

In this section, we show that neural representations may not be beneficial over raw inputs when the trainable network is a *linearized* neural network through presenting a sample complexity lower bound for this method in the infinite width limit.

More concretely, we consider NTK-Neural, which learns a model $f_{\mathbf{W}}^L$ of the form

$$\text{(NTK-Neural)} \quad f_{\mathbf{W}}^L(\mathbf{x}) := \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \phi' \left(\mathbf{w}_{0,r}^\top \mathbf{g}(\mathbf{x}) / \sqrt{D} \right) \left(\mathbf{w}_r^\top \mathbf{g}(\mathbf{x}) / \sqrt{D} \right),$$

where $\mathbf{g}(\mathbf{x}) := [\sigma(\mathbf{v}_1^\top \mathbf{x} + b_1), \dots, \sigma(\mathbf{v}_D^\top \mathbf{x} + b_D)]^\top \in \mathbb{R}^D$ are the neural random features (same as in (3)), and the $1/\sqrt{D}$ factor rescales $\mathbf{g}(\mathbf{x})$ to $O(1)$ norm on average.

⁵For example, in the high-dimensional setting when $\epsilon = \Theta(1)$ and d is large [27].

Infinite-width limit: a kernel predictor Model (NTK-Neural) is linear model with parameter \mathbf{W} , and can be viewed as a kernel predictor with a (finite-dimensional kernel) $H_{m,D} : \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$. In the infinite-width limit of $D, m \rightarrow \infty$, we have $H_{m,D} \rightarrow H_\infty$, where

$$H_\infty(\mathbf{x}, \mathbf{x}') := \mathbb{E}_{(u,v) \sim \mathcal{N}(\mathbf{0}, \Sigma(\mathbf{x}, \mathbf{x}'))} [\phi'(u)\phi'(v)] \cdot \Sigma_{12}(\mathbf{x}, \mathbf{x}'), \text{ and}$$

$$\Sigma(\mathbf{x}, \mathbf{x}') = \begin{pmatrix} \mathbb{E}_{\mathbf{v}, b} [\sigma(\mathbf{v}^\top \mathbf{x} + b)^2] & \mathbb{E}_{\mathbf{v}, b} [\sigma(\mathbf{v}^\top \mathbf{x} + b)\sigma(\mathbf{v}^\top \mathbf{x}' + b)] \\ \mathbb{E}_{\mathbf{v}, b} [\sigma(\mathbf{v}^\top \mathbf{x} + b)\sigma(\mathbf{v}^\top \mathbf{x}' + b)] & \mathbb{E}_{\mathbf{v}, b} [\sigma(\mathbf{v}^\top \mathbf{x}' + b)^2] \end{pmatrix},$$

(see e.g. [35, 20] for the derivation). Motivated by this, we consider kernel predictors of the form

$$\hat{f}_\lambda = \operatorname{argmin}_f \sum_{i=1}^n \ell(f(\mathbf{x}_i), y_i) + \lambda \|f\|_{H_\infty}^2 \quad (8)$$

as a proxy for (NTK-Neural), where $\|\cdot\|_{H_\infty}^2$ denotes the RKHS (Reproducing Kernel Hilbert Space) norm associated with kernel H_∞ . This set of predictors is a reliable proxy for the (NTK-Neural) method: for example, taking $\lambda \rightarrow 0_+$, it recovers the solution found by gradient descent (with a small stepsize) on the top layer of a wide three-layer network [20].

We now present a lower bound for the predictor \hat{f}_λ , adapted from [27, Theorem 3].

Theorem 3 (Lower bound for NTK-Neural). Suppose the input distribution is $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$, and $y_* = f_*(\mathbf{x})$ where $f_* \in L_2(\text{Unif}(\mathbb{S}^{d-1}))$ consists of polynomials of degree at least p^6 . Assume the sample size $n \leq O(d^{p-\delta})$ for some $\delta > 0$. Then for any fixed $\epsilon \in (0, 1)$, as $d \rightarrow \infty$, the predictor \hat{f}_λ defined in (8) suffers from the following lower bound with high probability (over $\{(\mathbf{x}_i, y_i)\}$):

$$\mathbb{E}_{\mathbf{x}} \left[\inf_{\lambda > 0} (\hat{f}_\lambda(\mathbf{x}) - f_*(\mathbf{x}))^2 \right] \geq (1 - \epsilon) \mathbb{E}_{\mathbf{x}} [f_*(\mathbf{x})]^2,$$

that is, any predictor of the form (8) will not perform much better than the trivial zero predictor.

No improvement over NTK-Raw; benefits of neural representations Theorem 3 shows that the infinite width version (8) of the NTK-Neural method requires roughly at least $\Omega(d^p)$ samples in order to learn any degree- p polynomial up to a non-trivial accuracy (in squared error). Crucially, this lower bound implies that NTK-Neural *does not* improve over NTK-Raw (i.e. NTK with the raw input) in the infinite width limit—the infinite width NTK-Raw already achieves sample complexity *upper bound* of $O(d^p)$ for learning a degree- p polynomial $y = f_*(\mathbf{x})$ when $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$ [27]. This is in stark contrast with our Theorem 2 which shows that Quad-Neural improves over Quad-Raw, suggesting that neural representations are perhaps only beneficial when fed into a sufficiently complex model.

6 Related work

Approximation theory and depth separation. Extensive efforts have been made on the expressivity of neural networks and the benefits of increased depth. Two separate focuses were pursued: 1) Universal approximation theory for approximating dense function classes, e.g., Sobolev and squared integrable functions [17, 31, 10, 34, 25, 16, 41, 44, 43]; 2) depth separation theory demonstrating the benefits of increased depth on expressing certain structured functions, e.g., saw-tooth functions [29, 18, 48, 49, 24]. More recently, the recent work [54] merged the two focuses by studying unbounded-depth ReLU networks for approximating Sobolev functions. In all these work, the network parameters are constructed in potentially weird ways, and it is unclear whether such networks can be efficiently found using gradient-based optimization.

Neural tangent kernels and beyond A growing body of recent work show the connection between gradient descent on the full network and the Neural Tangent Kernel (NTK) [35], from which one can prove concrete results about neural network training [42, 21, 20, 3, 56] and generalization [7, 4, 11]. Despite such connections, these results only show that neural networks are as powerful as shallow learners such as kernels. The gap between such shallow learners and the full neural network has been established in theory by [52, 1, 55, 26, 53, 23] and observed in practice [6, 39, 14]. Higher-order expansions of the {network, training dynamics} such as Taylorized Training [8, 9] and the Neural

⁶That is, $\|\mathbb{P}_{<p} f_*\|_{L_2} = 0$, where $\mathbb{P}_{<p}$ denotes the L_2 projection onto the space of degree $< p$ polynomials.

Tangent Hierarchy [32] have been recently proposed towards closing this gap. Finally, recent work by Allen-Zhu and Li [2] shows that there exists a class of polynomials that can be efficiently learned by a deep network but not any “non-hierarchical” learners such as kernel methods or neural tangent kernels, thereby shedding light on how representations are learned hierarchically.

Learning low-rank polynomials in high dimension In [47] and [45], the authors propose a tensor unfolding algorithm to estimate a rank k order p tensor with $(d)^{p/2}k$ samples. Under Gaussian input data, [13] propose a Grassmanian manifold optimization algorithm with spectral initialization to estimate a polynomial over k -dimensional subspace of variables of degree p with $O_{k,p}(d \log^d p)$ samples, where $O_{k,p}$ suppresses unknown (super)-exponential dependence on k and p . However, these methods explicitly use knowledge about the data distribution. Neural networks can often learn polynomials in distribution-free ways. [5, 7] show that wide two-layer networks that simulate an NTK require $\tilde{O}(d^p)$ samples to learn a degree- p polynomial. [27] show that $\Omega(d^p)$ samples is also asymptotically necessary for any rotationally invariant kernel. [8] show that a randomized wide two-layer network requires $\tilde{O}(d^{p-1})$ samples instead by coupling it with the quadratic Taylor model. Our algorithm belongs to this class of distribution-free methods, but achieve an improved sample complexity when the distribution satisfies a mild condition.

7 Conclusion

This paper provides theoretical results on the benefits of neural representations in deep learning. We show that using a neural network as a representation function can achieve improved sample complexity over the raw input in a neural quadratic model, and also show such a gain is not present if the model is instead linearized. We believe these results provide new understandings to hierarchical learning in deep neural networks. For future work, it would be of interest to study whether deeper representation functions are even more beneficial than shallower ones, or what happens when the representation is fine-tuned together with the trainable network.

Acknowledgment

We thank the anonymous reviewers for the suggestions. We thank Song Mei for the discussions about the concentration of long-tailed covariance matrices. JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0303, the Sloan Research Fellowship, and NSF CCF 2002272.

Broader impact

This paper extensively contributes to the theoretical frontier of deep learning. We do not foresee direct ethical or societal consequences. Instead, our theoretical finding reduces the gap between the theory and practice, and is in sharp contrast to existing theories, which cannot show any advantage of deep networks over the shallow ones. In viewing of a notably increasing trend towards establishing a quantitative framework using deep neural networks in diverse areas, e.g., computational social science, this paper will provide an important theoretical guideline for practitioners.

References

- [1] Z. Allen-Zhu and Y. Li. What can resnet learn efficiently, going beyond kernels? *arXiv preprint arXiv:1905.10337*, 2019.
- [2] Z. Allen-Zhu and Y. Li. Backward feature correction: How deep learning performs deep learning. *arXiv preprint arXiv:2001.04413*, 2020.
- [3] Z. Allen-Zhu, Y. Li, and Z. Song. A convergence theory for deep learning via overparameterization. *arXiv preprint arXiv:1811.03962*, 2018.
- [4] Z. Allen-Zhu, Y. Li, and Y. Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6155–6166, 2019.

- [5] A. Andoni, R. Panigrahy, G. Valiant, and L. Zhang. Learning polynomials with neural networks. In *International conference on machine learning*, pages 1908–1916, 2014.
- [6] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8139–8148, 2019.
- [7] S. Arora, S. S. Du, W. Hu, Z. Li, and R. Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- [8] Y. Bai and J. D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rk11GyBFPH>.
- [9] Y. Bai, B. Krause, H. Wang, C. Xiong, and R. Socher. Taylorized training: Towards better approximation of neural network training at finite width. *arXiv preprint arXiv:2002.04010*, 2020.
- [10] A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- [11] Y. Cao and Q. Gu. Generalization error bounds of gradient descent for learning overparameterized deep relu networks. *arXiv preprint arXiv:1902.01384*, 2019.
- [12] M. Chen, H. Jiang, W. Liao, and T. Zhao. Efficient approximation of deep relu networks for functions on low dimensional manifolds. In *Advances in Neural Information Processing Systems*, pages 8172–8182, 2019.
- [13] S. Chen and R. Meka. Learning polynomials of few relevant dimensions. *arXiv preprint arXiv:2004.13748*, 2020.
- [14] L. Chizat and F. Bach. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 8, 2018.
- [15] L. Chizat, E. Oyallon, and F. Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.
- [16] C. K. Chui and X. Li. Approximation by ridge functions and neural networks with one hidden layer. *Journal of Approximation Theory*, 70(2):131–141, 1992.
- [17] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- [18] O. Delalleau and Y. Bengio. Shallow vs. deep sum-product networks. In *Advances in neural information processing systems*, pages 666–674, 2011.
- [19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [20] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685, 2019.
- [21] S. S. Du, J. D. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*, 2018.
- [22] S. S. Du, X. Zhai, B. Póczos, and A. Singh. Gradient descent provably optimizes overparameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- [23] E. Dyer and G. Gur-Ari. Asymptotics of wide networks from feynman diagrams. *arXiv preprint arXiv:1909.11304*, 2019.
- [24] R. Eldan and O. Shamir. The power of depth for feedforward neural networks. In *Conference on learning theory*, pages 907–940, 2016.

- [25] K.-I. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural networks*, 2(3):183–192, 1989.
- [26] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. Limitations of lazy training of two-layers neural network. In *Advances in Neural Information Processing Systems*, pages 9108–9118, 2019.
- [27] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. Linearized two-layers neural networks in high dimension. *arXiv preprint arXiv:1904.12191*, 2019.
- [28] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [29] J. Håstad. Computational limitations of small-depth circuits. 1987.
- [30] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [31] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [32] J. Huang and H.-T. Yau. Dynamics of deep neural networks and neural tangent hierarchy. *arXiv preprint arXiv:1909.08156*, 2019.
- [33] K. Huang, Y. Wang, M. Tao, and T. Zhao. Why do deep residual networks generalize better than deep feedforward networks?—a neural tangent kernel perspective. *arXiv preprint arXiv:2002.06262*, 2020.
- [34] B. Irie and S. Miyake. Capabilities of three-layered perceptrons. In *IEEE International Conference on Neural Networks*, volume 1, page 218, 1988.
- [35] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- [36] C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan. Stochastic gradient descent escapes saddle points efficiently. *arXiv preprint arXiv:1902.04811*, 2019.
- [37] A. T. Kalai, A. R. Klivans, Y. Mansour, and R. A. Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008.
- [38] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [39] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. 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.
- [40] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *Conference on learning theory*, pages 1246–1257, 2016.
- [41] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6): 861–867, 1993.
- [42] Y. Li and Y. Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.
- [43] Y. Makovoz. Random approximants and neural networks. *Journal of Approximation Theory*, 85(1):98–109, 1996.
- [44] H. N. Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural computation*, 8(1):164–177, 1996.

- [45] A. Montanari and N. Sun. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71(11):2381–2425, 2018.
- [46] A. S. Morcos, D. G. Barrett, N. C. Rabinowitz, and M. Botvinick. On the importance of single directions for generalization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1iuQjxCZ>.
- [47] C. Mu, B. Huang, J. Wright, and D. Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In *International conference on machine learning*, pages 73–81, 2014.
- [48] B. Rossman, R. A. Servedio, and L.-Y. Tan. An average-case depth hierarchy theorem for boolean circuits. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1030–1048. IEEE, 2015.
- [49] M. Telgarsky. Benefits of depth in neural networks. *arXiv preprint arXiv:1602.04485*, 2016.
- [50] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [51] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [52] C. Wei, J. D. Lee, Q. Liu, and T. Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9709–9721, 2019.
- [53] B. Woodworth, S. Gunasekar, J. D. Lee, E. Moroshko, P. Savarese, I. Golan, D. Soudry, and N. Srebro. Kernel and rich regimes in overparametrized models. *arXiv preprint arXiv:2002.09277*, 2020.
- [54] D. Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94: 103–114, 2017.
- [55] G. Yehudai and O. Shamir. On the power and limitations of random features for understanding neural networks. *arXiv preprint arXiv:1904.00687*, 2019.
- [56] D. Zou, Y. Cao, D. Zhou, and Q. Gu. Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*, 2018.

Supplementary Materials for “Towards Understanding Hierarchical Learning: Benefits of Neural Representations”

A Proofs for Section 3

A.1 Proof of Optimization in Theorem 1

We first derive the gradient and Hessian of empirical risk $\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)$, which will be used throughout the rest of the proof. For a better presentation, we denote $\langle \cdot, \cdot \rangle$ as inner product and

$$f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \langle \mathbf{x}\mathbf{x}^\top, \mathbf{W}\mathbf{D}(\mathbf{x})\mathbf{W}^\top \rangle \text{ for } \mathbf{D}_{rr}(\mathbf{x}) = a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x})).$$

We compute the gradient and Hessian of $\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)$ along a given direction \mathbf{W}_* .

$$\begin{aligned} \nabla_{\mathbf{W}} \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) &= \frac{2}{n} \sum_{i=1}^n \ell'(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) \frac{1}{2\sqrt{m}} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{W} \mathbf{D}(\mathbf{x}_i) \quad \text{and} \\ \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] &= \frac{2}{n} \sum_{i=1}^n \ell'(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) \cdot \underbrace{\frac{1}{2\sqrt{m}} \langle \mathbf{x}_i \mathbf{x}_i^\top, \mathbf{W}_* \mathbf{D}(\mathbf{x}_i) \mathbf{W}_*^\top \rangle}_{f_{\mathbf{W}_*}^Q(\mathbf{x}_i)} \\ &\quad + \frac{4}{n} \sum_{i=1}^n \ell''(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) \cdot \left(\underbrace{\frac{1}{2\sqrt{m}} \langle \mathbf{x}_i \mathbf{x}_i^\top, \mathbf{W} \mathbf{D}(\mathbf{x}_i) \mathbf{W}^\top \rangle}_{\tilde{y}_i} \right)^2 \\ &= \underbrace{\frac{2}{n} \sum_{i=1}^n \ell'(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) f_{\mathbf{W}_*}^Q(\mathbf{x}_i)}_{\text{I}} + \underbrace{\frac{4}{n} \sum_{i=1}^n \ell''(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) \tilde{y}_i^2}_{\text{II}}. \end{aligned}$$

We denote $\widehat{\mathcal{D}}$ as the empirical data distribution, and bound I and II separately.

