

Neural Hilbert Ladders: Multi-Layer Neural Networks in Function Space

Zhengdao Chen

Google Research

Mountain View, CA 94043

ZHENGDAO.C3@GMAIL.COM

Editor: Lorenzo Rosasco

Abstract

To characterize the function space explored by neural networks (NNs) is an important aspect of learning theory. In this work, noticing that a multi-layer NN generates implicitly a hierarchy of reproducing kernel Hilbert spaces (RKHSs)—named a neural Hilbert ladder (NHL)—we define the function space as an infinite union of RKHSs, which generalizes the existing Barron space theory of two-layer NNs. We then establish several theoretical properties of the new space. First, we prove a correspondence between functions expressed by L -layer NNs and those belonging to L -level NHLs. Second, we prove generalization guarantees for learning an NHL with a controlled complexity measure. Third, we derive a non-Markovian dynamics of random fields that governs the evolution of the NHL which is induced by the training of multi-layer NNs in an infinite-width mean-field limit. Fourth, we show examples of depth separation in NHLs under the ReLU activation function. Finally, we perform numerical experiments to illustrate the feature learning aspect of NN training through the lens of NHLs.

Keywords: function space of neural networks, feature learning, approximation bound, Rademacher complexity, depth separation, mean-field limit

1. Introduction

The empirical success of deep learning has inspired many efforts to understand neural networks (NNs) theoretically, from aspects including approximation, optimization and generalization. In the basic setting of supervised learning, NNs can be seen as parameterizing a particular family of functions on the input domain, from which a solution can be found through training. Then, to explain what is special about NNs, it is crucial to investigate the space of functions (a.k.a. hypothesis class) they represent.

As modern NNs often involve a huge number of parameters, it is worthy of understanding the space of functions that can be represented by NNs with *unlimited* width. As a foundational result, the universal approximation theorem (e.g., Cybenko, 1989; Hornik et al., 1989) shows that, given enough width, NNs are capable of approximating virtually all reasonable functions, which suggests the vastness of this space. A more interesting question, though, is to also find a *complexity measure of functions* that quantify their representation cost in terms of the rate of approximation error, which would yield insights on what kind of functions are more naturally represented by NNs. This question has been studied fruitfully in the literature for shallow (a.k.a. two-layer) NNs (Barron, 1993; Bengio et al., 2005; Bach, 2017a; E et al., 2019), but remains mostly open for multi-layer NNs.

Meanwhile, to study the sample complexity of learning NNs, prior works have proved generalization guarantees that are based on *not* the number of parameters but certain norms of them (e.g., Bartlett, 1998; Neyshabur et al., 2015), which could serve as a complexity measure of NNs from the *generalization* point of view. Then, an important question is whether there is a complexity measure associated with width-unlimited multi-layer NNs that unifies the two perspectives of approximation and generalization.

Another central aspect of deep learning is the training of NNs, which involves a non-convex optimization problem but can often be solved sufficiently well by variants of gradient descent (GD). Although remarkable progress has been made to prove optimization guarantees for various settings, it remains intriguing what kind of exploration in *function space* is induced by the training of NNs. The Neural Tangent Kernel (NTK) analysis provides a candidate theory via a linearized approximation of NN training (Jacot et al., 2018), which implies that NNs represent functions in a particular reproducing kernel Hilbert space (RKHS) determined by their initialization. However, the NTK theory is unable to model the *feature learning* that occurs in the training of actual NNs (Chizat et al., 2019; Woodworth et al., 2020), which is crucial to the success of deep learning.

Hence, this work is motivated by the following questions that are critical yet largely open:

- *How to characterize the function space explored by the training of multi-layer NNs with arbitrary width?*
- *Can we associate with it a complexity measure of functions that governs both approximation and generalization?*

In this work, by viewing an L -layer NN as a ladder of RKHSs with L -levels, we propose a function space $\mathcal{F}^{(L)}$ and a complexity measure $\mathcal{E}^{(L)}$ that satisfy the following properties:

- i. *Width unlimited:* $\mathcal{F}^{(L)}$ contains all functions that can be represented by an L -layer NN with arbitrarily-wide hidden layers;
- ii. *Approximation guarantee:* Any function f in $\mathcal{F}^{(L)}$ can be approximated by an L -layer NN with a cost that depends on $\mathcal{E}^{(L)}(f)$;
- iii. *Generalization guarantee:* Generalization errors can be upper-bounded for learning in $\mathcal{F}^{(L)}$ with $\mathcal{E}^{(L)}$ under control;
- iv. *Depth dependence:* There exist choices of L and the activation function under which $\mathcal{F}^{(L)}$ is strictly smaller than $\mathcal{F}^{(L+1)}$.
- v. *Feature learning:* Gradient descent training of L -layer NNs in a feature-learning regime corresponds to a learning dynamics in $\mathcal{F}^{(L)}$;

To our knowledge, this is the first proposal satisfying all of the properties above—or even just (i) and (v) together—thus opening up a new perspective in the understanding of deep NNs.

The rest of the paper is organized as follows. In Section 3, we introduce the neural Hilbert ladder (NHL) model and the NHL spaces and NHL complexities that it gives rise to. In Section 4, we prove static correspondences between multi-layer NNs and NHLs, verifying (i) and (ii). In Section 5, we prove generalization bounds for learning NHLs via a Rademacher

complexity analysis, verifying (iii). In Section 6, we show examples of depth separation in the NHL spaces under the ReLU activation function, verifying (iv). In Section 7, we show that the training of multi-layer NNs in the mean-field limit translates to a particular learning dynamics of the NHL, verifying (v), and it is further illustrated by numerical results on synthetic tasks. Prior literature is discussed in Section 8.

2. Background

2.1 Basic Notations and Definitions

We use bold lower-case letters (e.g. \mathbf{x} and \mathbf{z}) to denote vectors and bold upper-case letters (e.g. \mathbf{U} and \mathbf{H}) to denote random variables or random fields. $\forall m \in \mathbb{N}_+$, we write $[m] := \{1, \dots, m\}$. When the indices i, j, t and s and the variables \mathbf{x} and \mathbf{x}' appear without being further defined, they are considered by default as under the universal quantifiers “ $\forall i, j \in [m]$ ”, “ $\forall t, s \geq 0$ ” and “ $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$ ”.

Suppose \mathcal{U} is a measurable space, and we let $\mathcal{P}(\mathcal{U})$ denote the the space of probability measures on \mathcal{U} . $\forall \mu \in \mathcal{P}(\mathcal{U})$, we let $L^2(\mathcal{U}, \mu)$ denote the space of square-integrable functions on \mathcal{U} with respect to μ , and $\forall \xi \in L^2(\mathcal{U}, \mu)$, we write $\|\xi\|_{L^2(\mathcal{U}, \mu)} := (\int |\xi(u)|^2 \mu(du))^{1/2}$. If \mathbf{U} is a \mathcal{U} -valued random variable, we let $\text{Law}(\mathbf{U}) \in \mathcal{P}(\mathcal{U})$ denote its law and let $\mathbb{E}[\phi(\mathbf{U})] = \int \phi(u)[\text{Law}(\mathbf{U})](du)$ denote the expectation of any measurable function $\phi : \mathcal{U} \rightarrow \mathbb{R}$ applied to \mathbf{U} . Suppose additionally that \mathcal{U} is equipped with a norm (or quasi-norm) $\|\cdot\|_{\mathcal{U}}$. We define $\mathbb{B}(\mathcal{U}; M) := \{u \in \mathcal{U} : \|u\|_{\mathcal{U}} \leq M\}$ for $M > 0$, write $\mathbb{B}(\mathcal{U}) := \mathbb{B}(\mathcal{U}; 1)$ for the unit ball in \mathcal{U} , and let $\hat{\mathcal{U}} := \{u \in \mathcal{U} : \|u\|_{\mathcal{U}} = 1\}$ denote the unit sphere in \mathcal{U} .

For $N \in \mathbb{N}_+$, we let $\text{Lip}(\mathbb{R}^N)$ denote the space of functions on \mathbb{R}^N with Lipschitz constant at most 1. We say a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is *non-negatively homogeneous with degree 1* (or *1-homogeneous* for short) if $\forall u \in \mathbb{R}, a \geq 0, f(au) = af(u)$. An example is the *regularized linear unit (ReLU)*, defined as $\text{ReLU}(u) = \max\{u, 0\}$, a popular choice for the activation function in NNs.

Let \mathcal{X} denote the *input domain*, which we assume throughout the paper to be a compact subset of \mathbb{R}^d . We let \mathcal{C} denote the space of continuous functions on \mathcal{X} equipped with the sup-norm $\|f\|_{\infty} := \sup_{\mathbf{x} \in \mathcal{X}} |f(\mathbf{x})|$ and the associated Borel sigma-algebra. For $\nu \in \mathcal{P}(\mathcal{X})$ and $f \in \mathcal{C}$, we write $\mathcal{E}_{\mathbf{x} \sim \nu} \{f(\mathbf{x})\} := \int_{\mathcal{X}} f(\mathbf{x}) \nu(d\mathbf{x})$. The $d \times d$ identity matrix is denoted by I_d .

2.2 Multi-Layer Neural Network (NN)

We consider an L -layer (*fully-connected*) NN with width m as expressing a function on \mathcal{X} of the following form:

$$f_m(\mathbf{x}) := \frac{1}{m} \sum_{i=1}^m a_i \sigma(h_i^{(L-1)}(\mathbf{x})) , \quad (1)$$

where we define $h_i^{(1)}(\mathbf{x}) := \mathbf{z}_i^{\top} \cdot \mathbf{x} = \sum_{j=1}^d z_{i,j} x_j$, and $\forall l \in [L-2]$,

$$h_i^{(l+1)}(\mathbf{x}) := \frac{1}{m} \sum_{j=1}^m W_{i,j}^{(l)} \sigma(h_j^{(l)}(\mathbf{x})) . \quad (2)$$

Here, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the *activation function*, and each $z_{i,j}, W_{i,j}^{(l)}$ and a_i is a weight parameter of the input layer, the l th middle layer and the output layer, respectively. For simplicity, we

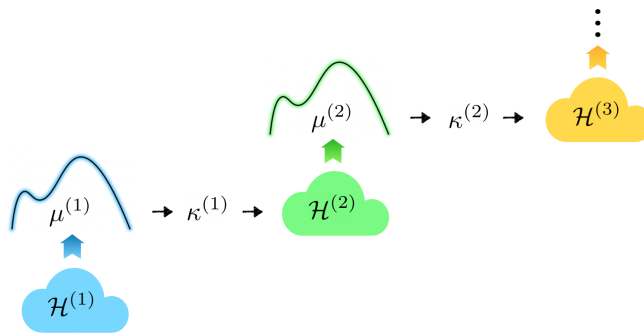


Figure 1: **Illustration of an NHL**, as defined in Definition 2. Each $\mathcal{H}^{(l)}$ is an RKHS; each $\mu^{(l)}$ is a probability measure on $\mathcal{H}^{(l)}$; each kernel function $\kappa^{(l)}$ is defined by $\mu^{(l)}$ through (4) and, in turn, defines $\mathcal{H}^{(l+1)}$ as its RKHS.

will omit the bias term in the main text and discuss its inclusion in Appendix B.1. We refer to $h_i^{(l)}$ as the *pre-activation function* represented by the i th neuron in the l th hidden layer.

The $1/m$ factor in (2) is often called the *mean-field scaling*, which allows large m limits to be considered while the parameters stay scale-free. Unlike the NTK scaling, the mean-field scaling allows feature learning to occur, including in the infinite-width limit (Yang and Hu, 2021). The comparison with NTK and other scaling choices are further discussed in Section 8.

2.3 Distributional View of Shallow NN

When $L = 2$, the model defined by (1) reduces to the following expression of a shallow NN:

$$f(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m a_i \sigma(\mathbf{z}_i^\top \cdot \mathbf{x}) ,$$

which can be equivalently represented as $f(\mathbf{x}) = \int a \sigma(\mathbf{z}^\top \cdot \mathbf{x}) \mu_m(da, d\mathbf{z})$, if we define $\mu_m(da, d\mathbf{z}) := \frac{1}{m} \sum_{i=1}^m \delta_{a_i}(da) \delta_{\mathbf{z}_i}(d\mathbf{z})$ as a probability measure on $\mathbb{R} \times \mathbb{R}^d$ that is the average of m Dirac measures. More generally, we may consider functions that can be represented in a distributional form as

$$f(\mathbf{x}) = \int a \sigma(\mathbf{z}^\top \cdot \mathbf{x}) \mu(da, d\mathbf{z}) , \quad (3)$$

for some $\mu \in \mathcal{P}(\mathbb{R} \times \mathbb{R}^d)$. In particular, the *Barron norm* (E et al., 2022) of a function f can be defined as

$$\|f\|_{\mathcal{B}} := \inf_{\mu} \int |a| \|\mathbf{z}\| \mu(da, d\mathbf{z}) ,$$

with the infimum taken over all $\mu \in \mathcal{P}(\mathbb{R} \times \mathbb{R}^d)$ such that (3) holds. One can further define the *Barron space* to contain all functions with a bounded Barron norm, which is equivalent to the \mathcal{F}_1 space considered by Bach (2017a) if σ is 1-homogeneous.

Favorable theoretical properties of the Barron space have been derived in prior literature, including approximation guarantees and Rademacher complexity bounds (Barron, 1993; Bach, 2017a; E et al., 2022). Meanwhile, the training dynamics of a two-layer NN can be understood through the Wasserstein gradient flow of the underlying probability measure

μ (Mei et al., 2018; Chizat and Bach, 2018; Rotskoff and Vanden-Eijnden, 2018; Sirignano and Spiliopoulos, 2020). Hence, through the distributional perspective, prior works have developed a promising picture of shallow NNs by viewing it as learning in the Barron space. A central motivation of this work is to extend this theory to NNs with more than two layers.

2.4 Reproducing Kernel Hilbert Space (RKHS)

A *Hilbert space* is a vector space equipped with an inner product, $\langle \cdot, \cdot \rangle$, such that the norm defined through $\| \cdot \| := \langle \cdot, \cdot \rangle$ makes it a complete metric space. Of particular interest to learning theory is a type of Hilbert spaces consisting of functions on the input domain, which we define below.

Let $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a symmetric and positive semi-definite function, which we call a *kernel function*. It is associated with a particular Hilbert space on \mathcal{X} , whose definition, existence and uniqueness are given by the following foundational result, often called the Moore-Aronszajn theorem (Aronszajn, 1950; Cucker and Smale, 2002):

Lemma 1 (Moore-Aronszajn) *There exists a unique Hilbert space, \mathcal{H} , consisting of functions on \mathcal{X} and equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, which satisfies the following properties:*

1. $\forall \mathbf{x} \in \mathcal{X}, \kappa(\mathbf{x}, \cdot) \in \mathcal{H}$;
2. $\forall f \in \mathcal{H}, \forall \mathbf{x} \in \mathcal{X}, \langle f, \kappa(\mathbf{x}, \cdot) \rangle_{\mathcal{H}} = f(\mathbf{x})$;
3. *the span of the set $\{\kappa(\mathbf{x}, \cdot)\}_{\mathbf{x} \in \mathcal{X}}$ is dense in \mathcal{H} .*

The Hilbert space that satisfies these properties is called the *Reproducing Kernel Hilbert Space* (RKHS) associated with (a.k.a. reproducing) the kernel function κ . The RKHS plays an important role in classical learning theory as well as in mathematics and physics, and we refer the readers to Cucker and Smale (2002); Mohri et al. (2018) for further background. In Section 8, we will review the prior efforts in understanding NNs through kernels and RKHSs.

3. Neural Hilbert Ladder (NHL)

We begin by introducing a way to create an RKHS from a *distribution of functions* on \mathcal{X} . If $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function of interest and $\mu \in \mathcal{P}(\mathcal{C})$, we define $\kappa_{\mu} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ by

$$\kappa_{\mu}(\mathbf{x}, \mathbf{x}') := \int \sigma(h(\mathbf{x}))\sigma(h(\mathbf{x}'))\mu(dh) , \tag{4}$$

which is symmetric and positive semi-definite. Hence, by Lemma 1, there is a unique RKHS on \mathcal{X} associated with the kernel function κ_{μ} , which we denote by \mathcal{H}_{μ} .

Leveraging this observation, we will present a recipe to construct a hierarchy of RKHSs iteratively. At the first level, we define $\mathcal{H}^{(1)} := \{\mathbf{x} \mapsto \mathbf{z}^{\top} \cdot \mathbf{x} : \mathbf{z} \in \mathbb{R}^d\}$ to be the space of linear functions on \mathbb{R}^d . Through the canonical isomorphism with \mathbb{R}^d , $\mathcal{H}^{(1)}$ inherits an inner product from the Euclidean inner product on \mathbb{R}^d , which makes $\mathcal{H}^{(1)}$ the RKHS associated with the kernel function $\kappa^{(0)}(\mathbf{x}, \mathbf{x}') := \mathbf{x}^{\top} \cdot \mathbf{x}'$. Then, for $L \geq 2$, we define an *L-level neural Hilbert ladder* (NHL) as follows:

Definition 2 Suppose each of $\mathcal{H}^{(2)}, \dots, \mathcal{H}^{(L)}$ is an RKHS on \mathcal{X} , and $\forall l \in [L-1]$, there exists $\mu^{(l)} \in \mathcal{P}(\mathcal{H}^{(l)})$ such that $\mathcal{H}^{(l+1)} = \mathcal{H}_{\mu^{(l)}}$, which is the RKHS associated with $\kappa^{(l)} := \kappa_{\mu^{(l)}}$. Then, we say that $(\mathcal{H}^{(l)})_{l \in [L]}$ is an L -level NHL induced by the sequence of probability measures, $(\mu^{(l)})_{l \in [L-1]}$. In addition, we say that a function f on \mathcal{X} belongs to the NHL if $f \in \mathcal{H}^{(L)}$.

Put differently, to define an NHL, at each level l we choose a probability measure supported on $\mathcal{H}^{(l)}$ (which can be viewed as the law of a random field on \mathcal{X}) to generate $\kappa^{(l)}$ via (4), which then determines $\mathcal{H}^{(l+1)}$. Thus, an NHL is a ladder of RKHSs constructed by interleaving them with random fields and kernel functions, as illustrated in Figure 1.

This concept allows us to define complexity measures of functions as well as function spaces quantitatively, as we introduce next.

3.1 Complexity Measures and Function Spaces

If $\mathcal{H}, \mathcal{H}'$ are two RKHSs on \mathcal{X} , then for $p > 0$ or $p = \infty$, we can define

$$\mathcal{D}_p(\mathcal{H}, \mathcal{H}') := \inf_{\mu \in \mathcal{P}(\mathcal{H}), \mathcal{H}_\mu = \mathcal{H}'} \|\mu\|_{\mathcal{H}, p},$$

where $\|\mu\|_{\mathcal{H}, p} := (\int \|h\|_{\mathcal{H}}^p \mu(dh))^{1/p}$ for $p > 0$ and $\|\mu\|_{\mathcal{H}, \infty}$ denotes the essential supremum of the function $\|\cdot\|_{\mathcal{H}}$ on \mathcal{H} with respect to μ . Note that the infimum is taken over all $\mu \in \mathcal{P}(\mathcal{H})$ such that κ_μ equals the kernel function of \mathcal{H}' . Then, given an RKHS \mathcal{H} on \mathcal{X} , we define

$$\mathcal{D}_p^{(L)}(\mathcal{H}) := \inf_{\mathcal{H}^{(2)}, \dots, \mathcal{H}^{(L-1)}} \prod_{l=1}^{L-1} \mathcal{D}_p(\mathcal{H}^{(l)}, \mathcal{H}^{(l+1)}), \quad (5)$$

with the infimum taken over all choices of $\mathcal{H}^{(2)}, \dots, \mathcal{H}^{(L-1)}$ as RKHSs. Heuristically speaking, it quantifies the least *price* (whose meaning regarding approximation and generalization will be discussed later) to be paid for arriving at \mathcal{H} as the L th-level of an NHL—note that one is free to choose the RKHSs and the probability measures in the middle of the ladder as long as Definition 2 is preserved. Then, we define the (L, p) -NHL complexity of a function f as:

$$\mathcal{C}_p^{(L)}(f) := \inf_{\mathcal{H}} \left(\|f\|_{\mathcal{H}} \cdot \mathcal{D}_p^{(L)}(\mathcal{H}) \right), \quad (6)$$

with the infimum taken over all RKHS \mathcal{H} . Finally, we define the (L, p) -NHL space, $\mathcal{F}_p^{(L)}$, to contain all functions with a finite (L, p) -NHL complexity:

$$\mathcal{F}_p^{(L)} := \{f : \mathcal{C}_p^{(L)}(f) < \infty\} = \bigcup_{\mathcal{D}_p^{(L)}(\mathcal{H}) < \infty} \mathcal{H}. \quad (7)$$

Note that unlike in the kernel theories of NNs (see Section 8), the function space $\mathcal{F}_p^{(L)}$ is not a *single* RKHS but an infinite union of them.

Some properties of $\mathcal{F}_p^{(L)}$ and $\mathcal{C}_p^{(L)}$ are in order:

Theorem 3 For $L \geq 2$, it holds that:

(a) For $p \in (0, \infty]$, $\mathcal{F}_p^{(L)}$ is a vector space and $\mathcal{C}_p^{(L)}$ is non-decreasing in p .

- (b) If $p \in [2, \infty]$ and σ is J_σ -Lipschitz with $\sigma(0) = 0$ (e.g. $J_\sigma = 1$ and σ is ReLU or tanh), then $\forall f \in \mathcal{F}_p^{(L)}$, it holds that $|f(\mathbf{x})| \leq (J_\sigma)^{L-1} \mathcal{C}_p^{(L)}(f) \|\mathbf{x}\|$, $\forall \mathbf{x} \in \mathcal{X}$, and f satisfies the Lipschitz condition on \mathcal{X} with constant $(J_\sigma)^{L-1} \mathcal{C}_p^{(L)}(f)$. Hence, $\mathcal{F}_p^{(L)} \subseteq \mathcal{C}$.
- (c) If σ is 1-homogeneous, then $\mathcal{C}_p^{(L)}$ is equal for all $p \in [2, \infty]$. Moreover, for each such p , $\mathcal{F}_p^{(L)}$ is a quasi-Banach space with $\mathcal{C}_p^{(L)}$ as the quasi-norm (see Definition 26).

These results are proved in Appendix A. Specially, when σ is 1-homogeneous (e.g. ReLU), Theorem 3(c) allows us to define $\mathcal{C}^{(L)} := \mathcal{C}_p^{(L)}$ and $\mathcal{F}^{(L)} := \mathcal{F}_p^{(L)}$ for any $p \in [2, \infty]$, which we will call the L -level NHL complexity and the L -level NHL space, respectively.

Remark 4 If σ is Lipschitz and $\sigma(0) = 0$, Theorem 3(b) implies that every $f \in \mathcal{F}^{(L)}$ needs to satisfy $f(0) = 0$. This corresponds to the omission of the bias terms in the NN and can be changed by either accounting for the bias terms in the definitions above (see Appendix B.1) or redefining the input space to have an extra constant dimension via a map $\mathbf{x} \mapsto (\mathbf{x}, 1) \in \mathbb{R}^{d+1}$.

3.1.1 EXAMPLE: $L = 2$

When σ is 1-homogeneous and $L = 2$, we see from (6) and Lemma 25 in Appendix A.3 that

$$\mathcal{C}^{(2)}(f) = \inf_{\mu^{(1)} \in \mathcal{P}(\mathcal{H}^{(1)})} \|f\|_{\mathcal{H}^{(2)}} ,$$

where the infimum is taken over all probability measures $\mu^{(1)}$ supported within the unit sphere of $\mathcal{H}^{(1)}$. Meanwhile, it is known (E and Wojtowytsch, 2022) that the Barron norm of a function f on \mathcal{X} can be rewritten as

$$\|f\|_{\mathcal{B}} = \inf_{\xi, \rho} \left(\int |\xi(\mathbf{z})|^2 \rho(d\mathbf{z}) \right)^{1/2} ,$$

where the infimum is taken over all $\rho \in \mathcal{P}(\mathbb{S}^{d-1})$ and all measurable functions $\xi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ such that

$$f(\mathbf{x}) = \int \xi(\mathbf{z}) \sigma(\mathbf{z}^\top \cdot \mathbf{x}) \rho(d\mathbf{z}) . \quad (8)$$

Noticing the analogous roles played by $\mu^{(1)}$ and ρ , we see an equivalence between $\mathcal{C}^{(2)}$ and the Barron norm in light of Lemma 6 below:

Proposition 5 If σ is 1-homogeneous, then $\|f\|_{\mathcal{B}} = \mathcal{C}^{(2)}(f)$.

Lemma 6 Let σ be Lipschitz and $\mu \in \mathcal{P}(\mathcal{C})$ with $\int \|h\|_\infty^2 \mu(dh) < \infty$. Then, a function f on \mathcal{X} belongs to \mathcal{H}_μ if and only if $\exists \xi \in L_2(\mathcal{C}, \mu)$ such that

$$f(\mathbf{x}) = \int \xi(h) \sigma(h(\mathbf{x})) \mu(dh) , \quad \forall \mathbf{x} \in \mathcal{X} . \quad (9)$$

Moreover, $\|f\|_{\mathcal{H}_\mu} = \inf_\xi \|\xi\|_{L^2(\mathcal{C}, \mu)}$, with the infimum over all $\xi \in L^2(\mathcal{C}, \mu)$ satisfying (9).

Lemma 6 generalizes prior insights on the duality between RKHS and integral transforms (Rahimi and Recht, 2008a; Bach, 2017a,b). While those results focus on cases where the basis functions are indexed by a compact set, here the basis functions $\{\sigma(h(\cdot))\}_{h \in \mathcal{C}}$ are indexed by an infinite-dimensional function space. In Appendix A.5, we prove this lemma based on a general result by Saitoh (1997).

Proposition 5 therefore implies that $\mathcal{F}^{(2)}$ is identical to the Barron space. In fact, when $L = 2$, (7) reduces to $\mathcal{F}^{(2)} = \bigcup_{\rho \in \mathcal{P}(\mathbb{R}^d)} \mathcal{H}_\rho$, where we define \mathcal{H}_ρ as the RKHS associated with the kernel function $\kappa_\rho(\mathbf{x}, \mathbf{x}') := \int \sigma(\mathbf{z}^\top \cdot \mathbf{x}) \sigma(\mathbf{z}^\top \cdot \mathbf{x}') \rho(d\mathbf{z})$. This agrees with the decomposition of the Barron space as a union of RKHSs (E et al., 2019).

Thus, $\mathcal{F}^{(L)}$ can be seen as generalizing the Barron space to the cases where $L > 2$.

3.2 Alternative Form via Coupled Random Fields

For each $l \in [L - 1]$, $\mu^{(l)}$ can be interpreted as the law of a random field, $\mathbf{H}^{(l)}$, whose sample paths—as functions on \mathcal{X} —belong to $\mathcal{H}^{(l)}$. In fact, one could define the random fields on a common probability space to derive an alternative formulation of the NHL that will be relevant later:

Proposition 7 *Given an L -level NHL as in Definition 2, there exist $L - 1$ random fields on \mathcal{X} , $\mathbf{H}^{(1)}, \dots, \mathbf{H}^{(L-1)}$, as well as $L - 1$ scalar random variables, $\Xi^{(1)}, \dots, \Xi^{(L-2)}$, and \mathbf{A} , which are defined on a common probability space and satisfy the following properties:*

- $\mathbf{H}^{(1)}, \dots, \mathbf{H}^{(L-1)}$ are mutually independent, and $\forall l \in [L - 1], \mu^{(l)} = \text{Law}(\mathbf{H}^{(l)})$;
- $\forall l \in [L - 2], \Xi^{(l)}$ is measurable with respect to the sigma-algebra generated by $\mathbf{H}^{(l)}$ and $\mathbf{H}^{(l+1)}$. Moreover,

$$\mathbf{H}^{(l+1)}(\mathbf{x}) = \mathbb{E}[\Xi^{(l)} \sigma(\mathbf{H}^{(l)}(\mathbf{x})) | \mathbf{H}^{(l+1)}], \quad (10)$$

where $\mathbb{E}[\cdot | \cdot]$ denotes the conditional expectation, and $\|\mathbf{H}^{(l+1)}\|_{\mathcal{H}^{(l+1)}}^2 = \mathbb{E}[(\Xi^{(l)})^2 | \mathbf{H}^{(l+1)}]$.

- \mathbf{A} is measurable with respect to $\mathbf{H}^{(L-1)}$. Moreover,

$$f(\mathbf{x}) = \mathbb{E}[\mathbf{A} \sigma(\mathbf{H}^{(L-1)}(\mathbf{x}))], \quad (11)$$

and $\|f\|_{\mathcal{H}^{(L)}}^2 = \mathbb{E}[\mathbf{A}^2]$.

The proof of Proposition 7 is given in Appendix A.4 and also based on Lemma 6.