$$\begin{aligned} \text{I} &= 2\mathbb{E}_{\widehat{\mathcal{D}}}[\ell'(f_{\mathbf{W}}^Q(\mathbf{x}), y) f_{\mathbf{W}_*}^Q(\mathbf{x})] \\ &= 2\mathbb{E}_{\widehat{\mathcal{D}}}[\ell'(f_{\mathbf{W}}^Q(\mathbf{x}), y) f_{\mathbf{W}}^Q(\mathbf{x})] + 2\mathbb{E}_{\widehat{\mathcal{D}}}[\ell'(f_{\mathbf{W}}^Q(\mathbf{x}), y) (f_{\mathbf{W}_*}^Q(\mathbf{x}) - f_{\mathbf{W}}^Q(\mathbf{x}))] \\ &\stackrel{(i)}{\leq} \langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \rangle + 2\mathbb{E}_{\widehat{\mathcal{D}}}[\ell(f_{\mathbf{W}_*}^Q(\mathbf{x}), y) - \ell(f_{\mathbf{W}}^Q(\mathbf{x}), y)] \\ &= \langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \rangle - 2(\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{W}_*}^Q)), \end{aligned}$$

where (i) follows directly by computing $\langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \rangle$ and the convexity of ℓ . For II, with $\ell'' \leq 1$, we have

$$\begin{aligned} \text{II} &\leq \frac{4}{n} \sum_{i=1}^n \sum_{r \leq m} \tilde{y}_i^2 = \mathbb{E}_{\widehat{\mathcal{D}}} \left[\frac{2}{m} \sum_{r \leq m} \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}))^2 (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x}))^2 (\mathbf{w}_{*,r}^\top \mathbf{h}(\mathbf{x}))^2 \right] \\ &\leq C^2 \mathbb{E}_{\widehat{\mathcal{D}}} \left[\frac{1}{m} \sum_{r \leq m} (\mathbf{w}_r^\top \mathbf{h}(\mathbf{x}))^2 (\mathbf{w}_{*,r}^\top \mathbf{h}(\mathbf{x}))^2 \right] \\ &\leq \frac{1}{m} C^2 B_h^4 \sum_{r \leq m} \|\mathbf{w}_r\|_2^2 \|\mathbf{w}_{*,r}\|_2^2 \\ &\leq m^{-1} C^2 B_h^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2, \end{aligned}$$

where the last step used Cauchy-Schwarz on $\{\|\mathbf{w}_r\|_2\}$ and $\{\|\mathbf{w}_{*,r}\|_2\}$, and the constant C is the uniform upper bound on ϕ'' . Putting terms I and II together, we have

$$\nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] \leq \left\langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{W}^*}^Q)) + m^{-1} C^2 B_h^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2. \quad (9)$$

Proof of Theorem 1, Optimization Part. We denote $\mathbf{W}^* = \operatorname{argmin}_{\|\mathbf{W}\|_{2,4} \leq B_{w,*}} \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)$ and let its risk $\widehat{\mathcal{R}}(f_{\mathbf{W}^*}^Q) = M$. We begin by choosing the regularization strength as

$$\lambda = \lambda_0 B_{w,*}^{-4},$$

where λ_0 is a constant to be determined.

We argue that any second order stationary point $\widehat{\mathbf{W}}$ has to satisfy $\|\widehat{\mathbf{W}}\|_{2,4} = O(B_{w,*})$. We have for any \mathbf{W} that

$$\begin{aligned} \left\langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle &= \mathbb{E}_{\mathcal{D}} \left[\ell'(f_{\mathbf{W}}^Q(\mathbf{x}), y) \cdot 2f_{\mathbf{W}}^Q(\mathbf{x}) \right] \\ &= 2\mathbb{E}_{\mathcal{D}} \left[\ell'(f_{\mathbf{W}}^Q(\mathbf{x}), y) \cdot (f_{\mathbf{W}}^Q(\mathbf{x}) - f_{\mathbf{0}}^Q(\mathbf{x})) \right] \\ &\stackrel{(i)}{\geq} 2(\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{0}}^Q)) \\ &\stackrel{(ii)}{\geq} -2, \end{aligned}$$

where (i) uses convexity of ℓ and (ii) uses the assumption that $\ell(0, y) \leq 1$ for all $y \in \mathcal{Y}$.

Combining with the fact that $\left\langle \nabla_{\mathbf{W}}(\|\mathbf{W}\|_{2,4}^4), \mathbf{W} \right\rangle = 4\|\mathbf{W}\|_{2,4}^4$, we have simultaneously for all \mathbf{W} that

$$\begin{aligned} \left\langle \nabla \widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle &\geq \left\langle \nabla_{\mathbf{W}}(\lambda \|\mathbf{W}\|_{2,4}^4), \mathbf{W} \right\rangle + \left\langle \nabla_{\mathbf{W}} \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle \\ &\geq 4\lambda \|\mathbf{W}\|_{2,4}^4 - 2. \end{aligned}$$

Therefore we see that any stationary point \mathbf{W} has to satisfy

$$\|\mathbf{W}\|_{2,4} \leq (2\lambda)^{-1/4}.$$

Choosing

$$\lambda_0 = \frac{1}{36}(2\tau M + \epsilon),$$

we get $36\lambda B_{w,*}^4 = 2\tau M + \epsilon$. The Hessian of $\widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q)$ along direction \mathbf{W}_* is

$$\begin{aligned} \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] &= \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] + \lambda \nabla_{\mathbf{W}}^2 \|\mathbf{W}\|_{2,4}^4 [\mathbf{W}_*, \mathbf{W}_*] \\ &= \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] + 4\lambda \sum_{r \leq m} \|\mathbf{w}_r\|_2^2 \|\mathbf{w}_{*,r}\|_2^2 + 2 \langle \mathbf{w}_r, \mathbf{w}_{*,r} \rangle^2 \\ &\leq \left\langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) - M) + m^{-1} C^2 B_h^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2 \\ &\quad + 12\lambda \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2 \\ &\stackrel{(i)}{\leq} \left\langle \nabla \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) - M) + m^{-1} C^2 B_h^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2 \\ &\quad + \lambda \|\mathbf{W}\|_{2,4}^4 + 36\lambda \|\mathbf{W}_*\|_{2,4}^4 \\ &\leq \left\langle \nabla \widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q) - M) + m^{-1} C^2 B_h^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2 \\ &\quad - \lambda \|\mathbf{W}\|_{2,4}^4 + 36\lambda \|\mathbf{W}_*\|_{2,4}^4. \end{aligned}$$

We used the fact $12ab \leq a^2 + 36b^2$. For a second order-stationary point $\widehat{\mathbf{W}}$ of $\widehat{\mathcal{R}}_{\lambda}(f_{\mathbf{W}}^Q)$, its gradient vanishes and the Hessian is positive definite. Therefore, we have

$$0 \leq -2(\widehat{\mathcal{R}}_{\lambda}(f_{\widehat{\mathbf{W}}}^Q) - M) + m^{-1} C^2 B_h^4 \|\widehat{\mathbf{W}}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2 - \lambda \|\widehat{\mathbf{W}}\|_{2,4}^4 + 36\lambda \|\mathbf{W}_*\|_{2,4}^4.$$

We choose $m = \epsilon^{-1}(2\lambda_0)^{-1/2}C^2B_h^4B_{w,*}^4 \geq \epsilon^{-1}C^2B_h^4\|\widehat{\mathbf{W}}\|_{2,4}^2\|\mathbf{W}_*\|_{2,4}^2$ and the above inequality implies

$$\begin{aligned} 2(\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q) - M) &\leq 2\tau M + \epsilon + \epsilon \\ \implies \widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q) &\leq (1 + \tau)M + \epsilon. \end{aligned}$$

The proof is complete. \square

A.2 Proof of Generalization in Theorem 1

Proof of Theorem 1, Generalization Part. Using symmetrization, we have

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) \right| \right] \leq 2\mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \left| \frac{1}{n} \sum_{i=1}^n \xi_i \ell(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) \right| \right],$$

where ξ is i.i.d. Rademacher random variables. The above Rademacher complexity can be bounded using the contraction theorem [51, Chapter 5]:

$$\begin{aligned} &\mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{n} \sum_{i=1}^n \xi_i \ell(y_i, f_{\mathbf{W}}^Q(\mathbf{x}_i)) \right] \\ &= \mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \max \left\{ \frac{1}{n} \sum_{i=1}^n \xi_i \ell(y_i, f_{\mathbf{W}}^Q(\mathbf{x}_i)), -\frac{1}{n} \sum_{i=1}^n \xi_i \ell(y_i, f_{\mathbf{W}}^Q(\mathbf{x}_i)) \right\} \right] \\ &\leq \mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{n} \sum_{i=1}^n \xi_i \ell(y_i, f_{\mathbf{W}}^Q(\mathbf{x}_i)) + \sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{n} \sum_{i=1}^n -\xi_i \ell(y_i, f_{\mathbf{W}}^Q(\mathbf{x}_i)) \right] \\ &\leq 4\mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{n} \sum_{i=1}^n \xi_i f_{\mathbf{W}}^Q(\mathbf{x}_i) \right] + 2\mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\frac{1}{n} \sum_{i=1}^n \xi_i \ell(0, y_i) \right] \\ &\leq 4\mathbb{E}_{\mathbf{x}_i, \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{\sqrt{m}} \sum_{r \leq m} \left\langle \frac{1}{n} \sum_{i=1}^n \xi_i a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top, \mathbf{w}_r \mathbf{w}_r^\top \right\rangle \right] + \frac{2}{\sqrt{n}} \\ &\leq 4\mathbb{E}_{\mathbf{x}_i, \xi} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \max_{r \in [m]} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \cdot \frac{1}{\sqrt{m}} \sum_{r \leq m} \|\mathbf{w}_r \mathbf{w}_r^\top\|_* \right] + \frac{2}{\sqrt{n}} \\ &\leq 4\mathbb{E}_{\mathbf{x}_i, \xi} \left[\max_{r \in [m]} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right] \cdot \underbrace{\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \frac{1}{\sqrt{m}} \sum_{r \leq m} \|\mathbf{w}_r\|_2^2}_{\leq B_w^2} + \frac{2}{\sqrt{n}}, \end{aligned}$$

where the last step used the power mean (or Cauchy-Schwarz) inequality on $\{\|\mathbf{w}_r\|_2\}$ and $\|\cdot\|_*$ denotes the matrix nuclear norm (sum of singular values). Now it only remains to bound the expected max operator norm above. We apply the matrix concentration lemma Bai and Lee [8, Lemma 8] to

deduce that

$$\begin{aligned}
& \mathbb{E}_{\mathbf{x}_i, \xi} \left[\max_{r \in [m]} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right] \\
& \leq 4\sqrt{\log(2Dm)} \cdot \mathbb{E}_{\mathbf{x}_i} \left[\sqrt{\max_{r \in [m]} \left\| \frac{1}{n^2} \sum_{i=1}^n \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i))^2 \|\mathbf{x}_i\|_2^2 \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}}} \right] \\
& \leq 4B_h \sqrt{\frac{\log(2Dm)}{n}} \cdot \mathbb{E}_{\mathbf{x}_i} \left[\sqrt{\max_{r,i} \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i))^2 \cdot \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}}} \right] \\
& \leq 4B_h \sqrt{\frac{\log(2Dm)}{n}} \left(\mathbb{E}_{\mathbf{x}_i} \left[\max_{r,i} \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i))^2 \right] \cdot \mathbb{E}_{\mathbf{x}_i} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right] \right)^{1/2} \\
& \leq 4C^2 B_h \sqrt{\frac{\log(2Dm)}{n}} \mathbb{E}_{\mathbf{x}_i} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right]^{1/2}.
\end{aligned}$$

Combining all the ingredients and substituting $M_{h,\text{op}} = B_h^{-1} \mathbb{E}_{\mathbf{x}_i} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right]^{1/2}$, the generalization error is bounded by

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\sup_{\|\mathbf{W}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\mathbf{W}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{W}}^Q) \right| \right] \leq \tilde{O} \left(\frac{B_h^2 B_w^2 M_{h,\text{op}}}{\sqrt{n}} \sqrt{\log(Dm)} + \frac{1}{\sqrt{n}} \right).$$

□

B Results on Feature Covariance

B.1 Technical tool

We first present a Lemma for relating the covariance of nonlinear random features to the covariance of certain polynomial bases, adapted from [27, Proposition 2].

Lemma 1 (Covariance through polynomials). Let $\mathbf{v}_i \stackrel{\text{iid}}{\sim} \text{Unif}(\mathbb{S}^{d-1})$ be random unit vectors for $i \in [D]$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_D]^\top \in \mathbb{R}^{D \times d}$.

(a) For any $k \geq 0$, suppose $D \leq O(d^{k+1-\delta})$ for some $\delta > 0$, then we have with high probability as $d \rightarrow \infty$ that

$$\lambda_{\min} \left((\mathbf{V}\mathbf{V}^\top)^{\odot(k+1)} \right) \geq \frac{1}{2},$$

where $(\cdot)^{\odot k}$ is the Hadamard product: $(\mathbf{A}^{\odot k})_{ij} = \mathbf{A}_{ij}^k$.

(b) In the same setting as above, let $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$ and define $\Sigma \in \mathbb{R}^{D \times D}$ with

$$\Sigma_{ij} = \mathbb{E}_{\mathbf{x}} \left[\mathbb{1} \{ \mathbf{v}_i^\top \mathbf{x} \geq 0 \} \mathbb{1} \{ \mathbf{v}_j^\top \mathbf{x} \geq 0 \} \right],$$

then we have

$$\lambda_{\min}(\Sigma) \geq \frac{1}{2} \|\mathbf{P}_{\geq k+1} \sigma_d\|_{L_2}^2,$$

where $\sigma_d : \text{Unif}(\mathbb{S}^{d-1}(\sqrt{d})) \rightarrow \mathbb{R}$ is defined as $\sigma_d(\mathbf{x}) := \mathbb{1} \{x_1 \geq 0\}$, and $\mathbf{P}_{\geq k+1}$ denotes the projection onto degree $\geq (k+1)$ polynomials under the base measure $\text{Unif}(\mathbb{S}^{d-1}(\sqrt{d}))$.

B.2 Lower bound on population covariance

We first present a lower bound when $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$ is uniform on the sphere, and when the features are the biasless indicator random features.

Lemma 2 (Lower bound of population covariance). Let $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})$ and suppose we sample $\mathbf{v}_i \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ for $1 \leq i \leq D$, where $D \leq O(d^K)$. Let $\Sigma \in \mathbb{R}^{D \times D}$ be the (population) covariance matrix of the random features $\{\mathbb{1}\{\mathbf{v}_i^\top \mathbf{x} \geq 0\}\}_{i \in [D]}$, that is,

$$\Sigma_{ij} := \mathbb{E}_{\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})} [\mathbb{1}\{\mathbf{v}_i^\top \mathbf{x} \geq 0\} \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} \geq 0\}],$$

then we have $\lambda_{\min}(\Sigma) \geq c > 0$ with high probability as $d \rightarrow \infty$, where $c = c_K$ is a constant that depends on K (and the indicator activation) but not d .

Proof. Let $\tilde{\mathbf{v}}_i := \mathbf{v}_i / \|\mathbf{v}_i\|_2$ denote the normalized version of \mathbf{v}_i , then $\tilde{\mathbf{v}}_i \stackrel{\text{iid}}{\sim} \text{Unif}(\mathbb{S}^{d-1})$ due to the spherical symmetry of $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Further using the positive homogeneity of $t \mapsto \mathbb{1}\{t \geq 0\}$ yields that

$$\begin{aligned} \Sigma_{ij} &= \mathbb{E}_{\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1})} [\mathbb{1}\{\tilde{\mathbf{v}}_i^\top \mathbf{x} \geq 0\} \mathbb{1}\{\tilde{\mathbf{v}}_j^\top \mathbf{x} \geq 0\}] \\ &= \mathbb{E}_{\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1}(\sqrt{d}))} [\mathbb{1}\{\tilde{\mathbf{v}}_i^\top \mathbf{x} \geq 0\} \mathbb{1}\{\tilde{\mathbf{v}}_j^\top \mathbf{x} \geq 0\}]. \end{aligned}$$

This falls into the setting of Lemma 1(b), applying which implies that with high probability (as $d \rightarrow \infty$) we have

$$\lambda_{\min}(\Sigma) \geq \frac{1}{2} \|\mathbb{P}_{\geq K+1} \sigma_d\|_{L_2}^2,$$

where $\sigma_d : \mathbb{S}^{d-1}(\sqrt{d}) \rightarrow \mathbb{R}$ is defined as $\sigma_d(\mathbf{x}) = \mathbb{1}\{x_1 \geq 0\}$, $\mathbb{P}_{\geq K+1}$ denotes the projection onto degree- $\geq K+1$ polynomials under the base measure $\text{Unif}(\mathbb{S}^{d-1}(\sqrt{d}))$, and the L_2 norm is under the same base measure. For large enough d , as $x_1 | \mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1}(\sqrt{d})) \Rightarrow \mathcal{N}(0, 1) := \gamma$ (where \Rightarrow denotes convergence in distribution), this is further lower bounded by

$$\frac{1}{4} \|\mathbb{P}_{\geq K+1} \mathbb{1}\{\cdot \geq 0\}\|_{L_2(\gamma)}^2 = c_K > 0 \quad (10)$$

as the indicator function is not a polynomial of any degree (so that its L_2 projection onto polynomials of degree $\leq K$ is not itself for any $K \geq 0$).

Decay of Eigenvalue Lower Bounds with Uniform Data.

We now provide a lower bound for the quantity $\|\mathbb{P}_{\geq K+1} \mathbb{1}\{\cdot \geq 0\}\|_{L_2(\gamma)}^2$, thereby giving a lower bound on c_K defined in (10). Indeed, we have

$$\|\mathbb{P}_{\geq K+1} \mathbb{1}\{\cdot \geq 0\}\|_{L_2(\gamma)}^2 = \sum_{j=K+1}^{\infty} \hat{\sigma}_j^2,$$

where

$$\mathbb{1}\{z \geq 0\} \stackrel{L_2(\gamma)}{=} \sum_{j=0}^{\infty} \hat{\sigma}_j h_j(z)$$

is the Hermite decomposition of $\mathbb{1}\{z \geq 0\}$. By [37], we know that

$$\hat{\sigma}_0 = \frac{1}{2}, \quad \hat{\sigma}_{2i+1} = (-1)^i \sqrt{\frac{1}{2\pi(2i+1)!} \frac{(2i)!}{2^{2i} i!}}, \quad \hat{\sigma}_{2i+2} = 0 \quad \text{for all } i \geq 0.$$

We now calculate the decay of $\hat{\sigma}_{2i+1}^2$. By Stirling's formula, we have

$$\hat{\sigma}_{2i+1}^2 = \frac{1}{2\pi(2i+1)!} \cdot \frac{(2i)!^2}{2^{2i}(i!)^2} \asymp \frac{1}{2\pi(2i+1)} \cdot \frac{\sqrt{2\pi \cdot 2i} (2i/e)^{2i}}{2^{2i} \cdot 2\pi i \cdot (i/e)^{2i}} \asymp C i^{-3/2}$$

for some absolute constant $C > 0$. This means that for all $i \geq 0$ we have $\hat{\sigma}_{2i+1}^2 \geq C i^{-3/2}$ for some (other) absolute constant $C > 0$, which gives

$$\sum_{j=K+1}^{\infty} \hat{\sigma}_j^2 \geq \sum_{i: 2i+1 \geq K+1} C i^{-3/2} \geq C K^{-1/2}.$$

Therefore we have $c_K \geq \Omega(K^{-1/2})$ for all K . \square

Covariance Lower Bounds for Non-uniform Data.

We further show that Assumption 2 can hold fairly generally when \mathbf{x} is no longer uniform on \mathbb{S}^{d-1} . Recall we choose $D \leq O(d^K)$ for some constant K .

We begin with assuming there exists a positive definite matrix $\mathbf{S} \in \mathbb{R}^{d \times d}$ such that \mathbf{x} is equal in distribution as $\mathbf{S}^{1/2}\mathbf{z} / \|\mathbf{S}^{1/2}\mathbf{z}\|_2$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. In other words, \mathbf{x} is distributed as a rescaled version of a d -dimensional Gaussian with *arbitrary covariance*, a fairly expressive set of distributions which can model the case where \mathbf{x} is far from uniform over the sphere. We show in this case that Assumption 2 holds, whenever \mathbf{S} has a bounded condition number (i.e. $\lambda_{\min}(\mathbf{S})/\lambda_{\max}(\mathbf{S}) \geq 1/\kappa$ where $\kappa > 0$ does not depend on d).

Indeed, we can deduce

$$\mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} \geq 0\} = \mathbb{1}\left\{\mathbf{v}_j^\top \frac{\mathbf{S}^{1/2}\mathbf{z}}{\|\mathbf{S}^{1/2}\mathbf{z}\|_2} \geq 0\right\} = \mathbb{1}\left\{(\mathbf{S}^{1/2}\mathbf{v}_j)^\top \mathbf{z} \geq 0\right\} = \mathbb{1}\left\{\frac{\mathbf{S}^{1/2}\mathbf{v}_j^\top}{\|\mathbf{S}^{1/2}\mathbf{v}_j\|_2} \mathbf{z} \geq 0\right\}.$$

Here the equality denotes two random variables following the same distribution. We apply the Hermite decomposition of indicator function to decompose the covariance matrix Σ :

$$\begin{aligned} \lambda_{\min}(\Sigma) &= \min_{\|\mathbf{u}\|_2=1} \mathbb{E}_{\mathbf{x}} [\mathbf{u}^\top \mathbf{g}(\mathbf{x}) \mathbf{g}(\mathbf{x})^\top \mathbf{u}] \\ &= \min_{\|\mathbf{u}\|_2=1} \mathbb{E}_{\mathbf{x}} \left[\sum_{i,j} \mathbb{1}\{\mathbf{v}_i^\top \mathbf{x} \geq 0\} \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} \geq 0\} u_i u_j \right] \\ &= \min_{\|\mathbf{u}\|_2=1} \sum_{i,j} (\mathcal{T}_1 + \mathcal{T}_2) u_i u_j, \end{aligned} \tag{11}$$

where \mathcal{T}_1 and \mathcal{T}_2 are given as follows,

$$\begin{aligned} \mathcal{T}_1 &= \sum_{\ell=0}^{\infty} \hat{\sigma}_\ell^2 \mathbb{E}_{\mathbf{z}} \left[h_\ell \left(\frac{\mathbf{S}^{1/2}\mathbf{v}_j^\top}{\|\mathbf{S}^{1/2}\mathbf{v}_j\|_2} \mathbf{z} \right) h_\ell \left(\frac{\mathbf{S}^{1/2}\mathbf{v}_i^\top}{\|\mathbf{S}^{1/2}\mathbf{v}_i\|_2} \mathbf{z} \right) \right] \\ &= \sum_{\ell=0}^{\infty} \hat{\sigma}_\ell^2 \left(\frac{\mathbf{v}_i^\top \mathbf{S} \mathbf{v}_j}{\|\mathbf{S}^{1/2}\mathbf{v}_j\|_2 \|\mathbf{S}^{1/2}\mathbf{v}_i\|_2} \right)^\ell, \\ \mathcal{T}_2 &= \sum_{\ell \neq k} \hat{\sigma}_\ell \hat{\sigma}_k \mathbb{E}_{\mathbf{z}} [h_\ell(\mathbf{v}_i^\top \mathbf{z}) h_k(\mathbf{v}_j^\top \mathbf{z})] = 0, \end{aligned}$$

where $\hat{\sigma}_\ell$ is the coefficient of Hermite decomposition of the indicator function, and \mathcal{T}_2 vanishes, due to the orthogonality of probabilistic Hermite polynomials. We proceed to bound the minimum singular value of Σ :

$$\lambda_{\min}(\Sigma) = \min_{\|\mathbf{u}\|_2=1} \sum_{i,j} \sum_{\ell=0}^{\infty} \hat{\sigma}_\ell^2 \left(\frac{\mathbf{v}_i^\top \mathbf{S} \mathbf{v}_j}{\|\mathbf{S}^{1/2}\mathbf{v}_j\|_2 \|\mathbf{S}^{1/2}\mathbf{v}_i\|_2} \right)^\ell u_i u_j.$$

Note that in the above decomposition, Σ is the sum of an infinite series of positive semidefinite matrices. To show Σ has a lower bounded smallest singular value, it suffices to show that there exists a summand in the infinite series being positive definite. We confirm this by analyzing the $\ell = K + 1$ summand. In fact, we show the following matrix $\Sigma_{K+1} \in \mathbb{R}^{D \times D}$ is positive definite and its smallest singular value is lower bounded by some constant independent of d .

$$[\Sigma_{K+1}]_{ij} = \hat{\sigma}_{K+1}^2 \left(\frac{\mathbf{v}_i^\top \mathbf{S} \mathbf{v}_j}{\|\mathbf{S}^{1/2}\mathbf{v}_j\|_2 \|\mathbf{S}^{1/2}\mathbf{v}_i\|_2} \right)^{K+1} \quad \text{for } i, j = 1, \dots, D.$$

We denote the normalized \mathbf{v}_j as $\tilde{\mathbf{v}}_j = \mathbf{v}_j / \|\mathbf{v}_j\|_2$, and derive a lower bound on the singular value of $\tilde{\Sigma}_{K+1}$ with

$$[\tilde{\Sigma}_{K+1}]_{ij} = (\tilde{\mathbf{v}}_i^\top \mathbf{S} \tilde{\mathbf{v}}_j)^{K+1}.$$

Using the tensor product notation, we rewrite $\tilde{\Sigma}_{K+1}$ as

$$\tilde{\Sigma}_{K+1} = \tilde{\mathbf{V}}^{*(K+1)} \mathbf{S}^{\otimes(K+1)} \left(\tilde{\mathbf{V}}^{*(K+1)} \right)^\top,$$

where $\tilde{\mathbf{V}} = [\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_D]^\top \in \mathbb{R}^{D \times d}$, $\tilde{\mathbf{V}}^{*(K+1)} \in \mathbb{R}^{D \times d^{K+1}}$ is the Khatri-Rao product, and $\mathbf{S}^{\otimes(K+1)} \in \mathbb{R}^{d^{K+1} \times d^{K+1}}$ denotes the Kronecker product. Then we know

$$\tilde{\Sigma}_{K+1} \succeq \lambda_{\min}^{K+1}(\mathbf{S}) \tilde{\mathbf{V}}^{*(K+1)} \left(\tilde{\mathbf{V}}^{*(K+1)} \right)^\top.$$

Moreover, using Lemma 1(a), we have $\lambda_{\min} \left(\tilde{\mathbf{V}}^{*(K+1)} \left(\tilde{\mathbf{V}}^{*(K+1)} \right)^\top \right) \geq 1/2$ as we picked $D \leq O(d^K)$. Substituting into Σ_{K+1} , we have

$$\begin{aligned} \lambda_{\min}(\Sigma_{K+1}) &\geq \frac{1}{2} \hat{\sigma}_{K+1}^2 \lambda_{\min}^{K+1}(\mathbf{S}) \left(\frac{\|\mathbf{v}_i\|_2 \|\mathbf{v}_j\|_2}{\|\mathbf{S}^{1/2} \mathbf{v}_i\|_2 \|\mathbf{S}^{1/2} \mathbf{v}_j\|_2} \right)^{K+1} \\ &\geq \frac{1}{2} \hat{\sigma}_{K+1}^2 \underbrace{\lambda_{\min}^{K+1}(\mathbf{S}) \lambda_{\max}^{-(K+1)}(\mathbf{S})}_{\kappa^{-(K+1)}}. \end{aligned}$$

Therefore the smallest singular value of Σ is lower bounded by $\Omega(\hat{\sigma}_{K+1}^2 \kappa^{-(K+1)})$, which is a constant only depending on K but not d . This finishes the proof. \square

We remark that Assumption 2 can hold much more generally than rescaled Gaussian distributions, provided the following set of conditions holds: There exists some random variable $\mathbf{z} \in \mathbb{R}^d$ such that \mathbf{x} is equal in distribution to $\mathbf{z}/\|\mathbf{z}\|_2$, a set of univariate polynomials $\{h_k^{(i)} : \mathbb{R} \rightarrow \mathbb{R}\}$ for $i = 1, \dots, D$ and $k \geq 0$ where each $h_k^{(i)}$ is a degree- k polynomial ‘‘assigned’’ to \mathbf{v}_i , and corresponding coefficients $\hat{\sigma}_k^{(i)} \in \mathbb{R}$, such that

$$\mathbb{1} \{ \mathbf{v}_i^\top \mathbf{z} \geq 0 \} \stackrel{L_2}{=} \sum_{k=0}^{\infty} \hat{\sigma}_k^{(i)} h_k^{(i)}(\mathbf{v}_i^\top \mathbf{z}).$$

Let $\Sigma_k, \Sigma_{k,>k}, \Sigma_{>k} \in \mathbb{R}^{D \times D}$ be defined as

$$\begin{aligned} [\Sigma_k]_{ij} &:= \mathbb{E}_{\mathbf{z}} \left[\hat{\sigma}_k^{(i)} \hat{\sigma}_k^{(j)} h_k^{(i)}(\tilde{\mathbf{v}}_i^\top \mathbf{z}) h_k^{(j)}(\tilde{\mathbf{v}}_j^\top \mathbf{z}) \right] \\ [\Sigma_{k,>k}]_{ij} &:= \mathbb{E}_{\mathbf{z}} \left[\hat{\sigma}_k^{(i)} h_k^{(i)}(\tilde{\mathbf{v}}_i^\top \mathbf{z}) \cdot \sum_{\ell>k} \hat{\sigma}_\ell^{(j)} h_\ell^{(j)}(\tilde{\mathbf{v}}_j^\top \mathbf{z}) \right] \\ [\Sigma_{>k}]_{ij} &:= \mathbb{E}_{\mathbf{z}} \left[\sum_{\ell>k} \hat{\sigma}_\ell^{(i)} h_\ell^{(i)}(\tilde{\mathbf{v}}_i^\top \mathbf{z}) \cdot \sum_{\ell>k} \hat{\sigma}_\ell^{(j)} h_\ell^{(j)}(\tilde{\mathbf{v}}_j^\top \mathbf{z}) \right], \end{aligned}$$

where $\tilde{\mathbf{v}}_i := \mathbf{v}_i / \|\mathbf{v}_i\|_2$. We assume we have

- The not-too-correlated condition: there exists some $\epsilon \in (0, 1]$ such that for any $k \geq 0$, we have

$$\Sigma_{\geq k} := \Sigma_k + \Sigma_{k,>k} + \Sigma_{k,>k}^\top + \Sigma_{>k} \succeq \epsilon(\Sigma_k + \Sigma_{>k}).$$

- For all large d and $D \leq O(d^k)$, we have

$$\Sigma_{k+1} \succeq c_{k+1} (\tilde{\mathbf{V}} \tilde{\mathbf{V}}^\top)^{\odot(k+1)},$$

where $c_{k+1} > 0$ is a constant that depends on k but not d .

In this case, we can deduce that for $D = O(d^K)$, Assumption 2 holds with $\lambda_K = \Omega(\epsilon^{K+2} c_{K+1})$. To see this, observe that we have the expansion $\mathbb{1} \{ t \geq 0 \} = \hat{\sigma}_0^{(i)} h_0^{(i)}(t) + \sum_{\ell>0} \hat{\sigma}_\ell^{(i)} h_\ell^{(i)}(t)$ for all $i = 1, \dots, D$, and thus

$$\Sigma = \Sigma_0 + \Sigma_{0,>0} + \Sigma_{0,>0}^\top + \Sigma_{>0,>0} \stackrel{(i)}{\succeq} \epsilon(\Sigma_0 + \Sigma_{>0}) \succeq \epsilon \Sigma_{>0},$$

where (i) applied the not-too-correlated condition. Repeating the above process for K times leads to

$$\Sigma \succeq \epsilon^{K+1} \Sigma_{>K} \succeq \epsilon^{K+2} (\Sigma_{K+1} + \Sigma_{>(K+1)}) \succeq \epsilon^{K+2} \Sigma_{K+1} \succeq \epsilon^{K+2} c_{K+1} (\tilde{\mathbf{V}} \tilde{\mathbf{V}}^\top)^{\odot(K+1)}.$$

Combining with existing lower bound $\lambda_{\min}((\tilde{\mathbf{V}} \tilde{\mathbf{V}}^\top)^{\odot(K+1)}) \geq 1/2$ (Lemma 1(a)), we see Assumption 2 holds with $\lambda_K = \Omega(\epsilon^{K+2} c_{K+1})$, a constant that depends on K and independent of d .

We further note that the above two conditions are all satisfied by the rescaled Gaussian distributions: Choosing $h_k^{(i)} \equiv h_k$ (the k -th Hermite polynomial) for all $i = 1, \dots, D$, the first condition holds with $\epsilon = 1$ since $\Sigma_{k,>k} = \mathbf{0}$, and the second condition holds with $c_{k+1} = (\lambda_{\min}(\mathbf{S})/\lambda_{\max}(\mathbf{S}))^{k+1}$ (as shown earlier). Combining with the fact they only assume things about the moments of \mathbf{x} (or \mathbf{z} ; since $h_k^{(i)}$ are polynomials), we see that they are indeed moment-based assumptions that contain Gaussian distributions with arbitrary covariances, and thus can be fairly general.

We also remark that while we have verified Assumption 2 for random features without biases, our analyses can be straightforwardly generalized to the case with bias by looking at the augmented input $[\mathbf{x}^\top, 1]^\top \in \mathbb{R}^{d+1}$ and analyzing its distributions in similar fashions.

B.3 Relative concentration of covariance estimator

Lemma 3 (Relative concentration of covariance estimator). Let $\{\mathbf{g}(\mathbf{x}_i) \in \mathbb{R}^D\}_{i=1}^n$ be i.i.d. random vectors such that $\|\mathbf{g}_1\|_2 \leq B_g$ almost surely and $\mathbb{E}[\mathbf{g}_1 \mathbf{g}_1^\top] = \Sigma \succeq \lambda_{\min} \mathbf{I}_D$. Let $\hat{\Sigma} := \frac{1}{n} \sum_{i=1}^n \mathbf{g}(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)^\top$ denote the empirical covariance matrix of $\{\mathbf{g}(\mathbf{x}_i)\}$. For any $\epsilon \in (0, 1)$, as soon as $n \geq C \epsilon^{-2} \lambda_{\min}^{-1} B_g^2 \log(n \vee D)$, we have

$$\mathbb{E} \left[\left\| \Sigma^{-1/2} \hat{\Sigma} \Sigma^{-1/2} - \mathbf{I}_D \right\|_{\text{op}} \right] \leq \epsilon.$$

Further, when $n \geq C \delta^{-2} \epsilon^{-2} \lambda_{\min}^{-1} B_g^2 \log(n \vee D)$ we have with probability at least $1 - \delta$ that

$$\left\| \Sigma^{-1/2} \hat{\Sigma} \Sigma^{-1/2} - \mathbf{I}_D \right\|_{\text{op}} \leq \epsilon,$$

where $C > 0$ is a universal constant. On the same event, we have the relative concentration

$$(1 - \epsilon) \Sigma \preceq \hat{\Sigma} \preceq (1 + \epsilon) \Sigma.$$

Proof. The first statement directly yields the second by the Markov inequality. To see how the second statement implies the third, we can left- and right- multiply the matrix inside by $\Sigma^{1/2} \mathbf{v}$ for any $\mathbf{v} \in \mathbb{R}^D$ and get that

$$\left| (\mathbf{v}^\top \Sigma^{1/2}) \Sigma^{-1/2} \hat{\Sigma} \Sigma^{-1/2} (\Sigma^{1/2} \mathbf{v}) - \mathbf{v}^\top \Sigma \mathbf{v} \right| = \left| \mathbf{v}^\top \hat{\Sigma} \mathbf{v} - \mathbf{v}^\top \Sigma \mathbf{v} \right| \leq \epsilon \mathbf{v}^\top \Sigma \mathbf{v},$$

which implies that $(1 - \epsilon) \Sigma \preceq \hat{\Sigma} \preceq (1 + \epsilon) \Sigma$.