4. Realization and Approximation by Neural Networks

For $p \in [2, \infty]$, we shall define quantities $M_{m,p}^{(1)}, \dots, M_{m,p}^{(L)} \geq 0$ associated with the NN defined in Section 2.2 by

$$M_{m,p}^{(1)} := \left(\frac{1}{m} \sum_{i=1}^m \|\mathbf{z}_i\|_2^p \right)^{\frac{1}{p}}, \quad M_{m,p}^{(L)} := \left(\frac{1}{m} \sum_{i=1}^m a_i^2 \right)^{\frac{1}{2}},$$

$$M_{m,p}^{(l+1)} := \left(\frac{1}{m} \sum_{i=1}^m \left(\frac{1}{m} \sum_{j=1}^m |W_{i,j}^{(l)}|^2 \right)^{\frac{p}{2}} \right)^{\frac{1}{p}}, \quad \forall l \in [L-2],$$

with the usual convention of $(\sum_{i=1}^m |u_i|^\infty)^{1/\infty} = \sup_{i \in [m]} |u_i|$. Modulo the $1/m$ scaling factor, they coincide with the per-layer Frobenius norms of the weight parameters when $p = 2$, and more generally, they appear in the *group norm* of finite-width NNs defined in Neyshabur et al. (2015). Meanwhile, with the $1/m$ scaling factor, these quantities admit width-independent upper bounds if each parameter is sampled i.i.d.

4.1 NN as NHL

First, we show that every L -layer NN represents a function in $\mathcal{F}_p^{(L)}$, verifying property (i).

Theorem 8 $\forall p \in [2, \infty]$, $f_m \in \mathcal{F}_p^{(L)}$ with $\mathcal{E}_p^{(L)}(f_m) \leq \prod_{l=1}^L M_{m,p}^{(l)}$. In particular, f_m belongs to the NHL of $(\mathcal{H}_m^{(l)})_{l \in [L]}$, where we define $\mathcal{H}_m^{(1)} := \mathcal{H}^{(1)}$, and $\forall l \in [L-1]$, $\mathcal{H}_m^{(l+1)} := \mathcal{H}_{\mu_m^{(l)}}$ with $\mu_m^{(l)} := \frac{1}{m} \sum_{i=1}^m \delta_{h_i^{(l)}} \in \mathcal{P}(\mathcal{C})$ being the empirical measure (on functional space) of the pre-activation functions of the neurons in the l th hidden layer.

The proof is given in Appendix B.2 and leverages Lemma 6. Thus, a multi-layer NN can indeed be viewed as representing an NHL whose complexity is controlled by the layer-wise norms of the NN parameters. In particular, we see that the series of random fields can be constructed based on the pre-activation functions in the respective hidden layers.

4.2 NHL can be Approximated by NN

Conversely, we can show that any function in $\mathcal{F}_\infty^{(L)}$ can be approximated relatively efficiently in squared error by an L -layer NN:

Theorem 9 Suppose σ is J_σ -Lipschitz with $\sigma(0) = 0$ and $\nu \in \mathcal{P}(\mathcal{X})$ satisfies $\mathcal{E}_{\mathbf{x} \sim \nu} \{\|\mathbf{x}\|^2\} = M_\nu < \infty$. Given any $f \in \mathcal{F}_\infty^{(L)}$, there exists a function f_m represented by an L -layer NN with width m such that $\mathcal{E}_{\mathbf{x} \sim \nu} \{|f_m(\mathbf{x}) - f(\mathbf{x})|^2\} \leq (L-1)^2 (J_\sigma)^{2L-2} M_\nu (\mathcal{E}_\infty^{(L)}(f))^2 / m$.

This result is proved in Appendix B.3, where we use an inductive-in- L argument to show that a randomized approximation strategy based on sampling each $\mu^{(l)}$ independently can achieve a low approximation error in expectation.

Theorem 9 implies that a function in $\mathcal{F}_\infty^{(L)}$ can be approximated within a squared error of ϵ by an L -layer NN using $O(L^5/\epsilon^4)$ number of parameters in total. In comparison, functions in the *neural tree space* defined by E and Wojtowytsch (2020) require above the order of $(L/\epsilon)^L$ parameters, or in other words, the approximation error bound decays *exponentially* slowly as the depth increases. The contrast highlights a property of multi-layer NNs—that all neurons in one hidden layer share the same pre-synaptic neurons of the preceding layer—which is captured by the NHL but not by the neural tree space representation, whose branching structure incurs the exponential dependence on the depth.

We note that for general activation functions, the current proof technique for Theorem 9 only succeeds when $p = \infty$. Nonetheless, when σ is 1-homogeneous, Theorems 8 and 9 are both applicable to the space $\mathcal{F}^{(L)}$ and establish a two-way correspondence between L -layer

NNs and functions in this space with the approximation cost governed by $\mathcal{C}^{(L)}$, thus verifying both criteria (i) and (ii).

5. Generalization Theory

To study the learning aspect of the function space defined above, we consider the task of fitting a target function on the input domain, $f^* : \mathcal{X} \rightarrow \mathcal{Y}$ with \mathcal{Y} being a bounded subset of \mathbb{R} , using $\mathcal{F}^{(L)}$ as the hypothesis class (and we assume in this section that σ is 1-homogeneous). Concretely, we desire for a function f from $\mathcal{F}^{(L)}$ that minimizes the *population risk*, defined as $\mathcal{R}(f) := \mathbb{E}_{\mathbf{x} \sim \nu} \{l(f(\mathbf{x}), f^*(\mathbf{x}))\}$, where $\nu \in \mathcal{P}(\mathcal{X})$ is an underlying data distribution on \mathcal{X} and l is a loss function on $\mathbb{R} \times \mathcal{Y}$ (e.g., in least squares regression, $l(\hat{y}, y) = \frac{1}{2}(\hat{y} - y)^2$). In supervised learning, instead of having access to ν directly, we are typically given a training set of size n , $S = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathcal{X}$, sampled i.i.d. from ν , and we write $\nu_n = \frac{1}{n} \sum_{k=1}^n \delta_{\mathbf{x}_k} \in \mathcal{P}(\mathcal{X})$. A standard strategy is to find a function from the given function space that achieves a low *empirical risk*, defined as $\mathcal{R}_n(f) := \mathbb{E}_{\mathbf{x}} \{l(f(\mathbf{x}), f^*(\mathbf{x}))\}$, where we write $\mathbb{E}_{\mathbf{x}}$ for $\mathbb{E}_{\mathbf{x} \sim \nu_n}$ for simplicity. Then, the question of generalization concerns how fast the discrepancy between \mathcal{R} and \mathcal{R}_n decays as n increases.

Classical learning theory suggests that one can prove uniform upper bounds on the discrepancy through estimating the *Rademacher complexity* of $\mathcal{F}^{(L)}$. Recall that the *empirical Rademacher complexity* of a function space \mathcal{F} with respect to the set S is defined by $\widehat{\text{Rad}}_S(\mathcal{F}) := \mathbb{E}_{\boldsymbol{\tau}} \left[\frac{1}{n} \sup_{f \in \mathcal{F}} \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right]$, where $\boldsymbol{\tau} = [\tau_1, \dots, \tau_n]$ is a vector of i.i.d. Rademacher random variables. Note that while the Rademacher complexity of RKHS is known (Mendelson, 2003; Bartlett et al., 2005), $\mathcal{F}^{(L)}$ is not *one* RKHS but a union of infinitely many of them (and hence richer), which calls for a different analysis.

Our main result in this section is the following:

Theorem 10 *If $\mathcal{X} \subseteq \mathbb{B}(\mathbb{R}^d)$ and σ is 1-homogeneous, then $\forall M > 0$, for any training set S of size n , $\widehat{\text{Rad}}_S(\mathbb{B}(\mathcal{F}^{(L)}; M)) \leq M(\sqrt{2L \log(2J_\sigma)} + 1)/\sqrt{n}$, where $J_\sigma := \max\{|\sigma(1)|, |\sigma(-1)|\}$.*

The proof is given in Appendix C.1. In particular, we carry out an inductive-in- L argument inspired by Neyshabur et al. (2015) combined with a strategy proposed by Golowich et al. (2018) for reducing the dependency of the bound on L from exponential to $O(\sqrt{L})$. Note that the same $O(1/\sqrt{n})$ rate is obtained as in the Rademacher complexity bounds for RKHSs.

Combining this result with Theorem 3(b) and classical generalization bounds based on Rademacher complexity (e.g. Mohri et al., 2018), we derive the following generalization guarantee for learning in the space of NHLs under the ReLU activation, thus verifying (iii):

Corollary 11 *Suppose that $\mathcal{X} \subseteq \mathbb{B}(\mathbb{R}^d)$, $\mathcal{Y} \subseteq [-1, 1]$, σ is ReLU, and $\forall y \in \mathcal{Y}$, the function $\hat{y} \mapsto l(\hat{y}, y)$ is J_l -Lipschitz on $[-1, 1]$. Then, $\forall \delta > 0$, with probability at least $1 - \delta$ over the i.i.d. sampling of a sample S of size n in \mathcal{X} , it holds for all functions f with $\mathcal{C}^{(L)}(f) \leq 1$ that $\mathcal{R}(f) \leq \mathcal{R}_n(f) + 2J_l(\sqrt{2L \log 2} + 1)/\sqrt{n} + 3\sqrt{\log(2/\delta)}/\sqrt{2n}$.*

6. Effect of Depth

A rich body of work has studied the benefit of depth in NNs in terms of whether certain functions can be approximated much more efficiently by deeper NNs than shallow ones, where the efficiency is usually measured by the width required. Meanwhile, it remains interesting

whether the effect of depth can be properly characterized for *width-unlimited* NNs in terms of the respective function spaces. Under the NTK theory, for example, the function spaces corresponding to NNs with different depths are actually equivalent (Bietti and Bach, 2021; Chen and Xu, 2021). By contrast, below we demonstrate an example where the NHL space is strictly larger as we increase the depth.

We consider the case where σ is the ReLU function. On one hand, since a single hidden layer is capable of expressing the identity function using a pair of ReLU neurons with opposite weights, it is straightforward to show that $\mathcal{F}^{(L)}$ does not shrink as L increases. In fact, we can obtain a slightly sharper bound on $\mathcal{C}^{(L+1)}$ relative to $\mathcal{C}^{(L)}$ by further exploiting the properties of ReLU:

Lemma 12 *If σ is ReLU, then $\forall L \geq 2$, $\mathcal{F}^{(L)} \subseteq \mathcal{F}^{(L+1)}$ and $\mathcal{C}^{(L+1)}(f) \leq \sqrt{2}\mathcal{C}^{(L)}(f)$, $\forall f \in \mathcal{F}^{(L)}$.*

The proof is given in Section E.1. On the other hand, Ongie et al. (2020) proved that the function $f(\mathbf{x}) = \max\{1 - \|\mathbf{x}\|_1, 0\}$ on \mathbb{R}^d cannot be represented by any infinite-width shallow ReLU NN with a finite Barron norm, and hence it does not belong to $\mathcal{F}^{(2)}$. Meanwhile, it can be represented by a finite-width 3-layer NN, and thus, as a consequence of Theorem 9, it belongs to $\mathcal{F}^{(3)}$. Therefore, we establish a strict inclusion of $\mathcal{F}^{(2)}$ within $\mathcal{F}^{(3)}$:

Proposition 13 *When σ is ReLU, $\mathcal{F}^{(2)} \subsetneq \mathcal{F}^{(3)}$.*

Prior results on depth separation (reviewed in Section 8) often involve functions whose approximation cost—in terms of the number of parameters—suffers from the curse of dimensionality only when the depth is small. In our framework, with the cost measured by the NHL complexity, we can show an analogous result via an example studied by Wojtowysch (2024), as we briefly recount below. Fixing $\epsilon_0 \in (0, 1)$, we define $f_{\epsilon_0} : \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$f_{\epsilon_0}(\mathbf{x}) = \begin{cases} 1, & \text{if } \|\mathbf{x}\| \leq \epsilon_0, \\ 1 - \frac{\|\mathbf{x}\| - \epsilon_0}{1 - \epsilon_0}, & \text{if } \epsilon_0 < \|\mathbf{x}\| \leq 1, \\ 0, & \text{if } \|\mathbf{x}\| > 1, \end{cases}$$

When the activation function is ReLU and the bias term is considered, it is shown that f_{ϵ_0} can be written as the composition of two functions whose Barron norms depend sublinearly on the input dimension d , thus implying a similar bound on $\mathcal{C}^{(3)}(f_{\epsilon_0})$. Meanwhile, the Barron norm of f_{ϵ_0} itself, $\mathcal{C}^{(2)}(f_{\epsilon_0})$, is exponentially large in d (Wojtowysch, 2024). Hence, we arrive at the following type of depth separation result for NHL spaces:

Proposition 14 *When σ is ReLU and the bias term is included, there exists a function f on \mathbb{R}^d such that $\mathcal{C}^{(3)}(f)$ depends sublinearly on d while $\mathcal{C}^{(2)}(f)$ is exponentially large in d .*

7. Training Dynamics under Gradient Flow (GF)

Given the correspondence between NNs and NHLs shown in Section 4.2, we may regard the training of NNs as a instantiating a particular learning rule for solving empirical risk minimization within $\mathcal{F}^{(L)}$, whose mechanism will be further analyzed in this section. Typically in practice, we initialize an NN by sampling its parameters randomly and train it by performing variants of GD on the parameters with respect to the empirical risk. We assume below that

Assumption 15 σ is differentiable and its derivative σ' is Lipschitz and bounded.

Assumption 16 l is differentiable with respect to the first argument and the derivative $\partial_{\hat{y}}l(\hat{y}, y)$ is Lipschitz in \hat{y} for $y \in \mathcal{Y}$.

Assumption 17 At $t = 0$, each $W_{i,j,0}$, $a_{i,0}$ and $z_{i,0}$ is sampled independently from ρ_W , $\rho_a \in \mathcal{P}(\mathbb{R})$ and $\rho_z \in \mathcal{P}(\mathbb{R}^d)$, respectively. Moreover, ρ_W and ρ_a have zero mean, ρ_W has a finite fourth-moment, ρ_z has a finite covariance, and ρ_a is bounded.

Note that Assumption 15 is standard in prior works on the mean-field theory of shallow NNs (e.g. Chizat and Bach, 2018) though it is not satisfied when σ is ReLU, and Assumption 16 is satisfied if l is the squared loss or the logistic loss, for instance.

For simplicity, we consider GD dynamics in the continuous-time limit—also called the *gradient flow* (GF)—where the parameters evolve over time t (added as a subscript to the quantities in (1) and (2)) according to a system of ordinary differential equations:

$$\begin{aligned} \frac{d}{dt} z_{i,t} &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) q_{i,t}^{(1)}(\mathbf{x}) \sigma'(h_{i,t}^{(1)}(\mathbf{x})) \right\}, \\ \frac{d}{dt} a_{i,t} &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) \sigma(h_{i,i}^{(L-1)}(\mathbf{x})) \right\}, \end{aligned}$$

and $\forall l \in [L - 2]$,

$$\frac{d}{dt} W_{i,j,t}^{(l)} = -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) q_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(h_{i,t}^{(l+1)}(\mathbf{x})) \sigma(h_{j,t}^{(l)}(\mathbf{x})) \right\},$$

where $\zeta_{m,t}(\mathbf{x}) := \partial_{\hat{y}}l(\hat{y}, f^*(\mathbf{x}))|_{\hat{y}=f_{m,t}(\mathbf{x})}$, $q_{i,t}^{(L-1)}(\mathbf{x}) := a_i \sigma'(h_{i,t}^{(L-1)}(\mathbf{x}))$, and $\forall l \in [L - 2]$,

$$q_{j,t}^{(l)}(\mathbf{x}) := \frac{1}{m} \left(\sum_{i=1}^m W_{i,j,t}^{(l)} q_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(h_{i,t}^{(l+1)}(\mathbf{x})) \right).$$

Note that GF causes not only the output function $f_{m,t}$ but also the pre-activation functions in the hidden layers—and hence $\mu_{m,t}^{(l)} := \frac{1}{m} \sum_{i=1}^m \delta_{h_{i,t}^{(l)}}$ —to evolve, thus leading to a movement of the overall NHL represented by the model. While the dynamics of $(\mu_{m,t}^{(l)})_{l \in [L-1]}$ is not closed but depending intricately on the weight matrices, we will show below that their effect can be subsumed by a mean-field description once we consider the infinite-width limit.

7.1 Mean-Field NHL Dynamics

Several prior works have considered the infinite-width limits of multi-layer NNs in the mean-field scaling and derived equations that govern their GF dynamics (Araújo et al., 2019; Nguyen and Pham, 2023; Sirignano and Spiliopoulos, 2022). Following their insights, our goal is to uncover the learning dynamics of the NHL in the mean-field limit. To do so, we first write down the mean-field limit in the form of coupled random fields introduced in Section 3.2, as follows. Let \mathbf{Z}_0 and \mathbf{A}_0 be a d -dimensional random vector with law ρ_z and a scalar random variable with law ρ_a , respectively, distributed independently from each other.

Then, we introduce time-dependent random fields $\mathbf{H}_t^{(1)}, \dots, \mathbf{H}_t^{(L-1)}$ and $\mathbf{Q}_t^{(1)}, \dots, \mathbf{Q}_t^{(L-1)}$ on \mathcal{X} . At initial time, they are defined by

$$\mathbf{H}_0^{(1)}(\mathbf{x}) = \mathbf{Z}_0^\top \cdot \mathbf{x}, \quad \mathbf{Q}_0^{(L-1)}(\mathbf{x}) = \mathbf{A}_0,$$

and $\mathbf{H}_0^{(l+1)}(\mathbf{x}) = \mathbf{Q}_0^{(l)}(\mathbf{x}) = 0, \forall l \in [L-2]$. Their evolution in time is given by, $\forall l \in [L-1]$,

$$\mathbf{H}_t^{(l)}(\mathbf{x}) = \mathbf{H}_0^{(l)}(\mathbf{x}) - \int_0^t \mathcal{E}_{\mathbf{x}'} \left\{ \zeta_s(\mathbf{x}') \kappa_{t,s}^{(l-1)}(\mathbf{x}, \mathbf{x}') \mathbf{Q}_s^{(l)}(\mathbf{x}') \sigma'(\mathbf{H}_s^{(l)}(\mathbf{x}')) \right\} ds, \quad (12)$$

$$\mathbf{Q}_t^{(l)}(\mathbf{x}) = \mathbf{Q}_0^{(l)}(\mathbf{x}) - \int_0^t \mathcal{E}_{\mathbf{x}'} \left\{ \zeta_s(\mathbf{x}') \gamma_{t,s}^{(l+1)}(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{H}_s^{(l)}(\mathbf{x}')) \right\} ds, \quad (13)$$

where we define $\kappa_{t,s}^{(0)}(\mathbf{x}, \mathbf{x}') = \gamma_{t,s}^{(L)}(\mathbf{x}, \mathbf{x}') = 1$ and, $\forall l \in [L-1]$,

$$\begin{aligned} \kappa_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') &:= \mathbb{E} \left[\sigma(\mathbf{H}_t^{(l)}(\mathbf{x})) \sigma(\mathbf{H}_s^{(l)}(\mathbf{x}')) \right], \\ \gamma_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') &:= \mathbb{E} \left[\mathbf{Q}_t^{(l)}(\mathbf{x}) \mathbf{Q}_s^{(l)}(\mathbf{x}') \sigma'(\mathbf{H}_t^{(l)}(\mathbf{x})) \sigma'(\mathbf{H}_s^{(l)}(\mathbf{x}')) \right]. \end{aligned}$$

Lastly, we set $f_t(\mathbf{x}) = \mathbb{E}[\mathbf{Q}_t^{(L-1)}(\mathbf{x}) \sigma(\mathbf{H}_t^{(L-1)}(\mathbf{x}))]$ (though, from the definition above, it can be seen that $\mathbf{Q}_t^{(L-1)}(\mathbf{x})$ is independent of \mathbf{x}) and $\zeta_t(\mathbf{x}) = \partial_{\hat{y}} l(\hat{y}, f^*(\mathbf{x}))|_{\hat{y}=f_t(\mathbf{x})}$.

We can show that, by the law of large numbers, the dynamics of f_t defined via the above is the infinite-width limit of the GF training of NNs considered in Section 7.

Theorem 18 *Suppose Assumptions 15 – 17 hold. Then $\forall t \geq 0$, as $m \rightarrow \infty$,*

1. $f_{m,t}(\mathbf{x}) \xrightarrow{a.s.} f_t(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$;
2. $\forall l \in [L-1]$, the probability measure $\mu_{m,t}^{(l)}$ converges weakly to $\mu_t^{(l)} := \text{Law}(\mathbf{H}_t^{(l)})$ in all finite distributions, that is, $\forall N \in \mathbb{N}_+, \forall \mathbf{x}'_1, \dots, \mathbf{x}'_N \in \mathcal{X}$,

$$\sup_{g \in \text{Lip}(\mathbb{R}^N)} \left| \int g(h(\mathbf{x}'_1), \dots, h(\mathbf{x}'_N)) \mu_{m,t}^{(l)}(dh) - \int g(h(\mathbf{x}'_1), \dots, h(\mathbf{x}'_N)) \mu_t^{(l)}(dh) \right| \xrightarrow{a.s.} 0.$$

Remark 19 *If $L \geq 4$, then for $2 \leq l \leq L-2$, the random field $\mathbf{H}_t^{(l)}$ defined through the above is actually deterministic for all $t \geq 0$, indicating a type of degeneracy in deep NNs under the mean-field scaling (Araújo et al., 2019; Nguyen and Pham, 2023). Randomness can be restored if we add a bias term to each layer that is randomly initialized (see Appendix B.1.2).*

Theorem 18 is proved in Appendix D.1, including in the more general case where the bias terms are added. The proof consists of two steps. First, we rewrite the limiting dynamics in a differential form using the formulation introduced in Section 3.2. Second, we prove the convergence of the dynamics as $m \rightarrow \infty$ via a propagation-of-chaos-type argument (McKean, 1967; Braun and Hepp, 1977), which also appears in several prior works on the mean-field limit of multi-layer NNs (Araújo et al., 2019; Pham and Nguyen, 2021; Sirignano and Spiliopoulos, 2022). Note that our result here is meant not to improve upon these prior works in techniques, but rather to demonstrate that the training dynamics in the mean-field

limit can be interpreted as a learning dynamics of the NHL. Specifically, as a corollary of Lemma 6, we see that the limiting dynamics indeed defines an evolving NHL (hence named the *mean-field NHL dynamics*), with each $\mathcal{H}_t^{(l+1)}$ being the RKHS associated with the kernel function $\kappa_t^{(l)} := \kappa_{t,t}^{(l)}$:

Proposition 20 *Suppose Assumptions 15 – 17 hold. $\forall t \geq 0$, f_t belongs to the NHL of $(\mathcal{H}_t^{(l)})_{l \in [L]}$, where we define $\mathcal{H}_t^{(1)} := \mathcal{H}^{(1)}$, and $\forall l \in [L - 1]$, $\mathcal{H}_t^{(l+1)} := \mathcal{H}_{\mu_t}^{(l)}$ as the RKHS associated with $\kappa_t^{(l)}$. Moreover, $\forall l \in [L - 1]$, as $m \rightarrow \infty$, $\kappa_{m,t}^{(l)}$ converges to $\kappa_t^{(l)}$ almost surely.*

Remark 21 *Unlike the mean-field limit of shallow NNs or a usual gradient descent dynamics, the dynamics of each $\mathbf{H}_t^{(l)}$ is notably no longer Markovian when $L > 2$. The memory effect emerges because we track the evolution of the neurons’ pre-activation functions rather than that of their weight parameters. Similar phenomena are also seen in an alternative infinite-width limit studied by Yang and Hu (2021); Bordelon and Pehlevan (2022),*

From the mean-field NHL dynamics, we derive that the output function satisfies

$$\frac{d}{dt} f_t(\mathbf{x}) = \mathcal{E}_{\mathbf{x}'} \left\{ \zeta_t(\mathbf{x}') \theta_t(\mathbf{x}, \mathbf{x}') \right\} , \quad (14)$$

where we define

$$\theta_t(\mathbf{x}, \mathbf{x}') := \sum_{l=1}^L \kappa_t^{(l-1)}(\mathbf{x}, \mathbf{x}') \gamma_t^{(l)}(\mathbf{x}, \mathbf{x}') ,$$

and we set $\kappa_t^{(0)}(\mathbf{x}, \mathbf{x}') := \mathbf{x}^\top \cdot \mathbf{x}'$, $\gamma_t^{(L)}(\mathbf{x}, \mathbf{x}') := 1$ and $\gamma_t^{(l)} := \gamma_{t,t}^{(l)}$, $\forall l \in [L - 1]$. Hence, one can view f_t as evolving according to a *functional gradient flow*—also called the residual dynamics in the mean-field theory of shallow NNs (Rotskoff and Vanden-Eijnden, 2018)—under a *time-varying* and *data-dependent* kernel function θ_t . It is analogous to the Neural Tangent Kernel (NTK) that governs the training dynamics of infinite-width NNs under a difference choice of scaling (Jacot et al. (2018); see the discussions in Section 8). However, a crucial difference is that the NTK remains *fixed* during training and agnostic to the training data—and hence the name “lazy training” (Chizat and Bach, 2018)—whereas the model here exhibits feature learning as θ_t evolves during training based on the data, as summarized in criterion (v).

7.1.1 EXAMPLE 1: $L = 2$

When $L = 2$, the mean-field NHL dynamics is Markovian and reduces to the following ODEs of the random variables \mathbf{A}_t and \mathbf{Z}_t :

$$\begin{aligned} \frac{d}{dt} \mathbf{A}_t &= \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \sigma(\mathbf{Z}_t^\top \cdot \mathbf{x}) \right\} , \\ \frac{d}{dt} \mathbf{Z}_t &= \mathbf{A}_t \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \sigma'(\mathbf{Z}_t^\top \cdot \mathbf{x}) \mathbf{x} \right\} . \end{aligned}$$

In this case, one can fully characterize the NHL dynamics by the evolution of the joint law of \mathbf{A}_t and \mathbf{Z}_t , which follows a Wasserstein gradient flow (WGF) in $\mathcal{P}(\mathbb{R}^{d+1})$, *without* an underlying probability space for defining the random variables. This allows us to recover the mean-field theory of shallow NNs, which we review in Section 8. In particular, under the

same assumptions, the global convergence guarantees of the WGF (Nitanda and Suzuki, 2017; Chizat and Bach, 2018; Mei et al., 2018; Rotskoff and Vanden-Eijnden, 2018; Wojtowytsch, 2020) also apply to the mean-field NHL dynamics at $L = 2$.

Note that when $L > 2$, the probability-space-free representation becomes less convenient. On one hand, the random fields $\mathbf{H}_t^{(l)}$ can no longer be reduced to finite-dimensional random variables when $l > 1$. On the other hand, the mean-field NHL dynamics (12) becomes non-Markovian. For these reasons, inspired by the approaches taken by Nguyen and Pham (2023); Fang et al. (2021), we choose to formulate the general mean-field dynamics through the random field representation.

7.1.2 EXAMPLE 2: DEEP LINEAR NN

When σ is the identity function, the model becomes a linear NN, whose output function can be written as $f_t(\mathbf{x}) = \mathbf{v}_t^\top \cdot \mathbf{x}$. Moreover, for all $l \in \{2, \dots, L\}$, $\mathcal{H}_t^{(l)}$ always contains the same set of functions, namely, the linear functions on \mathbb{R}^d , except that their norms in $\mathcal{H}_t^{(l)}$, which are governed by the kernel function $\kappa_t^{(l-1)}$, can differ with l and evolve over time.

Below, we show the mean-field NHL dynamics reduces in this case to a finite-dimensional system. We consider the setting of fitting a linear target function $f^*(\mathbf{x}) = (\mathbf{v}^*)^\top \cdot \mathbf{x}$ with least squares regression, and we define $\Sigma := \frac{1}{n} \sum_{k=1}^n \mathbf{x}_k \cdot \mathbf{x}_k^\top$ and $\zeta_t := \Sigma \cdot (\mathbf{v}_t - \mathbf{v}^*)$. We assume for simplicity that ρ_a has zero mean and unit variance while ρ_z has the zero vector as its mean and the identity matrix as its covariance.

Thanks to the linearity, each $\kappa_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}')$ is bilinear in \mathbf{x} and \mathbf{x}' while $\gamma_{l,t}^{(s)}(\mathbf{x}, \mathbf{x}')$ does not depend on \mathbf{x} or \mathbf{x}' . In other words, $\exists K_{t,s}^{(l)} \in \mathbb{R}^{d \times d}$ and $c_{t,s}^{(l)} \in \mathbb{R}$ such that $\kappa_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\top \cdot K_{t,s}^{(l)} \cdot \mathbf{x}'$ and $\gamma_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') = c_{t,s}^{(l)}$. Then, (14) reduces to

$$\frac{d}{dt} \mathbf{v}_t = \left(\sum_{l=1}^L c_t^{(l)} K_t^{(l-1)} \right) \cdot \Sigma \cdot (\mathbf{v}_t - \mathbf{v}^*) ,$$

where we set $K_t^{(0)} := I_d$, $c_t^{(L)} := 1$, and each $K_t^{(l-1)} := K_{t,t}^{(l-1)}$, $c_t^{(l)} := c_{t,t}^{(l)}$. Furthermore, the mean-field NHL dynamics reduces to closed equations in $K_{t,s}^{(l)}$ and $c_{t,s}^{(l)}$. For example, $\forall l \in [L - 2]$, it holds that

$$K_{t,s}^{(l+1)} = \int_0^t \int_0^s c_{r,p}^{(l+1)} K_{t,r}^{(l)} \cdot \zeta_r \cdot \zeta_p^\top \cdot K_{p,s}^{(l)} dp dr , \quad (15)$$

$$c_{t,s}^{(l)} = \int_0^t \int_0^s c_{t,r}^{(l+1)} c_{s,p}^{(l+1)} \zeta_r^\top \cdot K_{r,p}^{(l)} \cdot \zeta_p dp dr , \quad (16)$$

The full system of equations is derived in Appendix D.2. We see that although f_t is always a linear function (of its input), its training dynamics is nonlinear and non-Markovian, which is in contrast with, for example, the GF dynamics of plain linear regression:

$$\frac{d}{dt} \mathbf{v}_t = \Sigma \cdot (\mathbf{v}_t - \mathbf{v}^*) .$$




Figure 2: **Learning trajectories of linear 3-layer NN versus the NHL dynamics.** *Solid*: 3-layer linear NNs trained by GD with width 64 and 8192. *Dashed*: numerical integration of the NHL dynamics derived in Section 7.1.2. *Dot-dashed*: linear regression (LR) under population loss.