We now prove the first statement, which builds on the following Rudelson's inequality for controlling expected deviation of heavy-tailed sample covariance matrices:

Lemma 4 (Restatement of Theorem 5.45, [50]). Let $\{\mathbf{a}_i \in \mathbb{R}^D\}_{i=1}^n$ be independent random vectors with $\mathbb{E}[\mathbf{a}_i \mathbf{a}_i^\top] = \mathbf{I}_D$. Let $\Gamma := \mathbb{E}[\max_{i \in [n]} \|\mathbf{a}_i\|_2^2]$. Then there exists a universal constant $C > 0$ such that letting $\delta := C \Gamma \log(n \vee D)/n$, we have

$$\mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{a}_i \mathbf{a}_i^\top - \mathbf{I}_D \right\|_{\text{op}} \right] \leq \delta \vee \sqrt{\delta}.$$

We will apply Lemma 4 on the whitened random vectors $\mathbf{h}(\mathbf{x}_i) := \Sigma^{-1/2} \mathbf{g}(\mathbf{x}_i)$ (Here we slightly abuse the notation to denote $\mathbf{h}(\mathbf{x})$ as the whitened feature using the population covariance matrix).

Clearly, $\mathbb{E}_{\mathbf{x}}[\mathbf{h}(\mathbf{x}_i)\mathbf{h}(\mathbf{x}_i)^\top] = \mathbb{E}_{\mathbf{x}}[\boldsymbol{\Sigma}^{-1/2}\mathbf{g}(\mathbf{x}_i)\mathbf{g}(\mathbf{x}_i)^\top\boldsymbol{\Sigma}^{-1/2}] = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1/2} = \mathbf{I}_D$. Further, we have

$$\|\mathbf{h}(\mathbf{x}_i)\|_2^2 = \mathbf{g}(\mathbf{x}_i)^\top\boldsymbol{\Sigma}^{-1}\mathbf{g}(\mathbf{x}_i) \leq \lambda_{\min}^{-1}\|\mathbf{g}(\mathbf{x}_i)\|_2^2 \leq \lambda_{\min}^{-1}B_g^2$$

almost surely, and thus $\Gamma := \mathbb{E}_{\mathbf{x}}[\max_{i \in [n]}\|\mathbf{h}(\mathbf{x}_i)\|_2^2] \leq \lambda_{\min}^{-1}B_g^2$. Therefore, $\{\mathbf{h}(\mathbf{x}_i)\}$ satisfy the conditions of Lemma 4, from which we obtain

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}}\left[\left\|\boldsymbol{\Sigma}^{-1/2}\widehat{\boldsymbol{\Sigma}}\boldsymbol{\Sigma}^{-1/2} - \mathbf{I}_D\right\|_{\text{op}}\right] \\ &= \mathbb{E}_{\mathbf{x}}\left[\left\|\frac{1}{n}\sum_{i=1}^n\mathbf{h}(\mathbf{x}_i)\mathbf{h}(\mathbf{x}_i)^\top - \mathbf{I}_D\right\|_{\text{op}}\right] \leq \frac{C\lambda_{\min}^{-1}B_g^2\log(n \vee D)}{n} \vee \sqrt{\frac{C\lambda_{\min}^{-1}B_g^2\log(n \vee D)}{n}}. \end{aligned}$$

Therefore, setting $n \geq C\epsilon^{-2}\lambda_{\min}^{-1}B_g^2\log(n \vee D)$, we get that

$$\mathbb{E}_{\mathbf{x}}\left[\left\|\boldsymbol{\Sigma}^{-1/2}\widehat{\boldsymbol{\Sigma}}\boldsymbol{\Sigma}^{-1/2} - \mathbf{I}_D\right\|_{\text{op}}\right] \leq \epsilon.$$

This finishes the proof. \square

C Proof of Warm Up Example

We denote $\{Y_{\ell,j}^d(\mathbf{x})\}_j$ as collections of spherical harmonics of degree ℓ on unit sphere \mathbb{S}^{d-1} . Subscript j indexes the harmonics, and $j = 1, \dots, B(d, \ell)$ with $B(d, \ell) = \frac{2\ell+d-2}{\ell} \binom{\ell+d-3}{\ell-1}$. We omit the superscript d when it is clear from the context. The spherical harmonics form an orthonormal basis of $V(d, \ell)$, where $V(d, \ell)$ denotes the span of restrictions of degree- ℓ homogeneous harmonic polynomials to the unit sphere \mathbb{S}^{d-1} . In particular, we have

$$\int_{\mathbb{S}^{d-1}} Y_{\ell,i}(\mathbf{x})Y_{\ell,j}(\mathbf{x})\tau_{d-1}d\mathbf{x} = \delta_{ij}\delta_{\ell k} \quad \text{with } \tau_{d-1} \text{ the uniform probability measure on } \mathbb{S}^{d-1}.$$

For a clear presentation, we consider a single nonzero component in $B(r_*, p_{\max}, A)$ denoted by $\alpha(\boldsymbol{\beta}^\top \mathbf{x})^p$ with $\|\boldsymbol{\beta}\|_2 < \infty$. Fix $k = p/2$. We choose $h(\mathbf{x}) \in \mathbb{R}^{B(d,k)}$ by taking $\{Y_{k,i}(\mathbf{x})\}$ as entries with $j = 1, \dots, B(d, k)$ for $k = 0, \dots, p/2$. Since $\{Y_{k,i}(\mathbf{x})\}_{i=1}^{B(d,k)}$ is a basis of $V(d, k)$, there exists $\boldsymbol{\theta} \in \mathbb{R}^{B(d,k)}$ satisfying $\boldsymbol{\theta}^\top h(\mathbf{x}) = (\boldsymbol{\beta}^\top \mathbf{x})^k$ for any $\mathbf{x} \in \mathbb{S}^{d-1}$. Moreover, each entry of $\boldsymbol{\theta}$ is given by

$$\theta_i = \int_{\mathbb{S}^{d-1}} Y_{k,i}(\mathbf{x})(\boldsymbol{\beta}^\top \mathbf{x})^k \tau_{d-1}d\mathbf{x}.$$

We can also check that $\|\boldsymbol{\theta}\|_2 \leq \|\boldsymbol{\beta}\|_2^k$ by observing $\int_{\mathbb{S}^{d-1}} \boldsymbol{\theta}^\top h(\mathbf{x})h(\mathbf{x})^\top \boldsymbol{\theta} \tau_{d-1}d\mathbf{x} = \int_{\mathbb{S}^{d-1}} (\boldsymbol{\beta}^\top \mathbf{x})^{2k} \tau_{d-1}d\mathbf{x}$.

Given $\boldsymbol{\theta}$ in the previous paragraph, we further choose $a(\mathbf{w}_0) = 2\alpha$ and $\mathbf{w}^* = \boldsymbol{\theta}$. By the symmetry in \mathbf{w}_0 , it yields

$$\mathbb{E}_{\mathbf{w}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [a(\mathbf{w}_0)\mathbb{1}\{\mathbf{w}_0^\top h(\mathbf{x}) \geq 0\}((\mathbf{w}^*)^\top h(\mathbf{x}))^2] = \alpha(\boldsymbol{\beta}^\top \mathbf{x})^p.$$

We now pick some m sufficiently large, and let $\mathbf{w}_r^* = \sqrt{2\alpha}m^{1/4}\mathbf{w}^*$ for $r = 1, \dots, m$. We check that $\{\mathbf{w}_r^*\}_{r=1}^m$ can achieve a small regularized empirical risk. Indeed, we have

$$\begin{aligned} \frac{1}{n}\sum_{i=1}^n \ell(y_i, f_{\mathbf{w}^*}(\mathbf{x}_i)) &\leq \frac{1}{n}\sum_{i=1}^n |f_{\mathbf{w}^*}(\mathbf{x}_i) - \alpha(\boldsymbol{\beta}^\top \mathbf{x}_i)^p| \\ &\leq \sup_{\mathbf{x} \in \mathbb{S}^{d-1}} |f_{\mathbf{w}^*}(\mathbf{x}) - \mathbb{E}_{\mathbf{w}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [a(\mathbf{w}_0)\mathbb{1}\{\mathbf{w}_0^\top h(\mathbf{x}) \geq 0\}((\mathbf{w}^*)^\top h(\mathbf{x}))^2]| \\ &\leq \sup_{\mathbf{x}} 2\alpha(\boldsymbol{\beta}^\top \mathbf{x})^p \left| \frac{1}{m}\sum_{r=1}^m \mathbb{1}\{\mathbf{w}_{0,r}^\top h(\mathbf{x}) \geq 0\} - \frac{1}{2} \right|. \end{aligned}$$

By Markov's inequality, for any $t > 0$, we derive

$$\begin{aligned} \mathbb{P} \left(\left| \frac{1}{m} \sum_{r=1}^m \mathbb{1}\{\mathbf{w}_{0,r}^\top h(\mathbf{x}) \geq 0\} - \frac{1}{2} \right| \geq t \right) &\leq \frac{\mathbb{E} \left[\frac{1}{m} \sum_{r=1}^m \mathbb{1}\{\mathbf{w}_{0,r}^\top h(\mathbf{x}) \geq 0\} - \frac{1}{2} \right]^2}{t^2} \\ &= \frac{1}{4mt^2}. \end{aligned}$$

On the other hand, the norm of $\{\mathbf{w}_r^*\}_{r=1}^m$ satisfies

$$\sum_{r=1}^m \|\mathbf{w}_r^*\|_2^4 = \sum_{r=1}^m \frac{4\alpha^2}{m} \|\mathbf{w}^*\|_2^4 = 4\alpha^2 \|\boldsymbol{\theta}\|_2^4 \leq 4\alpha^2 \|\boldsymbol{\beta}\|_2^p.$$

Therefore, for a fixed optimization accuracy $\epsilon > 0$, we choose $m = O(\epsilon^{-3})$ and λ some constant. This yields $\mathcal{R}_{\ell,\lambda}(f_{\mathbf{w}^*}) \leq \epsilon$.

Considering any second order stationary point $\{\widehat{\mathbf{w}}_r\}_{r=1}^m$, Theorem ?? gives rise to the following generalization bound.

$$\mathbb{E}_{\mathbf{w}_0} [\mathcal{R}_\ell(f_{\widehat{\mathbf{w}}}) - \mathbb{E}_{\mathbf{x},y}[\mathcal{R}_\ell(f_{\widehat{\mathbf{w}}})]] \leq \tilde{O} \left(\frac{B_{\mathbf{w}}^2 M_{h(\mathbf{x}),\text{op}}}{\sqrt{n}} + \frac{1}{\sqrt{n}} \right),$$

It suffices to find $B_{\mathbf{w}}$ and $M_{h(\mathbf{x}),\text{op}}$. We already checked $B_{\mathbf{w}}^4 = \sum_{r=1}^m \|\mathbf{w}_r^*\|_2^4 \leq 2\alpha \|\boldsymbol{\beta}\|_2^p = O(1)$. For $M_{h(\mathbf{x}),\text{op}}$, we expand it as

$$M_{h(\mathbf{x}),\text{op}}^2 = \mathbb{E}_{\mathbf{x}} \left[\left\| \frac{1}{n} \sum_{i=1}^n h(\mathbf{x}_i) h(\mathbf{x}_i)^\top h(\mathbf{x}_i) h(\mathbf{x}_i)^\top \right\|_{\text{op}} \right].$$

By the orthogonality of spherical harmonics, we know

$$\mathbb{E}_{\mathbf{x}} [h(\mathbf{x}) h(\mathbf{x})^\top h(\mathbf{x}) h(\mathbf{x})^\top] = B(d, k) \mathbf{I}.$$

Hence, we have $M_{h(\mathbf{x}),\text{op}}^2 = O(B(d, k)) = \tilde{O}(d^k)$. Substituting into the generalization bound, we deduce

$$\mathbb{E}_{\mathbf{w}_0} [\mathcal{R}_\ell(f_{\widehat{\mathbf{w}}}) - \mathbb{E}_{\mathbf{x},y}[\mathcal{R}_\ell(f_{\widehat{\mathbf{w}}})]] \leq \tilde{O} \left(\sqrt{\frac{d^k}{n}} + \frac{1}{\sqrt{n}} \right).$$

From a statistical recovery perspective, to achieve ϵ generalization gap, the sample complexity grows in the order of $d^{p/2}/\epsilon^2$.

To learn multiple polynomials in $B(r_*, p_{\max}, A)$, we dilate $h(\mathbf{x})$ as concatenating spherical harmonics up to order $k = p_{\max}/2$, i.e., $h(\mathbf{x}) = [\{Y_{0,j}(\mathbf{x})\}_{j=1}^{B(d,0)}, \dots, \{Y_{k,j}(\mathbf{x})\}_{j=1}^{B(d,k)}]^\top$. With a slight modification on the argument for learning a single polynomial, we can show that the sample complexity grows as $\tilde{O}(d^{p_{\max}/2}/\epsilon^2)$.

D Proofs for Section 4

This section devotes to the proof of Theorem 2. The proof consists of two main parts: expressivity of neural representation (Sections D.1 and D.2) and generalization property of Quad-Neural (Section D.3). Besides, Section D.5 presents that using data dependent regularizer also achieves improved sample complexity.

D.1 Expressivity of Neural Random Features

Lemma 5. For a given vector $\boldsymbol{\beta}$ and integer $k \geq 0$, we let $\mathbf{v} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_d)$ be a standard Gaussian vector and $b \sim \mathbf{N}(0, 1)$ independent of \mathbf{v} . Then there exists $a(\mathbf{v}, b)$ such that $\mathbb{E}_{\mathbf{v},b}[a(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\}] = (\boldsymbol{\beta}^\top \mathbf{x})^k$ holds for any $\mathbf{x} \in \mathbb{S}^{d-1}$.

Proof. We denote by $H_j(x)$ the j -th probabilistic Hermite polynomial. We pick

$$a(\mathbf{v}, b) = \begin{cases} c_k H_k(\mathbf{v}^\top \boldsymbol{\beta} / \|\boldsymbol{\beta}\|_2) \mathbb{1}\{0 < -b < 1/(2k)\}, & \text{if } k \text{ is even} \\ c_k H_k(\mathbf{v}^\top \boldsymbol{\beta} / \|\boldsymbol{\beta}\|_2) \mathbb{1}\{|b| < 1/(2k)\}, & \text{if } k \text{ is odd} \end{cases},$$

where c_k is a constant to be determined. For a fixed \mathbf{x} , we denote $z_1 = \mathbf{v}^\top \boldsymbol{\beta} / \|\boldsymbol{\beta}\|_2$ and $z_2 = \mathbf{v}^\top \mathbf{x}$. It is straightforward to check that z_1, z_2 is jointly Gaussian with zero mean and $\mathbb{E}[z_1 z_2] = \boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2$. We can now deduce that z_1 and $(\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2) z_2 + \sqrt{1 - (\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2)^2} z_3$ follow the same distribution, where z_3 is standard Gaussian independent of z_1 and z_2 . For an even k , we can check

$$\begin{aligned} & \mathbb{E}_{\mathbf{v}, b}[a(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\}] \\ &= c_k \mathbb{E}_{z_1, z_2, b}[H_k(z_1) \mathbb{1}\{z_2 + b \geq 0\} \mathbb{1}\{0 < -b < 1/(2k)\}] \\ &= c_k \mathbb{E}_{z_2, z_3, b} \left[H_k \left((\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2) z_2 + \sqrt{1 - (\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2)^2} z_3 \right) \mathbb{1}\{z_2 + b \geq 0\} \mathbb{1}\{0 < -b < 1/(2k)\} \right] \\ &= c_k \mathbb{E}_b \mathbb{E}_{z_2, z_3} \left[H_k \left((\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2) z_2 + \sqrt{1 - (\boldsymbol{\beta}^\top \mathbf{x} / \|\boldsymbol{\beta}\|_2)^2} z_3 \right) \mathbb{1}\{z_2 \geq -b\} \mathbb{1}\{0 < -b < 1/(2k)\} \mid b \right] \\ &\stackrel{(i)}{=} c_k q_k (\boldsymbol{\beta}^\top \mathbf{x})^k \|\boldsymbol{\beta}\|_2^{-k}, \end{aligned}$$

where $q_k = \mathbb{E}_b \left[(k-1)!! \frac{\exp(-b^2/2)}{\sqrt{2\pi}} \mathbb{1}\{0 < -b < 1/(2k)\} \sum_{j=1, \text{ odd}}^{k-1} \frac{(-1)^{(k-1+j)/2}}{j!!} \binom{k/2-1}{(j-1)/2} b^j \right]$.

The equality (i) invokes *Lemma A.6* in Allen-Zhu et al. [4]. Similarly, for an odd k , we have

$$\mathbb{E}_{\mathbf{v}, b}[a(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\}] = c_k q_k (\boldsymbol{\beta}^\top \mathbf{x})^k \|\boldsymbol{\beta}\|_2^{-k}$$

with $q_k = \mathbb{E}_b \left[(k-1)!! \frac{\exp(-b^2/2)}{\sqrt{2\pi}} \mathbb{1}\{|b| \leq 1/(2k)\} \sum_{j=1, \text{ even}}^{k-1} \frac{(-1)^{(k-1+j)/2}}{j!!} \binom{k/2-1}{(j-1)/2} b^j \right]$. Here we unify the notation to denote q_k as the coefficient for both the even and odd k 's. Using *Claim C.1* in Allen-Zhu et al. [4], we can lower bound p_k by $|p_k| \geq \frac{(k-1)!!}{200k^2}$. The proof is complete by choosing $c_k = 1/p_k$, and accordingly, $|c_k| \leq \frac{200k^2}{(k-1)!!} \|\boldsymbol{\beta}\|_2^k$. \square

From Expectation to Finite Neuron Approximation.

Lemma 6. For a given $\epsilon > 0$ and $\delta > 0$, we choose $D = 2 \times 200^2 k^5 \|\boldsymbol{\beta}\|_2^{2k} / (\epsilon^2 \delta)$ and independently generate $\mathbf{v}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $b_j \sim \mathcal{N}(0, 1)$ for $j = 1, \dots, D$. Then with probability at least $1 - \delta$, we have

$$\left\| \frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k \right\|_{L_2} \leq \epsilon.$$

Proof. The desired bound can be obtained by Chebyshev's inequality. We bound the second moment of the L_2 norm as

$$\begin{aligned} & \mathbb{E}_{\mathbf{v}, b} \left\| \frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k \right\|_{L_2}^2 \\ &= \mathbb{E}_{\mathbf{v}, b} \mathbb{E}_{\mathbf{x}} \left[\frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k \right]^2 \\ &= \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{v}, b} \left[\frac{1}{D} \sum_{j=1}^D \left(a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - \mathbb{E}_{\mathbf{v}_j, b_j} [a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\}] \right) \right]^2 \\ &= \frac{1}{D^2} \mathbb{E}_{\mathbf{x}} \left[\sum_{j=1}^D \mathbb{E}_{\mathbf{v}_j, b_j} \left[a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - \mathbb{E}_{\mathbf{v}_j, b_j} [a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\}] \right]^2 \right] \\ &= \frac{1}{D} \mathbb{E}_{\mathbf{x}, \mathbf{v}, b} [a(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k]^2. \end{aligned}$$

Using Lemma 5, we have

$$\begin{aligned}
\mathbb{E}_{\mathbf{x}, \mathbf{v}, b} [a(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k]^2 &= \mathbb{E}_{\mathbf{x}, \mathbf{v}, b} \left[a^2(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^{2k} \right] \\
&\leq \mathbb{E}_{\mathbf{x}, \mathbf{v}, b} \left[c_k^2 H_k^2(\boldsymbol{\beta}^\top \mathbf{v} / \|\boldsymbol{\beta}\|_2) + (\boldsymbol{\beta}^\top \mathbf{x})^{2k} \right] \\
&= c_k^2 \sqrt{2\pi} k! + \mathbb{E}_{\mathbf{x}} [(\boldsymbol{\beta}^\top \mathbf{x})^{2k}] \\
&\leq \frac{200^2 k^4}{(k-1)!!(k-1)!!} k! \|\boldsymbol{\beta}\|_2^{2k} + \|\boldsymbol{\beta}\|_2^{2k} \\
&\leq 2 \times 200^2 k^5 \|\boldsymbol{\beta}\|_2^{2k}.
\end{aligned}$$

The last inequality invokes the identity $\frac{k!}{((k-1)!!)^2} + 1 \leq \frac{k!}{(k-1)!} + 1 \leq 2 \frac{k!}{(k-1)!} = 2k$. Therefore, choosing $D = 2 \times 200^2 k^5 \|\boldsymbol{\beta}\|_2^{2k} / (\epsilon^2 \delta)$ gives rise to

$$\begin{aligned}
&\mathbb{P} \left(\left\| \frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k \right\|_{L_2} \geq \epsilon \right) \\
&\leq \epsilon^{-2} \mathbb{E}_{\mathbf{x}, b} \left[\left\| \frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}^\top \mathbf{x})^k \right\|_{L_2}^2 \right] \\
&\leq \delta.
\end{aligned}$$

This completes the proof. \square

From A Single Polynomial to A Sum of Polynomials.

Lemma 7. Given a function $f(\mathbf{x}) = \sum_{s=1}^{r_*} (\boldsymbol{\beta}_s^\top \mathbf{x})^{k_s}$ defined on $\mathbf{x} \in \mathbb{S}^{d-1}$, and positive constants $\epsilon > 0$ and $\delta > 0$, we choose $D \geq \frac{2 \times 200^2 r_*^3 \sum_{s=1}^{r_*} k_s^5 \|\boldsymbol{\beta}_s\|_2^{2k_s}}{\epsilon^2 \delta}$. Then there exists scalar $a(\mathbf{v}_j, b_j)$ for $j = 1, \dots, D$, such that with probability at least $1 - \delta$ over independently randomly sampled $\mathbf{v}_j \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_d)$ and $b_j \sim \mathbf{N}(0, 1)$ for $j = 1, \dots, D$, we have

$$\left\| \frac{1}{D} \sum_{j=1}^D a(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - f(\mathbf{x}) \right\|_{L_2} \leq \epsilon.$$

Proof. We apply Lemma 5 and Lemma 6 repeatedly for r_* times. Specifically, for each fixed $s \leq r_*$, Lemma 5 implies that there exists $a_s(\mathbf{v}, b)$ such that $\mathbb{E}_{\mathbf{v}, b} [a_s(\mathbf{v}, b) \mathbb{1}\{\mathbf{v}^\top \mathbf{x} + b \geq 0\}] = (\boldsymbol{\beta}_s^\top \mathbf{x})^{k_s}$. Then we choose $D_s \geq 2 \times 200^2 k_s^5 \|\boldsymbol{\beta}_s\|_2^{2k_s} r_*^3 / (\epsilon^2 \delta)$ so that with probability at least $1 - \delta / r_*$, the following L_2 bound holds

$$\left\| \frac{1}{D_s} \sum_{j=1}^{D_s} a_s(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}_s^\top \mathbf{x})^{k_s} \right\|_{L_2} \leq \epsilon / r_*.$$

To this end, we set $D = \sum_{s=1}^{r_*} D_s \geq \frac{2 \times 200^2 r_*^3 \sum_{s=1}^{r_*} k_s^5 \|\boldsymbol{\beta}_s\|_2^{2k_s}}{\epsilon^2 \delta}$ and define

$$\mathbf{g}(\mathbf{x}) = [\mathbf{g}_1(\mathbf{x})^\top, \dots, \mathbf{g}_{r_*}(\mathbf{x})^\top]^\top$$

with the j -th element of \mathbf{g}_s as $[\mathbf{g}_s(\mathbf{x})]_j = \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\}$ for $j = 1, \dots, D_s$. In other words, we construct a random feature vector $\mathbf{g}(\mathbf{x}) \in \mathbb{R}^D$ by stacking all the random features for approximating the k_s -degree polynomial. Similar to \mathbf{g} , we denote $\mathbf{a} = [\mathbf{a}_1^\top, \dots, \mathbf{a}_{r_*}^\top]^\top$ with the j -th element of \mathbf{a}_s as $[\mathbf{a}_s]_j = \frac{1}{D_s} a_s(\mathbf{v}_j, b_j)$ for $j = 1, \dots, D_s$. Then we can bound the L_2 distance between $f(\mathbf{x})$ and

$\mathbf{a}^\top \mathbf{g}(\mathbf{x})$:

$$\begin{aligned} \|\mathbf{a}^\top \mathbf{g}(\mathbf{x}) - f(\mathbf{x})\|_{L_2} &\leq \sum_{s=1}^{r_*} \|\mathbf{a}_s^\top \mathbf{g}_s(\mathbf{x}) - (\boldsymbol{\beta}_s^\top \mathbf{x})^{k_s}\|_{L_2} \\ &\leq \sum_{s=1}^{r_*} \left\| \frac{1}{D_s} \sum_{j=1}^{D_s} a_s(\mathbf{v}_j, b_j) \mathbb{1}\{\mathbf{v}_j^\top \mathbf{x} + b_j \geq 0\} - (\boldsymbol{\beta}_s^\top \mathbf{x})^{k_s} \right\|_{L_2} \\ &\leq \epsilon. \end{aligned}$$

The above inequality holds with probability $1 - \delta$ by the union bound. We complete the proof. \square

Lemma 7 showcases how to express a sum of polynomials by stacking neural random features for approximating individual polynomials. This technique will be extensively used in the remaining proofs.

D.2 Expressivity of Quad-h

We show Quad-Neural with neural representation \mathbf{h} can approximate any function f of the form

$$f(\mathbf{x}) = \sum_{s=1}^{r_*} \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}, \quad \text{where } |\alpha_s| \leq 1, \quad \|(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}\|_{L_2} \leq 1, \quad p_s \leq p \text{ for all } s. \quad (12)$$

To ease the presentation, we temporarily assume all the p_s are even. We extend to odd-degree polynomials in 9. Recall we denote

$$\mathbf{g}(\mathbf{x}) = [\mathbf{g}_1(\mathbf{x})^\top, \dots, \mathbf{g}_{r_*}(\mathbf{x})^\top]^\top \quad \text{with } \mathbf{g}_s(\mathbf{x}) \text{ being a collection of random indicator functions.}$$

We whiten $\mathbf{g}(\mathbf{x})$ by the estimated covariance matrix $\widehat{\boldsymbol{\Sigma}}$ to obtain $\mathbf{h}(\mathbf{x}) = \widehat{\boldsymbol{\Sigma}}^{-1/2} \mathbf{g}(\mathbf{x})$. Note that $\mathbf{h}(\mathbf{x})$ is a D -dimensional vector. The approximation of Quad-h is stated in the following lemma.

Lemma 8. For a given f in the form of (12) with all p_s even, and for small constants $\epsilon > 0$ and $\delta > 0$, we choose $D \geq \frac{4 \times 50^2 r_*^3 \sum_{s=1}^{r_*} p_s^5 \|\boldsymbol{\beta}_s\|_2^{p_s}}{\epsilon^2 \delta}$, and $m \geq \frac{54 r_* D (1 + \log \frac{8}{\epsilon})}{\epsilon^2} \log \frac{1}{\epsilon}$. Let $\mathbf{w}_{0,r} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$ and $a_r \stackrel{\text{iid}}{\sim} \text{Unif}(\{\pm 1\})$ for $r = 1, \dots, m$, then there exist proper $\{\mathbf{w}_r^*\}$ such that with probability at least $1 - \delta$, we have

$$\left\| \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^*)^\top \mathbf{h}(\mathbf{x}))^2 - f(\mathbf{x}) \right\|_{L_2} \leq 7r_* \epsilon.$$

Proof. By definition, f can be written as a sum of polynomials with leading coefficients α_s . We partition m neurons into two parts according to the sign of a_r . We will use the positive part to express those polynomials with positive coefficient α_s , and negative part to express those with negative coefficients. We first show for sufficiently large m , the number of positive a_r 's exceeds $\frac{1}{3}m$ with high probability. This follows from the tail bound of i.i.d. binomial random variables. By the Hoeffding's inequality, we have

$$\mathbb{P}(\text{Number of positive } a_r \leq k) \leq \exp(-2m(1/2 - (k/m)^2)).$$

Letting $k = \frac{1}{3}m$ and setting $\mathbb{P}(\text{Number of positive } a_r \leq k) \leq \delta$, we have $m \geq 2 \log \frac{1}{\delta}$. We denote $\mathcal{I}_1 = \{1, \dots, m/3\}$ and $\mathcal{I}_2 = \{m/3 + 1, \dots, 2m/3\}$. Without loss of generality, we assume $a_r = 1$ for $r \in \mathcal{I}_1$.

The remaining proof is built upon Lemma 7. We choose $D = \frac{50^2 r_*^3 \sum_{s=1}^{r_*} p_s^5 \|\boldsymbol{\beta}_s\|_2^{p_s}}{\epsilon^2 \delta}$, so that with probability at least $1 - \delta$, there exists \mathbf{a} with $\|\mathbf{a}^\top \mathbf{g}(\mathbf{x}) - \sum_{s=1}^{r_*} (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2}\|_{L_2} \leq \epsilon$. We further partition \mathcal{I}_1 into r_* consecutive groups of equal size m_0 , i.e., $r_* m_0 = m/3$. Within a group, we aim to approximate $\alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}$ with $\alpha_s > 0$ for some fixed $s \leq r_*$. Accordingly, we choose

$$\begin{aligned}
\mathbf{w}_r^{s,*} &= 2\sqrt{\alpha_s}(3r_*)^{1/4}m_0^{-1/4}\widehat{\Sigma}^{1/2}[\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top \text{ for } r = 1, \dots, m_0. \text{ We have} \\
&\left\| \frac{1}{2\sqrt{m}} \sum_{r=1}^{m_0} a_{0,r} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^{s,*})^\top \mathbf{h}(\mathbf{x}))^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&= \left\| \frac{1}{2\sqrt{3r_*}m_0} \sum_{r=1}^{m_0} 4\sqrt{3r_*} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} \alpha_s m_0^{-1/2} (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&= \left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\alpha_s \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2}. \tag{13}
\end{aligned}$$

We know $\mathbf{a}_s^\top \mathbf{g}_s$ well approximates $(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2}$. If $\sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\}$ concentrates around 1, then the above L_2 norm can be bounded by $O(\epsilon)$. We substantiate this reasoning by the following claim:

Claim 1. With probability at least $1 - 2\delta$, we have

$$\sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \left| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} - 1 \right| \leq 6\sqrt{\frac{D \log(3m_0) (1 + \log \frac{2}{\delta})}{m_0}}.$$

The proof of the claim is deferred to Appendix D.4. Based on the claim, we are ready to finish proving (13). By the triangle inequality, we deduce

$$\begin{aligned}
&\left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\alpha_s \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&= \left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\alpha_s \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\mathbf{a}_s^\top \mathbf{g}_s)^2 + \alpha_s (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&\leq \left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\alpha_s \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\mathbf{a}_s^\top \mathbf{g}_s)^2 \right\|_{L_2} + \left\| \alpha_s (\mathbf{a}_s^\top \mathbf{g}_s)^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&\leq \alpha_s \left\| (\mathbf{a}_s^\top \mathbf{g}_s)^2 \right\|_{L_2} \left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} - 1 \right\|_{L_2} \\
&\quad + \alpha_s \left\| \mathbf{a}_s^\top \mathbf{g}_s + (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} \left\| \mathbf{a}_s^\top \mathbf{g}_s - (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} \\
&\leq 6\alpha_s \left(\left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} + \epsilon \right) \sqrt{\frac{D \log(3m_0) (1 + \log \frac{2}{\delta})}{m_0}} + \alpha_s \epsilon \left(2 \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} + \epsilon \right).
\end{aligned}$$

The above upper bound holds with probability no smaller than $1 - 3\delta$. Taking

$$m_0 = \frac{18D (1 + \log \frac{2}{\delta})}{\epsilon^2} \log \frac{1}{\epsilon},$$

for a small $\epsilon < \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2}$, with probability at least $1 - 3\delta$, the following

$$\left\| \frac{1}{m_0} \sum_{r=1}^{m_0} a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^{s,*})^\top \mathbf{h}(\mathbf{x}))^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \leq 7\alpha_s \epsilon \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2}$$

holds true for the s -th group with $\alpha_s > 0$. When $\alpha_s < 0$, we simply set $\mathbf{w}_r^{s,*} = \mathbf{0}$. As a result, in \mathcal{I}_1 , we can express all the polynomial with a positive coefficient.