7.2 Numerical Illustrations

7.2.1 EXPERIMENT 1: LINEAR NN

To validate the NHL dynamics derived above for linear NNs, we compare its numerical solution with the GD training of an actual 3-layer linear NN on a least squares regression task of learning a linear target \mathbf{v}^* , as described in Section 7.1.2. We choose $d = 10$, $n = 50$ and $\nu = \mathcal{N}(0, I_d)$. In Figure 2, we plot the learning trajectories in the linear model space projected into the first two dimensions, i.e., $v_{t,1}$ and $v_{t,2}$. We see that the NHL dynamics solved by numerical integration closely predicts the actual GD dynamics when the width is large. Moreover, the NHL dynamics presents a *nonlinear* learning trajectory in the space of *linear* models, which is in contrast with the linear learning trajectory of performing linear regression under the population loss.

7.2.2 EXPERIMENT 2: RELU NN

To gain insights into feature learning and the evolution of the NHL through training, we perform GD on 3-layer NNs with the ReLU activation on a least squares regression task. We choose $d = 1$, $n = 20$, $m = 512$, the target function being $f^*(\mathbf{x}) = \sin(2\mathbf{x})$, and ν being the uniform distribution on $[0, 2\pi]$. All parameters in the model, including untrained bias terms, are sampled i.i.d. from $\mathcal{N}(0, 1)$ at initialization.

We see from Figure 3(b) that the pre-activation values across all neurons in the second hidden layer—which correspond to $\mathbf{H}_{m,t}^{(2)}$ and approximate $\mathbf{H}_t^{(2)}$ —move substantially through training, suggesting the occurrence of feature learning. In fact, as shown in Figure 3(c), the movement results in a learned kernel function $\kappa_t^{(2)}$ that bears the same periodicity as the target function, showing that the kernel function is *adaptive* through training. In particular, as measured by the Centered Kernel Alignment (CKA) score (Cortes et al., 2012), $\kappa_t^{(2)}$ becomes more *aligned* with the target function during training—an important notion in the literature of learning kernels (Cristianini et al., 2001; Canatar and Pehlevan, 2022)—and more so than $\kappa_t^{(1)}$. It suggests that the space $\mathcal{H}_t^{(L)}$ can move *closer* to the target function via training, though a theoretical explanation for the alignment phenomenon is still lacking.

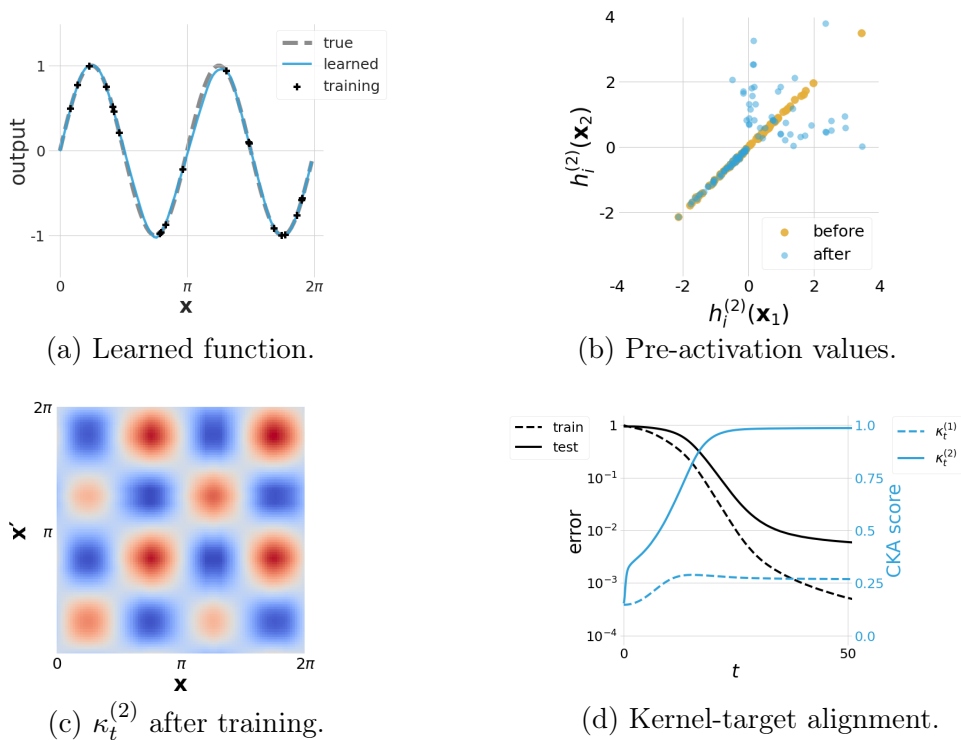


Figure 3: **Results of GD training of 3-layer NN with ReLU activation.** (a): Target versus learned function. (b): Pre-activation values across neurons in the second hidden layer on two training data points, \mathbf{x}_1 and \mathbf{x}_2 , before and after training. (c): The kernel function of the second hidden layer, $\kappa_t^{(2)}(\mathbf{x}, \mathbf{x}')$, after training (red means a higher value). (d): Training and test errors and the CKA scores of $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ with respect to the target function over time, averaged over 10 runs.

8. Related Works

Below, we further discuss the novelty and significance of our work relative to prior literature.

Function space of width-unlimited NNs As discussed in Section 2.3, prior works have proposed the function space of shallow NNs based on a type of variation norm, which is known to control the generalization and approximation errors (Bach, 2017a; E et al., 2019; Siegel and Xu, 2020; Chen et al., 2020b). Other works have also established the regularity, spectral properties and representer theorems (Savarese et al., 2019; Ongie et al., 2020; Parhi and Nowak, 2021; Siegel and Xu, 2023; Bartolucci et al., 2023; Unser, 2023; Ardeshir et al., 2023; Boursier and Flammarion, 2023) of this space. Hence, for shallow NNs, a relatively complete picture has been established that covers approximation, generalization and optimization.

For multi-layer NNs, however, a satisfactory theory for the function space and the complexity measure is missing. The neural tree space (E and Wojtowysch, 2020) is an interesting proposal, but it does not correspond directly with the training of NNs since the neurons in this model do *not* have shared preceding layers, which leads to approximation error bounds that grow exponentially in the depth, as discussed in Section 4.2. Another line of

studies including Lee et al. (2017); Sonoda and Murata (2017); Bartlett et al. (2018); Zou et al. (2020a); Lu et al. (2020); Zou et al. (2020b); Sonoda et al. (2021); Dutordoir et al. (2021); Littwin et al. (2021); E et al. (2022); Ding et al. (2022); Hayou (2022); Parhi and Nowak (2022) has focused on NNs with “bottleneck” layers that have fixed width, which can be viewed as expressing certain compositions or accumulations of functions in the Barron space. Though these models are interesting in their own right, the bottleneck assumption is not often obeyed in practice. In contrast, our work allows all layers of the NN to have unlimited width, resulting in a fundamentally different way of generalizing the Barron space—not by merely composing functions from the Barron space, but rather requiring a hierarchical generation of the RKHS. A thorough comparison between them is left for future work.

Mean-field theory of NNs In the mean-field scaling, shallow NNs under training are analogous to an interacting particles system (Rotskoff and Vanden-Eijnden, 2018, 2022). Hence, as described in Section 7.1.1, its infinite-width limit can be modeled as an integral over a probability measure on the parameter space (or, via an alternative form, a signed Radon measure), which evolves according to a Wasserstein GF during training (Mei et al., 2018; Chizat and Bach, 2018; Sirignano and Spiliopoulos, 2020). Notably, under suitable conditions, the Wasserstein GF can be proved to converge to global minimizers of the loss (Nitanda and Suzuki, 2017; Mei et al., 2018; Chizat and Bach, 2018; Rotskoff and Vanden-Eijnden, 2018; Wojtowytsch, 2020; Chen et al., 2022b; Boursier et al., 2022).

Several works have then proposed mean-field-type models for multi-layer NNs via probability measures defined in different ways (Nguyen, 2019; Araújo et al., 2019; Nguyen and Pham, 2023; Pham and Nguyen, 2021; Fang et al., 2021; Sirignano and Spiliopoulos, 2022). In particular, Araújo et al. (2019); Nguyen and Pham (2023) proved law-of-large-numbers results similar to Theorem 18 for the convergence of finite-width NNs to the mean-field limit, while Nguyen and Pham (2023); Fang et al. (2021) adopted a distributional view of the hidden-layer neurons that inspires our model. Nonetheless, these works have not addressed the function space and complexity measure associated with these models or proved error bounds. Chen et al. (2022a) proposed a mean-field theory on functional space to study the training dynamics of a type of partially-trained three-layer NNs in the infinite-width limit and defined a complexity measure based on optimal transport distances for distributions on functional space. The point of view therein sets the path for the current work.

NTK theory If we replace the $1/m$ factor by $1/\sqrt{m}$ in (1), we arrive at what is commonly called the *NTK scaling* of NNs. As shown by Jacot et al. (2018), if we initialize the NN randomly and take $m \rightarrow \infty$ under this scaling, then the pre-activation functions in the hidden layers barely move throughout training, and thus, the GF dynamics can be well-approximated by its linearization around the initialization, which is described by a kernel GF with a *fixed* kernel (that is the NTK). In other words, the evolution of the output function can also be written as (14) except that the kernel function θ_t is now independent of t . Thanks to this simplification, gradient descent is proved to converge to global minimum at a linear rate for over-parameterized NNs in the NTK regime (Du et al., 2019b,a; Allen-Zhu et al., 2019b; Zou et al., 2020a; Oymak and Soltanolkotabi, 2020). Furthermore, generalization guarantees can be proved for such models through the learning theory of RKHS (Arora et al., 2019b; Cao and Gu, 2019; E et al., 2020a).

However, the fact that the hidden layer neurons and hence the kernel function remain fixed to their initialization indicates a lack of *feature learning*. For this reason, the NTK limit is described as a regime of “*lazy training*” (Chizat et al., 2019; Woodworth et al., 2020), and the NTK theory does not satisfy criterion (v). Several works have studied the limitations of the NTK regime compared to feature-learning regimes, both theoretically (Ghorbani et al., 2019, 2020; Wei et al., 2019; Woodworth et al., 2020; Liu et al., 2020; Luo et al., 2021) and empirically (Geiger et al., 2020; Lee et al., 2020).

NNs as kernels Besides the NTK theory, a number of prior works have also explored the connection between NNs and kernels, by either proposing new kernel methods inspired by NNs (Cho and Saul, 2009; Mairal et al., 2014; Wilson et al., 2016; Bietti and Mairal, 2017; Shankar et al., 2020; Radhakrishnan et al., 2024) or by modeling NNs as kernels (Rahimi and Recht, 2008b; Montavon et al., 2011; Hazan and Jaakkola, 2015; Anselmi et al., 2015; Domingos, 2020; Aitchison et al., 2021; Amid et al., 2022; Seroussi et al., 2023; Yang et al., 2023), or both. As an exemplar of the latter category, Daniely et al. (2016) proposed the conjugate kernel model of multi-layer NNs together with a random feature scheme for approximating the kernel (Daniely et al., 2017) and a theoretical guarantee for learning with stochastic gradient descent (SGD) in polynomial time (Daniely, 2017b). However, like other efforts to understand NNs as kernel methods, it does not satisfy properties (i) or (v) as a theory for the function space of multi-layer NNs. In fact, under our framework, the conjugate kernel space can be seen as a particular *fixed* NHL determined by the random initialization of the weights, which is therefore quite restrictive. In contrast, the function space $\mathcal{F}^{(L)}$ is not *one* RKHS but an infinite collection of them and hence richer.

In the NTK scaling, a randomly-initialized NN in the infinite-width limit can also be viewed representing a function sampled from a Gaussian process whose covariance function is connected to the NTK, thus leading to a Bayesian interpretation (Neal, 1996; Williams, 1996; Lee et al., 2018; Matthews et al., 2018; Garriga-Alonso et al., 2019; Borovkykh, 2018; Novak et al., 2019; Agrawal et al., 2020). In particular, Lee et al. (2019) showed that SGD training corresponds to a linear dynamics of the Gaussian process and mimics Bayesian inference. However, like the NTK theory (and in contrast with ours), this analysis relies on a linear approximation of the training dynamics close to initialization and hence does not exhibit feature learning, either. Though our NHL framework also involves random fields, the fundamental difference can be seen through the fact that, in our case, while the *hidden layers* are modeled as random fields, the *output function* is always deterministic.

Complexity measures of NNs With large numbers of parameters, NNs in practice often have enough capacity to fit data with even random labels (Zhang et al., 2017). Hence, to derive meaningful generalization bounds, complexity measures of NNs that do not depend on the network size have been proposed. For example, several complexity measures based on certain norms of the parameters have been proposed, for both shallow (Bartlett, 1998; Koltchinskii and Panchenko, 2002; Bartlett and Mendelson, 2002; Rosset et al., 2007; Cho and Saul, 2009) and multi-layer ones (Neyshabur et al., 2015; Bartlett et al., 2017), which lead to generalization bounds that are independent of the number of parameters. In particular, the group norm in Neyshabur et al. (2015) is closely related to the NHL norm proposed in our work, as the NHL norm of the function represented by an NN can be bounded by the group norm of the NN. Thus, the NHL norm can be viewed as a generalization of the group norm

to the continuous, width-unlimited setup via the NHL model. In practice, regularizing the parameter norms through weight decay has long been proposed as a strategy to improve the performance of NNs (e.g., Hinton, 1987; Krogh and Hertz, 1991; Lee et al., 2020).

Beyond lazy training There have been efforts on extending the NTK analysis beyond the lazy training regime by considering higher-order Taylor expansions of the GD dynamics or corrections to the NTK due to a finite width or a large depth (Allen-Zhu et al., 2019a; Huang and Yau, 2020; Bai and Lee, 2020; Hanin and Nica, 2020; Yaida, 2020; Roberts et al., 2022; Hanin, 2022), though the function space implication of these proposals is not clear. Meanwhile, several works have studied the effect of different scaling choices on the behavior of the infinite-width limit (Golikov, 2020; Luo et al., 2021; Zhou et al., 2022). In particular, Yang and Hu (2021) proposed another choice called the maximum-update scaling, which exhibits feature learning while avoiding the degeneracy of the mean-field scaling mentioned in Remark 19. A few works have studied the training dynamics in the infinite-width limit under this scaling (Yang and Hu, 2021; Golikov and Yang, 2022; Hajjar et al., 2021; Ba et al., 2022; Bordelon and Pehlevan, 2022; Chizat et al., 2022), but the function space induced by this model is unclear except when only the penultimate layer is trained (Chen et al., 2022a).

Training dynamics of deep linear NNs Many prior studies have examined the GD or GF dynamics of deep linear NNs (Saxe et al., 2014; Jacot et al., 2021), including deriving their global convergence guarantees (Kawaguchi, 2016; Du and Hu, 2019; Eftekhari, 2020; Bah et al., 2022) and implicit bias (Gunasekar et al., 2017; Ji and Telgarsky, 2019; Arora et al., 2019a; Gidel et al., 2019; Li et al., 2021). The infinite-width limit of deep linear NNs under the maximum-update scaling have been studied in Bordelon and Pehlevan (2022); Chizat et al. (2022). In particular, the latter work derived the limiting dynamics and implicit bias rigorously, and our Figure 2 is inspired by the figures therein. Meanwhile, we are not aware of prior studies on deep linear NNs in the mean-field limit.

Depth separation A number of works have studied the benefit of depth in NN approximation, with examples including Telgarsky (2016); Eldan and Shamir (2016); Daniely (2017a); Safran and Shamir (2017); Yarotsky (2017); Vardi and Shamir (2020); Venturi et al. (2022). While most of these studies quantifies the price of approximation by the width required in the NN, a few have also considered NNs in the infinite-width limit. For example, Ongie et al. (2019) constructs a function that does not belong to the Barron space under ReLU activation but can be represented by a finite-width three-layer NN, which is directly used in our proof of Proposition 13. In addition, Ren et al. (2023) shows a function that cannot be approximated by shallow NNs but can be learned by an infinite-width three-layer NN, albeit of the type mentioned above with a bottleneck layer.

9. Conclusions

In this work, we propose to model multi-layer NNs as ladders of RKHSs, thereby defining the function space of multi-layer NNs as an infinite union of hierarchically-generated RKHSs. It is associated with a complexity measure of functions which depends on the number of layers and governs the approximation and generalization errors for functions in this space. Moreover, the training of multi-layer NNs in a feature-learning regime translates to a particular learning

dynamics of the NHLs. Hence, our proposal emerges as a desirable candidate for the function space of multi-layer NNs and establishes a rigorous foundation for their further studies.

There are several limitations of our work, including the various assumptions on the activation function σ —the approximation and generalization guarantees in Sections 4 and 5 are most tailored for 1-homogeneous activation functions (such as ReLU), whereas the analysis of the training dynamics in Section 7 assumes that the activation function is differentiable (thus excluding ReLU)—and that we ignore the effects of finite step sizes and stochasticity in practical NN training. Meanwhile, our work opens up interesting directions for future research, including additional properties of the NHL space, approximation theory in other metrics and formulations (e.g. Klusowski and Barron, 2018; E et al., 2020b), generalization guarantees beyond the analysis of Rademacher complexity, and long-time behaviors of the mean-field NHL dynamics when $L > 2$. Importantly, we hope it sets the stage for quantitative and rigorous probes into the popular idea that deep NNs perform hierarchical learning (Poggio et al., 2003, 2020; Allen-Zhu and Li, 2023; Chen et al., 2020a).

Acknowledgements The author is indebted to Joan Bruna and Eric Vanden-Eijnden for earlier discussions and thankful to Pengning Chao and the anonymous reviewers for feedbacks on the manuscript.

References

- Devanshu Agrawal, Theodore Papamarkou, and Jacob Hinkle. Wide neural networks with bottlenecks are deep gaussian processes. *The Journal of Machine Learning Research*, 21(1):7056–7121, 2020.
- Laurence Aitchison, Adam Yang, and Sebastian W Ober. Deep kernel processes. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 130–140. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/aitchison21a.html>.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep (hierarchical) learning. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 4598–4598. PMLR, 12–15 Jul 2023. URL <https://proceedings.mlr.press/v195/allen-zhu23a.html>.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL <https://proceedings.neurips.cc/paper/2019/file/62dad6e273d32235ae02b7d321578ee8-Paper.pdf>.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019b.

- Ehsan Amid, Rohan Anil, Wojciech Kotłowski, and Manfred K Warmuth. Learning from randomly initialized neural network features. *arXiv preprint arXiv:2202.06438*, 2022.
- Fabio Anselmi, Lorenzo Rosasco, Cheston Tan, and Tomaso Poggio. Deep convolutional networks are hierarchical kernel machines. *arXiv preprint arXiv:1508.01084*, 2015.
- Dyego Araújo, Roberto I Oliveira, and Daniel Yukimura. A mean-field limit for certain deep neural networks. *arXiv preprint arXiv:1906.00193*, 2019.
- Navid ArdesHIR, Daniel J. Hsu, and Clayton H. Sanford. Intrinsic dimensionality and generalization properties of the r -norm inductive bias. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 3264–3303. PMLR, 12–15 Jul 2023. URL <https://proceedings.mlr.press/v195/ardeshir23a.html>.
- Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL <https://proceedings.neurips.cc/paper/2019/file/c0c783b5fc0d7d808f1d14a6e9c8280d-Paper.pdf>.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019b.
- Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=akddwRG6EGi>.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017a.
- Francis Bach. On the equivalence between kernel quadrature rules and random feature expansions. *The Journal of Machine Learning Research*, 18(1):714–751, 2017b.
- Bubacarr Bah, Holger Rauhut, Ulrich Terstiege, and Michael Westdickenberg. Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. *Information and Inference: A Journal of the IMA*, 11(1):307–353, 2022.
- Yu Bai and Jason 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>.
- Zhidong Bai and Jack W Silverstein. *Spectral analysis of large dimensional random matrices*, volume 20. Springer, 2010.

- Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- Peter Bartlett, Dave Helmbold, and Philip Long. Gradient descent with identity initialization efficiently learns positive definite linear transformations by deep residual networks. In *International conference on machine learning*, pages 521–530. PMLR, 2018.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005. ISSN 00905364. URL <http://www.jstor.org/stable/3448616>.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/b22b257ad0519d4500539da3c8bcf4dd-Paper.pdf>.
- P.L. Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Transactions on Information Theory*, 44(2):525–536, 1998. doi: 10.1109/18.661502.
- Francesca Bartolucci, Ernesto De Vito, Lorenzo Rosasco, and Stefano Vigogna. Understanding neural networks with reproducing kernel banach spaces. *Applied and Computational Harmonic Analysis*, 62:194–236, 2023.
- Yoshua Bengio, Nicolas Le Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex Neural Networks. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, NIPS’05, pages 123–130, Cambridge, MA, USA, 2005. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2976248.2976264>. event-place: Vancouver, British Columbia, Canada.
- Alberto Bietti and Francis Bach. Deep equals shallow for reLU networks in kernel regimes. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=aDjoksTpXOP>.
- Alberto Bietti and Julien Mairal. Invariance and stability of deep convolutional representations. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/38ed162a0dbef7b3fe0f628aa08b90e7-Paper.pdf>.
- Blake Bordelon and Cengiz Pehlevan. Self-consistent dynamical field theory of kernel evolution in wide neural networks. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=sipwrPCrIS>.

- Anastasia Borovykh. A gaussian process perspective on convolutional neural networks. *arXiv preprint arXiv:1810.10798*, 2018.
- Etienne Boursier and Nicolas Flammarion. Penalising the biases in norm regularisation enforces sparsity. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=JtIqG47DAQ>.
- Etienne Boursier, Loucas Pillaud-Vivien, and Nicolas Flammarion. Gradient flow dynamics of shallow relu networks for square loss and orthogonal inputs. *Advances in Neural Information Processing Systems*, 35:20105–20118, 2022.
- Werner Braun and K Hepp. The vlasov dynamics and its fluctuations in the $1/n$ limit of interacting classical particles. *Communications in mathematical physics*, 56(2):101–113, 1977.
- Haim Brézis. *Functional analysis, Sobolev spaces and partial differential equations*, volume 2. Springer, 2011.
- Abdulkadir Canatar and Cengiz Pehlevan. A kernel analysis of feature learning in deep neural networks. In *2022 58th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1–8, 2022. doi: 10.1109/Allerton49937.2022.9929375.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *Advances in neural information processing systems*, 32, 2019.
- Lin Chen and Sheng Xu. Deep neural tangent kernel and laplace kernel have the same {rkhs}. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=vK9WrZ0QYQ>.
- Minshuo Chen, Yu Bai, Jason D Lee, Tuo Zhao, Huan Wang, Caiming Xiong, and Richard Socher. Towards understanding hierarchical learning: Benefits of neural representations. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 22134–22145. Curran Associates, Inc., 2020a. URL <https://proceedings.neurips.cc/paper/2020/file/fb647ca6672b0930e9d00dc384d8b16f-Paper.pdf>.
- Zhengdao Chen, Grant Rotskoff, Joan Bruna, and Eric Vanden-Eijnden. A dynamical central limit theorem for two-layer neural networks. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Zhengdao Chen, Eric Vanden-Eijnden, and Joan Bruna. A functional-space mean-field theory of partially-trained three-layer neural networks. *arXiv preprint arXiv:2210.16286*, 2022a.
- Zhengdao Chen, Eric Vanden-Eijnden, and Joan Bruna. On feature learning in shallow and multi-layer neural networks with global convergence guarantees. In *International Conference on Learning Representations*, 2022b. URL <https://openreview.net/forum?id=PQTW3iG4sC->.

- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems*, pages 3036–3046, 2018.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2937–2947, 2019.
- Lénaïc Chizat, Maria Colombo, Xavier Fernández-Real, and Alessio Figalli. Infinite-width limit of deep linear neural networks. *arXiv preprint arXiv:2211.16980*, 2022.
- Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 342–350. Curran Associates, Inc., 2009. URL <http://papers.nips.cc/paper/3628-kernel-methods-for-deep-learning.pdf>.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *The Journal of Machine Learning Research*, 13:795–828, 2012.
- Nello Cristianini, John Shawe-Taylor, André Elisseeff, and Jaz Kandola. On kernel-target alignment. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL <https://proceedings.neurips.cc/paper/2001/file/1f71e393b3809197ed66df836fe833e5-Paper.pdf>.
- Felipe Cucker and Steve Smale. On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39(1):1–49, 2002.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Amit Daniely. Depth separation for neural networks. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 690–696. PMLR, 07–10 Jul 2017a. URL <https://proceedings.mlr.press/v65/daniely17a.html>.
- Amit Daniely. Sgd learns the conjugate kernel class of the network. *Advances in Neural Information Processing Systems*, 30, 2017b.
- Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. *Advances in neural information processing systems*, 29, 2016.
- Amit Daniely, Roy Frostig, Vineet Gupta, and Yoram Singer. Random features for compositional kernels. *arXiv preprint arXiv:1703.07872*, 2017.
- Zhiyan Ding, Shi Chen, Qin Li, and Stephen J. Wright. Overparameterization of deep resnet: Zero loss and mean-field analysis. *Journal of Machine Learning Research*, 23(48):1–65, 2022. URL <http://jmlr.org/papers/v23/21-0669.html>.

- Pedro Domingos. Every model learned by gradient descent is approximately a kernel machine. *arXiv preprint arXiv:2012.00152*, 2020.
- Simon Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning*, pages 1655–1664. PMLR, 2019.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1675–1685. PMLR, 09–15 Jun 2019a. URL <http://proceedings.mlr.press/v97/du19c.html>.
- Simon S. Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=S1eK3i09YQ>.
- Vincent Dutordoir, James Hensman, Mark van der Wilk, Carl Henrik Ek, Zoubin Ghahramani, and Nicolas Durrande. Deep neural networks as point estimates for deep gaussian processes. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=sv1anLvYsTd>.
- Weinan E and Stephan Wojtowytsch. On the banach spaces associated with multi-layer relu networks: Function representation, approximation theory and gradient descent dynamics. *arXiv preprint arXiv:2007.15623*, 2020.
- Weinan E and Stephan Wojtowytsch. Representation formulas and pointwise properties for barron functions. *Calculus of Variations and Partial Differential Equations*, 61(2):1–37, 2022.
- Weinan E, Chao Ma, and Lei Wu. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5), 2019.
- Weinan E, Chao Ma, and Lei Wu. A comparative analysis of optimization and generalization properties of two-layer neural network and random feature models under gradient descent dynamics. *Science China Mathematics*, 63(7):1235–1258, 2020a.
- Weinan E, Chao Ma, Lei Wu, and Stephan Wojtowytsch. Towards a mathematical understanding of neural network-based machine learning: What we know and what we don’t. *CSIAM Transactions on Applied Mathematics*, 1(4):561–615, 2020b. ISSN 2708-0579. doi: <https://doi.org/10.4208/csiam-am.SO-2020-0002>. URL http://global-sci.org/intro/article_detail/csiam-am/18538.html.
- Weinan E, Chao Ma, and Lei Wu. The barron space and the flow-induced function spaces for neural network models. *Constructive Approximation*, 55(1):369–406, 2022.
- Armin Eftekhari. Training linear neural networks: Non-local convergence and complexity results. In *International Conference on Machine Learning*, pages 2836–2847. PMLR, 2020.

- Ronen Eldan and Ohad Shamir. The power of depth for feedforward neural networks. In *Conference on learning theory*, pages 907–940, 2016.
- Cong Fang, Jason Lee, Pengkun Yang, and Tong Zhang. Modeling from features: a mean-field framework for over-parameterized deep neural networks. In *Conference on learning theory*, pages 1887–1936. PMLR, 2021.
- Adrià Garriga-Alonso, Carl Edward Rasmussen, and Laurence Aitchison. Deep convolutional networks as shallow gaussian processes. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bklfsi0cKm>.
- Mario Geiger, Stefano Spigler, Arthur Jacot, and Matthieu Wyart. Disentangling feature and lazy training in deep neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(11):113301, 2020.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Limitations of lazy training of two-layers neural network. In *Advances in Neural Information Processing Systems*, pages 9111–9121, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *Advances in Neural Information Processing Systems*, 33:14820–14830, 2020.
- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Eugene Golikov. Towards a general theory of infinite-width limits of neural classifiers. In *International Conference on Machine Learning*, pages 3617–3626. PMLR, 2020.
- Eugene Golikov and Greg Yang. Non-gaussian tensor programs. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=AcHUIG2wA8->.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 297–299. PMLR, 06–09 Jul 2018. URL <https://proceedings.mlr.press/v75/golowich18a.html>.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pages 6151–6159, 2017.
- Karl Hajjar, Lénaïc Chizat, and Christophe Giraud. Training integrable parameterizations of deep neural networks in the infinite-width limit. *arXiv preprint arXiv:2110.15596*, 2021.
- Boris Hanin. Correlation functions in random fully connected neural networks at finite width. *arXiv preprint arXiv:2204.01058*, 2022.

- Boris Hanin and Mihai Nica. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJgndT4KwB>.
- Soufiane Hayou. On the infinite-depth limit of finite-width neural networks. *Transactions on Machine Learning Research*, 2022. URL <https://openreview.net/forum?id=RbLsYz1Az9>.
- Tamir Hazan and Tommi Jaakkola. Steps toward deep kernel methods from infinite neural networks. *arXiv preprint arXiv:1508.05133*, 2015.
- Geoffrey E Hinton. Learning translation invariant recognition in a massively parallel networks. In *International conference on parallel architectures and languages Europe*, pages 1–13. Springer, 1987.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. ISSN 0893-6080. doi: [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8). URL <https://www.sciencedirect.com/science/article/pii/0893608089900208>.
- Jiaoyang Huang and Horng-Tzer Yau. Dynamics of deep neural networks and neural tangent hierarchy. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4542–4551. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/huang201.html>.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Arthur Jacot, François Ged, Berfin Şimşek, Clément Hongler, and Franck Gabriel. Saddle-to-saddle dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. *arXiv preprint arXiv:2106.15933*, 2021.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HJflg30qKX>.
- Nigel Kalton. Chapter 25 - quasi-banach spaces. volume 2 of *Handbook of the Geometry of Banach Spaces*, pages 1099–1130. Elsevier Science B.V., 2003. doi: [https://doi.org/10.1016/S1874-5849\(03\)80032-3](https://doi.org/10.1016/S1874-5849(03)80032-3). URL <https://www.sciencedirect.com/science/article/pii/S1874584903800323>.
- Kenji Kawaguchi. Deep learning without poor local minima. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/f2fc990265c712c49d51a18a32b39f0c-Paper.pdf>.

- Jason M. Klusowski and Andrew R. Barron. Approximation by combinations of relu and squared relu ridge functions with ℓ^1 and ℓ^0 controls. *IEEE Transactions on Information Theory*, 64(12):7649–7656, 2018. doi: 10.1109/TIT.2018.2874447.
- Vladimir Koltchinskii and Dmitry Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- Anders Krogh and John A. Hertz. A simple weight decay can improve generalization. In *Proceedings of the 4th International Conference on Neural Information Processing Systems*, NIPS’91, page 950–957, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers Inc. ISBN 1558602224.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 1991.
- Holden Lee, Rong Ge, Tengyu Ma, Andrej Risteski, and Sanjeev Arora. On the ability of neural nets to express distributions. In *Conference on Learning Theory*, pages 1271–1296. PMLR, 2017.
- Jaehoon Lee, Jascha Sohl-dickstein, Jeffrey Pennington, Roman Novak, Sam Schoenholz, and Yasaman Bahri. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1EA-M-0Z>.
- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.
- Jaehoon Lee, Samuel S. Schoenholz, Jeffrey Pennington, Ben Adlam, Lechao Xiao, Roman Novak, and Jascha Sohl-Dickstein. Finite versus infinite neural networks: an empirical study. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/ad086f59924ffffe0773f8d0ca22ea712-Abstract.html>.
- Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=AH0s7Sm5H7R>.
- Etai Littwin, Omid Saremi, Shuangfei Zhai, Vimal Thilak, Hanlin Goh, Joshua M Susskind, and Greg Yang. Implicit acceleration and feature learning in infinitely wide neural networks with bottlenecks. *arXiv preprint arXiv:2107.00364*, 2021.
- Chaoyue Liu, Libin Zhu, and Misha Belkin. On the linearity of large non-linear models: when and why the tangent kernel is constant. *Advances in Neural Information Processing Systems*, 33:15954–15964, 2020.

- Yiping Lu, Chao Ma, Yulong Lu, Jianfeng Lu, and Lexing Ying. A mean field analysis of deep resnet and beyond: Towards provably optimization via overparameterization from depth. In *International Conference on Machine Learning*, pages 6426–6436. PMLR, 2020.
- Tao Luo, Zhi-Qin John Xu, Zheng Ma, and Yaoyu Zhang. Phase diagram for two-layer relu neural networks at infinite-width limit. *Journal of Machine Learning Research*, 22(71): 1–47, 2021.
- Julien Mairal, Piotr Koniusz, Zaid Harchaoui, and Cordelia Schmid. Convolutional kernel networks. *Advances in neural information processing systems*, 27, 2014.
- V A Marčenko and L A Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457, apr 1967. doi: 10.1070/SM1967v001n04ABEH001994. URL <https://dx.doi.org/10.1070/SM1967v001n04ABEH001994>.
- Alexander G. de G. Matthews, Jiri Hron, Mark Rowland, Richard E. Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=H1-nGgWC->.
- Henry P McKean. Propagation of chaos for a class of non-linear parabolic equations. *Stochastic Differential Equations (Lecture Series in Differential Equations, Session 7, Catholic Univ., 1967)*, pages 41–57, 1967.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33): E7665–E7671, 2018.
- Shahar Mendelson. On the performance of kernel classes. *J. Mach. Learn. Res.*, 4(null): 759–771, dec 2003. ISSN 1532-4435.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- Grégoire Montavon, Mikio L Braun, and Klaus-Robert Müller. Kernel analysis of deep networks. *Journal of Machine Learning Research*, 12(9), 2011.
- Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg, 1996. ISBN 0387947248.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401. PMLR, 2015.
- Phan-Minh Nguyen. Mean field limit of the learning dynamics of multilayer neural networks. *arXiv preprint arXiv:1902.02880*, 2019.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multilayer neural networks. *Mathematical Statistics and Learning*, 6(3):201–357, 2023.

- Atsushi Nitanda and Taiji Suzuki. Stochastic particle gradient descent for infinite ensembles. *arXiv preprint arXiv:1712.05438*, 2017.
- Roman Novak, Lechao Xiao, Yasaman Bahri, Jaehoon Lee, Greg Yang, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-dickstein. Bayesian deep convolutional networks with many channels are gaussian processes. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1g30j0qF7>.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A Function Space View of Bounded Norm Infinite Width ReLU Nets: The Multivariate Case. *arXiv:1910.01635 [cs, stat]*, October 2019. URL <http://arxiv.org/abs/1910.01635>. arXiv: 1910.01635.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width relu nets: The multivariate case. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H11NPxHKDH>.
- Samet Oymak and Mahdi Soltanolkotabi. Toward moderate overparameterization: Global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 1(1):84–105, 2020. doi: 10.1109/JSAIT.2020.2991332.
- Rahul Parhi and Robert D Nowak. Banach space representer theorems for neural networks and ridge splines. *J. Mach. Learn. Res.*, 22(43):1–40, 2021.
- Rahul Parhi and Robert D. Nowak. What kinds of functions do deep neural networks learn? insights from variational spline theory. *SIAM Journal on Mathematics of Data Science*, 4(2): 464–489, 2022. doi: 10.1137/21M1418642. URL <https://doi.org/10.1137/21M1418642>.
- Jiyoung Park, Ian Pelakh, and Stephan Wojtowytsch. Minimum norm interpolation by perceptrs: Explicit regularization and implicit bias. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=MlrFYNo1yc>.
- Huy Tuan Pham and Phan-Minh Nguyen. Global convergence of three-layer neural networks in the mean field regime. *ICLR*, 2021.
- Tomaso Poggio, Steve Smale, et al. The mathematics of learning: Dealing with data. *Notices of the AMS*, 50(5):537–544, 2003.
- Tomaso Poggio, Andrzej Banburski, and Qianli Liao. Theoretical issues in deep networks. *Proceedings of the National Academy of Sciences*, 117(48):30039–30045, 2020. doi: 10.1073/pnas.1907369117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1907369117>.
- Adityanarayanan Radhakrishnan, Daniel Beaglehole, Parthe Pandit, and Mikhail Belkin. Mechanism for feature learning in neural networks and backpropagation-free machine learning models. *Science*, 2024.
- Ali Rahimi and Benjamin Recht. Uniform approximation of functions with random bases. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 555–561. IEEE, 2008a.

- Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. *Advances in neural information processing systems*, 21, 2008b.
- Yunwei Ren, Mo Zhou, and Rong Ge. Depth separation with multilayer mean-field networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=uzFQpkEz0o>.
- Daniel A. Roberts, Sho Yaida, and Boris Hanin. *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks*. Cambridge University Press, 2022. doi: 10.1017/9781009023405.
- Saharon Rosset, Grzegorz Swirszcz, Nathan Srebro, and Ji Zhu. l_1 regularization in infinite dimensional feature spaces. In *International Conference on Computational Learning Theory*, pages 544–558. Springer, 2007.
- Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In *Advances in Neural Information Processing Systems*, pages 7146–7155, 2018.
- Grant Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of artificial neural networks: An interacting particle system approach. *Communications on Pure and Applied Mathematics*, 75(9):1889–1935, 2022. doi: <https://doi.org/10.1002/cpa.22074>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.22074>.
- Itay Safran and Ohad Shamir. Depth-width tradeoffs in approximating natural functions with neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2979–2987. JMLR. org, 2017.
- Saburo Saitoh. *Integral transforms, reproducing kernels and their applications*, volume 369. CRC Press, 1997.
- Saburo Saitoh and Masahiro Yamamoto. *Integral Transforms Involving Smooth Functions*, pages 149–163. Springer US, Boston, MA, 1999. ISBN 978-1-4757-2987-0. doi: 10.1007/978-1-4757-2987-0_14. URL https://doi.org/10.1007/978-1-4757-2987-0_14.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Conference on Learning Theory*, pages 2667–2690, 2019.
- A Saxe, J McClelland, and S Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *Int. Conf. on Learning Representations (ICLR)*, 2014.
- Inbar Seroussi, Gadi Naveh, and Zohar Ringel. Separation of scales and a thermodynamic description of feature learning in some cnns. *Nature Communications*, 14(1):908, 2023.
- Vaishal Shankar, Alex Fang, Wenshuo Guo, Sara Fridovich-Keil, Jonathan Ragan-Kelley, Ludwig Schmidt, and Benjamin Recht. Neural kernels without tangents. In *International Conference on Machine Learning*, pages 8614–8623. PMLR, 2020.

- Jonathan W Siegel and Jinchao Xu. Approximation rates for neural networks with general activation functions. *Neural Networks*, 128:313–321, 2020.
- Jonathan W Siegel and Jinchao Xu. Characterization of the variation spaces corresponding to shallow neural networks. *Constructive Approximation*, pages 1–24, 2023.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 47(1):120–152, 2022.
- Sho Sonoda and Noboru Murata. Double continuum limit of deep neural networks. In *ICML Workshop Principled Approaches to Deep Learning*, volume 1740, 2017.
- Sho Sonoda, Isao Ishikawa, and Masahiro Ikeda. Ghosts in neural networks: Existence, structure and role of infinite-dimensional null space. *arXiv preprint arXiv:2106.04770*, 2021.
- Matus Telgarsky. Benefits of depth in neural networks. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 1517–1539, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v49/telgarsky16.html>.
- Michael Unser. Ridges, neural networks, and the radon transform. *Journal of Machine Learning Research*, 24(37):1–33, 2023.
- Gal Vardi and Ohad Shamir. Neural networks with small weights and depth-separation barriers. *Advances in neural information processing systems*, 33:19433–19442, 2020.
- Luca Venturi, Samy Jelassi, Tristan Ozuch, and Joan Bruna. Depth separation beyond radial functions. *The Journal of Machine Learning Research*, 23(1):5309–5364, 2022.
- Colin Wei, Jason D. Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9709–9721, 2019.
- Christopher Williams. Computing with infinite networks. In M.C. Mozer, M. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9. MIT Press, 1996. URL <https://proceedings.neurips.cc/paper/1996/file/ae5e3ce40e0404a45ecacaaf05e5f735-Paper.pdf>.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378, 2016.
- Stephan Wojtowytsch. On the convergence of gradient descent training for two-layer relu-networks in the mean field regime. *arXiv preprint arXiv:2005.13530*, 2020.

- Stephan Wojtowytsch. Optimal bump functions for shallow relu networks: Weight decay, depth separation, curse of dimensionality. *Journal of Machine Learning Research*, 25(27): 1–49, 2024.
- Blake Woodworth, Suriya Gunasekar, Jason D. Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 3635–3673. PMLR, 09–12 Jul 2020. URL <https://proceedings.mlr.press/v125/woodworth20a.html>.
- Sho Yaida. Non-Gaussian processes and neural networks at finite widths. In Jianfeng Lu and Rachel Ward, editors, *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, volume 107 of *Proceedings of Machine Learning Research*, pages 165–192. PMLR, 20–24 Jul 2020. URL <https://proceedings.mlr.press/v107/yaida20a.html>.
- Adam X. Yang, Maxime Robeyns, Edward Milsom, Ben Anson, Nandi Schoots, and Laurence Aitchison. A theory of representation learning gives a deep generalisation of kernel methods. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 39380–39415. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/yang23k.html>.
- Greg Yang and Edward J. Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11727–11737. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21c.html>.
- Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.
- Hanxu Zhou, Zhou Qixuan, Zhenyuan Jin, Tao Luo, Yaoyu Zhang, and Zhi-Qin Xu. Empirical phase diagram for three-layer neural networks with infinite width. *Advances in Neural Information Processing Systems*, 35:26021–26033, 2022.
- Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep relu networks. *Machine Learning*, 109(3):467–492, 2020a.
- Difan Zou, Philip M. Long, and Quanquan Gu. On the global convergence of training deep linear resnets. In *International Conference on Learning Representations*, 2020b. URL <https://openreview.net/forum?id=HJxEhREKDH>.

Appendix A. Supplementary Materials for Section 3

A.1 Proof of Theorem 3(a)

For any function f , it is obvious by the property of L^p norms that $\mathcal{E}_p^{(L)}(f) \geq \mathcal{E}_{p'}^{(L)}(f)$ if $p \geq p'$.

Let $p \in (0, \infty]$. Suppose f is a function in $\mathcal{F}_p^{(L)}$. By definition, $\exists \mu^{(1)}, \dots, \mu^{(L-1)}$ such that $\|f\|_{\mathcal{H}^{(L)}} < \infty$ and $\forall l \in [L-1]$, $\|\mu^{(l)}\|_{\mathcal{H}^{(l)}, p} < \infty$ and $\mathcal{H}^{(l+1)} := \mathcal{H}_{\mu^{(l)}}$. Given any $c \in \mathbb{R}$, the function $cf \in \mathcal{H}^{(L)}$ with $\|cf\|_{\mathcal{H}^{(L)}} = |c|\|f\|_{\mathcal{H}^{(L)}}$, which implies that cf belongs to the same NHL as f and $\mathcal{E}_p^{(L)}(cf) \leq |c|\mathcal{E}_p^{(L)}(f) < \infty$. This shows that $\mathcal{F}_p^{(L)}$ is closed under scalar multiplication. Meanwhile, $\mathcal{E}_p^{(L)}(f) = \mathcal{E}_p^{(L)}(c^{-1}cf) \leq |c|^{-1}\mathcal{E}_p^{(L)}(cf)$. As a result, $\mathcal{E}_p^{(L)}(cf) = |c|\mathcal{E}_p^{(L)}(f)$. This proves the absolute homogeneity of $\mathcal{E}_p^{(L)}$.

Let f' be another function in $\mathcal{F}_p^{(L)}$. Similarly, by definition, $\exists \mu^{(1)'}, \dots, \mu^{(L-1)'}$ such that $\|f'\|_{\mathcal{H}^{(L)'}} < \infty$, $\mathcal{H}^{(1)'} = \mathcal{H}^{(1)}$ and $\forall l \in [L-1]$, $\|\mu^{(l)'}\|_{\mathcal{H}^{(l)', p}} < \infty$ and $\mathcal{H}^{(l+1)'} := \mathcal{H}_{\mu^{(l)'}}$. Then, to show that $\mathcal{F}_p^{(L)}$ is a vector space, we need an upper bound on $\mathcal{E}_p^{(L)}(f + f')$.

For $l \in [L-1]$, we define $\tilde{\mu}^{(l)} := \frac{1}{2}\mu^{(l)} + \frac{1}{2}\mu^{(l)'}$. Thus, $\tilde{\mu}^{(l)}$ is supported within $\mathcal{H}^{(l)} \cup \mathcal{H}^{(l)'}$. We define $\tilde{\mathcal{H}}^{(1)} := \mathcal{H}^{(1)}$ and $\forall l \in [L-1]$, $\tilde{\mathcal{H}}^{(l+1)} := \mathcal{H}_{\tilde{\mu}^{(l)}}$, and we will show that $f + f'$ belongs to the NHL formed by $(\tilde{\mathcal{H}}^{(l)})_{l \in [L]}$. To do so, we need the following lemma:

Lemma 22 *Let $\mu_1, \mu_2 \in \mathcal{P}(\mathcal{C})$. Suppose μ_1 is absolutely continuous with respect to μ_2 and the Radon-Nikodym derivative $\frac{d\mu_1}{d\mu_2}$ is upper-bounded by $M > 0$ on the support of μ_1 . Then it holds that $\mathcal{H}_{\mu_1} \subseteq \mathcal{H}_{\mu_2}$ with $\|\cdot\|_{\mathcal{H}_{\mu_2}} \leq \sqrt{M}\|\cdot\|_{\mathcal{H}_{\mu_1}}$.*

This lemma is proved in Appendix A.1.1. For each $l \in [L-1]$, noticing that $\mu^{(l)}$ (respectively, $\mu^{(l)'}$) are absolutely continuous with respect to $\tilde{\mu}^{(l)}$ with $\frac{d\mu^{(l)}}{d\tilde{\mu}^{(l)}} \leq 2$ on $\mathcal{H}^{(l)}$ (respectively, $\frac{d\mu^{(l)'}}{d\tilde{\mu}^{(l)}} \leq 2$ on $\mathcal{H}^{(l)'}$), we can apply Lemma 22 to obtain that,

$$\begin{aligned} \|\tilde{\mu}^{(l)}\|_{\tilde{\mathcal{H}}^{(l)}, p}^p &= \int \|h\|_{\tilde{\mathcal{H}}^{(l)}}^p \tilde{\mu}^{(l)}(dh) \\ &= \frac{1}{2} \int \|h\|_{\tilde{\mathcal{H}}^{(l)}}^p \mu^{(l)}(dh) + \frac{1}{2} \int \|h\|_{\tilde{\mathcal{H}}^{(l)}}^p \mu^{(l)'}(dh) \\ &\leq 2^{p/2-1} \int \|h\|_{\tilde{\mathcal{H}}^{(l)}}^p \mu^{(l)}(dh) + 2^{p/2-1} \int \|h\|_{\tilde{\mathcal{H}}^{(l)}}^p \mu^{(l)'}(dh), \end{aligned}$$

and in addition,

$$\begin{aligned} \|\tilde{\mu}^{(l)}\|_{\tilde{\mathcal{H}}^{(l)}, \infty} &= \max \left\{ \|\mu^{(l)}\|_{\tilde{\mathcal{H}}^{(l)}, \infty}, \|\mu^{(l)'}\|_{\tilde{\mathcal{H}}^{(l)}, \infty} \right\} \\ &\leq \sqrt{2} \max \left(\|\mu^{(l)}\|_{\mathcal{H}^{(l)}, \infty} + \|\mu^{(l)'}\|_{\mathcal{H}^{(l)'}, \infty} \right). \end{aligned}$$

Moreover, $\|f + f'\|_{\tilde{\mathcal{H}}^{(L)}} \leq \|f\|_{\tilde{\mathcal{H}}^{(L)}} + \|f'\|_{\tilde{\mathcal{H}}^{(L)}} \leq \sqrt{2}(\|f\|_{\mathcal{H}^{(L)}} + \|f'\|_{\mathcal{H}^{(L)'}})$. Therefore, $\forall p > 0$,

$$\begin{aligned} (\mathcal{C}_p^{(L)}(f + f'))^p &\leq \left(\prod_{l=1}^{L-1} \|\tilde{\mu}^{(l)}\|_{\tilde{\mathcal{H}}^{(l),p}}^p \right) \|f + f'\|_{\tilde{\mathcal{H}}^{(L)}}^p \\ &\leq 2^{pL/2-L+1} \left(\prod_{l=1}^{L-1} \left(\|\tilde{\mu}^{(l)}\|_{\mathcal{H}^{(l),p}}^p + \|\mu^{(l)'}\|_{\tilde{\mathcal{H}}^{(l),p}}^p \right) \right) (\|f\|_{\mathcal{H}^{(L)}} + \|f'\|_{\mathcal{H}^{(L)'}})^p \\ &< \infty, \end{aligned} \tag{17}$$

and in addition,

$$\begin{aligned} \mathcal{C}_\infty^{(L)}(f + f') &\leq \left(\prod_{l=1}^{L-1} \|\tilde{\mu}^{(l)}\|_{\tilde{\mathcal{H}}^{(l),\infty}} \right) \|f + f'\|_{\tilde{\mathcal{H}}^{(L)}} \\ &\leq 2^{L/2} \left(\prod_{l=1}^{L-1} \left(\|\mu^{(l)}\|_{\mathcal{H}^{(l),\infty}} + \|\mu^{(l)'}\|_{\tilde{\mathcal{H}}^{(l),\infty}} \right) \right) (\|f\|_{\tilde{\mathcal{H}}^{(L)}} + \|f'\|_{\tilde{\mathcal{H}}^{(L)}}) \\ &< \infty. \end{aligned}$$

Hence, $f + f' \in \mathcal{F}_p^{(L)}$, and this concludes the proof that $\mathcal{F}_p^{(L)}$ is a vector space, $\forall p \in (0, \infty]$.

A.1.1 PROOF OF LEMMA 22

Suppose $f \in \mathcal{H}_{\mu_1}$. By Lemma 6, $\exists \xi \in L^2(\mathcal{C}, \mu_1)$ such that $f = \int \xi(h) \sigma(h(\cdot)) \mu_1(dh)$ and $\|f\|_{\mathcal{H}_{\mu_1}} = \|\xi\|_{L^2(\mathcal{C}, \mu_1)}$. By the definition of Radon-Nikodym derivative, it then holds for all $\mathbf{x} \in \mathcal{X}$ that

$$\begin{aligned} f(\mathbf{x}) &= \int \xi(h) \sigma(h(\mathbf{x})) \frac{d\mu_1}{d\mu_2}(h) \mu_2(dh) \\ &= \int \left(\xi(h) \frac{d\mu_1}{d\mu_2}(h) \right) \sigma(h(\mathbf{x})) \mu_2(dh). \end{aligned}$$

Hence, there is

$$\begin{aligned} \|f\|_{\mathcal{H}_{\mu_2}}^2 &\leq \int \left| \xi(h) \frac{d\mu_1}{d\mu_2}(h) \right|^2 \mu_2(dh) \\ &= \int |\xi(h)|^2 \frac{d\mu_1}{d\mu_2}(h) \mu_1(dh) \\ &\leq M \|\xi\|_{L^2(\mathcal{C}, \mu_1)}^2 \\ &= M \|f\|_{\mathcal{H}_{\mu_1}}^2. \end{aligned}$$

A.2 Proof of Theorem 3(b)

In light of Theorem 3(a), it suffices to consider the case $p = 2$. The basic regularity properties of functions in RKHS given below is simple to show using the reproducing property and the Cauchy-Schwartz inequality:

Lemma 23 *Let \mathcal{H} be an RKHS on \mathcal{X} associated with the kernel function κ . Then $\forall f \in \mathcal{H}$,*

$$\begin{aligned} |f(\mathbf{x})| &\leq \|f\|_{\mathcal{H}}(\kappa(\mathbf{x}, \mathbf{x}))^{1/2}, \quad \forall \mathbf{x} \in \mathcal{X}, \\ |f(\mathbf{x}) - f(\mathbf{x}')| &\leq \|f\|_{\mathcal{H}}d_{\kappa}(\mathbf{x}, \mathbf{x}'), \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \end{aligned}$$

where we define $d_{\kappa}(\mathbf{x}, \mathbf{x}') := (\kappa(\mathbf{x}, \mathbf{x}) + \kappa(\mathbf{x}', \mathbf{x}') - 2\kappa(\mathbf{x}, \mathbf{x}'))^{1/2}$.

Another central observation is the following lemma, which is proved in Appendix A.2.1:

Lemma 24 *If σ is J_{σ} -Lipschitz and $\sigma(0) = 0$, then $\forall L \in \mathbb{N}$,*

$$\begin{aligned} \left(\kappa^{(L)}(\mathbf{x}, \mathbf{x})\right)^{1/2} &\leq (J_{\sigma})^L \|\mathbf{x}\| \prod_{l=1}^L \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2}, \quad \forall \mathbf{x} \in \mathcal{X}, \\ d_{\kappa^{(L)}}(\mathbf{x}, \mathbf{x}') &\leq (J_{\sigma})^L \|\mathbf{x} - \mathbf{x}'\| \prod_{l=1}^L \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2}, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \end{aligned}$$

Suppose $f \in \mathcal{F}^{(L)}$, and let $(\mathcal{H}^{(l)})_{l \in [L]}$ be an NHL to which it belongs. Then, Lemmas 23 and 24 allow us to derive that

$$|f(\mathbf{x})| \leq (J_{\sigma})^{L-1} \|\mathbf{x}\| \|f\|_{\mathcal{H}^{(L)}} \prod_{l=1}^{L-1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2}, \quad \forall \mathbf{x} \in \mathcal{X}, \quad (18)$$

$$|f(\mathbf{x}) - f(\mathbf{x}')| \leq (J_{\sigma})^{L-1} \|\mathbf{x} - \mathbf{x}'\| \|f\|_{\mathcal{H}^{(L)}} \prod_{l=1}^{L-1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2}, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}. \quad (19)$$

Theorem 3(b) thus follows if we minimize the right-hand side of (18) and (19) over all NHLs.

A.2.1 PROOF OF LEMMA 24

We can prove Lemma 24 inductively in L . By the definition of $\kappa^{(0)}$, there is

$$\begin{aligned} \kappa^{(0)}(\mathbf{x}, \mathbf{x}) &= \|\mathbf{x}\|^2, \quad \forall \mathbf{x} \in \mathcal{X}, \\ d_{\kappa^{(0)}}(\mathbf{x}, \mathbf{x}') &= \left(\|\mathbf{x}\|^2 + \|\mathbf{x}'\|^2 - 2\mathbf{x}^{\top} \cdot \mathbf{x}'\right)^{1/2} \\ &= \|\mathbf{x} - \mathbf{x}'\|, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}. \end{aligned}$$

Next, suppose that the statements of Lemma 24 hold for a certain $L \in \mathbb{N}$. Then, for $L + 1$, Lemma 23 together with the assumptions on σ implies that, $\forall \mathbf{x} \in \mathcal{X}$,

$$\begin{aligned}
\kappa^{(L+1)}(\mathbf{x}, \mathbf{x}) &\leq \int |\sigma(h(\mathbf{x}))|^2 \mu^{(L+1)}(dh) \\
&\leq (J_\sigma)^2 \int |h(\mathbf{x})|^2 \mu^{(L+1)}(dh) \\
&\leq (J_\sigma)^2 \left(\int \|h\|_{\mathcal{H}^{(L+1)}}^2 \mu^{(L+1)}(dh) \right) \kappa^{(L)}(\mathbf{x}, \mathbf{x}) \\
&\leq (J_\sigma)^2 \left(\int \|h\|_{\mathcal{H}^{(L+1)}}^2 \mu^{(L+1)}(dh) \right) (J_\sigma)^{2L} \|\mathbf{x}\|^2 \prod_{l=1}^L \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2}^2 \\
&= (J_\sigma)^{2L+2} \|\mathbf{x}\|^2 \prod_{l=1}^{L+1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2}^2,
\end{aligned}$$

and moreover, $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$,

$$\begin{aligned}
d_{\kappa^{(L+1)}}(\mathbf{x}, \mathbf{x}')^2 &\leq \int |\sigma(h(\mathbf{x})) - \sigma(h(\mathbf{x}'))|^2 \mu^{(L+1)}(dh) \\
&\leq (J_\sigma)^2 \int |h(\mathbf{x}) - h(\mathbf{x}')|^2 \mu^{(L+1)}(dh) \\
&\leq (J_\sigma)^2 \left(\int \|h\|_{\mathcal{H}^{(L+1)}}^2 \mu^{(L+1)}(dh) \right) (d_{\kappa^{(L)}}(\mathbf{x}, \mathbf{x}'))^2 \\
&\leq (J_\sigma)^2 \left(\int \|h\|_{\mathcal{H}^{(L+1)}}^2 \mu^{(L+1)}(dh) \right) (J_\sigma)^{2L} \|\mathbf{x} - \mathbf{x}'\|^2 \prod_{l=1}^L \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2}^2 \\
&\leq (J_\sigma)^{2L+2} \|\mathbf{x} - \mathbf{x}'\|^2 \prod_{l=1}^{L+1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2}^2.
\end{aligned}$$

which proves the statements for $L + 1$.

A.3 Proof of Theorem 3(c)

First, we show the equivalence for different $p \in [2, \infty]$. For any functions f , it is obvious that $\mathcal{E}_p^{(L)}(f) \geq \mathcal{E}_{p'}^{(L)}(f)$ if $p \geq p'$. Then, to prove Theorem 3(c), it suffices to show that for all f such that $\mathcal{E}_2^{(L)}(f) < \infty$, we can find an NHL $(\mathcal{H}^{(l)})_{l \in [L]}$ induced by $(\mu^{(l)})_{l \in [L-1]}$ such that $\|f\|_{\mathcal{H}^{(L)}} = \mathcal{E}_2^{(L)}(f)$ and $\forall l \in [L-1]$, $\mu^{(l)}$ is supported on $\hat{\mathcal{H}}^{(l)}$, the unit-norm sphere of $\mathcal{H}^{(l)}$ (in which case, $\|\mu^{(l)}\|_{\mathcal{H}^{(l)},p} = 1$, $\forall p \in [2, \infty]$).