To express polynomials with negative coefficients, we use \mathcal{I}_2 analogously. By evenly partitioning \mathcal{I}_2 into r_* consecutive groups, for a fixed $s \leq r_*$ and $\alpha_s < 0$, we choose $\mathbf{w}_r^{s,*} = 2\sqrt{|\alpha_s|}(3r_*)^{1/4}m_0^{-1/4}\widehat{\Sigma}^{1/2}[\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top$. Using exactly the same argument in \mathcal{I}_1 , with probability at least $1 - 3\delta$, for $\alpha_s < 0$, we also have

$$\left\| \frac{1}{m_0} \sum_{r=1}^{m_0} a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^{s,*})^\top \mathbf{h}(\mathbf{x}))^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \leq 7|\alpha_s| \epsilon \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2}.$$

The last step for proving Lemma 8 is to combine \mathcal{I}_1 and \mathcal{I}_2 together and choose the remaining weight parameters \mathbf{w}_r^* identically $\mathbf{0}$ for $r \geq 2m/3 + 1$. Substituting into the Quad-h model, with probability at least $1 - 4\delta$, we deduce

$$\begin{aligned}
& \left\| \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^*)^\top \mathbf{h}(\mathbf{x}))^2 - f(\mathbf{x}) \right\|_{L_2} \\
&= \left\| \frac{1}{2\sqrt{m}} \sum_{r \in \mathcal{I}_1 \cup \mathcal{I}_2} a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^*)^\top \mathbf{h}(\mathbf{x}))^2 - f(\mathbf{x}) \right\|_{L_2} \\
&\leq \sum_{s=1}^{r_*} \left\| \frac{1}{2\sqrt{m}} \sum_{r \in \mathcal{I}_1 \cup \mathcal{I}_2} a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^{s,*})^\top \mathbf{h}(\mathbf{x}))^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&\leq \sum_{s=1}^{r_*} \left\| \frac{1}{m_0} \sum_{r=1}^{m_0} 2\alpha_s \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} (\mathbf{a}_s^\top \mathbf{g}_s(\mathbf{x}))^2 - \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2} \\
&\leq 7\epsilon \sum_{s=1}^{r_*} |\alpha_s| \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} \\
&\leq 7\epsilon \sqrt{\sum_{s=1}^{r_*} \alpha_s^2 \sum_{s=1}^{r_*} \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \right\|_{L_2}} \\
&\leq 7r_* \epsilon.
\end{aligned}$$

The width m satisfies $m = 3r_* m_0 \geq \frac{54r_* D(1+\log \frac{2}{\delta})}{\epsilon^2} \log \frac{1}{\epsilon}$. Replacing $\delta = \delta/4$ completes the proof. \square

Expressivity with Odd-Degree Polynomials.

Quad-h model can also efficiently express odd-degree polynomials. We rely on the following decomposition trick. Let k be an integer. We rewrite a $(2k + 1)$ -degree polynomial as

$$(\boldsymbol{\beta}^\top \mathbf{x})^{2k+1} = \left(\frac{(\boldsymbol{\beta}^\top \mathbf{x})^{k+1} + (\boldsymbol{\beta}^\top \mathbf{x})^k}{2} \right)^2 - \left(\frac{(\boldsymbol{\beta}^\top \mathbf{x})^{k+1} - (\boldsymbol{\beta}^\top \mathbf{x})^k}{2} \right)^2.$$

Since QuadNTK can naturally implement the quadratic function, we only require that the neural representation $\mathbf{h}(\mathbf{x})$ can approximate $(\boldsymbol{\beta}^\top \mathbf{x})^{k+1} \pm (\boldsymbol{\beta}^\top \mathbf{x})^k$. This is true since random indicator functions can approximate $(\boldsymbol{\beta}^\top \mathbf{x})^{k+1}$ and $(\boldsymbol{\beta}^\top \mathbf{x})^k$ due to Lemma 5. We denote $\mathbf{a}_1^\top \mathbf{g}_1(\mathbf{x}) \approx (\boldsymbol{\beta}^\top \mathbf{x})^{k+1}$ in L_2 , and $\mathbf{a}_2^\top \mathbf{g}_2(\mathbf{x}) \approx (\boldsymbol{\beta}^\top \mathbf{x})^k$ in L_2 . Then by stacking \mathbf{g}_1 and \mathbf{g}_2 , we have $[\mathbf{a}_1^\top, \pm \mathbf{a}_2^\top][\mathbf{g}_1^\top, \mathbf{g}_2^\top]^\top \approx (\boldsymbol{\beta}^\top \mathbf{x})^{k+1} \pm (\boldsymbol{\beta}^\top \mathbf{x})^k$ in L_2 . Therefore, we only need to augment the dimension D of the neural representation to approximate odd-degree polynomials. We concretize this argument in the following lemma.

Lemma 9. For a given f in the form of (12), and small constants $\epsilon > 0$ and $\delta > 0$, we choose $D \geq \frac{8 \times 50^2 r_*^3 \sum_{s=1}^{r_*} p_s^5 \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil}}{\epsilon^2 \delta^5}$, and $m \geq \frac{54r_* D(1+\log \frac{8}{\delta})}{\epsilon^2} \log \frac{1}{\epsilon}$. Let $\mathbf{w}_{0,r} \stackrel{\text{iid}}{\sim} \mathbf{N}(\mathbf{0}, \mathbf{I}_D)$ and $a_r \stackrel{\text{iid}}{\sim} \text{Unif}(\{\pm 1\})$ for $r = 1, \dots, m$, then there exist proper $\{\mathbf{w}_r^*\}$ such that with probability at least $1 - \delta$, we have

$$\left\| \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^*)^\top \mathbf{h}(\mathbf{x}))^2 - f(\mathbf{x}) \right\|_{L_2} \leq 7r_* \epsilon.$$

Proof. We can write

$$\begin{aligned}
f(\mathbf{x}) &= \sum_{s=1}^{r_*} \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \mathbb{1}\{p_s \text{ is even}\} + \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \mathbb{1}\{p_s \text{ is odd}\} \\
&= \sum_{s=1}^{r_*} \alpha_s \left((\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right)^2 \mathbb{1}\{p_s \text{ is even}\} \\
&\quad + \alpha_s \left[\left(\frac{(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2+1} + (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2-1}}{2} \right)^2 - \left(\frac{(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2+1} - (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2-1}}{2} \right)^2 \right] \mathbb{1}\{p_s \text{ is odd}\}.
\end{aligned}$$

Applying Lemma 5 once, there exists \mathbf{a}_s such that $\|\mathbf{a}_s^\top - (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2}\|_{L_2} \leq \epsilon/r_*$, when p_s is even and the corresponding $D_s \geq \frac{4 \times 50^2 r_*^3 \sum_{s=1}^{r_*} p_s^5 \|\boldsymbol{\beta}_s\|_2^{2p_s}}{\epsilon^2 \delta}$.

For an odd p_s , we apply the technique in Lemma 6. There exist $\mathbf{a}_{s,+}$ and $\mathbf{a}_{s,-}$ with corresponding random indicator features $\mathbf{g}_{s,+}(\mathbf{x})$ and $\mathbf{g}_{s,-}(\mathbf{x})$ such that

$$\begin{aligned}
&\left\| [\mathbf{a}_{s,+}^\top, \pm \mathbf{a}_{s,-}^\top] [\mathbf{g}_{s,+}^\top(\mathbf{x}), \mathbf{g}_{s,-}^\top(\mathbf{x})]^\top - \left((\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2+1} \pm (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2-1} \right) \right\|_{L_2} \\
&\leq \left\| \mathbf{a}_{s,+}^\top \mathbf{g}_{s,+}(\mathbf{x}) - (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} + \left\| \mathbf{a}_{s,-}^\top \mathbf{g}_{s,-}(\mathbf{x}) - (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2-1} \right\|_{L_2} \\
&\leq \epsilon/r_*.
\end{aligned}$$

The corresponding neural representation dimension is $D_s \geq \frac{50^2 r_*^3 (p_s+1)^5 \|\boldsymbol{\beta}_s\|_2^{p_s+1} + (p_s-1)^5 \|\boldsymbol{\beta}_s\|_2^{p_s-1}}{\epsilon^2 \delta}$. Combining the even and odd degrees together, we can choose

$$D \geq \frac{4 \times 50^2 r_*^3}{\epsilon^2 \delta} \left(\sum_{s=1}^{r_*} p_s^5 \|\boldsymbol{\beta}_s\|_2^{p_s} \mathbb{1}\{p_s \text{ is even}\} + 2(p_s+1)^5 \|\boldsymbol{\beta}_s\|_2^{p_s+1} \mathbb{1}\{p_s \text{ is odd}\} \right).$$

Lemma 9 now follows from Lemma 8 by merging $\mathbf{g}_{s,+}$, $\mathbf{g}_{s,-}$ as a single feature \mathbf{g}_s , and $\mathbf{a}_{s,+}$, $\mathbf{a}_{s,-}$ as a single weight vector \mathbf{a}_s so that $\mathbf{w}_r^{s,*}$ can be chosen accordingly. Unifying the notation for even and odd degree polynomials, we have

$$D \geq \frac{8 \times 50^2 r_*^3}{\epsilon^2 \delta} \left(\sum_{s=1}^{r_*} (p_s+1)^5 \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil} \right).$$

□

D.3 Generalization of Quad-h

Lemma 8 and 9 construct a proper weight matrix $\mathbf{W}^* = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top$ such that Quad-Neural can well approximate f_* of form (12) in the L_2 sense. Now we show that for sufficiently large m , the empirical risk $\widehat{\mathcal{R}}_\ell(f_{\mathbf{W}^*})$ is comparable to that of f_* . By the Lipschitz property of the loss function, we derive

$$\begin{aligned}
\frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{W}^*}(\mathbf{x}_i), y_i) &\leq \frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{W}^*}(\mathbf{x}_i), y_i) - \ell(f_*(\mathbf{x}_i), y_i) + \ell(f_0(\mathbf{x}_i), y_i) \\
&\leq \frac{1}{n} \sum_{i=1}^n |f_{\mathbf{W}^*}(\mathbf{x}_i) - f_*(\mathbf{x}_i)| + \widehat{\mathcal{R}}(f_*).
\end{aligned}$$

For a given $\epsilon_0 > 0$, using Chebyshev's inequality, we have

$$\begin{aligned}
\mathbb{P}\left(|\widehat{\mathcal{R}}(f_\star) - \mathcal{R}(f_\star)| \geq \epsilon_0/2\right) &\leq \frac{4\mathbb{E}[(\widehat{\mathcal{R}}(f_\star) - \mathcal{R}(f_\star))^2]}{\epsilon_0^2} \\
&= \frac{4\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\ell(f_\star(\mathbf{x}), y) - \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\ell(f_\star(\mathbf{x}), y)]]^2}{n_2\epsilon_0^2} \\
&\leq \frac{4\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\ell(0, y) + |f_\star(\mathbf{x})| - \text{OPT}]^2}{n\epsilon_0^2} \\
&\leq \frac{8\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[|f_\star(\mathbf{x})|]^2 + 8(1 + \text{OPT})^2}{n\epsilon_0^2} \\
&\leq \frac{8(1 + \text{OPT})^2 + 8\mathbb{E}_{\mathbf{x}}\left[\sum_{s=1}^{r_\star} \alpha_s^2 \sum_{s=1}^{r_\star} (\boldsymbol{\beta}_s^\top \mathbf{x})^{2p_s}\right]}{n\epsilon_0^2} \\
&\leq \frac{8(1 + \text{OPT})^2 + 8r_\star^2}{n\epsilon_0^2}.
\end{aligned}$$

Choosing $n \geq \frac{8(1+\text{OPT})^2+8r_\star^2}{\delta\epsilon_0^2}$, $\widehat{\mathcal{R}}(f_\star) - \mathcal{R}(f_\star) \leq \epsilon_0/2$ holds with probability at least $1 - \delta$. We further invoke Lemma 8 and Chebyshev's inequality again on $\frac{1}{n} \sum_{i=1}^n |f_{\mathbf{W}^*}(\mathbf{x}_i) - f_\star(\mathbf{x}_i)|$:

$$\begin{aligned}
\mathbb{P}_{\mathbf{x}}\left(\frac{1}{n} \sum_{i=1}^n |f_{\mathbf{W}^*}(\mathbf{x}_i) - f_\star(\mathbf{x}_i)| \geq \epsilon_0/2\right) &\leq \frac{4\mathbb{E}_{\mathbf{x}}\left[\frac{1}{n^2} \left(\sum_{i=1}^n |f_{\mathbf{W}^*}(\mathbf{x}_i) - f_\star(\mathbf{x}_i)|\right)^2\right]}{\epsilon_0^2} \\
&\leq \frac{\frac{4}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{x}}\left[|f_{\mathbf{W}^*}(\mathbf{x}_i) - f_\star(\mathbf{x}_i)|^2\right]}{\epsilon_0^2} \\
&\leq \frac{196\epsilon^2 r_\star^2}{\epsilon_0^2},
\end{aligned}$$

where the last inequality holds with probability at least $1 - \delta$. We set $\frac{196r_\star^2\epsilon^2}{\epsilon_0^2} \leq \delta$, which implies $\epsilon^2 \leq \frac{\delta\epsilon_0^2}{196r_\star^2}$. Accordingly, the number of neurons in the top layer needs to be at least

$$m \geq \frac{10584r_\star^3 D \left(1 + \log \frac{8}{\delta}\right) \sum_{s=1}^{r_\star} \left\|(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}\right\|_{L_2}}{\delta\epsilon_0^2} \log \frac{7r_\star}{\sqrt{\delta}\epsilon_0},$$

and the dimension of the neural representation is

$$D = \frac{50^2 \times 392r_\star^5}{\delta\epsilon_0^2} \left(\sum_{s=1}^{r_\star} (p_s + 1)^5 \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil}\right). \quad (14)$$

This gives us that with probability at least $1 - 3\delta$ over the randomness of data and initialization⁷, the empirical risk satisfies

$$\frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{W}^*}(\mathbf{x}_i), y_i) \leq \text{OPT} + \epsilon_0.$$

Applying Theorem 1 part (2), for any second-order stationary point $\widehat{\mathbf{W}}$ and proper regularization parameter λ , we have

$$\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau)(\text{OPT} + \epsilon_0) + \epsilon_0 \leq (1 + \tau_0)\text{OPT} + \epsilon_0.$$

Bounding $B_{w,\star}$.

⁷To achieve probability $1 - \delta$, we replace δ with $\delta/3$, which only introduce a multiplicative constant in the size of m and D .

Towards establishing the generalization bound of $f_{\widehat{\mathbf{W}}}$, we first find $B_{w,\star}$:

$$\begin{aligned} \sum_{r=1}^m \|\mathbf{w}_r^*\|_2^4 &= \sum_{s=1}^{r_\star} \sum_{r \in \mathcal{I}_1 \cup \mathcal{I}_2} \|\mathbf{w}_r^{s,\star}\|_2^4 = 48r_\star \sum_{s=1}^{r_\star} \sum_{r=1}^{m_0} \alpha_s^2 m_0^{-1} \left\| \widehat{\Sigma}^{1/2} [\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top \right\|_2^4 \\ &= 48r_\star \sum_{s=1}^{r_\star} \alpha_s^2 \left\| \widehat{\Sigma}^{1/2} [\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top \right\|_2^4. \end{aligned}$$

To bound $\left\| \widehat{\Sigma}^{1/2} [\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top \right\|_2$, we first replace $\widehat{\Sigma}$ with Σ . We denote $\boldsymbol{\theta}_s = \Sigma^{1/2} [\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top$, and observe $\boldsymbol{\theta}_s$ is the optimal solution to the following least square problem

$$\boldsymbol{\theta}_s = \underset{\mathbf{u}_1}{\operatorname{argmin}} \left\| F(\mathbf{x}) - \mathbf{u}_1^\top \Sigma^{-1/2} \mathbf{g}(\mathbf{x}) \right\|_{L_2}^2 \quad \text{with} \quad F(\mathbf{x}) = \mathbf{a}_s^\top \mathbf{g}_s^\top(\mathbf{x}).$$

The optimal solution is $\boldsymbol{\theta}_s = \mathbf{u}_1^* = \Sigma^{-1/2} \mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})]$. Similarly, the optimal solution to the following least square problem

$$\min_{\mathbf{u}_2} \left\| F(\mathbf{x}) - \mathbf{u}_2^\top \mathbf{g}(\mathbf{x}) \right\|_{L_2}^2 \quad \text{with} \quad F(\mathbf{x}) = \mathbf{a}_s^\top \mathbf{g}_s(\mathbf{x})$$

is $\mathbf{u}_2^* = \Sigma^{-1} \mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})]$. The residual of \mathbf{u}_2^* is

$$\left\| F(\mathbf{x}) - (\mathbf{u}_2^*)^\top \mathbf{g}(\mathbf{x}) \right\|_{L_2}^2 = \|F(\mathbf{x})\|_{L_2}^2 - \mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})^\top] \Sigma^{-1} \mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})] \geq 0.$$

This implies

$$\|\boldsymbol{\theta}_s\|_2 = \sqrt{\mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})^\top] \Sigma^{-1} \mathbb{E}_{\mathbf{x}}[F(\mathbf{x})\mathbf{g}(\mathbf{x})]} \leq \|F(\mathbf{x})\|_{L_2} \leq \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2} + \epsilon,$$

where the last inequality follows from Lemma 5. This gives rise to

$$\sum_{s=1}^{r_\star} \alpha_s^2 \|\boldsymbol{\theta}_s\|_2^4 \leq \sum_{s=1}^{r_\star} \alpha_s^2 \left\| (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s/2} \right\|_{L_2}^4 \leq r_\star.$$

To switch back to $\widehat{\Sigma}$, we invoke Lemma 3 on the concentration of $\widehat{\Sigma}$ to Σ . Specifically, with probability at least $1 - \delta$, choosing $n_0 \geq 4c\delta^{-2} \lambda_{[p/2]}^{-1} D \log D$ for some constant c , we have

$$\frac{1}{2} \Sigma \preceq \widehat{\Sigma} \preceq \frac{3}{2} \Sigma.$$

Consequently, by denoting $\widehat{\boldsymbol{\theta}}_s = \widehat{\Sigma}^{1/2} [\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top$, we have

$$\left\| \widehat{\boldsymbol{\theta}}_s \right\|_2^2 = \left\| \widehat{\Sigma}^{1/2} \Sigma^{-1/2} \boldsymbol{\theta}_s \right\|_2^2 = \boldsymbol{\theta}_s^\top \Sigma^{-1/2} \widehat{\Sigma} \Sigma^{-1/2} \boldsymbol{\theta}_s \leq \left\| \Sigma^{-1/2} \widehat{\Sigma} \Sigma^{-1/2} \right\|_{\text{op}} \|\boldsymbol{\theta}_s\|_2^2 \leq \frac{3}{2} \|\boldsymbol{\theta}_s\|_2^2.$$

Plugging into $\sum_{r=1}^m \|\mathbf{w}_r^*\|_2^4$, we have

$$\sum_{r=1}^m \|\mathbf{w}_r^*\|_2^4 \leq 108r_\star \sum_{s=1}^{r_\star} \alpha_s^2 \|\boldsymbol{\theta}_s\|_2^4 \leq 108r_\star^2.$$

Therefore, we can set $B_{w,\star}^4 = 108r_\star^2$. Note that $B_{w,\star}$ is independent of the width m .