To this end, we introduce the following lemma, which is proved in Appendix A.3.1:

Lemma 25 *Suppose that σ is 1-homogeneous. Let \mathcal{H} be a Hilbert space of functions on \mathcal{X} , and let $\hat{\mathcal{H}}$ denote the unit-norm sphere of \mathcal{H} . Given any $\mu \in \mathcal{P}(\mathcal{H})$ such that $\|\mu\|_{\mathcal{H},2} = 1$, there exists $\tilde{\mu} \in \mathcal{P}(\hat{\mathcal{H}})$ such that $\|\cdot\|_{\mathcal{H}_{\tilde{\mu}}} \leq \|\cdot\|_{\mathcal{H}_\mu}$.*

Suppose $\mathcal{E}_2^{(L)}(f) < \infty$. By definition, we can find $\mu^{(1)}, \dots, \mu^{(L-1)}$ such that $\mathcal{E}_2^{(L)}(f) = \|f\|_{\mathcal{H}^{(L)}} \prod_{l=1}^{L-1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2}$ and $\forall l \in [L-1]$, $\mathcal{H}^{(l+1)} = \mathcal{H}_{\mu^{(l)}}$. Thanks to the homogeneity of

σ , we may assume without loss of generality that $\|f\|_{\mathcal{H}^{(l)}} = \mathcal{C}_2^{(L)}(f)$ while $\forall l \in [L-1]$, $\|\mu^{(l)}\|_{\mathcal{H}^{(l)},2} = 1$. Then, assume for contradiction that for some $l \in [L-1]$, $\mu^{(l)}$ is not supported entirely within $\hat{\mathcal{H}}^{(l)}$. But Lemma 25 implies that there exists a probability measure $\tilde{\mu}^{(l)}$ supported within $\hat{\mathcal{H}}^{(l)}$ such that if we define $\tilde{\mathcal{H}}^{(l+1)} := \mathcal{H}_{\tilde{\mu}^{(l)}}$, then $\forall f' \in \mathcal{H}^{(l+1)}$, $\|f'\|_{\tilde{\mathcal{H}}^{(l+1)}} \leq \|f'\|_{\mathcal{H}^{(l+1)}}$. In particular, this implies that $\|\mu^{(l+1)}\|_{\tilde{\mathcal{H}}^{(l+1)},2} \leq \|\mu^{(l+1)}\|_{\mathcal{H}^{(l+1)},2}$. Hence, by replacing $\mathcal{H}^{(l+1)}$ with $\tilde{\mathcal{H}}^{(l+1)}$, we obtain another NHL which contains f and also realizes the minimization problem in $\mathcal{C}_2^{(L)}(f)$.

Applying this argument to each l , we see that $\mu^{(1)}, \dots, \mu^{(L-1)}$ can be chosen such that $\forall l \in [L-1]$, $\mu^{(l)}$ is supported within the unit-norm sphere of $\mathcal{H}^{(l)}$, and hence $\|\mu^{(l)}\|_{\mathcal{H}^{(l)},p} = \|\mu^{(l)}\|_{\mathcal{H}^{(l)},2} = 1, \forall p \in [2, \infty]$. This proves that $\forall p \in [2, \infty]$, $\mathcal{C}_p^{(L)} = \mathcal{C}_2^{(L)}$.

Having shown the equivalence of $\mathcal{C}_p^{(L)}$ for $p \in [2, \infty]$ if σ is 1-homogeneous, below we only need to consider the case $p = 2$. Let us first recall the definition of a quasi-Banach space (e.g., Kalton, 2003):

Definition 26 *Let \mathcal{U} be a vector space. A functional $\|\cdot\|_{\mathcal{U}} : \mathcal{U} \rightarrow [0, \infty)$ is called a quasi-norm on \mathcal{U} with a constant $K \geq 1$ if the following properties are satisfied:*

- *Positive-definiteness:* $\forall f \in \mathcal{U}, \|f\|_{\mathcal{U}} > 0$ if $f \neq 0$;
- *Absolute homogeneity:* $\forall f \in \mathcal{U}, \forall c \in \mathbb{R}, \|cf\|_{\mathcal{U}} = |c|\|f\|_{\mathcal{U}}$;
- *Modified triangle inequality:* $\forall f, g \in \mathcal{U}, \|f + g\|_{\mathcal{U}} \leq K(\|f\|_{\mathcal{U}} + \|g\|_{\mathcal{U}})$.

Then, \mathcal{U} is called a quasi-Banach space if it is equipped with a quasi-norm $\|\cdot\|_{\mathcal{U}}$ and is complete with respect to the topology induced by $\|\cdot\|_{\mathcal{U}}$.

The positive-definiteness of $\mathcal{C}_2^{(L)}$ is implied by Theorem 3(b) and its absolute homogeneity has been shown in Appendix A.1. To show that $\mathcal{C}_2^{(L)}$ satisfies the modified triangle inequality, we follow the construction in Appendix A.1. Given f and $f' \in \mathcal{F}_2^{(L)}$, since σ is 1-homogeneous, we may assume without loss of generality that $\|f\|_{\mathcal{H}^{(L)}} = \mathcal{C}_2^{(L)}(f)$, $\|f'\|_{\mathcal{H}^{(L)'}} = \mathcal{C}_2^{(L)}(f')$, and $\forall l \in [L-1]$, $\|\mu^{(l)}\|_{\mathcal{H}^{(l)},2} = \|\mu^{(l)'}\|_{\mathcal{H}^{(l)'},2} = 1$. Then, (17) can be tightened to yield

$$\mathcal{C}_2^{(L)}(f + f') \leq 2^{\frac{L}{2}} (\|f\|_{\mathcal{H}^{(L)}} + \|f'\|_{\mathcal{H}^{(L)'}}) = 2^{\frac{L}{2}} (\mathcal{C}_2^{(L)}(f) + \mathcal{C}_2^{(L)}(f')) ,$$

which means that $\mathcal{C}_2^{(L)}$ is a quasi-norm on $\mathcal{F}_2^{(L)}$ (with a constant $2^{\frac{L}{2}}$).

Hence, to prove that $\mathcal{F}_2^{(L)}$ is a quasi-Banach space, it remains to show that it is complete under $\mathcal{C}_2^{(L)}$ as the quasi-norm. Let $(f_k)_{k=1}^{\infty}$ be a Cauchy sequence in $\mathcal{F}_2^{(L)}$. By definition, $\forall \epsilon > 0, \exists N(\epsilon) \in \mathbb{N}_+$ such that $\forall n_1, n_2 > N(\epsilon), \mathcal{C}_2^{(L)}(f_{n_1} - f_{n_2}) < \epsilon$. For each $k \in \mathbb{N}_+$, we define $\epsilon_k := 2^{-(k+1)L}$ and $n_k := \lceil N(\epsilon_k) \rceil$. We then define $g_0 := f_{n_1}$ and $g_k := f_{n_{k+1}} - f_{n_k}$ for each $k \in \mathbb{N}_+$. Thus, by construction, it holds that $\mathcal{C}_2^{(L)}(g_k) < \epsilon_k, \forall k \in \mathbb{N}_+$. By the argument above, $\forall k \in \mathbb{N}$, there exists an NHL $(\mathcal{H}_k^{(l)})_{l \in [L]}$ induced by $(\mu_k^{(l)})_{l \in [L-1]}$ such that $\|g_k\|_{\mathcal{H}_k^{(L)}} = \mathcal{C}_2^{(L)}(g_k)$ and $\mu_k^{(l)}$ is supported within $\hat{\mathcal{H}}_k^{(l)}, \forall l \in [L-1]$.

By Theorem 3(b), when evaluated at any $\mathbf{x} \in \mathcal{X}$, the sequence $(f_k(\mathbf{x}))_{k=1}^{\infty}$ is a Cauchy sequence in \mathbb{R} and hence admits a limit, which allows us to define a function f_{∞} on \mathcal{X} by $f_{\infty}(\mathbf{x}) := \lim_{k \rightarrow \infty} f_k(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$. Thus, it remains to show that $f_{\infty} \in \mathcal{F}_2^{(L)}$.

$\forall l \in [L - 1]$, we define

$$\mu^{(l)} := (2^l - 1) \sum_{k=0}^{\infty} 2^{-(k+1)l} \mu_k^{(l)} .$$

It is straightforward to show that the series converges in total variation as a measure on \mathcal{C} . In addition, as $\sum_{k=0}^{\infty} 2^{-(k+1)l} = \frac{1}{2^l - 1}$, it further holds that

$$\mu^{(l)}(\mathcal{C}) = (2^l - 1) \sum_{k=0}^{\infty} 2^{-(k+1)l} \mu_k^{(l)}(\mathcal{H}_k^{(l)}) = 1.$$

This implies that $\mu^{(l)} \in \mathcal{P}(\mathcal{C})$ and allows us to define $\mathcal{H}^{(l+1)} := \mathcal{H}_{\mu^{(l)}}$ for each l . Thus, to prove that $f_{\infty} \in \mathcal{F}_2^{(L)}$, it suffices to show that $\|f_{\infty}\|_{\mathcal{H}^{(L)}} < \infty$ and $\forall l \in [L - 1]$, $\|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2} < \infty$.

Notice that each $\mu_k^{(l)}$ is absolutely continuous with respect to $\mu^{(l)}$ with the Radon-Nikodym derivative upper-bounded by $2^{(k+1)l}/(2^l - 1)$ on $\mathcal{H}_k^{(l)}$. Thus, Lemma 22 implies that $\|\cdot\|_{\mathcal{H}^{(l+1)}}^2 \leq 2^{(k+1)l} \|\cdot\|_{\mathcal{H}_k^{(l+1)}}^2 / (2^l - 1)$. Hence, it holds for each $l \in [L - 2]$ that

$$\begin{aligned} \|\mu^{(l+1)}\|_{\mathcal{H}^{(l+1)}, 2}^2 &= \int \|h\|_{\mathcal{H}^{(l+1)}}^2 \mu^{(l+1)}(dh) \\ &= (2^{l+1} - 1) \sum_{k=0}^{\infty} 2^{-(k+1)(l+1)} \int \|h\|_{\mathcal{H}^{(l+1)}}^2 \mu_k^{(l+1)}(dh) \\ &\leq \frac{2^{l+1} - 1}{2^l - 1} \sum_{k=0}^{\infty} 2^{-(k+1)} \|\mu_k^{(l+1)}\|_{\mathcal{H}_k^{(l+1)}}^2 \\ &\leq 3 . \end{aligned}$$

Meanwhile,

$$\begin{aligned} \|\mu^{(1)}\|_{\mathcal{H}^{(1)}, 2}^2 &= \int \|h\|_{\mathcal{H}^{(1)}}^2 \mu^{(1)}(dh) \\ &= \sum_{k=0}^{\infty} 2^{-(k+1)} \int \|h\|_{\mathcal{H}^{(1)}}^2 \mu_k^{(1)}(dh) \\ &= 1 . \end{aligned}$$

Therefore, $\forall l \in [L-1]$, it holds that $\|\mu^{(l)}\|_{\mathcal{H}^{(l)},2} < \infty$. Meanwhile, $\forall k \in \mathbb{N}$, Lemma 22 implies that $\|\cdot\|_{\mathcal{H}^{(L)}}^2 \leq 2^{(k+1)(L-1)} \|\cdot\|_{\mathcal{H}_k^{(L)}}^2 / (2^{L-1} - 1)$, and hence

$$\begin{aligned} \|f_\infty\|_{\mathcal{H}^{(L)}} &= \left\| \sum_{k=0}^{\infty} g_k \right\|_{\mathcal{H}^{(L)}} \\ &\leq \sum_{k=0}^{\infty} \|g_k\|_{\mathcal{H}^{(L)}} \\ &\leq \left(2^{(L-1)/2} \mathcal{C}_2^{(L)}(f_{n_1}) + 2^{(k+1)(L-1)/2} \sum_{k=1}^{\infty} \mathcal{C}_2^{(L)}(g_k) \right) / \sqrt{2^{L-1} - 1} \\ &\leq \left(2^{(L-1)/2} \mathcal{C}_2^{(L)}(f_{n_1}) + \sum_{k=1}^{\infty} 2^{-(k+1)} \right) / \sqrt{2^{L-1} - 1} \\ &< \infty . \end{aligned}$$

This finishes the proof that $f_\infty \in \mathcal{F}_2^{(L)}$.

A.3.1 PROOF OF LEMMA 25

Without loss of generality, we assume that $\mu(\{0\}) = 0$ (since otherwise we can replace μ with an $\mu' \in \mathcal{P}(\mathcal{H})$ such that $\mu'(\{0\}) = 0$, $\|\mu'\|_{\mathcal{H},2} \leq 1$ and $\|\cdot\|_{\mathcal{H}_{\mu'}} \leq \|\cdot\|_{\mathcal{H}_\mu}$). Let \mathbf{H} be any \mathcal{H} -valued random variable with law μ . Note that we can define a bijection between $\mathbb{R}_+ \times \hat{\mathcal{H}}$ and $\mathcal{H} \setminus \{0\}$ via the map $(c, \hat{h}) \mapsto c\hat{h}$, and we let $(\mathbf{C}, \hat{\mathbf{H}})$ denote the image of \mathbf{H} under the inverse of this map, which is a pair of random variables supported on $\mathbb{R}_+ \times \hat{\mathcal{H}}$. We first see that $\mathbb{E}[\mathbf{C}^2] = \mathbb{E}[\|\mathbf{H}\|_{\mathcal{H}}^2] = 1$.

We choose a $\hat{\mathcal{H}}$ -valued random variable, $\tilde{\mathbf{H}}$, whose law has a Radon-Nikodym derivative of $\mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}}]$ with respect to the law of $\hat{\mathbf{H}}$, i.e., $\forall \hat{h} \in \hat{\mathcal{H}}$, $[\text{Law}(\tilde{\mathbf{H}})](d\hat{h}) = \mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}} = \hat{h}] [\text{Law}(\hat{\mathbf{H}})](d\hat{h})$. We can verify that $\text{Law}(\tilde{\mathbf{H}})$ defined as such is indeed a probability measure on $\hat{\mathcal{H}}$ since $\mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}} = \hat{h}] \geq 0$, and moreover,

$$[\text{Law}(\tilde{\mathbf{H}})](\hat{\mathcal{H}}) = \int_{\hat{\mathbf{H}}} \mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}} = \hat{h}] [\text{Law}(\hat{\mathbf{H}})](d\hat{h}) = \mathbb{E}[\mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}}]] = \mathbb{E}[\mathbf{C}^2] = 1 .$$

Consider any function $f \in \mathcal{H}_\mu$. By Lemma 6 and the bijection between $\mathbb{R}_+ \times \hat{\mathcal{H}}$ and $\mathcal{H} \setminus \{0\}$, there exists a function $\xi : \mathbb{R}_+ \times \hat{\mathcal{H}} \rightarrow \mathbb{R}$ such that

$$f(\mathbf{x}) = \mathbb{E}[\xi(\mathbf{C}, \hat{\mathbf{H}})\sigma(\mathbf{H}(\mathbf{x}))] = \mathbb{E}[\mathbf{C}\xi(\mathbf{C}, \hat{\mathbf{H}})\sigma(\hat{\mathbf{H}}(\mathbf{x}))] ,$$

and $\mathbb{E}[\xi(\mathbf{C}, \hat{\mathbf{H}})^2] = \|f\|_{\mathcal{H}_\mu}^2$. Then, defining a function $\tilde{\xi} : \hat{\mathcal{H}} \rightarrow \mathbb{R}$ by $\forall \hat{h} \in \hat{\mathcal{H}}$,

$$\tilde{\xi}(\hat{h}) := \frac{\mathbb{E}[\mathbf{C}\xi(\mathbf{C}, \hat{\mathbf{H}}) | \hat{\mathbf{H}} = \hat{h}]}{\mathbb{E}[\mathbf{C}^2 | \hat{\mathbf{H}} = \hat{h}]} ,$$

we see that

$$\begin{aligned}
f(\mathbf{x}) &= \mathbb{E} \left[\mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \sigma(\hat{\mathbf{H}}(\mathbf{x})) \mid \hat{\mathbf{H}} \right] \right] \\
&= \mathbb{E} \left[\frac{\mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} \right]}{\mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} \right]} \sigma(\hat{\mathbf{H}}(\mathbf{x})) \mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} \right] \right] \\
&= \int_{\hat{\mathcal{H}}} \tilde{\xi}(\hat{h}) \sigma(\hat{h}(\mathbf{x})) \mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} = \hat{h} \right] [\text{Law}(\hat{\mathbf{H}})](d\hat{h}) \\
&= \int_{\hat{\mathcal{H}}} \tilde{\xi}(\hat{h}) \sigma(\hat{h}(\mathbf{x})) [\text{Law}(\tilde{\mathbf{H}})](d\hat{h}) \\
&= \mathbb{E} \left[\tilde{\xi}(\tilde{\mathbf{H}}) \sigma(\tilde{\mathbf{H}}(\mathbf{x})) \right].
\end{aligned}$$

Thus, there is

$$\begin{aligned}
\|f\|_{\mathcal{H}_{\text{Law}(\tilde{\mathbf{H}})}}^2 &\leq \mathbb{E} \left[|\tilde{\xi}(\tilde{\mathbf{H}})|^2 \right] \\
&= \int_{\hat{\mathcal{H}}} \left| \frac{\mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} = \hat{h} \right]}{\mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} = \hat{h} \right]} \right|^2 [\text{Law}(\tilde{\mathbf{H}})](d\hat{h}) \\
&= \int_{\hat{\mathcal{H}}} \left| \frac{\mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} = \hat{h} \right]}{\mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} = \hat{h} \right]} \right|^2 \mathbb{E}[\mathbf{C}^2 \mid \hat{\mathbf{H}} = \hat{h}] [\text{Law}(\hat{\mathbf{H}})](d\hat{h}) \\
&= \int_{\hat{\mathcal{H}}} \frac{\left| \mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} = \hat{h} \right] \right|^2}{\mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} = \hat{h} \right]} [\text{Law}(\hat{\mathbf{H}})](d\hat{h}) \\
&= \mathbb{E} \left[\frac{\left| \mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} \right] \right|^2}{\mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} \right]} \right]
\end{aligned}$$

By the Cauchy-Schwartz inequality, $\left| \mathbb{E} \left[\mathbf{C} \xi(\mathbf{C}, \hat{\mathbf{H}}) \mid \hat{\mathbf{H}} \right] \right|^2 \leq \mathbb{E} \left[\mathbf{C}^2 \mid \hat{\mathbf{H}} \right] \mathbb{E} \left[|\xi(\mathbf{C}, \hat{\mathbf{H}})|^2 \mid \hat{\mathbf{H}} \right]$. Hence,

$$\|f\|_{\mathcal{H}_{\text{Law}(\tilde{\mathbf{H}})}}^2 \leq \mathbb{E} \left[\mathbb{E} \left[|\xi(\mathbf{C}, \hat{\mathbf{H}})|^2 \mid \hat{\mathbf{H}} \right] \right] = \mathbb{E} \left[|\xi(\mathbf{C}, \hat{\mathbf{H}})|^2 \right] = \|f\|_{\mathcal{H}_\mu}^2.$$

A.4 Proof of Proposition 7

For each $l \in [L-1]$, let $\mathbf{H}^{(l)}$ be a $\mathcal{H}^{(l)}$ -valued random variable with law $\mu^{(l)}$, and let $\mathbf{H}^{(1)}, \dots, \mathbf{H}^{(L-1)}$ be distributed independently on a common probability space.

First, we define \mathbf{A} as follows. By Lemma 6, there exists $\xi \in L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)})$ such that $f(\mathbf{x}) = \int \xi(h) \sigma(h(\mathbf{x})) \mu^{(L-1)}(dh)$ and $\|f\|_{\mathcal{H}^{(L)}} = \|\xi\|_{L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)})}$. Then, we define $\mathbf{A} := \xi(\mathbf{H}^{(L-1)})$, which is measurable with respect to $\mathbf{H}^{(L-1)}$. Hence, (11) as well as the equality $\|f\|_{\mathcal{H}^{(L)}}^2 = \mathbb{E}[\mathbf{A}^2]$ are implied.

Next, $\forall l \in [L-2]$, we define $\Xi^{(l)}$ as follows. By Lemma 6, $\forall h \in \mathcal{H}^{(l+1)}$, $\exists \xi_h \in L^2(\mathcal{H}^{(l)}, \mu^{(l)})$ such that

$$h(\mathbf{x}) = \int \xi_h(h') \sigma(h'(\mathbf{x})) \mu^{(l)}(dh') , \quad (20)$$

and

$$\|h\|_{\mathcal{H}^{(l+1)}} = \|\xi_h\|_{L^2(\mathcal{H}^{(l)}, \mu^{(l)})} . \quad (21)$$

We denote the map $h \mapsto \xi_h$ by $\Xi^{(l)}$, i.e., $[\Xi^{(l)}(h)](h') := \xi_h(h')$ for $h \in \mathcal{H}^{(l+1)}$ and $h' \in \mathcal{H}^{(l)}$, and finally define $\Xi^{(l)} := [\Xi^{(l)}(\mathbf{H}^{(l+1)})](\mathbf{H}^{(l)})$, which is, by definition, measurable with respect to the sigma-algebra generated by $\mathbf{H}^{(l)}$ and $\mathbf{H}^{(l+1)}$. Then, (10) and the relation $\|\mathbf{H}^{(l+1)}\|_{\mathcal{H}^{(l+1)}}^2 = \mathbb{E}[(\Xi^{(l)})^2 | \mathbf{H}^{(l+1)}]$ are implied by (20) and (21).

A.5 Proof of Lemma 6

We will take advantage of the following general result on integral transforms and RKHS (Saitoh, 1997; Saitoh and Yamamoto, 1999):

Lemma 27 (Saitoh) *Given a Hilbert space \mathcal{H}_0 and a map $\varphi : \mathcal{X} \rightarrow \mathcal{H}_0$, we define the kernel function $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ by $\kappa(\mathbf{x}, \mathbf{x}') = \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle_{\mathcal{H}_0}$, and let \mathcal{H} denote the RKHS associated with κ . Then, a function f on \mathcal{X} belongs to \mathcal{H} if and only if $\exists \xi \in \mathcal{H}_0$ such that*

$$f(\mathbf{x}) = \langle \xi, \varphi(\mathbf{x}) \rangle_{\mathcal{H}_0} , \quad \forall \mathbf{x} \in \mathcal{X} , \quad (22)$$

and moreover,

$$\|f\|_{\mathcal{H}} = \inf \{ \|\xi\|_{\mathcal{H}_0} : \xi \in \mathcal{H}_0 \text{ such that (22) holds} \} ,$$

with the infimum achieved at a unique $\xi^* \in \mathcal{H}_0$.

For our purpose, we would like to apply Lemma 27 with $\mathcal{H}_0 = L^2(\mathcal{C}, \mu)$ and $\varphi : \mathcal{X} \rightarrow L^2(\mathcal{C}, \mu)$ defined such that $\varphi(\mathbf{x})$ is the function on \mathcal{C} defined by $h \mapsto \sigma(h(\mathbf{x}))$. Indeed, $L^2(\mathcal{C}, \mu)$ is equipped with the inner product $\langle \xi_1, \xi_2 \rangle_{L^2(\mathcal{C}, \mu)} := \int \xi_1(h) \xi_2(h) \mu(dh)$ and complete (Brézis 2011, Theorem 4.8), and hence it is a Hilbert space. Moreover, $\forall \mathbf{x} \in \mathcal{X}$, as σ is Lipschitz, there exists $J_\sigma > 0$ such that $|\sigma(h(\mathbf{x}))| \leq |\sigma(0)| + J_\sigma |h(\mathbf{x})| \leq |\sigma(0)| + J_\sigma \|h\|_\infty$, and hence $\varphi(\mathbf{x}) \in L^2(\mathcal{C}, \mu)$ by the assumption on μ (note that $\varphi(\mathbf{x})$ is a measurable function on \mathcal{C} since the evaluation functionals are measurable with respect to the Borel sigma-algebra on \mathcal{C}). Thus, we are able to prove Lemma 6 by applying Lemma 27 with \mathcal{H}_0 and φ defined as above.

Appendix B. Supplementary Materials for Section 4

B.1 Including the Bias Terms

Below, we provide a way to include the bias term into the framework proposed in the main text. We note that alternative ways to do so are also possible, such as by generalizing the homogeneous Barron space construction proposed by Park et al. (2023).

B.1.1 MULTI-LAYER NN

By including the bias term, we mean replacing (2) in the definition of the multi-layer NN by

$$h_i^{(l+1)}(\mathbf{x}) := b_i^{(l+1)} + \frac{1}{m} \sum_{j=1}^m W_{i,j}^{(l)} \sigma(h_j^{(l)}(\mathbf{x})) .$$

In the GF dynamics, the bias terms evolve according to the following ODE:

$$\frac{d}{dt} b_{i,t}^{(l)} = -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t} q_{i,t}^{(l)}(\mathbf{x}) \right\} ,$$

where β denotes the learning rate of the bias parameters relative to the weight parameters. If $\beta = 0$, for example, it corresponds to having untrained bias terms.

B.1.2 NHL

To incorporate the bias term, we replace Definition 2 of the NHL in the following way. We consider the map

$$\begin{aligned} \mathcal{C} \times \mathbb{R} &\rightarrow \mathcal{C} \\ (h, b) &\mapsto h(\cdot) + b , \end{aligned} \tag{23}$$

If $\mu \in \mathcal{P}(\mathcal{C} \times \mathbb{R})$, we let $\mu_+ \in \mathcal{P}(\mathcal{C})$ denote its push-forward under this map.

Definition 28 *Suppose each of $\mathcal{H}^{(2)}, \dots, \mathcal{H}^{(L)}$ is an RKHS on \mathcal{X} , and $\forall l \in [L-1]$, there exists $\mu^{(l)} \in \mathcal{P}(\mathcal{H}^{(l)} \times \mathbb{R})$ such that $\mathcal{H}^{(l+1)}$ is the RKHS associated with the kernel function*

$$\kappa^{(l)}(\mathbf{x}, \mathbf{x}') := \int \sigma(h(\mathbf{x}) + b) \sigma(h(\mathbf{x}') + b) \mu(dh, db) .$$

In other words, we can write $\mathcal{H}^{(l+1)} = \mathcal{H}_{\mu_+^{(l)}}$. Then, we say that $(\mathcal{H}^{(l)})_{l \in [L]}$ is an L -level NHL (with the bias terms included) induced by the sequence of probability measures, $(\mu^{(l)})_{l \in [L-1]}$; in addition, a function f on \mathcal{X} belongs to the NHL if $f \in \mathcal{H}^{(L)}$.

If \mathcal{H} is a Hilbert space and $\mu \in \mathcal{P}(\mathcal{H} \times \mathbb{R})$, we write $\|\mu\|_{\mathcal{H}, p, +} := (\int \|h\|_{\mathcal{H}}^p + |b|^p \mu(dh, db))^{1/p}$ for $p \in (0, \infty)$, and analogously for $p = \infty$. Given an RKHS \mathcal{H} , we then define

$$\mathcal{D}_p(\mathcal{H}, \mathcal{H}') := \inf_{\mu \in \mathcal{P}(\mathcal{H} \times \mathbb{R}), \mathcal{H}_\mu = \mathcal{H}'} \|\mu\|_{\mathcal{H}, p, +} .$$

Next, we can define the (L, p) -NHL complexity of a function similarly as through (5) and (6).

In addition, in place of Proposition 7, the representation of an NHL through coupled random fields can be defined in the following way.

Proposition 29 *Given an L -level NHL as in Definition 28, there exist $L-1$ random fields on \mathcal{X} , $\mathbf{H}^{(1)}, \dots, \mathbf{H}^{(L-1)}$, as well as $2L-2$ scalar random variables, $\mathbf{B}^{(1)}, \dots, \mathbf{B}^{(L-1)}$, $\Xi^{(1)}, \dots, \Xi^{(L-2)}$, and \mathbf{A} , which are defined on a common probability space and satisfy the following properties:*

- *The pairs $(\mathbf{H}^{(1)}, \mathbf{B}^{(1)}), \dots, (\mathbf{H}^{(L-1)}, \mathbf{B}^{(L-1)})$ are mutually independent, and $\forall l \in [L-1]$, $\mu^{(l)} = \text{Law}(\mathbf{H}^{(l)}, \mathbf{B}^{(l)})$;*
- *$\forall l \in [L-2]$, $\Xi^{(l)}$ is measurable with respect to the sigma-algebra generated by $\mathbf{H}^{(l)}, \mathbf{H}^{(l+1)}, \mathbf{B}^{(l)}$, and $\mathbf{B}^{(l+1)}$. Moreover,*

$$\mathbf{H}^{(l+1)}(\mathbf{x}) = \mathbb{E}[\Xi^{(l)} \sigma(\mathbf{H}^{(l)}(\mathbf{x}) + \mathbf{B}^{(l)}) | \mathbf{H}^{(l+1)}] , \tag{24}$$

where $\mathbb{E}[\cdot | \cdot]$ denotes the conditional expectation, and $\|\mathbf{H}^{(l+1)}\|_{\mathcal{H}^{(l+1)}}^2 = \mathbb{E}[(\Xi^{(l)})^2 | \mathbf{H}^{(l+1)}]$;

- \mathbf{A} is measurable with respect to the sigma-algebra generated by $\mathbf{H}^{(L-1)}$ and $\mathbf{B}^{(L-1)}$.
Moreover,

$$f(\mathbf{x}) = \mathbb{E}[\mathbf{A}\sigma(\mathbf{H}^{(L-1)}(\mathbf{x}) + \mathbf{B}^{(L-1)})] , \quad (25)$$

$$\text{and } \|f\|_{\mathcal{H}^{(L)}}^2 = \mathbb{E}[\mathbf{A}^2].$$

Furthermore, in the mean-field dynamics, the evolution of the bias term is governed by

$$\frac{d}{dt}\mathbf{B}_t^{(l)} = -\beta\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x})\mathbf{Q}_t^{(l)}(\mathbf{x})\sigma'(\mathbf{H}_t^{(l)}(\mathbf{x}) + \mathbf{B}_t^{(l)}) \right\}$$

B.2 Proof of Theorem 8

First, by the definition of $\mathcal{H}_m^{(1)}$, there is $\|\mu_m^{(1)}\|_{\mathcal{H}_m^{(1)},p} = M_{m,p}^{(1)}$. For $l \in [L-2]$, Lemma 6 implies $\|h_i^{(l+1)}\|_{\mathcal{H}_m^{(l+1)}}^2 \leq \frac{1}{m} \sum_{j=1}^m |W_{i,j}^{(l)}|^2$, and so $\|\mu_m^{(l+1)}\|_{\mathcal{H}_m^{(l+1)},p}^p = \frac{1}{m} \sum_{i=1}^m \|h_i^{(l+1)}\|_{\mathcal{H}_m^{(l+1)}}^p \leq (M_{m,p}^{(l+1)})^p$. Finally, Lemma 6 implies $\|f\|_{\mathcal{H}_m^{(L)},p} \leq M_{m,p}^{(L)}$. Together, they prove the theorem.

B.3 Proof of Theorem 9

Let f be a function on \mathcal{X} with $\mathcal{C}_{\infty}^{(L)}(f) < \infty$. By definition, there exist probability measures $\mu^{(1)}, \dots, \mu^{(L-1)}$ and deterministic functions $\Xi^{(1)}, \dots, \Xi^{(L-1)}$ satisfying the conditions specified in Appendix A.4. Our strategy will be to consider a random approximation of f using an L -layer NN with width m that achieves a low approximation error in expectation.

$\forall l \in [L-1]$, we let $\{\mathbf{H}_i^{(l)}\}_{i \in [m]}$ be m independent samples from $\mu^{(l)}$ in $\mathcal{H}^{(l)}$. For $l = 1$, we let $\bar{\mathbf{H}}_i^{(1)} := \mathbf{H}_i^{(1)}$. For $l \in [L-2]$, writing $\bar{\mathbf{W}}_{i,j}^{(l)} := \Xi^{(l)}(\mathbf{H}_i^{(l+1)}, \mathbf{H}_j^{(l)})$, we iteratively define

$$\bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) := \frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x})) ,$$

and finally, writing $\bar{\mathbf{A}}_i := \Xi^{(L-1)}(\mathbf{H}_i^{(L-1)})$, we define

$$\mathbf{F}_m(\mathbf{x}) := \frac{1}{m} \sum_{i=1}^m \bar{\mathbf{A}}_i \sigma(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x})) ,$$

and will bound its difference with f in expectation. We will use the two following lemmas:

Lemma 30 $\forall l \in [L-1], \forall i \in [m]$, almost surely,

$$\mathcal{E}_{\mathbf{x} \sim \nu} \left\{ \mathbb{E} \left[(\bar{\mathbf{H}}_i^{(l)}(\mathbf{x}) - \mathbf{H}_i^{(l)}(\mathbf{x}))^2 \mid \mathbf{H}_i^{(l)} \right] \right\} \leq \frac{(l-1)^2}{m} (J_{\sigma})^{2l-2} M_{\nu} \prod_{\nu'=1}^l \|\mu^{(\nu')}\|_{\mathcal{H}^{(\nu')},\infty}^2 . \quad (26)$$

Lemma 31 $\forall a, b \in \mathbb{R}, \forall \lambda > 0$, there is $(a+b)^2 \leq (1+1/\lambda)a^2 + (1+\lambda)b^2$.

Lemma 30 is proved in Appendix B.3.1. Lemma 31 is easy to see since $2ab = 2(a/\sqrt{\lambda})(b\sqrt{\lambda}) \leq a^2/\lambda + \lambda b^2$. Compared to the more naive bound $(a+b)^2 \leq 2(a^2 + b^2)$, this lemma allows us to obtain a sharper bound when a and b are imbalanced in magnitude by adjusting λ .

$\forall \mathbf{x} \in \mathcal{X}$, we can write

$$\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x}) = \left(\mathbf{F}_m(\mathbf{x}) - \frac{1}{m} \sum_{i=1}^m \bar{\mathbf{A}}_i \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})) \right) + \left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{A}}_j \sigma(\mathbf{H}_j^{(L-1)}(\mathbf{x})) - f(\mathbf{x}) \right)$$

Thus, using Lemma 31, we see that, $\forall \mathbf{x} \in \mathcal{X}$, $\forall \lambda > 0$,

$$\mathcal{E}_{\mathbf{x} \sim \nu} \{ \mathbb{E} [|\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x})|^2] \} \leq (1 + 1/\lambda) \cdot \mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(I)} \} + (1 + \lambda) \cdot \mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(II)} \} ,$$

where we define

$$\begin{aligned} \text{(I)} &:= \mathbb{E} \left[\left(\mathbf{F}_m(\mathbf{x}) - \frac{1}{m} \sum_{i=1}^m \bar{\mathbf{A}}_i \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})) \right)^2 \right] , \\ \text{(II)} &:= \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{A}}_j \sigma(\mathbf{H}_j^{(L-1)}(\mathbf{x})) - f(\mathbf{x}) \right)^2 \right] , \end{aligned}$$

and we will bound these two terms separately. For (I), there is

$$\begin{aligned} \text{(I)} &= \mathbb{E} \left[\left(\frac{1}{m} \sum_{i=1}^m \bar{\mathbf{A}}_i (\sigma(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x})) - \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x}))) \right)^2 \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{m} \sum_{i=1}^m (\bar{\mathbf{A}}_i)^2 \right) \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x})) - \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})))^2 \right) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left(\frac{1}{m} \sum_{i=1}^m (\bar{\mathbf{A}}_i)^2 \right) \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x})) - \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})))^2 \right) \middle| \{ \mathbf{H}_{i'}^{(L-1)} \}_{i' \in [m]} \right] \right] \\ &\leq (J_\sigma)^2 \mathbb{E} \left[\left(\frac{1}{m} \sum_{i=1}^m (\bar{\mathbf{A}}_i)^2 \right) \left(\frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x}) - \mathbf{H}_i^{(L-1)}(\mathbf{x}))^2 \middle| \mathbf{H}_i^{(L-1)} \right] \right) \right] \end{aligned}$$

where on the second line, we use the Cauchy-Schwartz inequality; on the third line, we use the tower property of conditional expectation; and on the fourth line, we use that 1) each $\bar{\mathbf{A}}_i$ is measurable with respect to $\mathbf{H}_i^{(L-1)}$, 2) each $\bar{\mathbf{H}}_i^{(L-1)}$ is independent from $\mathbf{H}_{i'}^{(L-1)}$ when $i \neq i'$, and 3) σ is J_σ -Lipschitz. Then, using Lemma 30, we derive that

$$\begin{aligned} &\mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(I)} \} \\ &\leq (J_\sigma)^2 \mathbb{E} \left[\left(\frac{1}{m} \sum_{i=1}^m (\bar{\mathbf{A}}_i)^2 \right) \left(\frac{1}{m} \sum_{i=1}^m \mathcal{E}_{\mathbf{x} \sim \nu} \left\{ \mathbb{E} \left[(\bar{\mathbf{H}}_i^{(L-1)}(\mathbf{x}) - \mathbf{H}_i^{(L-1)}(\mathbf{x}))^2 \middle| \mathbf{H}_i^{(L-1)} \right] \right\} \right) \right] \\ &\leq \frac{(L-2)^2}{m} (J_\sigma)^{2L-2} M_\nu \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^m (\Xi^{(L-1)}(\mathbf{H}_i^{(L-1)}))^2 \right] \left(\prod_{l=1}^{L-1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, \infty}^2 \right) \\ &\leq \frac{(L-2)^2}{m} (J_\sigma)^{2L-2} M_\nu \|f\|_{\mathcal{H}^{(L)}}^2 \prod_{l=1}^{L-1} \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, \infty}^2 . \end{aligned} \tag{27}$$

For (II), there is

$$\begin{aligned}
 \text{(II)} &:= \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{A}}_j \sigma(\mathbf{H}_j^{(L-1)}(\mathbf{x})) - f(\mathbf{x}) \right)^2 \right] \\
 &= \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\left(\Xi^{(L-1)}(\mathbf{H}_i^{(L-1)}) \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})) - \mathbb{E} \left[\Xi^{(L-1)}(\mathbf{H}_i^{(L-1)}) \sigma(\mathbf{H}_i^{(L-1)}(\mathbf{x})) \right] \right)^2 \right] \\
 &\leq \frac{1}{m} \mathbb{E} \left[\left(\Xi^{(L-1)}(\mathbf{H}^{(L-1)}) \right)^2 \left(\sigma(\mathbf{H}^{(L-1)}(\mathbf{x})) \right)^2 \right],
 \end{aligned}$$

where we use the independence among $\{\mathbf{H}_i^{(L-1)}\}_{i \in [m]}$ and their equivalence in law to $\mathbf{H}^{(L-1)}$. By Theorem 3(b), $\forall \mathbf{x} \in \mathcal{X}$, there is

$$|\mathbf{H}^{(L-1)}(\mathbf{x})| \leq (J_\sigma)^{L-2} \|\mathbf{x}\| \mathcal{C}_2^{(L-1)}(\mathbf{H}_i^{(L)}) \leq (J_\sigma)^{L-2} \|\mathbf{x}\| \|\mathbf{H}^{(L-1)}\|_{\mathcal{H}^{(L-1)}} \prod_{l'=1}^{L-2} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, 2},$$

and hence, by the assumptions on σ , it holds that

$$\left(\sigma(\mathbf{H}^{(L-1)}(\mathbf{x})) \right)^2 \leq (J_\sigma)^2 (\mathbf{H}^{(L-1)}(\mathbf{x}))^2 \leq (J_\sigma)^{2L-2} \|\mathbf{x}\|^2 \|\mathbf{H}^{(L-1)}\|_{\mathcal{H}^{(L-1)}}^2 \prod_{l'=1}^{L-2} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, 2}^2,$$

which is almost surely bounded by $(J_\sigma)^{2L-2} \|\mathbf{x}\|^2 \prod_{l'=1}^{L-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2$. Thus,

$$\begin{aligned}
 \mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(II)} \} &\leq \frac{1}{m} (J_\sigma)^{2L-2} M_\nu \left(\prod_{l'=1}^{L-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \right) \mathbb{E} \left[\left(\Xi^{(L-1)}(\mathbf{H}^{(L-1)}) \right)^2 \right] \\
 &\leq \frac{1}{m} (J_\sigma)^{2L-2} M_\nu \|f\|_{\mathcal{H}^{(L)}}^2 \prod_{l'=1}^{L-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2.
 \end{aligned} \tag{28}$$

Together, combining (27) and (28) with Lemma 31 and choosing $\lambda = L - 2$, we derive that,

$$\begin{aligned}
 &\mathcal{E}_{\mathbf{x} \sim \nu} \{ \mathbb{E} [|\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x})|^2] \} \\
 &\leq ((1 + 1/\lambda)(L - 2)^2 + (1 + \lambda)) (J_\sigma)^{2L-2} M_\nu \|f\|_{\mathcal{H}^{(L)}}^2 \prod_{l'=1}^{L-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \\
 &= \frac{(L - 1)^2}{m} (J_\sigma)^{2L-2} M_\nu \|f\|_{\mathcal{H}^{(L)}}^2 \prod_{l'=1}^{L-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \\
 &= \frac{(L - 1)^2}{m} (J_\sigma)^{2L-2} M_\nu \mathcal{C}_\infty^{(L)}(f)^2.
 \end{aligned}$$

Interchanging the expectations, we get

$$\begin{aligned}
 \mathbb{E} [\mathcal{E}_{\mathbf{x} \sim \nu} \{ |\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x})|^2 \}] &= \mathcal{E}_{\mathbf{x} \sim \nu} \{ \mathbb{E} [|\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x})|^2] \} \\
 &\leq \frac{(L - 1)^2}{m} (J_\sigma)^{2L-2} M_\nu \mathcal{C}_\infty^{(L)}(f)^2.
 \end{aligned}$$

Thus, as a consequence of Markov's inequality, there exists a realization of $(\mathbf{H}_i^{(l)})_{l \in [L-1], i \in [m]}$ under which

$$\mathcal{E}_{\mathbf{x} \sim \nu} \{ |\mathbf{F}_m(\mathbf{x}) - f(\mathbf{x})|^2 \} \leq \frac{(L-1)^2}{m} (J_\sigma)^{2L-2} M_\nu \mathcal{C}_\infty^{(L)}(f)^2.$$

B.3.1 PROOF OF LEMMA 30

For $l = 1$, there is $\bar{\mathbf{H}}_i^{(1)} = \mathbf{H}_i^{(1)}$, and hence $\bar{\mathbf{H}}_i^{(1)}(\mathbf{x}) - \mathbf{H}_i^{(1)}(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{X}$, thus validating (26) when $l = 1$. Next, we suppose that the statement holds for some $l \in [L-2]$ and prove that it holds for $l+1$ as well.

For level $l+1$, $\forall i \in [m]$ and $\mathbf{x} \in \mathcal{X}$, we can write

$$\begin{aligned} \bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) &= \left(\bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) - \frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) \right) \\ &\quad + \left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) \right) \end{aligned}$$

Hence, by Lemma 31, it holds for any $\lambda > 0$ that

$$\begin{aligned} &\mathcal{E}_{\mathbf{x} \sim \nu} \left\{ \mathbb{E} \left[\left(\bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right] \right\} \\ &\leq (1 + 1/\lambda) \cdot \mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(I)} \} + (1 + \lambda) \cdot \mathcal{E}_{\mathbf{x} \sim \nu} \{ \text{(II)} \}, \end{aligned}$$

where we define

$$\begin{aligned} \text{(I)} &:= \mathbb{E} \left[\left(\bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) - \frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right], \\ \text{(II)} &:= \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right], \end{aligned}$$

and we will bound these two terms separately. For (I), there is

$$\begin{aligned} \text{(I)} &= \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} (\sigma(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x})) - \sigma(\mathbf{H}_j^{(l)}(\mathbf{x}))) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m (\bar{\mathbf{W}}_{i,j}^{(l)})^2 \right) \left(\frac{1}{m} \sum_{j=1}^m (\sigma(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x})) - \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})))^2 \right) \middle| \mathbf{H}_i^{(l+1)} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m (\bar{\mathbf{W}}_{i,j}^{(l)})^2 \right) \left(\frac{1}{m} \sum_{j=1}^m (\sigma(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x})) - \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})))^2 \right) \right. \right. \\ &\quad \left. \left. \middle| \mathbf{H}_i^{(l+1)}, \{\mathbf{H}_j^{(l)}\}_{j \in [m]} \right) \middle| \mathbf{H}_i^{(l+1)} \right] \\ &\leq (J_\sigma)^2 \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m (\bar{\mathbf{W}}_{i,j}^{(l)})^2 \right) \left(\frac{1}{m} \sum_{j=1}^m \mathbb{E} \left[(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x}) - \mathbf{H}_j^{(l)}(\mathbf{x}))^2 \middle| \mathbf{H}_j^{(l)} \right] \right) \middle| \mathbf{H}_i^{(l+1)} \right] \end{aligned}$$

where on the second line, we use the Cauchy-Schwartz inequality; on the third line, we use the tower property of conditional expectations; and on the last line, we use that 1) $\bar{\mathbf{W}}_{i,j}^{(l)}$ is measurable with respect to (the sigma-algebra generated by) $\mathbf{H}_i^{(l+1)}$ and $\mathbf{H}_j^{(l)}$, 2) $\bar{\mathbf{H}}_j^{(l)}$ and $\mathbf{H}_j^{(l)}$ are independent from $\mathbf{H}_i^{(l+1)}$, 3) each $\bar{\mathbf{H}}_j^{(l)}$ is independent from $\mathbf{H}_{j'}^{(l)}$ when $j \neq j'$, and 4) σ is J_σ -Lipschitz. Then, using the inductive hypothesis, we derive that

$$\begin{aligned}
 & \mathcal{E}_{\mathbf{x} \sim \nu} \{(\text{I})\} \\
 & \leq (J_\sigma)^2 \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m (\bar{\mathbf{W}}_{i,j}^{(l)})^2 \right) \left(\frac{1}{m} \sum_{j=1}^m \mathcal{E}_{\mathbf{x} \sim \nu} \left\{ \mathbb{E} \left[(\bar{\mathbf{H}}_j^{(l)}(\mathbf{x}) - \mathbf{H}_j^{(l)}(\mathbf{x}))^2 \mid \mathbf{H}_j^{(l)} \right] \right\} \right) \mid \mathbf{H}_i^{(l+1)} \right] \\
 & \leq \frac{(l-1)^2}{m} (J_\sigma)^{2l} M_\nu \left(\prod_{l'=1}^l \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \right) \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m (\bar{\mathbf{W}}_{i,j}^{(l)})^2 \mid \mathbf{H}_i^{(l+1)} \right] \\
 & \leq \frac{(l-1)^2}{m} (J_\sigma)^{2l} M_\nu \left(\prod_{l'=1}^l \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \right) \|\mathbf{H}_i^{(l+1)}\|_{\mathcal{H}^{(l+1)}}^2 \\
 & \leq \frac{(l-1)^2}{m} (J_\sigma)^{2l} M_\nu \left(\prod_{l'=1}^{l+1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \right),
 \end{aligned}$$

For (II), there is

$$\begin{aligned}
 (\text{II}) & := \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \bar{\mathbf{W}}_{i,j}^{(l)} \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) \right)^2 \mid \mathbf{H}_i^{(l+1)} \right] \\
 & = \mathbb{E} \left[\left(\frac{1}{m} \sum_{j=1}^m \Xi^{(l)}(\mathbf{H}_i^{(l+1)}, \mathbf{H}_j^{(l)}) \sigma(\mathbf{H}_j^{(l)}(\mathbf{x})) \right. \right. \\
 & \quad \left. \left. - \mathbb{E} \left[\Xi^{(l)}(\mathbf{H}_i^{(l+1)}, \mathbf{H}^{(l)}) \sigma(\mathbf{H}^{(l)}(\mathbf{x})) \mid \mathbf{H}_i^{(l+1)} \right] \right)^2 \mid \mathbf{H}_i^{(l+1)} \right] \\
 & \leq \frac{1}{m} \mathbb{E} \left[\left(\Xi^{(l)}(\mathbf{H}_i^{(l+1)}, \mathbf{H}^{(l)}) \right)^2 \left(\sigma(\mathbf{H}^{(l)}(\mathbf{x})) \right)^2 \mid \mathbf{H}_i^{(l+1)} \right],
 \end{aligned}$$

where we use the independence among $\{\mathbf{H}_j^{(l)}\}_{j \in [m]}$ and their equivalence in law to $\mathbf{H}^{(l)}$. By Theorem 3(b), $\forall \mathbf{x} \in \mathcal{X}$, there is

$$|\mathbf{H}_j^{(l)}(\mathbf{x})| \leq (J_\sigma)^{l-1} \|\mathbf{x}\| \mathcal{C}_2^{(l)}(\mathbf{H}_j^{(l)}) = \|\mathbf{x}\| \|\mathbf{H}_j^{(l)}\|_{\mathcal{H}^{(l)}} \prod_{l'=1}^{l-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, 2},$$

and hence, as σ is J_σ -Lipschitz with $\sigma(0) = 0$, it holds that

$$(\sigma(\mathbf{H}_j^{(l)}(\mathbf{x})))^2 \leq (J_\sigma)^2 (\mathbf{H}_j^{(l)}(\mathbf{x}))^2 \leq (J_\sigma)^{2l} \|\mathbf{x}\|^2 \|\mathbf{H}_j^{(l)}\|_{\mathcal{H}^{(l)}}^2 \prod_{l'=1}^{l-1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, 2}^2,$$

which is bounded almost surely by $(J_\sigma)^{2l} \|\mathbf{x}\|^2 \prod_{l'=1}^l \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2$. Thus,

$$\begin{aligned} \mathcal{E}_{\mathbf{x} \sim \nu} \{(\text{II})\} &\leq \frac{1}{m} (J_\sigma)^{2l} M_\nu \mathbb{E} \left[\left(\Xi^{(l)}(\mathbf{H}_i^{(l+1)}, \mathbf{H}^{(l)}) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right] \prod_{l'=1}^l \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \\ &\leq \frac{1}{m} (J_\sigma)^{2l} M_\nu \|\mathbf{H}_i^{(l+1)}\|_{\mathcal{H}^{(l+1)}}^2 \prod_{l'=1}^l \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2, \end{aligned}$$

and the right-hand side is bounded almost surely by $\frac{1}{m} (J_\sigma)^{2l} M_\nu \prod_{l'=1}^{l+1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2$.

Therefore, combining the bounds for (I) and (II) with Lemma 31, and further choosing $\lambda = l - 1$, we get

$$\begin{aligned} &\mathcal{E}_{\mathbf{x} \sim \nu} \left\{ \mathbb{E} \left[\left(\bar{\mathbf{H}}_i^{(l+1)}(\mathbf{x}) - \mathbf{H}_i^{(l+1)}(\mathbf{x}) \right)^2 \middle| \mathbf{H}_i^{(l+1)} \right] \right\} \\ &\leq ((1 + 1/\lambda)(l - 1)^2 + (1 + \lambda))(J_\sigma)^{2l} M_\nu \prod_{l'=1}^{l+1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \\ &\leq \frac{l^2}{m} (J_\sigma)^{2l} M_\nu \prod_{l'=1}^{l+1} \|\mu^{(l')}\|_{\mathcal{H}^{(l')}, \infty}^2 \end{aligned}$$

which proves the inductive hypothesis at level $l + 1$.

Remark 32 Note that using Lemma 31 and optimizing for λ as a function of l —compared to just using the basic inequality $(a + b)^2 \leq 2(a^2 + b^2)$ —reduces the explicit dependency of the right-hand side of (26) on l from exponential to quadratic.

Appendix C. Supplementary Materials for Section 5

C.1 Proof of Theorem 10

When σ is 1-homogeneous, we see that $\mathcal{C}^{(L)}$ can be alternatively expressed as

$$\begin{aligned} \mathcal{C}^{(L)}(f) &= \inf_{\mu^{(1)}, \dots, \mu^{(L-1)}} \|f\|_{\mathcal{H}^{(L)}} \\ \text{s.t. } &\|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2} = 1, \quad \forall l \in [L - 1]. \end{aligned}$$

In the following, for brevity, we will write $\sup_{\mu^{(l)}}$ and \sup_ξ for

$$\sup_{\substack{\mu^{(l)} \in \mathcal{P}(\mathcal{H}^{(l)}) \\ \|\mu^{(l)}\|_{\mathcal{H}^{(l)}, 2} \leq 1}} \quad \text{and} \quad \sup_{\substack{\xi \in L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)}) \\ \|\xi\|_{L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)})} \leq 1}},$$

respectively. By the definition of the empirical Rademacher complexity,

$$\widehat{\text{Rad}}_S(\mathbb{B}(\mathcal{F}^{(L)})) := \mathbb{E}_\tau \left[\frac{1}{n} \sup_{\mathcal{C}^{(L)}(f) \leq 1} \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right].$$

For any $\lambda > 0$, we consider the function $g_\lambda : \mathbb{R} \rightarrow \mathbb{R}$ defined by $g_\lambda(u) = \exp(\lambda u)$, which is positive, monotonically increasing and convex. Thus, using Jensen's inequality, we can write

$$\begin{aligned} n \widehat{\text{Rad}}_S(\mathbb{B}(\mathcal{F}^{(L)})) &= \frac{1}{\lambda} \log \left(g_\lambda \left(\mathbb{E}_\tau \left[\sup_{\mathcal{C}^{(L)}(f) \leq 1} \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right] \right) \right) \\ &\leq \frac{1}{\lambda} \log \left(\mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L)}(f) \leq 1} \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right) \right] \right) \\ &\leq \frac{1}{\lambda} \log \mathcal{M}_\lambda^{(L)}, \end{aligned}$$

where we define, $\forall l \in [L]$, $\mathcal{M}_\lambda^{(l)} := \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(l)}(f) \leq 1} \left| \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right| \right) \right]$.

Lemma 33 *Under the assumptions on σ in Theorem 10,*

$$\mathcal{M}_\lambda^{(L)} \leq (2J_\sigma)^{L-1} \mathbb{E}_\tau \left[g_\lambda \left(\left\| \sum_{k=1}^n \tau_k \mathbf{x}_k \right\| \right) \right].$$

This lemma is proved in Appendix C.1.1. Then, if we choose $\lambda = \frac{\sqrt{2(L-1)\log(2J_\sigma)}}{\sqrt{\sum_{k=1}^n \|\mathbf{x}_k\|_2^2}}$, we can generalize the argument used by Golowich et al. (2018) slightly to show that

$$\frac{1}{\lambda} \log \left(2^{L-1} \mathbb{E}_\tau \left[g_\lambda \left(\left\| \sum_{k=1}^n \tau_k \mathbf{x}_k \right\| \right) \right] \right) \leq (\sqrt{2L \log(2J_\sigma)} + 1) \sqrt{\sum_{k=1}^n \|\mathbf{x}_k\|_2^2},$$

which yields the desired result for the case $M = 1$. The extension to general $M > 0$ is easy to see because of the linearity of Hilbert space norms.

C.1.1 PROOF OF LEMMA 33

We see that

$$\begin{aligned} \sup_{\mathcal{C}^{(L)}(f) \leq 1} \left| \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right| &\leq \sup_{\mu^{(1)}, \dots, \mu^{(L-1)}, \xi} \left| \sum_{k=1}^n \tau_k \int \xi(h) \sigma(h(\mathbf{x}_k)) \mu^{(L-1)}(dh) \right| \\ &\leq \sup_{\mu^{(1)}, \dots, \mu^{(L-1)}, \xi} \left| \int \sum_{k=1}^n \tau_k \frac{\sigma(h(\mathbf{x}_k))}{\|h\|_{\mathcal{H}^{(L-1)}}} \xi(h) \|h\|_{\mathcal{H}^{(L-1)}} \mu^{(L-1)}(dh) \right|. \end{aligned}$$

By the Cauchy-Schwartz inequality and the homogeneity of σ , there is

$$\begin{aligned} &\left| \int \sum_{k=1}^n \tau_k \frac{\sigma(h(\mathbf{x}_k))}{\|h\|_{\mathcal{H}^{(L-1)}}} \xi(h) \|h\|_{\mathcal{H}^{(L-1)}} \mu^{(L-1)}(dh) \right| \\ &\leq \left(\sup_{h \in \mathcal{H}^{(L-1)}} \left| \sum_{k=1}^n \tau_k \frac{\sigma(h(\mathbf{x}_k))}{\|h\|_{\mathcal{H}^{(L-1)}}} \right| \right) \int |\xi(h)| \|h\|_{\mathcal{H}^{(L-1)}} \mu^{(L-1)}(dh) \\ &\leq \left(\sup_{\|\hat{h}\|_{\mathcal{H}^{(L-1)}} \leq 1} \left| \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right| \right) \left(\int |\xi(h)|^2 \mu^{(L-1)}(dh) \right)^{1/2} \left(\int \|h\|_{\mathcal{H}^{(L-1)}}^2 \mu^{(L-1)}(dh) \right)^{1/2}, \end{aligned}$$

and hence

$$\begin{aligned} \sup_{\mathcal{C}^{(L)}(f) \leq 1} \left| \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right| &\leq \sup_{\substack{\|\hat{h}\|_{\mathcal{H}^{(L-1)}} \leq 1 \\ \mu^{(1)}, \dots, \mu^{(L-2)}}} \left| \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right| \\ &= \begin{cases} \sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} \left| \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right|, & \text{if } L \geq 3, \\ \sup_{\|\hat{h}\|_{\mathcal{H}^{(1)}} \leq 1} \left| \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right|, & \text{if } L = 2. \end{cases} \end{aligned}$$

Notice that, since g_λ is positive, there is $g_\lambda(|u|) \leq g_\lambda(u) + g_\lambda(-u)$. In addition, the assumptions on σ imply that it is J_σ -Lipschitz. Therefore, when $L \geq 3$,

$$\begin{aligned} \mathcal{M}_\lambda^{(L)} &\leq \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L)}(f) \leq 1} \left| \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right| \right) \right] \\ &= \mathbb{E}_\tau \left[\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} g_\lambda \left(\left| \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right| \right) \right] \\ &\leq \mathbb{E}_\tau \left[\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} g_\lambda \left(\sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right) \right] + \mathbb{E}_\tau \left[\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} g_\lambda \left(- \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right) \right] \\ &= \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right) \right] + \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} \sum_{k=1}^n (-\tau_k) \sigma(\hat{h}(\mathbf{x}_k)) \right) \right] \\ &= 2\mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} \sum_{k=1}^n \tau_k \sigma(\hat{h}(\mathbf{x}_k)) \right) \right] \\ &\leq 2J_\sigma \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\mathcal{C}^{(L-1)}(\hat{h}) \leq 1} \sum_{k=1}^n \tau_k \hat{h}(\mathbf{x}_k) \right) \right], \end{aligned}$$

where for the second and the fourth lines we use the monotonicity of g_λ ; for the fifth line we use the symmetry of the Rademacher distribution; and for the sixth line we apply the *contraction lemma* (Equation 4.20 in Ledoux and Talagrand, 1991) by using the Lipschitzness of σ and the monotonicity and convexity of g_λ . Hence, we derive that

$$\mathcal{M}_\lambda^{(L)} \leq 2J_\sigma \mathcal{M}_\lambda^{(L-1)}.$$

Thus, by induction, it holds that

$$\begin{aligned} \mathcal{M}_\lambda^{(L)} &\leq (2J_\sigma)^{L-1} \mathcal{M}_\lambda^{(1)} \\ &= (2J_\sigma)^{L-1} \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\|f\|_{\mathcal{H}^{(1)}} \leq 1} \left| \sum_{k=1}^n \tau_k f(\mathbf{x}_k) \right| \right) \right] \\ &= (2J_\sigma)^{L-1} \mathbb{E}_\tau \left[g_\lambda \left(\sup_{\|z\|_2 \leq 1} \left| \sum_{k=1}^n \tau_k z^\top \cdot \mathbf{x}_k \right| \right) \right] \\ &\leq (2J_\sigma)^{L-1} \mathbb{E}_\tau \left[g_\lambda \left(\left\| \sum_{k=1}^n \tau_k \mathbf{x}_k \right\| \right) \right], \end{aligned}$$

which proves the lemma.