Bounding $M_{h,\text{op}}$ and B_h .

The remaining ingredients are $M_{h,\text{op}}$ and $\|\mathbf{h}(\mathbf{x})\|_2$. Conditioned on the event $\frac{1}{2} \Sigma \leq \widehat{\Sigma} \leq \frac{3}{2} \Sigma$, we know $\widehat{\Sigma}^{-1} \leq 2\Sigma^{-1}$. Therefore, we have

$$\|\mathbf{h}(\mathbf{x})\|_2^2 = \mathbf{g}(\mathbf{x})^\top \widehat{\Sigma}^{-1} \mathbf{g}(\mathbf{x}) \leq 2\mathbf{g}(\mathbf{x})^\top \Sigma^{-1} \mathbf{g}(\mathbf{x}) \leq 2\lambda_{[p/2]}^{-1} D.$$

Note that the norm of $\mathbf{h}(\mathbf{x})$ is in the order of \sqrt{D} according to Assumption 2.

Lastly, we bound $M_{h,\text{op}}$ as

$$\begin{aligned}
B_h^2 M_{h,\text{op}}^2 &= \mathbb{E}_{\mathbf{x}} \left[\left\| \frac{1}{n_2} \sum_{i=n_1+1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right] \\
&= \mathbb{E}_{\mathbf{x}} \left[\left\| \widehat{\Sigma}^{-1/2} \left(\frac{1}{n_2} \sum_{i=n_1+1}^n \mathbf{g}(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)^\top \right) \widehat{\Sigma}^{-1/2} \right\|_{\text{op}} \right] \\
&\leq \mathbb{E}_{\mathbf{x}} \left[\left\| \widehat{\Sigma}^{-1/2} \Sigma^{1/2} \Sigma^{-1/2} \left(\frac{1}{n_2} \sum_{i=n_1+1}^n \mathbf{g}(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)^\top \right) \Sigma^{-1/2} \Sigma^{1/2} \widehat{\Sigma}^{-1/2} \right\|_{\text{op}} \right] \\
&\leq \left\| \widehat{\Sigma}^{-1/2} \Sigma^{1/2} \right\|_{\text{op}} \mathbb{E}_{\mathbf{x}} \left[\left\| \Sigma^{-1/2} \left(\frac{1}{n_2} \sum_{i=n_1+1}^n \mathbf{g}(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)^\top \right) \Sigma^{-1/2} \right\|_{\text{op}} \right] \left\| \Sigma^{1/2} \widehat{\Sigma}^{-1/2} \right\|_{\text{op}} \\
&\leq \frac{3}{2} \left\| \widehat{\Sigma}^{-1/2} \Sigma^{1/2} \right\|_{\text{op}}^2.
\end{aligned}$$

The last inequality holds, due to Lemma 3 and $\widehat{\Sigma}$ is obtained using independent samples. Conditioned on the same event $\frac{1}{2}\Sigma \leq \widehat{\Sigma} \leq \frac{3}{2}\Sigma$, we have

$$\left\| \widehat{\Sigma}^{-1/2} \Sigma^{1/2} \right\|_{\text{op}}^2 = \left\| \Sigma^{1/2} \widehat{\Sigma}^{-1} \Sigma^{1/2} \right\|_{\text{op}} \leq 2.$$

Therefore, $M_{h,\text{op}}^2 \leq 3B_h^{-2}$. Putting all the ingredients together and applying Theorem 1, by choosing

$$m \geq \max \left\{ \frac{10584r_*^3 D (1 + \log \frac{8}{\delta}) \sum_{s=1}^{r_*} \|(\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s}\|_{L_2}}{\delta \epsilon_0^2} \log \frac{7r_*}{\sqrt{\delta} \epsilon_0}, \frac{108C^2 D^2 r_*^2}{\epsilon_0 \sqrt{2\lambda_0}} \right\},$$

we establish for any SOSP $\widehat{\mathbf{W}}$, the generalization error bounded by:

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\left| \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \right| \right] \leq \tilde{O} \left(\frac{\|\mathbf{h}(\mathbf{x})\|_2^2 B_{w,*}^2 M_{h,\text{op}}}{\sqrt{n}} \right) = \tilde{O} \left(\sqrt{\frac{2\lambda_{[p/2]}^{-1} D r_*^2}{n}} \right).$$

Using Markov's inequality, we have

$$\begin{aligned}
\mathbb{P} \left(\left| \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \right| \geq \epsilon_0 \right) &\leq \frac{\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\left| \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \right| \right]}{\epsilon_0} \\
&\leq \tilde{O} \left(\epsilon_0^{-1} \sqrt{\frac{2\lambda_{[p/2]}^{-1} D r_*^2}{n}} \right).
\end{aligned}$$

We set the above probability upper bounded by δ , which requires

$$n = \tilde{O} \left(\frac{\lambda_{[p/2]}^{-1} r_*^7}{\epsilon_0^4 \delta^3} \left(\sum_{s=1}^{r_*} (p_s + 1)^5 \|\boldsymbol{\beta}_s\|_2^{2[p_s/2]} \right) \right).$$

We can now bound $\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q)$ as

$$\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) = \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) + \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau_0) \text{OPT} + 2\epsilon_0,$$

which holds with probability at least $1 - \delta$.

Taking $\|\boldsymbol{\beta}_s\|_2 = \sqrt{\delta}$, the sample size n grows in the order of $\tilde{O} \left(\frac{d^{[p/2]} \lambda_{[p/2]}^{-1} r_*^8 p^5}{\epsilon_0^4 \delta^3} \right)$. On the other hand, estimating covariance matrix Σ requires $n_0 = \tilde{O} \left(4\delta^{-2} \lambda_{[p/2]}^{-1} D \log D \right)$ samples, which is in the order of $\tilde{O} \left(\frac{d^{[p/2]} \lambda_{[p/2]}^{-1} r_*^6 p^5}{\epsilon_0^2 \delta^2} \right)$. Adding n_1, n_2 together, the sample complexity n grows in the order of $\tilde{O} \left(\frac{d^{[p/2]}}{\epsilon_0^4} \text{poly}(r_*, p, \delta^{-1}) \right)$.

D.4 Proof of Claim 1

Proof. To show $\sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\}$ is well concentrated, we observe that by symmetry, the following holds

$$\begin{aligned} \sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} &= \sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) / \|\mathbf{h}(\mathbf{x})\|_2 \geq 0\} \\ &\leq \sup_{\mathbf{y} \in \mathbb{S}^{D-1}} \frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\}. \end{aligned}$$

For a given \mathbf{y} , each $2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\}$ is bounded in $[0, 2]$, hence it is sub-Gaussian with variance proxy 1. Using the Hoeffding's inequality, for every \mathbf{y} , we have

$$\mathbb{P}\left(\left|\frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\} - 1\right| \geq t\right) \leq 2 \exp(-m_0 t^2 / 2).$$

To bound the supremum, we discretize the unit sphere. Let $\{\bar{\mathbf{y}}_i\}_{i=1}^{\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2)}$ be a γ -covering of \mathbb{S}^{D-1} with $\gamma < 1$, where $\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2)$ denotes the covering number. By the volume ratio argument, we bound $\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2) \leq \left(\frac{3}{\gamma}\right)^D$. Applying the union bound, we derive

$$\begin{aligned} \mathbb{P}\left(\max_{\mathbf{y} \in \{\bar{\mathbf{y}}_i\}_{i=1}^{\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2)}} \left|\frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\} - 1\right| \geq t\right) &\leq 2\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2) \exp(-m_0 t^2 / 2) \\ &\leq 2 \exp\left(-m_0 t^2 + D \log \frac{3}{\gamma}\right). \end{aligned}$$

Taking $t = \sqrt{\frac{D \log \frac{3}{\gamma} (1 + \frac{1}{D} \log \frac{2}{\delta})}{m_0}}$, with probability at least $1 - \delta$, we have

$$\max_{\mathbf{y} \in \{\bar{\mathbf{y}}_i\}_{i=1}^{\mathcal{N}(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2)}} \left|\frac{1}{m_0} \sum_{r=1}^{m_0} 2\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\} - 1\right| \leq \sqrt{\frac{D \log \frac{3}{\gamma} (1 + \frac{1}{D} \log \frac{2}{\delta})}{m_0}}. \quad (15)$$

By the definition of γ -covering, for any given $\mathbf{y} \in \mathbb{S}^{D-1}$, there exists $\bar{\mathbf{y}}$ such that $\|\bar{\mathbf{y}} - \mathbf{y}\|_2 \leq \gamma$. We evaluate how many pairs $\mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\}, \mathbb{1}\{\mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \geq 0\}$ taking different values, which is equivalent to $(\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}})$ having opposite signs. Using the Hoeffding's inequality again, with probability at least $1 - \delta$, we have

$$\begin{aligned} &\left|\frac{1}{m_0} \sum_{r=1}^{m_0} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}\}\right. \\ &\quad \left. - \mathbb{E}\left[\frac{1}{m_0} \sum_{r=1}^{m_0} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}\}\right]\right| \leq \sqrt{\frac{\log(2/\delta)}{2m_0}}. \end{aligned}$$

To bound the expectation, we observe that $(\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}})$ is jointly Gaussian with zero mean and the covariance matrix

$$\begin{pmatrix} 1 & \mathbf{y}^\top \bar{\mathbf{y}} \\ \mathbf{y}^\top \bar{\mathbf{y}} & 1 \end{pmatrix}.$$

Therefore, we find the following probability

$$\begin{aligned} \mathbb{P}(\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ opposite signs}) &= 2\mathbb{P}(\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \leq 0) \\ &= 2 \int_0^\infty \int_{-\infty}^0 \frac{1}{2\pi} (1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2)^{-1/2} \exp\left(-\frac{u^2 - (\mathbf{y}^\top \bar{\mathbf{y}})uv + v^2}{2(1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2)}\right) dudv \\ &\stackrel{(i)}{\leq} 2 \int_0^\infty \int_{-\infty}^0 \frac{1}{2\pi} (1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2)^{-1/2} \exp\left(-\frac{u^2 + v^2}{2(1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2)}\right) dudv \\ &= 2 \int_0^\infty \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{v^2}{2(1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2)}\right) dv \\ &= \sqrt{1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2}, \end{aligned}$$

where inequality (i) holds since $uv < 0$. We further bound $1 - (\mathbf{y}^\top \bar{\mathbf{y}})^2 = 1 - (1 + \mathbf{y}^\top (\bar{\mathbf{y}} - \mathbf{y}))^2 \leq 1 - (1 - \gamma)^2 \leq 2\gamma$. Consequently, we deduce $\mathbb{P}(\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}) \leq \sqrt{2\gamma}$. Taking $\gamma = m_0^{-1} \log 1/\delta$, we have

$$\mathbb{E} \left[\frac{1}{m_0} \sum_{r=1}^{m_0} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}\} \right] \leq \sqrt{\frac{2 \log(2/\delta)}{m_0}}.$$

This implies with probability at least $1 - \delta$,

$$\sqrt{\frac{\log(2/\delta)}{2m_0}} \leq \frac{1}{m_0} \sum_{r=1}^{m_0} \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}\} \leq \sqrt{\frac{9 \log(2/\delta)}{2m_0}}. \quad (16)$$

Combining (15) and (16) together, with probability at least $1 - 2\delta$, we deduce

$$\begin{aligned} & \sup_{\mathbf{y} \in \mathbb{S}^{d-1}} \left| \frac{1}{m_0} \sum_{r=1}^{m_0} 2 \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\} - 1 \right| \\ & \leq \sup_{\|\mathbf{y} - \bar{\mathbf{y}}\|_2 \leq m_0^{-1} \log 1/\delta} \left| \frac{1}{m_0} \sum_{r=1}^{m_0} 2 \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y}, \mathbf{w}_{0,r}^\top \bar{\mathbf{y}} \text{ having opposite signs}\} \right| \\ & \quad + \max_{\mathbf{y} \in \{\bar{\mathbf{y}}_i\}_{i=1}^{N(\gamma, \mathbb{S}^{D-1}, \|\cdot\|_2)}} \left| \frac{1}{m_0} \sum_{r=1}^{m_0} 2 \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{y} \geq 0\} - 1 \right| \\ & \leq \sqrt{\frac{9 \log \frac{2}{\delta}}{2m_0}} + \sqrt{\frac{D \log(3m_0) (1 + \frac{1}{D} \log \frac{2}{\delta})}{m_0}} \\ & \leq 6 \sqrt{\frac{D \log(3m_0) (1 + \log \frac{2}{\delta})}{m_0}}. \end{aligned}$$

As a result, we know

$$\sup_{\mathbf{x} \in \mathbb{S}^{d-1}} \left| \frac{1}{m_0} \sum_{r=1}^{m_0} 2 \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}) \geq 0\} - 1 \right| \leq 6 \sqrt{\frac{D \log(3m_0) (1 + \log \frac{2}{\delta})}{m_0}}$$

holds with probability at least $1 - 2\delta$. \square

D.5 Learning in Quad-Neural with Data Dependent Regularizer

We consider using data dependent regularizer for learning with unwhitened features $\mathbf{g}(\mathbf{x})$, which also yields improved sample complexity. The full learning algorithm is described Algorithm 2.

Note that Quad-g shares the same QuadNTK model as Quad-h, and only replaces the neural representation \mathbf{h} with \mathbf{g} . The superscript on $\widehat{\mathcal{R}}_\lambda^{\text{dreg}}$ stands for data dependent regularization. We show Quad-g enjoys a similarly nice optimization landscape and good generalization properties as Quad-h.

Theorem 1' (Optimization landscape and generalization of Quad-g). Suppose Assumption 2 holds.

(1) (Optimization) Given any $\epsilon > 0$ and $\delta > 0$, $\tau = \Theta(1)$, and some radius $B_{w,*} > 0$, suppose the width $m \geq \tilde{O}(D^2 B_{w,*}^4 \epsilon^{-1})$, sample size $n_0 = \tilde{O}(\delta^{-2} D)$, and we choose a proper regularization coefficient $\lambda > 0$. Then with probability $1 - \delta$ over $\tilde{\mathcal{S}}_{n_0}$, any second-order stationary point $\widehat{\mathbf{W}}$ of the regularized risk $\widehat{\mathcal{R}}_\lambda^{\text{dreg}}(f_{\widehat{\mathbf{W}}}^Q)$ satisfies $\|\widehat{\mathbf{W}} \widehat{\Sigma}^{1/2}\|_{2,4} \leq O(B_{w,*})$, and achieves

$$\widehat{\mathcal{R}}_\lambda^{\text{dreg}}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau) \min_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_{w,*}} \widehat{\mathcal{R}}(f_{\mathbf{w}}^Q) + \epsilon.$$

(2) (Generalization) For any radius $B_w > 0$, we have with high probability (over $(\mathbf{a}, \mathbf{W}_0, \tilde{\mathcal{S}}_{n_0})$) that

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\mathbf{w}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{w}}^Q) \right| \right] \leq \tilde{O} \left(\frac{B_g^2 B_w^2 M_{g,\text{op}}}{\sqrt{n}} + \frac{1}{\sqrt{n}} \right),$$

Algorithm 2 Learning with Unwhitened Neural Random Features

Input: Labeled data S_n , unlabeled data \tilde{S}_{n_0} , initializations $\mathbf{V} \in \mathbb{R}^{D \times d}$, $\mathbf{b} \in \mathbb{R}^D$, $\mathbf{W}_0 \in \mathbb{R}^{m \times D}$, parameters (λ, ϵ) .

Step 1: 1) Construct model $f_{\mathbf{W}}^Q$ as

$$\text{(Quad-g)} \quad f_{\mathbf{W}}^Q(\mathbf{x}) = \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{g}(\mathbf{x})) (\mathbf{w}_r^\top \mathbf{g}(\mathbf{x}))^2,$$

where $\mathbf{g}(\mathbf{x}) = [\mathbb{1}\{\mathbf{v}_1^\top \mathbf{x} + b_1 \geq 0, \dots, \mathbf{v}_D^\top \mathbf{x} + b_D \geq 0\}]^\top$ is the neural random features.

2) Use \tilde{S}_{n_0} to estimate the covariance matrix of $\mathbf{g}(\mathbf{x})$, i.e., $\hat{\Sigma} = \frac{1}{n_0} \sum_{i=1}^{n_0} \mathbf{g}(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)^\top$.

Step 2: Find a second-order stationary point $\hat{\mathbf{W}}$ of the data dependent regularized empirical risk (on the data S_n):

$$\hat{\mathcal{R}}_\lambda^{\text{dreg}}(f_{\mathbf{W}}^Q) = \frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{W}}^Q(\mathbf{x}_i), y_i) + \lambda \left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4}^4.$$

where $M_{g,\text{op}}^2 = B_g^{-2} \mathbb{E}_{\mathbf{x}} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right]$.