Appendix D. Supplementary Materials for Section 7

D.1 Proof of Theorem 18

Notations If $(u_m)_{m \in \mathbb{N}_+}$ and $(u'_m)_{m \in \mathbb{N}_+}$ are two sequences of non-negative random variables, we write $u_m = o_{\mathbb{P}}(u'_m)$ if it holds almost surely that $\forall \epsilon > 0, \exists M > 0$ such that $\forall m > M, u_m \leq \epsilon u'_m$.

Preliminaries We start by defining the limiting dynamics in an alternative way in terms of flow maps and will show its equivalence to the definition given in Section 7.1. We will also generalize the analysis by taking into account the bias terms as introduced in Appendix B.1. For $t \geq 0$, let $\mathbf{A}_t, \Xi_t^{(1)}, \dots, \Xi_t^{(L-2)}, \mathbf{B}_t^{(1)}, \dots,$ and $\mathbf{B}_t^{(L-1)}$ be random variables, let \mathbf{Z}_t be a d -dimensional random vector, and let $\mathbf{H}_t^{(1)}, \dots, \mathbf{H}_t^{(L-1)}$ and $\mathbf{Q}_t^{(1)}, \dots, \mathbf{Q}_t^{(L-1)}$ be random fields on \mathcal{X} , which all depend on time and are defined as follows. At initial time, we let $\mathbf{A}_0, \mathbf{Z}_0, \mathbf{B}_0^{(1)}, \dots, \mathbf{B}_0^{(L-1)}$ be distributed independently with laws $\rho_a, \rho_z, \rho_{b(1)}, \dots, \rho_{b(L-1)}$, respectively. Then, we define $\mathbf{H}_t^{(l)}, \mathbf{Q}_t^{(l)}, \Xi_t^{(l)}, \mathbf{B}_t^{(l)}$ and \mathbf{Z}_t alternatively as:

$$\begin{aligned}
 \mathbf{H}_t^{(l)}(\mathbf{x}) &= \begin{cases} H_t^{(1)}(\mathbf{x}, \mathbf{Z}_0, \mathbf{B}_0^{(1)}), & l = 1 \\ H_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}), & l \in \{2, \dots, L-2\} \\ H_t^{(L-1)}(\mathbf{x}, \mathbf{A}_0, \mathbf{B}_0^{(L-1)}), & l = L-1 \end{cases} \\
 \mathbf{Q}_t^{(l)}(\mathbf{x}) &= \begin{cases} Q_t^{(1)}(\mathbf{x}, \mathbf{Z}_0, \mathbf{B}_0^{(1)}), & l = 1 \\ Q_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}), & l \in \{2, \dots, L-2\} \\ Q_t^{(L-1)}(\mathbf{x}, \mathbf{A}_0, \mathbf{B}_0^{(L-1)}), & l = L-1 \end{cases} \\
 \Xi_t^{(l)} &= \begin{cases} \Xi_t^{(1)}(\mathbf{Z}_0, \mathbf{B}_0^{(1)}, \mathbf{B}_0^{(2)}), & l = 1 \\ \Xi_t^{(l)}(\mathbf{B}_0^{(l)}, \mathbf{B}_0^{(l+1)}), & l \in \{2, \dots, L-2\} \\ \Xi_t^{(L-2)}(\mathbf{A}_0, \mathbf{B}_0^{(L-2)}, \mathbf{B}_0^{(L-1)}), & l = L-2, \end{cases} \quad (29) \\
 \mathbf{B}_t^{(l)} &= \begin{cases} B_t^{(1)}(\mathbf{Z}_0, \mathbf{B}_0^{(1)}), & l = 1 \\ B_t^{(l)}(\mathbf{B}_0^{(l)}), & l \in \{2, \dots, L-2\} \\ B_t^{(L-1)}(\mathbf{A}_0, \mathbf{B}_0^{(L-1)}), & l = L-1 \end{cases} \\
 \mathbf{Z}_t &= Z_t(\mathbf{Z}_0, \mathbf{B}_0^{(1)}), \\
 \mathbf{A}_t &= A_t(\mathbf{A}_0, \mathbf{B}_0^{(L-1)}).
 \end{aligned}$$

by introducing the following (deterministic) functions:

$$\begin{aligned}
 \mathbf{H}_t^{(l)}, \mathbf{Q}_t^{(l)} : & \begin{cases} \mathcal{X} \times \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}, & l = 1, \\ \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}, & l \in \{2, \dots, L-2\}, \\ \mathcal{X} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}, & l = L-1, \end{cases} \\
 \Xi_t^{(l)} : & \begin{cases} \mathbb{R}^d \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}, & l = 1, \\ \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}, & l \in \{2, \dots, L-3\}, \\ \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}, & l = L-2, \end{cases}
 \end{aligned}$$

$$\begin{aligned}
B_t^{(l)} &: \begin{cases} \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R} , & l = 1 , \\ \mathbb{R} \rightarrow \mathbb{R} , & l \in \{2, \dots, L-2\} , \\ \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} , & l = L-1 , \end{cases} \\
Z_t &: \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d , \\
A_t &: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} ,
\end{aligned}$$

They can be interpreted as flow maps and are defined as follows: $\forall t \geq 0$,

- For $l \in [L-1]$, $H_t^{(l)}$ is defined by, $\forall \mathbf{x} \in \mathcal{X}$, $\forall \mathbf{z} \in \mathbb{R}^d$, $\forall a, b \in \mathbb{R}$,

$$\begin{aligned}
H_t^{(1)}(\mathbf{x}, \mathbf{z}, b) &= Z_t(\mathbf{z}, b)^\top \cdot \mathbf{x} , \\
H_t^{(2)}(\mathbf{x}, b) &= \mathbb{E} \left[\Xi_t^{(1)}(\mathbf{Z}_0, \mathbf{B}_0^{(1)}, b) \sigma(H_t^{(1)}(\mathbf{x}, \mathbf{Z}_0, \mathbf{B}_0^{(1)}) + B_t^{(1)}(\mathbf{Z}_0, \mathbf{B}_0^{(1)})) \right] , \\
H_t^{(l+1)}(\mathbf{x}, b) &= \mathbb{E} \left[\Xi_t^{(l)}(\mathbf{B}_0^{(l)}, b) \sigma(H_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) + B_t^{(l)}(\mathbf{B}_0^{(l)})) \right] , \quad \forall l \in \{2, \dots, L-3\} , \\
H_t^{(L-1)}(\mathbf{x}, a, b) &= \mathbb{E} \left[\Xi_t^{(L-2)}(a, \mathbf{B}_0^{(L-2)}, b) \sigma(H_t^{(L-2)}(\mathbf{x}, \mathbf{B}_0^{(L-2)}) + B_t^{(L-2)}(\mathbf{B}_0^{(L-2)})) \right] .
\end{aligned}$$

- For $l \in [L-1]$, $Q_t^{(l)}$ is defined by, $\forall \mathbf{x} \in \mathcal{X}$, $\forall \mathbf{z} \in \mathbb{R}^d$, $\forall a, b \in \mathbb{R}$,

$$\begin{aligned}
Q_t^{(L-1)}(\mathbf{x}, a, b) &= A_t(a, b) , \\
Q_t^{(L-2)}(\mathbf{x}, b) &= \mathbb{E} \left[\Xi_t^{(L-2)}(\mathbf{A}_0, b, \mathbf{B}_0^{(L-1)}) Q_t^{(L-1)}(\mathbf{x}, \mathbf{B}_0^{(L-1)}) \sigma'(H_t^{(L-1)}(\mathbf{x}, \mathbf{A}_0, \mathbf{B}_0^{(L-1)}) \right. \\
&\quad \left. + B_t^{(L-1)}(\mathbf{A}_0, \mathbf{B}_0^{(L-1)})) \right] , \\
Q_t^{(l-1)}(\mathbf{x}, b) &= \mathbb{E} \left[\Xi_t^{(l-1)}(b, \mathbf{B}_0^{(l)}) Q_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) \sigma'(H_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) + B_t^{(l)}(\mathbf{B}_0^{(l)})) \right] , \\
&\quad \forall l \in \{1, \dots, L-2\} , \\
Q_t^{(1)}(\mathbf{x}, \mathbf{z}, b) &= \mathbb{E} \left[\Xi_t^{(1)}(\mathbf{z}, b, \mathbf{B}_0^{(2)}) Q_t^{(2)}(\mathbf{x}, \mathbf{B}_0^{(2)}) \sigma'(H_t^{(2)}(\mathbf{x}, \mathbf{B}_0^{(2)}) + B_t^{(2)}(\mathbf{B}_0^{(2)})) \right] .
\end{aligned}$$

- For $l \in [L-2]$, $\Xi_t^{(l)}$ is defined by, $\forall \mathbf{z} \in \mathbb{R}^d$, $\forall a, b \in \mathbb{R}$,

$$\begin{aligned}
\frac{d}{dt} \Xi_t^{(1)}(\mathbf{z}, b, b') &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(2)}(\mathbf{x}, b') \sigma'(H_t^{(2)}(\mathbf{x}, b') + B_t^{(2)}(b')) \right. \\
&\quad \left. \sigma(H_t^{(1)}(\mathbf{x}, \mathbf{z}, b) + B_t^{(1)}(\mathbf{z}, b)) \right\} , \\
\frac{d}{dt} \Xi_t^{(l)}(b, b') &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(l+1)}(\mathbf{x}, b') \sigma'(H_t^{(l+1)}(\mathbf{x}, b') + B_t^{(l+1)}(b')) \right. \\
&\quad \left. \sigma(H_t^{(l)}(\mathbf{x}, b) + B_t^{(l)}(b)) \right\} , \quad \forall l \in \{2, \dots, L-2\} , \\
\frac{d}{dt} \Xi_t^{(L-2)}(a', b, b') &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(L-1)}(\mathbf{x}, a', b') \sigma'(H_t^{(L-1)}(\mathbf{x}, a', b') + B_t^{(L-1)}(a', b')) \right. \\
&\quad \left. \sigma(H_t^{(l)}(\mathbf{x}, b) + B_t^{(l)}(b)) \right\} ,
\end{aligned}$$

together with the initial conditions

$$\begin{aligned}\Xi_0^{(1)}(\mathbf{z}, b, b') &= 0, \\ \Xi_0^{(l)}(b, b') &= 0, \quad \forall l \in \{2, \dots, L-3\}, \\ \Xi_0^{(L-2)}(a', b, b') &= 0.\end{aligned}$$

- For $l \in [L-1]$, $B_t^{(l)}$ is defined by, $\forall \mathbf{z} \in \mathbb{R}^d, \forall a, b \in \mathbb{R}$,

$$\begin{aligned}\frac{d}{dt}B_t^{(1)}(\mathbf{z}, b) &= -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(1)}(\mathbf{x}, \mathbf{z}, b) \sigma'(H_t^{(1)}(\mathbf{x}, \mathbf{z}, b) + B_t^{(1)}(\mathbf{z}, b)) \right\}, \\ \frac{d}{dt}B_t^{(l)}(b) &= -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(l)}(\mathbf{x}, b) \sigma'(H_t^{(l)}(\mathbf{x}, b) + B_t^{(l)}(b)) \right\}, \\ &\quad \forall l \in \{2, \dots, L-2\}, \\ \frac{d}{dt}B_t^{(L-1)}(a, b) &= -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(L-1)}(\mathbf{x}, a, b) \sigma'(H_t^{(L-1)}(\mathbf{x}, a, b) + B_t^{(L-1)}(a, b)) \right\},\end{aligned}$$

together with the initial conditions

$$\begin{aligned}B_0^{(1)}(\mathbf{z}, b) &= b, \\ B_0^{(l)}(b) &= b, \quad \forall l \in \{2, \dots, L-2\}, \\ B_0^{(L-1)}(a, b) &= b.\end{aligned}$$

- Z_t is defined by, $\forall \mathbf{z} \in \mathbb{R}^d, \forall b \in \mathbb{R}$,

$$\frac{d}{dt}Z_t(\mathbf{z}, b) = -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) Q_t^{(1)}(\mathbf{x}, \mathbf{z}, b) \sigma'(H_t^{(1)}(\mathbf{x}, \mathbf{z}, b) + B_t^{(1)}(\mathbf{z}, b)) \mathbf{x} \right\},$$

together with the initial condition

$$Z_0(\mathbf{z}, b) = \mathbf{z}.$$

- A_t is defined by, $\forall a, b \in \mathbb{R}$,

$$\frac{d}{dt}A_t(a, b) = -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \sigma(H_t^{(l)}(\mathbf{x}, a, b) + B_t^{(l)}(a, b)) \right\},$$

together with the initial condition

$$A_0(a, b) = a.$$

Lastly, f_t and ζ_t are defined in the same way as in Section 7.1.

Defined in this way, one can verify that $\Xi_0^{(l)} := 0, \forall l \in [L-2]$, and furthermore, for $t \geq 0$, the dynamics of $\mathbf{A}_t, \mathbf{Z}_t, \Xi_t^{(1)}, \dots, \Xi_t^{(L-2)}, \mathbf{B}_t^{(1)}, \dots, \mathbf{B}_t^{(L-1)}$ satisfy the following equations:

$$\begin{aligned}\frac{d}{dt}\mathbf{A}_t &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \sigma(\mathbf{H}_t^{(L-1)}(\mathbf{x}) + \mathbf{B}_t^{(L-1)}) \right\}, \\ \frac{d}{dt}\mathbf{Z}_t &= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \mathbf{Q}_t^{(1)}(\mathbf{x}) \sigma'(\mathbf{H}_t^{(1)}(\mathbf{x}) + \mathbf{B}_t^{(1)}) \mathbf{x} \right\}, \\ \frac{d}{dt}\mathbf{B}_t^{(l)} &= -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \mathbf{Q}_t^{(l)}(\mathbf{x}) \sigma'(\mathbf{H}_t^{(l)}(\mathbf{x}) + \mathbf{B}_t^{(l)}) \right\}, \quad \forall l \in [L-1],\end{aligned}$$

and $\forall l \in [L - 2]$,

$$\frac{d}{dt} \Xi_t^{(l)} = -\mathcal{E}_x \left\{ \zeta_t(\mathbf{x}) \mathbf{Q}_t^{(l+1)}(\mathbf{x}) \sigma'(\mathbf{H}_t^{(l+1)}(\mathbf{x}) + \mathbf{B}_t^{(l+1)}) \sigma(\mathbf{H}_t^{(l)}(\mathbf{x}) + \mathbf{B}_t^{(l)}) \right\}, \quad (30)$$

and moreover, the random fields satisfy $\mathbf{H}_t^{(1)}(\mathbf{x}) = \mathbf{Z}_t^\top \cdot \mathbf{x}$, $\mathbf{Q}_t^{(L-1)}(\mathbf{x}) = \mathbf{A}_t$, and $\forall l \in [L - 2]$,

$$\mathbf{H}_t^{(l+1)}(\mathbf{x}) = \mathbb{E} \left[\Xi_t^{(l)} \sigma(\mathbf{H}_t^{(l)}(\mathbf{x}) + \mathbf{B}_t^{(l)}) \middle| \mathbf{H}_t^{(l+1)} \right]. \quad (31)$$

$$\mathbf{Q}_t^{(l)}(\mathbf{x}) = \mathbb{E} \left[\Xi_t^{(l)} \mathbf{Q}_t^{(l+1)}(\mathbf{x}) \middle| \mathbf{H}_t^{(l)} \right]. \quad (32)$$

From (30), we derive that, $\forall l \in [L - 2]$,

$$\Xi_t^{(l)} = - \int_0^t \mathcal{E}_x \left\{ \zeta_s(\mathbf{x}) \mathbf{Q}_s^{(l+1)}(\mathbf{x}) \sigma'(\mathbf{H}_s^{(l+1)}(\mathbf{x}) + \mathbf{B}_s^{(l+1)}) \sigma(\mathbf{H}_s^{(l)}(\mathbf{x}) + \mathbf{B}_s^{(l)}) \right\} ds. \quad (33)$$

When $\rho_{b^{(1)}} = \dots = \rho_{b^{(L-1)}} = \delta_0$ and $\beta = 0$, there is $\mathbf{B}_t^{(l)} = 0$, $\forall t \geq 0$, $\forall l \in [L - 1]$. Then, substituting (33) into (31) and (32), we obtain (12) and (13). Hence, the definitions given by (29) are indeed consistent with the mean-field dynamics described in Section 7.1.

To simplify the propagation-of-chaos argument in the remaining proof, we also define

$$\begin{aligned} \tilde{z}_{i,t} &= Z_t(z_{i,0}, b_{i,0}^{(1)}) \\ \tilde{a}_{i,t} &= A_t(a_{i,0}, b_{i,0}^{(L-1)}) \\ \tilde{b}_{i,t}^{(l)} &= \begin{cases} B_t^{(1)}(z_{i,0}, b_{i,0}^{(1)}), & l = 1 \\ B_t^{(l)}(b_{i,0}^{(l)}), & l \in \{2, \dots, L-2\} \\ B_t^{(L-1)}(a_i, b_{i,0}^{(L-1)}), & l = L-1 \end{cases} \\ \tilde{h}_{i,t}^{(l)}(\mathbf{x}) &= \begin{cases} H_t^{(1)}(\mathbf{x}, z_{i,0}, b_{i,0}^{(1)}), & l = 1 \\ H_t^{(l)}(\mathbf{x}, b_{i,0}^{(l)}), & l \in \{2, \dots, L-2\} \\ H_t^{(L-1)}(\mathbf{x}, a_i, b_{i,0}^{(L-1)}), & l = L-1 \end{cases} \\ \tilde{q}_{i,t}^{(l)}(\mathbf{x}) &= \begin{cases} Q_t^{(1)}(\mathbf{x}, z_{i,0}, b_{i,0}^{(1)}), & l = 1 \\ Q_t^{(l)}(\mathbf{x}, b_{i,0}^{(l)}), & l \in \{2, \dots, L-2\} \\ Q_t^{(L-1)}(\mathbf{x}, a_i, b_{i,0}^{(L-1)}), & l = L-1 \end{cases} \\ \tilde{W}_{i,j,t}^{(l)} &= \begin{cases} \Xi_t^{(1)}(z_{j,0}, b_{j,0}^{(1)}, b_{i,0}^{(2)}), & l = 1 \\ \Xi_t^{(l)}(b_{j,0}^{(l)}, b_{i,0}^{(l+1)}), & l \in \{2, \dots, L-3\} \\ \Xi_t^{(L-2)}(a_i, b_{j,0}^{(L-2)}, b_{i,0}^{(L-1)}), & l = L-2. \end{cases} \end{aligned}$$

Main proof Given a function $g : \mathbb{R}^N \rightarrow \mathbb{R}$, by the triangle inequality, there is

$$\begin{aligned} & \left| \int g(h(\mathbf{x}'_1), \dots, h(\mathbf{x}'_N)) \mu_{m,t}^{(l)}(dh) - \mathbb{E} \left[g(\mathbf{H}_t^{(l)}(\mathbf{x}'_1), \dots, \mathbf{H}_t^{(l)}(\mathbf{x}'_N)) \right] \right| \\ &= \left| \frac{1}{m} \sum_{i=1}^m g(h_{i,t}^{(l)}(\mathbf{x}'_1), \dots, h_{i,t}^{(l)}(\mathbf{x}'_N)) - \mathbb{E} \left[g(\mathbf{H}_t^{(l)}(\mathbf{x}'_1), \dots, \mathbf{H}_t^{(l)}(\mathbf{x}'_N)) \right] \right| \\ &\leq \text{(I)} + \text{(II)}, \end{aligned}$$

where

$$\begin{aligned}
 \text{(I)} &:= \left| \frac{1}{m} \sum_{i=1}^m g(\tilde{h}_{i,t}^{(l)}(\mathbf{x}'_1), \dots, \tilde{h}_{i,t}^{(l)}(\mathbf{x}'_N)) - \mathbb{E} \left[g(\mathbf{H}_t^{(l)}(\mathbf{x}'_1), \dots, \mathbf{H}_t^{(l)}(\mathbf{x}'_N)) \right] \right| \\
 \text{(II)} &:= \left| \frac{1}{m} \sum_{i=1}^m g(h_{i,t}^{(l)}(\mathbf{x}'_1), \dots, h_{i,t}^{(l)}(\mathbf{x}'_N)) - \frac{1}{m} \sum_{i=1}^m g(\tilde{h}_{i,t}^{(l)}(\mathbf{x}'_1), \dots, \tilde{h}_{i,t}^{(l)}(\mathbf{x}'_N)) \right|
 \end{aligned}$$

For the first term, there is

$$\text{(I)} = \begin{cases} \left| \frac{1}{m} \sum_{i=1}^m g(H_t^{(1)}(\mathbf{x}'_1, \mathbf{z}_{i,0}, b_{i,0}^{(1)}), \dots, H_t^{(1)}(\mathbf{x}'_N, \mathbf{z}_{i,0}, b_{i,0}^{(1)})) \right. \\ \quad \left. - \mathbb{E} \left[g(H_t^{(1)}(\mathbf{x}'_1, \mathbf{Z}_0, \mathbf{B}_0^{(1)}), \dots, H_t^{(1)}(\mathbf{x}'_N, \mathbf{Z}_0, \mathbf{B}_0^{(1)})) \right] \right|, & l = 1 \\ \left| \frac{1}{m} \sum_{i=1}^m g(H_t^{(l)}(\mathbf{x}'_1, b_{i,0}^{(l)}), \dots, H_t^{(l)}(\mathbf{x}'_N, b_{i,0}^{(l)})) \right. \\ \quad \left. - \mathbb{E} \left[g(H_t^{(l)}(\mathbf{x}'_1, \mathbf{B}_0^{(l)}), \dots, H_t^{(l)}(\mathbf{x}'_N, \mathbf{B}_0^{(l)})) \right] \right|, & l \in \{2, \dots, L-2\} \\ \left| \frac{1}{m} \sum_{i=1}^m g(H_t^{(L-1)}(\mathbf{x}'_1, a_i, b_{i,0}^{(L-1)}), \dots, H_t^{(L-1)}(\mathbf{x}'_N, a_i, b_{i,0}^{(L-1)})) \right. \\ \quad \left. - \mathbb{E} \left[g(H_t^{(L-1)}(\mathbf{x}'_1, \mathbf{A}_0, \mathbf{B}_0^{(L-1)}), \dots, H_t^{(L-1)}(\mathbf{x}'_N, \mathbf{A}_0, \mathbf{B}_0^{(L-1)})) \right] \right|, & l = L-1 \end{cases}$$

Since for each $i \in [m]$, $b_{i,0}^{(l)}$, $\mathbf{z}_{i,0}$ and a_i are independent realizations of $\mathbf{B}_0^{(l)}$, \mathbf{Z}_0 and \mathbf{A} , and moreover, each $H_t^{(l)}$ is a Lipschitz function at any finite $t \geq 0$, we know from the law of large numbers that $\text{(II)} = o_{\mathbb{P}}(1)$.

For the second term, if $g \in \text{Lip}(\mathbb{R}^N)$, then

$$\begin{aligned}
 \text{(II)} &= \left| \frac{1}{m} \sum_{i=1}^m g(h_{i,t}^{(l)}(\mathbf{x}'_1), \dots, h_{i,t}^{(l)}(\mathbf{x}'_N)) - \frac{1}{m} \sum_{i=1}^m g(\tilde{h}_{i,t}^{(l)}(\mathbf{x}'_1), \dots, \tilde{h}_{i,t}^{(l)}(\mathbf{x}'_N)) \right| \\
 &\leq \left(\frac{1}{N} \sum_{k=1}^N \left| \Delta h_{m,t}^{(l)}(\mathbf{x}'_k) \right|^2 \right)^{1/2},
 \end{aligned}$$

where we define, $\forall l \in [L-1]$, $\forall \mathbf{x} \in \mathcal{X}$,

$$\Delta h_{m,t}^{(l)}(\mathbf{x}) := \left(\frac{1}{m} \sum_{j=1}^m |h_{j,t}^{(l)}(\mathbf{x}) - \tilde{h}_{j,t}^{(l)}(\mathbf{x})|^2 \right)^{\frac{1}{2}}.$$

Lemma 34 $\frac{1}{N} \sum_{k=1}^N \left| \Delta h_{m,t}^{(l)}(\mathbf{x}'_k) \right|^2 = o_{\mathbb{P}}(1)$.

This lemma is proved in Appendix D.1.1 using a propagation-or-chaos argument (Braun and Hepp, 1977), and it implies that $\text{(II)} = o_{\mathbb{P}}(1)$. This concludes this proof of Theorem 18.

D.1.1 PROOF OF LEMMA 34

We additionally define

$$\begin{aligned}
\Delta\zeta_{m,t}(\mathbf{x}) &:= |\zeta_{m,t}(\mathbf{x}) - \zeta_t(\mathbf{x})| \\
\Delta\mathbf{z}_{m,t} &:= \left(\frac{1}{m} \sum_{j=1}^m |\mathbf{z}_{j,t} - \tilde{\mathbf{z}}_{j,t}|^2 \right)^{\frac{1}{2}}, \\
\Delta a_{m,t} &:= \left(\frac{1}{m} \sum_{i=1}^m |a_{i,t} - \tilde{a}_{i,t}|^2 \right)^{\frac{1}{2}}, \\
\Delta b_{m,t}^{(l)} &:= \left(\frac{1}{m} \sum_{i=1}^m |b_{i,t}^{(l)} - \tilde{b}_{i,t}^{(l)}|^2 \right)^{\frac{1}{2}}, \quad \forall l \in [L-1], \\
\Delta q_{m,t}^{(l)}(\mathbf{x}) &:= \left(\frac{1}{m} \sum_{j=1}^m \left| q_{i,t}^{(l)}(\mathbf{x}) \sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)}) - \tilde{q}_{i,t}^{(l)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) \right|^2 \right)^{\frac{1}{2}}, \quad \forall l \in [L-1], \\
\Delta W_{m,t}^{(l)} &:= \left(\frac{1}{m^2} \sum_{i,j=1}^m |W_{i,j,t}^{(l)} - W_{i,j,0}^{(l)} - \tilde{W}_{i,j,t}^{(l)}|^2 \right)^{\frac{1}{2}}, \quad \forall l \in [L-2] \\
&= \left(\frac{1}{m^2} \|W_t^{(l)} - W_0^{(l)} - \tilde{W}_t^{(l)}\|_{\mathbb{F}}^2 \right)^{\frac{1}{2}} \\
&\geq \left(\frac{1}{m^2} \|(W_t^{(l)} - W_0^{(l)} - \tilde{W}_t^{(l)})^\top (W_t^{(l)} - W_0^{(l)} - \tilde{W}_t^{(l)})\|_2 \right)^{\frac{1}{2}},
\end{aligned}$$

and finally,

$$\begin{aligned}
\Delta_{m,t} &= \sup_{k \in [n]} \Delta\zeta_{m,t}(\mathbf{x}_k) + \Delta\mathbf{z}_{m,t} + \Delta a_{m,t} \\
&\quad + \sum_{l=1}^{L-1} \left(\sup_{k \in [n]} \Delta h_{m,t}^{(l)}(\mathbf{x}_k) + \sup_{k \in [n]} \Delta q_{m,t}^{(l)}(\mathbf{x}_k) + \Delta b_{m,t}^{(l)} \right) + \sum_{l=1}^{L-2} \Delta W_{m,t}^{(l)},
\end{aligned}$$

At initial time, we see that $\Delta_{m,0} = 0$. For $t \geq 0$, we will bound its growth by examining each term on the right-hand side.