Proof of Theorem 1', Optimization Part. We recall the second-order directional derivative of $\hat{\mathcal{R}}(f_{\mathbf{W}}^Q)$ satisfies

$$\nabla_{\mathbf{W}}^2 \hat{\mathcal{R}}(f_{\mathbf{W}}^Q)[\mathbf{W}_*, \mathbf{W}_*] \leq \left\langle \nabla \hat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle - 2(\hat{\mathcal{R}}(f_{\mathbf{W}}^Q) - \hat{\mathcal{R}}(f_{\mathbf{W}^*}^Q)) + m^{-1} B_g^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_*\|_{2,4}^2,$$

which is established in (9) and \mathbf{W}^* is any given matrix. Note we have replaced B_h with $B_g = \|\mathbf{g}(\mathbf{x})\|_2$, and B_g is upper bounded by \sqrt{D} . Similar to the proof A.1, we specialize \mathbf{W}^* to be the optimizer $\mathbf{W}^* = \operatorname{argmin}_{\|\mathbf{W} \hat{\Sigma}^{1/2}\|_{2,4} \leq B_{w,*}} \hat{\mathcal{R}}(f_{\mathbf{W}}^Q)$ and denote its risk $\hat{\mathcal{R}}(f_{\mathbf{W}^*}^Q) = M$. We choose the regularization coefficient as

$$\lambda = \lambda_0 B_{w,*}^{-4},$$

where λ_0 is to be determined. We argue that any second-order stationary point $\hat{\mathbf{W}}$ has to satisfy $\left\| \hat{\mathbf{W}} \hat{\Sigma}^{1/2} \right\|_{2,4} = O(B_{w,*})$. We already know from proof A.1 that for any \mathbf{W} , $\left\langle \nabla \hat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle \geq -2$ holds.

Combining with the fact

$$\left\langle \nabla_{\mathbf{W}} \left(\left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4}^4 \right), \mathbf{W} \right\rangle = 4 \left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4}^4,$$

we have simultaneously for all \mathbf{W} that

$$\begin{aligned} \left\langle \nabla \hat{\mathcal{R}}_\lambda^{\text{data}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle &\geq \left\langle \nabla_{\mathbf{W}} (\lambda \left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4}^4), \mathbf{W} \right\rangle + \left\langle \nabla_{\mathbf{W}} \hat{\mathcal{R}}(f_{\mathbf{W}}^Q), \mathbf{W} \right\rangle \\ &\geq 4\lambda \left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4}^4 - 2. \end{aligned}$$

Therefore we see that any stationary point \mathbf{W} has to satisfy

$$\left\| \mathbf{W} \hat{\Sigma}^{1/2} \right\|_{2,4} \leq (2\lambda)^{-1/4}.$$

Choosing

$$\lambda_0 = \frac{1}{36} (2\tau M + \epsilon),$$

we get $36\lambda B_{w,\star}^4 = 2\tau M + \epsilon$. The second-order directional derivative of $\widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q)$ along direction \mathbf{W}_\star is upper bounded by

$$\begin{aligned}
\nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q)[\mathbf{W}_\star, \mathbf{W}_\star] &= \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q)[\mathbf{W}_\star, \mathbf{W}_\star] + \lambda \nabla_{\mathbf{W}}^2 \left\| \mathbf{W} \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 [\mathbf{W}_\star, \mathbf{W}_\star] \\
&= \nabla_{\mathbf{W}}^2 \widehat{\mathcal{R}}^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q)[\mathbf{W}_\star, \mathbf{W}_\star] + 4\lambda \sum_{r \leq m} \left(\mathbf{w}_{\star,r} \widehat{\Sigma} \mathbf{w}_{\star,r} \right) \left(\mathbf{w}_r \widehat{\Sigma} \mathbf{w}_r \right) \\
&\quad + 8\lambda \sum_{r \leq m} \left\langle \mathbf{w}_r \widehat{\Sigma}^{1/2}, \mathbf{w}_{\star,r} \widehat{\Sigma}^{1/2} \right\rangle^2 \\
&\leq \left\langle \nabla \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) - M) + m^{-1} B_g^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_\star\|_{2,4}^2 \\
&\quad + 12\lambda \left\| \mathbf{W} \widehat{\Sigma}^{1/2} \right\|_{2,4}^2 \left\| \mathbf{W}_\star \widehat{\Sigma}^{1/2} \right\|_{2,4}^2 \\
&\stackrel{(i)}{\leq} \left\langle \nabla \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) - M) + m^{-1} B_g^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_\star\|_{2,4}^2 \\
&\quad + \lambda \left\| \mathbf{W} \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 + 36\lambda \left\| \mathbf{W}_\star \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 \\
&\leq \left\langle \nabla \widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q), \mathbf{W} \right\rangle - 2(\widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q) - M) \\
&\quad + m^{-1} B_g^4 \|\mathbf{W}\|_{2,4}^2 \|\mathbf{W}_\star\|_{2,4}^2 - \lambda \left\| \mathbf{W} \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 + 36\lambda \|\mathbf{W}_\star\|_{2,4}^4.
\end{aligned}$$

We used the fact $12ab \leq a^2 + 36b^2$. For a second order-stationary point $\widehat{\mathbf{W}}$ of $\widehat{\mathcal{R}}_\lambda(f_{\widehat{\mathbf{W}}}^Q)$, its gradient vanishes and the Hessian is positive definite. Therefore, we have

$$0 \leq -2(\widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q) - M) + m^{-1} B_g^4 \left\| \widehat{\mathbf{W}} \right\|_{2,4}^2 \|\mathbf{W}_\star\|_{2,4}^2 - \lambda \left\| \widehat{\mathbf{W}} \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 + 36\lambda \left\| \mathbf{W}_\star \widehat{\Sigma}^{1/2} \right\|_{2,4}^4.$$

By Assumption 2, we have $\lambda_{\min}(\Sigma) \geq \lambda_k$. Moreover, by Lemma 3, when $n_0 = O(\delta^{-2} D \log D)$, with probability at least $1 - \delta$, we have the following relative concentration of $\widehat{\Sigma}$:

$$\frac{1}{2} \Sigma \preceq \widehat{\Sigma} \preceq \frac{3}{2} \Sigma.$$

Combining these two ingredients together, we deduce

$$\begin{aligned}
\|\mathbf{W}\|_{2,4}^4 &= \sum_{r=1}^m \left\| \widehat{\Sigma}^{-1/2} \widehat{\Sigma}^{1/2} \mathbf{w}_r \right\|_2^4 \leq \sum_{r=1}^m \left\| \widehat{\Sigma}^{-1/2} \right\|_{\text{op}}^4 \left\| \widehat{\Sigma}^{1/2} \mathbf{w}_r \right\|_2^4 \\
&\leq \sum_{r=1}^m 4\lambda_k^{-2} \left\| \widehat{\Sigma}^{1/2} \mathbf{w}_r \right\|_2^4 = 4\lambda_k^{-2} \left\| \mathbf{W} \widehat{\Sigma} \right\|_{2,4}^4.
\end{aligned}$$

Exactly the same argument yields $\|\mathbf{W}_\star\|_{2,4}^4 \leq 4\lambda_k^{-2} \left\| \mathbf{W}_\star \widehat{\Sigma} \right\|_{2,4}^4$. Therefore, we choose $m = 4\epsilon^{-1} \lambda_k^{-2} (2\lambda_0)^{-1/2} C^2 B_g^4 B_{w,\star}^4 \geq \epsilon^{-1} C^2 B_g^4 \left\| \widehat{\mathbf{W}} \right\|_{2,4}^2 \|\mathbf{W}_\star\|_{2,4}^2$ and the above inequality implies

$$\begin{aligned}
2(\widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q) - M) &\leq 2\tau M + \epsilon + \epsilon \\
\implies \widehat{\mathcal{R}}_\lambda^{\text{data}}(f_{\widehat{\mathbf{W}}}^Q) &\leq (1 + \tau)M + \epsilon.
\end{aligned}$$

Plugging in the naive upper bound $\|\mathbf{g}(\mathbf{x})\|_2 \leq \sqrt{D}$ in B_g , the proof is complete. \square

Proof of Theorem 1', Generalization Part. Built upon the proof A.2, we have

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\sup_{\|\mathbf{W} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \right| \right] \leq 2\mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\sup_{\|\mathbf{W} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \left| \frac{1}{n} \sum_{i=1}^n \xi_i \ell(f_{\widehat{\mathbf{W}}}^Q(\mathbf{x}_i), y_i) \right| \right],$$

where ξ is i.i.d. Rademacher random variables. Recall that the whitened feature is $\mathbf{h}(\mathbf{x}) = \widehat{\Sigma}^{-1/2} \mathbf{g}(\mathbf{x})$. We further have

$$\begin{aligned}
& \mathbb{E}_{(\mathbf{x}_i, y_i), \xi} \left[\left\| \sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \frac{1}{n} \sum_{i=1}^n \xi_i \ell(y_i, f_{\mathbf{w}}^Q(\mathbf{x}_i)) \right\| \right] \\
& \leq 4 \mathbb{E}_{\mathbf{x}_i, \xi} \left[\left\| \sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \frac{1}{\sqrt{m}} \sum_{r \leq m} \left\langle \frac{1}{n} \sum_{i=1}^n \xi_i a_r \phi''(\mathbf{w}_{0,r}^\top \mathbf{g}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top, \widehat{\Sigma}^{1/2} \mathbf{w}_r \mathbf{w}_r^\top \widehat{\Sigma}^{1/2} \right\rangle \right\| \right] \\
& \quad + \frac{2}{\sqrt{n}} \\
& \leq 4 \mathbb{E}_{\mathbf{x}_i, \xi} \left[\left\| \sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \max_{r \in [m]} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right\| \right] \\
& \quad \times \frac{1}{\sqrt{m}} \sum_{r \leq m} \left\| \widehat{\Sigma}^{1/2} \mathbf{w}_r \mathbf{w}_r^\top \widehat{\Sigma}^{1/2} \right\|_* + \frac{2}{\sqrt{n}} \\
& \leq 4 \mathbb{E}_{\mathbf{x}_i, \xi} \left[\left\| \max_{r \in [m]} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \phi''(\mathbf{w}_{0,r}^\top \mathbf{h}(\mathbf{x}_i)) \mathbf{h}(\mathbf{x}_i) \mathbf{h}(\mathbf{x}_i)^\top \right\|_{\text{op}} \right\| \right] \\
& \quad \times \underbrace{\sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \frac{1}{\sqrt{m}} \sum_{r \leq m} \left\| \widehat{\Sigma}^{1/2} \mathbf{w}_r \right\|_2^2}_{\leq B_w^2} + \frac{2}{\sqrt{n}},
\end{aligned}$$

Consequently, the generalization error is still bounded by

$$\mathbb{E}_{(\mathbf{x}_i, y_i)} \left[\left\| \sup_{\|\mathbf{w} \widehat{\Sigma}^{1/2}\|_{2,4} \leq B_w} \left| \mathcal{R}(f_{\mathbf{w}}^Q) - \widehat{\mathcal{R}}(f_{\mathbf{w}}^Q) \right| \right\| \right] \leq \tilde{O} \left(\frac{B_h^2 B_w^2 M_{h,\text{op}}}{\sqrt{n}} \sqrt{\log(Dm)} + \frac{1}{\sqrt{n}} \right).$$

□

When using Quad-g to learn low-rank polynomials in the form of

$$f_\star(\mathbf{x}) = \sum_{s=1}^{r_\star} \alpha_s (\boldsymbol{\beta}_s^\top \mathbf{x})^{p_s} \quad \text{defined in (5),}$$

we derive the following sample complexity bound.

Theorem 2' (Sample complexity of Quad-g). Suppose Assumption 2 holds, and there exists some f_\star that achieves low risk: $\mathcal{R}(f_\star) \leq \text{OPT}$. Then for any $\epsilon, \delta \in (0, 1)$ and $\tau = \Theta(1)$, choosing

$$D = \Theta \left(\text{poly}(r_\star, p) \sum_s \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil} \epsilon^{-2} \delta^{-1} \right), \quad m \geq \tilde{O}(\text{poly}(r_\star, D) \epsilon^{-2} \delta^{-1}), \quad (17)$$

$n_0 = \tilde{O}(D \delta^{-2})$, and a proper $\lambda > 0$, Algorithm 2 achieves the following guarantee: with probability at least $1 - \delta$ over the randomness of data and initialization, any second-order stationary point $\widehat{\mathbf{W}}$ of $\widehat{\mathcal{R}}_\lambda^{\text{dreg}}(f_{\widehat{\mathbf{W}}}^Q)$ satisfies

$$\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau) \text{OPT} + \epsilon + \tilde{O} \left(\sqrt{\frac{\text{poly}(r_\star, p, \delta^{-1}) \lambda_{\lceil p/2 \rceil}^{-1} \epsilon^{-2} \sum_{s=1}^{r_\star} \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil}}{n}} \right).$$

In particular, for any $\epsilon > 0$, we can achieve $\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau) \text{OPT} + 2\epsilon$ with sample complexity

$$n_0 + n \leq \tilde{O} \left(\text{poly}(r_\star, p, \lambda_{\lceil p/2 \rceil}^{-1}, \epsilon^{-1}, \delta^{-1}) \sum_{s=1}^{r_\star} \|\boldsymbol{\beta}_s\|_2^{2\lceil p_s/2 \rceil} \right). \quad (18)$$

Proof. The proof reproduces that for Quad- \mathbf{h} in Sections D.1, D.2, and D.3. Specifically, following the same argument in Lemma 9, we can establish the expressivity of Quad- \mathbf{h} , where for $r = 1, \dots, m_0$, we only need to choose

$$\mathbf{w}_r^{s,*} = \begin{cases} 2\sqrt{\alpha_s}(3r_*)^{1/4}m_0^{-1/4}[\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top, & \text{for the } s\text{-th group in } \mathcal{I}_1 \text{ with } \alpha_s > 0 \\ 2\sqrt{|\alpha_s|}(3r_*)^{1/4}m_0^{-1/4}[\mathbf{0}^\top, \dots, \mathbf{a}_s^\top, \dots, \mathbf{0}^\top]^\top, & \text{for the } s\text{-th group in } \mathcal{I}_2 \text{ with } \alpha_s < 0 \end{cases}$$

Remember $\mathcal{I}_1 = \{1, \dots, m/3\}$ where $a_r = 1$ for $r \in \mathcal{I}_1$ and $\mathcal{I}_2 = \{m/3 + 1, 2m/3\}$ with $a_r = -1$. Compared to using whitened representation $\mathbf{h}(\mathbf{x})$, we remove the multiplicative factor $\widehat{\Sigma}^{1/2}$ in \mathbf{w}_r^* (see Lemma 8). The corresponding representation dimension $D = \frac{8 \times 50^2 r_*^3 \sum_{s=1}^{r_*} p_s^5 \|\beta_s\|_2^{2\lceil p_s/2 \rceil}}{\epsilon^2 \delta}$ and the width $m \geq \frac{54r_* D(1 + \log \frac{8}{\delta})}{\epsilon^2} \log \frac{1}{\epsilon}$ remain unchanged. Then with probability $1 - \delta$, we have

$$\left\| \frac{1}{2\sqrt{m}} \sum_{r=1}^m a_r \mathbb{1}\{\mathbf{w}_{0,r}^\top \mathbf{g}(\mathbf{x}) \geq 0\} ((\mathbf{w}_r^*)^\top \mathbf{g}(\mathbf{x}))^2 - f(\mathbf{x}) \right\|_{L_2} \leq 7r_*\epsilon.$$

The rest of the proof follows Section D.3, where we need to upper bound $M_{g,\text{op}}$, $B_{w,*}$, and B_g , respectively. We use the naive upper bound on $B_g \leq \sqrt{D}$, since each entry of $\mathbf{g}(\mathbf{x})$ is bounded by 1. By definition, we have

$$B_g^2 M_{g,\text{op}}^2 = B_h^2 M_{h,\text{op}}^2 \leq 3.$$

Lastly, observe $B_{w,*}^4 = \left\| \mathbf{W}^* \widehat{\Sigma}^{1/2} \right\|_{2,4}^4 = \sum_{r=1}^m \left\| \widehat{\Sigma}^{1/2} \mathbf{w}_r^* \right\|_2^4$. An upper bound has been already derived in Section D.3, which is $108r_*^2$. As can be seen, quantities $M_{g,\text{op}}$, $B_{w,*}$, and B_g all retain the same order as using the whitened neural representation $\mathbf{h}(\mathbf{x})$ (with possibly different absolute constants). Therefore, in order to achieve

$$\mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) = \mathcal{R}(f_{\widehat{\mathbf{W}}}^Q) - \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) + \widehat{\mathcal{R}}(f_{\widehat{\mathbf{W}}}^Q) \leq (1 + \tau)\text{OPT} + 2\epsilon,$$

the sample complexity needs to satisfy

$$n = \widetilde{O} \left(\frac{\lambda_{\lceil p/2 \rceil}^{-1} r_*^7}{\epsilon_0^4 \delta^3} \left(\sum_{s=1}^{r_*} (p_s + 1)^5 \|\beta_s\|_2^{2\lceil p_s/2 \rceil} \right) \right),$$

and n_0 stays the same for the covariance estimation. This yields the same sample complexity (again with a potentially different absolute constant) as using the whitened representation $\mathbf{h}(\mathbf{x})$. \square