(1) $\Delta h_{m,t}^{(l)}$

When $l = 1$,

$$\Delta h_{m,t}^{(1)}(\mathbf{x}) = O(\Delta\mathbf{z}_{m,t} + \Delta b_{m,t}^{(1)})$$

For $l \in \{2, \dots, L-3\}$,

$$\begin{aligned}
 h_{i,t}^{(l+1)}(\mathbf{x}) - \tilde{h}_{i,t}^{(l+1)}(\mathbf{x}) &= (b_{i,t}^{(l+1)} - \tilde{b}_{i,t}^{(l+1)}) + \frac{1}{m} \sum_{j=1}^m W_{i,j,0}^{(l)} \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \\
 &\quad + \frac{1}{m} \sum_{j=1}^m (W_{i,j,0}^{(l)} - W_{i,j,0}^{(l)} - \tilde{W}_{i,j,t}^{(l)}) \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \\
 &\quad + \left(\frac{1}{m} \sum_{j=1}^m \tilde{W}_{i,j,t}^{(l)} \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) - \frac{1}{m} \sum_{j=1}^m \tilde{W}_{i,j,t}^{(l)} \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) \right) \\
 &\quad + \left(\frac{1}{m} \sum_{j=1}^m \tilde{W}_{i,j,t}^{(l)} \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) - H_t^{(l+1)}(\mathbf{x}, b_{i,0}^{(l+1)}) \right)
 \end{aligned}$$

By the Marchenko-Pastur law of the eigenvalues of sample covariance matrices (Marčenko and Pastur, 1967; Bai and Silverstein, 2010), under the assumption that ρ_W has a finite fourth moment, $\frac{1}{m} \|(W_0^{(l)})^\top W_0^{(l)}\|$ converges almost surely to some finite number, and hence

$$\begin{aligned}
 &\frac{1}{m} \sum_{i=1}^m \left| \frac{1}{m} \sum_{j=1}^m W_{i,j,0}^{(l)} \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \right|^2 \\
 &= O\left(\frac{1}{m} \left(1 + \Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_{m,t}^{(l)}\right)^2 \left(\frac{1}{m} \|(W_0^{(l)})^\top W_0^{(l)}\|\right)\right) \\
 &= o_{\mathbb{P}}(1 + (\Delta_{m,t})^2).
 \end{aligned}$$

In addition,

$$\begin{aligned}
 &\frac{1}{m} \sum_{i=1}^m \left| \frac{1}{m} \sum_{j=1}^m (W_{i,j,0}^{(l)} - W_{i,j,0}^{(l)} - \tilde{W}_{i,j,t}^{(l)}) \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \right|^2 \\
 &\leq \frac{1}{m^2} \left\| (W_t^{(l)} - W_0^{(l)} - \tilde{W}_t^{(l)})^\top (W_t^{(l)} - W_0^{(l)} - \tilde{W}_t^{(l)}) \right\|_2 \frac{1}{m} \sum_{j=1}^m |\sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)})|^2 \\
 &= O((\Delta W_{m,t}^{(l)})^2 (\Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_{j,t}^{(l)})^2)
 \end{aligned}$$

Moreover, since the deterministic maps $H_t^{(l)}$ and $\Xi_t^{(l)}$ are Lipschitz at any finite $t \geq 0$, we can deduce from the law of large numbers that $\forall i \in [m]$,

$$\begin{aligned}
 &\left| \frac{1}{m} \sum_{j=1}^m \tilde{W}_{i,j,t}^{(l)} \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) - H_t^{(l+1)}(\mathbf{x}, b_{i,0}^{(l+1)}) \right| \\
 &= \left| \frac{1}{m} \sum_{j=1}^m \Xi_t^{(l)}(b_{i,0}^{(l+1)}, b_{j,0}^{(l)}) \sigma(H_t^{(l)}(\mathbf{x}, b_{j,0}^{(l)}) + B_t^{(l)}(b_{j,0}^{(l)})) \right. \\
 &\quad \left. - \mathbb{E} \left[\Xi_t^{(l)}(b_{i,0}^{(l+1)}, \mathbf{B}_0^{(l)}) \sigma(H_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) + B_t^{(l)}(\mathbf{B}_0^{(l)})) \right] \right| \\
 &= o_{\mathbb{P}}(1).
 \end{aligned}$$

Thus,

$$\begin{aligned} (\Delta h_{m,t}^{(l+1)}(\mathbf{x}))^2 &= O((\Delta b_t^{(l+1)})^2 + (\Delta W_t^{(l)})^2 + (\Delta h_{m,t}^{(l)}(\mathbf{x}))^2 + (\Delta b_t^{(l)})^2) \\ &\quad + o_{\mathbb{P}}(1 + (\Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_t^{(l)})^2) \end{aligned}$$

and so

$$\Delta h_{m,t}^{(l+1)}(\mathbf{x}) = O(\Delta b_{m,t}^{(l+1)} + \Delta W_{m,t}^{(l)} + \Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_{m,t}^{(l)}) + o_{\mathbb{P}}(1 + \Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_{m,t}^{(l)}) .$$

With a similar argument, we can obtain the same bound for $\Delta h_{m,t}^{(2)}$ and $\Delta h_{m,t}^{(l+1)}$.

So, by induction, $\forall l \in [L-1], \forall \mathbf{x} \in \mathcal{X}$,

$$\begin{aligned} \Delta h_{m,t}^{(l)}(\mathbf{x}) &= O\left(\Delta \mathbf{z}_{m,t} + \sum_{l'=1}^l \Delta b_{m,t}^{(l')} + \sum_{l'=1}^{l-1} \Delta W_t^{(l')}\right) + o_{\mathbb{P}}(1 + \Delta h_{m,t}^{(l)}(\mathbf{x}) + \Delta b_{m,t}^{(l)}) \quad (34) \\ &= O(\Delta_{m,t}) + o_{\mathbb{P}}(\Delta_{m,t}) . \end{aligned}$$

(2) $\Delta q_{m,t}^{(l)}$

For $l = L-1$,

$$\begin{aligned} &|q_{i,t}^{(L-1)}(\mathbf{x})\sigma'(h_{i,t}^{(L-1)}(\mathbf{x}) + b_{i,t}^{(L-1)}) - \tilde{q}_{i,t}^{(L-1)}(\mathbf{x})\sigma'(\tilde{h}_{i,t}^{(L-1)}(\mathbf{x}) + \tilde{b}_{i,t}^{(L-1)})| \\ &= |a_{i,t}\sigma'(h_{i,t}^{(L-1)}(\mathbf{x}) + b_{i,t}^{(L-1)}) - \tilde{a}_{i,t}\sigma'(\tilde{h}_{i,t}^{(L-1)}(\mathbf{x}) + \tilde{b}_{i,t}^{(L-1)})| , \end{aligned}$$

and hence

$$\Delta q_{m,t}^{(L-1)}(\mathbf{x}) = O(\Delta a_{m,t} + \Delta h_{m,t}^{(L-1)}(\mathbf{x}) + \Delta b_{m,t}^{(L-1)}) .$$

For $l \in \{3, \dots, L-2\}$,

$$\begin{aligned} &q_{j,t}^{(l-1)}(\mathbf{x})\sigma'(h_{j,t}^{(l-1)}(\mathbf{x}) + b_{j,t}^{(l-1)}) - \tilde{q}_{j,t}^{(l-1)}(\mathbf{x})\sigma'(\tilde{h}_{j,t}^{(l-1)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l-1)}) \\ &= \left(\frac{1}{m} \sum_{i=1}^m W_{i,j,0}^{(l-1)} q_{i,t}^{(l)}(\mathbf{x})\sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)})\right) \sigma'(h_{j,t}^{(l-1)}(\mathbf{x}) + b_{j,t}^{(l-1)}) \\ &\quad + \left(\frac{1}{m} \sum_{i=1}^m (W_{i,j,t}^{(l-1)} - W_{i,j,0}^{(l-1)} - \tilde{W}_{i,j,t}^{(l-1)}) q_{i,t}^{(l)}(\mathbf{x})\sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)})\right) \sigma'(h_{j,t}^{(l-1)}(\mathbf{x}) + b_{j,t}^{(l-1)}) \\ &\quad + \left(\frac{1}{m} \sum_{i=1}^m \tilde{W}_{i,j,t}^{(l-1)} (q_{i,t}^{(l)}(\mathbf{x})\sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)}) - \tilde{q}_{i,t}^{(l)}(\mathbf{x})\sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}))\right) \sigma'(h_{j,t}^{(l-1)}(\mathbf{x}) + b_{j,t}^{(l-1)}) \\ &\quad + \left(\frac{1}{m} \sum_{i=1}^m \tilde{W}_{i,j,t}^{(l-1)} \tilde{q}_{i,t}^{(l)}(\mathbf{x})\sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)})\right) (\sigma'(h_{j,t}^{(l-1)}(\mathbf{x}) + b_{j,t}^{(l-1)}) - \sigma'(\tilde{h}_{j,t}^{(l-1)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l-1)})) \\ &\quad + \left(\frac{1}{m} \sum_{i=1}^m \tilde{W}_{i,j,t}^{(l-1)} \tilde{q}_{i,t}^{(l)}(\mathbf{x})\sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) - \tilde{q}_{j,t}^{(l-1)}(\mathbf{x})\right) \sigma'(\tilde{h}_{j,t}^{(l-1)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l-1)}) . \end{aligned}$$

Note that $\forall j \in [m]$, by the Lipschitzness of the deterministic maps at finite t and the law of large numbers,

$$\begin{aligned}
 & \left| \frac{1}{m} \sum_{i=1}^m \tilde{W}_{i,j,t}^{(l-1)} \tilde{q}_{i,t}^{(l)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) - \tilde{q}_{j,t}^{(l-1)}(\mathbf{x}) \right| \\
 & \leq \left| \frac{1}{m} \sum_{i=1}^m \Xi_t^{(l-1)}(b_{i,0}^{(l)}, b_{j,0}^{(l-1)}) Q_t^{(l)}(\mathbf{x}, b_{i,0}^{(l)}) \sigma'(H_t^{(l)}(\mathbf{x}, b_{i,0}^{(l)}) + B_t^{(l)}(b_{i,0}^{(l)})) \right. \\
 & \quad \left. - \mathbb{E} \left[\Xi_t^{(l-1)}(\mathbf{B}_0^{(l)}, b_{j,0}^{(l-1)}) Q_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) \sigma'(H_t^{(l)}(\mathbf{x}, \mathbf{B}_0^{(l)}) + B_t^{(l)}(\mathbf{B}_0^{(l)})) \right] \right| \\
 & = o_{\mathbb{P}}(1) .
 \end{aligned}$$

Via other techniques analogous to those used in part (1), we see that

$$\begin{aligned}
 \Delta q_{m,t}^{(l-1)}(\mathbf{x}) & = O(\Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta h_{m,t}^{(l-1)}(\mathbf{x}) + \Delta b_{m,t}^{(l-1)} + \Delta W_{m,t}^{(l)} \\
 & \quad + (\Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta h_{m,t}^{(l-1)}(\mathbf{x}) + \Delta b_{m,t}^{(l-1)} + \Delta W_{m,t}^{(l)})^2) + o_{\mathbb{P}}(1 + \Delta q_{m,t}^{(l)}(\mathbf{x}))
 \end{aligned}$$

Similarly, we can obtain the same bound for $\Delta q_{m,t}^{(L-2)}(\mathbf{x})$ and $\Delta q_{m,t}^{(1)}(\mathbf{x})$. Thus, by induction,

$$\Delta q_{m,t}^{(l)} = O\left(\Delta_{m,t} + (\Delta_{m,t})^{2^{L-l-1}}\right) + o_{\mathbb{P}}\left(1 + \Delta_{m,t} + (\Delta_{m,t})^{2^{L-l-2}}\right)$$

(3) $\Delta b_{m,t}^{(l)}$

For $l \in [L-1]$,

$$\begin{aligned}
 \frac{d}{dt}(b_{i,t}^{(l)} - \tilde{b}_{i,t}^{(l)}) & = -\beta \left(\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) q_{i,t}^{(l)}(\mathbf{x}) \sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)}) \right\} \right. \\
 & \quad \left. - \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \tilde{q}_{i,t}^{(l)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) \right\} \right) \\
 & = -\beta \mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) \left(q_{i,t}^{(l)}(\mathbf{x}) \sigma'(h_{i,t}^{(l)}(\mathbf{x}) + b_{i,t}^{(l)}) - \tilde{q}_{i,t}^{(l)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) \right) \right\} \\
 & \quad - \beta \mathcal{E}_{\mathbf{x}} \left\{ (\zeta_{m,t}(\mathbf{x}) - \zeta_t(\mathbf{x})) \tilde{q}_{i,t}^{(l)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l)}) \right\}
 \end{aligned}$$

Thus,

$$\left(\frac{1}{m} \sum_{i=1}^m \left| \frac{d}{dt}(b_{i,t}^{(l)} - \tilde{b}_{i,t}^{(l)}) \right|^2 \right)^{\frac{1}{2}} = O\left(\mathcal{E}_{\mathbf{x}} \left\{ \Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta \zeta_{m,t}(\mathbf{x}) + \Delta \zeta_{m,t}(\mathbf{x}) \Delta q_{m,t}^{(l)}(\mathbf{x}) \right\}\right) ,$$

which implies that

$$\begin{aligned}
 \frac{d}{dt} b_{m,t}^{(l)} & = O\left(\mathcal{E}_{\mathbf{x}} \left\{ \Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta \zeta_{m,t}(\mathbf{x}) + \Delta \zeta_{m,t}(\mathbf{x}) \Delta q_{m,t}^{(l)}(\mathbf{x}) \right\}\right) \\
 & = O\left(\Delta_{m,t} + (\Delta_{m,t})^2\right)
 \end{aligned}$$

(4) $\Delta W_{m,t}^{(l)}$

For $l \in [L - 2]$,

$$\begin{aligned}
& \frac{d}{dt}(W_{i,j,t}^{(l)} - \tilde{W}_{i,j,t}^{(l)}) \\
&= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) q_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(h_{i,t}^{(l+1)}(\mathbf{x}) + b_{i,t}^{(l+1)}) \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \right\} \\
&\quad + \mathcal{E}_{\mathbf{x}} \left\{ \zeta_t(\mathbf{x}) \tilde{q}_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l+1)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l+1)}) \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) \right\} \\
&= -\mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) \left(q_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(h_{i,t}^{(l+1)}(\mathbf{x}) + b_{i,t}^{(l+1)}) - \tilde{q}_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l+1)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l+1)}) \right) \right. \\
&\quad \left. \cdot \sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) \right\} \\
&\quad - \mathcal{E}_{\mathbf{x}} \left\{ \zeta_{m,t}(\mathbf{x}) \tilde{q}_{i,t}^{(l+1)}(\mathbf{x}) \sigma'(\tilde{h}_{i,t}^{(l+1)}(\mathbf{x}) + \tilde{b}_{i,t}^{(l+1)}) \left(\sigma(h_{j,t}^{(l)}(\mathbf{x}) + b_{j,t}^{(l)}) - \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) \right) \right\} \\
&\quad - \mathcal{E}_{\mathbf{x}} \left\{ (\zeta_{m,t}(\mathbf{x}) - \zeta_t(\mathbf{x})) \tilde{q}_{i,t}^{(l+1)}(\mathbf{x}) \sigma(\tilde{h}_{j,t}^{(l)}(\mathbf{x}) + \tilde{b}_{j,t}^{(l)}) \right\}
\end{aligned}$$

Thus,

$$\begin{aligned}
& \frac{1}{m^2} \sum_{i,j=1}^m \left(\frac{d}{dt}(W_{i,j,t}^{(l)} - \tilde{W}_{i,j,t}^{(l)}) \right)^2 \\
&= O \left(\mathcal{E}_{\mathbf{x}} \left\{ (1 + \Delta \zeta_{m,t}(\mathbf{x})) (\Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta h_{m,t}^{(l-1)}(\mathbf{x}) + \Delta b_{m,t}^{(l-1)})^2 + \Delta \zeta_{m,t}(\mathbf{x}) \right\} \right)
\end{aligned}$$

and so

$$\begin{aligned}
\frac{d}{dt} \Delta W_{m,t}^{(l)} &= O \left(\mathcal{E}_{\mathbf{x}} \left\{ (1 + \Delta \zeta_t(\mathbf{x})) (\Delta q_{m,t}^{(l)}(\mathbf{x}) + \Delta h_{m,t}^{(l-1)}(\mathbf{x}) + \Delta b_{m,t}^{(l-1)}) + \Delta \zeta_{m,t}(\mathbf{x}) \right\} \right) \\
&= O \left(\Delta_{m,t} + (\Delta_{m,t})^2 \right) .
\end{aligned}$$

(5) $\Delta z_{m,t}$

$$\begin{aligned}
\frac{d}{dt} \Delta z_{m,t} &= O \left(\mathcal{E}_{\mathbf{x}} \left\{ \Delta \zeta_{m,t}(\mathbf{x}) + \Delta q_{m,t}^{(1)}(\mathbf{x}) + \Delta \zeta_{m,t}(\mathbf{x}) \Delta q_{m,t}^{(1)}(\mathbf{x}) \right\} \right) \\
&= O \left(\Delta_{m,t} + (\Delta_{m,t})^2 \right) .
\end{aligned}$$

(6) $\Delta a_{m,t}$

$$\begin{aligned}
& \frac{d}{dt} \Delta a_{m,t} \\
&= O \left(\mathcal{E}_{\mathbf{x}} \left\{ \Delta \zeta_{m,t}(\mathbf{x}) + \Delta h_{m,t}^{(L-1)}(\mathbf{x}) + \Delta b_{m,t}^{(L-1)} + \Delta \zeta_{m,t}(\mathbf{x}) (\Delta h_{m,t}^{(L-1)}(\mathbf{x}) + \Delta b_{m,t}^{(L-1)}) \right\} \right) \\
&= O \left(\Delta_{m,t} + (\Delta_{m,t})^2 \right) .
\end{aligned}$$

(7) $\Delta \zeta_{m,t}$

$$\begin{aligned}
\Delta \zeta_{m,t}(\mathbf{x}) &= O \left(\mathcal{E}_{\mathbf{x}} \left\{ \Delta a_{m,t} + \Delta h_{m,t}^{(L-1)}(\mathbf{x}) + \Delta b_{m,t}^{(L-1)} + \Delta a_{m,t} (\Delta h_{m,t}^{(L-1)}(\mathbf{x}) + \Delta b_{m,t}^{(L-1)}) \right\} \right) \\
&= O \left(\Delta_{m,t} + (\Delta_{m,t})^2 \right) .
\end{aligned}$$

Therefore, for the duration in which $\Delta_{m,t} \leq 1$, it holds that

$$\frac{d}{dt} \Delta_{m,t} = O(\Delta_{m,t}) + o_{\mathbb{P}}(1) ,$$

and hence, by Grönwall's inequality,

$$\Delta_{m,t} = o_{\mathbb{P}}(1) . \quad (35)$$

Thus, for any finite $t \geq 0$, when m is large enough, we can always ensure that $\Delta_{m,t} \leq 1$. This implies that (35) holds for all finite $t \geq 0$.

Finally, applying (34) to each $\mathbf{x} \in \{\mathbf{x}'_1, \dots, \mathbf{x}'_N\}$, we arrive at Lemma 34.

D.2 Derivation of the Training Dynamics of Deep Linear NNs

In the linear NN case,

$$\begin{aligned} & \kappa_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') \\ &= \mathbb{E} \left[\mathbf{H}_t^{(l)}(\mathbf{x}) \mathbf{H}_s^{(l)}(\mathbf{x}') \right] \\ &= \mathbb{E} \left[\left(\mathbf{H}_0^{(l)}(\mathbf{x}) - \int_0^t \mathcal{E}_{\mathbf{x}''} \left\{ \zeta_r(\mathbf{x}'') \kappa_{t,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \mathbf{Q}_r^{(l)}(\mathbf{x}'') \right\} dr \right) \right. \\ & \quad \left. \left(\mathbf{H}_0^{(l)}(\mathbf{x}') - \int_0^s \mathcal{E}_{\mathbf{x}''} \left\{ \zeta_r(\mathbf{x}'') \kappa_{s,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \mathbf{Q}_r^{(l)}(\mathbf{x}'') \right\} dr \right) \right] \\ &= \kappa_{0,0}^{(l)}(\mathbf{x}, \mathbf{x}') + \int_0^t \int_0^s \mathcal{E}_{\mathbf{x}'', \mathbf{x}'''} \left\{ \zeta_r(\mathbf{x}'') \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l)}(\mathbf{x}'', \mathbf{x}''') \kappa_{t,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \kappa_{s,p}^{(l-1)}(\mathbf{x}', \mathbf{x}''') \right\} dr dp \\ & \quad - \int_0^t \mathcal{E}_{\mathbf{x}''} \left\{ \zeta_r(\mathbf{x}'') \kappa_{t,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \mathbb{E}^{(l)} \left[\mathbf{H}_0^{(l)}(\mathbf{x}') \mathbf{Q}_r^{(l)}(\mathbf{x}'') \right] \right\} dr \\ & \quad - \int_0^s \mathcal{E}_{\mathbf{x}''} \left\{ \zeta_r(\mathbf{x}'') \kappa_{s,r}^{(l-1)}(\mathbf{x}', \mathbf{x}'') \mathbb{E}^{(l)} \left[\mathbf{H}_0^{(l)}(\mathbf{x}) \mathbf{Q}_r^{(l)}(\mathbf{x}'') \right] \right\} dr \end{aligned}$$

Since

$$\begin{aligned} \mathbb{E} \left[\mathbf{H}_0^{(l)}(\mathbf{x}) \mathbf{Q}_r^{(l)}(\mathbf{x}'') \right] &= \mathbb{E} \left[\mathbf{H}_0^{(l)}(\mathbf{x}') \int_0^r \mathcal{E}_{\mathbf{x}'''} \left\{ \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l+1)}(\mathbf{x}'', \mathbf{x}''') \mathbf{H}_p^{(l)}(\mathbf{x}''') \right\} dp \right] \\ &= \int_0^r \mathcal{E}_{\mathbf{x}'''} \left\{ \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l+1)}(\mathbf{x}'', \mathbf{x}''') \kappa_{0,p}^{(l)}(\mathbf{x}, \mathbf{x}''') dp \right\} , \end{aligned}$$

we then have

$$\begin{aligned} & \kappa_{t,s}^{(l)}(\mathbf{x}, \mathbf{x}') \\ &= \kappa_{0,0}^{(l)}(\mathbf{x}, \mathbf{x}') + \int_0^t \int_0^s \mathcal{E}_{\mathbf{x}'', \mathbf{x}'''} \left\{ \zeta_r(\mathbf{x}'') \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l)}(\mathbf{x}'', \mathbf{x}''') \kappa_{t,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \kappa_{s,p}^{(l-1)}(\mathbf{x}', \mathbf{x}''') \right\} dr dp \\ & \quad - \int_0^t \mathcal{E}_{\mathbf{x}'', \mathbf{x}'''} \left\{ \zeta_r(\mathbf{x}'') \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l+1)}(\mathbf{x}'', \mathbf{x}''') \kappa_{t,r}^{(l-1)}(\mathbf{x}, \mathbf{x}'') \kappa_{0,p}^{(l)}(\mathbf{x}', \mathbf{x}''') \right\} dr \\ & \quad - \int_0^s \mathcal{E}_{\mathbf{x}'', \mathbf{x}'''} \left\{ \zeta_r(\mathbf{x}'') \zeta_p(\mathbf{x}''') \gamma_{r,p}^{(l+1)}(\mathbf{x}'', \mathbf{x}''') \kappa_{s,r}^{(l-1)}(\mathbf{x}', \mathbf{x}'') \kappa_{0,p}^{(l)}(\mathbf{x}, \mathbf{x}''') \right\} dr . \end{aligned}$$

Notice that $\kappa_{0,p}^{(l)}(\mathbf{x}, \mathbf{x}') = 0, \forall l > 1, \forall p \geq 0$ while $\kappa_{0,0}^{(1)}(\mathbf{x}, \mathbf{x}') = \kappa_{0,p}^{(0)}(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\top \cdot \mathbf{x}', \forall p \geq 0$. Thus, using the linearity, we can derive (15) for $l \in [L-2]$, and moreover,

$$\begin{aligned} K_{t,s}^{(1)} &= 1 + \int_0^t \int_0^s c_{r,p}^{(1)} \zeta_r \cdot \zeta_p^\top dp dr \\ &\quad - \int_0^t \int_0^r c_{r,p}^{(2)} \zeta_r \cdot \zeta_p^\top \cdot K_{p,0}^{(1)} dp dr \\ &\quad - \int_0^s \int_0^r c_{r,p}^{(2)} K_{0,p}^{(1)} \cdot \zeta_p \cdot \zeta_r^\top dp dr . \end{aligned}$$

With a similar argument, we can derive (16) for $l \in [L-2]$, and moreover,

$$\begin{aligned} c_{t,s}^{(L-1)} &= 1 + \int_0^t \int_0^s \zeta_r^\top \cdot K_{r,p}^{(L-1)} \cdot \zeta_p^\top dp dr \\ &\quad - \int_0^t \int_0^r c_{p,0}^{(L-1)} \zeta_r^\top \cdot K_{r,p}^{(L-2)} \cdot \zeta_p^\top dp dr \\ &\quad - \int_0^s \int_0^r c_{p,0}^{(L-1)} \zeta_r^\top \cdot K_{r,p}^{(L-2)} \cdot \zeta_p^\top dp dr . \end{aligned}$$

Appendix E. Supplementary Materials for Section 6

E.1 Proof of Lemma 12

For any function $f \in \mathcal{F}^{(L)}$, we let $(\mathcal{H}^{(l)})_{l \in [L]}$ be an L -level NHL that it belongs to. Hence, there exist $\mu^{(1)}, \dots, \mu^{(L-1)}$ that satisfy the condition given in Definition 2, and furthermore, there exists $\xi \in L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)})$ such that

$$f(\mathbf{x}) = \int_{\mathcal{H}^{(L-1)}} \xi(h) \sigma(h(\mathbf{x})) \mu^{(L-1)}(dh) , \quad \forall \mathbf{x} \in \mathcal{X} . \quad (36)$$

We can decompose $\xi = \xi_+ - \xi_-$, where ξ_+ and $\xi_- \in L^2(\mathcal{H}^{(L-1)}, \mu^{(L-1)})$ are both non-negative on $\mathcal{H}^{(L-1)}$, and then define

$$\begin{aligned} f_+(\mathbf{x}) &:= \int_{\mathcal{H}^{(L-1)}} \xi_+(h) \sigma(h(\mathbf{x})) \mu^{(L-1)}(dh) , \quad \forall \mathbf{x} \in \mathcal{X} , \\ f_-(\mathbf{x}) &:= \int_{\mathcal{H}^{(L-1)}} \xi_-(h) \sigma(h(\mathbf{x})) \mu^{(L-1)}(dh) , \quad \forall \mathbf{x} \in \mathcal{X} . \end{aligned}$$

Since ReLU has non-negative outputs, we see that f_+ and f_- are both non-negative functions on \mathcal{X} . Furthermore, they satisfy $f = f_+ - f_-$, $\|f_+\|_{\mathcal{H}^{(L)}}^2 \leq \int_{\mathcal{H}^{(L-1)}} |\xi_+(h)|^2 \mu^{(L-1)}(dh)$ and $\|f_-\|_{\mathcal{H}^{(L)}}^2 \leq \int_{\mathcal{H}^{(L-1)}} |\xi_-(h)|^2 \mu^{(L-1)}(dh)$, and hence

$$\|f_+\|_{\mathcal{H}^{(L)}}^2 + \|f_-\|_{\mathcal{H}^{(L)}}^2 \leq \int_{\mathcal{H}^{(L-1)}} |\xi(h)|^2 \mu^{(L-1)}(dh) .$$

Thus, we can further define $\mu^{(L)} := \frac{1}{2} \delta_{f_+} + \frac{1}{2} \delta_{f_-} \in \mathcal{P}(\mathcal{H}^{(L)})$ as well as $\bar{\xi} \in L^2(\mathcal{H}^{(L)}, \mu^{(L)})$ as $\bar{\xi} = 2\mathbf{1}_{f_+} - 2\mathbf{1}_{f_-}$, and let $\mathcal{H}^{(L+1)} = \mathcal{H}_{\mu^{(L)}}$. It holds that

$$\int_{\mathcal{H}^{(L)}} \bar{\xi}(h) \sigma(h(\mathbf{x})) \mu^{(L)}(dh) = \frac{1}{2} \cdot 2\sigma(f_+(\mathbf{x})) - \frac{1}{2} \cdot 2\sigma(f_-(\mathbf{x})) = f(\mathbf{x}) .$$

and therefore,

$$\begin{aligned}
 \|f\|_{\mathcal{H}^{(L+1)}}^2 \|\mu^{(L)}\|_{\mathcal{H}^{(L)},2}^2 &\leq \left(\int_{\mathcal{H}^{(L)}} |\bar{\xi}(h)|^2 \mu^{(L)}(dh) \right) \|\mu^{(L)}\|_{\mathcal{H}^{(L)},2}^2 \\
 &\leq 4 \cdot \frac{1}{2} (\|f_+\|_{\mathcal{H}^{(L)}}^2 + \|f_-\|_{\mathcal{H}^{(L)}}^2) \\
 &\leq 2 \int_{\mathcal{H}^{(L-1)}} |\xi(h)|^2 \mu^{(L-1)}(dh) .
 \end{aligned}$$

By optimizing over all ξ such that (36) holds, we derive that the left-hand side can be upper-bounded by $2\|f\|_{\mathcal{H}^{(L)}}^2$. Therefore, f also belongs to the $(L+1)$ -level NHL, $(\mathcal{H}^{(l)})_{l \in [L+1]}$, with $\mathcal{E}^{(L+1)}(f) \leq \sqrt{2}\mathcal{E}^{(L)}(f)$.